Evidential Relevance and Expressiveness of Digital Traces: An Investigative Perspective

Relevanz und Aussagekraft digitaler Spuren:
Eine ermittlungsorientierte Perspektive

Der Technischen Fakultät der
Friedrich-Alexander-Universität
Erlangen-Nürnberg
zur
Erlangung des Doktorgrades Dr.-Ing.
vorgelegt von

Jan Gruber

Als Dissertation genehmigt von
der Technischen Fakultät der
Friedrich-Alexander-Universität
Erlangen-Nürnberg

Tag der mündlichen Prüfung: 4. April 2024

Gutachter: Prof. Dr.-Ing. F. C. Freiling
Prof. Dr. Z. J. M. H. Geradts

Abstract

In this day and age, almost any criminal investigation deals with some pieces of digital evidence. Given the wealth of digital data stored on both end-user devices and cloud infrastructure, a tremendous challenge for investigators and prosecutors is to determine the relevant pieces to solve the case; however, given an investigative question, there exists no straightforward method to find "sufficient digital evidence" to do so. Thus, the present thesis leaps to improve the understanding and interpretation of digital traces for criminal investigations on a foundational level. As a unifying result, we propose the *Cyber-traceological Model*, which provides a general way to translate investigative hypotheses to relevant traces—both in idealized and real-world scenarios. The model is grounded in formal definitions of when traces are generally relevant and how they can be expressive on a conceptual level. Building up on these concepts, we are able to define an investigative knowledge base in a precise manner. For digital systems, we then show how relevance can be determined to fill the knowledge base by calculating necessary and sufficient evidence in state machine representations. We use these concepts to refer to rigorous notions of different classes of reconstructability that investigators can use to uncover and comprehend past events. We expressed the concepts of necessity and sufficiency of digital traces in temporal logic and employed a model checker to calculate traces of those classes based on a model of the system under investigation to demonstrate practical feasibility. Since this necessitates the availability of a representation of the system under investigation as a transition system, which is often hard to achieve in real-world scenarios, we additionally investigate ways of collecting, representing, and using phenomenon-specific knowledge of criminal phenomena to establish a notion of evidential relevance from a more holistic and realistic perspective. Using cognitive maps as a particular form to express node-link relationships, we show how this phenomenon-specific knowledge can build a bridge from abstract process models to case-specific concretizations by constituting a meso-level abstraction supporting the quest to find relevant traces more pragmatically. We vividly illustrate the construction of an instance of such a phenomenon-specific knowledge base and its applicability in the example of botnet crime. Lastly, we study how expressiveness of digital traces could be hampered by undetected contamination effects. Here, we provide a novelly universal definition of evidence contamination—applicable both for physical and digital evidence—and aim to substantiate and validate the proposed definition by presenting examples, counterexamples, and edge cases of contamination of digital evidence to build the grounds for future research improving the understanding of contamination. In essence, the results of this dissertation are aggregated in the proposed Cyber-traceological Model, which systematically sketches out how to translate case-related hypotheses into relevant traces. It aims to span the arc from abstract considerations to concrete investigative work, thus hinting at the potential to solidify the practical application by insights gained from theoretical considerations of fundamental attributes of digital evidence.

Zusammenfassung

Digitale Beweismittel spielen heutzutage in beinahe jedem Ermittlungsverfahren eine zunehmend bedeutsame Rolle. Angesichts der Fülle an digitalen Daten, die sowohl auf Endgeräten als auch in der Cloud-Infrastruktur gespeichert sind, stehen Ermittler vor der enormen Herausforderung, die für die Lösung des Falles relevanten Spuren zu ermitteln. Bislang gibt es allerdings keine effektive Methode, um systematisch diejenigen digitalen Spuren zu identifizieren, die zur Beantwortung einer bestimmten Ermittlungsfrage benötigt werden. Die vorliegende Arbeit zielt deshalb darauf ab, das Verständnis und die Interpretation digitaler Spuren für forensische Ermittlungen auf einer grundlegenden Ebene zu verbessern. Als integrierendes Ergebnis stellen wir hierzu das digital-spurenkundliche Modell vor, das eine allgemeinen Methode zur Übersetzung von Ermittlungshypothesen in relevante Spuren bereitstellt, die sowohl in idealisierten als auch in realen Umgebungen angewandt werden kann. Im Grundsatz basiert dieses Modell auf den formalen Definitionen der Relevanz und Aussagekraft von Spuren, was zunächst auf konzeptioneller Ebene ausgedrückt wird. Darauf aufbauend kann eine investigative Wissensbasis auf formale Weise konstruiert werden. Um diese zu füllen, zeigen wir dann was unter notwendigen und hinreichenden Beweisen zu verstehen ist und wie diese anhand der Repäsentation digitaler Systeme als Zustandsautomaten berechnet werden können. Diese Spurenklassen werden mittels temporaler Logik, genauer gesagt *Linear-time Temporal Logic*, ausgedrückt und können mit Hilfe eines Model Checkers berechnet werden, was die praktische Anwendbarkeit der Methode illustriert. Hierbei ist allerdings zu beachten, dass dies die Verfügbarkeit eines Modells des zu untersuchenden Systems erfordert, was die reale Anwendung erschwert. Demnach untersuchen wir die Nutzung phänomenspezifischen Wissens als ergänzende Möglichkeit, um die Relevanz digitaler Spuren aus einer ganzheitlicheren und inbesondere auf reale Anwendungen ausgerichteten Perspektive zu erheben, zu beschreiben und zur Verfügung zu stellen. Unter Verwendung von kognitiven Karten als einer speziellen Form zur verknüpften Darstellung von Information zeigen wir, wie dieses phänomenspezifische Wissen eine Brücke von abstrakten Prozessmodellen zur konkreten Fallarbeit schlagen kann, was die Suche nach relevanten Spuren auf pragmatischere Weise unterstützt. Zur praktischen Illustration beschreiben wir die Erstellung und den Aufbau einer solchen phänomenspezifischen Wissensbasis sowie deren Nutzung am Beispiel der Botnetzkriminalität, um die Anwendbarkeit des Vorgehens zum Auffinden aussagekräftiger Spuren zu veranschaulichen. Schließlich untersuchen wir, wie die Aussagekraft digitaler Spuren durch übersehene Kontaminationseffekte beeinträchtigt wird. Hierzu wird eine universelle Definition der Kontamination von Beweismitteln erarbeitet, die sowohl für physische als auch für digitale Beweismittel anwendbar ist. Die vorgeschlagene Definition wird durch die Betrachtung von Beispielfällen, Gegenbeispielen und Grenzfällen der Kontamination digitaler Beweismittel dargelegt und validiert, was eine Grundlage für künftige Forschungsarbeiten zur Verbesserung des Verständnisses des Wesens von Kontamination schaffen soll. Letztlich werden die im Rahmen dieser Dissertation erarbeiteten Ergebnisse im sogenannten digital-spurenkundlichen Modell integriert, das eine systematische Methode skizziert, wie fallbezogene Hypothesen in relevante Spuren übersetzt werden können. Dies soll den Bogen von abstrakten Überlegungen zu grundlegenden Eigenschaften digitaler Spuren hin zur konkreten Ermittlungsarbeit spannen und damit das Potential darlegen, die forensische Praxis durch eine Fundierung in theoretischen Überlegungen zu verbessern.

Acknowledgments

First and foremost, I thank my doctoral advisor, Felix Freiling, for his always warm-hearted and compassionate mentoring during my dissertation. By offering to work at his IT Security Infrastructures Lab, he made it possible for me to take another career path, conduct research, and pursue a doctoral degree. This opportunity surely has been a very rewarding path, on which my *Doktorvater*'s thoughtful supervision, continuous support, and positive outlook were invaluable. Looking back, I am very grateful for the many interesting and insightful conversations with him, which fundamentally influenced my view of problem-solving, research, and knowledge.

I thank Zeno Geradts for his valuable remarks on my dissertation plan and his willingness to be the second examiner of this thesis. Furthermore, I express gratitude for collaborating with the co-authors of my preceding papers, Lutz Schröder, Christopher Hargreaves, Merlin Humml, Lena Voigt, and Zinaida Benenson. I am especially grateful for the fruitful collaboration with the before-mentioned Merlin Humml from the Chair of Theoretical Computer Science, who eased my approach to theoretical topics and became a good friend beyond that. Moreover, I thank the members of the forensic computing group for exchanging ideas, discussing problems, and proofreading papers. The research training group 2475 "Cybercrime and Forensic Computing" stimulated my research endeavors as well; the lively discussions during workshops and seminars provided valuable impetus for further considerations with a broadened perspective. I also want to thank all my colleagues at the lab for making the lunch and coffee breaks very worthwhile.

More specifically, I express my gratitude to the kind alumni, colleagues, or friends who commented on papers on which this thesis is based or even earlier versions of this dissertation: Andreas Dewald, Céline Vanini, Christian Lindenmeier, Christian Riess, Gaston Pugliese, Konstantin Bayreuther, Kevin Gomez Buquerin, Lena Voigt, Merlin Humml, Janine Schneider, Julian Geus, Ralf Moll, and Üsame Cengiz—thank you for your effort!

Going back in time, I reckon that the step to take up the endeavor of academic research was strongly influenced by my brother, who has been devoting himself to science so far; I thank him for implicitly pointing to this perspective and being my first friend since the early days of our childhood. Finally, I thank my parents for their continuous encouragement and incessant support of my dissertation project as well as my entire life. Thank you!

This work was partially supported by *Deutsche Forschungsgemeinschaft* (DFG, German Research Foundation) as part of the Research and Training Group 2475 "Cybercrime and Forensic Computing" (grant number 393541319/GRK2475/1-2019).

Contents

Acronyms

AE action-induced evidence
CE characteristic evidence
CM cognitive map
CSAM child sexual abuse material
CXE characteristic counter evidence
DCO device configuration overlay
DFA deterministic finite automaton
DF digital forensics
EDR endpoint detection and response
FSM finite-state machine
FoCC Facet-oriented Criminalistic Cycle
GCL Guarded Command Language
GRP general reconstruction problem
HPA host protected area
IoT Internet of Things
JFIF JPEG File Interchange Format
LTL linear-time temporal logic
LTS labeled transition system
ME merged evidence
NE necessary evidence
OS operating system
POS point of sale
RAM random-access memory
RCE remote code execution
RDP Remote Desktop Protocol
SE sufficient evidence
SOC security operations center
SOP standard operating procedure
SRP specific reconstruction problem
TLA temporal logic of actions
Tor The Onion Router
TTPs tactics, techniques, and procedures
VPN virtual private network
XMPP Extensible Messaging and Presence Protocol

List of Figures

List of Tables

Listings

1 Introduction

"What happened and who did it?"—questions regarding the course of the deed and the perpetration are central to any forensic investigation. Given the general paradigm that some cause implies a specific effect, forensic science and especially criminalistics are involved in reasoning backward—from some observed effect to the respective cause [15, 198], as illustrated by Fig. 1.1.

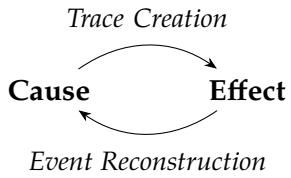

Figure 1.1: Simple yet universal paradigm of trace creation in the context of criminalistics.

Since the early days of forensic science, society has been witnessing several groundbreaking advances to attribute effects to their cause and originator—starting with Karl Landsteiner's discovery of blood groups [146], Edmond Locard's *exchange principle* [152] to Alec John Jeffreys' discovery of DNA fingerprinting [132]. Each of these breakthroughs has vastly improved forensic science, making the world a more just place. Looking at the subjects of these advances, it becomes clear that reasoning from effect to cause is relatively straightforward when dealing with material evidence, such as bloodstains, shoe marks, and the like. With the proliferation of digital evidence, however, a whole new category of challenges occurred. First and foremost, we have to note that the quest to find case-related traces constituting effects of causes in the digital realm with literally trillions of digital objects of various kind aggravates the criminalistic task. Second, questions in regard to the evidential value of this sort of (at least relatively) "new" evidence, its analysis, and its aggregated usage to achieve solid attribution and to solve the "problem of identity" [141] are not even considered yet. Hence, the present thesis leaps to improve the understanding and interpretation of digital traces on a foundational level.

1.1 Motivation

Imagine the following situation: Law enforcement detectives arrive at a crime scene. Being inside the perimeter, they grasp the situation. In their mind, there is a reciprocation of receiving effects and finding hypotheses about their causes and their relation to the deed to match their observations to the various explanation attempts.

It becomes apparent that it is essential for this task, commonly named forensic event reconstruction, to have a clear understanding of traces and their meaning. This requirement has been identified early on as the works of influential researchers of the discipline, such as its founding father Hans Gross [98] or the influentlial forensic scientist Paul Kirk [141], corroborate. Both took on the elaborate quest of cataloging physical evidence and looked at various kinds of traces in different criminal contexts and the respective analysis options available at that time. Although they had written these works multiple decades ago, in which technological analysis methods vastly improved, they essentially described many important (classes of) traces that still seem partly applicable for archaic offenses such as violent crimes even today.

In most cases nowadays, crime scenes have a digital dimension due to the widespread use of computing devices for all sorts of purposes, including personal communication, e-commerce or other business purposes, smart home applications, and many more. Not too rarely, the findings "making the case" are even predominantly comprised of digital evidence—especially when dealing with the ever-increasing cybercrime offenses [80, p. 6].

Given the wealth of data available and its increasing complexity, the question of which evidence is relevant to prove which course of actions or which offense is highly pressing. However, with the advent of the vast amount of digital traces in criminal proceedings, we can see a distinct difference between physical and digital evidence: While the former is based on universally valid laws of nature—even when dealing with human-made analogous items, the latter is a result of human minds designing a piece of software instructing a machine, which creates certain artifacts used as evidence. Though practically relevant, building an encyclopedia of digital traces, as Hans Gross [98] or Kirk [141] started for physical evidence, seems only partly meaningful from a research perspective since it will always be several steps behind the newest app or the latest version of an operating system available. Given the fast-paced technological advances, such an endeavor—nonetheless practically vital—is doomed to be incomplete; thus, the present work tackles more foundational questions.

Such a foundational task is the translation of investigative (and hereby connected legal) demands to digital evidence. Given an investigative question, there exists no straightforward method to find "sufficient digital evidence" based on that. It is easy to grasp that this is a problem of considerable complexity since there are numerous factual and procedural aspects of the legal domain that must be coordinated with highly technical questions of the field of forensic computing—most of them being not tackled yet. We argue that this distant goal should be divided into two subproblems: The first one is the quest to put up hypotheses of pertinence for the investigation, and the second one is to find relevant digital traces that can be used to assess those previously identified case-related hypotheses. Both subproblems are unsolved yet. To approach this, we deal with the second subproblem and related questions, such as the sufficiency and necessity of digital traces for event reconstruction, the impedance of their expressiveness, and the interplay of investigative findings in the present thesis. The reader probably noticed the numerous use of the adjective "digital" in the preceeding paragraphs; however, if one steps back and re-reads this section, the same questions are still applicable for physical traces and have

not been exhaustively answered by traditional forensic science yet—an observation gently hinting at a unification of the disciplines.

1.2 A Unification of the Forensic Disciplines?

Digital forensic science has advanced rapidly over the past decades due to the increasing reliance on technology and digital devices in everyday life; as a result, many techniques and methodologies have been developed to investigate and analyze digital evidence. In the following paragraphs, we look at the relationship between traditional and digital forensic science and point out synergies as well as possible potentials to generally apply model-based approaches, as employed in this thesis in the field of digital forensic science, in the traditional branch.

1.2.1 Synergies Stemming from Digital Technology

We have observed several synergy effects with other branches of traditional forensic science back and forth. Of course, digital forensics is based on the foundational principles of traditional forensics science (such as the preservation of evidence, the establishment of the chain of custody, and the unambiguous analysis of data), and large parts of the research were thus concerned with accommodating the forensics requirements when working with digital data. By applying these principles, digital investigators aim to ensure that evidence is collected and analyzed in a manner that is scientifically sound and defensible in court. But the branch of digital forensics also contributed back to other branches of the superordinate discipline. A way in which traditional forensics has benefited from digital forensics is through the use of digital technology to aid in investigations. There is a myriad of examples available. Most apparent is the use of electronic tools like 3D scanners, dome cameras, and drones to capture and document crime scenes more thoroughly and reconstruct those using virtual reality techniques. Still, there is more: We generally observe a trend toward "digitized crime scene analysis". This umbrella term is used to refer to the collection, processing, and examination of traces in a digitized form incorporating digitized latent fingerprint forensics, digitized fiber forensics, digitized firearm toolmark forensics, digitized locksmith toolmark forensics, and so on [163]. Besides that, forensic computer scientists have developed a set of skills and techniques that are not traditionally used in conventional forensic investigations, given that digital forensics requires a deep understanding of computer systems, networking, and software engineering. Since it relies heavily on the use of algorithms and data analysis techniques to uncover hidden information and patterns in large sets of data, some of these same techniques and statistical approaches from data science can be applied to actual casework or applications of forensic intelligence [193] to sift through large amounts of case data and to help investigators identify patterns and evidence in there that may not be immediately conspicuous.

1.2.2 Specific Characteristics of Digital Forensic Science

The salient question here is, however, if research in digital forensic science is able to generate scientific insight going beyond mere application from which traditional forensic science benefits. Looking at the history of computer science and the nature of digital systems, we observe a trend that might suggest that.

Digital evidence is commonly considered to be comprised of multiple abstraction layers on top of a physical carrier [30, 87]. On the lowest level, the physical coding follows the laws of nature and (quantum) physics. For instance, this could be the silicon in flash storage, the magnetic platters of hard drives, and similar objectification. From this layer onwards (or upwards), everything is digitized—meaning that arbitrary abstractions can be stacked on top. Controversial opinions even suggest that one may imagine that the universality of computation could construct basically everything in the analogous world and much more in those upper layers—an argument which is also put forward by the supporters of the simulation theory by Bostrom [22]. While this appears to be mere speculation at this point, digital forensics is of increasing importance without contention. Obviously, it still is a single branch of the vast reservoir of disciplines that are sorted into forensic science; however, we observe that the forensic problems posed in the digital realm point out largely general problems and require a foundational approach to tackle those.

1.2.3 Potentials of Model-based Approaches in Forensic Science

While most forensic disciplines operate on concrete, tangible evidence, the nature of evidence in digital forensics is more abstract. Computer scientists in general are used to work in intangible environments; hence, they commonly simplify structures for two reasons: Firstly, they impose restrictions on the models to prevent any ambiguity or confusion, as marginal phenomena can be excluded by deliberately limiting the model. Secondly, simplifying the model to focus on the essential components, using symbolic notations and removing any unnecessary details can help to gain insights into complex connections that are hard to grasp otherwise. Restrictions of the matter may lead to an expansion of the understanding; vagueness is eliminated on the one hand, and implicitity is turned outward on the other.

By creating a simplified representation of a complex system or process, researchers may discover that there are key factors that they were previously unaware of, that are not well understood, or that require further investigation. Therefore, this characteristic can help to identify gaps in knowledge and highlight areas where further research is needed. Additionally, such models can be used to communicate complex ideas and concepts in a more concise yet accessible way. Furthermore, models can be used to facilitate the development and testing of new theories and ideas. By creating a simplified representation of a complex system or process, researchers can look at it in a focused way, test different hypotheses, and explore the potential outcomes of different scenarios. This can help to identify new patterns and relationships, leading to the development of new theories and ideas.

By using model-driven approaches, we suspect that they have the potential to provide universal insights and—with bold fantasy—one might even imagine that the theories developed by digital forensics could rise and become metatheories for other disciplines of the field. Such a development would then constitute an inversion of the parent-child relationship; of course, this may be a somewhat far-fetched idea, but there seems to be tentative potential that other branches of forensic science could profit from the results of reasoning about the abstract nature of things using models as it is natural maybe even imposed when working with digital systems—an assumption that is gently hinted at in various places in the further course of the thesis since we employ such a model-based approach to tackle the research questions that are discussed in the next section.

1.3 Research Questions

In Section 1.1, we already mapped out the underlying question in criminalistics, i.e., which traces can be used to prove which offense; we then inferred the connected subproblems of (1) finding appropriate case-related hypotheses and (2) finding traces relevant to those pre-determined hypotheses—each standing for its own yet being connected. Here, we focus on the meaningfulness and usefulness of digital evidence from a principal viewpoint to address that second subproblem. While questions of trustworthiness and integrity of digital evidence have been examined deeply in the past, we identify and investigate essential qualities of (digital) evidence that are more outcome-oriented though not well understood yet. Those are the necessity and sufficiency as well as the relevance and expressiveness of digital traces. Ultimately, we aim to determine what constitutes strong digital evidence and how it can be used effectively by focusing on these qualities. Concretely, we pose the following research questions:

1. When is digital evidence considered to be relevant and expressive yet reliable?

2. What is necessary and sufficient evidence in forensic event reconstruction?

Furthermore, we investigated more practical implications of these analyses:

3. How is telling evidence best used in real-world investigations to achieve the criminalistic goals?

4. What hampers expressiveness of digital evidence?

Before addressing these research questions, we revisit and scrutinize the terminology commonly encountered in this field, which will provide the common ground to establish and embed our contributions presented next.

1.4 Contributions

Based on the results stemming from answering the previously presented research questions (Section 1.3), we propose a structured method to identify relevant traces in digital investigations. A simplified version is illustrated in Fig. 1.2 to provide an initial overview. Taken separately, we can summarize the contributions as follows:

- We develop formal notions of the concepts of relevance and expressiveness of (digital) evidence in relation to investigative hypotheses. Building up on these concepts, we are able to define an investigative knowledge base in a precise manner. Besides improving the theoretical understanding of these critical aspects of the overall investigative process and the nature of digital evidence, the concepts enable us to derive evidential reliability properties in a rigorous manner (published as [101]).

- To clearly define relevant evidence on a technical level, i.e., answering the question of what is necessary and sufficient to conclude which past events happened, we establish foundational notions of event reconstructability by introducing different classes of evidence expressed in temporal logic. By doing so, we clearly delineate what is sufficient and necessary to conclude the execution of certain actions given a system model. Moreover, the practical feasability of calculating evidence sets corresponding to these classes has been demonstrated by the use of model checking software (published as [105]). This ties in with the previous chapter and materializes the concepts elaborated, thereby establishing the relevance relation as the basis for determining expressiveness and constructing an investigative knowledge base.

- Given that the concepts of necessary and sufficient evidence are only applicable in idealized settings so far, we additionally propose a practically oriented method to incorporate phenomenon-specific knowledge into investigations in a structured way to improve real-world investigations. This aims at bridging the abstraction gap from general process models to case-specific concretizations by using a meso-level of abstraction in the form of cognitive maps to capture phenomenon-specific knowledge supporting the quest to find relevant traces in a more pragmatic way. We demonstrate the method and its applicability in the example of botnet crime (published as [103]).

- Lastly, we show how expressiveness could be hampered by contamination effects during the acquisition or analysis phases. To do so, we provide a universal definition of evidence contamination—applicable both for physical and digital evidence. Aiming to substantiate and validate the definition, we provide examples, counterexamples as well as edge cases of contamination of digital evidence (published as [104]).

The thoughtful combination of these insights and their integration into the bigger picture of the criminalistic task led to the development of the *Cyber-traceological Model* that provides a structured method of translating investigative demands to the respective relevant traces.

The core component of this model is the *investigative knowledge base*. Since its construction is meticulous, that part and the underlying concepts will be the core of the present thesis.

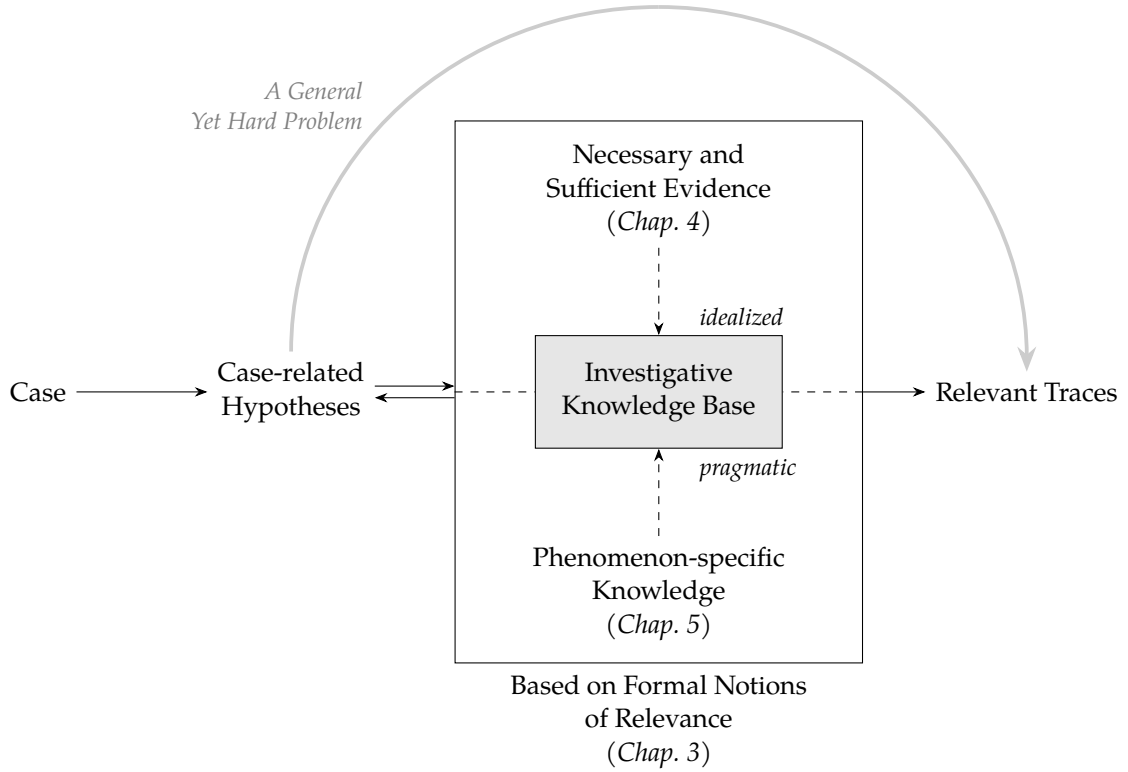

Figure 1.2: Initial overview of the *Cyber-traceological Model*. This simplified version of the model depicts the overall process and the basic building blocks. The figure provides an overview of their interrelations and locates these in their respective chapter.

Given that three of our previously published works were concentrating on various aspects of finding relevant traces, it is apparent that this task is not trivial to solve. In short, we develop both formal and pragmatic ways to encode the understanding of digital traces and their interpretation in an investigative knowledge base. It could be constructed in two different ways: In an idealized setting where a model of the system is available, one could use an automata-theoretic approach resorting to evidential concepts expressed in temporal logic that are calculated using a model checker [105]. For more comprehensive real-world scenarios, the knowledge base could be constructed by mining phenomenon-specific knowledge with expert interviews or document analyses—a process where investigative measures, the resulting digital traces, and their meaning for the case are mapped to one another [103]. This construct provides the investigators with a mechanism to translate investigative demands to the relevant traces that will help fulfill these demands and eventually solve the case. Starting with some suspicion and a set of case-related hypotheses, investigators can then consult the investigative knowledge base to identify and consider (only) those digital traces that are relevant to the predetermined hypotheses, i.e., traces that either refute or support these [101].

As indicated in the previous section, it becomes obvious that—although we call it the *Cyber*-traceological Model and focus on digital evidence—it is not necessarily specific to the digital domain. The need to translate investigative demands to relevant traces seems to be more general. Maybe one could even reach so far and name the task of

finding relevant traces—digital or physical—based on the investigative hypotheses the "holy grail of criminalistics". This again brings up the question of whether the determined confrontation with the cyber-dimension might even contribute to the advance of the general field as well—sparking a discussion on this topic would be an unexpected overachievement of this thesis.

1.5 Publications

Parts of the scientific results presented in this thesis have already been published and, hence, validated by peer reviews. Those articles are listed below:

[103] Jan Gruber, Lena L. Voigt, Zinaida Benenson, and Felix C. Freiling. Foundations of cybercriminalistics: From general process models to case-specific concretizations in cybercrime investigations. *Forensic Sci. Int. Digit. Investig.*, 43 (Supplement):301438, 2022. ISSN 2666-2817. doi: 10.1016/J.FSIDI.2022.301438

[104] Jan Gruber, Christopher J. Hargreaves, and Felix C. Freiling. Contamination of digital evidence: Understanding an underexposed risk. *Forensic Sci. Int. Digit. Investig.*, 44(Supplement):301501, 2023. ISSN 2666-2817. doi: 10.1016/j.fsidi.2023.301501

[105] Jan Gruber, Merlin Humml, Lutz Schröder, and Felix C. Freiling. Formal Verification of Necessary and Sufficient Evidence in Forensic Event Reconstruction. In Edita Bajramovic and Ricardo J. Rodríguez, editors, *Proceedings of the Digital Forensics Research Conference Europe (DFRWS EU)*, pages 1–11, Bonn, 3 2023. dfrws.org

[101] Jan Gruber and Merlin Humml. A Formal Treatment of Expressiveness and Relevance of Digital Evidence. *Digital Threats*, 7 2023. ISSN 2692-1626. doi: 10.1145/3608485

The article "Foundations of cybercriminalistics" [103] concerning the use of phenomenon-specific knowledge for cybercriminalistics was awarded with the *Best Student Paper Award* at the "Digital Forensics Research Conference APAC '22" and also received a special award from the registered society for criminalistics in Germany, "Deutsche Gesellschaft für Kriminalistik e.V.". Furthermore, the article "Contamination of digital evidence" [104] received the *Best Student Paper Award* at the "Digital Forensics Research Conference EU '23". These and the other two articles published at international conferences concerned with the topic of IT forensics, hence, constitute the cornerstones of this work and are used throughout the thesis as follows:

Parts of Chapter 2 and the entirety of Chapter 3 include direct excerpts of the article "A Formal Treatment of Expressiveness and Relevance of Digital Evidence" [101], published in the journal *Digital Threats: Research and Practice* by ACM under the Creative Commons 4.0 license. Chapter 4 is based mainly on the conference contribution "Formal Verification of Necessary and Sufficient Evidence in Forensic Event Reconstruction" [105], which has been published in the proceedings of the *Digital Forensic Research Conference EU '23* as

an open-access publication. Chapter 5 and Chapter 6 contain, to a large degree, direct excerpts of the articles "Foundations of cybercriminalistics" [103] and "Contamination of digital evidence" [104], respectively—both were published in the journal *Forensic Science International: Digital Investigation* by Elsevier under the Creative Commons 4.0 license. For all these referenced works, the authors either retained their copyrights or obtained scholarly communication rights, allowing the use of these articles in a subsequent compilation, such as this dissertation.

Besides these main publications, the author of this thesis contributed to the following publications during his dissertation that were also concerned with various aspects of digital investigations, such as cryptocurrency investigations [62], the acquisition of technical cyber threat intelligence [100, 99] and cybercriminalistic strategies [102, 106], i.e., adressing questions of responsibility of police forces to prevent digital hazards and factors determining successes investigating cybercrime:

[62] Dominic Deuber, Jan Gruber, Merlin Humml, Viktoria Ronge, and Nicole Scheler. Argumentation Schemes for Blockchain Deanonymisation. *FinTech*, 3:236–248, 2024. doi: 10.3390/fintech3020014

[100] Jan Gruber and Felix Freiling. Fighting Evasive Malware. *Datenschutz und Datensicherheit - DuD*, 46(5):284–290, 5 2022. doi: 10.1007/s11623-022-1604-9

[102] Jan Gruber, Dominik Brodowski, and Felix C. Freiling. Die polizeiliche Aufgabe und Pflicht zur digitalen Gefahrenabwehr. *Zeitschrift für das gesamte Sicherheitsrecht (GSZ)*, 5:171–176, 2022. ISSN 2567-3823

[99] Jan Gruber. Identifizierung von Malware-Infrastruktur mittels verteilter Spamtrap-Systeme. In Albrecht Ude, editor, *Sicherheit in vernetzten Systemen: 30. DFN-Konferenz*, pages A1–A27. BoD–Books on Demand, Hamburg, 02 2023. ISBN 3756881393

[106] Jan Gruber, Lena L. Voigt, and Felix C. Freiling. Faktoren erfolgreicher Cybercrime-Ermittlungen. *Kriminalistik*, 77:266–271, 5 2023. ISSN 0023-4699

Those evolved—directly or indirectly—from our research and have rounded off our view of the field and the efforts to gain scientific insights into digital investigations and cybercriminalistics. At a more abstract level, all these articles address efforts to make certain aspects of digital investigations more effective, as do the core publications on which we base this thesis, whose outline is presented next.

1.6 Statement and Outline

To conclude this introduction, the present dissertation focuses

(a) on improving the understanding of foundational aspects of digital investigations and
(b) on the effective as well faultless use of digital evidence

using a model-based approach in the context of forensic computing as a research discipline on the one hand and cybercriminalistics as an applied craft of practical importance on the other.

Approaching these two topics, the present dissertation is divided into two distinct yet connected parts: The first part, up to and including Chapter 4, establishes a theoretical basis by providing formal notions of relevance and expressiveness as well as reconstructability classes, thus dealing with a core topic of forensic computing. The second part deals with the practical application and the potential implications by looking at the effective use of digital evidence from a pragmatic viewpoint, which is more related to digital forensics and cybercriminalistics.

Structure of this work. At a more fine-grained level the monograph is structured in seven chapters. After having introduced and motivated the importance of the topic in the present introduction (Chapter 1), we convey essential terms, conceptions, and core topics of the field of digital investigations in Chapter 2 by revisiting, scrutinizing, and relating those to one another, where necessary and helpful. Since the remaining chapters deal with substantially specialized topics, respective related work and background knowledge, as necessary, will be provided within the chapters in separate sections.

Chapter 3 deals with the question when a finding is relevant and what aspects determine its relevance in regard to the investigation to establish formal notions of those trace qualities that can be used to construct an investigative knowledge base for actual use. In Chapter 4, we concisely define what constitutes necessary and sufficient evidence by resorting to temporal logic. This is the basis to define what constitutes relevant traces (in the sense of the previous chapter) for idealized scenarios, where a model of the digital system under investigation is available. In Chapter 5, we leave the abstract and formal realm to show how expressive evidence can be identified and chained together to construct an investigative knowledge base of practical use. Here, we provide a method to navigate through the chaos of real-world investigations and achieve overarching investigative goals based on the inclusion of phenomenon-specific knowledge in investigative knowledge bases—yet in accordance with the formal notions developed earlier. After grasping these concepts, we investigate how expressiveness of digital traces could be hampered by contamination effects in Chapter 6. Lastly, Chapter 7 provides a summary of the findings and integrates the previously developed building blocks in a detailed version of the *Cyber-traceological Model* that has previously been introduced in a simplified rendition. Then, we raise open questions and give an outlook on further potentials for future work.

Ways of reading this work. Fig. 1.3 illustrates the structure of the present thesis visually and depicts possible reading paths of this work depending on the reader's interests. Chapter 2 refreshes overall background knowledge and introduces important related concepts to put those in connection to each other; hence, it is suggested to read up to this point. The following chapters can then be read independently to a large extend. However, it is advised to approach those chapters that are assigned to the theoretical or practical plane respectively in order of appearance, e.g., Chapter 4 takes up concepts developed

in Chapter 3. To capture the chapters' essence, short synopses are provided at the end of each chapter for the hasty reader.

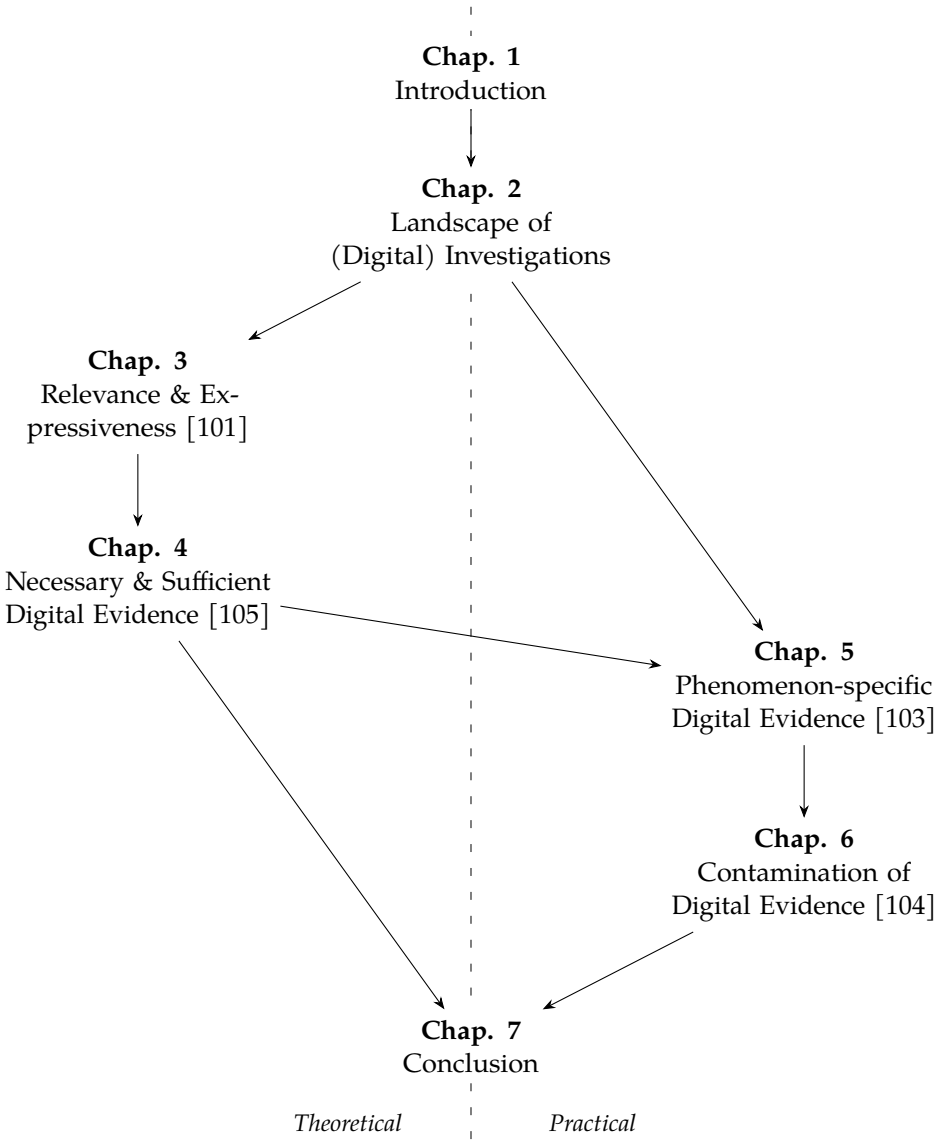

Figure 1.3: Structure of the thesis and possible ways of reading the work as indicated by the arrow-head edges; numerical references in brackets denote the peer-reviewed articles on which the respective chapter is based. Chapters 3 and 4 are primarily concerned with theoretical considerations to build a formal foundation of the concepts, whereas Chapters 5 and 6 consider practical aspects of the endeavor of effective use of digital traces in real-world investigations. Chapter 1, Chapter 2, and Chapter 7 constitute an integrative frame around the main results and have both theoretical and practical parts.

2 The Landscape of (Digital) Investigations

2.1 Introduction

Events, traces, hypotheses, evidence—Say, what are we talking about? Depending on who is your dialogue partner, there might be slight differences in the assumed meaning of these central terms of criminal investigations. Criminalistics is itself a rather broad occupation, as Kirk discovered already in 1963, when he stated that it is one "that has all of the responsibility of medicine, the intricacy of the law, and the universality of science" [140, p. 238]. Given that a diverse set of professions, i.e., jurists, law enforcement officers, and scientists, work in this field, the translation of criminalistic hypotheses to legal demands, then scientific assessments to gain technical results and vice versa can be frictional. This is because the previously mentioned practitioners of those different professions have slightly different viewpoints: Jurists are primarily concerned with the delicacies of procedural and material law, police officers consider operative issues and forensic scientists fixate on the development of a deeper understanding. Strongly abbreviated, their foci could be polarizingly summarized as follows: lawyers and prosecutors respectively aim at winning the case, police officers aim at the general reconstruction of the deed as well as catching the perpetrator, and scientists aim at providing expert witness of excellent quality. Though having these different foci, on a deeper level all of them aim to clarify the situation to uncover the truth by investigating the past events. Hence, we now develop a commonground of terminology geared toward investigative methodology. In our opinion, such a restriction to the investigative perspective is sensible because investigations are at their core, i.e., on a meta-level, detached from legal frameworks specifying procedural requirements or strict testing regimes of scientific assessments.

2.1.1 Contribution of the Chapter

In this chapter, we paint a colorful backdrop of this thesis by conveying broader background knowledge while we aim at a terminological integration. Furthermore, this chapter is considered to function as a playfield for stimulating and elaborating conceptual thoughts connected to cybercriminalistics and its intersection with forensic science. Given that there is a wealth of vocabulary to describe the disciplines, components of investigative thinking, and various aspects of (digital) evidence, we sharpen the terminology and their interrelationship. We infer several interrelationships and contribute a visual representation of the terms, as later shown in Figs. 2.1 and 3.1. We employ an investigative perspective and acknowledge the various uses of digital evidence in different settings; hence, we recourse to general notions of the terms and do not regard strictly legal considerations,

such as the specific interpretation of the US Rules of Evidence [164]. We view this helpful, even necessary, to develop and later relate the notions of relevance and expressiveness as well as necessity and sufficiency as we will introduce those in the further course of this thesis. Besides that, we focus on the theory of investigations. To do so, we recapitulate the scientific method as well as the Criminalistic Cycle to solve the criminalistic task and discuss the structured generation of hypotheses. Here, we transfer the thoughts connected to the so-called hierarchy of propositions by Cook et al. [51] to digital investigations.

2.1.2 Chapter Outline

The chapter starts by looking at the various disciplines concerned with collecting, examining, and interpreting physical and digital evidence to discern those critically by definitional efforts (Section 2.2). Having established a concise distinction (for the use in this thesis), we provide a brief historical retrospection on digital investigations supplemented by a collection of observations and thoughts (Section 2.3). After that, we focus on the investigative core and the central paradigms of investigative case work (Section 2.4). Here, we take up rather recent developments in the field such as the need to focus on the trace as postulated by Margot [158] and the formalized trace model by Jaquet-Chiffelle and Casey [131]. Then, we provide a clarification of important terms related to the description of features that are indispensable to form (digital) evidence, which are visually presented in a conceptual network. Lastly, we focus on the role and the structured generation of hypotheses in digital investigations (Section 2.5) before we summarize the main findings of this chapter (Section 2.6).

2.2 The Relation of the Disciplines

"Criminal science" (in German "Kriminalwissenschaften") is an umbrella term that is, at least in German-speaking countries, commonly used to refer to several branches of science concerned with crime in the broadest sense. On the one hand, it contains the jurisprudence concerned with criminal material law and criminal procedural law as well as related legal disciplines, e.g., police law. On the other hand, several non-legal research disciplines can be assembled under this umbrella term. Those are criminology, criminalistics, and forensic science, i.e., another collective term to refer to the application of various scientific disciplines in forensic contexts. Criminology is primarily concerned with the causes, phenomena and theories of crime thereby employing a sociological perspective. In contrast, criminalistics is targeted at the means and methods of combatting criminal offenses [137, pp. 57 ff.]. The jurisprudence as well as the sociologically focused branch named criminology are effortlessly distinguishable from both one another and criminalistics. The relation of criminalistics and forensic science, however, is not so explicit and maybe even debatable. In any case, it is clear that this has changed over time and is still being negotiated in the scientific discourse [194]. Given the investigative perspective employed in this dissertation, we now take up this discussion and approach the relations.

2.2.1 Forensic Science vs. Criminalistics

The above section has already indicated that forensic science is a broad field, which encompasses many different areas of study. Generally speaking, it is the application of scientific methods and techniques to the investigation of crimes in the context of the justice system. It is important to underline that forensic scientists strive to extract probative facts from evidence, i.e., physical matter in the analog world or bit patterns in the digital domain [21, p. 92], using scientifically sound methods. At a very foundational level, the extraction of probative facts is always based on trace creation, which in turn is based on Locard's *exchange principle* [152], which states that there is inevitably a transfer of matter, which has been later generalized to a *transfer of traits* [126], each time when two entities interact.

Given the vast amount of types and categories of physical traces, it is not surprising that forensic science is comprised of many different (sub)disciplines: psychiatry, pathology, odontology, entomology, chemistry, engineering, computing, and other fields. Given the complexity of each discipline, it is common for forensic scientists to specialize in a particular area of study, such as ballistics, the study of corpse colonizing insects, or DNA analysis. They either work in forensic laboratories, sometimes as expert witnesses, or are associated with law enforcement agencies. This raises the question whether they are just scientists working in this specific context or whether they have certain unique characteristics.

While scientists are interested in gaining new insights to build up and evolve fundamental knowledge, *forensic* scientists take the generally accepted knowledge in the respective field and utilize it to "assess the value of observations as evidence of an individual past event from evidence or data that has actually occurred and is possibly non-replicable" [195, p. 168]. Hence, in addition to their technical expertise, forensic scientists of all disciplines are distinguished by their specialization in forensic subtleties. Those subleties are related to the special tasks of the collection of traces at crime scenes, their analysis, the interpretation of the results, and the communication of the findings related to the case-related hypotheses [195, p. 166]. In essence, these statements are summarized well by the unifying definition recently provided by Ristenbatt III et al. [194]:

Definition 2.2.1 (Forensic science [194, p. 29])**.** *Forensic Science* is the "application of scientific principles and techniques relating to the examination and analysis of physical data/findings in civil and criminal justice matters."

After having grasped forensic science, we now need to relate criminalistics to forensic science. The term "criminalistics" probably goes back to the Austrian legal practitioner Hans Gross, who used the German word *Kriminalistik* [125, p. 10] to describe a systematized form of investigative work by consulting scientific methods as one of the first in the field [96, pp. 399 f.], which is why he is considered a pioneer if not the founding father of this discipline. In strong demarcation of purely legal issues, he considered the discipline to be concerned with combatting crime by law enforcement agencies in the reality of life. In the original wording, he used the phrase "Lehre von der [...] Bekämpfung der Kriminalität durch die Strafverfolgungsorgane [...] in der Lebenswirklichkeit" [98, p. 5]. His view

encompassed four areas: techniques of crime commission, forensics (describing the collection and examination of physical evidence using scientific methods), criminalistic tactics, and the organization of crime fighting [98, pp. 7 ff.]. In the original wording the terms used were "Verbrechenstechnik", "Kriminaltechnik", "Kriminaltaktik", and "Organisation der Verbrechensbekämpfung" [98, p. 7 f.], which described a largely practically oriented discipline [1, p. 31].

Interestingly, there has been a notable divergence between the German and US-American notion of the term "criminalistics". The influential scientist Paul L. Kirk, a professor of biochemistry, who was considered an identification expert in criminal investigations, proposed to use the term "criminalist" for referring to an "examiner of physical evidence" in criminal investigations [138, p. 167]:

> Unless we are to proceed through life labelled as examiners of this and examiners of that, or technical experts in the field of thus and so, it would seem desirable for some standardized nomenclature to be adopted. In the absence of better terms than "criminology" and "criminalistics" in their European sense, and of "criminologist" and "criminalist" to designate the individual practicing in these fields, the tentative suggestion is offered that they be regularly adopted in this country in the sense indicated.

By doing so, he dropped the phenomenological, tactical as well as organizational aspects of combatting crime and put a stronger emphasis on the scientific aspects of crime investigation, which leads actually to a closer alignment with the term "forensic scientist". Several decades later, Kirk's conception has been still established: The California Association of Criminalists, for instance, defined criminalistics as the "profession and scientific discipline directed to the recognition, identification, individualization, and evaluation of physical evidence by the application of the natural sciences to matters of the law" [55, p. 202].[1] This corresponds to a change in the scope of expertise: On the one hand, forensic scientists are considered to be specialists within a rather confined area of knowledge. On the other hand, criminalists are considered to be generalists in the whole spectrum of the scientific evaluation of physical traces [194, p. 29]. In the meantime, however, Ristenbatt III et al. noted a "conflation of the terms *forensic science* and *criminalistics*", which becomes especially apparent when referring to the definitions provided by *The National Institute of Standards and Technology Organization of Scientific Area Committees* (OSAC) lexicon [194, p. 28 ff.]. Thus, Ristenbatt III et al. proposed to use the term "traceology" as a more precise and not biased version of "criminalistics". This term goes back to the 1920s, when early criminalists at the Humboldt-Universität Berlin coined the term "traceology"[2] [156, p. 97] to make obvious that the trace and its study aiming to reconstruct past events is at the center of the interest—a conceptual development originally sparked by Margot [156, 157]. They underlined that they envision a traceologist—their proposed replacement for the term "criminalist"—to be a generalist with a broader scientific knowledge here, who is concerned with the "study of event traces created during an event" [194, p. 29] in a holistic way. They defined the term "traceology" as follows:

[1] This has originally been elaborated by De Forest et al. [56, p. 4].
[2] The original term in German is "Spurenkunde".

Definition 2.2.2 (Traceology [194, p. 29 f.])**.** *Traceology* is the "[s]tudy of event traces created during an event, which encompasses the detection, recognition, identification, process of individualization toward source attribution, and evaluation of the physical record created (be it an item, pattern, or signal)."

This differentiation from the broader field of forensic science and the placement of a strong focus on the trace, as already suggested by Margot [156, 158], seems to be sensible and helpful for clarifying the terminology, albeit it does not introduce a new concept. Interestingly, such a conflation is inexistent in the German speaking area. Going back to the notion of Hans Gross, one could argue that the terminological focus on the trace, as proposed by Ristenbatt III et al., leads to a conceptual loss because other elements, such as the phenomenological aspects of crime commission, criminalistic tactics, organizational and strategic concerns, of Gross' initial notion are excluded by this narrow definition. Hence, we propose to use the term "traceology" only for the object area that is concerned with the collection and examination of pieces of evidence using scientific methods, what Gross named "Kriminaltechnik". By doing so, the originally broad orientation of the term "criminalistics" remains but the other rather interdisciplinary aspects as described above are *not* discarded. Criminalistic tactics describing the case-by-case oriented procedures to solve crimes and other operative concerns, for example, are increasingly important for combatting organized crime and crime of terror, as observed by Brodowski [26, pp. 488 f.]. In addition, phenomenological knowledge is an absolute requirement to combat crime, and organizational as well as strategic factors are equally crucial components of criminalistic work influencing its success.

In view of the (re-)discovered focus on traceology, we therefore propose to adjust the original definition of the term "criminalistics" by Gross [98] as follows:

Definition 2.2.3 (Criminalistics)**.** *Criminalistics* is the profession and scientific discipline of combatting crime. It is comprised of the study of the organization and strategy of crime fighting, the study of criminal phenomena, and their factual investigation using traceology and criminalistic tactics.

On the one hand, this definition follows the long-established definition going back to the original notion of Hans Gross [98] to a large degree and is conformant with more contemporary definitions [29, 58]; on the other hand, it takes up recent (but also historically rooted) developments of studying traces. So, we argue that criminalistics is more than a subdiscipline of forensic science for two reasons: First, it aggregates other subdisciplines of forensic science by drawing on the knowledge, methods, and techniques developed by those. By doing so, the practitioners of criminalistics apply those in a more holistic manner at the forefront of actual investigations and not just in a lab. Secondly, it contains aspects that are not concerned by other forensic sciences, such as the above mentioned criminalistic tactics and operative aspects.

2.2.2 Digital Forensic Science vs. Cybercriminalistics

After having clarified the relation between forensic science and criminalistics in the previous section, we now shift our focus to the digital dimension. Unsurprisingly, the terminology in the cyber-dimension of the field can also appear to be conflated and unclear in some ways. "Digital forensic science", "forensic computing", "digital investigations", and "cybercriminalistics" are often used interchangeably but a closer look reveals notable differences in their respective meaning, the scope of application, and their connotation, which will be clarified in the following paragraphs.

Early on, a definition for digital forensic science has been proposed by Palmer [180] in 2001:

Definition 2.2.4 (Digital forensic science [180, p. 16])**.** *Digital forensic science* is "[t]he use of scientifically derived and proven methods toward the preservation, collection, validation, identification, analysis, interpretation, documentation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events found to be criminal, or helping to anticipate unauthorized actions shown to be disruptive to planned operations."

This proposal goes back to a research workshop with discussions of attendees from academia, military, law enforcement, and the private sector. Notable are two aspects: First, the term is not strictly confined to the courtroom—something to be terminologically expected—and, second, the definition includes also a preemptive dimension. Even 20 years of rapid technological advancement later, this definition still seems to be largely unchallenged and applicable. The main difference to other branches of traditional forensic science is the restriction of the domain of evidence in question: *Digital* forensic science is restricted to explore *digital* evidence [21, p. 91] stemming from *digital* sources.

This restriction enables us to further separate subfields of the branch of digital forensic science based on the respective subject, i.e., the specific type of digital evidence, that is analyzed as it has been proposed by Böhme et al. [21], who identified computer forensics and multimedia forensics as different subfields. They argued that multimedia forensics has to be considered an "independent" subfield [21] because it deals with image and video data generated by some capturing device, i.e., a light sensing chip or a digital microphone. In view of recent advances, however, this discrimation becomes blurry when incorporating the development and use of generative machine learning models (e.g., in the case of assessing deep fakes). Nevertheless it seems sensible to seperate subfields. Here, we see three different possibilities to achieve such a differentiation of digital forensic science: Those are the methods employed, the type of item (e.g., smartphone, desktop computer, Internet of Things (IoT) devices, etc.), or the type of traces encountered (e.g., network traces, multimedia data, main memory, etc.). Depending on the context, i.e., organizational aspects in law enforcement divisions, academic research, expert witness specialization, different subfield differentiations might be more apt than others. However, while this definition of digital forensic science and its division into subfields appears to be solid,

there are several other related terms that do not fit in this classification and need to be clarified as well.

One of those is the term "forensic computing" that began to circulate in the forensics community early on as well. Opposingly to "computer forensics", this term actually avoids the "syntactical mess that uses the noun *computer* as an adjective and the adjective *forensics* as a noun" [36, p. 31] and can hence be considered to be more apt and precise. As one of the first researchers, McKemmish [160] referred to the term "forensic computing" to describe the activities concerned with digital evidence and information technology in the forensic context. He identified four key topics with which the field is concerned in his perception: Those were the identification of digital evidence, its preservation, its analysis, and its presentation. Comparing this to the definition of digital forensic science, one clearly sees a significant overlap.

Tackling this overlap, Dewald and Freiling [65] scrutinized the various questions, tasks, and contents that are subject to the broad (and fuzzy) sphere of computer forensics what has been considered to be a synonym of digital forensics (DF) by them [65, p. 5]. In contrast to the subject-focused differentiation into subfields, they employed a task- or focus-oriented division: Based on their encounters, they proposed a distinct separation of the discipline "into a rigorous scientific part on the one hand, and a more general methodology of searching and seizing digital evidence and conducting digital investigations on the other" [65]. They argued that "forensic computing", the term they coined for the scientific part, is at its very core—just as other branches of traditional forensic science—concerned with associations using the fundamental notions of transfer (of traits) in the digital realm. In essence, Dewald and Freiling [65] proposed the following convincing definition:

Definition 2.2.5 (Forensic computing [65, p. 9])**.** *Forensic computing* is the application of computer science and its methods to deal with the search for and establishment of associations of digital traces.

Their proposal leads to considering forensic computing to be strictly focused on the scientific interpretation of digital evidence. They elaborated their definition by discussing several (investigative) statements that are typically the subject of the work of expert witnesses, who are tasked to prove or disprove these statements. Given this focus, we grasp digital forensic science to be a true branch of forensic science since it is grounded on the same universal principles. Opposingly to the rigorous scientific part, Dewald and Freiling considered "[g]uiding an investigation and formulating such questions in the process of an investigation is the 'digital investigation' part of computer forensics" [65, p. 9]. Hence, they used the term "digital investigations" to refer to the part that is concerned with the entire process of handling and processing digital evidence for forensic purposes, including but not limited to recovery as well as inspection of and search for evidence and its management [65, p. 6]. This has a notable overlap with the definition provided by Hargreaves [111], who grasps digital investigations as "a process that formulates and tests hypotheses using digital evidence, where the results could be admitted to a court of law" [111, p. 14]. Given that the evidential value depends on all upstream or accompanying procedural aspects as well—and not only the correctness of scientifically rigorous associations—, we

argue that those activities of digital investigation should also be considered to be part of the wider field of digital forensic science. In our opinion, these activities have also to be grounded in sound and scientifically validated procedures, even when they can be considered to be "preparatory" or even "mechanical" tasks, which are (primarily) used by investigators in the field.

In view of the comprehensive interpretation of the term "digital forensic science", there arises a need to demarcate it from "cybercriminalistics"—two terms that are unfortunately often used interchangeably, although there are arguably big differences in their meanings. Reminiscing the definitions provided above as well as the discussion of the separation of traditional forensic science and criminalistics, we recognize an expansion of the scope of the general terms by the incorporation of the digital domain. It becomes apparent that the newly emerged digital dimension of several aspects of Definition 2.2.3, which is closely aligned with the view of Gross and Geerds [98], can be grouped and subsumed by the term "cybercriminalistics". Hence, we propose to define the field of cybercriminalistics as follow:

Definition 2.2.6 (Cybercriminalistics)**.** *Cybercriminalistics* is the digital dimension of the profession and scientific discipline of combatting crime. It is comprised of the study of the organization and strategy of cybercrime fighting and specifically digital investigations, the study of cybercriminal phenomena and their investigation using traceology of digital traces and digital criminalistic tactics.

Contextualisation of the Definition of Cybercriminalistics. One building block of cybercriminalistics is the study of cybercriminal phenomena both in the narrower sense, i.e., *core cybercrime*, and in the wider sense, i.e., *cyber-enabled crime*, which includes every criminal offense where a computing device has been used as means of action [119, pp. 22f.]. Another building block are digital investigations, those are—much like in the sense of Dewald and Freiling [65, p. 9]—all the procedural aspects of working with digital evidence to solve investigative questions. To underline the holistic approach that is needed, we refer to the term traceology (Definition 2.2.2). This appears to be sensible because digital traces are often linked to events in the physical world in one way or another, which requires such a holistic view. At this point, it has to be noted that traceological activities related to digital evidence are included in the notions of Dewald and Freiling [65] and Hargreaves [111] of digital investigations. The building block of digital criminalistic tactics is concerned with operative aspects of fighting cybercriminal offenses and conducting digital investigative measures. Think of delicate questions related to the execution of IP address tracking, stingray operations, wiretaps at the source, realization of live forensics of a server system, and much more. Besides these technically complex investigative measures, tactical considerations are also related to the holistic practice of collecting and analyzing the wealth of digital evidence, such as triage, prioritization, and examination strategies. Lastly, the study of cybercriminalistics should also include the specifically digital aspects of the organization and strategy of crime fighting, i.e., the execution of specifically digital investigations and matters concerned with cybercrime offenses. These matters include responsibility issues, resource allocation, questions of efficient cooperation, legal regulation, and so on.

Delimitation and Demarcation of the Definition of Cybercriminalistics. Definition 2.2.6 is inclusive and far-reaching; however, it is not concerned with the general digitization of investigations. It is important to underline that we do not consider the digitization of investigative processes to be included. An example for this is the use of digital devices when executing classical examinations as in the previously mentioned digitized crime scene analysis, i.e., the collection, processing, and examination of traces in a digitized form [163] ranging from digitized latent fingerprint scanning to virtual reality reproductions of crime scenes. Also not included is the increasingly digitized reporting and case-processing or the use of big data platforms for linking traditional crime information and forensic intelligence. Despite this demarcation, the definition of cybercriminalistics covers a wide range of criminalistic work. However, this wide coverage can be seen analogously to the commonly employed differentation of core cybercrime and cybercrime in the wider sense, that is equally far-reaching. So, in this day and age, investigators will be more often than not confronted with offenses that have (at least) some digital component to them. Given the pervasiveness of computing, i.e., the spread of digital devices of all sorts in smart home, automotive, or manufacturing contexts, we boldly forecast that cases that do not exhibit a cybercriminalistics dimension will slowly diminish. In view of the rapidly progressing digitization, we, therefore, anticipate a comparably rapid increase in the importance of cybercriminalistics—but this will not affect the significance of traditional criminalistics.

Collaboration between digital forensic scientist and criminalistic investigators. According to Dewald and Freiling "[i]nvestigators investigate, scientists associate" [65, p. 9] and, actually, it must be pointed out that the distinction between digital forensic science and digital investigations can be observed in real-world cases which has been reflected by the research community. Much like with traditional forensic science, where "[b]oth scientists and investigators have important synergistic roles to play at crime scenes" [55, p. 201], digital evidence has to be made accessible to case investigators since the advantages stemming from case knowledge and tactical skills "far outweigh the risks" [18, pp. 107 f.]. This is because case investigators know the case, hence, it is most sensible when DF experts are consulted to support the assessment of delicate questions. Van Beek et al. [232] identified that such a "collaboration between detectives and digital investigators is crucial in understanding digital evidence" [232, p. 22]. Therefore, we propose—comparably to classical crime scenes, where "[t]he scientist and the scene investigator" had to "learn to appreciate the other's strengths, establish a partnership, and work together" [55, p. 200]—that the holistically oriented and educated forensic science practitioners called (cyber-)traceologists and investigators work closely together and consult deepened expert knowledge where needed to provide investigative results of the best possible quality.

To sum up, this section aimed to provide terminological clarity in regard to the various designators of the different disciplines and their relation. It becomes clear that despite providing delineated definitions, the borders are not completely clear-cut, and several intersections exist. These intersections and the relations as put up in the discussed or developed definitions are visualized in Fig. 2.1. In addition, it must be stressed that the interrelations and the intermeshing of activities lead to the necessity of keeping an overview on the one hand and a balanced collaboration between the involved personnel of the different disciplines on the other hand.

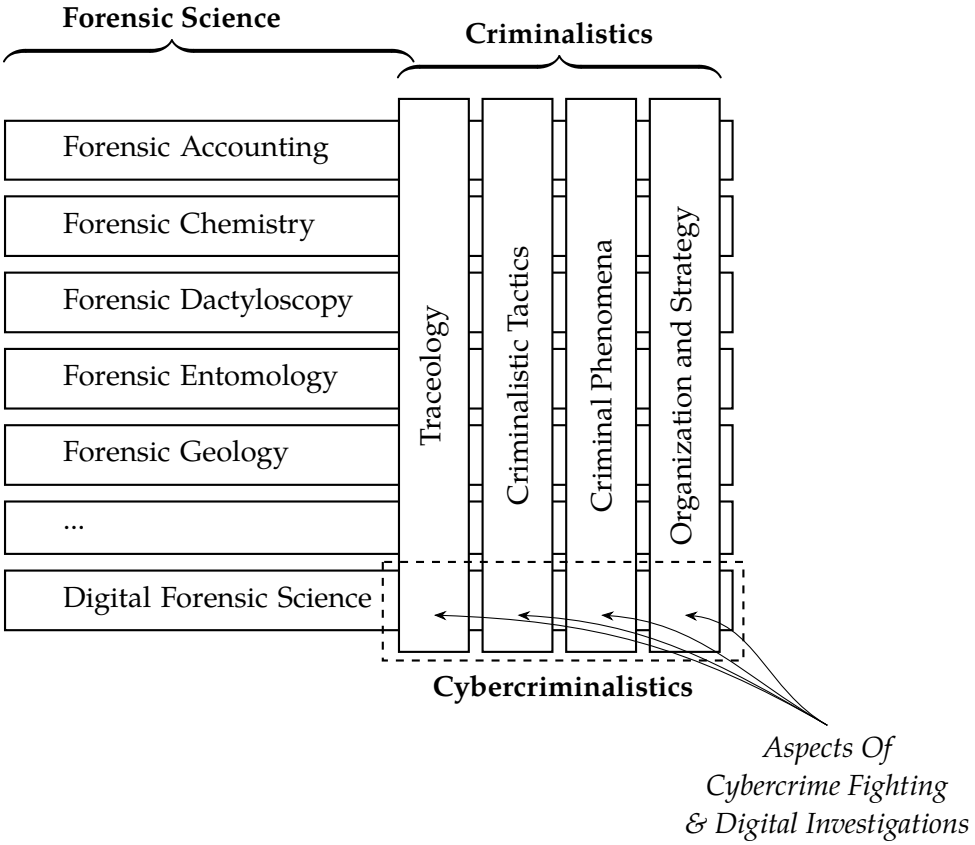

Figure 2.1: Proposed view of the relation of the intertwined disciplines of forensic science, criminalistics, and cybercriminalistics. Referring to Definitions 2.2.1 to 2.2.4 and 2.2.6, we consider grouping the single branches of forensic science as traceology, i.e., the holistic study of traces. Traceology is considered part of criminalistics, which additionally combines tactics, phenomenology, organization, and strategy. Cybercriminalistics then comprises the digital dimension of these fields and is primarily concerned with digital traces, digital investigations, and combatting cybercrime.

2.3 A Brief Retrospection of the Scientific Discourse of Cybercriminalistics

Having grasped the intersections and differences of digital forensic science and cybercriminalistics, we now provide a brief retrospection of the scientific discourse concerning cybercriminalistics. Since the early endeavors of Hans Gross the science of solving crimes evolved and progressed by the utilization of various scientific achievements, which had an impact on the tactics and procedures of this discipline from a technical point of view. One important change, however, is that information technology is now not only increasingly used as a tool for fighting crime, illustrated by the trend of digitized crime scene analysis mentioned in Section 1.2, but has increasingly gained importance as a source of traces and even as an instrument of offense—an impactful development that sparked research in various fields.

2.3.1 Initial Approaches

Technical analyses and phenomenological studies of (components of) cybercrime offenses have thus been published partly even at top tier security conferences [e.g., 5, 214, 223] and the perpetrators' use of information technology as means of crime is also reflected in the scientific discussion of various procedural models for the digital investigative process. They were compared in detail by Pollitt [185] and later Selamat et al. [204]. A notable contribution was the comparison and integration of the physical crime scene procedures into the digital investigative process by Carrier and Spafford [33]. Still all those important works are concerned with the overall procedures of a digital forensic scientist and the coordination of the various steps of a forensic examination, so they neglect many operational aspects of investigations related to cybercriminalistics.

Actually, the term "cybercriminalistics", as we defined it in the previous section, has not been a dedicated subject in its entirety in scientific literature—an initial foray has been our own publication [103]; however, there are works conceptually concerned with the investigation of cybercrime offenses. A rare instance of a research article dealing both with concrete technical and investigative questions has been published by de Graaf et al. [57], who elaborated a forensic investigation model intertwined with operative investigative measures that has been applied to a real-world botnet case [57, p. 312]. From an abstract point of view, Hunton critically analyzed cybercrime phenomena and their investigation by law enforcement to kick off a more substantial discourse of the topic [121, 122, 123].

2.3.2 Hunton's Cybercrime Investigation Framework

With a focus on cybercrime investigations, Hunton carefully reviewed iterative process models and found that while some of them "predominately focus [...] on digital evidence recovery", others "consider the rigorous and scientific processes of digital forensic examination", but all of them lack the requirement of a broader understanding of criminal

investigations [122, p. 386]. Hunton criticized further that those models are either too abstract or restricted to the examination of digital evidence and forensic recovery, which led to his proposal of the *Cybercrime Investigation Framework*. This framework extends the *Core Investigative Doctrine*, as proposed by the Association of Chief Police Officers [6, p. 48], by identifying eight stages, which can be typically found in every cyber investigation. Those are supported by the so-called *Cybercrime Execution Stack* [122, pp. 389 f.], which we now recapitulate.

2.3.2.1 The Cybercrime Execution Stack

The conceptual model common to all cybercrime offenses was proposed by Hunton [121]. It contains the seven components of the execution of a cybercrime offense illustrated in Fig. 2.2. According to Hunton, those need to be considered alongside their respective characteristics in the context of any cybercrime investigation.

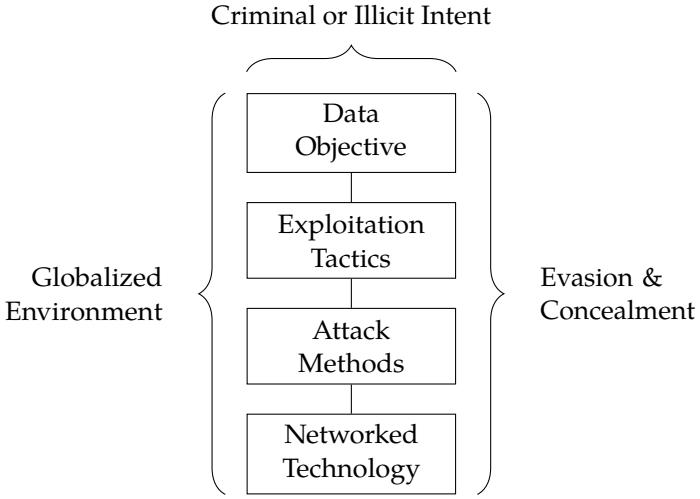

Figure 2.2: Cybercrime Execution Stack [121, pp. 531 ff.].

The starting point is always the criminal or illicit intent of the perpetrator, which leads to the *data objectives*. Those are—at least to a certain extent—necessary intermediate goals and describe the collection, the exchange, the transmission, or the use of data by the perpetrator [121, pp. 531 f.]. The flow of data, respectively its usage, constructs the connection between the perpetrators' intent and the *exploitation tactics*. On the one hand, those can be technical attack vectors, social engineering, or illegal agreements [121, p. 532]. The *attack methods*, on the other hand, represent specific expressions of the *exploitation tactics*. Hunton [121] explained that trojans, worms, and viruses form technical attack vectors, while social engineering can be executed by the use of email, blogs, social networks, and the like. This modeling aims at the deconstruction of the deed into specific technical execution methods, which could be employed by the perpetrator. The observation of the *networked technology* that was used for the respective attack method ultimately leads to the relevant trace material. The perpetrators need some sort of access device to connect to the network. Then, they interact with various network resources via well-defined communication channels, which are enabled by infrastructure components like access

points, routers, DNS-servers, firewalls, and other devices. Modeling those interactions on the network layer should allow the investigator to identify relevant trace material in a structured manner [121, p. 533]. In addition, the *globalized environment* is a distinguishing feature of cybercrime investigations and forms a separate dimension on its own. Ultimately, the possibilities of anonymization—called *evasion and concealment*—represent a substantial opportunity for cybercriminals, which is the reason why their potential and the associated tactics and technologies must always be taken into account when forming hypotheses regarding the perpetration and the sequence of actions [121, p. 533].

According to the author, the Cybercrime Execution Stack aims at supporting detectives and prosecutors during the planning of cybercrime proceedings that exhibit high complexity. This is accomplished because the proposed thinking model enables the investigators to take all relevant technological aspects into account and, thus, simplifies the identification of needed expert knowlege as well as other capabilities required for conducting a successful investigation.

2.3.2.2 The Stages of Cybercrime Investigations

The "logical model of the distinct elements consistently found" when investigating cybercrime offenses deemed Cybercrime Execution Stack [122, p. 389], which has been explained above, seeks to guide the key stages relevant and specific to cybercrime investigations as they are now explained and additionally illustrated in Fig. 2.3 [123, p. 63].

Regardless of the initiation method of the investigation, the investigators need to outline the technical and logical elements involved in the cyber-related offense in question to gain an initial understanding of the case [123, pp. 63 f.]. Hunton [123] suggests relying on the Cybercrime Execution Stack to accomplish this and model the case. The next stage is the structured assessment as a formal review of the current state of the investigation and its elements modeled before, which builds the foundation for the formulation of hypotheses [123, p. 64]. Considering the impact and risk in the next stage of the model should assist the choice and prioritization of investigative actions, which are then identified and planned. The "tools"-stage should ensure that the identified actions are performed with the "most valid and appropriate" tooling to guarantee the reliability of the evidence acquisition and its examination [123, p. 65]. Those tools are used to actually conduct the actions in compliance with technical and judicial requirements, while resorting to appropriate low-level examination models, if needed. Finally, an evaluation and contextualization of the outcome have to be conducted, either to integrate it into a broader law enforcement process or to begin the next iteration of the presented stages [123, pp. 65 f.].

Although this model aims to the consideration of specific issues of cybercrime and tries to be a blueprint schedule for cybercrime investigations, it still appears rather general and cyberphenomenon-agnostic.

Therefore, the fundamental nature of Hunton is now further abstracted to trace the essence of investigative thinking, which aim to prepare for Chapters 3 and 4.

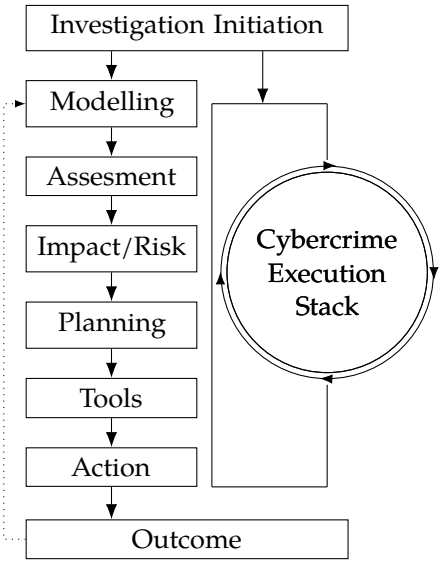

Figure 2.3: Stages of cybercrime investigation [123, p. 63].

2.4 The Investigative Core of Criminalistics

We now revisit various areas and elements of investigative reasoning to fuse them, which should serve as a starting point for the presentation of the obtained results of this thesis.

2.4.1 The Paradigm of Forensic Science

To begin with, we refer to Inman and Rudin [126] and their discussion of the high-level paradigm of practicioning forensic science. According to their compilation of previous works and their continuation, solving cases consists of five building blocks [126, pp. 1, 15 f.]:

1. Transfer
2. Identification
3. Individualization
4. Association
5. Reconstruction

The whole practice of forensic science is based on Locard's *exchange principle*. This long established postulate describes "the dynamics of an offense cause traces to be exchanged between the offender(s), victim(s) and surrounding environments (physical and virtual)" [131, p. 2] causing a *transfer* of traits.

At first, the forensic science practitioner has to identify evidence by placing items in a class or group (*identification*) and hence determine what it is. In order to narrow down, an *individualization* step has to be performed as originally described by Kirk [140], which

means that a common source or origin is inferred for the items in question. Afterward, an *association* between two items could be established. Inman and Rudin [126] define association "as an inference of contact between the 'source' of the evidence and a 'target' "[126, p. 15]. In this step, it is assessed whether the source and target item were actually in contact. Finally, it is attempted to *reconstruct* past events. In the notion of Inman and Rudin [126], this step is "the ordering of associations in space and time" [126, p. 16]. So, by building up on this universal concept of associations, spatially and temporally ordered links are established that can be used to deduce actions of persons and ultimately reconstruct the deed.

Figure 2.4: A section of the universal paradigm of forensic science by Inman and Rudin [126, Fig. 4]

In essence, forensic science is backward-oriented [15], also called retrodictive [198]; therefore, practitioners put traces as cues of past events in the absolute center of their reasoning [158, p. 30] since they have to be considered "fundamental vectors of information" produced by some activity or presence by principle [198, p. 3].

2.4.2 From Traces to Facets

Given the traces' importance and their delicacies, it is not surprising that the concept of traces has been dissected further. In view of the increasing digitization, previously proposed definitions and concepts were revisited and sharpened by the so-called formalized trace model proposed by Jaquet-Chiffelle and Casey [131] in 2021. The intuition of this model is that a trace denotes a difference Δ at a certain point in time t between a scene $\mathbb{S}_{R,\mathbb{W}_E}$ bounded by a region R situated in a world $\mathbb{W}_E$ in which a certain event E happened and another (hypothetical) scene $\mathbb{S}_{R,W_{\neg E}}$ in an abstract world $W_{\neg E}$ in which it did not. The event E is defined as "a complete collection of related things that have happened (or are happening) in a World within a specific closed interval of time" [131, p. 4]. According to Jaquet-Chiffelle and Casey, the *theoretical* trace is then the "full modification" of that bounded scene by the event after its occurrence—perceptible or not.[3] Opposingly, the *abstract* trace is more confined and is comprised of "the perceptible modification of the Scene of Investigation resulting from the completed Event of interest and subsequent Intrinsic events (internal dynamics of the Scene of Investigation)" [131, p. 8] according to some tolerance relation at precision p. Given the fact that extrinsic events—which can be regarded as pollution [158] or evidence dynamics [44]—happen beside intrinsic events in real-world environments as well, Jaquet-Chiffelle and Casey refer to a tangible world Ω to grasp the *tangible trace*. They define the *tangible trace* at a point in time $t > t_0(E)$ after the beginning of event E by calculating the difference $\widetilde{\Delta_p}$ with respect to the above mentioned

[3] This is modeled by the so-called exact difference operator Δ that takes "[e]ven infinitesimal modifications, at sub-atomic level for example," into account [131, p. 7].

tolerance relation p between the tangible world Ω and a hypothetical one, where the event did not happen ($\Omega_{\neg E}$), as follows [131, p. 8 f.]:

$$\textbf{Tangible Trace}(t,\ p) = \widetilde{\Delta_p}(\mathbb{S}_{R,\Omega}(t);\ \mathbb{S}_{R,\Omega_{\neg E}}(t)), t > t_0(E)\,, \tag{2.1}$$

Here, it must be noted that such tangible traces are not always directly observable by the human means of the forensic examiners themselves. Often some sort of technological aid in the form of an observation instrument is needed to make the differences perceptible. Catchy examples are the profiling of so-called trace DNA ([162, p. 434]) and Luminol tests to reveal blood remnants [p. 10][131]. Jaquet-Chiffelle and Casey propose to call those perceptible differences *facets*, which are defined as follows:

Definition 2.4.1 (Facet [131, p. 10])**.** A *facet* denotes a perceivable part of a tangible trace (2.1), which could be actually detected, recorded, viewed, and examined by some sort of observation instrument.

When studying a tangible trace only certain facets are considered while others are not taken into consideration. In practice, a particular observed facet is often used to refer to and describe the trace itself although it is just the observed part of it. The respective employed perspective is then reflected in the designation of the facet, which describes either the properties ("biological trace, micro trace, paint trace, digital trace") or the presumed entity ("shoe trace, tool trace") [131, p. 10].

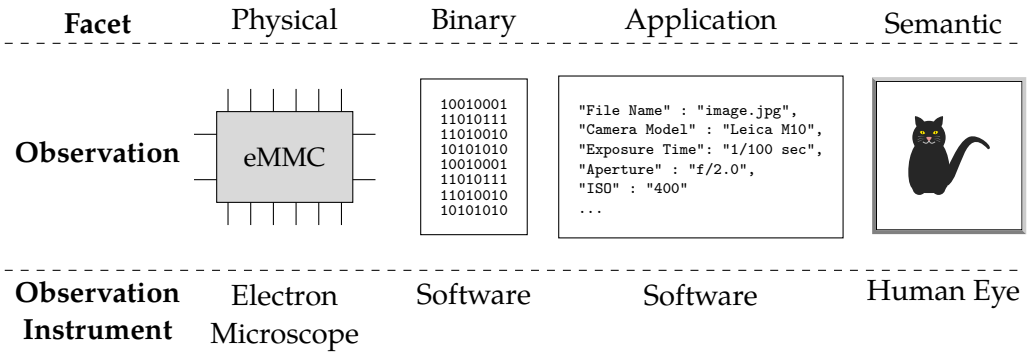

Figure 2.5: The different observation levels of a digital tangible trace of an image according to Jaquet-Chiffelle and Casey [131, p. 11]; Depending on the observation instrument used by the examiner a different facet of the tangible trace is perceived. On the physical level, the examiner can analyze a storage medium, e.g, an eMMC-chip, via an electron microscope. On the binary level, the image could be examined by a hex viewer. Using file format-specific software, the examiner could view (meta)data or pixel values on the application level. Lastly, the digital tangible trace could be examined via the human eye on the semantic level.

Fig. 2.5 exemplifies the different facets stemming from the same tangible trace that are subject of the examination depending on which observation instrument is used. It is shown that an image stored on a digital storage medium, such as an eMMC-chip, could be observed on physical level by using an electron microsocope. Using a piece of software,

e.g., a hex viewer, it could be observed on the binary level. When a file format-specific software is employed, the pixel data and the corresponding metadata can be observed on the application level. Lastly, when the image is presented visually, it can be observed via the human eye on the semantic level.

Fig. 2.5 gave a first insight into digital tangible traces. Jaquet-Chiffelle and Casey defined the digital tangible trace as follows:

Definition 2.4.2 (Digital tangible trace [131, p. 10])**.** A *digital tangible trace* denotes "the modifications of the Scene, subsequently perceptible in binary form, resulting from the Event of interest and subsequent intrinsic events."

So this definition and the preceeding explanations provide a solid ground for reasoning more precisely about what has been previously described as vestiges and remnants of previous events [158, pp. 32 f.] and, thus, allows approaching the question of what makes a facet inferred from a tangible traces evidential.

2.4.3 From Facets to (Digital) Evidence

A crime scene is full of tangible traces that could be potentially observed to find entry into the criminal procedure. However, most of those are caused by "normal" circumstances and are hence not particularly helpful (some even useless) for the investigation at hand [158, p. 33]. Therefore, the present subsection will focus on when we consider a facet to be evidence. To do so, we now look at the process of finding evidence, aspects of traditional and digital evidence, its features, and required reliability before we have a first glimpse at relevance.

2.4.3.1 The Process

For a successful criminal investigation, "[a]sking relevant questions and defining the problems to be solved are the crucial first steps" according to De Forest [55, p. 199]. To do so effectively, a structured and systematic procedure is necessary as shown below.

The Scientific Method. Criminalistics researchers proposed to address posing and answering these questions in a structured manner by employing a scientific method as elaborated by De Forest [55, Tab. 2], whose approach geared toward criminalistic application is visually depicted in Fig. 2.6.

There are four stages: At first, data is gathered by the major activity of observation, then the findings are processed by identifying connections and interrelations to infer a hypothesis. This inferred hypothesis "is checked against the data and observations" (what actually constitutes facets as observable parts of tangible traces) to refine or even discard it after

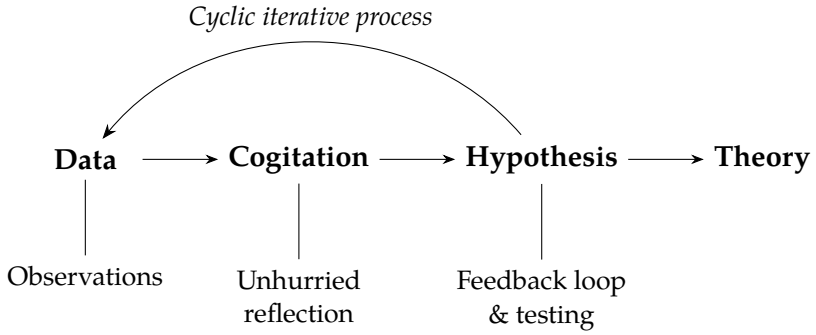

Figure 2.6: Scientific method employed in criminalistics as proposed by De Forest [55].

falsification; if it has been verified, the investigators will arrive at a theory explaining the observations and, hence, the course of actions of the investigated deed. It is crucial, however, to consider the possibility of faulty observations due to flaws in the collection and analysis methods, which necessitates seeking both support and refuting evidence potentially leading to alternative hypotheses [37, p. 6]. The approach outlined in Fig. 2.6 aims to be sufficiently systematic on the one hand and yet flexible enough to be applicable to the "unique aspects of each case" on the other [55, p. 200]. In essence, this method is still considered to be "one of the most powerful tools available to forensic examiners fulfilling our responsibility to provide accurate evidence relating to an investigation in an objective manner", as Casey attested [37, p. 5 f.].

The Criminalistic Cycle. The scientific method, as set out above, remains quite elusive. Hence, the Suisse criminal jurists Walder and Hansjakob came up with a comparable yet restricted approach, which should be easier to apply in day-to-day investigations. They defined the criminalistic task as finding the truth so that the questions of factuality, illegality, and guilt can be answered [240, pp. 10 ff.]. Numerous scholars and practitioners have tried to accomplish this by the repeated introduction of generally applicable models, like the British *Core Investigative Doctrine* [6, p. 48] or the *Criminal Case Analysis*[4] by the German criminalist Ackermann [2]. However, Walder and Hansjakob emphasized the procedural aspect of the criminalistic task and created the five-step *Criminalistic Cycle* presented in Fig. 2.7, which is based on the FBI's intelligence cycle, a thinking model regularly applied in the context of intelligence gathering [181, p. 3]. The approach of Walder and Hansjakob [240] generalizes iterative process models and aims to be applicable to the full range of crimes. Due to the simple generics of this cycle on the one hand and the simultaneous accuracy of the iterative phases on the other hand, this thinking model is convincing, although the parallelism of the activities, which should frequently occur in practice, is not obvious. The iterative nature of this process is a central feature of the investigative process, which matches the conceptions of other authors, such as Stelfox [213, p. 153].

The cycle can be summarized as follows: After becoming suspicious as a consequence of present information or a bundle of that (step 1) the provisional data has to be analyzed (step 2), in order to be able to derive hypotheses describing the potential courses of

[4] Originally called *Kriminalistische Fallanalyse* by Ackermann.

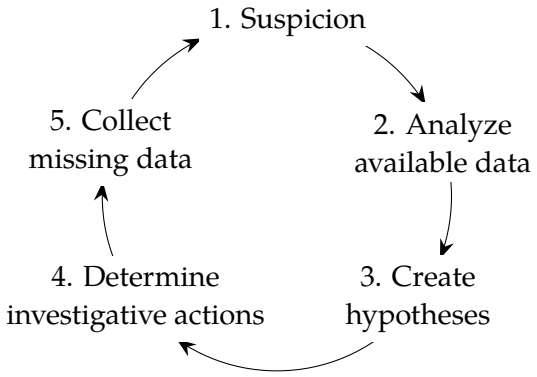

Figure 2.7: Criminalistic Cycle by Walder and Hansjakob [240, p. 93].

action of the deed (step 3). For each and every potential hypothesis, it has to be determined which data has to be collected to verify or falsify the hypothesis in question aiming to prove the objective and subjective elements of the offense (step 4). Afterward, the investigative actions determined in the previous step will be executed (step 5). If their results verify the hypothesis, the criminalistic task is solved. If this is not the case, the initial suspicion must be adapted (back at step 1) to run through this cycle again [240, pp. 90 ff.].

So the collection and analysis of data related to hypotheses formed during the investigation are key activities. These observations recorded as data stem from or form evidence items; thus, we now look at term "evidence" and its qualities in the next section.

2.4.3.2 Physical and Digital Evidence

In view of the scientific method and the more practically oriented Criminalistic Cycle set out above, observations, or in this context more precisely facets, can back hypotheses and, hence, can act as evidence for the hypothesis in question. This process of a tangible trace becoming a piece of evidence is called the *analytical chain of evidence*, which is characterized by transforming data (in its broadest sense as set out in the scientific method by De Forest [55]) into information [114, p. 208 ff.]. According to Hazard [114], this transformation process is divided into three distinct stages that a piece of evidence goes through:

1. physical trace,
2. clue, and
3. evidence.

A "physical trace is 'a mark, a signal or an object' " existing "without any given meaning" on the first stage.[5] A clue is understood as a physical trace (better a facet) that is recognized by an investigator as an indicative sign for presence or action on the second stage. Evidence is then a clue that affects the probability of a proposition [114, p. 209], which denotes the final stage of the analytical chain of evidence.

[5]We think this aptly corresponds to a facet, which would be a more precise yet universal term.

Physical Evidence. Consulting the *Oxford English Dictionary*, we see that evidence is generally considered to be "testimony or facts tending to prove or disprove any conclusion" [66, II.5]. More specifically in forensic contexts, evidence can be defined as "[i]nformation, whether in the form of personal testimony, the language of documents, or the production of material objects, that is given in a legal investigation, to establish the fact or point in question" [66, III.6]. Looking at the procedural law in some jurisdictions, one sees that evidence admitted in the main hearing might be even more restricted; according to the German Code of Criminal Procedure,[6] for instance, strict proof can only be provided by expert witness testimony (§§ 72 ff. GCCP), documents (§ 249 GCCP), judicial inspection (§ 86 GCCP), testimony provided by the defendant (§ 136 GCCP), and witness testimony (§§ 48 ff. GCCP). Given that we deliberately employ an investigative viewpoint here as used by criminalists to remain detached from specific jurisdictions (as well as from specific uses for law enforcement, intelligence gathering, or corporate investigations), we resort to an established general definition provided by Miller and refer to evidence as "information used to decide whether disputed propositions are true" [165, p. 21]). It must be emphasized that this information must be comprised of actual facts; those are understood as "external events, conditions, circumstances, contexts, conditions, experiences or interpersonal mental facts such as motives, plans, views, thoughts" [206, translated by the author] but actually it is "no more or less than that which tends to persuade the court to a particular conclusion" [210, p. 6].

Digital Evidence. While the previously presented definitions are applicable to digital objects as well, we now look at the specifics of *digital* evidence, with which cybercriminalistics is primarily concerned. There exist numerous definitions (for an entrypoint refer to the dissertation of Hargreaves [111, pp. 15 ff.]). Obviously, notable definitions focus on the digitized nature of the data: For instance, the Association of Chief Police Officers [7] grasps computer-based evidence as "information and data of investigative value that is stored on or transmitted by a computer." The term *investigative value* is very broad. Casey [36] provides a more specific definition and considers "any data stored or transmitted using a computer that support or refute a theory of how an offense occured or that address critical elements of the offense such as intent or alibi" to be digital evidence, by doing so the author deviates from the definitions provided by the standard bodies, the Standard Working Group on Digital Evidence (SWGDE) and the International Organziation of Computer Evidence (IOCE), who focus "too heavily on proof and neglect data that simply further an investigation" [36, p. 12]. Hence, the investigative value of the definition proposed by the Association of Chief Police Officers [7] seems to be acknowledged by him as well. A similar but slightly more precise definition is provided by Carrier [32]: in his dissertation, he defines digital evidence as "digital data that supports or refutes a hypothesis about digital events or the state of digital data" [32, p. 11]. While this is perfectly sensible for the scope of his theoretical work concerned with event reconstruction in digital systems, Hargreaves [111] critized the restriction regarding "digital events". He comprehensibly argues that those are merely representations of events in the physical world, which enables the investigator to also assess hypotheses regarding real-world events by this kind of evidence [111, p. 16]. A previous definition by Carrier [31] does not include this overly strict restriction and refers to "a digital object that contains reliable information

[6] In German "Strafprozessordnung".

that supports or refutes a hypothesis" [31, p. 4]. He makes clear that not any but reliable information can be used to assess hypotheses—both concerning digital and real-world events. Thus, this definition is both convincing and preferential [111, pp. 16 f.]. To gain a better understanding of Carrier's notion, we now scruitinize the components and start with the specific features of the "digital object" being the central subject of the definition.

Features of Digital Evidence. The special nature of digital traces is worth noting because it largely impacts collection, examination as well as the later assessment in the big picture of a case. These are topics that are primarily relevant for Chapters 5 and 6. Already in 1997, Sommer [208] described how computer evidence differs from classical evidence. There, the author focused on rapid changes within a moment, e.g., during transmission, easy alteration and modifications during the collection process, and the constant need for interpretation [208, p. 140]. Dewald [63] refined the features of digital evidence and provided an in-depth review of these specifics, i.e., technical (un)avoidability, volatility, manipulability, copyability, and semantics, in his dissertation [63, pp. 39–46]. These can be briefly outlined as follows: Volatility classifies digital traces according to the time frame, in which they are available for collection. They can be either persistent, semi-persistent or volatile. Persistent traces are accessible for a long time since they remain in their state without power supply. Semi-persistent traces are only accessible as long as the power supply is available otherwise they disappear after a short period. Volatile traces are temporary and constantly change during regular operation [63, pp. 39 f.]. The basis for technical (un)avoidability has been identified by the seminal work on file system forensics by Carrier [31]. Certain datastructures are *essential* for a flawless operation of the file system. Dewald coined these traces "technically unavoidable" and defined that they are inevitably created. The counterpart are technically avoidable traces, which are "created for their own sake", hence, not impacting the correct functioning of the system [63, pp. 40 f.]. That class of traces is especially vulnerable for manipulation. In 2015, Freiling and Gruhn [85] refined this black and white categorization by splitting it into a pure and an application-dependent notion. Using the latter approach, they categorized data fields as strongly essential, weakly essential, or non-essential depending on whether the data field is necessary for correct operation for all, some or none application, which, hence, has impact on the trustworthiness of the data field when used as evidence. Weakly and non-essential data fields are more prone to manipulation; thus, manipulability is another feature of digital traces. While there is a risk for incomprehensible manipulation when analyzing a "closed system" [63, p. 42], recent research by Freiling and Hösch [86] and Schneider et al. [201] suggests that creating convincing forgeries is demanding. An advantageous feature, however, is the copyability of digital data. Artifacts coded in discrete form, i.e., binary data, allows us to create perfect copies that are indistinguishable of the original [63, p. 43]. Lastly, (differing) semantics on different abstraction levels complicate the interpretation. Carrier described "each abstraction layer [...] as a function of inputs and outputs" by applying a translation rule set describing how the input data should be transformed and processed respectively, which might introduce a margin of error [30, p. 3 f.]. On the application layer, Dewald further classified the data occuring: He differentiated between primary data, secondary data, program data, configuration data, and log data. Primary data are those data for whose processing an application has been designed. Secondary data ease this processing by providing caching, indexing, journaling or other related

functionality [63, pp. 43 ff.]. The naming of the latter of the three can be considered more or less self-explanatory: Program data comprises the (machine or byte) code as well as source code to run an application. Configuration data contains parameters influencing the program's execution. Log data—if available—might be largely helpful in understanding past program behavior.

Since the man-made abstraction is a basic principle and a central aspect in IT-systems (refer for example to computer networks [219]), we observe an alienation of the digital evidence from natural laws. While traditional evidence is based on universally valid laws of nature—even when dealing with man made analogous items, the latter is a result of human minds designing a piece of software instructing a machine, which creates certain artifacts used as evidence. This results in changing evidential items [149, p. 2], because they are largely dependent of the system design—meaning that the artifacts resulting in digital evidence are fast moving targets opposed to the items of concern of the branches of traditional forensic science.

As already indicated, these specifics of the digital realm and the evidence acquired there also influences its reliability. Drawing on the definition of digital evidence provided by Carrier [31], we learnt that the "digital object" needs to provide "reliable information" to be considered evidential in the forensic context. Hence, we now introduce the notion of reliability of digital evidence in the following paragraph, which is further scrutinized later in Chapter 3.

Reliability of Digital Evidence. Using information stemming from highly complex digital systems for answering investigative and potentially eventually legal questions with stark consequences obviously bears risks. Hence, the reliability of such evidence was in scope of researchers early on. In 1992, Miller [165] was one of the very first researchers to discuss and define reliability for "machine-generated evidence", e.g., information stemming from a speedometer, a breathalyzer, and the like. There, he subdivided reliability into authenticity, accuracy, and completeness to point out general problems in establishing and assessing those features for machine-generated evidence [165, pp. 24 ff.]. Sommer [208] took up Miller's considerations from a more technical point of view and applied it to digital evidence in general by presenting various relevant artifacts and notable criteria to consider by the trier of fact. He concisely paraphrased and listed these key features as follows [208, p. 139]:

Authentic
: Specifically linked to the circumstances and persons alleged.

Accurate
: Free from any reasonable doubt about the quality of procedures used to collect the material, analyse the material if that is appropriate and necessary and finally to introduce it into court – and produced by someone who can explain what has been done. In the case of Exhibits which themselves contain statements – a letter or other document, for example – 'accuracy' must also encompass accuracy of content; and that normally requires the

> document's originator to make a Witness Statement and be available for cross-examination.
>
> **Complete**
>
> Tells within its own terms a complete story of a particular set of circumstances or events [emphases, closing punctuations, and line breaks added for readability].

By resorting to the properties deviating from physical traces, he presented key tests to check whether remotely-acquired digital evidence meets traditional evidence standards. Those include the assessment of whether the remote computer worked correctly and whether the provenance has been both recorded and tested. This requires a comprehensible acquisition process and the comprehensibility of the chain of custody, viz., the transparent continuity of evidence. Furthermore, he underlined the assessment of the linkage of computer data to the accused party and documented as well as sound processing of the evidence for presentation, which he compiled into a pragmatic and informal set of questions [208, p. 141]. Roughly a decade later, Hargreaves revisited Miller's proposal in his dissertation, inferred apt definitions for digital evidence with a focus on live analysis scenarios, and compared those with Sommer's understanding [111, pp. 60–68].[7] In this comparison, he concluded that "for the purposes of assessing reliability of digital evidence from live investigations, the proposed requirements provide more explicit criteria describing what is necessary for digital evidence to be considered reliable" [111, p. 68]. His refined requirements for establishing reliability of digital evidence from live investigations can be summarized as follows: First, authenticity is the ability to demonstrate the origin of the digital evidence to provide traceability to some physical piece of evidence and sampling or collection process [111, p. 63]. Second, accuracy demands that the acquisition, processing, and interpretation errors should be assessable and stay within acceptable bounds for the investigation at hand [111, p. 64]. Lastly, completeness necessitates that the maximum amount of digital evidence related to the investigation should be preserved while being able to assess which is lost [111, p. 65]. To deal with the concept of digital evidence, as we do in this section, this understanding is sufficient so far; however, while the characteristic of authenticity of digital evidence is well researched [e.g., 161], completeness and accuracy need more discussion because the former is very much case-dependent, and the latter seems to be hard to nail down which is why we will later propose a formal approach in the following Chapter 3.

A More Precise Definition of Digital Evidence. In view of the sharpened terminology and understanding provided by Jaquet-Chiffelle and Casey [131] and the reliability properties of Sommer [210], which have been reconsidered by Hargreaves [111], we can even refine the definition of digital evidence to make it more precise. Based on these precursing works, we propose to define the term as follows:

Definition 2.4.3 (Digital evidence)**.** *Digital evidence* is formed by an authentic observed facet of a digital tangible trace (Definition 2.4.2) that accurately supports or refutes a case-related hypothesis.

[7] Hargreaves [111, pp. 66 ff.] referred to another publication of Sommer [209] but his remarks are in complete agreement.

Differing from the definition proposed by Carrier [31], i.e., "a digital object that contains reliable information that supports or refutes a hypothesis", two aspects are refined now: First, the aspect of "reliability" is broken down into the authenticity and accuracy of the evidence. Authenticity und accuracy are used according to Hargreaves's understanding, viz., that the facet's provenance is undoubtedly traceable and that the errors during its processing stay within acceptable bounds. The completeness-property, however, has been dropped since this feature is related to an evidence collection, i.e., multiple facets in a case, and not to a single piece of evidence (i.e., a single exhibit) because we consider the accuracy-property to be violated if the observation-process of a tangible trace, in which the facet should be recorded, does not produce a sensible and in itself complete facet. Second, we now use the precise term "observed facet of a digital tangible trace" instead of referring to "a digital object that contains [...] information". This is avantageous for several reasons: It is precisely defined in mathematical terms representing Locard's *exchange principle*. Referring to this term supports a unified understanding and eases communication across other forensic disciplines. Furthermore, this definition distinguishes between digitally observed facets of *non-digital* tangible traces, e.g., a digitally scanned latent fingerprint, and an observed facet of a *digital* tangible trace, e.g., the binary representation extracted from the memory cells of flash memory. Finally, this abstracts from an object, such as a file, a socket descriptor, and so on, to a precise description of the modification of the digital dimension of the scene.

A Model of (Digital) Evidence by Freiling and Sack [84]. When employing Definition 2.4.3 and working with "observed facets of a digital tangible trace" as evidence, some confusion might arise what actually is the *piece* of evidence and also how it can be introduced into the main hearing of the court. Is the piece of evidence only the seized device? Is it the storage media on which a forensic container is stored? Is it the printout of an incriminated image? While the question of how the evidence is introduced into the main hearing is largely subject to the respective legal system, the question of what constitutes the piece of evidence can be generally answered from an investigative point of view.

To clarify that and unify the terminology employed when talking about evidence despite the particularities of physical and digital evidence, Freiling and Sack [84] made an effort to introduce an abstract model of evidence items—the so-called *CSI model*, an acronym made up of its components, i.e., *claim*, *support*, and *information*. To illustrate the model, let us consider the example of a smartphone that has been seized, then safely sealed in an evidence bag, and tagged with all the necessary information on a label, such as the file number of the case, the evidence number, the date and location of the seizure, the seizing officer, and so on. Much like as in this example, a piece of evidence is commonly comprised of the physical object, a label (usually a tag on the evidence bag), and the actual meaning in the case [84, p. 326]. This triple of basic parts (C, S, I) was further abstracted to generally define a *piece of evidence* as shown in Definition 2.4.4.

Definition 2.4.4 (Piece of evidence [84, pp. 326 f.])**.** A *piece of evidence*[8] consists of three components:

1. the *claim* of what the piece of evidence pretends to be,
2. the *support* functioning as a carrier, and
3. the *information* of the actual trace.

It is easy to see that the *claim* corresponds to the evidence label stuck on the bag. The *support* corresponds to the content within the evidence bag, whereas the support includes both the carrier and the trace in the notion of Freiling and Sack [84]. Lastly, the *information* comprises the actual trace (in an abstract manner). It is best understood as the (case-relevant) knowledge that can be extracted from the *support*. Note that multiple different kinds of information, e.g., DNA-particles, a fingerprint, and gunshot residue, might be extracted from a single *support* and that this *information* can (potentially) exist actually detached from the support [199, pp. 58 ff.].

Obviously, the last definitional building block of *information* remains rather elusive in this definition and Freiling and Sack [84] vaguely thought of it as an overapproximation by including all features of the object and their mental interpretation by the investigators. However, even though the component *information* is not easy to grasp, employing this model, enables a clear distinction of which part constitutes the vital part and is actually of (main) concern for the trier of fact—independently of handling physical matter or a facet of a digital tangible trace.

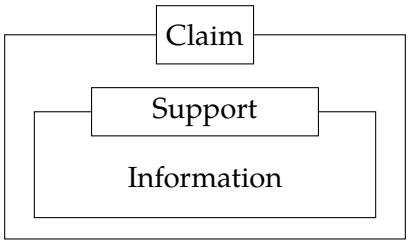

Figure 2.8: The CSI model of (digital) evidence, as proposed by Freiling and Sack [84], illustrating the inclusive relationship of the components.

Freiling and Sack [84] clarified that pieces of evidence have a lifecycle and their claim, support, and information might—and during analysis often—change [84, pp. 330 ff.]. An easy-to-grasp example might be the genetic analysis of dried blood found on a textile piece. So a piece of evidence, represented by a triple (C, S, I), is transformed into another representation (C^*, S^*, I^*) during the analysis phase. Freiling and Sack described that new pieces of evidence may emerge from the original during this phase [199, p. 330]. To describe such transformations, they used the formal notion of a directed graph $G = (P, E)$, whereas P is the set of vertices (all potential pieces of evidence) and $E \subseteq P \times P$ is the set of edges denoting the transitions.[9] The lifecycle of a piece of evidence is then depicted by

[8] Freiling and Sack [84] used originally the German word "Beweismittel".

[9] Note that we use the symbol P for "piece of evidence" instead of the originally used symbol B for the German term "Beweismittel" to reflect the English terms in the explanations.

a specific path through G. To illustrate the transformation of a piece of evidence, refer to the simple exemplary case shown in Example 2.4.5.

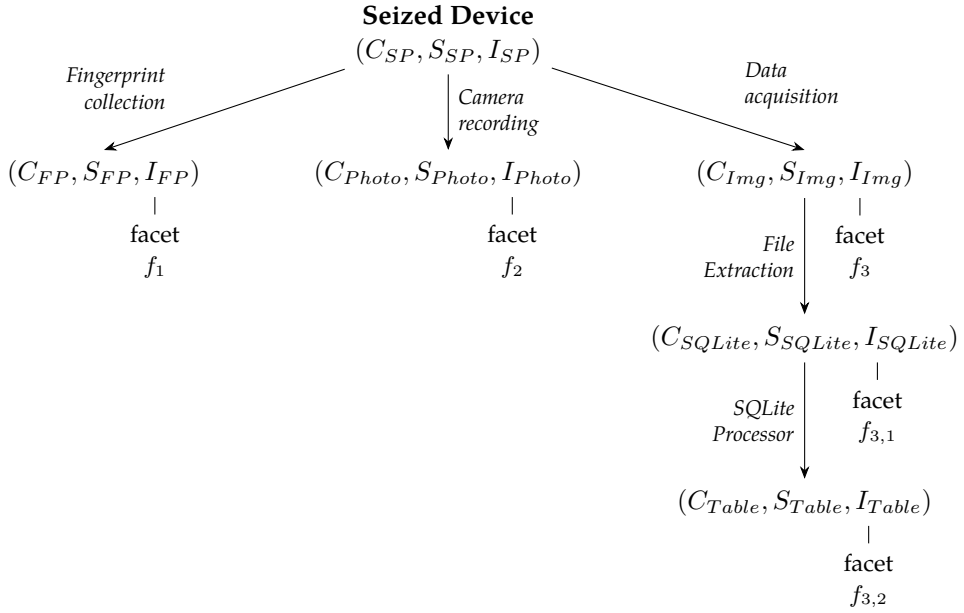

Figure 2.9: Derived pieces of evidence from a seized smartphone using the CSI model; Transformations of the initial (C_{SP}, S_{SP}, I_{SP}) triple denoting a seized smartphone in an exemplary case of blackmailing (Example 2.4.5) and the subsequently derived pieces of evidence.

Example 2.4.5 (Transformations of pieces of evidence according to the CSI model of Freiling and Sack [84] in an exemplary case of blackmailing)**.** Let us assume a suspect threatened a victim to publish intimate recordings of them by sending messages via a messaging app: It is easy to see that the suspect's smartphone constitutes a vital piece of evidence to solve the case; hence, investigators would seize the smartphone (C_{SP}, S_{SP}, I_{SP}). The claim identifies and signifies the item's origin, the time and location of seizure, and some additional context of the case. The smartphone itself constitutes the support for multiple information components. There might be latent fingerprints of the suspect constituting one part of a physical tangible trace, which can be observed using magnetic powder that is collected on adhesive foil to form a facet with a refined claim and a different support (C_{FP}, S_{FP}, I_{FP}). A photographic image taken of the display showing the extorting messages by the investigators right during the search forms another facet with another refined claim and a new support $(C_{Photo}, S_{Photo}, I_{Photo})$. Additionally, the data stored on flash-storage of the smartphone contains a digital tangible trace, which can be acquired using a forensic acquisition toolsuite to preserve the storage's content as an image in a forensic container format forming another facet on the binary level linked to a refined claim and a different support $(C_{Img}, S_{Img}, I_{Img})$. Of course, not the whole block of data is relevant, so the SQLite-database of the messaging app in question might be extracted forming a facet on the application level and stored on some storage media $(C_{SQLite}, S_{SQLite}, I_{SQLite})$. From this database, the relevant chat with the victim can be extracted using the SQLite Command Line Shell sending the respective statements to form a new facet on the semantic level that is presented in tabular form on a paper printout to the prosecutor and

the court $(C_{Table}, S_{Table}, I_{Table})$. We can compactly visualize this in Fig. 2.9 and see how it helps to precisely view the transformations a piece of evidence might go through during its lifecycle.

Relating the Models. While the CSI model may initially seem abstract, the definitions of digital evidence (Definition 2.4.3) and a piece of evidence (Definition 2.4.4) integrate well—the former can be contained in the latter. The formalized study of the trace by Jaquet-Chiffelle and Casey [131] and the CSI model by Freiling and Sack [84] employ substantially different approaches and thus provide a different view of the matter. However, we also see the potential that the formalized study of the trace by Jaquet-Chiffelle and Casey [131] and the CSI model by Freiling and Sack [84] can complement each other. As indicated in Example 2.4.5, we argue that the concept of facets roughly aligns with the combination of the *support* and *information* components. The *support* component either constitutes the observation of a tangible trace itself, e.g., a bloody knife, or serves to record the observation, e.g., a DNA profile as a result of the analysis of some biological material. The *information* component includes the observation and the inherently linked conclusion of the event of interest; however, that has yet to be scrutinized in Section 3.6.2. In order to be helpful in an investigation, the observed facets have to be attached to a *claim*, thus forming a piece of evidence according to the CSI model.

These model-based approaches to reasoning about evidence enable investigators and forensic scientists to clearly differentiate the elements of traces and understand the lifecycle of a piece of evidence to answer investigative questions, which are focused next.

2.5 A Closer Look at Investigative Hypotheses

"*Hypotheses are nets: only he who casts will catch.*"—not exclusively in research but also in a criminalistic context this statement by Novalis (cited after Popper [187, p. xiv]) appears to be literally accurate since the formation of hypotheses constitutes a central aspect of criminalistic case work [25, p. 300]. In a general research process, hypotheses are built upon observations. Those as well as additional observations can form evidence that supports or refutes a particular hypothesis. In a criminalistic context, however, we need to ask what the specifics are when dealing with hypotheses and their generation. This is a question at which we will now look more closely after briefly revisiting the basics.

2.5.1 Terminological Consideration

Looking at the etymological roots, the term *hypothesis* stems from an ancient greek word meaning *basis of an argument*, literally "a placing under" (thesis = a placing, hypo = under) [70]. So, the observations are placed under the assumption of a specific explanation. De facto, this is a tentative conjecture explaining the observations while it is testable, i.e., verifiable or falsifiable, what constitutes an essential feature of (scientific) hypotheses [186,

Chap. 1]. In the field of forensic science some authors, e.g., Cook et al. [51] resort to the term "proposition", either generally or to distinctly deem the remaining hypotheses after an investigation [193, p. 139]. In the absence of explicit explanations for this deviation, we assume that this should take up the distinction of testability in the terminology used; So, this term is used (a) to suggest a link between two concepts, i.e., the assumed causal relationship between the event and the observed trace, and (b) to denote that this assumed link, which constitutes an explanation of the observations related to the deed, cannot be verified by experiment with absolute certainty. This underlines that forensic scientists cannot prove the proposition with absolute scientific rigor; however, their repeated failure of falsification despite the use of appropriate methods developed by sound research can suggest the proposition's validity [63, p. 14]. Given the widespread use of the term *hypothesis*, e.g., Casey and Rose define that term in the context of forensic science as "an informed supposition to explain the observations in the data gathering phase of a forensic analysis" [40, p. 48], we do not explicitely discriminate between the terms *proposition* and *hypothesis* in the further course of this thesis, but we acknowledge that the fact in question cannot be ultimately proven in real-world forensic encounters but may lead to a sufficient conviction of the involved parties regarding the resulting theory of (elements of) the deed. Although we could dive deeply into a research-theoretic and epistemological excursion here, we now shift our focus to criminalistic use.

In addition to the term *hypothesis* in investigations, some German criminalists refer to the locution *criminalistic version*. While some of these authors use the two terms interchangeably in broad accordance with the meaning of the term *hypothesis* set out above, e.g., Brodag [25, p. 300] and Roll [196, p. 382], others, e.g., Gundlach as well as Walder and Hansjakob [240], differentiate between those. The authors of the latter group argue that (event) versions provide explanations of the end result of a criminal deed ("the what"); hypotheses (of the deed) in turn provide downstream explanations of how the course of actions occured ("the how") [240, p. 171]. So, the different terms are used to discriminate presumptions of differing subject matters and relate them in a certain hierarchical order—a feature that is encountered in the next paragraph as well.

2.5.2 The Formation of Hypotheses

Solving cases involves forming possible explanations on the course of events and the commission of the assumed offense in form of versions and hypotheses—a process that is a central component of criminalistic thinking. Thereby, it is crucial that the criminalists and forensic examiners act as a "neutral finders of fact" [37, p. 5] and avoid any prejudice. The effectivness and success is largely dependent of the objectivity of the investigators [63, p. 13]. Forming hypotheses influenced on subjective assumptions, i.e., prejudices, not backed by objective observations, is one of the most common mistakes in this process [40, p. 48]. The formation of hypotheses and their assessment is the criminalist's tool to reconstruct events in the exterior world, physical or digital, and even the inner world of a suspect (at least to a certain degree). When chained together, those hypotheses will form a version of the deed and, hence, allow for the clarification of the case by assessing each hypothesis against a counter-hypothesis.

Despite that all hypotheses are based on traces observed at some criminalistically relevant location, that is usually a (primary or secondary) crime scene [25, p. 301], there is also a subjective component to hypotheses generation: It is crucial to build upon (experiental) knowledge of specific criminal phenomena and even imagination [25, p. 303 f.]. Imagination requires a great deal of creativity, experience, and mental openess to develop apt hypotheses, facilitating both intuitive and reflexive thinking methods in combination [240, pp. 173 ff.].

For a successful investigation, it is generally advised to start with many different versions of the possible course of actions to begin with, so that no premature determination is made by the investigators [25, p. 300]. The amount of versions corresponds to the breadth of the investigation, as depicted in Fig. 2.10. Based on these versions, the investigators put up more detailed hypotheses in order to falsify or refute those. The ones that are not excluded by the assessments serve as starting point for refined, more fine-grained hypotheses [63, pp. 13 f.]. The amount of subsequent hypotheses corresponds to the depth of the investigation. The more of those subsequent hypotheses have unsuccessfully been tried to falsify in a flawless manner, the more probable the resulting version of the deed is and the more thourough the investigation has been conducted.

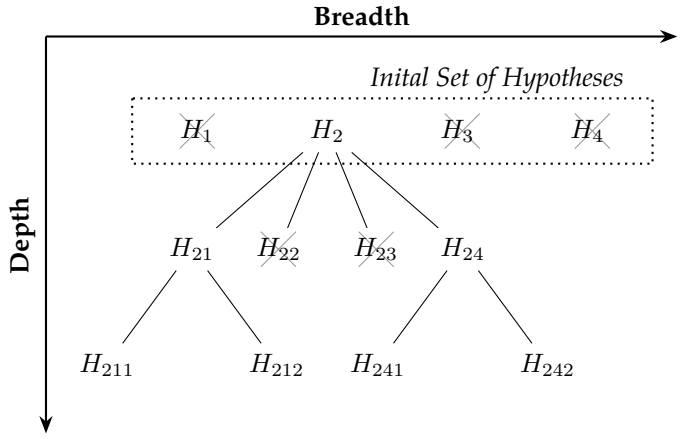

Figure 2.10: The attributes *breadth* and *depth* of the hypotheses formation process based on Dewald [63].

2.5.2.1 The *Hierarchy of Propositions* by Cook et al. [51]

In 1998, Cook et al. [51] proposed a classification of the formulated hypotheses, which they referred to as propositions, that can guide forensic scientists in assessing "the evidential weight associated with the presence or absence of transferred material" [51, p. 238]. The authors argued that the forensics scientist has to consider a pair of propositions corresponding to the prosecution and defense positions, in order to assess the evidential weight of a transfer of matter (following Locard's *exchange principle* [152]). They noticed that the propositions in question can be sorted into a hierarchy. Hence, Cook et al. distinguished between three levels on which hypotheses of the deed and perpertration should be discussed. These are the *source level* (1), the *activity level* (2), and lastly the *offense level* (3). Source level propositions are concerned with establishing an association of transfer material, i.e.,

comparing "samples and some kind of population of alternative sources" [51, p. 233]. The propositions of the second level deal with activities, they require more circumstantial information and are "not necessarily dependent on transfer material being found". Lastly, the offense level propositions might also deal with events but in the light of the broader legal question. Here, the entire matter has to be taken into account and assessed in order to ascertain the presence of the features of the deed. Table 2.1 illustrates the concept using an example referenced in the article of Cook et al. [51], where pairs of propositions on the different hierarchy levels for a case of burglary, which involved breaking a window to make entry into the property, are formed.

Table 2.1
An exemplary instance of the hierarchy of propositions in a case of burglary; taken from Cook et al. [51, p. 2].

Level	Generic	Propositions
III	Offense	Mr. A committed the burglary
		Another person committed the burglary
II	Activity	Mr. A is the man who smashed window X
		Mr. A was not present, when smashing window X
I	Source	The glass fragments came from window X
		They came from some other broken glass

Scrutinizing this hierarchy of broad categories, Cook et al. concluded the following: "The higher the level of propositions that are addressed, the greater the reliance on the expertise of the scientist" and "the greater the value added by the scientific evidence" [51, p. 238 f.]. Nonetheless, the line between those levels is sometimes thin and the authors argued that they are not rigidly demarcated one from another. Moving up the hierarchy, however, more and more information and assessments of hypotheses of lower levels are required, what they call a "framework of circumstances", in order to verify or falsify the propositions on these upper levels [51, p. 238].

2.5.2.2 The Extended Hierarchy of Propositions in Forensic Genetics

In the field of forensic genetics, an additional level, the so-called *sub-source level*, has been introduced by Gill et al. [91] to accommodate the specifics of DNA evidence, which was taken up later by Vennemann et al. [235]. The main specificity goes back to the DNA's actual purpose of encoding the genetic information to enable evolution; hence, genetic information is highly characteristic and allows to individualize a person, animal, or plant just by assessing the alleles of the DNA particles found. Although that this is a highly specific feature of this type of evidence, the introduction of a separate level seems a bit contrived and raises the question, why the forensic geneticists have not covered the attribution of DNA particles to persons on the source level. In their extended hierarchical model, the source level is then solely used to assess the type of substance, e.g., blood, semen, flakes of shed skin, saliva, and so on [235, p. 397]. So, this extended hierarchy is not readily applicable to other types of physical evidence since the sub-source level has no equivalent there.

2.5.2.3 The Hierarchy of Propositions in Cybercriminalistics

From the perspective of forensic science practitioners, such a structured approach as originally proposed by Cook et al. [51] or the extended version by Gill et al. [91] seems to be vastly helpful. However, the existing publications dealt only with physical evidence and—judging by the absence of literature—the digital forensic science community has not taken up this approach yet. Nevertheless, it can be assumed that there is potential to apply this mental framework to digital evidence and cybercriminal phenomena as well. This assumption follows from the argumentation presented in Section 2.2.2 where—going back to Dewald and Freiling [65]—it has been stated that forensic computing is a true branch of forensic science since it is also concerned with the establishment of associations at its core by employing a hypothesis-driven scientific method. Hence, the three-level approach by Cook et al. [51] is applicable to structuredly discern hypotheses related to digital evidence as well. At first sight, one might even be tempted to project the sub-source level of Gill et al. [91] to digital evidence and consider the propositions on attribution (of using the computing device) on this level. However, despite biometric features being widely used in digital devices today, generally available and precisely individualizing features, such as DNA, have yet to be found in the digital context, which is why we refrain from applying the extended version and adhere to the original one. In order to tangibly apply the original hierarchy proposed by Cook et al., we now look at a fictious case of network intrusion and discuss a few hypotheses related to such a case (listed in Table 2.2).

Table 2.2
An example of the hierarchy of propositions in a case of data espionage.

Lv.	Generic	Exemplary Propositions
III	Offense	Actor T committed data espionage by penetrating the network
		Another actor $\neg T$ committed the data espionage
II	Activity	The binary provides remote access to workstation W_j via I
		The binary does not provide remote access to workstation W_j via I
I	Source	Workstation W_j regularily connected to the IP-address I
		Workstation W_j never connected to IP address I

An Exemplary Case of Network Intrusion. The analysts of a security operations center (SOC) identify packet transmissions in regular intervals from the observed network to a suspicious IP address I by a cursory statistical flow analysis based on their network security monitoring logs at their edge router. The classification as potential beaconing to a command-and-control server constitutes the initial suspicion of data espionage, so they inform the competent law enforcement agency to investigate the case further. Aiming to track down the IP address of the workstation responsible for the network packets through several subnets, the investigators capture TCP flow data. At the nearest managed switch, they initiate a full packet capture using its mirror port. These data could be used to assess the hypotheses on the source level h_1 denoting that workstation W_j is transmitting regular network packets to the suspicious IP address in question. Based on the observed regularity, the investigators might form the hypotheses h_2 on the activity level that it is some sort of beaconing to a command-and-control server enabling remote access into the network.

In order to assess this incident, e.g., to discern the cause and assess the severity, a live analysis of the previously identified workstation W_j might be conducted. The analysts would commonly acquire the main memory of the machine and scan for socket descriptors related to network connections. The find of such a socket descriptor corresponding to the remote IP address and port number can then be used determine the process which created (and whose process space houses) the descriptor. So, the hypothesis on the source level could be even refined to the responsible process running on the workstation. Based on the process, the running binary could be recovered.[10] Decompilation and static code analysis can then be used to gain insights into the functionality of the binary with a focus on the beaconing logic and the question of whether a command-and-control channel is established and, hence, the functionality of remote access is provided by the binary. The resulting insights of the static code analysis allow the assessment of the hypothesis on the activity level h_2 dealing with the ability of remote access. Assuming that h_2 is supported, the investigators would establish code similarity, investigate the command-and-control server with IP address I and track the administrators and operators of this server. By doing so, they would form and assess a variety of additional hypotheses on the source and activity level that are not explicated here. For the sake of the example, they eventually find cues that the binary and the server have been deployed by a certain threat group T. So, they can then finally assess the hypothesis on the offense level h_3 that states that the network has been penetrated by that threat group.

The fictious case study above illustrated that the hierarchy of propositions is easily applicable to cases involving digital evidence and cybercriminal phenomena. This example concludes the present section, in which we looked at the essence of investigative hypotheses and their formation. As sketched out in the introduction (Section 1.1), hypothesis formation is not of central concern to this thesis; still, we consider it highly relevant to convey the basics of structured investigative reasoning since we latently employ it as a starting point in the further chapters.

2.6 Summary

The landscape of (digital) investigations is seemingly meandering and untamed at first sight. Having thrown light on some essential landmarks, the present chapter aims to sketch out a map of the terrain. This includes the terminological clarification of key terms, the revisitation of procedural approaches, and a closer look at the main subject matter of investigations.

At first, we set out the relation of the disciplines of forensic science and criminalistics as well as their digital subdivisions, i.e., digital forensic science and cybercriminalistics. We demarcate forensic science and criminalistics by returning to the initial definition of Hans Gross [98] who employed a broader understanding of criminalistics, which has been sharpened in this chapter by taking up recent terminology. The former is focused on the scientific methods and principles to examine and analyze traces in criminal or justice law

[10] For the sake of simplicity, we assume that no code injection is involved here.

matters. In the notion of Gross, the latter is more than just the holistic use and aggregation of the branches of forensic science, which has contemporarily been termed traceology [158, 194], instead it should be understood as a broader profession and scientific discipline of combatting crime, including phenomenological, organizational, tactical, and strategic considerations besides traceology at its core. Looking at the digital sphere, we further clarify the relationship of digital forensic science and cybercriminalistics while we sorted additional terms, such as computer forensics, forensic computing, and digital investigations, into these broader disciplines. In agreement with other authors in the research discipline, we argue that digital forensic science is a true forensic science and hence a branch of the umbrella term. Digital forensic science can be further divided either into subfields depending on the examination subject, i.e., computer, network, multimedia, and so on, or the orientation of the focus—being either (more) associative or investigative. The latter differentiation separates forensic computing, which deals primarily with the establishment of associations in the true spirit of forensic science [65], from digital investigations, which are concerned with the whole process of handling and processing digital evidence for forensic purposes.

After having clarified the relation of the disciplines, we trace the history of cybercriminalistics research by presenting selected works. Here, we focus on the Cybercrime Investigation Framework proposed by Hunton [122] and its components, i.e., the Cybercrime Execution Stack and the stages of cybercrime investigations. The conceptual framework tries to address the challenges posed by cybercrime phenomena in a generic way. The Cybercrime Execution Stack, which puts abstracted components of cybercrime offense, i.e., data objective, exploitation tactics, attack methods, networked technology, into relation to each other, should enable the investigators to take all relevant technological aspects into account and, thus, simplify the identification of needed capabilities required for conducting a successful investigation. This is supplemented by setting out the stages of cybercrime investigations as a general blueprint schedule. One of the key drivers to unifying these components into the Cybercrime Investigation Framework is to conflate technical aspects of examining digital evidence and the legal processes associated with law enforcement investigations by aligning the cybercrime-specific components with the generalized stages of an investigation, as defined by the Association of Chief Police Officers [6]. However, these works seem to exhibit a largely organizational focus remaining detached from actual investigative concerns.

In the main part of the chapter the investigative core of criminalistics is discussed, which focuses on the trace and its use as evidence. To begin with, the universal forensic paradigm by Inman and Rudin [126] as the frame of retrodictive thinking is revisited: Based on the transfer of traits, investigators place a found trace in a class or group, then they individualize it by determining a common source, which potentially allows the establishment of an association of the trace with the source and serves as the basis for the reconstruction of an event. Given the importance of the trace, its formal study by Jaquet-Chiffelle and Casey [131] is then summarized. Using this concept, the rather elusive vestiges and remnants of previous events can be precisely grasped as facets, which are perceivable parts of a tangible trace detected, recorded, viewed, and examined by some observation instrument. On this basis, the topic of how to move from facets to evidence is tackled. Here, we describe established processes, such as the general scientific method and the

Criminalistic Cycle, to assess investigative questions in a structured way by collecting facets of tangible traces. Since these facets form evidence, we have a closer look at how digital evidence can be grasped definitionally by referring to definitions in existing research literature. After that, we revisit features of digital evidence as put up by Dewald [63], i.e., technical (un)avoidability, volatility, manipulability, copyability and semantics, and have a first glimpse at its reliability and relevance, which will be later formally scrutinized in Chapter 3. Based on these considerations, we provide an improved definition and consider digital evidence to be an authentic observed facet of a digital tangible trace that accurately supports or refutes a case-related hypothesis. Furthermore, we go over the CSI model by Freiling and Sack [84] to use cleaned up terminology not only when talking about the tangible trace but also about the exhibit, which is comprised of a *claim* of what the exhibit pretends to be, the *support* on which the trace is contained, and the extractable *information* according to the CSI model.

The last part of the chapter is focused on investigative hypotheses, their nature, and structured formation. It is laid out how a structured formation process influences the breadth and depth of an investigation. We revisit the hierarchy of propositions by Cook et al. [51] and conclude that it is not limited to physical evidence, which is shown by an exemplary application to a fictious (but nonetheless realistic) cybercrime case.

To sum up, the present chapter aims to paint a rich backdrop for the scientific results to be presented in the following chapters by compiling and preparing preliminary (and partly additional) background information.

3 Relevant and Expressive Digital Evidence

3.1 Introduction

"Trace 4334 has led to the perpetrator" [54]—statements or headlines like this are rather common to accompany news of the resolution of a serious violent crime. One such instance is the murder of Carolin G., who became a victim of a fatal sexual assault while jogging in a wooded area near her home in Baden-Württemberg, Germany. Meticulous work of the task force "SOKO Erle" set up by the criminal investigative department in Freiburg, Germany led to the impressive identification of the perpetrator: The murder weapon could be identified as an iron rod that belonged to hydraulic jacking equipment of a specific truck model. By analyzing and filtering digital toll data of commercial trucks with this information, law enforcement identified the freight forwarding company, where the perpetrator worked. The investigators were then able to identify the suspect by matching the employee data records of this company with connection data of the radio cell servicing the area around the crime scene at the time of the murder to ultimately convict him using DNA samples [54, 216].

Successful investigations, like the one of "SOKO Erle", illustrate that it is extremely difficult to determine beforehand which traces will solve the case. Hence, task forces working on serious crime cases often employ a meticulously methodical approach and follow traces whose importance seems far-fetched at the time of processing. Actually, this could be seen as a trawling method of investigation, because there is—apart from some experiential knowledge—no clear understanding of what traces are of critical relevance in the concrete case. The seemingly infinite amount of digital traces on a system and the circumstance that almost every investigation deals with some pieces of digital evidence—often found on multiple devices—even aggravates the situation. The quickly increasing number of computing devices paired with ever-rising storage capacities renders this approach, if not inapplicable, at least inefficient, so the quest to improve that trawling method becomes even more pressing.

In Chapter 1, we introduced the overarching criminalistic task and named its two subproblems, i.e., finding hypotheses and then identifying traces to assess those. Furthermore, we posed the question of what can be regarded as "sufficient digital evidence". In Chapter 2, we looked at definitions of (digital) evidence and came across the reference to the hypotheses. Regarding the investigative core of (cyber)criminalistics, the relevance of certain traces remained implicit but already shimmered under the surface. Despite the fact that there exists a multitude of studies in digital forensics where specific traces of specific applications were analyzed to determine their meaning, little research has been done to

define what actually is sufficient digital evidence on a foundational level. However, when talking about traces and the observed facets, one cannot avoid investigating the underlying concepts more deeply. In this chapter, we hence explore the question of expressiveness and relevance of digital traces in a formal and universal way. Knowing what story the digital trace might tell—and, thus, how expressive it is—could be helpful in assessing and selecting digital traces; by grasping their relevance, they enable the targeted assessments of case-related hypotheses and could then even be delimited to the most useful ones. In our view, this is significant since such a clarification is needed to work towards the remote vision of exactly knowing where to look to solve a specific offense. This, however, requires a foundational understanding of relevance and expressiveness.

3.1.1 Contribution of the Chapter

In the present chapter, we give formal definitions of expressiveness and relevance of forensic traces. By defining those two largely neglected characteristics, we can precisely infer the semantic quality of traces in regard to the investigative goals. We do not only supplement the (controlled) vocabulary but also relate our definitions of these concepts to other quality criteria of digital evidence concerning its reliability. Using the newly established formal notions, we define the attributes of reliability, i.e., completeness, accuracy, and authenticity, in a rigorous way. Having this improved understanding allows us to reason more precisely about the tenets of investigations because it enables us to answer key questions like how to grasp relevance of traces, especially digital ones, how to determine which ones are most expressive out of all potentially relevant ones, and lastly at what point enough traces have been collected to solve the case. In order to demonstrate this and the general applicability of the formal notions, we formalize and refine the so-called *Criminalistic Cycle*, a thinking model for crime investigations, using our concepts and, hence, propose the *Facet-oriented Criminalistic Cycle* (Fig. 3.2). Additionally, we introduce a measure for expressiveness, which will be later put into practice in Chapter 4.

3.1.2 Chapter Outline

The remainder of this chapter is structured as follows: First, we revisit previous works that are directly or indirectly connected to the topic of relevance and expressiveness of (digital) traces in Section 3.2 to get an initial understanding of what forensic investigators, as well as legal practitioners, consider to be relevant. Then, we propose a certain notion of the terms, relevance and expressiveness, and develop a formalization in Section 3.3. Afterward, we establish a link between those formalized concepts and the notion of reliability in Section 3.4 by inferring accuracy and completeness properties. Based on these insights, we interrelate the terms in a larger context by putting up and explaining a conceptual network to clarify how facets, relevance, expressiveness, hypotheses, and reliability criteria are linked to one another. Using the newly established concepts, we demonstrate the practical use in Section 3.6 by taking the newly established concepts to formalize and thus improve two criminalistic thinking models, the Criminalistic Cycle and the CSI Model, before discussing

the insights and the proposed formalizations in Section 3.7, which leads to outlining a remaining research gap. Finally, we summarize the chapters' insights in Section 3.8.

3.2 Related Work

3.2.1 The Beginnings

Beginning in the early days of forensic science, practitioners realized that having a clear understanding of traces and their meaning is essential to solve the criminalistic task. Hence, influential researchers of the discipline, such as the founding father of criminalistics Hans Gross [98] or the biochemistry professor and pioneer of forensic science Paul Kirk [141], took on the elaborate quest of cataloging physical evidence and looked at various kinds of traces in various criminal contexts and the respective analysis options (available at that time). By doing so, they documented what is relevant for an investigation from a pragmatic point of view. Such approaches are also represented in rather recent research where scientists analyzed the use, the analysis options, interpretation possibilities, and the meaning of specific traces, e.g., gunshot residues [166]. This rather technical viewpoint is oriented towards what is scientifically possible and demanded by the legal practitioners.

3.2.2 The Legal Perspective

In the legal profession, there are various conceptions regarding what constitutes *relevant* evidence in court. Generally, evidence must be relevant in order to be admissible, so this discrimination between relevant and irrelevant items is, in fact, an exclusionary mechanism found in different judicial systems [115, p. 6]. In the UK, for instance, judicial relevance is assumed when the evidence in question makes a fact more probable or when it helps to assess another piece of uncorroborated evidence.[11] Robertson et al. generalized this statement [195, p. 168]:

> The fundamental principle of the law of evidence is that evidence which is relevant is admissible unless it is excluded by some other rule or its probative value is outweighed by its prejudicial effect. The first question to be asked of any scientific evidence therefore is whether it is relevant. We have argued [...] that it is relevant if it helps to distinguish between appropriate hypotheses.

Put simply, if a trace presented to the court helps the jury to decide on the issue, it is considered relevant [23, p. 10]. A strictly forensic statistics approach would define evidence as relevant if it exhibits a likelihood ratio other than 1.0 for two competing hypotheses [195, p. 168], on which the trier of fact has to decide; however, this method is mostly, if not exclusively used in expert witness testimonies. Furthermore, it is important to underline

[11] In the jurisdiction of the UK according to case law by the House of Lords; DPP v Kilbourne [1973] A.C. 729 (31 January 1973).

that relevance is not only restricted to disputing propositions directly. Relevance can also be assumed if the information helps to assess the reliability of the source of the information constituting the evidence or information supplied by another source.[12] Nonetheless, it is important to note here that relevance does not unconditionally imply admissibility because the respective legal process might impose further restrictions regarding admissibility, probative value, weight, and so on. Those, however, are dependent on procedural aspects of the jurisdiction that come on top of the solely investigative perspective that criminalists (and private investigators) may employ to support a decision-making process—potentially initiated by a prosecutor or court in a forensics context [193, p. 138], as we do in the present thesis.

3.2.3 The Criminalistic Perspective

Hazard [115] identified that relevance is subjective to the role: She distinguished between "a forensic and a legal relevancy, showing different, but at the same time complementary perceptions of the same dimension that could be useful to conduct a criminal case" [115, p. 8]. The former notion constitutes a more criminalistically focused view of the relevance of physical traces. Hazard expressed the following understanding of this concept [114, p. 210]:

> The detected physical trace is perceived as being relevant because (1) on a factual and objective point of view, a link has been recognized between the discovered physical trace and the questioned (criminal) activity and (2) it is subjectively appropriate to collect and analyze it since there is a perception of its informative value by the investigators for the case at hand.

The above-mentioned recognition is a reference to semiotic considerations and the theory of sign and signification. Employing such a viewpoint, Hazard tied relevance closely to "the perception of trace-objects [sic] in a specific context" to signify those as clues that are used as evidence [115, p. 7].

By being more broad and general, Hazard accommodated investigative reasoning in the early stages of the criminalistic process where many details of the deed are still indistinct and the evaluation of relevance is ongoing. This also highlights that relevance is a cognitive concept, whose evaluation underlies a certain subjectiveness. While Hazard's statement seems vastly helpful and perfectly valid, we see additional value in differentiating this definition. Given changes in general conditions that come with the use of information technology—both as helpful tools for collection and analysis as well as actual evidence—there are assumed benefits to breaking this down and eliminating certain subjectivity. Hence, we develop tenets in the present chapter and try to lift it from a practical orientation containing certain vagueness to a rigorous formal notion.

[12] "It is relevant if it tends, directly or indirectly, to support or undermine a disputed proposition, or if it relates to the reliability of other information supplied by the source or to the reliability of information supplied by another source of information." [165, p. 22]. More on the topic of reliability in Section 3.4.

3.3 Formal Definition of Expressiveness and Relevance of Facets

In the previous section, we elaborated that the observation of a trace, or a facet to use a more precise terminology, is relevant "if it helps to distinguish between appropriate hypotheses" [195, p. 168]. Note that it is not about *any* hypotheses but *appropriate* ones. So, this means that facets need to have the quality of supporting (maybe even "proving") or refuting a hypothesis that is itself related to a factually relevant event to be considered themselves as factually relevant because they then contribute to the achievement of investigative objectives. That restriction to "appropriate hypothesis" in regard to relevance highlights two different flavors of it, i.e., case-related and hypothesis-related relevance: Hypotheses might exhibit case-related relevance, while facets might exhibit hypothesis-related relevance. Facets that exhibit such relevance regarding multiple hypotheses can be considered to be more expressive. It is sensible for investigators to iteratively search for those more expressive facets tending to prove or refute multiple hypotheses at once.

To further clarify the understanding of *relevance* and *expressiveness* of facets expressed in natural language, we now formalize those concepts to create a solid understanding of their nature and the factors influencing these features. Note that we solely refer to the variant of hypothesis-related relevance of facets—unless explicitly stated otherwise—when we elaborate on the term in the following sections.

To give a clear definition of these terms, we refer to a set of hypotheses

$$H \coloneqq \{h_1, h_2, h_3, \ldots, h_n\},$$

of which each h_i provides some explanation of (the nature of characteristics of) the past events. This allows us to map both direct relevance in terms of disputing propositions and the extended notion of relevance in terms of the reliability of an information source, i.e., the reliability of another facet, by Miller [165]. In addition, we refer to a set of perceivable facets (Definition 2.4.1) of tangible traces,

$$F \coloneqq \{f_1, f_2, f_3, \ldots, f_n\},$$

where each f_i was retrieved by using some observation instrument. Dealing with digital tangible traces (Definition 2.4.2), we consider these set elements to be digital objects on a deliberate abstraction level—ranging from single bits to more abstracted objects, such as file system superblocks or certain entries in a table of an SQLite-database.

Furthermore, we define a relation between facets and hypotheses in an *investigative knowledge base*.

Definition 3.3.1 (Investigative knowledge base)**.** An *investigative knowledge base* $(H, F, \mathsf{supports}, \mathsf{refutes})$ consists of a set H of hypotheses where each element provides a possible explanation of the facets, a set F of perceivable facets of the tangible traces potentially present at crime scenes, together with two relations, $\mathsf{supports} \subseteq F \times H$ and $\mathsf{refutes} \subseteq F \times H$, relating facets to hypotheses with the expected meanings.

We consider a knowledge base to be *consistent*, if no facet both supports and refutes the same hypothesis, i.e., the two relations are disjoint:[13]

$$\mathsf{supports} \cap \mathsf{refutes} \;=\; \varnothing$$

Example 3.3.2. Take for example a case involving a JPEG image. The set of facets would contain, inter alia, the encoded information data structures with metadata (application level) and the visual content (semantic level).[14] A facet on the application level would then be constituted by the JPEG File Interchange Format (JFIF) data structures, encompassing supplementary details such as the *IFD0*, *ExifIFD*, *GPS IFD*, and beyond. These data structures comprise timestamp information, GPS coordinates, and other technical information. One hypothesis supported by this facet could be that the capturing device was located at the specified coordinates at the given date and time.

We fix a consistent investigative knowledge base $(H, F, \mathsf{supports}, \mathsf{refutes})$ for the following definitions. The hypothesis-related relevance of a facet can then be defined as supporting or refuting a hypothesis.

Definition 3.3.3 (Relevance)**.** A facet $f \in F$ *is relevant to* a hypothesis $h \in H$ if f either supports or refutes h. Formally, this defines a relation $\mathsf{relevant} \subseteq F \times H$:

$$f \;\mathsf{relevant}\; h \coloneqq f \;\mathsf{supports}\; h \cup f \;\mathsf{refutes}\; h$$

Given a hypothesis $h \in H$ and a set of facets $F' \subseteq F$, we denote by $F'|_h$ the set of facets in F' that relate to h:

$$F'|_h \coloneqq \{f \in F' \mid f \;\mathsf{relevant}\; h\}$$

In other words, $F'|_h$ captures those facets that are relevant to $h \in H$, whereas the "$|$"-symbol is used to signify a filter-operation.

Similarly, when given a set $H' \subseteq H$, we denote all those facets that are relevant to the hypotheses in H' by writing $F'|_{H'}$ for the union of all $F'|_h$ for $h \in H'$:

$$F'|_{H'} \coloneqq \bigcup_{h \in H'} F'|_h$$

The notion of relevance gives rise to the *expressiveness* that can then be grasped intuitively as the set of supported or refuted hypotheses.

Definition 3.3.4 (Expressiveness)**.** For a facet $f \in F$ and a set of hypotheses $H' \subseteq H$, we denote the *expressiveness* by $H'|_f$, i.e., the set of hypotheses in H' that relate to f:

$$H'|_f \coloneqq \{h \in H' \mid f \;\mathsf{relevant}\; h\}$$

[13] Using set quantors, we can state an equivalent definition of the consistence of an investigative knowledge base: $\forall f \in F : \forall h \in H : \neg(f \;\mathsf{supports}\; h \cap f \;\mathsf{refutes}\; h)$.

[14] This classification is based on the explications of Jaquet-Chiffelle and Casey [131, p. 11, Fig. 1], visually depicted in Fig. 2.5.

In the same way as above, we define the *expressiveness of a set* $F' \subseteq F$ of facets, as follows:

$$H'|_{F'} \coloneqq \bigcup_{f \in F'} H'|_f$$

A quick way to internalize this notation is to interpret the "|"-symbol and its subscript as the filtering of the outer set via the relevant relation. E.g., the expressiveness $H|_F$ denotes all hypotheses, which are related to any facet in F. Similarly, the relevance $F|_H$ is the collection of facets, which are related to any hypothesis in H. Intuitively, relevance captures the notion of a facet contributing to the assessment of hypothesis h, and the expressiveness of f informs the investigator what hypotheses could be answered when a facet f is discovered.

Example 3.3.5. Continuing Example 3.3.2, we see that facet f denoting the JFIF data structures, including the header data, is relevant to a hypothesis revolving around the whereabout of the capturing device at the specified coordinates because of its support for that. Furthermore, we could add a second hypothesis to the expressiveness $H|_f$ of that facet f, since it might be used to perform source identification to identify the software stack of the capturing smartphone as well [168]. For instance, resorting to these data, one could additionally refute the hypothesis that the JPEG file in question was captured with the smartphone of the suspect. In this way, the same facet can be utilized to draw a richer picture than by only looking at time and location using the data fields in the *ExifIFD* and the *GPS IFD*.

3.3.1 Relation of Expressiveness and Relevance

In Definition 3.3.3, we expounded that the expressiveness $H|_f$ determines whether the facet could be of use to prove or disprove some hypotheses in H. Here, we observe an implication that if a facet f is relevant, then the hypothesis h must be necessarily part of the set $H|_f$ denoting its expressiveness.

$$f \text{ relevant } h \implies h \in H|_f$$

However, the contraposition gives more insight since it determines that if a hypothesis h is not part of $H|_f$, then it cannot be of relevance by any means.

$$h \notin H|_f \implies \neg(f \text{ relevant } h)$$

Therefore, the knowledge of the expressiveness of a facet enables the investigator to collect and analyze only relevant facets and discard others.

3.3.2 Derivation of Metrics

The availability of an investigative knowledge base (Definition 3.3.1) enables us to derive several metrics regarding facets and hypotheses; the knowledge of various properties of

facets and hypotheses seems advantageous when guiding an investigation. To quantify how expressive a specific facet is, we introduce the *expressiveness ratio*:

Definition 3.3.6 (Expressiveness ratio)**.** The *general expressiveness ratio* of a facet $f \in F$ is defined as the ratio of hypotheses that can be decided by discovering f in an investigation:[15]

$$\mathsf{expr}(\mathsf{f}) := \frac{|(H|_f)|}{|H|}$$

This definition of the general expressiveness ratio captures the intuition that the more conclusions we can draw from a facet, i.e., the larger the cardinality of the set of assessed hypotheses, the more expressive it is. Given that a facet could be very expressive, while not necessarily being helpful for answering a specific set of questions posed in a particular case, we also have the notion of the *relative expressiveness ratio* of a facet f, which takes only the actually case-relevant hypotheses $\widetilde{H} \subseteq H$ into account:

$$\mathsf{expr}_{\widetilde{\mathsf{H}}}(\mathsf{f}) := \frac{|(\widetilde{H}|_f)|}{|\widetilde{H}|}$$

Investigators seeking to decide on the hypotheses $\widetilde{H}$ can limit their analysis to those facets that have a relative expressiveness ratio greater than zero w.r.t. $\widetilde{H}$ given that all other facets are not relevant to the assessment of any hypotheses from $\widetilde{H}$. This is the case because given a facet $f \in F$ and a hypothesis $h \in \widetilde{H}$ such that f is relevant to h, we know that by Definition 3.3.4 $h \in \widetilde{H}|_f$ and thus $\mathsf{expr}_{\widetilde{\mathsf{H}}}(\mathsf{f}) > 0$.[16]

Note that the relative expressiveness ratio w.r.t. the complete set of all possible hypotheses $\mathsf{expr}_{\mathsf{H}}(\mathsf{f})$ is exactly $\mathsf{expr}(\mathsf{f})$, which reveals that the former is indeed a special case of the latter. Additionally, one can put the hypotheses at the center of the consideration and determine their specificity to describe how easy or hard it is to find facets to assess those.

Definition 3.3.7 (Specificity ratio)**.** The *specificity ratio* of a hypothesis $h \in H$ is defined as the ratio of facets that can decide h by their assessment:

$$\mathsf{spec}(\mathsf{h}) := 1 - \frac{|(F|_h)|}{|F|}$$

This definition denotes the *general* case: The more facets are relevant to a hypothesis, i.e., the larger the cardinality of the set of contributing facets in the numerator, the less specific the hypothesis. Note that we see no sensible application of a *relative* metric regarding a restricted set of facets, i.e., a collection of facets, since this would merely state how scarce the investigators were by collecting facets in that instance.

[15] Note that $\mathsf{expr}(\mathsf{f})$ is neither defined for an empty nor for an infinite set H of hypotheses since these cases are practically neglectable.

[16] The metric assumes that $\widetilde{H}$ is finite; however, in practice, an investigator will not seek to decide an infinite amount of hypotheses.

To conclude this section, we sum up the intuitions of the established concepts, i.e.:

Relevance
is given if a facet supports or refutes a hypothesis.
Expressiveness
denotes the set of hypotheses that can be assessed by the facet in question.
Specificity
determines how many facets can be used to assess a specific hypothesis in question.

3.4 Relation to Reliability

Relevance of facets is one side of the coin; the other is reliability. Recalling the definitions of digital evidence in Section 2.4.3.2, we already learned that it is not only required that the information is directly or indirectly linked to the investigative questions, but it also needs to be reliable at the same time in order to be considered evidential [31, 111, 165, 208]. Since the reliability of digital evidence is commonly assessed indirectly by determining whether the traces are authentic, accurate, and complete [111, 165, 208], we now turn our attention to this second necessity, scrutinize the relation of the notion of expressiveness and relevance with reliability, and formally express *accuracy* and *completeness*. However, since the set-theoretic approach presented before does not depict the provenance of facets, we will not address the third feature, authenticity, here but integrate it later with the already established CSI model by Freiling and Sack [84] to trace provenance and thus establish authenticity as well (Section 3.6.2).

3.4.1 A Formal Notion of Accuracy

Accuracy demands that the quality of the employed procedures (starting at collection and ending with the introduction into court) is "[f]ree from any reasonable doubt about the quality" [208, p. 139] and that "amount of error should be acceptable in the context of the current investigation" [111, p. 66]. In practice, investigators might make errors—either during acquisition, processing, analysis, or even documentation—, which could ultimately lead to a false interpretation of the facets at hand. To model this attribute of reliability, we define the notion of *accuracy* of an investigative knowledge base w.r.t. another investigative knowledge base capturing the true interpretation of facets.

Definition 3.4.1 (Accuracy of a knowledge base)**.** Given an investigative knowledge base $(H_\top, F_\top, \mathsf{supports}_\top, \mathsf{refutes}_\top)$ representing ground truth, we consider another investigative knowledge base $(H, F, \mathsf{supports}, \mathsf{refutes})$, where $H \subseteq H_\top$ and $F \subseteq F_\top$, to be *accurate* when both the supports and the refutes relation are included in the $\mathsf{supports}_\top$ and $\mathsf{refutes}_\top$ relation, respectively.

$$\mathsf{supports} \subseteq \mathsf{supports}_\top$$
$$\mathsf{refutes} \subseteq \mathsf{refutes}_\top$$

Thus, accuracy is achieved, if the investigators' conceptions are consistent with the reality and they thus draw the correct conclusions. Note, however, that accuracy does not prevent the investigator from overlooking facets or their relevance but just prevents drawing incorrect conclusions.

3.4.2 A Formal Notion of Completeness

In terms of reliability, the collection of evidence in a case is considered to be complete, when "the maximum amount of digital evidence relevant to the investigation" was preserved, while it is also possible to assess what information is lost [111, p. 65]. Starting with this statement, we define an evidence collection to describe two different completeness properties.

Definition 3.4.2 (Evidence collection)**.** An *evidence collection* is a tuple $(F_{\mathsf{found}}, F_{\mathsf{not}})$ consisting of a set F_{found} of *found facets* that have been successfully collected and a disjoint set F_{not} of *unretrievable facets* that could not be collected.

The maximal notion of completeness, loosely matching the intuition of Hargreaves [111, p. 65], can then be captured by what we call *exhaustive completeness*.

Definition 3.4.3 (Exhaustive completeness)**.** An evidence collection $(F_{\mathsf{found}}, F_{\mathsf{not}})$ is *exhaustively complete w.r.t. a set H of hypotheses* when all facets relevant to H have either been recovered or could not be recovered.

$$(F_{\mathsf{found}}, F_{\mathsf{not}})\ \mathsf{complete}_{\mathsf{E}}\ H \iff F|_H \subseteq (F_{\mathsf{found}} \cup F_{\mathsf{not}})$$

An exhaustively complete evidence collection contains every possible facet that is relevant to the case regardless of whether fewer facets would have sufficed to conclude the case or whether it could not be solved at all, since every facet potentially contributing to the assessment of hypotheses has been regarded. To capture the notion of being sufficient to decide all hypotheses of the case, we define *decisive completeness*.

Definition 3.4.4 (Decisive completeness)**.** An evidence collection $(F_{\mathsf{found}}, F_{\mathsf{not}})$ is *decisively complete w.r.t. a set H of hypotheses* when for each hypothesis in H a relevant facet is contained in the set F_{found}.

$$(F_{\mathsf{found}}, F_{\mathsf{not}})\ \mathsf{complete}_{\mathsf{D}}\ H \iff H \subseteq H|_{F_{\mathsf{found}}}$$

In other words, this means that the hypotheses H in question is a subset of all the hypotheses assessable by the facets in F_{found}. There might be cases where decisive completeness cannot be achieved, since facets listed in the knowledge base are not available or, for whatever reason, could not be recovered. To account for those cases, decisive completeness can not be used as the only stopping criterion for evidence collection. We therefore define a combined stopping criterion that we call *completeness*.

Definition 3.4.5 (Completeness)**.** An evidence collection $(F_{\mathsf{found}}, F_{\mathsf{not}})$ is *complete w.r.t. a set H of hypotheses* when $(F_{\mathsf{found}}, F_{\mathsf{not}})$ is decisively complete or exhaustively complete w.r.t. H.

$$\mathsf{complete} := \mathsf{complete}_{\mathsf{D}} \cup \mathsf{complete}_{\mathsf{E}}$$

The intuition behind this definition of completeness is to stop the collection when *enough* facets have been recovered to assess all case-related hypotheses (decisively complete) but try to collect *all* facets relevant to the hypotheses if no subset suffices (exhaustively complete).

3.4.3 Insights into the Reliability of Digital Forensic Science

Modeling the investigative knowledge and the associated properties in the way presented in the previous section enables us additionally to underpin more universal insights into the reliability of DF fieldwork.

It is well-known that certain areas of digital forensics can be considered precisely fathomed. Other areas, however, are poorly understood, resulting in fewer findings and less accurate analyses. Examples of the first category are the analysis of long-established and thoroughly understood file systems, like NTFS and ext4, or the examination of systems used for the online distribution of child sexual abuse material—a criminal mass phenomenon with a long-standing history. Examples of the second category could be the analysis of new and highly sophisticated exploits or very recent file systems. This observation poses the question of how to discriminate between those two areas. Furthermore, it is unclear which factors are involved in such discrimination. In these questions, modeling investigative knowledge, accuracy, and completeness in the formal sense comes in helpful.

It is quite obvious that categorizing one of these areas depends on the scope of the available knowledge base (Definition 3.3.1) and its accuracy (Definition 3.4.1) from which the examiner can draw. For a simplified analogy imagine a good expert witness here, whose amount of experience constitutes the amount of known facets as well as the supportsand refutesrelations. Analyses based on a "good" knowledge base—regardless of whether it is implemented in a tool or whether it exists only in the investigator's head or an encyclopedia—will likely produce reliable, i.e., highly accurate and decisively complete results. The proposition here is that the quality of the knowledge base regarding a specific investigation is primarily related to the degree of understanding of the programs' behaviors, which were responsible for creating the traces in question in the first place that are encountered in that investigation. To deduce that, think of a program in the broadest sense (including user applications and system software like filesystem drivers) that is poorly understood. When dealing with traces of such a program, examiners might be unable to answer the most basic questions because they do not know where to look for and how to interpret the findings. Expressed in another way, the examiners arrive at exhaustive completeness too soon in such a case. Additionally, their interpretations of the findings might not be accurate as well. In that case, the examiners might reach decisive completeness, however, they draw wrong conclusions since their understanding of the

facet interpretation is flawed. On the contrary, if a program and its inner workings are well understood because it is either open-source and precisely specified or has been analyzed with various reverse engineering methods, such as black box testing as well as dynamic and static code analysis, in the course of several digital forensic investigations, then examiners know which action results in which digital tangible trace. In such a case, they are able to collect the observable facets of it and can infer their meaning for past events. When the meaning of each detail, e.g., some bitflags in some database field, is well-known, they can precisely infer which facet to expect from which action and vice versa. Put differently, the investigators have a complete understanding of the expressiveness of the facets encountered. In essence, we argue that an in-depth understanding of a program's behavior opens up building a knowledge base that is well-filled, detailed, and accurate in regard to facets of this particular program P. Concretely, it means that the cardinality of the set of facets related to the program $F_P \subseteq F$ contained in the knowledge base is rather large, and the relations supports $\subseteq F_P \times H$ and refutes $\subseteq F_P \times H$, relating facets of this particular program to hypotheses, are both accurate and filled with many facet-hypotheses-pairs. Effectively, such a knowledge base enables the examiners to arrive at a decisively complete evidence collection without misinterpretations more probably because they know where to look at, interpret the observations well, reason precisely about hypotheses, and reconstruct the most likely course of events. Ultimately, we can indirectly conclude from this observation that the degree of program understanding determines the division into reasonably and less well-understood areas of digital forensics.

3.5 A Conceptual Network of Relevant and Expressive Evidence

After the formal effort in the previous section, we temporarily step back and turn our heads to the intuitive interrelation of the terms and describe how they fit into the investigative core of criminalistics, as it has been presented and discussed in Section 2.4. To do so, we will briefly revisit these and put them piecemealy into context by presenting and verbalizing a conceptual network shown in Fig. 3.1, i.e., a hypergraph explicating the connections between the terms of digital evidence. So, this section aims to serve as a short and refreshing yet clarifying interlude before we show the actual application of the new understanding.

When investigators arrive at a crime scene—be it solely physical, solely digital, or a combination—they are tasked with building a retrodictive model of the events related to the deed at hand that happened in the past. To tackle this task, they will resort to the means and techniques of traceology. Based on Locard's foundational discovery mentioned several times, events will lead to a transfer of matter or traits and, hence, result in traces, i.e., the distinguishable difference between the present tangible world and a hypothetical one in which the event did not take place. Employing some observation instrument, they begin to observe facets as the perceivable parts of tangible traces. Those facets might carry more or less information. The notion of facets as a "vector of information" [158, p. 33] links facets to specific hypotheses since the information content is characterized by the possibility of assessing hypotheses linked to the event that ultimately created the tangible trace from which the facet was taken. If such a link can be established, the facet exhibits hypothesis-related

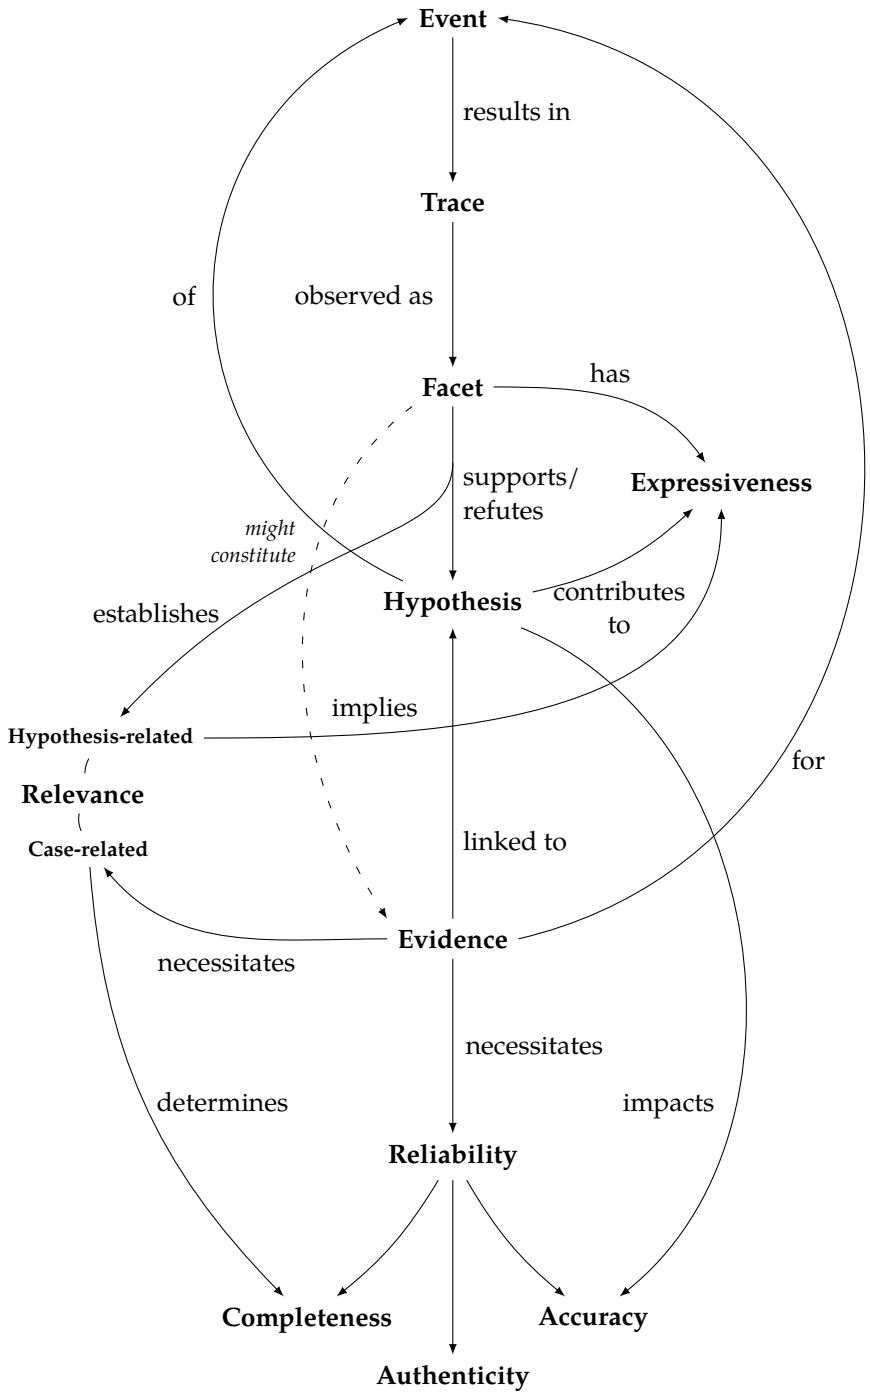

Figure 3.1: Conceptual network illustrating the relation of the terms *expressiveness* and *relevance* to reliable evidence. Events produce traces, which can be observed by examiners as facets. Facets can become evidence for an event, which is indicated by the dashed connection, if they support or refute a hypothesis that contributes to deciding on the matter of the specific case given their reliability, i.e., their authenticity, accuracy, and the completeness of the evidence collection.

relevance. Facets that tend to support or refute multiple (non-conflicting) hypotheses are considered to be more expressive. Investigators iteratively search for expressive facets tending to prove or refute their hypotheses [240, p. 93]. It then constitutes evidence for the hypothesis in question. However, evidence necessitates that the assessed hypothesis exhibits case-related relevance, i.e., the hypothesis is concerned with some event related to the deed. In that case, facets form evidence of a hypothesis, as indicated by the dashed edge in Fig. 3.1. Note that this reflects a preliminary, investigative perspective employed by law enforcement officers, crime scene examiners, and so on since the trier of fact, who is involved a lot later in the process though, might judge the case-related relevance differently. Adititonally, evidence needs to be reliable: This means that the provenance of a facet must be tractable to the source so that it can be considered authentic. The examination and interpretation of the facet in relation to the assessed hypothesis have to be within certain error bounds in order to be accurate. Lastly, the collection of facets has to be complete regarding the case-related hypotheses, i.e., the different imaginable versions of the deed.

By guiding the look to the bigger picture, we aimed to clarify how these formal notions fit into the tradition of criminalistic reasoning and solidify the intuitions. The overview provided by the conceptual network now serves as a starting point to turn to the question of how our formal notions can be actually applied in practice.

3.6 Application of the Concepts

Having both gained a rigorous understanding and placed the terms into a conceptual network, we now bring the formal description to practice in this section. We see three different possibilities to do so: a procedural application, a merely conceptual application, and a material application. First, we extend an existing criminalistic thinking model using our formalized notion of expressiveness, relevance, and completeness in Section 3.6.1. Second, we show how it can be integrated into the CSI model by Freiling and Sack [84] in Section 3.6.2. Third, we tangibly demonstrate the concrete calculation of expressiveness using a model of a digital system. However, since we build up on the concepts of sufficient and necessary evidence as discussed in Chapter 4, we will introduce this type of application there after all needed prerequisites have been conveyed (Section 4.8.2.2).

3.6.1 Integration into the Criminalistic Cycle

Investigations are commonly modeled as iterative processes [213, p. 153]. However, since general process models aim to be universally valid, they necessarily remain rather vague [103]. Hence, we show how we could incorporate and thereby improve the preciseness of one such universal model by using our formalized notions of relevance, expressiveness, and completeness.

One convincing instance of those thinking models is the *Criminalistic Cycle* proposed by the Suisse criminal jurists Walder and Hansjakob [240], who emphasized the procedural

aspect of the investigative task. In Section 2.4.3.1, we already described their modeling of the criminalistic thinking process. While the approach by Walder and Hansjakob [240] is quite persuasive, the single steps, however, appear to be rather imprecise. Given the formal notions of facets' expressiveness and relevance, we are able to nail the steps down and extend the thinking model to create an improved version, which we call the *Facet-oriented Criminalistic Cycle* (*FoCC*), as illustrated in Fig. 3.2.

To keep track of the investigative effort when processing a case, we introduce the notion of an investigative system:

Definition 3.6.1 (Investigative system)**.** An *investigative system* $(KB, \widetilde{H}, E)$ is comprised of the investigators' knowledge base *KB*, a subset $\widetilde{H}$ of the hypotheses in *KB* denoting *case-related hypotheses*, and an *evidence collection* E.

The investigators' knowledge base guides the investigation and leads to iterative updates of the set of case-related hypotheses and continually adds facets to the evidence collection structure until the investigator determines the investigation to be complete. Here, the major difference to the original and rather coarse-grained version of Walder and Hansjakob [240, p. 93] is the introduction of formal notions, an explicit termination condition, and the more precise design of the individual steps.[17] As a first step, the investigators analyze the initially available facets, e.g., some witness testimony. Based on these, a set of case-related hypotheses $\widetilde{H}$ is derived. The follow-up step in the original model by Walder and Hansjakob, i.e., determining the next investigative actions, can be split into two steps: It has to be checked if the evidence collection is considered complete, i.e., decisively complete or, alternatively, exhaustively complete. Both cases lead to termination since the case is either solved or unsolvable at all, otherwise the cycle is continued and follow-up iterations have to be conducted until completeness is achieved.

If the evidence collection is not complete yet, the set of unanswered hypotheses $\widetilde{H}_U$ has to be assessed, which can be formed by removing those hypotheses from the set of case-related hypotheses $\widetilde{H}$ that can be answered by the already available facets:

$$\widetilde{H}_U = \widetilde{H} \setminus H|_{F_{\mathsf{found}}}$$

Then, the investigator picks a (non-empty) set F_C of facets to collect in this iteration from the missing facets.

$$F_C \subseteq (F|_{\widetilde{H}_U} \setminus F_{\mathsf{not}})$$

Note that this also excludes found facets as those would have answered any hypothesis they are relevant for and hence be excluded from $F|_{\widetilde{H}_U}$, which we have already defined in Definition 3.4.2. Here, several strategies could be employed: For instance, one could opt to determine the smallest set of facets by picking those that would make the evidence collection decisively complete w.r.t. $\widetilde{H}_U$ if all were to be found. Alternatively, one might (inefficiently) choose to consider all facets relevant to the set of undecided hypotheses $\widetilde{H}_U$, i.e., pick the whole set $F|_{\widetilde{H}_U} \setminus F_{\mathsf{not}}$.

[17] As stated in Section 3.1, we want to remind that the quest to find *case-relevant* hypotheses is out of scope in this chapter.

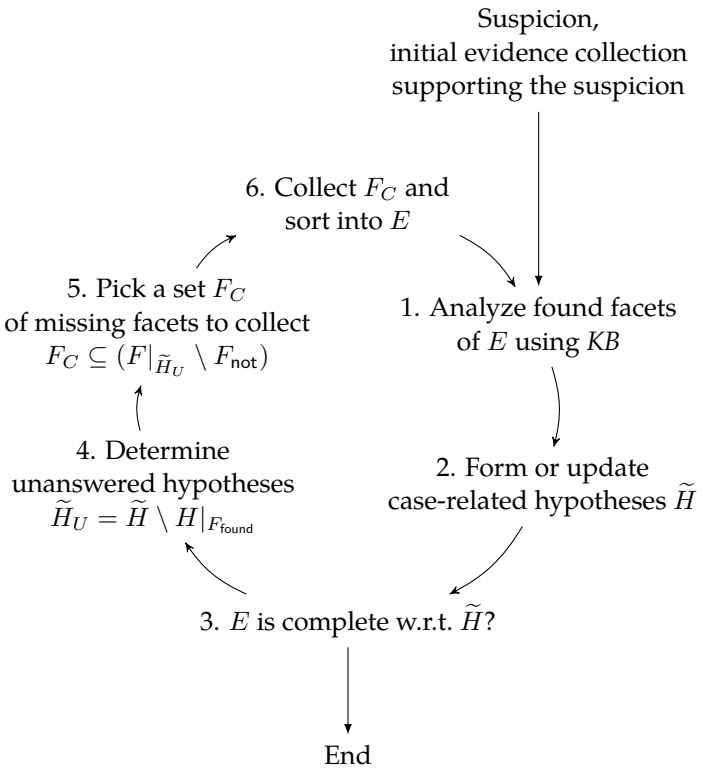

Figure 3.2: The Facet-oriented Criminalistic Cycle (inspired by the original version of Walder and Hansjakob [240, p. 93]). The investigation is initiated with some suspicion and an initial set of found facets as part of an evidence collection. Based on these, the first step is to analyze the found facets (step 1). The findings enable the investigators to derive case-related hypotheses (step 2). The next step is to scrutinize completeness, i.e., decisive completeness that is achieved if all hypotheses can already be assessed by the found facets, or exhaustive completeness, which is, in turn, achieved if all facets relevant to the hypotheses, both found or not recoverable, have already been considered (step 3). If neither of these properties is fulfilled, then the next step is to determine the unanswered hypotheses (step 4), which are then used to pick a set of missing facets (step 5). Those facets are collected and added to the sets of found facets or unretrievable facets, respectively (step 6), before another iteration of the cycle is started.

Once the missing facets have been determined, all of them or a subset F_C can be collected to extend the evidence collection by sorting into the sets of found facets and unretrievable facets depending on whether they were collected or could not be collected. In view of an updated evidence collection E, the investigators have to assess these newly collected facets and, in the process, spawn another iteration of the cycle with possibly new hypotheses $\widetilde{H}$.

By putting the Criminalistic Cycle on a formal basis, we improve the preciseness of the steps and incorporate an exact termination condition. Furthermore, it becomes obvious that the result of the investigation is not only dependent on the available facets but also on the accuracy and consistency of the investigators' knowledge base as those guide the collection of evidence and determine when the investigator will stop collecting further evidence. In principle, the *FoCC* can incorporate updates of the investigators' knowledge base at any point in the cycle to reflect learning during the investigation by simply replacing the used knowledge base with a new, updated one. Note that even in the presence of an arbitrarily changing knowledge base the evidence collection will grow in a monotone fashion leading to a terminating investigation.[18]

These improvements aim to demonstrate the usefulness of the newly introduced concepts to provide procedural clarity, while we opted for a simplistic yet precise formalization of the cycle.

3.6.2 Integration into the CSI Model

As outlined in Section 2.4.3.2, the CSI model describes the components and the lifecycle of pieces of evidence. During the initial presentation of the reasoning by Freiling and Sack behind this model, it has already been indicated that the components *claim* and *support* are solid and tangible but the component *information* appears rather elusive. This elusiveness can be considered both an impediment and a striking feature at the same time since that vagueness is required for describing relationships between pieces of evidence [84, p. 330] and let the model remain open for the adaptability and extensibility of novel methods to look at and work with traces.

Clarification of The *Information* Component. To illustrate the nature of this component, Freiling and Sack referred to a knife, which served as a weapon of crime in a homicide and contains blood accumulation, as an example. *Information* in the sense of to the CSI model could then be grasped according to Freiling and Sack [84, p. 329, translated by the author] as follows:

- The length and width of the knife to compare it with the victim's puncture marks,
- DNA information of the adhering blood, in order to be able to assign it to a person,

[18] This assumes that a finite amount of facets suffices to conclude the investigation.

- the assignment of the knife to a person, which can be derived from a fingerprint, if available, or
- the location where the knife was found.

Sack [199, p. 59] underlined that the *information* is about the interpretation of the features of the object in the investigators' minds by setting those into relation to the course of action, which leads to the distinction of *general* and *relevant information* contained by the *support* of the piece of evidence.

Thinking about these statements, we see potential to integrate the thoughts of Freiling and Sack expressed in natural language into our formal concepts to refine the latently existing intuition in there: The features of the object, as Freiling and Sack [84] put it, correspond to facets as defined by Jaquet-Chiffelle and Casey [131], i.e., the results of the observation and analysis of the *support* to gain or extract certain *information* stemming from a tangible trace. Furthermore, we see that the differentation into *general* and *relevant information* is strictly dependent on the case. So, if a certain feature of the piece of evidence can help in assessing a case-related proposition, it constitutes *relevant information*. Put differently, a facet exhibiting hypothesis-related relevance for a case-related hypothesis corresponds to the understanding of the component *information* of a piece of evidence in the CSI model.

Having an investigative knowledge base $(H, F, \mathsf{supports}, \mathsf{refutes})$ at hand, we can model this component of the CSI model precisely yet remain adaptable and extensible: To do so, we refer to the set of facets that can be directly observed or derived by further processing from the piece of evidence P in question as F_P and the set of case-related hypotheses as $\widetilde{H}$. Hence, we can grasp the (case-dependent) *information* component of the piece of evidence as

$$I_P = \{(f, h) \in F_P \times \widetilde{H} \mid f \ \mathsf{relevant}\ h\},$$

which constitutes a compact yet precise notation. It is important to note that both the investigative knowledge base and the set of case-related hypotheses $\widetilde{H}$ can and will be updated over time. The former is necessary when novel methods of examination and analysis are developed that can be used to derive facets of a new kind, the latter happens, when new findings are available. Both adaptions might cause the *information* component I_P of the piece of evidence to change, so the proposed formalization retains that originally intended feature of adaptability and extensibility.

We now illustrate the application of the rigorous concretization using the original example of the bloody knife as provided by Freiling and Sack.

Example 3.6.2. We begin with rephrasing the above list concerning the "features" of the bloody knife:

- Measuring the dimensions of the knife yields a facet f_{Dim}, which is relevant to the hypothesis h_1 that the victim's puncture marks have been caused by that knife.
- Employing the polymerase chain reaction method to gather a DNA-profile from blood adhering to the knife yields a facet f_{DNA}, which is relevant to

the hypothesis h_2 that the knife has been in contact with the blood of a specific person.

- Collecting fingerprints using fingerprint powder and a foil yields a facet f_{FP}, which is relevant to the hypothesis h_3 that the suspect held the knife.
- Documenting the location where the knife was laying around photographically yields a facet f_{Loc}, which is relevant to the hypothesis h_4 that the suspect's escape route went through that location.

So, we name specific facets, phrase (more or less) precise hypotheses, and define that the respective facet is relevant to a hypothesis by either supporting or refuting it. It is rather obvious to see that this information corresponds to (an excerpt of) the investigative knowledge base since it depicts the investigators' conception of "bloody knives". In the given example, the set of facets that can be derived from the concrete instance knife of a bloody knife, which constitutes the piece of evidence in question, can be described as

$$F_{\mathsf{knife}} = \{f_{\mathsf{Dim}}, f_{\mathsf{DNA}}, f_{\mathsf{FP}}, f_{\mathsf{Loc}}, \dots\}\,.$$

The exemplary set of case-related hypotheses in this homicide is

$$\widetilde{H} = \{h_1, h_2, h_3, h_4, \dots\}\,.$$

Having the knowledge base as well as these two sets available, lets us then infer the relevant information of the bloody knife

$$\begin{aligned} I_{\mathsf{knife}} &= \{(f,h) \in F_{\mathsf{knife}} \times \widetilde{H} \mid f \text{ relevant } h\} \\ &= \{(F_{\mathsf{Dim}}, h_1), (F_{\mathsf{DNA}}, h_2), (F_{\mathsf{FP}}, h_3), (F_{\mathsf{Loc}}, h_4), \dots\}\,. \end{aligned}$$

Scrutinizing the CSI model in view of the newly developed formal notions of relevance and expressiveness has led to a more precise understanding of the *information* component of a piece of evidence P. It becomes apparent that the relevance relation relevant, which is taken from an investigative knowledge base, determines the general information content of the piece of evidence given the current state of knowledge. The case-related hypotheses $\widetilde{H}$ can then be used to distill the relevant information I_P. This means that the application of the formal notions helps to explicate and, hence, better grasp the question of what actually makes up a piece of evidence.

Potential to Trace Provenance. By integrating the set-theoretic notions of relevance and expressiveness, which are based on an investigative knowledge base (Definitions 3.3.1, 3.3.3 and 3.3.4) into the CSI model, we conceive reciprocal benefits: We achieve not only a sharpening of the intuition of the CSI model but gain the means to trace provenance as well. When we defined the reliability properties (Section 3.4), we could only define accuracy and completeness but were unable to formalize the third reliability property, i.e., authenticity, because we do not model the provenance of the facets in an evidence collection.

Freiling and Sack model a directed graph, where the set of all potential pieces of evidence, i.e., (C, S, I) triples, form the vertices and the transformations from one piece to another

denote the edges. Hence, they could trace the provenance of a piece of evidence and define authenticity as well as integrity by looking at the triple's transformations that are depicted by the edges of a path through the directed graph. Deviating from the authenticity notions of Sommer [208] and Hargreaves [111] (Section 2.4.3.2), they consider a piece of evidence to be *authentic*, on the one hand, if the *support*-component originates immediately from the seizure, i.e., $s_n = s_1$. On the other, they consider it to be *integer*, if the relevant information remains preserved, which constitutes a definition that agrees with the concept of integrity in cryptography [84, pp. 323 ff.]. Therefore, their understanding of authenticity implies the integrity of a piece of evidence, while both attributes are based on the correctness of the claim [84, p. 336].

Comparing this approach with the understanding of Sommer [208] and Hargreaves [111], it becomes obvious that the conception of authenticity by these two authors is less restricting than the sole focus on the *support* component of Freiling and Sack, thus, it is more related to general traceability of the provenance, i.e., the correctness of the established link of the claim and the information; integrity, as understood by Freiling and Sack [84], provides a more precise expression of that traceability, while it additionally might be helpful in identifying (specific) violations of accuracy and completeness properties. It seems that they strictly had images of hard disk drives, solid state drives, memory cards, USB sticks, and the like in mind, when they defined integrity. Both Freiling and Sack [84] as well as Sack [199], however, leave open how the reader has to understand and apply the term in case of physical evidence.

Since we showed that the *information* component contains pairs of facets and hypotheses, we gain the possibility to describe authenticity (and integrity) of pieces of evidence—hence also facets—in the strict notions of Freiling and Sack [84] or in the more lenient versions of Sommer [208] or Hargreaves [111], if required. The latter is needed because neither authenticity nor integrity as defined by Freiling and Sack [84] is applicable, when one thinks about derived (C, S, I) triples, as presented in the Examples 2.4.5 and 3.6.2, since those would be neither authentic due to the changing support nor integer due to the respectively extracted information.

With the goal of defining the authenticity of derived (C, S, I) triples, we model a directed acyclic graph of pieces of evidence in the spirit of Freiling and Sack [84] using a slightly refined formalism to be able to express certain attributes in a more precise way:

Definition 3.6.3 (Evidence graph)**.** An *evidence graph* (P, E) consists of a finite set P forming the vertices representing pieces of evidence and a set $E \subseteq P \times P$ of edges, which is non-reflexive and loop-free, representing tansformations of pieces of evidence.

We project the (C, S, I) triples to the nodes in the graph (P, E) and access the respective component using the "."-notation. Using this definition, we can express authenticity in the notion of Freiling and Sack [84, p. 336] in a more precise version, as follows:

Definition 3.6.4 (Strict authenticity)**.** A piece of evidence p_n is *strictly authentic* iff the path $P = (p_1, \ldots, p_n)$ in the evidence graph (P, E) is known in its entirety, all claims of the

pieces of evidence along the path $p_i.C$ are correct, and the support of the original piece of evidence p_1 and the piece of evidence p_n remain unchanged, i.e., $p_1.S = p_n.S$.

So, a piece of evidence is only strictly authentic if the support stems directly from the seizure. Contrary to the original definition of Freiling and Sack, we include the correctness of the claim and the knowledge of the whole path in the definition itself, which they considered a general prerequisite. However, it is obvious that strict authenticity is often not achievable if we deal with derived pieces of evidence—both digital, such as selective data extractions, or physical. Hence, we define a lenient version of authenticity.

Definition 3.6.5 (Lenient authenticity)**.** A piece of evidence p_n is *leniently authentic* iff the path $P = (p_1, \ldots, p_n)$ in the evidence graph (P, E) is known in its entirety, all claims of the pieces of evidence along the path $p_i.C$ are correct, and the relevant information $p_n.I$ has already been contained in the original piece of evidence $p_1.I$, i.e., $p_1.I \supseteq p_n.I$.

By verifying lenient authenticity of a piece of evidence, we can infer that the facets contained in the *information* component are leniently authentic as well and thus gain means to verify the authenticity of facets to complete the formalization of the attributes constituting reliability.

3.7 Discussion

In this section, we differentiate our proposed concepts from related work, address the limitations of the formalism, and derive implications of the gained understanding.

3.7.1 Differentation from Related Work

Reminiscing the beginnings of criminalistics, relevant evidence was described by the pioneers, such as Gross and Geerds [98] and Kirk [141], in a concrete manner by looking at certain traces and their meaning. Legal scholars reasoned in a more abstract way about relevance and put it in direct connection to disputing propositions [165], which is often intertwined with specific questions, such as probative value, reliability, and also admissibility. The rather recent criminalistic perspective on the matter by Hazard [114] is far broader and includes the subjective perception of informative value, which might be based on experiental knowledge.

Based on these thoughts, we looked at relevance and expressiveness from a different angle: The present chapter provides a formal basis for investigative knowledge bases, which have remained informal and intuitive in previous works. By using a formal approach, we try to lift it from a practical orientation containing certain vagueness to a rigorous formal notion based on a set-theoretic approach. Indeed, the formal description of the

concept of expressiveness and relevance enabled us to explicate insights into the nature of digital evidence in form of a conceptual network that have remained implicit so far. Furthermore, we relate relevance to the comparably important concepts of accuracy and completeness as part of reliability, which would have remained more opaque if expressed only in natural language. Having the rigorous notions available helps us grasp and describe the criminalistic thinking process in far more detail. We showed this by improving two instances of such thinking processes: On the one hand, the extension of the *Criminalistic Cycle*, the proposed FoCC, can point out which facets to collect next in order to solve the case and introduces an exact termination criterion. On the other hand, the further formalization of one essential component of the CSI model has provided means to grasp and generally describe the information in a piece of evidence and model its authenticity.

3.7.2 Differentiation from Probabilistic Reasoning in Digital Investigations

In traditional branches of forensic science, comparative analyses of trace and reference material are the most common types of examinations [195, p. 104]. Forensic examiners, especially in the fields of forensic genetics [211], forensic anthropology [16], and others, employ likelihood ratios to communicate their judgement of the value of evidence. Likelihood ratios are statistical tools to quantify the value of evidence by comparing the likelihood of observed evidence under two competing hypotheses: one favoring the prosecution's case and another favoring the explanation provided by the defense. They act as a numerical measure of the evidential value, helping the trier of fact to make a decision based on the Bayes' theorem by updating their prior odds with their own likelihood ratio (derived from the results presented by the expert witness) to determine the posterior odds [92, p. 3].

While reporting the value of digital evidence in the form of likelihood is (at least to date) uncommon, there have been several efforts to quantify probabilities linked to investigative questions for digital evidence. As one of the first, Kwan et al. [144] proposed to use Bayesian belief networks for reasoning about investigative hypotheses since such a computational representation, which is based on the Bayes theorem at its core, provides a "useful formalism for quantifying and propagating the strengths of investigative hypotheses and supporting evidence" [144, p. 287]. However, a major issue is the subjectivity when assigning probabilities that are used to relate hypotheses and certain findings. Building upon this work, Tse et al. [225] proposed two methods to assess the validity of the Bayesian belief networks of Kwan et al. [144].

In addition, Overill and Silomon [174] reviewed two approaches for quantifying the plausibility of digital traces: On the one hand, they looked at the quantification of the "probability of recovering the evidential traces E given that the hypothesis H is correct"—$Pr(E|H)$. On the other hand, they considered the probability that "the hypothesis H is correct given that the evidential traces E have been recovered"—$Pr(H|E)$ [174, p. 2]. Both measures have to deal with subjectivity and the quest for deriving apt a priori probabilities again; they necessitate some model to estimate the probability of observing the traces, e.g., the operational complexity model, which has been introduced by Overill et al. [175]. To show its usefulness, Overill and Silomon [173] presented the odds of the so-called Trojan horse

defense as explanations for recovered digital evidential traces in five common e-crime scenarios and additionally applied it to assess the defendant's statement of accidentally downloading child pornography in cases where relatively small numbers of child sexual abuse material (CSAM) images were found in a larger corpus of adult pornography images [176]. While the previous works dealt with the correctness of hypotheses, Overill and Chow [172] discussed the relative weight of single evidence items depicted in a Bayesian belief network. By determining items that have a high impact on the probability of the root hypotheses in that network, the investigator could prioritize the collection and examination of pieces that tend to have a favorable Return-on-Investment and Cost-Benefit-Ratio for the overall investigative goal.

However, none of the above approaches appears to be immediately and universally applicable since they require some kind of probability model, which is hard to acquire in practice. Actually, we consider the task of determining the connections, the respective probabilities between observed traces, and the conclusions to be very complex, which is underlined by the fact that it has not been tackled by the previously discussed approaches, which just estimated the probabilities by conducting expert interviews. Additionally, none of these approaches explicitly considered the question of expressiveness. However, Overill and Chow [172] illustrated how to infer the relative weight of single evidence items in a Bayesian belief network aiming to optimize the Cost-Benefit-Ratio of evidence collection. This constituted a first leap toward assessing the influence of facets on a given hypothesis—a feature that is inherently encoded in the Bayesian belief network as a probability.

In the present chapter, we chose a different route mainly for two reasons: First, we see the difficulty in developing and establishing probabilistic models when dealing with digital evidence, which is implicitly backed by the above-mentioned research. Compared to other forensic disciplines, such as DNA analysis or fingerprint comparison, no reliable statistical model has been established so far, making it challenging to assign sensible likelihood ratios consistently and reliably. Second, there seems to be a general complexity to employing the Bayesian method for arbitrary evidence outside of the comparative examination of a trace and reference specimen. Hence, we formally address the relevance of facets in regard to the hypotheses for which they are pertinent. This still contributes to a clear yet universal understanding, which can be used to unambiguously reason about criminalistic processes (Section 3.6.1) and characteristics of evidence (Section 3.6.2).

3.7.3 Limitations

Given the related work, and especially the probabilistic methods presented above, there is a need to discuss the limitations of our proposed concept of relevance and expressiveness.

First of all, we acknowledge the abstract nature of the formalization. Some may argue that this constitutes a strong point since it enables reasoning about facets' qualities without distractions, aiming to fuel discussions on collecting and assessing them. Others may consider it a drawback because it merely provides a foundation for reasoning but has no immediate applicability on its own, which is an argument that we revisit in Chapters 4 and 5.

Furthermore, the definitions in Section 3.3 and the considerations for applying the concept dealt only with strictly boolean cases. Hence, the concept provides no means to express uncertainty, i.e., to model situations where one facet supports a hypothesis to a certain degree while another refutes it. However, we see the potential to extend the proposed formalization in order to express probabilities. To do so, the supports and refutes relations could be defined as fuzzy relations so that facets get related to hypotheses with grades from the unit interval. Beyond that, a fitting fuzzy semantic for the used connectives can also be picked to adapt the remaining definitions. For the defined ratios, we envision tallying up the containment grades to achieve sensible measures in fuzzy cases. Given that all investigations and analyses almost always have to deal with a certain degree of uncertainty, we consider such an extension to be vastly helpful.

3.8 Summary

To a large degree, digital investigations aim toward the reconstruction of past events based on digital tangible traces produced by IT systems. Looking at digital traces, there are many concepts to describe their quality—most of them concerned with procedural aspects, i.e., authenticity and integrity for example. Besides that, there exist important concepts that have been largely overlooked by the digital forensics community: Two of those criteria are *relevance* and *expressiveness* of digital evidence, which are scrutinized in the present chapter. Unlike others, those are directly concerned with reaching the investigative goal. Given the ever-rising wealth of digital data, questions of how to determine the case-relevant pieces of digital evidence are especially pressing. Therefore, we approach these two overlooked concepts of digital evidence in a formal way and define those generically in Section 3.3.

To this end, we propose the notion of an investigative knowledge base, which models the understanding of facets, hypotheses, as well as the supports and refutes relations, which map facets to hypotheses in the expected meaning. On top of this, relevance can be defined in terms of these relations. The newly formed relevance relation, in turn, gives rise to grasping the expressiveness of a facet as the set of hypotheses that can be assessed by it. This opens up the possibility to derive metrics for rating the helpfulness of facets using the expressiveness ratio or the degree of difficulty to assess a hypothesis using the specificity ratio.

Additionally, it is shown how the established reliability criteria, accuracy and completeness, can be formally defined using these newly proposed notions in Section 3.4. Accuracy can be assumed when the investigators' assessment matches (past) reality and completeness has to be split into two variants, i.e., decisive and exhaustive completeness. In the case of the former variant, enough expressive facets have been recovered to form a decision. In the case of the latter variant, however, the investigators have searched for or looked at every possible facet contributing to solving the case without being able to do so. Interestingly, the reliability criteria lead to the substantiation of the observation that there are areas in digital forensics that are better understood than others. Naturally, this can be explained by looking at the investigative knowledge base and its filling level as well as its accuracy,

which is in turn largely dependent on the understanding of the programs' behavior on the encountered system.

Aiming to further exemplify these thoughts, a conceptual network of digital evidence is provided in Section 3.5, in which important features of digital evidence, including the concepts of expressiveness and relevance, are related to one another. To illustrate the usefulness of the concepts, we present two applications: First, we demonstrate that the notion of expressiveness and completeness can be used to guide investigations by presenting the *Facet-oriented Criminalistic Cycle* as a thinking model, which extends the well-established Criminalistic Cycle by Walder and Hansjakob [240] in Section 3.6.1. Here, we introduce formal criteria to help investigators in the decision-making process of which facets to collect and when. We provide formal means to determine the unanswered hypotheses and the facets to collect next based on their knowledge base and the current state of the evidence collection. Second, we show in Section 3.6.2 how the concepts help to improve the precision of the so-called CSI model by Freiling and Sack [84], which models pieces of evidence. Specifically, we propose to grasp the relevant information of a piece of evidence as a set of facet-hypothesis-tuples that are contained or can be derived from the piece of evidence in question.

Having presented two potential applications, we relate our approach to previous work in Section 3.7.1 and turn our focus to probabilistic reasoning in digital investigations, which is also implicitly concerned with the relevance of evidence, in Section 3.7.2. Lastly, we discuss the limitations of the proposed formalization. Here, we note that the concept remains largely abstract, i.e., elusive and not immediately applicable. Furthermore, the model is unable to deal with uncertainty or probabilities so far since it uses crisp logic and deals only with boolean cases of relevance.

We believe nevertheless that the presented formalization is valuable for several reasons: First, it contributes to an improved understanding of the problem of "sufficient digital evidence" because we clearly separated two subproblems—finding case-related hypotheses and determining facets relevant to these hypotheses. Second, we clearly define when a facet is considered relevant and thus expressive, i.e., when it can contribute to solving a given investigative question. Furthermore, it generalizes previous approaches using probabilistic approaches like the calculation of the relative weight of single evidence items in Bayesian belief networks. Third, we used the formal concepts to define the reliability criteria, accuracy, completeness, and authenticity, in a more precise way. Lastly, we showed how this understanding can be used to sharpen the criminalistic thinking process and even derive metrics. Accordingly, this chapter aims to improve the understanding of these critical aspects of the overall investigative process by unifying the generic view of facets' relevance and establishing the notion of expressiveness of facets.

4 Necessary and Sufficient Digital Evidence

4.1 Introduction

In order to be able to interpret digital traces, much research within the digital forensic science community follows an experimental or empirical approach [e.g., 19, 234]. Here, researchers conduct elaborate experiments in controlled laboratory environments to gain a practical understanding of which cause implies which specific observable effect—knowledge that can then be used to "reason backward". By doing so, investigators get empowered to approximate the reconstruction of past events in digital systems when confronted with a specific trace situation. While this is a very valid and extremely valuable approach, of course, the trace interpretation is merely based on an inference of causality, which gets increasingly difficult to establish when dealing with complex systems. So, it is a challenging task to gain an understanding of the relevance and expressiveness of digital traces, which we elaborated on in Chapter 3. It is worth noting, however, that the activities within a digital system are entirely defined by the program, the machine model, and the user input, so the question of determining causality becomes, in fact, well-posed and specific. Eventually, at least in theory, event reconstruction and the establishment of relevance of digital traces in the formal sense (Definitions 3.3.3 and 3.3.4), i.e., knowing which facet can help in assessing which hypotheses, should be possible, and it comes naturally to resort to formal methods as known from software and hardware verification to solve the arising reconstruction problems in a mathematically well-founded way. Hence, we investigate different classes of evidence and connected reconstructability properties, aiming to improve our understanding, capabilities, and reliability of digital evidence on a foundational yet abstract level.

4.1.1 Contribution of the Chapter

In the present chapter, we employ the established formalism of linear-time temporal logic (LTL) to develop a new approach to forensic event reconstruction for systems modeled as transition systems. The approach distinguishes between two different classes of evidence of specific actions: *sufficient evidence* and *necessary evidence*. Intuitively, *sufficient evidence* of an action σ is a state predicate whose observation guarantees that σ has actually happened before reaching the current state. Conversely, *necessary evidence* of an action σ is a state predicate that is always and persistently observed after the occurrence of σ, so that its negation can be used to refute the claim that σ has happened. By defining these classes, we provide general notions of *forensic reconstructability* that are not explicit in previous work. The ability to calculate evidence of these evidential classes allows us to infer the relevance

of facets, i.e., effects of actions in the automata in this case, in regard to hypotheses of past events.

Besides offering general insight, our approach also has other (more practical) advantages: After calculating the above evidence sets, the actual checking of whether an event has occurred or not boils down to checking a state predicate on the final state q, which is computationally easy. In a sense, our approach therefore factors out the computational complexity of the analysis problem into the calculation of these evidence sets for certain actions of interest.

Since we formalize the descriptions of necessary and sufficient evidence in LTL, we open the solution space enabling the use of highly optimized tools to calculate these evidence sets and thus profit from decades of research in this area (mostly within the formal verification community). We demonstrate this by utilizing the symbolic model checker NuSMV to calculate these evidence sets in a case study that has been the subject of multiple previous publications in the field of forensic event reconstruction. To do so, we use a prototypical implementation that we provide as open-source software[19] and show how to determine the expressiveness of observed traces using our approach.

4.1.2 Chapter Outline

The chapter is structured as follows: To begin with, we revisit related work regarding formal event reconstruction in digital forensics (Section 4.2). Having an initial understanding of previous research in the field, we provide background on the formal concepts involved, i.e., LTL, model checking, and the Guarded Command Language (GCL) notation, in Section 4.3. Next, we have a closer look at Dewald's specific reconstruction problem (SRP) and illustrate the limitations of his approach to solving it in Section 4.4. We then describe our method of reasoning about evidence using LTL in Section 4.6. In Section 4.7, we describe the implementation to bring our theoretical reasoning to practice. Afterward, we apply the developed procedure to Gladyshev's "ACME Manufacturing" benchmark example and illustrate the benefits of our approach in Section 4.8.1. In addition, we take up the concepts presented in Chapter 3 and describe how the notions presented there relate to the classes of evidence described in temporal logic. The integration is illustrated by the calculation of expressiveness in Section 4.8.2. Lastly, we provide further discussion in Section 4.9 and conclude the chapter in Section 4.10.

[19] `https://github.com/jgru/evidential-calculator`, commit `91afceb`.

4.2 Related Work

4.2.1 Pioneering Formal Event Reconstruction

Gladyshev and Patel [95] can be considered the pioneers of the area of forensic event reconstruction for digital systems. In 2004, they formulated the *event reconstruction problem* as follows: Using expert knowledge of a digital system, determine all possible sequences of events that have previously happened within the system from its final state and available clues of the system's behavior in the past [95, pp. 4 ff.]. Digital systems are represented as finite-state machines (FSMs), and evidence is formalized in the form of *evidential statements*, which are defined as a series of (assumed) observations to find possible scenarios that agree with the observations [94]. Actual event reconstruction is accomplished via *backtracing*, i.e., computing all possible computations leading to the state in question. The main challenge arising from this concept is the extremely large number of computations to consider even for fairly small systems.[20] Indeed, the number of even just the loop-free computations is exponential in the number of states, which in turn is already exponential in the number of variables, a phenomenon known as the state explosion problem.

4.2.2 Model Checking and Formal Event Reconstruction

To avoid the need to enumerate all computations of a system, we considered it natural to resort to temporal logic for the specification of investigative goals and use model checking for their verification. This removes roughly one exponential layer from the theoretical complexity (leaving "only" the state explosion problem). In an early approach of this kind, Rekhis and Boudriga [190] introduced a dedicated extension S-TLA of the *temporal logic of actions* (*TLA*) [145] to provide a formalism to enunciate hypotheses when details are missing. By doing so, they aimed to provide possible explanations by backward chaining, discussed in view of possible applications to reason about investigative scenarios. In a follow-up publication, Rekhis and Boudriga [191] discussed this approach in more detail by presenting a custom model checking algorithm for this formalism (called S-TLC) and applied that to a rather abstract and constructed case study.

Soltani and Hosseini-Seno [207] formalized event reconstruction in a modal μ-calculus, with model checking performed within the mCRL2 tool set; they proposed to address the state explosion problem by exploiting structural symmetries. Their approach, compared to ours in more detail in Section 4.9, was demonstrated on a simplified model of the FAT file system.

Building upon the specific formalization of Gladyshev and Patel [95], James et al. [130] proposed to transform the FSM model into a deterministic finite automaton (DFA) that encodes the set of system computations as a formal language, i.e., a set of strings. In order to answer an investigative question, they converted witness statements into regular

[20] Gladyshev states that "for many real-world systems the brute-force exploration of their exact finite state machine models is infeasible" [93, p. 146].

expressions and eventually into further DFAs, and used them to spot conflicting statements. Technically, this approach is based on taking products of DFAs to check consistency. Like model checking-based approaches, it avoids the doubly exponential complexity of the original backtracing approach and actually is quite related to the use of LTL model checking as pursued in the present chapter, as LTL model checking is also based on translating formulae into a suitable form of automata. An advantage of the use of LTL as advocated in our work is the more compact and readable mode of expression it offers in comparison to writing automata directly, and indeed the translation from LTL into automata incurs exponential blowup [124]. Also, the use of modern symbolic model checkers to some degree alleviates the exponential dependence of the state space on the number of variables, which, contrastingly, remains in full force in a direct automata-theoretic approach.

4.2.3 Set Theory and Formal Event Reconstruction

Since the reconstruction of all paths leading to the final state of a real-world system is computationally intense due to the so-called state explosion problem [93, p. 146], and often might not even be helpful in solving the case considering the size of the resulting set of possible explanations, Dewald [64] approached forensic event reconstruction differently. Hence, he formulated the specific reconstruction problem (SRP), which is geared toward the question of whether a *specific* event or action has occurred. Given a state q and a set of actions Σ, the question to ask is whether a specific action $\sigma_i \in \Sigma$ with relevance for solving the case necessarily happened before reaching q [64, p. 341]. To solve the SRP, Dewald defined the concept of *characteristic evidence* (*CE*). Intuitively, the *CE* of an action in regard to a set of other actions are those traits that are left by this particular action and none of the others. Therefore, the discovery of *CE* is sufficient to prove that an action has occurred, but the absence of *CE* proves nothing. To do so, he used set calculations that are (generally) computationally feasible. However, this comes at the price of losing precision in that there are widespread examples of sufficient evidence that cannot be detected using *CE*, as we will later show in Section 4.4.3.

4.3 Background

To facilitate access to our approach, we now provide a brief introduction to LTL, model checking, and Dijkstra's notation of guarded commands, which we use to describe programs.

4.3.1 Linear-time Temporal Logic

Temporal logic is a formalism to describe and reason about systems in terms of time. In 1977, Pnueli [184] established LTL, in which the nature of time is considered to be linear, i.e., the basic concept is to model time as a sequence of states, so-called *computation paths*,

that describe the evolution of a system over (discrete) time.[21] We briefly recall the syntax and semantics of LTL.

LTL is parametrized over a set `Atoms` of atomic facts of relevance for the underlying system and the task at hand. In a forensic context, such atoms might be, for example, *'mtime of* `/etc/shadow` *has been changed'*, *'socket descriptor 0xEF has been closed'*, or on some other apt abstraction level even propositions such as *'the email has been sent'*. Formulae of the logic are evaluated over a transition system that is a (simplified) model $\mathcal{M}$ of the underlying real-world system; we refer to $\mathcal{M}$ as a *transition system*, or briefly as a *model*. Formally, $\mathcal{M}$ is a triple $\mathcal{M} = (S, \to, L)$, where S is a set of *states*, $\to$ is a binary relation on S indicating *transitions* between states, and L is a labeling function $L\colon S \to \mathcal{P}(\texttt{Atoms})$, which assigns to each state $s \in S$ the set of atomic facts true at s. A (*computation*) *path* of $\mathcal{M}$ is then a sequence $s_1, s_2, s_3, \ldots$ of states such that $s_i \to s_{i+1}$ for all $i > 1$.

LTL provides a language in which to describe sets of computation paths of $\mathcal{M}$ in a rather compact and intuitive way using temporal operators. The simplest LTL formulae are atomic facts $p \in \texttt{Atoms}$. Such a formula is true for all computation paths where p is true in the first state. Temporal operators can then be used to describe the evolution of states along a computation path. For example, a formula of the shape $\Box p$ means that p must be true in the current and all future states of a computation path in the system.

Formally, the syntax of LTL (as we use it in this dissertation) is given by the following Backus Naur form

$$\phi, \psi ::= \bot \mid p \mid \neg\phi \mid \phi \wedge \psi \mid \bigcirc\phi \mid \Box\phi \mid \phi\,\mathcal{R}\,\psi$$

where, as indicated above, p ranges over `Atoms`. The semantics of the logic is given by specifying which paths of states $\pi = s_1, s_2, s_3, \ldots$ *satisfy* which formulae (in which case we also say that the formula *holds* for the path). Along π, we move into the future by taking suffixes of π: For $1 \leq i < n$, we denote by π^i the suffix $s_i, \ldots, s_n$ of π (in particular, $\pi^1 = \pi$). The interpretation of the propositional operators $\bot, \neg, \wedge$ is standard; for instance, no path satisfies $\bot$, and a path satisfies $\neg\phi$ if it does not satisfy ϕ. The semantics of the main temporal operators are given as follows:

- $\bigcirc\phi$ holds for a path π if ϕ holds in the *next* state, i.e., for π^2:

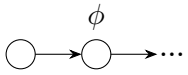

- As mentioned above, $\Box\phi$ holds for π if ϕ holds in the present and all future states of π, i.e., for all suffixes π^i:

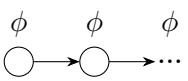

- $\phi\,\mathcal{R}\,\psi$ states that the property ψ must be true until and including the point in time when ϕ becomes true, so that ϕ *releases* ψ. (If ϕ never happens, then ψ is never

[21] For more information refer to the textbook by Huth and Ryan [124] as an extensive resource.

"released" and it degenerates to $\Box\psi$.) Formally, $\phi \mathcal{R} \psi$ holds for π if either ψ holds for all suffixes π^i, or there is $i \geq 1$ such that ϕ holds for π^i, and for all $j \leq i$, ψ holds for π^j:

Further propositional connectives $\vee$, $\rightarrow$, and $\top$ are defined from $\neg$, $\wedge$, $\bot$ in the usual way. Moreover, one can define further temporal operators *eventually* and *until* by duality from the above operators:

- $\Diamond\phi$ states that the property ϕ has to hold in some state on the subsequent path. Note that $\neg\Box\phi \equiv \Diamond\neg\phi$. Formally, $\Diamond\phi$ holds for π if there is some $i \geq 1$ such that ϕ holds for π^i:

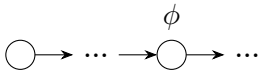

- $\phi \mathcal{U} \psi$ states that the property ϕ has to hold until ψ becomes true, which itself must hold at some future state along π^i. Note that $\mathcal{R}$ is the dual of $\mathcal{U}$ since $\phi \mathcal{R} \psi \equiv \neg(\neg\phi \mathcal{U} \neg\psi)$. Formally, $\phi \mathcal{U} \psi$ holds for π if there is some $i \geq 1$ such that ψ holds for π^i, and for all $j < i$, ϕ holds for π^j:

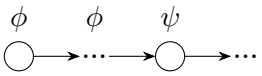

For our approach, it happens that we will only require $\bigcirc$, $\Box$, and $\mathcal{R}$, but for the derivation of the idea and its comparison, we also need $\Diamond$ and $\mathcal{U}$.

Finally, we have a notion of a formula ϕ being true in a *state* q_0 (rather than a computation path) in a finite-state transition system: ϕ holds at q_0 if ϕ holds for *all* paths that start at q_0. Note that at the level of states, saying that $\neg\phi$ holds at q_0 is *not* equivalent to saying that ϕ does not hold at q_0: The former means that no path starting at q_0 satisfies ϕ, while the latter means that not all paths starting at q_0 satisfy ϕ.

4.3.2 Model Checking

The process of checking whether ϕ holds at a specific state q_0 is known as *model checking* [10]—a term also applied more widely to checking satisfaction of formulae in other logical formalisms. Model checking is a technique that was originally developed in the domain of formal verification of hardware to determine whether a state machine meets a given specification. Nowadays, model checking is applicable to software programs as well. Numerous model checkers for various formalisms, like SPIN, TLA+/TLC, and NuSMV, have been released as open-source software. Model checkers generally face a combinatorial blow-up of the state space, however, in order to deal with the already mentioned state

explosion problem symbolic algorithms have been developed that avoid constructing the whole graph of the FSM. In our present approach, we use NuSMV [49] to model-check LTL formulae over finite-state transition systems using such symbolic algorithms.

4.3.3 Guarded Commands

Guarded commands are a programming notation proposed by Edsger Dijkstra in the 1970s [68] to facilitate formal reasoning about the correctness of computer programs. Its distinguishing feature is the use of so-called *guards*, i.e., Boolean expressions, associated with commands to express conditions that must be satisfied for the respective command to be able to execute. In essence, it is a compact notation to specify transition systems, and can thus also be used to reason about parallel programs [42]. Since GCL has been developed with a strong emphasis on formal methods and mathematical rigor, it is well-suited for reasoning about forensic event reconstruction. Hence, we briefly introduce the notation, as it was used by Chandy and Misra [42], using the program shown in Listing 4.1.

The state space of the program is defined by an initial set of variables that store values from a specific domain. For consistency with preceding works, we use Boolean values 0 and 1 as the range of all variables for the examples (except the case study) in this chapter. The second line defines the initial state of the program by assigning a specific value to each variable.

The program is formulated as a set of actions which each consist of a *name*, a *guard*, and a *command*. The name (e.g., `a0`) is merely used to refer to specific actions and is separated from the guard with a colon. The guard is a Boolean state predicate, e.g., $a = 0$, and the command is an assignment of values to variables, e.g., $a := 1$. Multiple assignments within a command are executed in parallel. The guard and the command are separated by an arrow. An empty guard stands for the predicate *true*.

Variables: $\{a, b\}$

Initial state: $\{a = 0, b = 0\}$

Actions:

`a0`: $a = 0 \quad \longrightarrow \quad a := 1$
`a1`: $b = 0 \quad \longrightarrow \quad b := 1$

Listing 4.1: Example program to illustrate the Guarded Commands Language.

The set of actions defines the state transition relation of the program in the following way: If the guard of an action evaluates to true in a given state, we say that the action is *enabled* in that state. For a given state, the set of all enabled actions defines the set of possible next

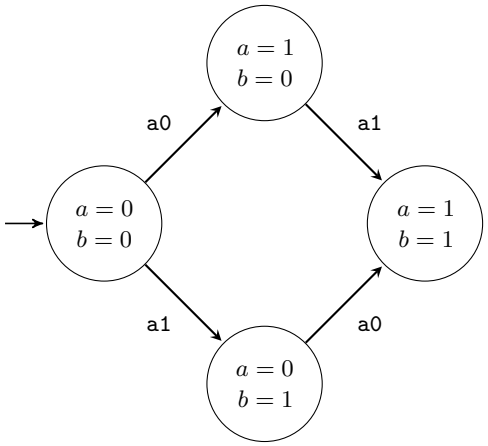

Figure 4.1: State transition diagram of Listing 4.1.

states. During the execution of the program, one enabled action is nondeterministically chosen, and the command of that action is executed, resulting in the next state.

The program shown in Listing 4.1 has two variables a and b and two actions `a0` and `a1`. In the initial state, both actions are enabled resulting in a nondeterministic choice of which action is executed. In this example, each action disables itself by falsifying its guard. So when either `a0` or `a1` is executed first, the remaining other action is then executed. In the final state where $a = 1$ and $b = 1$ holds, no action is enabled anymore, as shown in the program's state transition diagram (Listing 4.1).

If the range of the variables is bounded, every such program effectively defines a labeled finite-state transition system, with values of variables encoded by sufficiently many atoms. Since LTL needs an unlabeled transition system to operate on, we let this system include atoms that record the action that has been taken in the previous step, on the understanding that these atoms are not part of the real system and hence cannot be used as evidence, which instead can observe only the effect of the actions on the variables. De facto, we thereby encode the actions, represented by their edge labels, into the state where their respective effect unfolds.

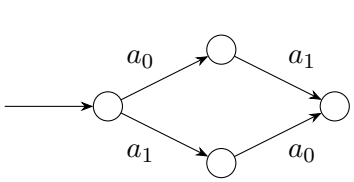

(**a**) Automaton with edge labels

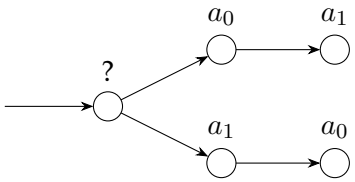

(**b**) Post-state edge encoding

Figure 4.2: The employed post-state label encoding for Listing 4.1. Edge labels denoting the identifier of the respective action in the labeled transition system (LTS), as shown on the left, are moved into the edge target where their respective effect unfolds, as shown on the right. Note that we regard those atoms induced by the post-state edge encoding as immaterial in our calculations.

4.4 Characteristic Evidence and Its Insufficiencies

In the previous chapters, it was deduced that an elementary task in any investigation is to reconstruct past events related to the deed based on the remaining clues, i.e., the facets as the perceivable parts of traces produced by events [131], as precisely as the facets allow.

4.4.1 The Specific Reconstruction Problem

Inspired by the ideas of Gladyshev and Patel [95], Dewald [64] has formalized the problem of forensic event reconstruction as the specific reconstruction problem (SRP): Given an initial state q_0 and an observed state q in a finite-state transition system, determine whether or not a specific action $\sigma \in \Sigma$ necessarily happened before reaching q, in the sense that every computation path that starts in q_0 and ends in q must contain a transition caused by the action σ.[22]

The SRP is illustrated in Fig. 4.3, which graphically depicts a finite-state transition system with states as circles and state transitions as labeled edges between states, where the labels indicate the actions that induce the transition (recall from Section 4.3.3 that the actions are represented as atoms in the formal model, which however do not feature in the real system). Assuming that the system is acquired in state q, an instance of the SRP would be to ask whether a specific action, e.g., σ_0 or σ_1, happened in the past. The answers to these questions can be easily derived from looking at the graph: While σ_0 definitely did not occur on the way to q, σ_1 undoubtedly occurred because there is no path from q_0 to q on which σ_1 does not happen. The SRP regarding action σ_2, however, is not so easy to answer since q can be reached with or without executing that action. This observation shows that there are always three possible answers regarding the SRP and some action σ: (1) yes, σ definitely happened on the way to q, (2) no, σ did not happen on the way to q, and (3) σ may or may not have happened on the way to q.

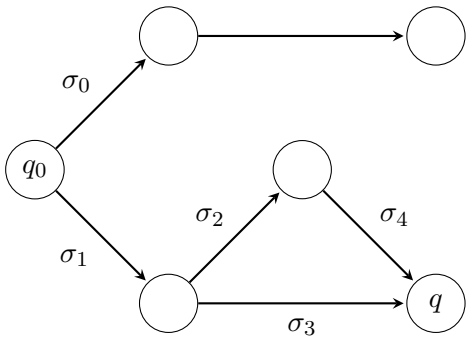

Figure 4.3: An exemplary LTS to illustrate the SRP according to Dewald [64, p. 341]. Confronted with the system in the state q, the investigators have the task of determining whether a specific action, e.g., q_2, has happened in the past.

[22] Note that in the formal verification community such a computation path is sometimes referred to as a *trace*, which is fundamentally different from the notion of that term in the forensic science community, as employed in this dissertation.

4.4.2 Dewald's Characteristic Evidence Method in Detail

To solve the SRP without such a computationally intense effort as the approach by Gladyshev and Patel [95] requires, Dewald [64] defined the concept of *characteristic evidence* (*CE*) using set theory. As mentioned previously, the *CE* of an action σ with respect to a set of other actions Σ' are those values of variables that are left only by σ and by no other action in Σ'. The discovery of *CE* in q is sufficient to prove that σ has occurred. Formally, $CE(\sigma, \Sigma')$ is the state predicate defined by all assignments of σ that are not performed by any other action in Σ'.

Expressed in a formal way, Dewald [63, pp. 86 ff.] defined an evidence set $E(\sigma)$ of an action $\sigma \in \Sigma$ by the powerset of all variable assignments performed by this action, i.e., $\mathcal{P}\big(\{[v = d] | [v = d] \in \sigma\}\big)$. Such evidence sets can be merged, which he defined as *merged evidence* (*ME*) of a set Σ of actions:

$$ME(\Sigma) = \bigcup_{\sigma \in \Sigma} \bigcup_{e \in E(\sigma)} e \tag{4.1}$$

Based on these definitions, he came up with the concept of CE, which is defined as follows [63, pp. 86 ff.]:

$$CE(\sigma, \Sigma') = E(\sigma) \setminus \Big(\mathcal{P}\big(ME(\Sigma') \cup ZE\big)\Big) \tag{4.2}$$

The characteristic evidence $CE(\sigma, \Sigma')$ of that action with respect to a set of comparative actions can be considered the evidence of the action σ without the unified evidence of all other actions in Σ', where $\sigma \notin \Sigma'$, and the so-called *zero evidence*, i.e., variable assignments that are already present in the initial state q_0. The result of this calculation is a set of $\langle name, value \rangle$-pairs. If a tuple that is included in the set $CE(\sigma, \Sigma')$ can be observed in the state q, then one can conclude that the target action σ must have happened sometime before reaching this state and thus the SRP is solved [64, p. 342].

For example, in Lst. 4.1, *CE* of action `a0` is the condition $a = 1$ since no other action sets variable a to that value. Hence, observing $a = 1$ implies that `a0` has happened. Similarly, the *CE* of action `a1` is $b = 1$, and observing that value implies that `a1` previously occurred.

The concept of *CE* has the advantage of being computationally feasible with careful implementation regarding the powerset while still remaining easily comprehensible. Several case studies exist that perform event reconstruction using this approach with filesystem metadata [134, 135] and entries from log files [147], which illustrate the principal applicability of this method in practice impressively.

4.4.3 Incompleteness of the Characteristic Evidence Method

While the approach of Dewald [64] is both intriguing by its simplicity and provenly helpful in real-world settings, especially for identifying files for a subsequent analysis and reconstructing past events in certain limits, we found several shortcomings that motivate the need for enhancing the notion of evidence in digital systems.

Dewald's approach draws its simplicity from ignoring the states of the system that are visited during a given execution. This leads to the fact that the *CE* sets can become very small in practice. This is also the reason why *CE* is not a complete characterization of whether or not an action has been executed. Furthermore, the set of *CE* becomes smaller the larger the set Σ' of other actions is, but a large set of comparative actions increases the precision of the concept. Overall, the concept of *CE* does not identify all conditions that can be used to conclude that a certain action must have happened in the past because it ignores the preconditions, i.e., the guards in the GCL programs, as we now illustrate using two examples.

4.4.3.1 Unreachable Action

Consider the program in Lst. 4.2 where action `a1` is never executed because a is never set to 1; hence, `a1` is unreachable, as the state transition diagram provided in Fig. 4.4 quickly shows.

Variables: $\{a, b\}$

Initial state: $\{a = 0, b = 0\}$

Actions:

`a0`: $a = 0 \quad \longrightarrow \quad b := 1$
`a1`: $a = 1 \quad \longrightarrow \quad b := 1$

Listing 4.2: Example program with an unreachable action.

If we compute *CE* for action `a0` with respect to $\Sigma = \{\texttt{a1}\}$ without any further precautions, we observe that `a0` and `a1` have the same effect and therefore cannot be distinguished in this respect. The formal calculation of *CE* for either action results in an empty set, as it is shown in (4.3)[23]. However, if $b = 1$ is observed, it is clear that only `a0` could have been executed since `a1` is unreachable.

$$\begin{aligned} & CE(a_0, \{a_1\}) \\ = \; & CE(a_1, \{a_0\}) \\ = \; & \big\{\{b = 1\}, \varnothing\big\} \setminus \mathcal{P}\Big(\{b = 1\} \cup \{a = 0, b = 0\}\Big) \\ = \; & \varnothing \end{aligned} \tag{4.3}$$

[23] While referencing formulae by plain numbering may seem uncommon for some readers, this referencing method follows the recommendation in *The Chicago Manual of Style* [227, p. 594].

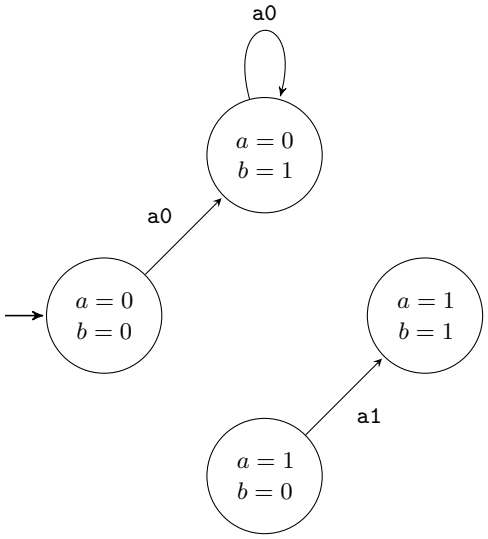

Figure 4.4: State transition diagram of Listing 4.2.

Dewald argues that in real systems, unreachable states and non-executable actions can be neglected; therefore, he prunes the system model beforehand by excluding those actions that are not executed on any path of the system [63, p. 78]. Following this approach, the result of $CE(a_0, \{a_1\})$ is, in fact, $\big\{\{b = 1\}\big\}$ when using a pruned system model as well. However, this pruning step requires checking each action on each computation path, which is why we consider this to be a drawback.

4.4.3.2 Action Guarded by a Semaphore

The second example is shown in Lst. 4.3 where the actions `a2` and `a3` can only be executed if a is set to 1, i.e., if action `a1` has been executed before to enable the semaphore, as apparent in Fig. 4.5. The calculation of *CE* of `a1` compared to $\Sigma' = \{\texttt{a0}, \texttt{a2}, \texttt{a3}\}$ observes only the immediate effect of `a1` and results in the characteristic evidence $\big\{\{a = 1\}\big\}$.

Variables: $\{a, b\}$
Initial state: $\{a = 0, b = 0\}$
Actions:

a0:	$\longrightarrow$	$a := 0$
a1:	$\longrightarrow$	$a := 1$
a2: $a = 1$	$\longrightarrow$	$b := 0$
a3: $a = 1$	$\longrightarrow$	$b := 1$

Listing 4.3: Example program with actions guarded by a semaphore.

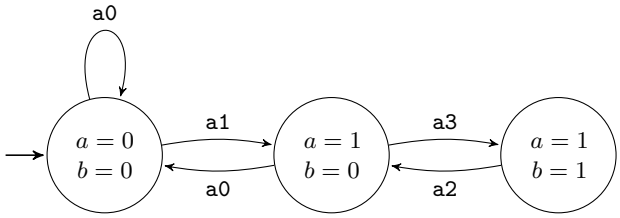

Figure 4.5: State transition diagram of Listing 4.3.

$$\begin{aligned} & CE(a_1, \{a_0, a_2, a_3\}) \\ &= \big\{\{a = 1, \varnothing\}\big\} \setminus \mathcal{P}\left(\{a = 0, b = 1, b = 0\} \cup \{a = 0, b = 0\}\right) \\ &= \big\{\{a = 1\}\big\} \end{aligned} \tag{4.4}$$

But while in fact the observation of $a = 1$ can be used to conclude that `a1` has happened in the past, this is not the only condition to allow that conclusion. Since action `a3` is dependent on action `a1`, the observation of `a3` ($b = 1$) implies that both `a3` and `a1` have been executed before, which again underlines that *CE* sets are incomplete.

Giving an outlook on future work in his dissertation, Dewald stated that, potentially, there might be weaker conditions that are sufficient to reconstruct events [63, p. 196]. By looking at the examples above, we roughly sketched out the nature of those weaker conditions stemming from the incorporation of state information. Hence, we will approach these conditions in the following sections to enhance the concept.

4.5 Solution of the SRP in LTL

The examples given in Listings 4.2 and 4.3 suggested to incorporate state information in order to find those reconstructibility conditions. The consideration of time can be accomplished by using temporal logics, such as LTL. This branch of temporal logic, as introduced in Section 4.3.1, can be considered an apt choice since it provides concise yet brief formulae and helpful backward modalities.

Fig. 4.6 shows a rough graphical visualization of what it means to solve the SRP: Specifically, it suffices to establish that the target action must have occurred in the computation by means of which the observed state q was reached from the initial state q_0; by use of well-developed formal verification techniques, the necessity of going through all possible computation paths or even to enumerate the entire state space explicitly may be avoided using symbolic algorithms so it suffices to consider those parts needed for our specification-based checks presented below.

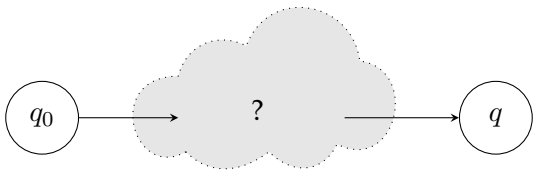

Figure 4.6: The intuition of the SRP. q_0 denotes the initial state and q the observed state. Finding out whether action σ happened between q_0 and q does not necessarily involve explicitly enumerating all computation paths or even all states of the system when using symbolic model verification techniques.

4.5.1 Temporal Logic Approach by Dewald

Dewald [64] aimed to solve the SRP using a set-theoretic approach. However, he pointed out that temporal logic could be used to describe and solve the general reconstruction problem (GRP) and SRP potentially. On the one hand, he formulated the GRP in LTL as "$\Diamond q$" [64, p. 345]. While this looks sensible on the first glimpse, the details of LTL's semantics lead to unfavorable results since we need a way to transform the model to single computation paths, in order to verify systems. Commonly, this is done by referring to all paths [124, p. 182]. This, however, leads to a meaning of the statement $\Diamond q$ that can be verbalized as follows: "All paths starting in the initial state will finally reach q". However, what should have been (but actually is not) expressed here is the referral to only those paths that lead to q. Thus, this formula will either be not satisfied or should not be indicative of anything. On the other hand, the SRP was then formulated as "$\neg(\neg a \,\mathcal{U}\, q) \wedge \Diamond q$" by Dewald [64, p. 345]—an approach, which suffers from the same intricacies as described above. Even using a different semantics than the common one described by Huth and Ryan [124, p. 182] in their seminal textbook to retrieve the computation paths, we have the only alternative to query existentially which does not lead to sensible results either, given that we (most probably) already know that a path exists to the observed state q or not. Hence, there is a need to refine these logical specifications, but nonetheless they were a pioneering effort pointing to potential improvements in his method.

4.5.2 Proposed Temporal Logic Approach

To solve the SRP, we need a characterization of conditions on q that allows concluding that some target action σ has previously happened. Such conditions are formalized as state predicates E which represent *evidence*. If E is true in some state q, we say that q *contains* or *provides* evidence E. So in essence we define a formula that takes a potential candidate of evidence and checks its evidential value.

Considering the intricacies of LTL semantics, we express the underlying but now refined concept of the characteristic evidence method proposed by Dewald [64], as follows:

$$\Box\big(\bigcirc(E \wedge \neg\sigma) \rightarrow E\big) \tag{4.5}$$

In (4.5), we state that if the execution of an action other than the target action σ results in a state where the evidence E is observable, then, the evidence must have already been present at the current state. Put differently, the presence of evidence E is an invariant of all actions other than the target action σ. From here, we can form the contraposition of (4.5):

$$\square\big(\neg E \rightarrow \bigcirc(\neg\sigma \rightarrow \neg E)\big) \tag{4.6}$$

Here, it becomes apparent that the absence of evidence of the target action is an invariant of all actions except the target action.

Based on this insight, we can resort to LTL extended with past-time modalities to formulate the intuition of solving the SRP in an easier-to-read version. Past operators do not increase LTL's expressivity but their usage makes the encoding of certain properties indeed more convenient [124, p. 221]. Hence, we use the past operator *once* ($\diamondminus$) to constrain the past. $\diamondminus\phi$ states that ϕ had to hold sometime in the past. Based on (4.6), we can formulate an easier-to-read version of it using the past-time modalities that captures more than Dewald's notion of CE:

$$\square(E \rightarrow \diamondminus\sigma) \tag{4.7}$$

The formula (4.7) states that always if evidence E is observable at the current state, then at some point in the past, the target action σ must have occurred at some point in the past.

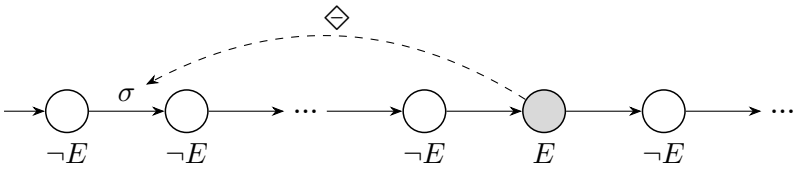

Figure 4.7: Intuition of the solution to the SRP expressed in LTL as given in (4.7). If evidence E is observable in the current state, then at some state in the past the target action σ must have been executed. Note that the evidence E may not be present immediately and that it might diminish again.

So, using one of these equivalent formulae, the SRP can be solved, if (and only if) there is evidence E observable in the current state that is (mediately) induced by the target action σ, or, considering state, some other action σ' that necessitates that σ happened before, as visualized in Fig. 4.7. Given a tool to check this formula on a system, it can be used to empirically test whether E is a witness of action σ. Practically, one could replace the actual evidence E with a conjunct of the elements of the set of observed evidence $E_{obs} \subseteq \mathcal{P}(AP)$, which needs not to be necessarily complete:

$$\mathcal{M} \models \square\Big(\big(\bigwedge_{e_i \in E_{obs}} e_i\big) \rightarrow \diamondminus\sigma\Big) \tag{4.8}$$

If (4.8) yields true, we can conclude, based on the observed state, that the target action σ must have been executed sometime before, so the SRP is solved. Conversely, when investigators have to rule out that a certain action was executed, they can check their hypothesis by simply negating the implication ($\neg\diamondminus\sigma$) in (4.8), stating that the target action σ could have never happened.

The observation from (4.8) suggests that it is possible to classify observable traces regarding their meaning for reconstruction problems. These different classes of observable evidence will be investigated in the next section.

4.6 Necessary and Sufficient Evidence

We now present a formally complete and practically more widely usable notion of "useful evidence" to solve the SRP and formulate it using LTL formulae.

4.6.1 Sufficient Evidence

We now turn to the first type of evidence that is useful for event reconstruction. Intuitively, *sufficient evidence* of an action σ is a state predicate (evidence) SE such that observing that predicate implies that σ has previously happened. This includes *any* state that is only reached after σ has happened, i.e., not only those states that are an immediate effect of executing σ itself but also those states that result from follow-up actions that are guarded by σ. The only restriction is that these states cannot be reached unless σ has happened. Since we are assuming that the post-state of the system is observed statically, we may restrict sufficient evidence formulae to be purely propositional, i.e., to consist only of atoms and propositional operators $(\bot, \neg, \wedge)$.

Given that the LTL past-time modalities are not supported by all model checkers, we now use the $\mathcal{R}$-operator of LTL to formalize the property of *sufficient evidence* (SE) of σ as follows:

$$(\bigcirc\sigma)\,\mathcal{R}(\neg SE) \tag{4.9}$$

The expression (4.9) holds for state predicates SE that are "released" by the action σ. It thus says that a state satisfying SE can only be reached after σ has been executed, which is illustrated in Fig. 4.8. In that sense, SE is optional in nature but definitive regarding the conclusion implied by its observation.

Note that the *next* operator $\bigcirc$ is needed for technical reasons because of the encoding of actions as atoms in post-states mentioned in Section 4.3.3 (Fig. 4.2).

When enumerating sufficient evidence in this way, properties of unreachable states end up in the set of SE as they never become true. This is due to the mechanics of the *release* operator $\phi\,\mathcal{R}\,\psi$, which is formally defined such that it also holds for a path π if ψ, i.e., in our case $\psi = \neg SE$, holds for all suffixes $1 \leq i < n$ of that path π^i. Hence, one may optimize the list of formulae obtained by pruning conditions that are not actually satisfied in any reachable state. To do so, one requires additionally that

$$\Box(\neg SE) \tag{4.10}$$

does *not* hold in the model (cf. the discussion in Section 4.3.1).[24]

[24] Note that the conjunct of the expressions (4.9) and (4.10) is not the same as checking them one after another.

We emphasize that one is interested in the weakest possible formula that still constitutes sufficient evidence; e.g., if both $a = 1$ and the disjunction $a = 1 \vee b = 1$ are *SE* for σ, then the latter formula is preferable as it allows establishing more easily that σ has happened.

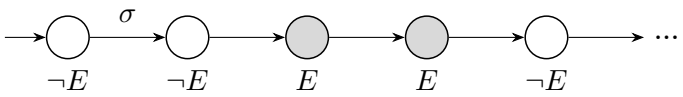

Figure 4.8: Visualization of a property E that constitutes sufficient evidence in the sense that $SE \coloneqq E$ satisfies (4.9). Note that E needs not actually occur immediately after $\neg E$ is released by σ (which would be one step earlier than shown in the above example) but may occur at some future state after σ is executed.

Overall, the resulting intuition behind the concept of *SE* can be summed up as follows:

> "Whenever evidence *SE* is observable, we can conclude that the target action σ must have been executed."

In Fig. 4.8, however, it also becomes apparent that sufficient evidence can be diminished by following transitions; hence, the reversed conclusion is invalid, so we develop the notion of *necessary evidence* next.

4.6.2 Necessary Evidence

As a counterpart to *SE*, we consider *necessary evidence* (*NE*) to be evidence which must be inevitably present in all subsequent states after the target action has been executed. The property of a formula *NE* being necessary evidence is formally expressed in (4.11), which states that except in the initial state (caused by the edge label encoding realized by excluding it via the *next* operator $\bigcirc$),the execution of the target action σ implies presence of the evidence *NE* in all future states:

$$\bigcirc\big(\Box(\sigma \rightarrow \Box NE)\big) \tag{4.11}$$

This property is illustrated in Fig. 4.9. The intuition behind the concept of *NE* might be verbalized as follows:

> "Whenever the target action σ has been executed, evidence *NE* is observable in all subsequent states."

Note that, in contrast to *SE*, from the presence of *NE* one cannot draw conclusions on the execution of the target action because *NE* may already hold before σ is executed. But since *NE* must hold until the final state, one can establish that σ has *not* happened using *NE*: If *NE* is not observable in the current state, then σ has not happened (yet). In opposition to the situation with sufficient evidence, we are interested in the *strongest* possible formulae when looking for *NE*. For instance, if both $a = 1$ and the conjunction $a = 1 \wedge b = 1$ are *NE* for σ, then the latter is preferable as it allows excluding the possibility that σ has happened more easily.

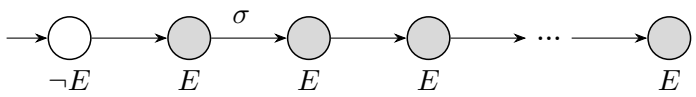

Figure 4.9: Visualization of a property E that constitutes necessary evidence in the sense that $NE \coloneqq E$ satisfies (4.11). Note that E *may* hold before σ is executed, but *must* hold in every state after σ is executed; that is, all occurrences of E shown above except the first are mandated by (4.11).

Our notion of necessary evidence thus enhances and extends Dewald's concept of *characteristic counter evidence* (*CXE*), which describes facets (values of variables) that exclude the execution of the target action: The negation of any formula constituting *CXE* is similar in spirit to necessary evidence in our sense. Like *CE*, *CXE* only considers changes directly induced by the target action, so not all necessary evidence can be obtained by negating characteristic counter evidence, as seen in separating examples similar to the ones shown above for sufficient evidence. Since *NE* is not limited to action-induced changes, the negation of it can be considered a more *universal* form of *CXE*.

4.6.3 Action-induced Evidence

Given the definitions of *SE* and *NE* above, one can ask how these relate to Dewald's notion of *CE*. We answer this question by formalizing *action-induced evidence* (*AE*) in LTL. The resulting concept relates to states that are introduced by the target action itself and no other action. Formally, a formula *AE* is action-induced evidence if the following holds:

$$\begin{aligned} &\neg AE \\ &\wedge \bigcirc\Box(\sigma \rightarrow AE) \\ &\wedge \Box\big(\textstyle\bigwedge_{\sigma' \in \Sigma \setminus \{\sigma\}} (\neg AE \rightarrow \bigcirc(\sigma' \rightarrow \neg AE))\big) \end{aligned} \tag{4.12}$$

(the big conjunction symbol $\bigwedge$ expresses a finite conjunction over all actions σ' other than σ). This captures the spirit of Dewald's definition: *AE* does not hold initially, is brought about by σ, and is not brought about by any other action σ', the latter in the sense that if *AE* does not hold before the execution of σ', then it does not hold afterward either. However, we differ from Dewald's *CE* on actions that can never be executed, a situation that is resolved in his original method by removing unreachable actions before the calculation of *CE* since it does not take guards into account at all. In our case, such a non-executable action would have precisely all unreachable evidence as *AE* and hence include facets that could never be witnessed. To avoid this unintuitive (although correct) result, one could employ the same pruning as for *SE* and check that $\Box(\neg\sigma)$ does not hold in the initial state and get an empty set of *AE* instead. On the one hand, one sees easily that action-induced evidence is indeed a subset of sufficient evidence. On the other hand, as we have seen above, the converse implication does not hold, i.e., sufficient evidence needs not be action-induced. The relations among some of the mentioned evidence classes are graphically summarized in Fig. 4.10.

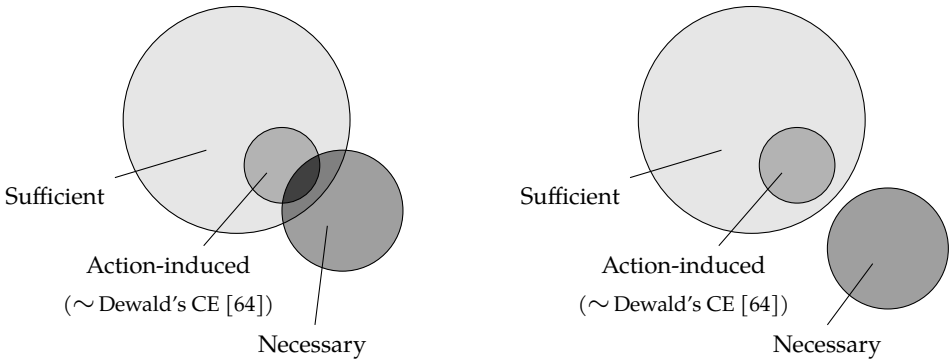

(**a**) Overlapping classes of evidence. (**b**) Separated classes of evidence.

Figure 4.10: Venn diagram of the classes of evidence illustrating their mutual relations. Most notably, the notion of sufficient evidence is strictly broader than that of action-induced evidence, which is a subset of the former and captures Dewald's notion of characteristic evidence [64]. Necessary evidence might also be part of sufficient evidence, e.g., a variable assignment by a future action guarded by the target action that will not change after its assignment. However, it might also be completely separate. Examples of separate necessary evidence are an effect shared by two actions or, alternatively, a mere guard for an action that does not change after the execution of the target action. Note that the necessity of evidence is orthogonal to these classes.

4.6.4 Examples

To illustrate the above concepts, let us look at some examples.[25] The program in Lst. 4.1 is one of the simplest cases. Here, variables a and b are "witness" variables for the execution of `a0` and `a1`, respectively: They have value 1 if and only if that action was executed. So for both actions, *SE*, *NE* and, *AE* are essentially the same (up to unnecessary strengthening or weakening), namely the conditions that $a = 1$ (for `a0`) or $b = 1$ (for `a1`), correspondingly.

For Lst. 4.2, there is no characteristic, viz., action-induced, evidence in the sense of Dewald for `a0`, as we already calculated in (4.3) since the (unreachable) action `a1` has the same effect as action `a0`. On the other hand, the condition $b = 1$ is both *SE* and *NE* for `a0`.

Considering Lst. 4.3, we observe that the condition $a = 1$ is *AE* for `a1`, but the weaker condition $a = 1 \vee b = 1$ is *SE* for `a1`. Since both a and b can switch between 0 and 1 unboundedly often, there is no *NE* for any action.

A more complex example is shown in Listing 4.4 with its corresponding state transition diagram provided in Fig. 4.11. The program has four variables and four actions. Like in Lst. 4.1, the actions have a specific variable that they exclusively set to 1.

[25] A presentation of the examples in literate programming style can be found at `https://github.com/jgru/evidential-calculator/tree/master/examples`

Variables: $\{a, b, c, d\}$

Initial state: $\{a = 0, b = 0, c = 0, d = 0\}$

Actions:

a0:	$\longrightarrow$	$a := 1$
a1: $a = 1$	$\longrightarrow$	$b := 1$
a2: $b = 1$	$\longrightarrow$	$c := 1; d := 0$
a3: $b = 1$	$\longrightarrow$	$d := 1; b := 0$

Listing 4.4: Example program to illustrate the evidence set calculation.

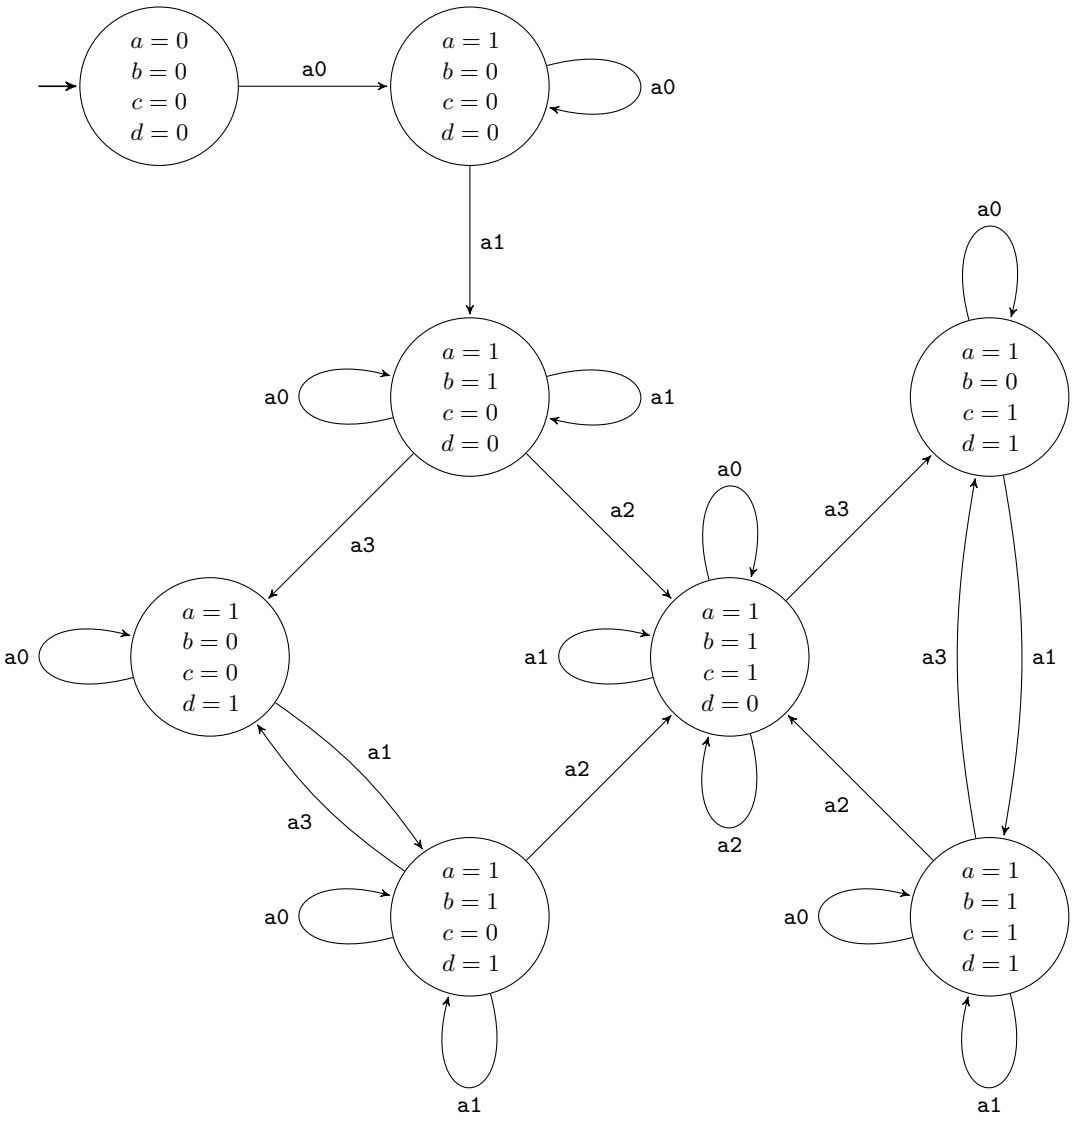

Figure 4.11: State transition diagram of Listing 4.4.

The action-induced evidence (in this case corresponding to Dewald's characteristic evidence) containing only the immediate effects of each action is easily calculated, with results shown in Table 4.1.

Sufficient evidence contains both the immediate effects of the action and the effects of subsequent actions that are guarded by the respective target action. The resulting state conditions are given in Table 4.2.

Values of variables are necessary evidence only if they do not change after the respective action has been executed. Since the variables b and d could change back to 0, the values of these variables are only included in combination with other variables. The resulting conditions are given in Table 4.3.

Table 4.1
Action-induced evidence (AE) set of the example program listed in Listing 4.4.

Action	Condition
`a0`	$a = 1$
`a1`	$b = 1$
`a2`	$c = 1$
`a3`	$d = 1$

Table 4.2
Sufficient evidence (SE) set of the example program listed in Listing 4.4.

Action	Condition
`a0`	$a = 1 \lor b = 1 \lor c = 1 \lor d = 1$
`a1`	$b = 1 \lor c = 1 \lor d = 1$
`a2`	$c = 1$
`a3`	$d = 1 \lor (b = 0 \land c = 1)$

Table 4.3
Necessary evidence (NE) set of the example program listed in Listing 4.4.

Action	Condition
`a0`	$a = 1 \land (b = 1 \lor c = 0 \lor d = 1)$
`a1`	$a = 1 \land (b = 1 \lor d = 1)$
`a2`	$a = 1 \land c = 1 \land (b = 1 \lor d = 1)$
`a3`	$a = 1 \land (b = 1 \lor d = 1) \land (c = 1 \lor d = 1)$

4.7 Implementation

In order to transfer the theoretical concepts presented above into practice, we have developed a prototype to calculate evidence sets automatically. Fig. 4.12 provides an overview of the components, inputs, and outputs of the tool. The source code of our prototypical implementation is publicly available under the GNU Lesser General Public License v3.0.[26]

4.7.1 Dependencies of the Prototype

Our implementation uses the established model checker NuSMV[27] [49]. In addition, we employ the Python library PyNuSMV[28] [27] to control the model checker conveniently and provide the specifications to check. Our implementation serves as a proof of concept; of course, a variety of other reasoners could be employed.

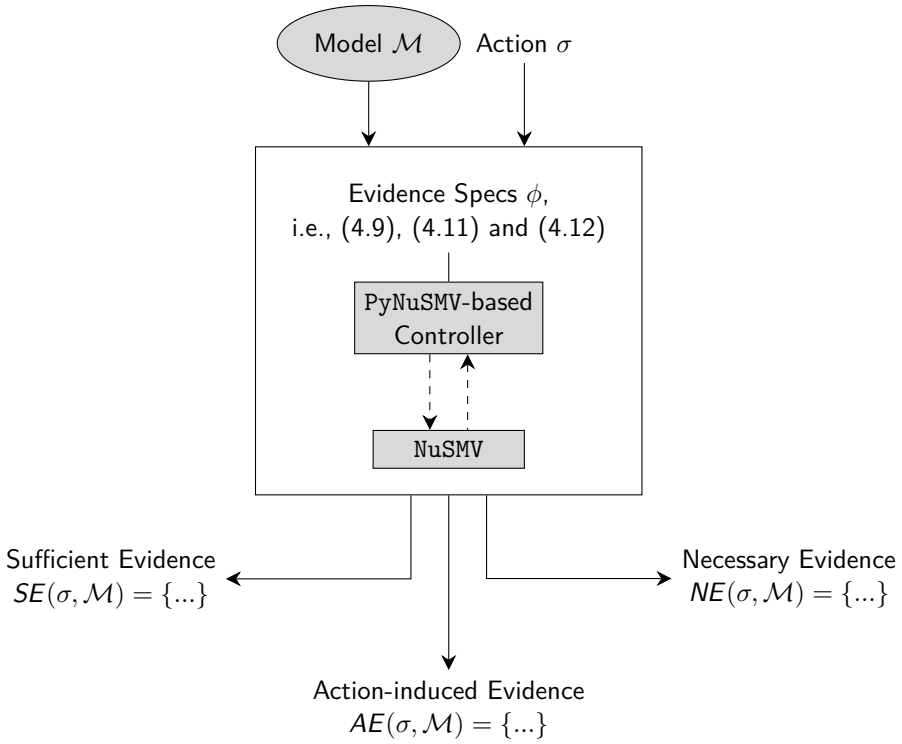

Figure 4.12: Overview of our prototypical implementation. The system is implemented in Python and builds on the PyNuSMV library, which provides bindings to control the model checker NuSMV. When a model $\mathcal{M}$ is provided in the form of NuSMV's input language, the evidence set for the specified action σ_i will be calculated based on the LTL formulae describing the classes of evidence.

[26] `https://github.com/jgru/evidential-calculator`, commit `91afceb`.
[27] `https://nusmv.fbk.eu/`, v2.6.0, accessed 13 Dec. 2023.
[28] `https://github.com/LouvainVerificationLab/pynusmv`, v1.0rc8, commit `e7a8e0a`.

4.7.2 Calculation of Evidence Sets

Our tool takes a model $\mathcal{M}$ describing the system under investigation in NuSMV's specification language and the identifier of the target action σ as inputs. Using these inputs, the program calculates and outputs the various evidence formulae. The idea of the algorithm is to enumerate all possible valuations of the variables, expressed as state predicates, and use the model checker to verify in each case whether the given state predicate is *SE* or *NE*, respectively, according to the LTL formula schemes discussed above. We phrase the algorithm as working with per-state values of variables (encoded using atoms as mentioned earlier). We need the notion of a *partial valuation* for a set V of variables with assigned ranges of values. Such a partial valuation is a finite conjunction of formulae of the form $a = v$ where a is a variable in V and v is a value in the range of a. In the general terminology proposed by Jaquet-Chiffelle and Casey [131] to describe observable parts of a trace (Section 2.4.2), a partial valuation can be considered a facet in this context. The negation $\neg p$ of a partial valuation p is formed by negating the formula representation of p. So the conjunction becomes a disjunction and all the equalities become inequalities. The set of all such partial valuations over V will be referred to as $\mathsf{PVal}(V)$. We require that each variable is mentioned at most once in a partial valuation; if a variable is not mentioned, its value is regarded as immaterial. The calculation proceeds then as follows:

1. Load the model $\mathcal{M}$ into the model checker.
2. Retrieve the set V of variables from the model.
3. For each partial valuation p of the variables in V, do the following:
   a) Form the LTL specification ϕ expressing that p belongs to the evidence class of interest w. r. t. the action σ (details are discussed below), and
   b) Check that ϕ holds in the model; if yes, accommodate p in the corresponding evidence set as specified below.

In Step 3a, we fill in a concrete partial valuation p as an evidence candidate in (4.9) and (4.11), in a manner that depends on whether we are looking for sufficient or necessary evidence. Specifically, we write $q_0 \models \phi$ if the initial state q_0 of the model $\mathcal{M}$ satisfies ϕ (this is checked in Step 3b). We denote the evidence sets computed by the algorithm by $SE(\sigma, \mathcal{M})$ and $NE(\sigma, \mathcal{M})$ respectively, which are formally defined, as follows:

Definition 4.7.1 (Set of sufficient evidence)**.** The *set of sufficient evidence* $SE(\sigma, \mathcal{M})$ of a target action σ with respect to a system model $\mathcal{M}$ is defined as the set of those partial valuations $p \in \mathsf{PVal}(V)$ for which the formula $q_0 \models (\bigcirc\sigma)\,\mathcal{R}(\neg p)$ holds at q_o:

$$SE(\sigma, \mathcal{M}) \coloneqq \big\{p \mid p \in \mathsf{PVal}(V), q_0 \models (\bigcirc\sigma)\,\mathcal{R}(\neg p)\big\} \tag{4.13}$$

Definition 4.7.2 (Set of necessary evidence)**.** The *set of necessary evidence* $NE(\sigma, \mathcal{M})$ of a target action σ with respect to a system model $\mathcal{M}$ is defined as the set of those partial valuations $p \in \mathsf{PVal}(V)$ for which the formula $\bigcirc\big(\Box(\sigma \rightarrow \Box(\neg p))$ holds q_o:

$$NE(\sigma, \mathcal{M}) \coloneqq \big\{\neg p \mid p \in \mathsf{PVal}(V), q_0 \models \bigcirc\big(\Box(\sigma \rightarrow \Box(\neg p))\big)\big\} \tag{4.14}$$

After their construction, these sets can be used to test hypotheses and solve the SRP. To this end, we read $SE(\sigma, \mathcal{M})$ disjunctively, and $NE(\sigma, \mathcal{M})$ conjunctively. Note that in the latter case, the formulae contained in $NE(\sigma, \mathcal{M})$ are negated descriptions of valuations, so while $SE(\sigma, \mathcal{M})$ is effectively computed as a disjunctive normal form, $NE(\sigma, \mathcal{M})$ constitutes a conjunctive normal form.[29]

We remark that the above algorithm is, of course, exponential in the number of variables; that is, it makes exponentially many calls to the model checker. This is due to the fact that the algorithm computes the optimal evidence formula, e.g., the weakest sufficient evidence. Indeed, the actual tool implements an optimization according to which small partial valuations are tried first, and partial valuations extending ones that are already included in the evidence set are disregarded; this leads to more compact evidence formulae as apparent in Tables 4.2 and 4.3. Alternatively, one may just call the model checker with some target formula that is hypothesized to contain sufficient evidence (for instance, a complete description of a specific observed state); in this approach, the computational cost is just that incurred by the model checker. Of course, due to the well-known state explosion problem, model checking is, in principle, already exponential in the number of variables, but modern symbolic model checkers will often perform more efficiently in practice (within the bounds of PSpace-complexity).

4.8 Application

After the contrived examples, we now turn our heads to a case study taken from the literature on formal event reconstruction. Afterward, we show how the classes of evidence, as described above, can be used to instantiate relevance to breathe more life into the topic of Chapter 3.

4.8.1 Case Study

To illustrate the helpfulness of our methods for actual case work, we now apply them to a case study that has been repeatedly discussed in previous works in the field.

4.8.1.1 The Investigation at ACME Manufacturing

Gladyshev and Patel [95] presented a fictitious example case concerning a made-up company called *ACME Manufacturing*, subsequently picked up by James et al. [130] as well as by Soltani and Hosseini-Seno [207].

[29] The employed dualization is needed to get disjunctions from the partial valuation in order to match the intuition of the strictest NE.

The underlying situation is described as follows: There is a local area network at the ACME Manufacturing company with two computers and a networked printer. Alice and Bob operate the network and share the costs. Alice, however, refuses to pay for the maintenance of the printer and claims to have never used it. Since Bob disagrees because he once saw Alice collecting printouts, an investigation of the facts has to be initiated to resolve the dispute. The functioning of the printer in question has been described by Gladyshev and Patel [95, p. 4] in the following way:

> According to the manufacturer, the printer works as follows:
>
> 1. When a print job is received from the user it is stored in the first unallocated directory entry of the print job directory.
> 2. The printing mechanism scans the print job directory from the beginning and picks the first active job.
> 3. After the job is printed, the corresponding directory entry is marked as "deleted", but the name of the job owner is preserved.
>
> The manufacturer also noted that
>
> 4. The printer can accept only one print job from each user at a time.
> 5. Initially, all directory entries are empty.

A forensic examination of the print job directory uncovers two processed print jobs of Bob; the rest of the directory was empty, as shown in Listing 4.5, illustrating the observed evidence E_{obs}. However, this finding does not provide a straightforward answer to the investigative question of interest. So, what should an analyst conclude based on this finding? Apparently, more reasoning is necessary to assess the investigative hypothesis that Alice has printed once.

```
first entry  = job from B (deleted)
second entry = job from B (deleted)
third entry  = empty
fourth entry = empty
...
nth entry    = empty
```

Listing 4.5: Observed evidence E_{obs} extracted from the print job directory of the printer in the ACME network.

4.8.1.2 Calculation of Sufficient Evidence

In view of the description of the inner workings of the printer provided by the manufacturer, we can model the system under investigation in NuSMV's specification language. For the sake of simplicity, we restrict the model to two entries in the print job directory, which is abundant to depict the situation. To solve the case, we then deploy our newly developed tool, which has already been presented in Section 4.7. By providing the model $\mathcal{M}$ and

Table 4.4
Set of sufficient evidence $SE(\texttt{add\_job\_a}, \mathcal{M})$ of the ACME Manufacturing example. The set is read disjunctively; therefore, observing one of its elements (which are partial valuations of the variables) is sufficient to prove that action `add_job_a`, which denotes the submission of a print job to the networked printer by Alice, has happened in the past. Note that if a variable is not mentioned, its value is regarded as immaterial. Furthermore, partial valuations extending ones that are already included in the evidence set are disregarded.

	Variable	=	Value
	first entry	=	job from A
$\vee$	first entry	=	job from A (deleted)
$\vee$	second entry	=	job from A
$\vee$	second entry	=	job from B
$\vee$	second entry	=	job from A (deleted)
$\vee$	second entry	=	job from B (deleted)

the action of interest—in this case, `add_job_a`—we can calculate the evidence sets.[30] The resulting set of sufficient evidence for action `add_job_a`, which encodes the submission of a print job to the networked printer by Alice, is presented in Table 4.4. Referring to the set $SE_a = SE(\texttt{add\_job\_a}, \mathcal{M})$, we see that there is at least one element $s \in SE_a$, e.g.,

$$(\text{second entry} = \text{job from B (deleted)}),$$

that is also included in the set E_{obs} of observed evidence, which has been acquired by looking at the print job directory and is illustrated in Listing 4.5. Therefore, an investigator must draw the conclusion that Alice had printed at least once.

Of course, this fact may be validated by manual reasoning as well, as we are dealing with a very simple case: As defined in the specification of the networked printer, a user can only submit one job at a time, and the entries in print job directory are used strictly in sequential order. Observing two deleted jobs of Bob implies that Bob must have submitted a print job once when there was another print job of Alice waiting to be processed.[31]

Besides these findings, we want to note that the submission of print jobs exhibits *AE*; this, however, is not helping to solve the investigative action. Furthermore, we cannot observe any *NE* in this case study since every entry in the print job directory could be potentially overwritten by follow-up print jobs.

[30] Conventiently, the investigative hypothesis corresponds with the execution of a single action in this case.
[31] For further reference and illustration purposes, we present the implemented solution of the ACME Manufacturing scenario using our tool in a literate programming style at `https://github.com/jgru/evidential-calculator/blob/master/examples/acme.org`.

4.8.1.3 Benefits of Our Approach

Given information on the inner workings of the networked printer, we see that it is obvious that adding print jobs is conditional in two regards: First, this operation is guarded by the number of possible print jobs per user (in this case, one) and secondly, the effect of the action (which entry in the print directory is populated) is state-dependent. In such a scenario, the method of Dewald [64] will not be able to produce sensible results. The approach proposed by Gladyshev and Patel [95] can solve the case but needs far more considerations and the introduction of specific and rather involved concepts like *evidential statements* which combine observations and a hypothesis. Moreover, their approach needs to employ a custom backtracing algorithm instead of relying on a capable off-the-shelf model checker.

Soltani and Hosseini-Seno [207] solve the ACME Manufacturing case study by employing a model checker as well, in their case mCRL2, which works with an action-based branching-time temporal logic, a fixpoint extension of Hennessy-Milner logic. Their formula describing the printer investigation is, for intrinsic reasons, rather longer than ours; for an additional comparison of the approaches, refer to Section 4.9.

Our solution directly translates the reconstruction problem to a computable property of the observable evidence, as shown above. The concisely stated and intuitive understanding of evidence specifies how a situation of facets has to be interpreted. Without the need to state complex interrelations, it is possible for the investigator to solve the SRP directly by looking at the acquired facets—in this case, the listing of deleted entries in the print job directory.

4.8.2 Determination of Relevance Based on NE/SE

Having an improved insight into the classes of reconstructability and the ability to determine facets belonging to a certain class of evidence, we now address the question of how the concepts of necessary and sufficient evidence relate to the formal notions of relevance and expressiveness developed in Chapter 3. Up to this point, the descriptions there remained largely abstract, since we did not look at how to sensibly construct an investigative knowledge base; however, using the concepts of necessary and sufficient evidence, we can change this because relevance, as we defined it (Definitions 3.3.3 and 3.3.4), and these evidence classes are closely interlinked and (almost) practically usable indeed, which should be obvious by the end of this section. We approach this in two steps: First, we instantiate the relevance relation using a system model, the previously provided LTL-formulae, and a model checker. Second, we derive the expressiveness of relevant facets based on this result.

4.8.2.1 Establishing the Relevance Relation Based on NE/SE

To facilitate recall, in Section 3.3 we established the term *relevance* formally in Definitions 3.3.3 and 3.3.4. As defined there, a facet must be considered relevant if it either supports or refutes a hypothesis h of investigative interest, which is reflected in the relations:

$$f \text{ relevant } h \coloneqq f \text{ supports } h \cup f \text{ refutes } h\,.$$

As defined in Section 4.6 of the present chapter, the notions of sufficiency, i.e., $(\bigcirc\sigma)\,\mathcal{R}(\neg SE)$, and necessity, i.e., $(\bigcirc\sigma)\,\mathcal{R}(\neg SE)$, of facets can be used to exactly determine the execution of an action in an automaton; hence, they can be used to support, respectively refute the hypothesis in regard to the execution of a certain action, as became already apparent during the fictitious investigation at ACME Manufacturing.

Since we now have both the concepts of sufficient and necessary evidence (Section 4.6) and a prototypical implementation (Section 4.7) available to derive sets of facets of these classes, $SE(\sigma, \mathcal{M})$ and $NE(\sigma, \mathcal{M})$, at hand, we have the means to establish a relevance-relation concerned with a specific system. For a given model of a deterministic automaton $\mathcal{M} = (S, \rightarrow, L)$, we can construct an investigative knowledge base $KB_{\mathcal{M}}$, as defined in Definition 3.3.1. For each action $\sigma \in \Sigma$ we generate a hypothesis h_σ capturing the intuition that "action σ was executed" forming our set of hypotheses in the investigative knowledge base. As the set of facets, we take the set $\mathsf{PVal}(V)$ of partial valuations of variables in the model. Then, we have to form the supports and refutes relations: To do so, we calculate sufficient evidence $SE(\sigma, \mathcal{M})$ for each action σ defined in the model $\mathcal{M}$ using the method described in Section 4.7.2 in the next step. In the case of sufficient evidence, the resulting evidence set has to be read disjunctively. If a facet f is included in $SE(\sigma, \mathcal{M})$, then it supports the hypothesis that σ happened, hence it generates an element in the supports relation such that f supports h_σ. Similarly, the elements in $NE(\sigma, \mathcal{M})$ generate the refutes relation. Put simply, we infer the relevance relation based on the evidence sets by looking in which evidence sets related to which action the facet occurs so that we can map facets to hypotheses. The result $KB_{\mathcal{M}}$ is an investigative knowledge base specific to the system model $\mathcal{M}$ containing the relevance relationship of facets and hypotheses practically and effectively, which remained rather elusive up to now.

4.8.2.2 Calculating Expressiveness Based on NE/SE

Given the construction of the investigative system described above, the expressiveness $H|_f$ of a facet, i.e., a partial valuation of the automaton's variables, is then simply the set of all hypotheses on action executions that can be proven or refuted by observing that facet, as defined in Definitions 3.3.3 and 3.3.4. As a quantitative metric, the expressiveness ratio then calculates what percentage of hypotheses in a given set H regarding action executions can be decided by retrieving a given facet.

To illustrate this, we again refer to the example program presented as Listing 4.4 and determine the respective expressiveness of the single facets that contribute to assessing at

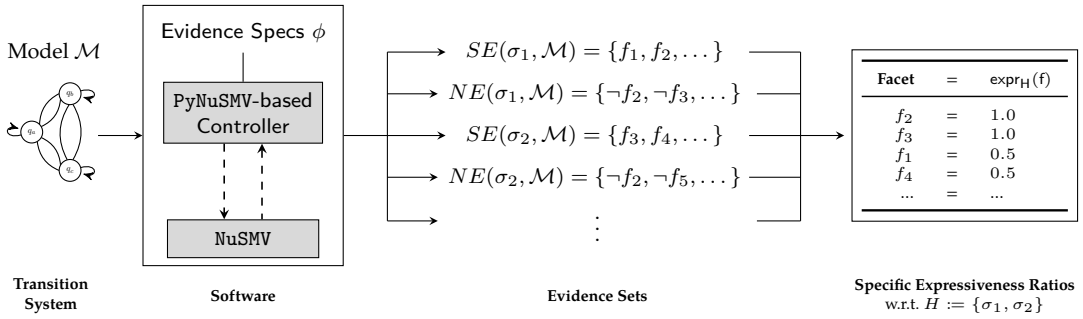

Figure 4.13: Overview of the calculation of expressiveness ratios of facets, i.e., partial valuations that are part of evidence sets as retrieved by the proposed NE/SE-approach.

least one hypothesis regarding the execution of an action in the automaton. We do so by counting their ocurrences in the *NE* or *SE* sets first and employing the formula presented in Definition 3.3.6, i.e.,

$$\mathsf{expr}_{\mathsf{H}}(\mathsf{f}) = \frac{|(H|_f)|}{|H|}\,,$$

afterward to calculate the relative expressiveness ratios, which are shown in Table 4.5.[32] There it becomes apparent that some facets, e.g., $c = 1$ or $d = 1$, exhibit a higher relative expressiveness ratio and, thus, might be generally more useful.

Table 4.5
Relative expressiveness ratios for facets of the example program Lst. 4.4. Only facets are considered that contribute to the assessment of at least one hypothesis, thus the expressiveness ratios $\mathsf{expr}_{\mathsf{H}}(\mathsf{f})$ of facets $F' = \{f \in F : |(H|_f)| \geq 1\}$ w.r.t. the set of hypotheses $H = \{\mathtt{a0}, \mathtt{a1}, \mathtt{a2}, \mathtt{a3}\}$ are listed.

Facet $f \in F'$	$\mathsf{expr}_{\mathsf{H}}(\mathsf{f})$	$\lvert(H\rvert_f)\rvert$
$a = 0$	1.00	4
$a = 1$	0.25	1
$b = 1$	0.50	2
$b = 0 \wedge d = 0$	0.75	3
$c = 1$	0.75	3
$c = 0$	0.25	1
$c = 1 \wedge b = 0$	0.25	1
$c = 1 \wedge b = 0 \wedge d = 0$	0.25	1
$d = 1$	0.75	3
$d = 0 \wedge c = 0$	0.25	1

This application of the concept shows how forensic scientists can quantify facets' expressiveness in a system with finite actions. Based on this knowledge, the investigator is empowered to consider those facets first that allow assessing multiple hypotheses at once,

[32] Note that we refer to the relative expressiveness ratios solely because we could additionally include the negations of the hypotheses regarding action executions, which would add unnecessary complexity to the example.

as indicated by their higher expressiveness ratios, which can be potentially enough to solve the case. Concerning the Facet-oriented Criminalistic Cycle (FoCC) Fig. 3.2, the availability of the expressiveness means requiring fewer iterations of the Facet-oriented Criminalistic Cycle until achieving the completeness-property (Definition 3.4.5) of the facet collection.

4.8.2.3 Integration with the Concept of Relevance

After having shown the practical interplay, we now briefly summarize how the *NE* and *SE* integrate into the notion of relevance on a conceptual level:

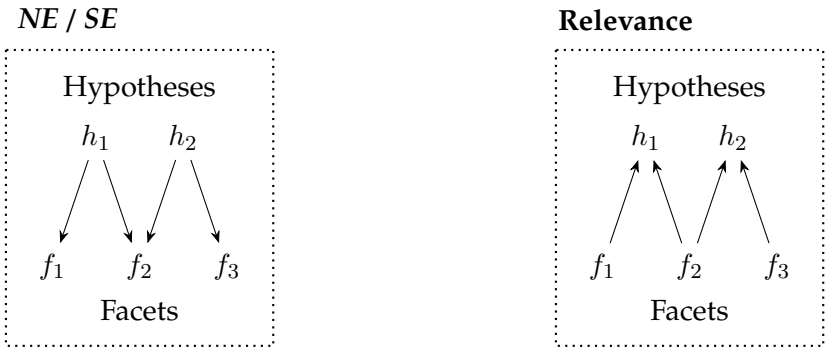

(**a**) *NE* and *SE* map hypotheses to facets (**b**) Relevance maps facets to hypotheses

Figure 4.14: Duality of necessary evidence and sufficient evidence regarding relevance. The figure illustrates the different views that the concepts enable: Given a hypothesis, the necessary and sufficient facet sets (*NE* / *SE*) allow to directly find relevant facets, while the relevance of facets allows to determine assessable hypotheses based on a given facet.

Given a system model, we derive the facets based on the partial valuations and the hypothesis based on the available actions in the system. *NE* and *SE* serve then as means to establish the supportsand refutesrelations. Those relations are the vital parts of the investigative knowledge base since the evidence sets determine which trace situations allow which conclusions in regard to past events. The newly established notion of *NE* and *SE* offers basically a dual view of the facets' relevance. On the one hand, *NE*/*SE* effectively maps (hypotheses of) actions to facets. On the other hand, the notion of relevance exhibits a duality and maps facets to hypotheses, as illustrated in Fig. 4.14. Then, the expressiveness of facets $H|_f$ can be considered a meta property that allows us to find the assessable hypotheses.

4.9 Discussion

The approach presented in the present chapter aims to unify and generalize the problem of event reconstruction. We contribute to a better understanding of digital evidence, which allows precise reasoning about the quality of traces using, for the first time, linear-time temporal logic. In regard to event reconstruction and the SRP, the proposed method is

actionable and practically usable, since it is solely concerned with observable parts of traces—variable assignments in an automaton constituting the facets. Instead of abstract statements about past states of FSMs, our method puts the SRP into the center, which is regularly of principal interest in forensic analyses.

4.9.1 Differentation from Related Work

By considering the state information, we have largely extended and improved the approach of Dewald [64], which is limited to action-induced evidence that is calculated using set theory neglecting state information. To deal with the state explosion problem, we resort to an off-the-shelf symbolic model checker that allows checking properties of a system without building the complete state graph.

Regarding direct automata-theoretic methods for forensic event reconstruction, we have already provided a technical comparison with the work of James et al. [130] in Section 4.2. Another approach that is closely related to ours, mentioned already in Section 4.2, uses an action-based form of the modal μ-calculus as the temporal specification language [207]. The notion of sufficient evidence remains implicit in the cited work, and necessary evidence is not considered. On a technical level, we have already noted that formulae specifying sufficient evidence in the mentioned flavor of the μ-calculus are inherently longer than our LTL formulae. This is partly due to a standard tension between labeling transitions or states; these forms of labeling are interconvertible by standard methods, which however incur blowup by a linear factor (indeed, recall that transition labels are encoded as atoms in our approach). Specifically, the formula templates given by Soltani and Hosseini-Seno are of linear size in the total number of actions, while our formula templates in LTL are of constant size. This is relevant insofar as model checking in either logic is (roughly) exponential in the formula size. Also, the μ-calculus is a branching-time logic, while it appears that for purposes of reconstruction of past events, linear-time formalisms that restrict attention to sequences of events, such as LTL, are inherently more suitable.

Following the strategy of separating concerns, we can swiftly determine the execution of an action once we calculated the evidence sets. In addition to that, the calculation of those sets beforehand provides clear guidance for an investigator on where to look to prove or refute the hypothesis of the execution of a certain action, and what to conclude on which observation.

4.9.2 Limitations

Although the presented method seems to be a helpful approach, there remain still various challenges. A severe drawback is the high time complexity of the calculation of the evidence

sets, which grows exponentially to the number of variables (the resulting number of partial valuations) and linearly to the number of actions (the number of sets to be computed):

$$O\left(\prod_{v \in \text{Variables}} (values(v) + 1) \times |\,\text{actions}\,|\right)$$

However, we optimize the calculation by trying small partial valuations first and discarding partial valuations extending ones that are already included in the evidence set.

Furthermore, technical limitations restrict the size of processable models. In our experiments with NuSMV, we were only able to handle up to 2^{1024} states—which sounds astronomic but effectively means one can merely use at most 1024 Boolean variables, illustrating that at present, there is no actual escape from the state explosion problem. Another hard problem is to infer apt models, i.e., state machines of the system under investigation. Currently, this involves human reasoning, a universally applicable and largely automated approach has not been proposed yet. In that regard, it is important to note that the results are only as good (and accurate) as the employed model itself. If the model is flawed, the results will most certainly be misleading as well. Thus, there is the possibility that involved parties, e.g., the defendant's lawyers, might challenge the model's correctness to undermine the conclusions.

4.10 Summary

In the previous chapters, we elaborated on the importance of event reconstruction as a salient step in every (digital) investigation. The stakeholders of forensic investigations require highly accurate results from analyzing (often) complex digital evidence. In the quest to fulfill this requirement, various approaches working towards a formal solution have been developed over the past 20 years. Nevertheless, the problem remained unsolved in practice, as we showed by scrutinizing the research efforts in this field in Section 4.2.

After introducing the necessary background material, i.e., linear-time temporal logic, model checking, and the guarded command notation, to ease the approach of the chapter's topic in Section 4.3, we describe the specific reconstruction problem and the approach by Dewald [64] to solve it. To do so, we had a closer look at his method of characteristic evidence calculation. While this provides one solution to the SRP and is (even) applicable in real-world scenarios, e.g., for fingerprinting file system metadata or log entries, it is incomplete. Its succinct set-theoretic approach has the advantage of easy understanding and effortless implementation but comes at the cost of incompleteness due to the neglect of state information. This insufficiency is shown in Section 4.4, where we provide two examples showing the incorporation of state, suggesting that there might be weaker conditions sufficient to reconstruct events. In Section 4.5, we tackle the solution of the SRP in LTL. First, we briefly critique preliminary thoughts and then derive our solution to the SRP using temporal logic.

Building upon this foundation, we present the classes of evidence by employing an automata-theoretic approach in Section 4.6. From a methodical point of view, our proposed

method of formalizing digital evidence with the help of LTL formulae generalizes forensic reconstructability and enables investigators to reason about an existing trace situation at a digital crime scene more concisely than before. There are three classes: Action-induced evidence only considers the *immediate* effects of an action and, hence, comes close to what Dewald [64] deemed *characteristic evidence*. Sufficient evidence, i.e., expressed in LTL $(\bigcirc\sigma)\,\mathcal{R}(\neg SE)$, allows to prove the execution of the target action σ in the automaton by containing all the facets that can only be observable if the action has been executed. Necessary evidence, i.e., expressed in LTL $\bigcirc\big(\Box(\sigma \rightarrow \Box NE)\big)$, allows to refute the execution of the target action by stating all the facets that must be observable in all subsequent states. The differences and improvements in comparison with the *CE* method are shown with three examples before we present the implementation of a prototypical tool that relies on the widely used model checker NuSMV to calculate evidence sets. To this end, the tool loads the model into the model checker, retrieves the set of variables from the model, and checks for each partial valuation of these variables if they constitute *NE* or *SE* w.r.t. to each action.

Aiming to illustrate its applicability, we apply the concept using our tool to Gladyshev's "ACME Manufacturing" benchmark example. It is shown that matching the observed facets with the calculated sufficient evidence set provides a straightforward solution to the investigative question (Section 4.8.1). Afterward, it is shown that the concepts of necessary and sufficient evidence can be used to establish the relevance relation and infer expressiveness of facets (Section 4.8.2) since it provides means to assign facets to the supports and refutes relations of an investigative knowledge base. This connection explicates the duality of the concepts behind the classes of evidence and the concept of relevance and expressiveness, as defined in the previous Chapter 3—the former maps hypotheses to facets while the latter maps facets to hypotheses. Lastly, we discussed our approach in the light of related work summarizing our improvements regarding the CE method by Dewald [64] and the strengths of our method regarding another formal event reconstruction approach [207], i.e., the implicitness of the SE concept and the absence of a notion of *NE* but also technical concerns such as the constant size of formulae templates, and general considerations that linear-time formalisms, as employed by us, seem inherently more suitable than branching-time logic. Besides that, we identify several limitations of the proposed method. First of all, these are the high runtime complexity, the limited model size, and the challenge of constructing apt models.

However, from a conceptual point of view, knowing about these classes and showing how to calculate the evidence sets provides means to assess the meaningfulness of the evidence and its inferrable implications in a general way. It defines under which circumstances what parts of the observable traces can be used to reconstruct or refute past activities. It illustrates a technical realization of event reconstruction based on state predicates, thus forming a translation of the reconstruction problem and the system under investigation. In particular, we introduce clearly defined evidence classes as a precise way of communicating the findings and their significance for the case. This improved notion of digital evidence and the technical realization presented in this chapter can support investigators to better reason about crucial questions of what happened and who did it.

5 Phenomenon-specific Digital Evidence

5.1 Introduction

Knowing about necessary and sufficient evidence is salient to reconstructing single events. These concepts constitute a leap toward understanding what comprises "sufficient digital evidence" when confronted with a computer program. For modern-day cybercrimalistics, however, it is not nearly enough to consider a single computer program and reconstruct single events but to chain several investigative actions together to identify perpetrators and provide evidence for their commission of the deeds. Hence, there is a need to develop an approach for uncovering relevant and expressive traces[33] geared toward the real-world requirements of cybercriminalistics, as we aim for in the present chapter.

The Combat Against Cybercrime. This is an especially pressing question since cybercrime is a growing domain of criminal activity with substantial economic and political impact [89, p. 44], and the fight against cybercrime has become a top priority in many national security policies like those of the United States [222] and the European Union [221, p. 3]. Consequently, considerable investments have been made into the anticipation of and response to cybercrime. Nonetheless, the global cybercrime industry still appears to flourish, developing new and innovative business models at an alarming rate. And while there may be spectacular stories of successful cyber operations by law enforcement agencies such as the arrest of the infamous carder Roman Seleznev in 2014 [61] or the takedown of *Emotet*'s botnet infrastructure in January 2021 [77], a considerable percentage of cases end with unsuccessful preliminary investigations against unknown individuals. For example, the statistics of the Federal Criminal Police Office in Germany show a significant rise of case numbers in cybercrime from 63,959 in 2012 up to 124,137 in 2021, while the clear-up rate (the percentage of crimes that have been "cleared up" by the police) stagnates at a low level of 29.9 % [80, p. 6].[34] The observation that Anderson et al. made in 2013, namely, "we are extremely inefficient in fighting cybercrime", apparently still holds today.

Looking at how national legal systems dealt with new criminal phenomena in the past shows however that the inefficiency in dealing with cybercrime is not a law of nature. Together with law enforcement, the scientific community must actively research and

[33] Note that the term "trace" is used in this chapter without distinction into tangible traces and observable facets, as discussed in Section 2.4.2 (Definition 2.4.1), because it improves the readability of the text. At the same time, the exact differentiation is quite negligible at this stage.

[34] Due to the change of capture modalities in the statistics, we resort here only up to 2021 in order to ensure comparability.

develop new methods that are more effective against these new types of crime. And indeed, the scientific research community has contributed a multitude of examination methods, mostly from applied natural sciences and medicine, establishing domains of *forensic* science. In contributing such methods, forensic scientists usually help to translate the legal questions of relevance for the court to a scientific question towards evidence that can be examined in a laboratory [125]. The answers to those scientific questions can be used by judges to determine guilt or innocence. However, as we discussed in Chapter 3, determining the relevant pieces of evidence is itself a difficult task, and using the analysis results of scientific experiments to assess a single hypotheses are only pieces within a larger puzzle. In fact, the (cyber)criminalists need to have a more holistic view since their task is much more affected by procedural aspects to solve the criminalistic task by apt decision making [213, p. 147].

The Challenges of Cybercriminalistics. Compared to rather primitive offenses, combatting high-tech crimes poses additional challenges for (cyber)criminalists. To support this endeavor, the scientific research community has contributed much insight into phenomenological, technological, and economic aspects of the underground economy. As we discussed in Section 2.2.2, computer scientists have also developed the branch of *digital* forensic science (Definition 2.2.4) dealing with the many delicacies of digital traces [38]. However, the whole field of *cybercriminalistics* (Definition 2.2.6), in particular using traceology of digital traces and digital criminalistic tactics concerned with process-related issues and operational aspects, is not well understood yet.

We identified two primary deficiencies in the literature, which should be preferably tackled and pave further research efforts in the field of cybercriminalistics: The first concerns the scientific production of abstract but (in a specific case) rather useless process models to perform a forensic investigation, see Pollitt [185] for an overview. All these models lack guidance on how to perform "the next step" in an investigation. If specific guidance is provided, then only a specific step or small area mainly of technical nature is addressed, like it is the case with Venter [236], who, for example, created process flow diagrams for first responders at electronic crime scenes. Nevertheless, to be able to conduct investigations of high-tech crimes, more direction is needed in general. Undisputedly, it is the intent of general models to act upon a high level of abstraction to provide a simplified view of the course of the investigation. But there appears to be a lack of mechanisms to bridge the gap spanning between abstract thoughts and the actual concretization of the model in a real-world investigation. A main factor for this insufficiency seems to be the missing conception of relevant technical traces and their actual use. As we showed in the previous chapters, it is incredibly challenging to identify relevant traces for answering investigative hypotheses. As long as these remain latent or obscure, it seems reasonable to say that the effectiveness of cybercriminalistics will be impeded.

The second deficiency refers to the lack of attention to *operational aspects* of an investigation: Many people think that most cases can practically be solved in isolation and after the crime is committed.[35] While this is accurate for many cases, there are still procedural

[35] Though, Berger [15] states that "[c]riminalistics is reasoning backwards"; his understanding of "criminalistics" seems to be more aligned with the understanding of Kirk [139, 140], who considered it to be a holistic

delicacies. Furthermore, it might not hold for ongoing crimes committed by transnational organized crime groups, as is the case with many cybercrime phenomena. Such criminal activities require an increased *operationalization* of the investigation, placing more emphasis on tactical considerations and proactive procedural management that can be regarded as an accompaniment of serious crime [26, pp. 488 f.].

The iterative manner of investigations stresses the importance of operational and temporal aspects. This means that it is not enough to examine one scene, but rather follow one lead to another, while performing different investigative measures and incorporating criminalistic reasoning and adapting the course of action to uncover secondary scenes. Here, the most salient aspect is to chain investigative measures according to the questions of interest. Thinking of the identified subproblems of the criminalistic task relates to the use of relevant traces to put up hypotheses of pertinence orderly for the investigation. None of the existing process models of digital forensic investigations gives any direction on the incorporation of such procedural aspects during the investigations. However, given the various anonymization and obfuscation measures used by cybercriminals and the volatile and transient qualities of digital evidence, there seems to be a need to focus on operative measures in the domain of cybercriminalistics.

5.1.1 Contribution of the Chapter

This chapter proposes a method to identify and use relevant traces and hence apply the theoretical considerations of Chapter 3 in practice. Furthermore, the approach aims to systematically bridge the gap between high-level process models found in literature and low-level operational measures in cybercriminalistics. To accomplish this, we suggest the use of phenomenon-specific knowledge as an intermediary level, which extends superordinate constructs and helps to apply an abstract model to the specific individual case (see Fig. 5.1).

Similar in spirit but different in application to Tanner and Dampier [220], we propose to utilize an extension of concept maps, namely *cognitive maps* (*CMs*), to encode the phenomenon-specific knowledge. The deviation from the formal investigative knowledge base, as defined in Definition 3.3.1, allows to relate traces, investigative measures and their outcome to another. Our approach, therefore, literally "maps" the digital crime scene of specific phenomena into representations that encode knowledge mined from literature or interviews with domain experts in a structured way. As a result, this procedure provides an easily accessible and visualizable thinking model. This enables the investigator better to understand the phenomenon-specific problem space, triage evidence items, and derive a prioritized plan of action for targeted acquisition and analysis of hypothesis-relevant and presumably promising artifacts in the bulk of digital trace material. This approach of encoding information enables investigators to derive potential chains as well as rings of evidence on the basis of an offense and find a solution to the criminalistic task in a structured almost guided manner. Operating on this meso-generic abstraction level means that the approach is both generic enough to be of help for all cases of the phenomenon

cross section of forensic science, which is not according to the well-founded Definitions 2.2.3 and 2.2.6.

and specific enough to derive the following actions and work towards the identification of a perpetrator.

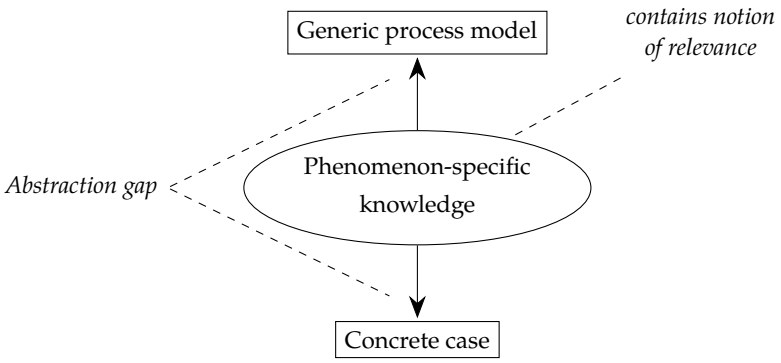

Figure 5.1: The abstraction hierarchy and the use of phenomenon-specific knowledge as an intermediary step between abstract process models and concrete investigative steps.

To illustrate the proposed method and the benefit of its results, we present an instance of a CM for the technical investigations in the field of botnet crime whose correctness and completeness have been validated, firstly, by interviews with domain experts, and secondly, by applying it to two real-world cases in this field. Since botnet crime is one of the major areas of organized cybercrime in which approaches of cyber attribution and classical investigations have to work hand in hand, having a unified view of the technical knowledge of this phenomenon seems therefore particularly helpful.

5.1.2 Chapter Outline

The remainder of this chapter is structured as follows: First, we look at related work regarding practical cybercrime investigations and the general use of visualization in investigations in Section 5.2. Second, we identify and describe an abstraction gap in case work in Section 5.3. In Section 5.4, we present the proposed approach of using phenomenon-specific knowledge bases to bridge the identified gap by having a conception of relevant traces. Then, we illustrate and evaluate our method by applying it to the phenomenon of botnet crime in Section 5.5 and discuss it in Section 5.6. Lastly, Section 5.7 summarizes the findings and gives an outlook to possible future work.

5.2 Related Work

Earlier in Chapter 2, we set the scene and presented many precoursing thoughts on (digital) investigations and cybercriminalistics, especially in Sections 2.2 and 2.3. These preceding explanations serve as core references to related work, which will now be supplemented by more practically oriented information associated with the chapter's content. First, we briefly summarize the cybercrime investigation process as presented in the practical textbook on cybercrime investigations by Bandler and Merzon [11]. Second, we take a

look at the use of visualization to aid investigations for paving the way of our approach to represent phenomenon-specific knowledge.

5.2.1 Practical Cybercrime Investigations

In general, investigations are described by the use of process models to provide a structured framework for understanding and conducting the flow, as well as the procedural steps of forensic investigations.

Most research referenced in Section 2.3 was fairly theoretical and remained detached from problems encountered in cybercrime investigations. A very practical orientation, however, is employed by Bandler and Merzon [11], former New York County District Attorneys. In their textbook, they provide actual guidance on investigating cybercrime offenses. They explain methods, procedures, and aspects of U.S. law in this field and also proposed a process model, called the *Cybercrime Investigation Process*, which is divided into five phases: the *inititation phase*, the *records phase* concerned with analyzing basic subscriber and usage data of relevant online accounts, the *data search phase*, potentially the *wiretap phase*, and lastly the *physical world phase* [11, pp. 232–240]. Within these phases, the investigations follow a cyclic approach consisting of exploration (*Explore & Hunt*), acquisition of data (*Obtain*), their analysis (*Analyze*), and preparation for presentation (*Hold and prepare as evidence*), to finally fuel the planning of the next iteration (*Plan next hunt*) [11, p. 180]. They describe how to generally carry out individual investigations aiming to identify perpetrators. In order to assign online identities, such as nicknames, email addresses, and so on, to an individual in the real world, they developed a so-called attribution process consisting of six steps, which is identified by the acronym *ID-PLUS*. The process is comprised of six phases, i.e., *Isolate* → *Determine* → *Pursue* → *Link* → *Uncover* → *Summarize and articulate*, which can be summarized as follows [11, p. 268]: Starting from a criminally relevant event to be isolated (*Isolate*), directly connected personal information, which is called pedigree by the authors is derived (*Determine*). Then, an fragmentary online identity is created by cyclically searching connections to further identifiers and the information generated from them (*Pursue*). At best, these indirectly connected profiles and accounts can be used to create a link to a person in the real world (*Link*). If this has been accomplished, investigators need to search for connections between the crime suspect and the crime. Furthermore, the remaining identifiers must be searched in the opposite direction to uncover more links (*Uncover*). Finally, the uncovered chains or rings of evidence must be articulated comprehensibly (*Summarize*) [11, pp. 268–273].

5.2.2 Visualization in (Digital) Investigations

Not only in project management but also in forensic science, visualization is considered to be helpful in order to keep the overview, manage tasks, grasp processes, and support the decision-making processes in forensic investigations. The *Forensic Field Map* proposed by van Beek [231] is a rather recent instance of this observation since it provides an abstract (visual) framework for discussing the broad forensic field by offering a structure

to visualize knowledge chains in the context of the criminal justice system. It incorporates methods, tools, procedures, and results over the different phases, ranging from preservation, acquisition, extraction, relation, to evaluation.

While this again provides an abstracted view of the matter and seems to constitute primarily a tool for managing and communicating, we focus on actual investigations: Traditionally, a large part of the criminalistic work is concerned with the collection, examination, and interpretation of the traces at a crime scene. To do so, the creation of sketches or maps of a physical scene is an integral part of the criminalistic procedure. Evidence items are drawn in relation to each other to document the scene [74, pp. 126 ff.], gain an understanding of it to form apt hypotheses, and, ultimately, to come up with versions of the course of the deed (see Section 2.5). Hence, it is not surprising that a visual representation of significant findings and the available case information has been found helpful in the related context of crime analysis [43, 242], as well as in investigative work [52, 220, 236] to aid criminalistic thinking. Link charts plot entities and their relations to illustrate the flow of illicit trafficking, complicity, geospatial references, and much more. By doing so, they ease to grasp connections and gather an overview of case data [197, p. 195]. Link analysis is thus a standard method to process investigative data efficiently and effectively. Its practical usefulness is suggested by the multiple discussions of those in application-oriented textbooks for investigators and especially crime analysts [72, pp. 35 ff.].

Since digital investigators operate on an elusive, abstracted crime scene with many distributed systems, where they cannot simply apply any "geographic" mapping, such node-link representations have also been profitably employed in the past: For example, Tanner and Dampier [220] proposed an extension of the DFRWS investigative process model [180, p. 17] by using *concept maps*, a hierarchic visualization method "with the most general, most inclusive concept at the top, and the most specific, least general concepts toward the bottom" [171, p. 177]. Tanner and Dampier augmented each of the six generic phases of the original process model with increasingly refined concepts that "clarify how the evidentiary items are related to each other" [220, p. 298]. The concept maps contain much more information than simple checklists and can be helpful in the general recollection of key points and their relation to one another within a single phase. However, the model is a mere collection of information in a hierarchical and condensed way and does not help in identifying the "next step" to be performed in an investigation. Furthermore, the crux of the matter remains to be the knowledge of the relevance and meaningfulness of the traces that may be encountered, which has not been addressed so far.

5.3 Investigative Knowledge Bases Bridging the Abstraction Gap

It appears to be the consensus that the criminalistic task has to be seen as a general cognition process rooted in the context of guilt and innocence, where a solid parallel to research work can be drawn. This is backed by the established application of the *scientific method for criminalistics* by De Forest [55], as explained in Section 2.4.3.1. To briefly summarize, we can say that some course of events has been observed, which sparks the need to answer the central case-specific questions. In order to answer these questions, hypotheses are formed

to be either falsified or verified later. With this in mind, the framework for the investigative process defined by various abstract models and the case-specific proceedings can be seen as two extremes, constituting the macro level on the one side, and the micro level, on the other. The notion of how hypotheses have to be strung together and which traces can be used for assessing those, as it has been captured in a knowledge base, comprises a meso-generic level of abstraction.

Necessity of an Intermediary Step. Regarding crime scene work, traditional criminalistics differentiates between two approaches. On the one hand, the systematic search [148, pp. 122 ff.] of a crime scene describes the seamless search by employing a specific search pattern without limiting the area to examine, as the name implies. A heuristic approach, on the other hand, narrows the search area after a mental reconstruction of the possible course of actions of the deed [215, pp. 27 f.]. In dealing with cybercrime offenses, investigators must deal with a widely (sometimes globally) spread and highly abstracted crime scene with first and secondary scenes [34, p. 6]. These scenes house a wealth of digital artifacts and open up various investigative possibilities, which renders a strictly systematic approach typically inapplicable since "the detailed examination of every possible piece of technology used would be an overwhelming and impossible task" [123, p. 65]. Given the existing process models' abundance of guidance on the actual concretization of investigations, an urgent need arises to incorporate an intermediary step.

Knowledge Bases as an Intermediary Step. As such a step, we imagine the use of investigative knowledge bases, as described in Definition 3.3.1, that can guide the investigators by providing hints on how to answer specific investigative hypotheses; hence, they constitute a meso-generic level of abstraction and should give direction on the question of where to start the search as well as how to continue the process from one finding to another.

This observation results in the following abstraction hierarchy for the investigative work:

1. A general process model
2. Investigative knowledge base
3. Case-specific concretization

Thus, a general process model like Hunton's Cybercrime Investigation Framework [122] provides the overall skeleton. It defines the key elements and general phases of the investigation. The utilization of knowledge bases then brings life to those phases, which eases the task of concretization by providing a meso-level abstraction. So, we understand them as being a bridge to span the identified gap. The actual casework on each unique offense will be steered in the right direction, which augments other strategies like case-based reasoning [142] to be employed on this level to solve the criminalistic task at hand.

Although we had developed a notion of a *formal* investigative knowledge base for the sake of gaining general insight in Chapter 3, two significant questions remain concerning the application to real-world investigations: First, how should we determine relevant and

expressive traces in real-world settings? Second, how should we encode the investigative knowledge to reflect the criminalistic proceedings? Both will be addressed in the next section.

5.4 Phenomenon-specific Knowledge Bases

The fact that various criminal phenomena exhibit certain consistent specifics we see the potential to infer relevant traces and hence, build an investigative knowledge base. However, we need to clarify first what we grasp as a *criminal phenomenon*: A criminal phenomenon appears to be best characterized by its goal and the tactics, techniques, and procedures (TTPs) employed by the perpetrator. The concept of TTPs constitutes a rather recent approach developed in terrorism studies and threat intelligence to record and describe behavioural patterns of criminals (in this context often referred to adversaries or threat actors) [133, p. 2]. They describe the actors' overall *tactics*, the specific *procedures* observed in the process, and the *techniques* used in specific incidents or deeds. So, TTPs depict the behavior required for the commission of the crime. Thus, we consider those to be a finer-grained version of or a supplement to the concept of what is traditionally deemed *modus operandi* [83, p. 189] in criminalistics [240, p. 187]; therefore, the concept of TTPs appears apt to grasp and group phenomena.

The knowledge of TTPs, on the one hand, paired with investigative experience and scientific research, on the other, constitutes criminalistic knowledge. The structured preparation and processing of this can then be used to gain an understanding of which traces are relevant to which investigative (sub)question. Having extracted this information, one can build an investigative knowledge base specific to the phenomenon in question, i.e., what we deem a *phenomenon-specific knowledge base*, that enables the introduction of a meso-level abstraction in investigations aiming to assist the investigators in the field effectively to understand and "process" the crime scene.

5.4.1 The Nature of Cognitive Maps of Crime

In Section 5.2, it was noted that node-link representations are commonly used in investigations. Given the fact that a formal investigative knowledge base is rather abstract, we propose to employ a conceptual mapping, remotely analogous to the mapping of crime scenes, to store investigative knowledge and, thus, to aid criminalistic thinking in the spirit of the before-mentioned visualization approaches. To recall, the task of a criminalist is to "analyze the problem and understand the principle in order to arrive at a correct interpretation of the criminal act", as identified by Kirk back in 1963 [139, p. 368]. While the view of the court and the forensic accessors on the matter resembles backward reasoning, the focus of the case investigators is much more affected by procedural aspects and explorative thinking to solve the problem by apt decision-making [213, p. 147]. Bringing those statements together leads to the argument that the criminalist has to be empowered

to ask the right questions at the right time and in the right order with the appropriate degree of granularity.

Axelrod [9] modeled the decision-making process of policy makers with the help of CMs in the form of signed digraphs, where the nodes represent concepts, which are connected by edges and denote causal links, whose sign encodes the causal influence [9, p. 5]. Such a causal map-based method has also been employed to help incident responders reason about intrusions [143], which was refined later to resemble probabilistic relations between collected evidence and events occurring during security incidents in a formal way [192]. These works, as well as the topic map approach of Tanner and Dampier [220], inspired us to encode relevant traces extracted from phenomenon-specific knowledge in a similar way and create a map of the digital crime scenes, where trace material, investigative actions, and resulting artifacts are interconnected. This approach differs from Tanner and Dampier's concept maps in the way that the resulting cognitive map is not hierarchically organized and does not intend to enhance a process model's phases by presenting graphical views of checklist activities. Instead, we create a knowledge base capturing a qualitative model of (a part of) the soft knowledge domain of cybercriminalistics as a node-link representation. This representation should serve much like cartographic material as an aiding tool to navigate through the chaos of real-world investigations since it provides a view of which traces are relevant to solve the case.

Such a CM storing investigative knowledge of relevant traces in the wider sense builds up on an extended version of the linkage theory [148, pp. 114 f.]. However, by the incorporation of investigative actions, the procedural associations are shown, and the relations between trace materials and the potentially resulting artifacts are enriched. This supports the linkage in the criminalist's mind of how things belong together and serves as a form of knowledge representation of relevant and expressive traces in a visual form—representations that have already been found helpful for related tasks in the domain [220, 236, 242], as elaborated in Section 5.2.2. The availability of such a concretized model aids the investigator in identifying and prioritizing investigative actions to uncover expressive traces and assessing their potential success and relevance by categorizing and linking sources of insight, the needed action and the expected results. It is not intended to serve as a procedural blueprint but a map of options to process or find relevant traces instead—explorability is an inherent feature. Such a presentation condenses phenomenon-specific knowledge of expressive traces, while it remains relatively easy to grasp. It supports scanning all investigative options, while the use of linear checklists is omitted in favor of placing the focus on interconnected tasks and objectives [12, p. 148], which stimulates an associative and radial style of thinking.

5.4.2 The Construction of Cognitive Maps of Crime

We identified two main challenges during the construction and refinement process of such a meso-generic CM of the crime scene: First, all relevant information has to be found, analyzed, and mined. Second, the phenomenon-specific knowledge has to be persisted while still being easily accessible to the investigators—preferrably both in a textual and visualized form.

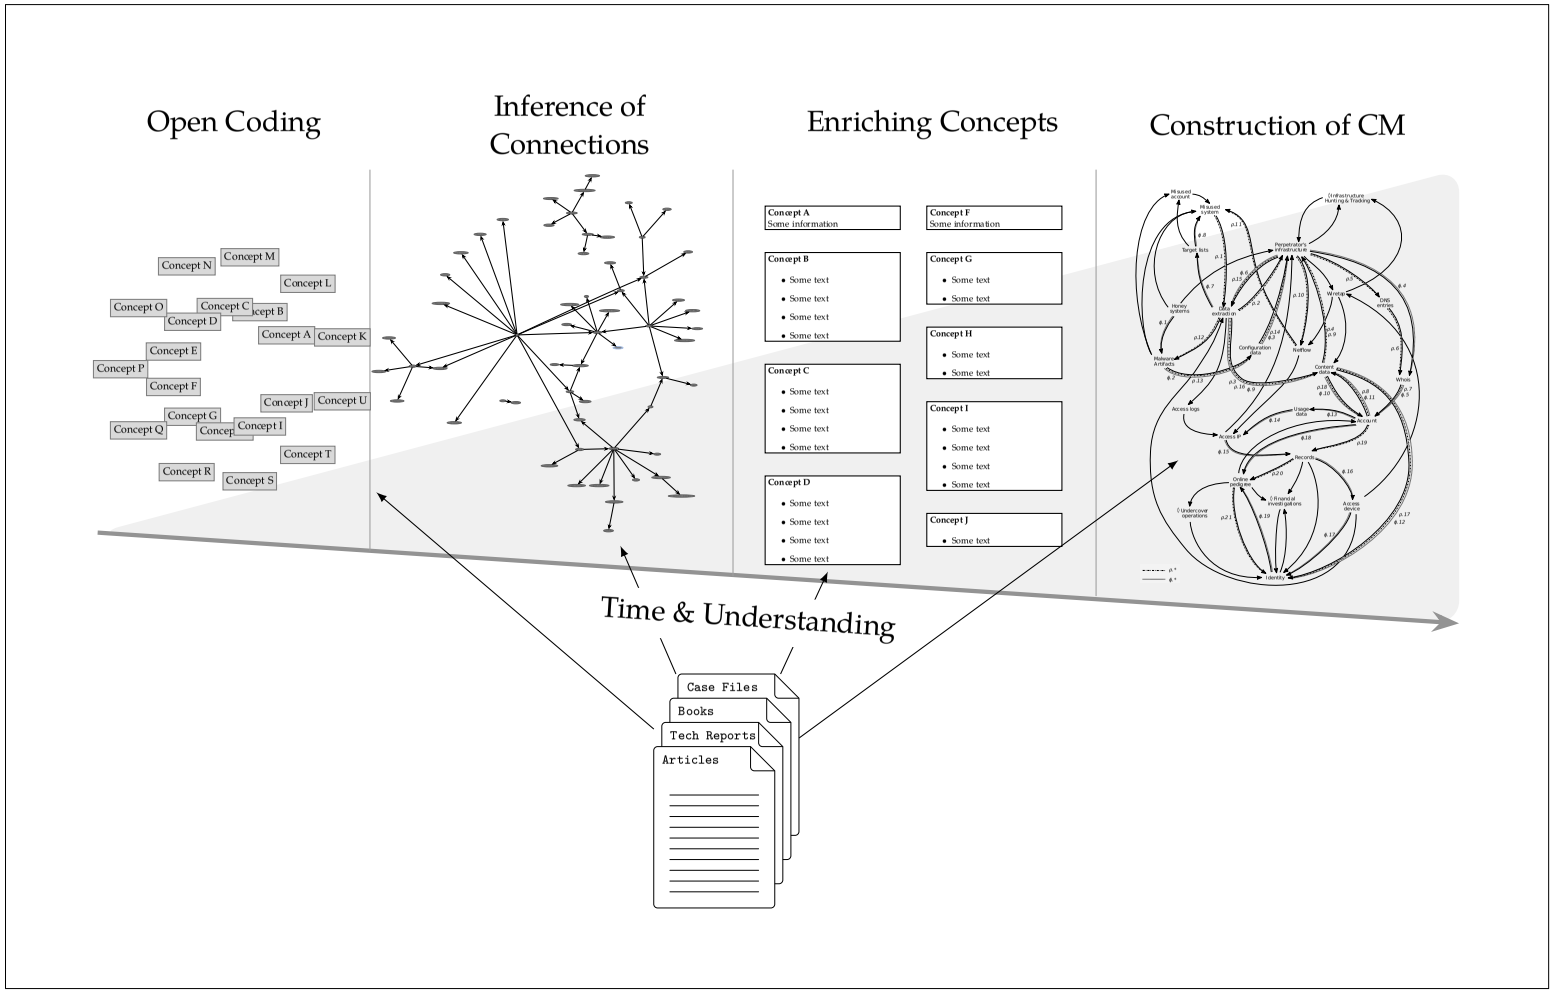

Figure 5.2: Construction process for the exemplary CM of botnet crime (Fig. 5.4);
In the first step, an open coding phase should be used to extract concepts. In a second step, we inferred the relation of the concepts to one another, and in a third step, we enriched the concepts, i.e., we gathered and documented additional information regarding each concept, such as documentation of how to interpret or analyze the trace in question or conduct an specific investigative measure. In a fourth step, we finally constructed and visualized the CM. Note that this visualization of the construction process was inspired by an illustration of the grounded theory method by Griffin and Richardson [97, p. 5].

5.4.2.1 Mining Cognitive Maps

The mining of CMs relies heavily on the availability of domain knowledge and logical reasoning. CMs can be obtained by four methods: (1) The modeler can question experts, (2) the experts could model the map themselves, (3) the modeler could analyze data that shows causal relationships, or (4) the modeler analyzes and evaluates documents [177, p. 47].

For the CM of botnet crime, which will be presented later as an example, we followed method (4), i.e., we evaluated scientific literature, textbooks, and technical reports on the attribution of cyber incidents. In addition, the knowledge base could be refined further by resorting to testimonies respectively guilty pleas, as well as the targeted questioning of informants, the analysis of underground forums[36], and other options to learn about the offenders' tactics, techniques, and procedures. To gather, analyze, and structure the information contained in theses sources, we propose to employ a four step procedure, as depicted in Fig. 5.2: The initial phase involves employing open coding to extract fundamental concepts. Following this, we move to a subsequent phase where we deduce the interconnections between these concepts. In a third phase, we enhance these concepts by acquiring and documenting supplementary information for each one. This supplementary data encompasses guidelines on interpreting or analyzing the relevant traces and conducting specific investigative actions. In the final stage, we proceed to construct and visually represent the CM.

Additionally, it is important to stress the necessity of the incorporation of a feedback loop at this point, so that the CM will be updated accordingly after being confronted with new documents or cases. By doing so, we ensure the map's relevance and actuality and aim to avoid the experience trap [205].

Besides that, we used methods (1) and (3), i.e., conducting interviews with domain experts and consulting publically available data on real cases, to validate our map, as we will show below.

5.4.2.2 Persisting and Handling Cognitive Maps

CMs follow a node-link representation to store information and encode the mental model. Inspired by ISO-standardized topic maps [189], we propose to place the CM upon a knowledge management system providing linkability—both between the nodes representing the concepts themselves, i.e., so-called *associations*, and additional material, i.e., *occurrences*. This enables us to enrich each node of the map with documents or multimedia resources, such as phenomenological background information, documentation, implementation plans, legal requirements, references to file numbers of past cases, etc. By doing so, investigators can immediately acquire additional information regarding necessary preconditions and the stealthiness or implementational difficulty of the investigative measure at hand.

[36] Based on other research, Holt [118, p. 522] concluded that much information is shared among criminals in forums.

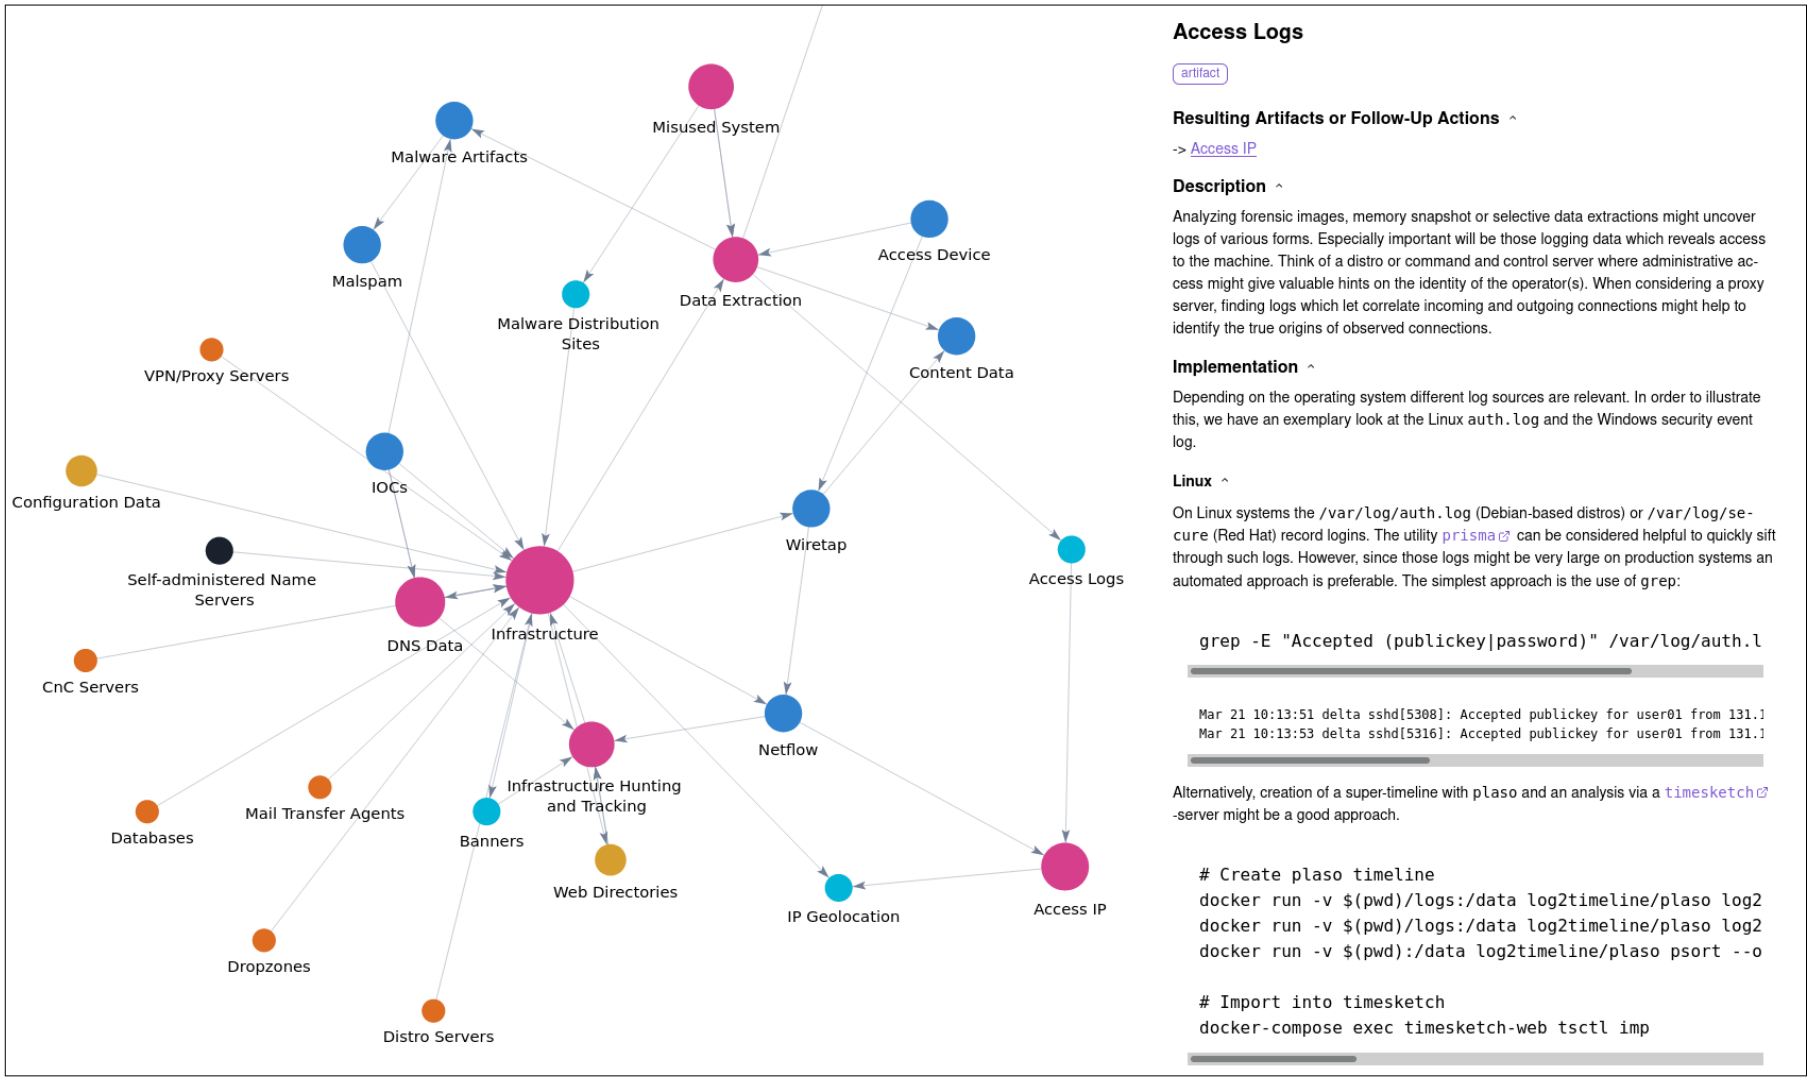

Figure 5.3: Excerpt of the browser-based visualization showing the deposited content of a specific node interactively.

Tool-wise, there are various options conceivable to construct such a flexible knowledge (data)base. In addition, having a machine-readable representation available opens up several possibilities for exploration and utilization of the phenomenon-specific data in various forms: One might think of programmatic and interactive 2D (see Fig. 5.3) or 3D representation as well as tailormade derivations for posters, coursebooks, or one's spatial CM as shown in Fig. 5.4.

5.4.2.3 Exemplary Illustration of a Knowledge Base

Building our prototype, we opted for a plain text solution and a light markup language to store the data underlying the CM. Aiming for interoperability, future-proofing, and ease of use, we went with the famous and established Emacs package `Org-mode`, which was complemented by its extension `Org-roam`[37]—both published as Free Software under the GNU General Public License version 3. `Org` is a note-taking and authoring system based on its own markup language, generally comparable to Markdown in its use but housing unique support for literate programming [202, 69]. In addition to that, `Org-roam` adds a thin wrapper around it providing backlinks by caching references in an SQLite database for networked thought. With this solution, we collected each concept—be it an investigative measure or some facet—as a separate plain text file, as shown in Listing 5.1. As illustrated there, each node is addressable by its ID (line 2). Content-wise, it houses a description of the concept (line 11) and links to resulting facets or follow-up actions (line 8). Furthermore, it can keep details on how to execute an action, acquire or process a facet (line 14) in executable code blocks, and reference additional information (line 21).

Storing the phenomenon-specific knowledge in this plain text node-link format has several advantages: First, we can easily export the content to various other formats, such as PDF. Second, it enables us to quickly navigate and sift through the knowledge base by relying on (fuzzy) searches. Moreover, we can document and execute command line-driven investigations in the style of literate programming. Lastly, we could resort to various visualization options. For example, we can create an interactive browser-based visualization of the graph using a patched version of the graphical frontend `org-roam-ui`[38], as it is illustrated in Fig. 5.3, or produce a printable overview of the map utilizing the widely acknowledged tool `graphviz` [75].

[37] See `https://github.com/org-roam/org-roam`, commit `c51cadf`.
[38] See `https://github.com/org-roam/org-roam-ui`, commit `df1f952`.

```
:PROPERTIES:
:ID: 00112233-4455-6677-8899-aabbccddeeff
:END:
#+title: Some Facet
#+filetags: :facet:

* Resulting Artifacts or Follow-Up Actions
- [[id:313657b9-5638-4c1b-887c-911e9c2b6935][Action A]]
- [[id:0d220dd0-0410-40eb-b51f-1b0832841807][Action B]]

* Description
<snip>

* Implementation
<snip>
#+begin_src shell
# Command documentation in the style of literate programming
<snip>
#+end_src

* References
- [[https://...][External Reference]]
```

Listing 5.1: Plain text data format representing the interlinked concepts in the exemplary `Org`-based knowledge repository.

5.5 A Phenomenon-specific Knowledge Base of Botnet Crime

To illustrate the previously described conceptual approach, we present the construction of a phenomenon-specific knowledge base for the technical aspects of botnet crime and largely related information stealer malware in form of a CM in Fig. 5.4. This phenomenon was chosen because the capabilities of botnets provide cybercriminals with a powerful and versatile toolset for fueling the underground economy and conducting a wide range of illicit activities at scale, often with anonymity and financial motivation. It is still a significant driver and enabler of many other cyber-related phenomena [4, pp. 288 f.], like identity theft, banking trojans, and ransomware; hence, botnets pose a significant challenge to law enforcement.

We employ a broad perspective of technical investigations and include crime scene work at primary as well as secondary scenes and consider other traces located at third parties, such as service providers. To focus on technicalities, we separated self-contained yet complex steps like financial investigations or undercover operations into their own maps. We keep references on those, which are denoted by the ◊-symbol. It must also be noted that not all conceivable but only the most relevant connections between the entities have been extracted from the underlying knowledge repository and drawn for reasons of clarity of the printed illustration. However, after having compared the map with the statements of domain experts as described below, we are confident that it presents the overall picture of botnet investigations well.

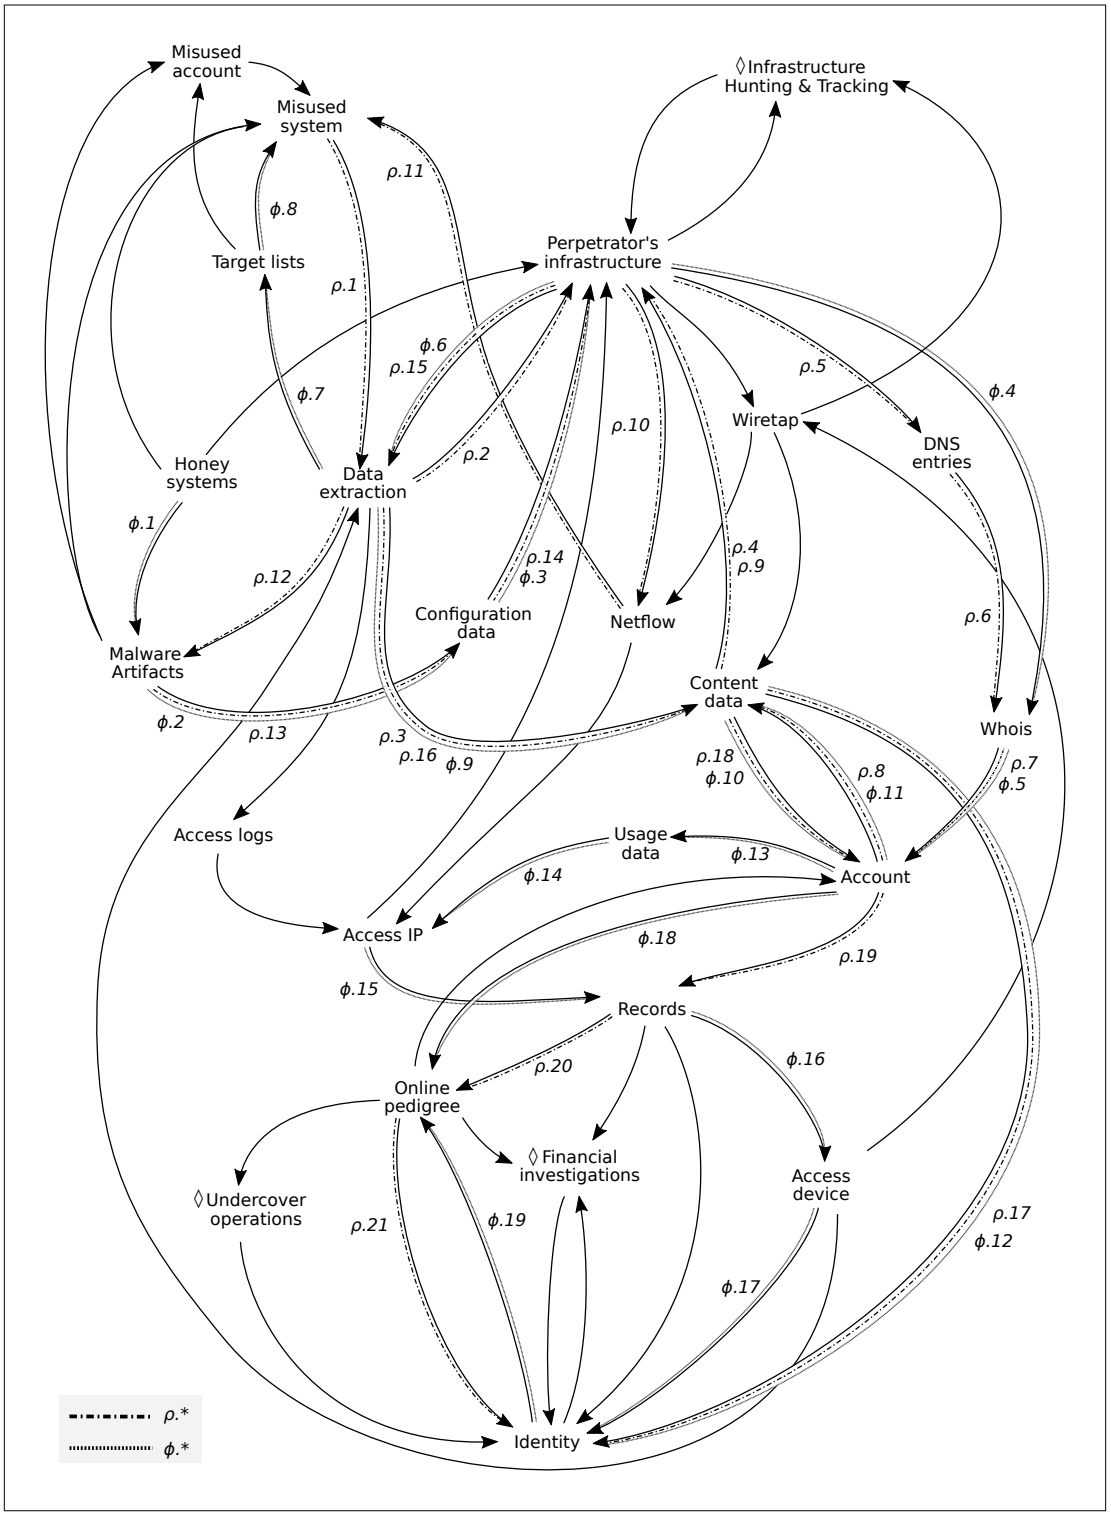

Figure 5.4: Cognitive map for the technical investigations in the field of botnet crime and information stealer malware. The edge labels $\rho.*$ and $\phi.*$ correspond to the real-world cases discussed in Section 5.5.3.2; The symbol ◊ denotes references on self-contained, but outsourced (sub)maps.

5.5.1 Approach of Building the Cognitive Map

Since the solution of a cybercrime offense in the narrower sense exhibits a strong parallel to the discipline of geopolitically motivated cyber attribution performed by (cyber threat) intelligence services, we consulted the academic publications in this field to compare procedures and trace materials with the crime scene of botnet crime. Primarily, we relied on the work of Wheeler and Larsen [241] and Steffens [212]. Already existing (yet unstructured and not consolidated) domain knowledge and the attribution methods described by the previously mentioned authors were supplemented by the technical analysis of botnets by Tiirmaa-Klaar et al. [224] as well as the practical textbook for cybercrime investigations by Bandler and Merzon [11]. Additional logical reasoning, based on Hunton's Cybercrime Execution Stack, and some visualization efforts led to the CM of botnet crime presented in Fig. 5.4.

5.5.2 Usage of the Map And Target Audience

If investigators are confronted with a cybercrime offense in the context of botnet or information stealer malware, they may refer to the investigative knowledge base. Subsequently, they can examine the observable traces at the specific crime scene and then identify the relevant elements within the knowledge base presented as a CM. This process allows them to deduce the available investigative options and their expected outcomes.

Each node on the CM serves as a point of reference, much like a physical map, enabling the derivation of clear directions when combined with the available case-specific findings. Consequently, it becomes clear which facets are relevant, i.e., constitute potential pieces of evidence, and how they fit into the broader investigative context.

To illustrate this quality, let us consider that the investigators identified a perpetrator's account at any service provider. They might look at the node "account" in Fig. 5.4 and infer their investigative options by looking at the outgoing edges. By doing so, they derive the options to request stock data to collect the "records", to acquire "content data" by a search warrant for data, or to acquire "usage data", which could be seen as a middle course. By following each outgoing edge of the subordinate nodes, they could find her way through the graph to ultimately solve the criminalistic task. So, when utilizing the knowledge base, it becomes clear to them what facets are relevant and expressive, which will be collected and analyzed next. Therefore, it eases the transferral of the distributed virtual crime scene into a manageable representation that supports a decision-making process because it captures a notion of relevance and expressiveness of traces —a crucial aspect of every criminal investigation.

Since various stakeholders with different responsibilities are involved in such investigations, it is only natural that the (visual) representation of the knowledge needs to provide differing levels of detail for different recipients. For the purpose of the chapter, we opted to present a map of rather coarse granularity, which might be used by decision-makers on a tactical level, like the command staff of an investigatory commission, who are responsible

for controlling the investigations but are not concerned with the actual execution and examination. However, given the availability of the underlying knowledge repository, it is rather easy to construct different views with deviating detail.

5.5.3 Validation of the Cognitive Map

In order to document the validity of the reasoning behind the phenomenon-specific mapping of the criminal phenomena of botnet crime presented in Fig. 5.4, we evaluate it in two ways: First, we draw on interviews with domain experts and then trace two real cases on the map.

5.5.3.1 Interviews with Domain Experts

After modeling the phenomenon-specific knowledge base that is presented as a CM based on literature reviews, we first aimed for validation of the results that were retrieved by the precursing document analyses. The overall process for this first part of the evaluation is illustrated in Fig. 5.5.

Study Design. In the field of knowledge engineering, three methods of knowledge collection are commonly distinguished [117]: Document analysis, observation studies, and interviewing experts. The first approach, i.e., analyzing documents, has been employed to create the knowledge base. The second approach, i.e., conducting observation studies, is extremely challenging to implement in the area of criminal prosecution for various reasons. The long duration of the investigative proceedings and the sensitivity, conspirativeness, and the classified material hinder the application of this method in that context. Therefore, we decided to conduct expert interviews to test our hypothesis that the presented cognitive map captures the relevant parts of the real-world investigative process sufficiently well. In previous studies, it has been demonstrated that expert interviews allow to acquire both the experts' technical and process knowledge, which is spontaneously available to them [20].

Sample Description. In view of the limited number of experts in the field, we resorted to a qualitative method and chose experts based on "their unique characteristics or their experience" [53], i.e., substantial involvement in a current or past large scale investigation in the field of botnet crime. So experts from various renowned institutions were contacted who have many years of experience in handling extensive investigative proceedings in the field of cybercrime and, in particular, botnet crime. In order to take different views into account, we decided for an interdisciplenary sample of four domain experts; choosing two specialized cybercrime prosecutors leading the proceedings and two technical experts in the field of digital forensics having more than 5 and up to 13 years of experience with investigative work at several renowned national institutions fighting cybercrime. For a thourough description of the sample, refer to Table 5.1.

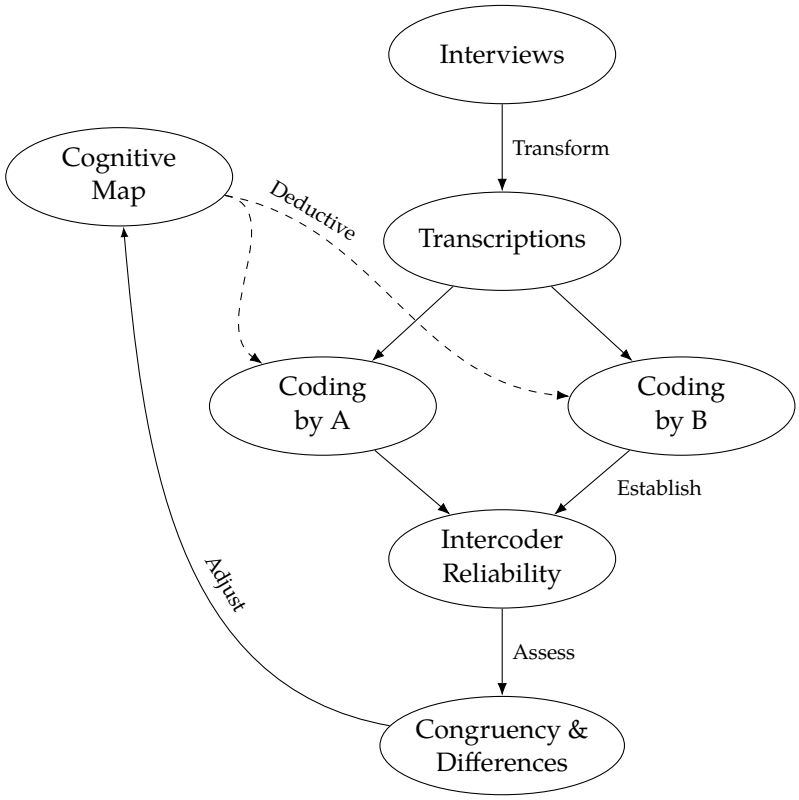

Figure 5.5: Illustration of the validation process of the CM. At first, interviews with four renowned domain experts have been conducted, which have been transcribed afterward. Then, the author of this thesis and a graduate student coded the interviews independently in a deductive fashion by ticking off the concepts present on the map. In the next step, the results were compared, and in case of deviation, the conflicting view was resolved by a discussion to achieve intercoder reliability.

Three of them were involved in remarkable botnet takedowns that have attracted worldwide attention. More details cannot be presented here to warrant the anonymity of the interviewees. Given the low number of successful botnet investigations, it must be noted that there are only a few experts in this specialized domain, which are, therefore, difficult to find and recruit.

Interview Process. We prepared an interview guide containing open questions to give direction, as shown in excerpts in Listing 5.2, and establish comparability between the interviews when retrieving the interviewees' experience regarding technical aspects in botnet investigations. The interviews themselves were conducted in a semi-structured manner, which allowed the direction of the interview to be adapted to the thought processes of the interviewee. We focused on the information that the expert considered important, allowing us to remain open and to avoid preconceived assumptions [13]. In addition, we facilitated elements of event-recall-interviews to ease the retrieval of relevant experiential information [117].[39] During the process of creating the guide, we talked to other criminal

[39] This claim is supported by the statement of Hoffman et al.: "Experts often have clear memories of tough or salient cases they have encountered [...]. Related to the utility of using test cases in task analysis, it can be

Table 5.1
Sample description of the interview study. Four renowned domain experts, all working at highly specialized organizations entrusted with botnet investigations, were interviewed.

Expert	Profession	Formal training	Experience in years	Size of org.
A	Prosecutor	Law	> 5	> 50
B	Tech. Analyst	Computer Science	> 13	> 90
C	Tech. Analyst	Computer Science	> 12	> 90
D	Prosecutor	Law	> 8	> 50

investigators and tried to gear the questions toward their working methods and thinking models. By conducting two interviews with test subjects who prepared themselves by reading real-world case files beforehand, we refined the guide to ensure acquiring in-depth information from the respondents.

In the actual interview, we asked the domain experts first to describe a memorable and relevant botnet crime case in detail in which they have been involved to set the scene and facilitate recall of the experiential knowledge. Then, we posed multiple additional questions to infer further investigative measures, the resulting artifacts and links between them. For example, this included questions regarding alternative approaches or sources of evidence as well as questions concerned with the underlying principles. By doing so, we were able to gather a more complete picture of their view on the field.

The duration of the interviews ranged from 40 minutes to over an hour. The interviews were audio recorded, after a consent form was signed. Furthermore, the interview protocol was approved by the data protection office of our university, where the data was stored on access-restricted servers in encrypted form.[40]

```
1  <snip>
2
3  *** Question Block 2: Approach to the investigation
4
5  1. How was the process initiated and what was the initial situation?
6      a. What evidence and leads were initially available?
7      b. What clarifications did you carry out first?
8
9  2. What investigative measures did you then take?
10     a. Why did you proceed in this way?
11     b. Which traces or findings were decisive for your further action and conclusions?
12     ...
13 <snip>
```

Listing 5.2: Exemplary excerpt of the interview guide for illustration purposes. We started to collect chronological information and posed additional questions to infer decisive connections.

useful to structure an interview by having the expert recall past events or cases, which can then be the focus of probe questions intended to facilitate recall [...]" [117, p. 136].

[40] Please note that the ethics commission at our university does not handle non-medical studies.

Analysis and Results. The transcriptions of the interviews were independently analyzed by two researchers, i.e., a co-author of the research article [103] on which this chapter is based on and the author of this thesis, employing a deductive approach. In this context, deductive coding means the respondents' statements were mapped to the concepts within our knowledge base if matched. Such a deductive approach is typically applicable "when the structure of analysis is operationalized on the basis of previous knowledge and the purpose of the study is theory testing" [76]. We coded the interviews separately to extract the concepts mentioned by the domain experts [200]. Then, a reconciliation of the results was conducted by analyzing each code and comparing it.[41] The few disagreements on the codes were discussed and could be resolved. We found that 22 out of 24 concepts in the map were mentioned at least once, as illustrated in Fig. 5.6. *Honey systems* and *WHOIS-data* were the only concepts that were not mentioned. At times, it was necessary to classify specializations of superordinate topics, e.g., when the interviewee mentioned "logs of logins to online accounts" as a special form of "usage data".

Furthermore, we refined our map by adding two missing links involving follow-up actions connected to "access devices" that are already included in Fig. 5.4. Given this result, we contend that the knowledge base and the derived CM captures the essence of botnet crime investigations. Considering the large degree of congruency with the experts' explanations, we further conclude that the map's intersubjective validity and coverage has been sufficiently substantiated to be of help for actual investigations.

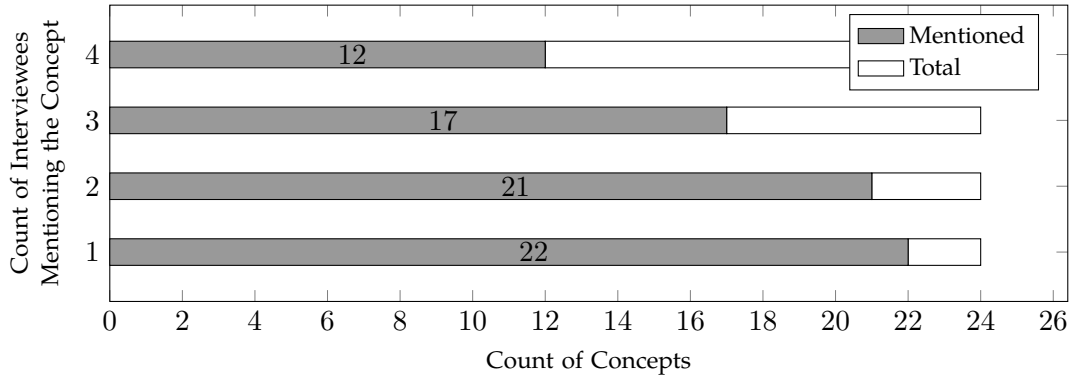

Figure 5.6: Congruency of the presented CM of botnet crime. $\simeq 91.7\%$ of the concepts in the CM presented as Fig. 5.4 were mentioned by at least one and $\simeq 87.5\%$ by at least two interviewees. 100% corresponds to all 24 concepts listed in Fig. 5.4.

5.5.3.2 Tracing Real-World Cases on the CM

With the intention of illustrating how the map can help to find relevant facets and navigate actual cases, its content will now be contextualised by considering two exemplary real-world investigations, whose relevant leads could be inferred from the map. This underlines the usefulness and plausibility of such an approach by uncovering valuable connections between investigative actions and findings.

[41] MaxQDA, a specialized software for performing qualitative data analysis (CAQDAS) [237], was used to facilitate the coding and categorization process.

The Case of Carding Kingpin R. Seleznev. First, we follow the identification of the already mentioned infamous carder Roman Seleznev on our map. Seleznev has to be considered one of the leading cybercriminals [153, ch. 1], member of the "Carder.su organization" and co-founder of CarderPlanet as well as one of the largest traffickers of credit card data between 2005 and 2014. He was found responsible for over 400 hacks of point of sale (POS) terminals, which he infected with his information stealer malware. In order to conduct those POS hacks, he scanned the internet for open Remote Desktop Protocol (RDP) ports serving as off-site admistrative access and exploited those remote desktop services by password brute-force or dictionary attacks [48]. Then, he installed his malicious software, which stealthily collected credit card data to send it back to collection servers, from where he put the credit card data in various flavors, i.e., "CVVs", "dumps", and "fullz",[42] on sale on so-called "dump shops" and underground forums [170, p. 63]. The investigators unveiled multiple chains of evidence to attribute ultimately 169 million US-dollar in credit card fraud to Roman Seleznev. Those steps are roughly outlined in the following sections, while references are drawn to the cognitive map in Fig. 5.4 by using the edge labels '$\rho.*$'.

The investigation against Roman Seleznev was kicked off with a complaint by the restaurant *Schlotzsky's Deli* in Idaho, reporting an issue with its point-of-sale system. Detective David Dunn, member of the Secret Service Electronic Crime Task Force, conducted the evidence collection [170, p. 63], which meant acquiring an image of the affected POS system as well as its main memory ($\rho.1$). Thereby, he could identify that it was beaconing to a Russian IP address ($\rho.2$). As it was outside of US jurisdiction, this finding was not helpful at all, unfortunately. However, a little while later, another lead came up from an independent case against a fraudster in Ohio who used credit card data stolen from the POS system of *Schlotzky's Deli*. The analysis of the forensic image of the fraudster's computer revealed ($\rho.3$) that he bought those credit card data from "dump shops" hosted at the domains named *track2.name* and *bulba.cc*. These domains were operated by the same individual as several chat messages found in the image suggested ($\rho.4$). Following this lead, detective Dunn then queried the domain registries ($\rho.5, \rho.6$) and linked the email-addresses *rubensamvelich@-yahoo.com* and *bulbacc@yahoo.com*, which were used to register the domains ($\rho.7$), to the suspect of the "*Schlotzky's Deli* intrusion". He provided an affidavit to the court with jurisdiction to obtain a search warrant for data aiming to seize those accounts. After receiving the stored communications, he uncovered ($\rho.8$) that the user of one of the email accounts rented a server operated by the host service provider *HopOne* located in Virginia ($\rho.9$). The capture of netflow data ($\rho.10$) obtained by a so-called pen/trap-order showed that this machine was used as a collection server to collect credit card data of various US-based restaurants running similar POS systems and therefore led to the identification of numerous other victims ($\rho.11$), which built up the case and justified an arrest. The extraction of the malware from the initial victim's system ($\rho.1, \rho.12$) and its static analysis revealed configuration data ($\rho.13$), which also proved that this server was used to exfiltrate credit card data ($\rho.14$). Afterward, a seizure of this server located on American soil was conducted ($\rho.15$) and the subsequent analysis of the acquired image brought up important attribution evidence from the contained content data ($\rho.16$), as Seleznev performed private internet browsing on the machine, where he made travel reservations providing his passport number among other things ($\rho.17$). However, this

[42] Those terms are fraudster slang for different variations of credit card information.

failure in Seleznev's operational security was not the only one leading to his identification. In addition, the user of the email-address *rubensamvelich@yahoo.com*, which was linked to the dump shop's domain *track2.name*, used this mailbox (already discussed as $\rho.7, \rho.8$) to open up a PayPal account ($\rho.18$). An additional records request was sent out ($\rho.19$), which brought up Seleznev's identity so that all accounts added up to a rich online pedigree ($\rho.20$), which could be pinned to him by multiple paths and provided a solid attribution ($\rho.21$) [48].

The Case of the Alleged *Jabber Zeus Crew* Member M. Yakubets. As a second case, we trace the identification of the elite cybercriminal Maksim Yakubets on our map, who is nowadays wanted by the FBI for orchestrating the operations of the cybercrime group *Evil Corp* [229].[43] However, here, we look at the activity related to the malware *Jabber Zeus* around 2009, when Yakubets was involved in the operation of this financial botnet by being concerned with financial theft and money laundering [228, p. 2].

The investigation was initiated after the FBI "received numerous complaints of fraudulent ACH [Automated Clearing House] transfers" in May 2009, where the perpetrators gained one-time PINs in an unauthorized way to initiate the transfers. Those defrauds were linked to the malware family *Zeus*, which was sold on various underground forums at that time [228, p. 10], but was modified to be able to send one-time passwords by utilizing the instant messaging protocol Extensible Messaging and Presence Protocol (XMPP). This circumstance was discovered only a few weeks later by researchers of a private sector threat intelligence company who gathered a modified *Zeus* sample ($\phi.1$). In the course of their analysis ($\phi.2$), they uncovered an XMPP server at the domain *incomeet.com*, which acted as communication endpoint of the discovered malware receiving those one-time PINs ($\phi.3$). This server was operated by a host service provider in New York ($\phi.4$) and rented by an inexisting person or a stoog ($\phi.5$) [228, p. 12], so that this lead turned out to be a dead end. Then, FBI agents conducted four searches of said server ($\phi.6$), on which extensive logs were stored, because the operators had configured it to record "ongoing logs of every chat message sent through the server". As a result, the investigators found XMPP messages containing banking credentials submitted by bots, which enabled the identification of further victims ($\phi.7, \phi.8$). In addition to that, the analysis of those chat messages reveiled Russian communications among the conspirators concerning the operation of the botnet and the recruitment of money mules ($\phi.9$) [228, pp. 13 ff.]. The further evaluation of chat logs showed that one of the conspirators with the moniker *aqua* used the email address *aquamo@mail.ru* ($\phi.10$) [228, pp. 19 f.]. Given this Russian mail service provider, a mutual legal assistance treaty was transmitted to Russian authorities to identify the user of the said mailbox. They gathered the mailbox' content ($\phi.11$) and found that there were messages addressed to *aqua* concerning cybercriminal activities as well as to a person named Maksim Yakubets, who had made travel arrangements with this email address and provided a telephone number for the delivery of a baby carriage, which was linked to the aforementioned individual ($\phi.12$). Furthermore, Russian investigators obtained usage data ($\phi.13$) containing login IP addresses ($\phi.14$), to which records could be obtained from the internet service provider ($\phi.15$). Those pointed to a specific address in Moscow, where Yakubets could be located while conducting a search warrant ($\phi.16, \phi.17$) [228, pp. 26 f.].

[43] Note that Maksim Yakubets is assumed to be innocent until he has been proven guilty.

The FBI agents secured the attribution by requesting records from a Skype account linked to *aquamo@mail.ru* named *maksim.ya* (ϕ.18), which unveiled that he was connected with a female, considered to be his spouse at that time. This finding enabled a backlink from the real-world identity to his online pedigree (ϕ.19). By looking at the date of birth of the common son, the agents found a correlation within chat logs from that day on the XMPP server at *incomeet.com*, where the user *aqua* informed a co-conspirator, that "[s]he gave birth" (supporting ϕ.12). Given these chains of evidence, the investigators concluded with high confidence that the moniker *aqua* was used by Maksim Yakubets [228, pp. 29 f.], who is suspected to be engaged in cybercrime on an almost unimaginable scale for over a decade [14].

Verdict of Tracing Real-World Cases. While today's cybercrime actors might not be as careless about their operational security as Seleznev and Yakubets were a few years back, it is especially important to consider those investigative actions as expedient, which are aiming to break the *evasion and concealment* measures (like The Onion Router (Tor), virtual private networks (VPNs), virtual currency mixing, and so on) and exploit their operational security failures, in order to be able to conduct the identification of these perpetrators. To accomplish this, each and every inadequacy or negligence in the anonymization measures taken by the perpetrator has to be recorded and exploited. By bundling the measures, a large spread of investigative pressure must be achieved in order to increase the probability of such a finding—the human factor in cybercrime. This is the only way to assume that the many obstacles, which result from various anonymization techniques, technological advancements, or the crime-as-a-service paradigm, can be circumvented, aiming to achieve the desired disguise of the perpetrator. The utilization of a phenomenon-specific knowledge base providing information on expressive and relevant traces, e.g., in the form of cognitive maps of the digital crime scene, can help to select those measures and conceive a thorough investigation planning.

5.6 Discussion

After having presented the nature of phenomenon-specific knowledge bases, their construction, and their design to get an understanding of relevance and expressiveness for specific criminal phenomena, we now place our proposal in the overall picture and take a critical look at it. First, we start with the limitations of the phenomenon-specific knowledge base at the example of botnet crime that was visualized as CM. Second, we take a step back, relate it to other approaches, and point out existing shortcomings.

5.6.1 Limitations of the Exemplary Cognitive Map

In Section 5.5.3, we demonstrated both the validity and the usefulness of the phenomenon-specific knowledge base of botnet crime in the form of a CM. However, it must be noted that the presented CM is limited to the technical parts of the investigations and provides the

information in a specificly chosen degree of abstraction. Deliberately, we omitted follow-up investigative measures after identifying a suspect to keep the content focused and concise. Also, more sensitive investigative measures commonly conducted in a highly conspirative manner were omitted for obvious reasons. Additionally, the complexity of the real world might necessitate the ability to dive into a fully meshed interactive representation for working at the operational level, as it has been created in the course of this research, which is exemplified in Fig. 5.3.

Regarding the validation of this exemplary study, it has to be noted that the interviewees whose statements were used for evaluation were exclusively German; hence, there is a need to validate the results with an international group despite the high amount of international cooperation in botnet investigations because the practical experience of German domain experts is also subject to the German legal system. Furthermore, we need to note that the analyzing researchers play an interpretative role in qualitative research; therefore, it is possible that our preexisting understanding of the phenomenon may have influenced the outcomes of the deductive analysis. However, to minimize this concern, we took steps to address it by establishing an intercoder agreement between two researchers during the analysis.

5.6.2 General Limitations and Open Questions

Critically reviewing the proposed approach, we identify three fields with the potential for improvement or additional insights. Those are related to the knowledge representation, knowledge engineering, and the usefulness of phenomenon-specific knowledge.

Representation. Thinking about the chosen representation of the phenomenon-specific knowledge base, we see significant concerns arise: First, the reader might wonder why we did not structure the phenomenon-specific knowledge using a full-fledged ontology approach. For example, one could think of the *PREVISION* ontology for crime investigations by Müller et al. [169], which is based on the *Unified Cyber Ontology* by Syed et al. [218] and uses the intelligence pentagram, as described by Dragos [71]. An alternative choice could have been the *CASE* ontology proposed by Casey et al. [41], which is also extending the *Unified Cyber Ontology*. Indeed, building on an ontology could provide a semantic framework, automated reasoning, and enhance both the overview and situational awareness of the case under investigation. However, in the present work, we focus on the main connections of traces, investigative measures, and their results to gain a foundational understanding of relevant and expressive evidence at a meso-level of abstraction targeted at human reasoning—an application where ontologies introduce great complexity and hinder human understanding in our perception. For exploring automated reasoning in the future, such approaches need to be considered and the knowledge base needs to be structured in ways the ontologies propose.

Second, the question arises why the phenomenon-specific knowledge base, as realized in this chapter, deviates from the formal investigative knowledge base as defined in Def-

inition 3.3.1. While we could have aimed for a projection of the phenomenon-specific knowledge by specifically posing hypotheses and mapping the identified traces to these, we refrained from doing so, since we did not see that this would add value for the investigator. On the contrary, we would lose the relations to investigative measures as an expression of law enforcement action. Nevertheless, we cannot identify noteworthy obstacles that hinder one from projecting the presented phenomenon-specific knowledge base in a formal representation in the sense of Definition 3.3.1.

Engineering. For the example of botnet crime, we relied on the analysis of documents of various kinds for extracting phenomenon-specific knowledge. Such an approach only applies to phenomena that are (or have been) in focus by researchers who publish their findings. For many other phenomena, such corpora of literature only exist incompletely or even not at all. In such cases, the phenomenon-specific knowledge is only fragmentarily available in the minds of experienced investigators. Hence, Voigt [238] developed in her Master's thesis, supervised by the author of this dissertation, an inductive expert interview approach. She proposed to derive all relevant concepts and infer their relation to one another from semi-structured interviews [238, pp. 51 f.], much like we did for the validation. This inductive domain expert interview approach is embedded into a five-step-procedure for creating phenomenon-specific knowledge representations [238, p. 54], contributing to a generalized method for collecting such investigative knowledge.

Usefulness. On the one hand, we are confident that the availability of an investigative knowledge base providing a catalog of relevant and expressive traces for a specific scope of the offense or phenomenon is a precious aid for investigators. On the other hand, which representation will be most effective while easy to adapt has yet to be empirically assessed. Although using node-link representations in general and CMs in specific for knowledge representation and decision-making is well founded, their efficacy and helpfulness in real-world investigations has yet to be assessed—potentially considering alternative possibilities.

5.7 Summary

In the previous chapters, we employed a theoretically driven view of the matter of evidence. However, after showing how we can determine relevance of traces using an automata-theoretic approach, it remained questionable, if and how the notion of relevance can be transferred into a more practical setting because the presented approach has limitations regarding the construction and processing of system models. Hence, we explore the use of phenomenon-specific knowledge to infer what constitutes relevant traces in real-world investigations.

We motivate the need by concluding that the combat against cybercrime is too inefficient for multiple reasons. Two outstanding deficiencies targeted in the present chapter are an abstraction gap, on the one hand, and a lack of attention to operational aspects, on the

other. By critically reviewing existing models for the investigative process and testing their applicability in real-world investigations of cybercrime offenses, we identified this gap in the level of abstraction (Section 5.2). While being universally applicable, existing models do not provide guidance on the concretization of the model in real-world investigations. The existing process models for cybercrime investigations commonly act on a macro level of abstraction, while the real-world case work is the opposite pole and is characterized by being very specific. This justifies the necessity of introducing an intermediary step. In order to bridge this gap and provide actual guidance for conducting investigations practically, we argue for the integration of an investigative knowledge base, which provides an understanding of the relevance of traces for specific investigative questions, into more general models (Section 5.3). Phenomenon-specific knowledge can be used to derive this information. We grasp this term as the experiential knowledge related to a specific criminal phenomenon, which is defined by the perpetrators' goals and the employed TTPs as a more detailed conception of the modus operandi. We propose incorporating node-link knowledge representations visualized as cognitive maps of relevant and expressive evidence related to investigative measures. Such a representation acts on the meso-generic level of abstraction and encodes a qualitative "mental model" of what we formalized in Chapter 3. We look into ways of constructing, persisting, handling, and storing such a phenomenon-specific knowledge base and present our choices in that matter, so that it can provide guidance for the concretization of abstract process models and effectively aid the investigator to better understand the investigative options and expressive traces available (Section 5.4).

To breathe life into the proposed approach and illustrate it vividly, we present an exemplary map for the phenomenon of botnet crime and information stealer malware. For its construction, we conducted a literature analysis, in which we first carried out an open coding phase, inferred connections, enriched the identified concepts, and finally visualized it both as a cognitive map as well as a browser-based graph representation. By interviewing proven experts in the field, we measured a high congruency of about 87.5% and, hence, can validate its intersubjective correctness and completeness. In addition, we traced two real-world cases on the created map to show how the use of phenomenon-specific knowledge bases can support the identification of relevant traces and navigate actual investigations aiming to demonstrate the usefulness and plausibility of phenomenon-specific knowledge encoding relevant evidence (Section 5.5).

Lastly, we discuss the limitations of both the general approach, as it is presented in this chapter, and the exemplary CM for illustration purposes. The exemplary map focuses on technical investigations and stops at the identification of the suspect. Thus, it does not represent the investigations regarding the phenomenon in question in its entirety. Furthermore, additional validation is possible, which should incorporate an international sample of domain experts. From a general perspective, we explained why we deviate from the formal investigative knowledge base, on the one hand, and do not employ an existing ontology to structure the phenomenon-specific knowledge base, on the other hand. Mainly, this is done to present the information for human reasoning and decision-making processes rather than for automated reasoning, which could be tackled in future work. We further address the need for developing a generalized method of collecting investigative knowledge to infer relevant and expressive traces regarding criminal phenomena where

the experiential knowledge is not as well documented as in the case of botnet crime. In addition, we articulate the necessity to evaluate the usefulness and efficacy of (different kinds of) phenomenon-specific knowledge bases for supporting investigators in real-world cases (Section 5.6).

However, besides these open questions, we see potential that the proposed method can improve the decision-making process of cybercrime investigations by making it more structured and effective while it integrates well with more general models, like the *Criminalistic Cycle* [240] or the *Cybercrime Investigation Framework* [122]. Having mapped a holistic view of the phenomenon-specific knowledge for botnet crime and illustrated its helpfulness by explaining two real-world cases, the necessity becomes obvious to create CMs for other cybercrime offenses, to gather, analyze, and document relevant traces and related investigative knowledge for those categories of offenses as well.

In essence, the present chapter transfers the solely theoretic considerations of relevance and expressiveness into practice. It proposes a method of storing phenomenon-specific knowledge, which encodes relevant traces, as cognitive maps to bridge the abstraction gap of digital investigations. Such improvements are vital to understanding crucial aspects of digital investigations to contain the growing domain of cybercriminal activity and lower its substantial economic and political impact.

6 Contamination of Digital Evidence

6.1 Introduction

"*Woman Without a Face*"—after the brutal murder of a police officer, German law enforcement chased an unknown female serial killer [110]. The "Phantom of Heilbronn" was suspected to have murdered five other victims between a prolonged period of time between 1993 and 2009. The hypothesis on the perpetrator was formed based on DNA evidence that has been found on all murder scenes. Back then, the investigations were pushed forward under high pressure by "SOKO Parkplatz", a task force at Baden-Württemberg state police, to find this high-profile killer. However, the chase came to a dramatic yet surprising end when it turned out that the DNA belonged to a woman who had been working at the factory producing the cotton swabs used to gather the DNA traces [67, 116]: The phantom was the result of evidence contamination, which hindered the expressiveness of the evidence to such an extent that it was effectively not only useless but severely misleading.

This story is just one of several notable examples in which physical evidence had been inadvertently altered by adding DNA from other sources. All these examples illustrate the risks of handling evidence that is invisible to the human eye. For decades now, forensic technicians have been well aware of (cross-)contamination when conducting classical crime scene work [148, pp. 56, 259]. So, for physical evidence, where "[c]ontamination is a fact of life for investigators" [90, p. 143], there appears to be an increasing awareness of the adverse effects since contamination "can have a significant negative impact on the investigation if the existence of the contamination is not known" [203]. Therefore, forensic science developed strict regulations and processes on how to act and proceed at a crime scene, an example being so-called "Police Elimination Databases" like those employed in Austria [182], which have been set up to help identifying crime scene contaminations by investigators. Still, while it is not known with certainty in how many cases a DNA sample is contaminated by investigators or examiners [82, p. 122], studies from Austria suggest that approximately 1–2% of biological traces are contaminated [182].

Since physical evidence is often considered to be fundamentally different from digital evidence, the question arises whether contamination is also possible when handling digital evidence. This question is becoming increasingly relevant as investigative methods get closer to running or "live" systems. The challenges of encryption and ever-increasing amounts of data have led incident responders and forensic investigators to develop and explore methods like triage [39], live analysis [3], and selective imaging [79]. However, interacting with a live system will inevitably result in changes that provide an increased

potential for inadvertent modifications; in addition, it can be shown that this holds true both for work at the crime scene but also for post-acquisition lab environments.

Given this chance of modifications to the examination object or derived evidence, there is an urgent need to discuss the phenomenon of contamination at digital crime scenes and digital forensic labs more clearly. Previous works [60, 150, 154] already identified contamination as a problem by name, but did not go into the details of this specific topic. So overall, contamination of digital evidence remains an elusive concept. After having an understanding of using relevant evidence in real-world investigations, we scrutinize contamination as one effect that can hamper the expressiveness and hence diminish the value of (digital) evidence in the present chapter.[44]

6.1.1 Contributions of the Chapter

In this chapter, we revisit works on traditional contamination published in other subdisciplines of forensic science and identify crucial properties of contamination. Based on these preceding works, we develop a generalized definition of contamination—valid both for physical and digital evidence—as a first contribution in Section 6.3, namely the "inadvertent transfer of traits to an object of relevance at any point in the forensic process". Building upon this harmonized definition, we relate contamination to the established concept of general evidence dynamics and show that it can be effectively considered a subset of it as a second contribution. As a third contribution, we present and classify nine examples of contamination of digital evidence to point out the risk posed by it in real-world encounters, as well as four counterexamples to demarcate the phenomenon. Specifically, we classify the trait transfers in the discussed examples in three dimensions: They can be actively or passively induced by the responsible party, result in evidence addition or subtraction, and might happen directly or indirectly with intermediate systems or items involved. By addressing the specifics of digital evidence in this context, we argue that our definition can be helpful in understanding the risks arising from contamination in this domain (Section 6.5). As a fourth contribution, we identify specifics, more profound reasons for digital contamination, and its broader effects. We show that such unwanted changes to the system under investigation after securing the scene stem from a complex interplay between system configuration, autonomous processes, and the analyst's actions.

Overall, our exposition is intended to serve as a proposal to the community to develop a shared understanding of an important yet underexposed phenomenon.

6.1.2 Chapter Outline

The remainder of this chapter is structured as follows: At first, we approach the subject matter and provide a brief history of contamination research to give an overview of related

[44] Given the explanations in Section 2.4.3, one could argue that a contaminated facet can never constitute "evidence" since its reliability is violated; however, we use the term "evidence" in this chapter without the substantial criteria developed earlier to improve the general readability of the text.

work regarding traditional contamination in Section 6.2. Here, we also discuss definitions of contamination in the analogous domain. In order to harmonize those previous approaches, we identify common properties and then develop a novel definition that is also applicable for digital evidence and crime scenes in Section 6.3.

In Section 6.4, we demonstrate the practical emergence of contamination by presenting selected examples of contamination in the context of digital forensics to further explore the aspects of the previously mentioned properties and to pin down essential details for understanding this phenomenon in-depth, its occurrence, and its adverse effects on the expressiveness of evidence. Aiming to delineate contamination from general evidence dynamics, we provide counterexamples and show edge cases to highlight real-world complexity when facing a conflict of objectives during acquisitions and analyses.

In Section 6.5, we have an even closer look at the phenomenon by highlighting its specifics and intricacies in the digital domain as part of discussing our results. By doing so, we identify several remaining challenges and outline a research gap, before we conclude the chapter in Section 6.6 by summarizing the findings.

6.2 Related Work

Based on decades of experience with physical evidence, researchers in forensic science appear to be well aware of contamination effects and have already proposed guidelines and best practices to avoid them. We now briefly revisit some milestones on this path and then characterize common elements of definitions of contamination. As contamination seems to be most relevant (and probably most dangerous) when dealing with microscopic particles, it is not suprising that most of the literature concerns handling of DNA evidence.

6.2.1 Overview of Physical Contamination

Already back in 2005, van Oorschot et al. [233] described that fingerprint brushes, gloves, and other tools routinely used during examinations could contaminate evidential items. This can happen either via direct DNA transfer initiated by the forensic technician or via an indirect DNA transfer by items being previously examined [188]. In the past, such shortcomings in contamination avoidance during lab work also led to severe consequences and false accusations in actual proceedings [see 167, as an example].

Those observations were then taken up by Meakin and Jamieson [162], who presented a review of the risk of transfer of so-called "trace DNA"—a term that describes a small number of DNA particles "that cannot be attributed to a particular biological source" but contains enough information to recover a full DNA profile and identify an individual person [162, p. 434]. They described scenarios of direct and indirect transfer of trace DNA and presented several factors affecting deposition, persistence, and analytical recovery.

Given the increased sensitivity of DNA analytics, Margiotta et al. [155] conducted experiments regarding the contamination risk by gloves in forensic casework and quantified a high risk of DNA transfer, which underlined the need for awareness, the necessity of DNA-free gloves and instruments, appropriate cleaning systems, and multiple layers of gloves. Additionally, Fonneløp et al. [82] experimentally analyzed the risk of contamination caused by police investigators and secondary DNA transfer from evidence bags and underlined that there is a need to evaluate existing practices, identify weaknesses in evidence handling, then suggest and implement improvements in the lab to finally demonstrate the effectiveness of contamination monitoring [82, p. 122]. As the main result of their work, Fonneløp et al. presented 12 practical guidelines that reduce the risk of contaminating a DNA sample, such as the frequent change of gloves and the wearing of double gloves, the separation of suspect and victim exhibits, or establishing national elimination databases [82, p. 128].

Pickrahn et al. [183] quantified the chance of DNA contamination in Austrian cases and underlined the importance of reference profiles stored in such databases since they identified that police contamination is a real issue [183].

6.2.2 Efforts to Grasp Physical Contamination

Besides laboratory studies that have been motivated by practical needs, Inman and Rudin [125] discuss the issue from a conceptual point of view and provide the following working definition of contamination [125, p. 211]:

> "Any substance inadvertently introduced into or onto an item of evidence after its recognition by a responsible party."

Their definition focuses on the physical matter ("substances") added to an evidence item *after* the crime scene has been identified and secured.

A comparable but differing definition is put up by Grübler [107, translated by the author]:

> "Unintentional contamination of traces and reference material with a similar material that is irrelevant for the creation of the evidence [...]. If the contamination is not recognized, the trace causation leads to false-positive results [...]."[45]

Grübler [107] also refers to the physical matter (calling it "material") but further specifies that contamination cannot occur with any material but only with the same kind of material. Moreover, both emphasize the unintentionality of the act. Conversely, Grübler [107] does not consider temporal aspects and omits to limit the time frame in which contamination could occur. Another difference is that he includes—in contrast to Inman and Rudin

[45] Original wording: "unbeabsichtigte Verunreinigung von Spuren und Vergleichsmaterial mit gleichartigem, aber für die Spurenentstehung irrelevantem Material [...]. Wird die Verunreinigung nicht erkannt, führt die Spurenverursachung zu falsch-positiven Ergebnissen (Trugspuren)." [107, pp. 362 f.]

[125]—the actual effects of undetected contamination ("false-positive results") in his definition.

Gehl and Plecas [90, p. 113] are even more explicit in their criminalistics textbook:

> "Contamination is the unwanted alteration of evidence that could affect the integrity of the original exhibit or the crime scene. This unwanted alteration of evidence can wipe away original evidence transfer, dilute a sample, or deposit misleading new materials onto an exhibit."

In their definition, they focus not only on the ultimate outcome, as Grübler [107] did, but the more subtle effects on the evidence. In a broader understanding, they include evidence destruction ("wipe away") as well. In agreement with the previous definitions, they also focus on the physical matter (indicated by the terms "materials" and "sample"). Since they use the adjective "unwanted", they make clear that the effects are undesired. In this definition, the temporal aspect, which is emphasized by Inman and Rudin [125], is also missing. Although Gehl and Plecas use the terms "exhibit" and "crime scene", what necessitates that the items and locations have to be recognized as such beforehand, their comments clarify, however, that they do not draw temporal distinctions, which makes it difficult to distinguish these effects from other concepts like evidence dynamics [44] in general.

Lastly, we have a look at an official definition of a standards body: The standard ISO 21043-1:2018(en) 'Forensic Sciences Part 1 Terms and Definitions' by ISO/TC 272 [129] is more brief. According to this document, contamination is defined as the

> "undesirable introduction of a substance to an item at any point in the forensic process ".

While the other definitions refer to evidence, the ISO's definition refers to "item[s]", which are in turn defined as any "object, substance or material that is collected, derived or sampled as part of the forensic process" [129, 3.19]. Using this definition, they limit the scope of contamination on relevant items. Again, the standardization body focuses here on physical matter. They regarded a temporal confinement since the forensic process must have been started, and the result is characterized by undesirability—only implicitely stating missing intent.

6.2.3 Aspects of a Common Definition of Physical Contamination

There exist several other definitions focusing on biological evidence. While Schwendener et al. [203, p. 518] distinguish between contamination and pollution, other publications describe similar aspects as those used in the definitions above [see, for example, 182, 226]. So while the above definitions differ from each other in detail, we can observe four intersecting aspects:

(a) an alteration through introduction (or transfer) of substance,
(b) an item of evidence in the forensic process,
(c) temporal confinement to the forensic process, and
(d) unintentionality.

While the focus of all definitions was physical evidence, interestingly, all but the first aspect are directly applicable also to digital evidence. We revisit these four aspects in the following section in which we develop a generalized definition of contamination that covers both physical and digital evidence.

6.3 A Generalized Definition of Contamination

We now revisit the four aspects of contamination definitions identified in the previous section and investigate their applicability to digital evidence in order to arrive at a more general definition of contamination. We begin with the arguably most important aspect.

6.3.1 Alteration through Transfer

The first aspect concerns the alteration of evidence through the introduction of physical matter. While it appears unreasonable to apply this to digital evidence, it has been observed that the transfer of substance, i.e., physical matter, is not the core concept underlying evidence modification [126]. For example, toolmarks leave traces at the crime scene without exchanging physical substance. The underlying principle rather is the more general notion of *transfer of traits* [126, p. 15], meaning the transfer of patterns that change the interpretation of what is found at the crime scene. It is straightforward to observe that such transfers can also happen in the domain of digital evidence. In general, such transfers exist in various manifestations, which can be classified as additive or subtractive, actively or passively induced, and direct or indirect, as Fig. 6.1 illustrates.

The addition or removal of evidence is considered on a semantic level. By using these terms, we differentiate between introducing new information or removing existing one. For some readers, it may appear somewhat surprising that this concept includes the destruction of traces. Still, one can understand the removal of traces as another transfer with a new pattern that contains less information of relevance. As an example in the physical domain, one might think of a footwear mark on a dusty surface, where investigators fail to protect it from heavy rain by setting up an appropriate cover. In the digital domain, we can imagine similar effects, like overwriting ring buffers or chunks of data that have been marked free before, hence introducing new patterns.

Moreover, trait transfers can be passive or active, as indicated in the above examples. We consider it to be active if it is linked to an investigator's action and, thus, its immediate effect. Conversely, we believe a transfer to be passive if it is a consequence of the refrainment to

take a necessary measure to avoid or stop it. In both cases, the investigators are responsible for the contamination that occurs.

Finally, the transfer may not necessarily be induced directly. New traits could be brought in via an intermediary tool, system, or person. For example, at a physical crime scene, an intermediary contaminating tool might be a fingerprint brush, an evidence bag, or another item used during the examination, as described by Poy and van Oorschot [188]. In the digital domain, failure to sanitize media prior to a disk clone, could be analogous.

It is worth noting that, at least in theory, such modifications—including relocation, obscuration, obliteration, and removal of evidence—are entirely avoidable.

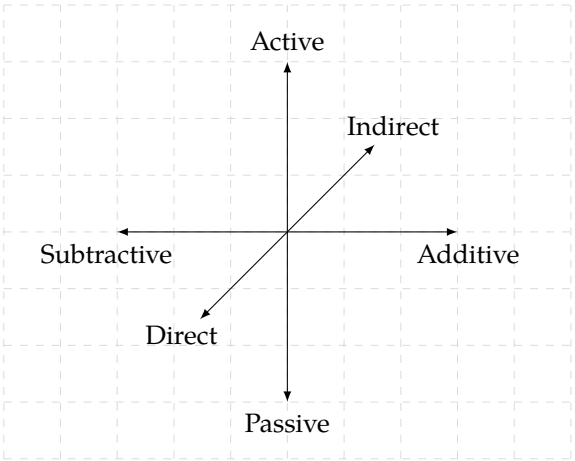

Figure 6.1: Three dimensions of transfer of traits for contamination. Contamination can stem from a measure taken by a responsible party (active) or the refrainment from taking some action (passive). By doing so, either new information is brought to the object of relevance (additive), or it could be removed from it (subtractive). Such a transfer could happen via another system involved (indirect) as well as without any intermediary (direct).

6.3.2 Object of Relevance

All previously mentioned definitions of physical contamination refer to items and further restrict those by speaking of evidence, traces, or—in the broadest sense—anything collected during the forensic process. By narrowing down this second aspect, we stress that it is not about the alteration of any object but an object that is considered ex ante of relevance for the investigation. To be contaminated, such an object must be subject to a transfer of traits, which results in a violation of evidence integrity,[46] thus altering the semantics of the evidence comprised of the physical or digital object at hand.

It is worth noting that contamination is defined regarding an object and not a specific location, such as the scene of the deed or a body location since there are several occasions for contamination to happen—not just at a crime scene but also in lab environments.

[46] Rare and rather theoretical edge cases, like an in situ replacement of an identical byte string, are intentionally not captured by our definition because it does not have any adverse effect.

6.3.3 Temporal Confinement

Besides the spatial aspect, the concept of contamination should be restricted to a confined period of time. It is paramount to underline that it can only occur after investigators of law enforcement or another empowered and responsible party have identified the object mentioned above as potentially relevant for the investigation, and initiated the forensic process by doing so. This commonly involves establishing necessary crime scene protocols and other precautions during seizure and lab work.

Drawing on the definitions of ISO/TC 272 [129] and Inman and Rudin [125], changes before the initiation of the forensic process are not considered contamination. The term *evidence dynamics*, a concept described initially by Chisum and Turvey that refers "to any influence that changes, relocates, obscures, or obliterates physical evidence, regardless of intent" [44, p. 7], captures such temporally preceeding changes [see also 45, p. 144].

Thus, contamination is directly related to the protective function of the police officers or any competent investigator, analyst, or examiner tasked with conducting the forensic process, which includes securing evidence from any alteration after locking the scene, and the subsequent analysis and examination activities.

6.3.4 Missing Intent

Lastly, the salient feature is missing intent regarding the transfer of traits. In crime scene and laboratory work, the central paradigm is to limit changes to evidence as much as possible; however, sometimes they are inevitable and, therefore, they are allowed as long as necessary, when the person is competent to decide and able to explain the relevance and implications [8, p. 6]. This stems from the trade-off between completeness of the evidence collection and intrusiveness, i.e., the number of induced changes, of the employed method [111]. With physical evidence, some analyses require the destruction of (parts of) the evidence item. Such trade-offs also exist in the digital domain, like the acquisition of the main memory of a running system, which necessitates an intrusive interaction and specific changes to it [120, pp. 2 f.] because the alternative is missing potentially crucial volatile memory.

Thus, the feature of unintentionality is characterized by the fact that the investigator does not consider that conducting a measure or its omission will result in unintended and unforeseen modifications. This feature distinguishes contamination sharply from deliberate changes. On the one hand, a conflict of investigative objectives might necessitate intentional changes that require accepting the unavoidable alteration. On the other hand, a badly intended change by the investigator with the goal to compromise the availability or usefulness of evidence to the forensics process [113] has to be considered as evidence tampering.

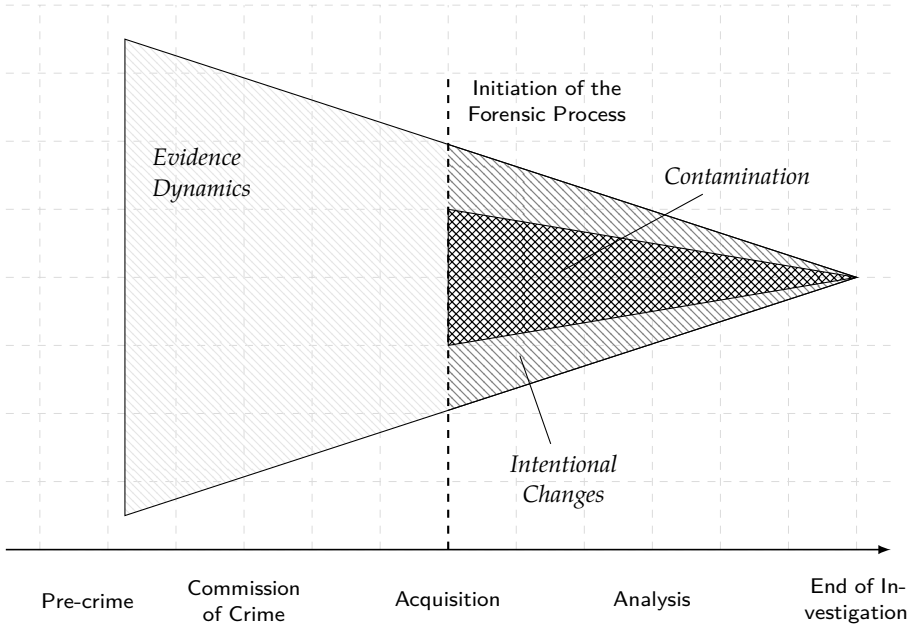

Figure 6.2: Relation of contamination to the broader concept of evidence dynamics described by Chisum and Turvey [44]. Most notable is that contamination can be considered a subset of evidence dynamics connected to the investigator's actions after identifying an object of relevance. It is important to strictly distinguish those from intentional changes necessitated by the analysis process itself or stemming from a conflict of investigative goals. The conical shape should indicate that there tends to be less opportunity for evidence dynamics and contamination as the process progresses toward the end of the investigation.

6.3.5 Definition

We now combine the four aspects presented above into a generalized definition of contamination:

Definition 6.3.1 (Contamination)**.** *Contamination* is any inadvertent transfer of traits to an object of relevance at any point in the forensic process.

In solid distinction to the concept of evidence dynamics, we emphasize temporal demarcation and unintentionality. Therefore, we consider contamination and its effects to be a subset of evidence dynamics, as illustrated in Fig. 6.2. A distinction between the more general evidence dynamics and contamination is helpful because the latter is directly affected by the investigators' actions. However, the exact point in time of its start might be subject to the forensic process model used.

Overall, all changes to relevant pieces of evidence after initiating the forensic process that violate its integrity without intent must be called contamination. We believe that the generalized definition provides a common ground by taking a new and more precise perspective on contamination because it abstracts from the specifics of physical traces and holds for both physical and digital evidence. While remaining applicable to physical

evidence, our new definition helps to characterize the existence of contamination in digital investigations more accurately, as illustrated by the examples provided in the following section.

6.4 An Example-guided Contemplation of Digital Evidence Contamination

With the goal to further improve the understanding of contamination in the digital domain and explore the applicability of the previously presented definition, we give several examples of DF fieldwork that result in contamination, and then we present non-contamination examples to delineate the phenomenon. Lastly, we examine two edge cases that illustrate that it is not always as easy to grasp and categorize the phenomenon.

6.4.1 Examples of Contamination

First, we discuss several occasions of contamination during live analyses; afterward, we refer to contamination in lab work. In both phases, we aim to capture the phenomenon's essence by highlighting and discussing the previously identified properties of possible and not even unlikely scenarios summarized in Table 6.1. This excursion should illustrate by example that adapting the definitions from the physical domain captures what we also intuitively understand as contamination.

6.4.1.1 Contamination During Live Responses

We now turn to several scenarios where a DF examiner interacts with running computer systems to acquire or analyze evidence.

Media Data Copying by Thumbnailing Service. To provide convenient file previews to the user, modern desktop environments, such as GNOME (as well as XFCE), run a thumbnailing service. GNOME, for example, relies on a D-Bus service called `Tumbler`, which provides thumbnails for various URIs and MIME types upon an application's request. To do so, it can resort to plugins, such as the `FFmpegThumbnailer`, `PopplerThumbnailer` and others.[47] Following the "Thumbnail Managing Standard" by Finke and Sessink [81], thumbnails are stored in the user's home directory. Conducting a live response using a storage medium containing some files of types that are considered by the thumbnailing service and its plugins can, therefore, lead to copying the file previews to the system under investigation if the investigator opens a directory containing those files via the file browser. The investigator who does not intentionally want to copy those data to the system under

[47] `https://gitlab.xfce.org/xfce/tumbler/-/tree/master/plugins`, commit `1d304f4`.

Table 6.1
Classification of the contamination examples discussed in Section 6.4 regarding the characteristics of the respectively observed transfer of traits and the phase of occurrence.

Scenario	Characteristics of the trait transfer			Phase
	Additive/Subtractive	*Direct/Indirect*	*Active/Passive*	
Media data copying by thumbnailing service	Additive	Direct	Active	Live response
File deletion by systemd-tmpfiles service	Subtractive	Direct	Passive	Live response
Malware artifact deletion by EDR action	Subtractive	Direct	Passive	Live response
Cloud storage synchronization	Additive/Subtractive	Indirect	Passive	Live response
Cache overwrite of smart home sensor	Additive/Subtractive	Direct	Active	Live response
Modifications by SQLite write-ahead log	Additive/Subtractive	Direct	Active	Lab work
Failing to employ software write blocking	Additive/Subtractive	Direct	Active	Lab work
Remote wipe of mobile device	Subtractive	Indirect	Passive	Lab work
Signal's RCE in Cellebrite UFED	Additive/Subtractive	Indirect	Passive	Lab work

investigation actively initiates a direct and additive transfer of traits of the thumbnail information to the object of relevance after the forensic process has been started. The use of an unclean storage medium in this example can be seen remotely analogous to a polluted cotton swab for collecting DNA trace material.

File Deletion by systemd-tmpfiles. Modern servers run many daemon processes to keep the system in good condition. One essential and on many Linux distributions already pre-installed service is `systemd-tmpfiles`, which automatically creates, deletes, and cleans up certain volatile and temporary files [59]. For long-running server systems, this service offers a convenient timer-based solution to clean up space routinely. During a live response, such a timer-based clean-up could be initiated by `systemd-tmpfiles`, which might delete crucial evidence from the system under investigation. The investigator's passive omission to disable this autonomous service of the operating system may lead to a direct and subtractive transfer of traits, although the forensic process has already been initiated. One might argue that digital deletion is not subtractive since it is just overwriting several data fields with zeros (or another value indicating its invalidity), which comprises an additive transfer. While this is plausible on a low level, it seems preferable to define this on the semantic level of evidence regarding the case-relevant available information. Furthermore, there is an increased risk of unwanted and irreversible removal of traces when working with devices that are capable of wear-leveling/trimming.

Malware Artifact Deletion by Endpoint Detection & Response. When a serious incident occurs, and the stress level rises, many processes run in parallel, and mitigation, as well as containment measures, might intervene more than required. An obvious example is an endpoint detection and response (EDR) agent doing its job after updating signatures. While this is desirable from the viewpoint of cyber defense, it may destroy crucial evidence, e.g., an implant stored as an obfuscated "one-liner PowerShell script" in an autostart-registry key, which must be considered an object of relevance for the investigation and should have been analyzed further. Such a deletion of malware artifacts by the EDR comprises a direct and subtractive trait transfer. When it happens before the forensic process has been kicked off, then it is certainly part of evidence dynamics; however, if it happens passively after initiating the forensic process due to omitting to synchronize between analysis and containment activities, we consider it to be a contamination, which might make the job unnecessarily hard and jeopardize the investigation.

Cross-Device Synchronization of Cloud Storage. Cloud storage services, like Dropbox, Google Drive, OneDrive, and iCloud, offer the ability to conveniently store and share files to colloborate. They synchronize seamlessly between different devices and provide the option to share specific data with certain users [78]. Unsurprisingly, such services have also been misused to share CSAM between offenders. In practice, investigators might conduct a live response when executing a search warrant in such a case to circumvent encryption. If the system under investigation is then still connected to the internet for some time, which might be necessary in some cases to lawfully acquire remote data accessible by the machine, background synchronization of one of the above-mentioned

cloud storage services might result in a passive transfer of traits. Given that some remote storage infrastructure is involved, this transfer is indirect in its nature. Imagine a case where an accomplice has been tipped off and tries to cover his tracks by deleting files from a shared folder, providing an excuse to the offender who might claim that he did not even want to possess the material and has already tried to delete it. Here, it is a decision that needs to be made while considering the benefits and risks of maintaining or removing network connectivity. The choice should be intentional, and informed; the discussion of contamination issues can feed into that decision.

Cache Modifications in Smart Home Devices. Another variant of a cache overwrite with a cyber-physical dimension can be found in the field of smart home devices. Certain lightbulbs are equipped with a motion sensor to trigger switching on the light. In the case of death investigations, for instance, timestamps related to motion sensor activations could be important evidence. However, some products, for example, specific motion sensors, cache such timestamps only for the last activation, thus having the potential to be easily overwritten [73, Fig. 10]. Imagining a missing person case, when the responding detectives enter the apartment to find a body, their movement inevitably leads to overwriting the motion sensor's last activation timestamp. This timestamp update comprises a direct and additive trait transfer, ultimately resulting in a loss of information. In such a scenario, detectives must consider this piece of data an object of relevance because it may be a crucial piece of evidence to narrow down the time of death.

6.4.1.2 Contamination During Lab Work

Looking at the following examples, we show that contamination cannot only occur when dealing with "live" systems but also in lab analyses.

Failing to Boot into Forensic Live OS. Nowadays, DF laboratories may resort to software write blockers in the form of bootable live operating systems (OSes), e.g., *Grml-Forensic*[48] or *TSURUGI Acquire*,[49] because non-destructive physical withdrawal of storage media is often not possible anymore when examining modern notebooks. Those forensic live OSes used for acquisition employ kernel-level write protection; however, if—for one reason or another—the examiner fails to hit the required key sequence to boot directly into the forensic live OS, various modifications, like updates, changed access timestamps, cache clean-ups, and other autonomous actions, happen. This inevitably results in contamination in the form of an actively induced direct transfer of traits—possibly a wild interplay of simultaneously adding and removing relevant information.

Remote Wiping of Mobile Devices. Not only the provider of tailored "crypto phones" like the infamous *EncroChat* devices but also the major manufacturers of mobile OSes, i.e.,

[48] `https://grml-forensic.org/`, accessed 13 Dec. 2023.
[49] `https://tsurugi-linux.org/`, accessed 13 Dec. 2023.

Apple and Android, provide their customers the ability to remotely conduct a factory reset of a smartphone linked to their account in case of loss or theft.[50] Criminals could misuse this feature to cover their traces after being targeted by investigating authorities who seized their devices. Therefore, a standard operating procedure (SOP) is to cut off the network connectivity of mobile devices. Nevertheless, there are various imaginable ways where remote wiping could occur, though: Obviously, one option is that the seizing officer might fail to activate the so-called "airplane mode". Another option could be that the seized device runs out of battery when lying on the backlog waiting to be processed. If the device is powered on with network connectivity enabled (e.g., a SIM still included, an eSIM, or Wi-Fi enabled), and a failure to use Faraday solutions occurs, the device might connect to the internet and receive the command to perform a factory reset. In both cases, we observe a passive, indirect, and subtractive transfer of traits, thus, leading to an inevitable loss of evidence because every potential digital object of relevance stored on the mobile device has been wiped.

Modifications by SQLite Write-ahead Log. SQLite databases are important datastores, especially common on mobile devices today. A journal, the so-called write-ahead log (WAL), is frequently used to provide atomicity and durability. Its task is to buffer changes made to the database; after a certain number of operations, they are executed at once [151, p. 561]. To conduct in-depth analyses, examiners often need to extract those file-based databases from previously acquired images. However, when regular SQLite database viewers are employed to view such databases, those tacitly apply changes recorded in a potentially provided write-ahead log, which may result in loosing important evidence by viewing only the "full up-to-date version" of the datastore [28]. While the original disk image's integrity remains untouched, the integrity of the derived piece of evidence is inadvertently violated since committing changes in the WAL constitutes an actively induced direct transfer of traits onto the object of relevance—the SQLite database. This can lead to an addition or subtraction of traces which is relevant if that derived copy is relied on as evidence. Think of a deleted entry that is committed just by opening the SQLite database.

Signal's RCE on Cellebrite UFED. In 2021, Signal's technical report, in which they documented how they have achieved remote code execution (RCE) on the forensic tools called Cellebrite UFED and Physical Analyzer, alerted many forensic practitioners. There, Marlinspike [159] described how a vulnerability in FFmpeg, which had been used for file parsing by these Cellebrite products, could lead to RCE during the processing of the acquired evidence and, hence, allows compromising the examination [159]. Though remaining a mind game, Marlinspike pointed out what we consider an interplay of anti-forensics and contamination. Hypothetically, an actual exploit payload could "seek to undetectably alter previous reports" or "compromise the integrity of future reports" [159].[51] Such an anti-forensic measure (undoubtedly constituting a criminal act, of course) would initiate a trait transfer after starting the forensic process during the DF examination and severely

[50] E.g., Google's " Find My Device App" for Android mobile devices.

[51] The term "report" is used to described the containers created by Cellebrite UFED's acquisition containing the device data.

hamper the integrity of the data acquisition, which constitutes the object of relevance here. The DF examiners have no intent for such a modification but are obliged to avoid it by using up-to-date and secure tools, much like they would be obliged to keep unauthorized persons out of the crime scene perimeter. This points out the risk of contamination stemming from anti-forensics.

6.4.2 Examples of Non-Contamination

To further delineate the phenomenon, we now present some counterexamples of contamination, i.e., where certain aspects of the definition are missing. This is critical to ensure that the definition does not go too far since the defined term has a clearly negative connotation.

Well-meaning IT-Support. Consider a compliance case, e.g., related to the disclosure of trade secrets, where high-level management commissions IT-support staff to inspect the respective employee's Windows system to substantiate a gut feeling. Basically, they are instructed to do some triage for e-discovery—a task for which they have never been trained. Unfortunately but unsurprisingly, they do it badly and browse the network share containing the trade secrets from the suspect's computer, hence, unwantedly initiating an additive trait transfer regarding all sorts of objects of relevance, like the "shellbags" and "recent files"-registry keys, Microsoft Edge's WebCache, "jump lists", and so on. While this scenario fulfills three out of four aspects of the definition presented in Section 6.3, the critical point is that the forensic process has not been initiated yet because the inspection aimed to verify a gut feeling and to establish initial suspicion potentially kicking off an investigation; therefore, we consider this an example of evidence dynamics with obviously very adverse effects. However, we want to stress here that the initiation of the forensic process does only necessitate a responsible party but not necessarily law enforcement authorities of some kind.

Accessing Hidden Disk Areas. The host protected area (HPA) and device configuration overlay (DCO) constitute disk areas whose access is prohibited by the disk controller [108]. When inspecting the storage devices in a case related to CSAM, examiners might consider making those particular disk sectors accessible to ensure they are not missing any relevant sectors. Using `hdparm`, they send ATA commands to modify the drive's configuration, which constitutes a transfer of traits. This example does not constitute contamination according to the definition above, since two definitional aspects are not fulfilled: The most obvious one is the absence of inadvertence; on the contrary, the action was purposeful and well-balanced since the examiners were aware of only changing the configuration related to the HPA and DCO to acquire more disk sectors. The second aspect ruling out contamination effects is to argue that the disk configuration parameters do not constitute an object of relevance here.

6.4.3 Examples of Edge Cases

Scrutinizing various scenarios, we present two examples of evidence acquisition whose categorization regarding contamination effects is debatable.

Jailbreak-enabled Mobile Data Acquisitions. Due to improved device security features, data acquisition of mobile devices is increasingly difficult. One common way to get access to the device data is the use of custom boot loaders. Since digital signatures secure boot chains on modern devices, boot ROM vulnerabilities have been exploited to gain more complete data acquisitions [88, p. 5]. An instance of this approach is the so-called `checkra1n` jailbreak tool for the mobile operating system iOS. Given that iOS runs in a restricted mode to impede access to internal functions or file system data, "jailbreaking" may be employed to gain root privileges and collect certain otherwise unaccessible pieces of data, such as system databases on the device [136]. Obviously, this is an invasive procedure entailing several changes to the device. By introducing a high amount of modifications that might not even be specifiable more closely since it is a closed-source binary, one could argue for the presence of an inadvertent trait transfer. However, in the absence of valid alternatives, examiners might decide to accept a relatively high amount of intrusiveness to improve the completeness of the evidence collection. Therefore, we refrain from calling this contamination since the trait transfers induced by exploiting the boot ROM vulnerability and the OS modifications are on purpose, and the modified objects are considered irrelevant, or at least less relevant than the ones gained; but this is certainly a decision that should be based on a full understanding of the effects of the modifications.

Page Smearing in RAM Acquisitions. In many investigations, examiners can find crucial evidence in main memory. Especially in intrusion analyses, random-access memory (RAM) acquisition and its subsequent analysis are needed to identify (fileless) malware artifacts. Still, in classical investigations where encryption is employed, it can be helpful to extract secrets for later use [109, 112].

However, on modern systems with more than 8 GB of RAM and heavy load, RAM acquisitions might suffer from so-called "page smearing". This term describes an inconsistency between the acquired page tables and the contents of the physical pages in the dump because they changed during the time needed to perform the complete acquisition. Besides losing potentially crucial data because of overwritten injected code or corrupted kernel data structures, this might also lead to a wrong assignment of memory pages to processes [35, p. 24] or corruptions of content data [178, p. 9:5] due to the duration, i.e., the temporal dimension, of the acquisition. The investigators intend to acquire an atomic, consistent, and correct snapshot [239] at the moment of running the main memory acquisition tool of their choice. Yet, there is actually a passive trait transfer stemming from the operating system's ongoing write operations after the forensic process has been started by deciding to acquire the RAM. The trait transfer here can be both additive or subtractive since evidence could be lost by inconsistent kernel data structures or added by some new information placed in memory pages. Though, it is uncertain if an object of relevance is concerned—a

question that is hard to answer. When dealing with a virtualized system, it is sensible to resort to the hypervisor and acquire the RAM by snapshotting the VM to avoid any smearing in the majority of cases. Nevertheless, if an investigator is confronted with a bare metal system, live smears are practically unavoidable; while not being optimal, collecting RAM with some smearing is definitely more efficacious than losing all the evidence.

Hence, we refrain from speaking in black-and-white terms and argue to apply the proposed definition with a sense of proportion.

6.5 Discussion

In view of the many practical examples and counterexamples provided in the previous section, we now broaden the scope to the overarching aspects and discuss the specifics and the intricacies of digital contamination.

6.5.1 Specifics of Digital Contamination

We identified three significant specialties impacting digital contamination, making it substantially different from the phenomenon in the physical sphere. Those are a direct cause of the features of the "abstract" cyber domain.

6.5.1.1 The Wealth of Autonomous Processes

Most notable is the existence of far more latent processes that might not be noticed by the investigators securing the crime scene. On a (strictly) physical scene, there are similar issues, e.g., insect activities on bodies, the volatilization of gaseous substances, or—most dramatically—the indirect transfer of substances containing DNA. Those, however, are limited in number and not nearly as numerous as arbitrarily running processes on digital systems. Given the real possibility of virtually arbitrary programs running on the system under investigation, the evidence and its meaning could be changed entirely by such processes during the examination; in the physical domain, the evidence could not just be dispersed or added arbitrarily. When working with digital evidence, it seems substantially harder to quantify the subtractive variant of contamination. In contrast, the additive variant might be detected and mitigated in some cases, e.g., thinking of a log entry caused by the investigator's action. However, it remains opaque, which data was potentially and unknowingly overwritten during the examination. Such opaqueness is a feature that is also reflected in the spatial dimension.

6.5.1.2 Spatial Detachment of Cause and Effect

Much like the commission of cybercrime, there is a spatial detachment of cause and effect, which is very specific to the cyber domain. We have illustrated several indirect trait transfers involving remote systems in the examples. Unlike with analogous evidence, the perpetrator or third parties might retain the ability to remotely intervene on a scene or impact seized evidence, although it has already been "secured" by competent responders. We can imagine several examples in that regard ranging from a still established command-and-control channel during a live analysis, a perpetrator who has not been taken into custody initiating remote wiping of a seized mobile device, to an anti-forensics measure such as the RCE exploiting a bug in forensic tools, as Marlinspike [159] described. That spatial decoupling indeed is a unique and delicate feature of the digital domain.

6.5.1.3 Complexity and Quantity Problems

A major difference is that (strictly) physical and digital crime scenes differ regarding volume, variety, and number of items that have to be considered by the investigators. Carrier [30] aptly named these specifics the *complexity* and *quantity problem*. With digital traces, it is thus often impossible to discriminate and evaluate their relevance at first. Furthermore, traces in digital systems can have various expressions and characteristics resembling their respective trace abstraction, basically the facet of the trace that is used for establishing an association. Since a tool-based translation of data is always needed to view data at a useful level of abstraction and actually make sense of the facets at hand, it is far more challenging here to identify the "evidence items" than it is in the physical world.

6.5.2 Intricacies of the Common Definition

Since our common definition (Section 6.3) mirrors physical contamination, we now want to critically review it in light of the specifics of digital evidence mentioned above. The *transfer of traits* and its connected qualities, the requirement of *missing intent*, and the temporal placement after the *forensic process initiation* are solid building blocks. However, discussing the *object of relevance* is very intricate because it is tremendously difficult to grasp in digital scenes.

One might argue to exclude this building block from the definition altogether; however, this would lead to an expansion, even a delimitation of covered situations, because it would capture alterations of objects which are clearly irrelevant to a case since they do not even remotely contribute to answering questions of factuality, guilt, and unlawfulness. We want to underline that if an irrelevant object is modified, it should not be called contamination because dropping such a restriction would mean that every real-world crime scene would be contaminated, even with perfect handling of the crime site.

When working with digital evidence, the problem of determining relevance becomes even more precarious: in many cases, there is only a vague idea of what might be relevant at first. During the analysis, examiners deal with different levels of abstraction, but on which level should we operate to determine contamination effects? To date, there is no method to adjudge the perfect level of abstraction to work on, and so there is none to determine the abstraction level to pin down contamination; however, we imagine that the previously developed notion of relevant evidence can help to confine this definitional building block—at least in regard to the provisional case-related hypotheses.

Linked to these thoughts, we assume a temporal aspect of relevance and put up for discussion, whether it is contamination or not, if an object is accidently, but incorrectly, deemed to be of relevance, and is inadvertently altered. We would argue to consider this to be contamination at that point, though it will be practically irrelevant. The other way round, an inadvertent alteration of an object initially identified as irrelevant but later considered as relevant is a comparably complex edge case. Here, the discussion revolves around whether the forensic process regarding this object has been initiated or not. This is easier to answer in the physical domain: If it has not been seized and packaged in an evidence bag, it will likely fall under the concept of evidence dynamics. In digital, however, we see the need for future discussions.

6.5.3 Implications of the Improved Understanding

Intending to tackle this problem, standardization organizations created tailored regularities in the analogous world: The ISO-standard 18385:2016 [128] regulates the requirements for products to collect, store, and analyze DNA evidence; FSR-G-206 [226] published by the British government provides guidance on how to control and avoid contamination in scene examination involving DNA evidence recovery. Still, regarding digital evidence, there exist only general guidelines [e.g., 127]. However, given its increasing importance, it seems imperative to improve awareness and procedures in this subdiscipline.

We can infer several learnings from the examples presented in Section 6.4. Though, we refrain from trying to derive any specific guidelines since those would be either too specific (e.g., "check when the `systemd-tmpfiles` timer triggers next") or would be far too broad and general (e.g., "improve the education of the examiners and limit direct interaction with the system under investigation"). A general consequence is that any inadvertent change to objects of relevance has to be avoided. If this is—for whatever reason—not possible, the alterations have to be detected and labeled as such to enable a correct interpretation given that knowledge. Overall, this boils down to having intent for any alteration—including its side effects. Purposeful alterations, however, inevitably require a solid understanding of the system, the action, and its implications.

Easing this task requires SOPs that are peer-reviewed. There is not only a requirement for evaluating the correctness of the results but for reviewing the potential contamination effects. SOPs derived from there will profit in soundness, while the provided definition is also helpful to evaluate non-SOPs.

Regarding traditional crime scene work, a lack of research in quality assurance and a unified understanding of quality as such has yet to be sought [47]. In the digital domain, this observation can be considered equally delicate [179]. In digital forensics, we can already resort to cryptographic hash functions to ensure data integrity and similarity hashing, e.g., by employing a piecewise approach [24], for determining the resemblance of data. Nevertheless, there is a strong need to develop methods and metrics for measuring acquisition quality and the intrusiveness of analysis tools and techniques [111]. Here, we need both simple metrics and a contextualized focus on the semantics of evidence.

In that regard, we want to highlight two points: First, it is essential to conduct rigorous experiments to collect leftover artifacts and quantify the impact of specific system interactions and forensic tools, which is a necessary prerequisite even to be able to have intent. Second, we would like to initiate the strive for precise documentation of contamination in live scenarios and lab environments, as some of our examples already did. We anticipate that there might be not only different types of circumstances but also distinct classes of contamination. An appropriate repository listing such occurrences can help rule out misinterpretations and identify fields to improve evidence acquisition, handling, and analysis. Based on such findings, the digital forensics research community should attempt to develop methods for identifying digital contamination and propose targeted countermeasures to minimize it or propose techniques to prevent contamination, much like it has already happened concerning DNA evidence in classical forensic science.

6.6 Summary

In the previous chapters, we focused on the expressiveness and relevance of (digital) evidence to reach investigative goals. To do so, availability, accuracy, completeness, and known provenance of the findings are always crucial to providing resilient and reliable digital evidence usable for solving cases, as sketched out before. Implicitly, we assumed here to always work with evidence of the best possible quality; however, the expressiveness of digital traces might be severely hindered in the course of their acquisition and analysis if not done correctly. Thus, in the present chapter, we scrutinize contamination as one effect that can do so, which is well-known in traditional forensic science but has been underexposed in digital forensic science.

To gain an understanding, we revisit several works on contamination when handling traditional evidence, especially DNA traces. Afterward, we look at existing definitions and identify that they are not applicable to digital evidence (Section 6.2). Hence, we propose a definition of evidence contamination (Definition 6.3.1) that not only covers physical evidence but also extends to digital evidence (Section 6.3). To recapitulate, the proposed notion of contamination is characterized by

(a) a transfer of traits to
(b) an object of relevance
(c) at any point in the forensic process
(d) without intent.

Furthermore, we integrate the proposed conception of contamination into the broader scheme of evidence dynamics and explain how we consider contamination a subset of it. We present several examples of partly severe contamination effects in digital forensics to shed further light on it. Those illustrate how easily digital evidence might be contaminated, resulting from a more or less convoluted trait transfer, classified as additive or subtractive, passive or active, and direct or indirect. Eventually, this leads either to a wrong assessment of investigative hypotheses or renders an assessment impossible because the evidence is corrupted. Even if contamination does not imply wrong conclusions regarding the guilt or innocence of an individual, contamination in the sense of our harmonized definition might cause serious misinterpretation errors of investigative hypotheses and, therefore, severely hamper the reconstruction of the deed (or incident). To delineate contamination from general evidence dynamics and other effects, we present counterexamples of contamination and point out where certain aspects of the definition are missing. Finally, we look at edge cases, i.e., jailbreaking of seized phones and page smearing during RAM acquisitions, where a precise classification is not apparent (Section 6.4).

Based on the presentation of examples, counterexamples, and edge cases, we discuss the newly proposed definition. We identify that the wealth of autonomous processes, the spatial detachment of cause and effect, as well as the complexity and quantity problems related to digital crime scenes, constitute specifics that favor this phenomenon in the cyber domain. Moreover, those features lead to the identification of several intricacies revolving around the many abstraction layers when dealing with digital evidence, making digital contamination, in some respects, much more complex and partly ambiguous compared to its physical counterpart. Based on these insights, we derive several implications of the new definition: Here, we identify the need for rigorous experiments to collect leftover artifacts, quantify the impact of specific system interactions, and understand the related mechanisms favoring contamination. Moreover, we formulate the necessity for precise measurement and documentation of contamination effects, as we started to do so in this chapter, as well as develop methods for identifying contamination and countermeasures to minimize its effects (Section 6.5).

To conclude, there is a real possibility that examiners contaminate digital evidence during the acquisition and analyses, which impacts the flawless assessment of hypotheses. Hence, it is necessary to raise awareness of that phenomenon in the community of practitioners and researchers as one effect to be studied that can hamper the expressiveness of digital evidence.

7 Conclusion

7.1 The Big Picture of Relevant and Expressive Evidence

For the past 20 years, forensic scientists and criminalists alike have been confronted with a steep rise in the amount and complexity of digital data in criminal cases. The task in (digital) investigations is to pose apt investigative hypotheses and look at the expressive yet relevant traces. However, both researchers and practitioners face rapid technological advancements but an incomplete understanding of basic principles in criminalistics; hence, the present work tackles the fundamental but complicated question of what constitutes "sufficient digital evidence" and how to find it and solve the criminalistic task eventually.

To this end, the present thesis focuses on developing an understanding of the foundational attributes of digital evidence and the effective yet faultless use of digital evidence in cybercriminalistics employing an investigative perspective. The former is primarily a question in the field of forensic computing. At the same time, the latter constitutes a core topic of cybercriminalistics as an applied craft of practical importance that builds up on the gained understanding. Hence, the present dissertation is divided into two parts: the first part, up to and including Chapter 4, uses a model-based approach to grasp the relevance and expressiveness of digital traces and understand what constitutes sufficient and necessary evidence. The second part deals with the practical application of the formerly established concepts, aiming to find and effectively chain relevant traces, avoiding contamination.

7.1.1 Accomplishments

We differentiated digital forensic science from cybercriminalistics and returned to the investigative core of criminalistics first. Approaching the research question of when digital evidence is considered to be relevant and expressive yet reliable, we derived formal notions of the concepts of *relevance* and *expressiveness of* (*digital*) *evidence* concerning investigative hypotheses—essential attributes that remained implicit in the field of digital forensic science but provide valuable insights regarding their meaningfulness for an investigation. Briefly expressed informally in natural language, we say that a facet is relevant to an investigative hypothesis if it either supports or refutes it. Expressiveness, in turn, is characterized by the wealth of hypotheses a facet can assess. These definitions enabled us to build an investigative knowledge base rooted in formal understanding. Going further, we improved the understanding of accuracy, exhaustive and decisive completeness, and authenticity by providing rigorous definitions. Besides improving the theoretical understanding of these

critical aspects of the overall investigative process and the nature of digital evidence, we showed how the formal definitions can be integrated into the existing thinking models known as the *Criminalistic Cycle* [240] and the *CSI model* [84] to sharpen the employed intuitions and build the ground for improving targeted selection and solidified assessment processes for facets to enhance reasoning of case-relevant hypotheses. As a first step, we demonstrated the application (a) by proposing the *Facet-oriented Criminalistic Cycle* and (b) by calculating the *expressiveness ratios* of facets.

Having an abstract understanding of evidential relevance, we approached the quest of finding such relevant evidence on a theoretical level that is either necessary or sufficient for concluding the occurrence of an event in the past—a task that is known as the specific reconstruction problem (SRP). We revisited Dewald's characteristic evidence method [64] as one solution to the SRP and pointed out its insufficiency by ignoring state information, i.e., preconditions and chains. Aiming to define the attributes of necessity and sufficiency of digital traces in forensic event reconstruction in a more complete sense, we then employed an automata-theoretic approach. We used linear-time temporal logic to define sufficient and necessary evidence: The former allows us to prove the execution of an action σ in the automaton by referring to those facets that are immediate or mediate effect of σ. The latter allows us to refute the execution of an action σ by stating all the facets that must be observable in all subsequent states after σ has happened. Using model checking software, these notions of general reconstructability classes can be practically used to calculate evidence sets of the evidence classes named, which materializes the concept of relevance and expressiveness as a basis to construct an investigative knowledge base.

The automata-theoretic approach is only applicable in idealized settings because of the difficulty in crafting models of real systems and the limits regarding the models' sizes. Unfortunately, it cannot directly be transferred to real-world investigations yet. This limitation raised the research question of how telling evidence can be found and documented for use in real-world investigations to achieve criminalistic goals. The work on this question led to the identification of a severe abstraction gap between general process models and case-specific concretizations in actual investigations. To introduce an intermediary step, we proposed the collection and use of phenomenon-specific knowledge to encode what is relevant evidence when investigating a specific criminal phenomenon. Using *phenomenon-specific investigative knowledge bases* in the form of cognitive maps allows the practitioner to bridge the abstraction gap mentioned above by introducing a meso-level of abstraction in the form of cognitive maps to capture phenomenon-specific knowledge, which supports the quest to find relevant traces in a more pragmatic setting. To demonstrate the applicability of this method, we presented an exemplary cognitive map of technical investigations in cases of botnet crime, the correctness and completeness of which have been validated by conducting interviews with domain experts.

Having a notion of expressive as well as relevant traces and the investigative measures to proceed in the investigation available in the form of a phenomenon-specific knowledge base empowers the investigator to solve the criminalistic task. However, there still exist effects that can hinder the expressiveness of digital traces. We scrutinized *contamination of digital evidence* as one underexposed instance of such a failure. Given the absence of any work on the contamination of digital evidence, we referred to existing definitions

in the sphere of traditional forensic science and derived a novel definition of evidence contamination—applicable both for physical and digital evidence. This definition has been substantiated and validated by discussing examples, counterexamples, and edge cases of contamination of digital evidence and led to identifying specifics, intricacies, and implications——essential to know about for both the research community and forensic practitioners.

7.1.2 Integration of Results

In the above summary of accomplishments, a common thread running through the individual results has already become apparent. The individual results can be considered components of a structured solution to the criminalistic task. Hence, their integration led to the development of the *Cyber-traceological Model*, which provides a unified yet structured method of generically translating investigative demands to the relevant traces. In Chapter 1, we already introduced a coarse and simplified version of it briefly as Fig. 1.2 to state the relation of the individual chapter contributions; now that the reader knows the details, we present and explain the full-fledged model, as illustrated in Fig. 7.1.

The Cyber-traceological Model. We place an investigative knowledge base at the core of the Cyber-traceological Model (and at the figure's center). This core component can be based on the formal notions of relevance (Definition 3.3.3), as exemplarily shown in the upper right-hand side of Fig. 7.1. In an idealized setting, where a model $\mathcal{M}$ of the relevant system under investigation is available, the investigative knowledge base can be filled using the evidential specifications, as put up in Chapter 4, to calculate sufficient and necessary evidence sets for the actions available in the automata with a model checker, as illustrated at the top of Fig. 7.1. A more pragmatic and real-world-oriented method to construct the investigative knowledge base using phenomenon-specific knowledge is shown at the bottom of Fig. 7.1. Here, the creation of cognitive maps is presented in order to infer relevant and expressive traces regarding a specific criminal phenomenon based on experiential knowledge from different sources, such as documents or domain experts, employing a mining process, that is constituted by coding, inference, and enriching. Then, the knowledge base, either constructed by employing the automata-theoretic approach or by using phenomenon-specific knowledge, is used to derive relevant evidence as the outcome shown on the right-hand side of Fig. 7.1. Hence, it provides a conception of which facets to collect in order to assess the set of case-related hypotheses.[52]

We propose using the newly developed Facet-oriented Criminalistic Cycle (FoCC) as a process model, which encloses the figure circularly, embedding the translation into an iterative procedure. When confronted with a concrete case, the investigators come up with

[52] It is apparent that the formalized investigative knowledge base, as it is shown for exemplary purpose in Fig. 7.1, contains the bare minimum of information to solve cases; it neither fully represents the knowledge of the model for the case of necessary and sufficient evidence nor the investigative links incorporated in phenomenon-specific knowledge. We decided not to reflect this formally because we suspect that the added value of blowing up the model is limited for human reasoning.

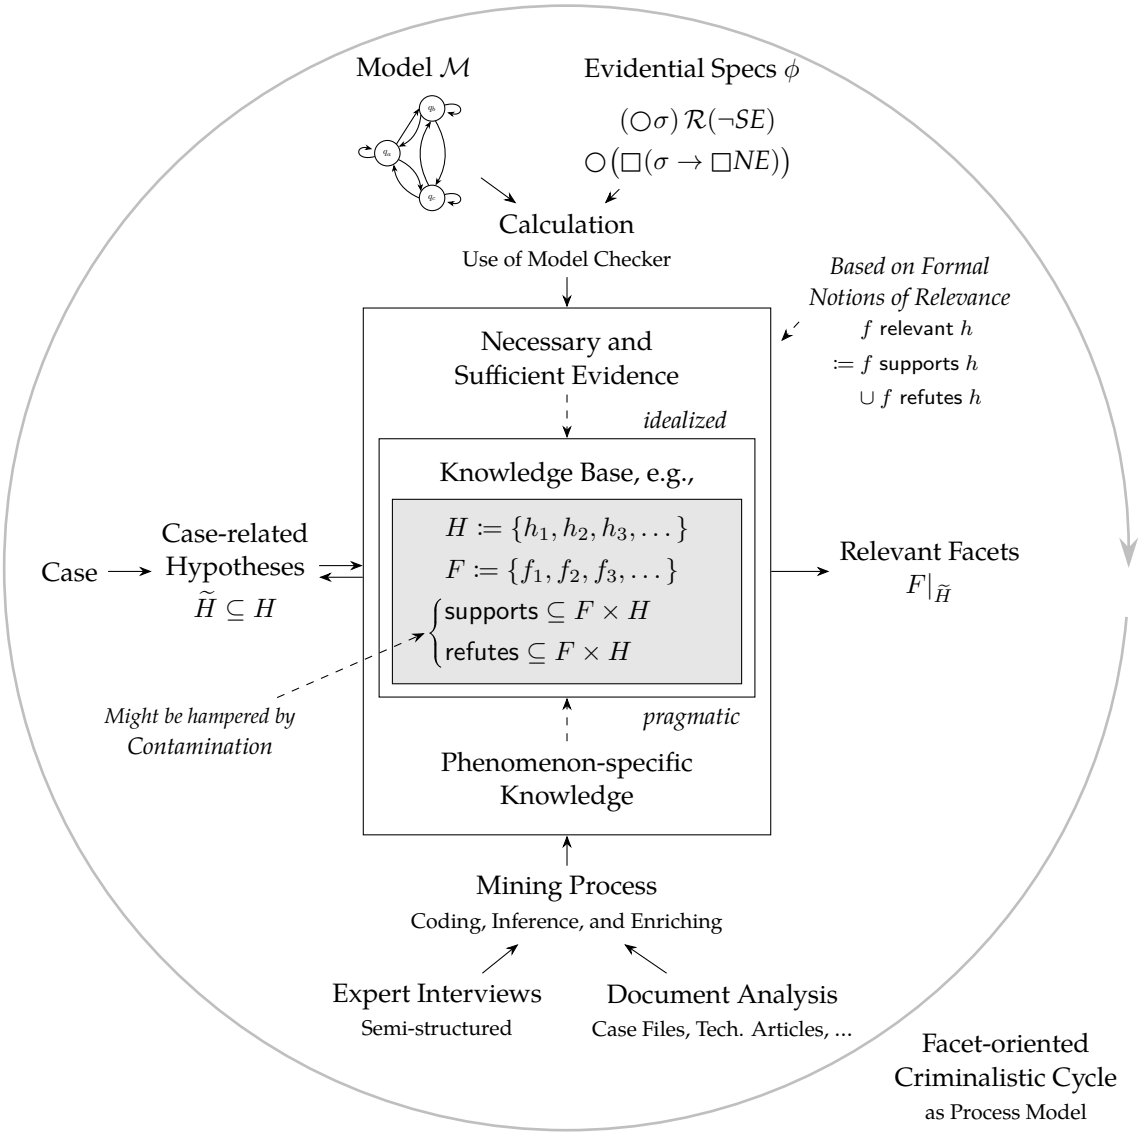

Figure 7.1: Detailed version of the Cyber-traceological Model. The full-fledged version of the model depicts the overall process and the basic building blocks to provide a structured method of translating investigative hypotheses to relevant facets by employing an investigative knowledge base that can be built using different approaches, as illustrated in this thesis.

an initial set of case-related hypotheses $\widetilde{H}$ about the deed. To answer these, they query the investigative knowledge base and bring the set of relevant facets $F|_{\widetilde{H}}$ to light.

This set of facets will help to assess (at least) the hypotheses in $\widetilde{H}$ according to the knowledge base and, hence, can be used as a basis for collecting certain individual facets. It is expected that this will lead to iterative updates of the set of case-related hypotheses, which in turn will lead to continually collecting facets until the investigator determines the investigation to be complete, as it has been defined as part of the FoCC (Section 3.6.1), which provides the procedural frame for the Cyber-traceological Model. Note that there is an intersection between the construction of case-related hypotheses and their answering, as indicated by the opposing arrowheads, since the knowledge base contains helpful or crucial information to form apt hypotheses. However, this connection has to be explored deeper in future work.

In this thinking model, we can even locate contamination effects using the definition of *accuracy* as one of the reliability criteria: Suppose contamination happens and remains undetected by the investigators. In that case, the accuracy (Definition 3.4.1) of the knowledge base as one reliability feature is hampered, i.e., the investigators will draw the wrong conclusion because their conceptions of what a facet means concerning a hypothesis do not match the reality. So, knowing about contamination effects prevents the investigators from drawing wrong conclusions.

Significance. Traceology, as a holistic combination of the branches of forensic science [158, 194], is primarily concerned with assessing hypotheses related to past events based on facets, the observable parts of traces. Recently, the trace and the formal study of its nature have been of increased interest in the forensic science community—aiming to unify traditional and digital branches [131]. Given the absence of any straightforward method to find "sufficient digital evidence", the present thesis took up this development. It extended the formal study of cyber-traceology, viz., the elaboration of the meaning of digital traces for investigative hypotheses. Improving the foundational understanding aims to contribute to solving the second subproblem of the criminalistic task, i.e., determining relevant digital traces that can be used to assess those previously identified case-related hypotheses. For doing so, it is salient to have an unambiguous understanding of the attributes of relevance, expressiveness, necessity, and sufficiency of evidence and the interplay of findings, as developed in this dissertation and reflected in the Cyber-traceological Model.

7.2 Future Directions

Both digital forensic science and cybercriminalistics are relatively new (sub)fields compared to the long-standing tradition of forensic science. Interestingly, the questions tackled in this thesis are, in their more profound nature, similar to what pioneers like Edmond Locard or Hans Gross had faced back in time. However, their convolution seems to be amplified by the features of the digital domain, especially by quantity and complexity problems of evidence in fast-changing IT environments. By facing these, the present thesis

worked at the foundation and contributed to the understanding of fundamental connections and attributes linked to digital evidence employing a model-based approach. Still, it also uncovered that many contemporary questions, more or less fundamental, remain to be answered; hence, this section points out several directions for further research that are grouped by chapters and directly or indirectly related to the respective results presented there.

Relevant and Expressive Digital Evidence Though we can use the concept of relevance and expressiveness in practice already today, we identify notable shortcomings related to the formalization. One drawback is that the employed model is unable to deal with uncertainty so far since it uses crisp logic and deals only with boolean cases of relevance. However, the use of likelihood ratios by expert witnesses is common to assess the defense and prosecution hypotheses [195]. Hence, there is a need to fully elaborate and adapt the definitions presented in this chapter for fuzzy cases and to employ those in probabilistic models. While adapting the supports and refutes relations to be fuzzy relations so that facets get related to hypotheses with grades from the unit interval seems to be straightforward. A fitting fuzzy semantic for the used connectives must be picked to adapt the definitions presented.

Furthermore, it must be underlined that the formalization depends on establishing the supports and refutes relations. As we showed, this can be accomplished either by estimates of domain experts as proposed by Kwan et al. [144] and the methods described in Chapter 5 to a less formal degree or by using rigorous mathematical approaches as suggested in Chapter 4. The former is mainly subjective, not necessarily representative, and thus prone to errors; the latter is computationally intense and requires modeling the system under investigation beforehand, as we will discuss in the following paragraph. So, we see a need to find novel approaches to systematically construct probabilistic models and bring forensic statistics to digital forensic science, on the one hand, and to combine research in theoretical computer science with methods of digital forensic science, on the other.

Necessary and Sufficient Digital Evidence For the proposed automata-theoretic approach, we identified the tedious construction, the limited model size, and the computational intensity as severe drawbacks regarding a more widespread practical application of the *NE/SE* method. Thus, future work should aim to determine suitable trace abstractions to model forensically relevant parts of a system (or reality) and strive to develop automated approaches for the generation of state machines. Here, we envision the application of model learning—a method developed in the research field of automata theory. Cho et al. [46] ventured a first foray into model learning with a forensic application. They showed that it is possible to infer protocol state machines for the analysis of botnet command and control protocols. However, there are still many research challenges regarding operations on data, the quality of models, and the lack of expressiveness of Mealy machines [230, p. 94 f.]. Nevertheless, we expect great benefit from learning black-box state machine models of a program under investigation, which would then enable investigators to determine necessary and sufficient evidence to improve forensic event reconstruction tremendously. Moreover, our *NE/SE* approach via linear-time temporal logic will, in principle, allow for

investigating more complex behavior composed of more than one action—definitely an auspicious possibility to explore. Finally, there is a need to explore more efficient ways of calculating evidence sets since the current algorithm exhibits a high time complexity in that it grows exponentially to the number of variables (the resulting number of partial valuations) and linearly to the number of actions (corresponding to the number of sets to be computed). Lastly, we suggest exploring other kinds of system models, especially non-deterministic state machines, to investigate their effect on the evidential concepts proposed in this work.

Phenomenon-specific Digital Evidence We proposed using phenomenon-specific knowledge to bridge the abstraction gap by providing the investigators with a concrete understanding of relevant traces. Voigt [238] continued the work and proposed a general procedure to acquire and represent phenomenon-specific cybercrime knowledge. While this provides the means to acquire phenomenon-specific knowledge inductively, the delicate design to capture relevance and expressiveness in its entirety has to be further researched besides mechanisms to update knowledge bases in order to continuously ensure their actuality. Aiming to ensure practicality, we see the need to assess the usability and usefulness of knowledge bases encoded as node-link relations in general and cognitive maps in specific. To measure the efficacy of this approach in supporting investigations, we propose collaboration with practitioners to conduct user studies and assess the impact of task-relevant information in the form of phenomenon-specific knowledge bases on the analysis results, in a similar way Sunde and Dror [217] did for investigating "biasability" of examiners by providing task-irrelevant information.

In the spirit of Gross and Geerds [98], we could imagine an ongoing collection of phenomenon-specific knowledge for all significant cybercrime phenomena, e.g., various types of online fraud (e.g., investment fraud, romance scams, and others), ransomware, CSAM, dark web narcotics trafficking, and many more, as we did for the phenomenon of botnet crime, to build up an encyclopedia of facets and the hypotheses assessed by them for real-world application—much like the vision of the founders of modern-day criminalistics. Lastly, the research community should develop ways to quickly understand new phenomena and identify relevant traces to identify and convict suspects of crime, where no experiential knowledge is yet available to acquire, which is probably the most challenging task.

Contamination of Digital Evidence The newly proposed definition opens up to threads of future work. On the one hand, there is the need to deepen the theoretical understanding of this phenomenon; on the other hand, there is a pile of applied follow-up work. From a theoretical point of view, further research has to be conducted to pin down the "object of relevance", where we expect that the foundational notions developed in Chapter 3 might be of help. Additionally, we suspect there are distinct classes of contamination, both handling digital and physical evidence, where contamination effects are rooted in common circumstances or mechanisms that should be investigated. From a more applied point of view, we see the need to systematically carry out experiments to identify and understand factors that favor contamination, enabling researchers to infer SOPs to avoid contamination.

Such efforts could eventually lead to the development of a standard for handling digital evidence, resembling the ISO standard 18385:2016 [128] for DNA handling in the physical world.

Landscape of Digital Investigations To the best of our knowledge, we were the first to define and delineate the branch of cybercriminalistics. Opposingly to "traditional" criminalistics, its cyber-branch is characterized by the fact that many, if not most, insights stem from academic research (outside of law enforcement agencies) due to the technicality of the matter. From a general perspective, researchers should hence strive to solidify technical investigations and cyber-traceology, improve digital criminalistic tactics, and empirically study cybercrime phenomena in a targeted manner, as the present work intended to do. Going back to the criminalistic task, which we divided into two subproblems, i.e., finding apt hypotheses and then finding facets of relevance for these hypotheses, we see the need to address the first subproblem of structured hypothesis generation in future work. While there exist methods aiming to understand and support hypothesis generation [50, 51], many open questions revolve around the quest to find ways that empower the investigator to systematically come up with apt investigative propositions, i.e., those hypotheses exhibiting case-related relevance, and to refine and relate those to one another, which can be considered complex problems in argumentation theory [17].

7.3 Emerging Prospects

"What happened and who did it?"—referring to the question central to criminalistics posed at the beginning, we have to affirm that as much as digitization permeates everyday life, as much significance has come to digital traces nowadays for answering this question to solve criminal cases. The present thesis leaps toward improving the understanding and interpretation of digital traces on a foundational level at the intersection of forensic computing and cybercriminalistics.

In essence, the developed notions of relevance, expressiveness, necessity, and sufficiency serve as a basis for work that aims to make investigations more precise and effective in practice eventually. In retrospect, we satisfactorily note that this has already been shown with the refined conception of contamination and the proposed method to overcome the abstraction gap in investigations using phenomenon-specific knowledge. Lastly, we see many questions of both digital and traditional forensic science that remain mysteries that need to be solved. Could model-based approaches, as employed in this thesis to gain fundamental insights and conceptions of digital evidence, also be of pertinence to traditional forensic science?

Bibliography

[1] Rolf Ackermann. Kriminalistik – Wissenschaft – Gesellschaft. In Heiko Artkämper and Horst Clages, editors, *Kriminalistik gestern-heute-morgen*, Schriftenreihe der Deutschen Gesellschaft für Kriminalistik e.V., chapter 1, pages 21–48. Richard Boorberg Verlag, Stuttgart, Germany, 1 edition, 7 2014.

[2] Rolf Ackermann. Methodik der kriminalistischen Fallbearbeitung - Kriminalistische Fallanalyse. In Horst Clages and Rolf Ackermann, editors, *Der rote Faden: Grundsätze der Kriminalpraxis*, pages 131–160. CF Müller GmbH, 13 edition, 2017.

[3] Frank Adelstein. Live forensics: diagnosing your system without killing it first. *Commun. ACM*, 49(2):63–66, 2006. doi: 10.1145/1113034.1113070.

[4] Ross J. Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J. G. van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the Cost of Cybercrime. In Rainer Böhme, editor, *The Economics of Information Security and Privacy*, pages 265–300. Springer, 2013. doi: 10.1007/978-3-642-39498-0_12.

[5] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. Understanding the Mirai Botnet. In Engin Kirda and Thomas Ristenpart, editors, *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, pages 1093–1110. USENIX Association, 2017. URL `https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis`.

[6] Association of Chief Police Officers. Practice Advice on The Core Investigative Doctrine, 2005.

[7] Association of Chief Police Officers. ACPO Good Practice Guide for Computer-Based Electronic Evidence, 2005.

[8] Association of Chief Police Officers. ACPO Good Practice Guide for Digital Evidence, 03 2012.

[9] Robert Axelrod. *Structure of Decision: The Cognitive Maps of Political Elites*. Princeton University Press, 1976. ISBN 978-1-400-87195-7.

[10] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008. ISBN 978-0-262-02649-9.

[11] John Bandler and Antonia Merzon. *Cybercrime Investigations: A Comprehensive Resource for Everyone*. CRC Press, 6 2020. doi: 10.1201/9781003033523.

[12] Nicole Beebe and Jan Guynes Clark. A hierarchical, objectives-based framework for the digital investigations process. *Digit. Investig.*, 2(2):147–167, 2005. doi: 10.1016/j.diin.2005.04.002.

[13] Emma Bell, Alan Bryman, and Bill Harley. *Business Research Methods*, volume 6. Oxford University Press, 2022.

[14] Brian A. Benczkowski. Remarks at the Yakubets Press Conference, 12 2019. URL `https://www.justice.gov/opa/speech/assistant-attorney-general-brian-benczkowski-delivers-remarks-yakubets-press-conference`.

[15] Charles Berger. Criminalistics is reasoning backwards. *Nederlands Juristenblad*, 85: 784–789, 2010.

[16] Charles E.H. Berger, Hans H. de Boer, and Mayonne van Wijk. Chapter 3.2 – Use of Bayes' theorem in data analysis and interpretation. In Zuzana Obertová, Alistair Stewart, and Cristina Cattaneo, editors, *Statistics and Probability in Forensic Anthropology*, pages 125–135. Academic Press, 2020. ISBN 978-0-12-815764-0. doi: 10.1016/B978-0-12-815764-0.00014-9.

[17] Floris Bex. Argumentation and Evidence. *Philosophical Foundations of Evidence Law*, page 183–198, 9 2021. doi: 10.1093/oso/9780198859307.003.0014.

[18] R. A. F. Bhoedjang, Alex van Ballegooij, Harm M. A. van Beek, J. C. van Schie, F. W. Dillema, Ruud B. van Baar, F. A. Ouwendijk, and M. Streppel. Engineering an online computer forensic service. *Digit. Investig.*, 9(2):96–108, 2012. doi: 10.1016/j.diin.2012.10.001.

[19] Mattheüs B. Blankesteijn, Aya Fukami, and Zeno.J.M.H. Geradts. Assessing data remnants in modern smartphones after factory reset. *Forensic Science International: Digital Investigation*, 46:301587, 2023. ISSN 2666-2817. doi: 10.1016/j.fsidi.2023.301587.

[20] Alexander Bogner, Beate Littig, and Wolfgang Menz. *Interviews mit Experten: eine praxisorientierte Einführung*, volume 1. Springer-Verlag, 2014.

[21] Rainer Böhme, Felix C. Freiling, Thomas Gloe, and Matthias Kirchner. Multimedia Forensics Is Not Computer Forensics. In Zeno J. M. H. Geradts, Katrin Franke, and Cor J. Veenman, editors, *Computational Forensics, Third International Workshop, IWCF 2009, The Hague, The Netherlands, August 13-14, 2009. Proceedings*, volume 5718 of *Lecture Notes in Computer Science*, pages 90–103. Springer, 2009. doi: 10.1007/978-3-642-03521-0_9.

[22] Nick Bostrom. Are we living in a computer simulation? *The philosophical quarterly*, 53(211):243–255, 2003.

[23] Owen Defries Brady. *Exploiting Digital Evidence Artefacts*. PhD thesis, King's College London, 2019.

[24] Frank Breitinger and Harald Baier. A Fuzzy Hashing Approach based on Random Sequences and Hamming Distance. *7th annual Conference on Digital Forensics, Security and Law (ADFSL)*, (TUD-CS-2012-0164):89–101, 5 2012. URL `http://tubiblio.ulb.tu-darmstadt.de/102073/`.

[25] Wolf-Dietrich Brodag. *Kriminalistik – Grundlagen der Verbrechensbekämpfung*. Kriminalistik und Kriminologie. Richard Boorberg Verlag, Stuttgart, Germany, 8 edition, 2001. ISBN 3-415-02756-2.

[26] Dominik Brodowski. *Verdeckte technische Überwachungsmaßnahmen im Polizei-und Strafverfahrensrecht: Zur rechtsstaatlichen und rechtspraktischen Notwendigkeit eines einheitlichen operativen Ermittlungsrechts*, volume 119 of *Tübinger Rechtswissenschaftliche Abhandlungen*. Mohr Siebeck, 2 2016.

[27] Simon Busard and Charles Pecheur. PyNuSMV: NuSMV as a Python Library. In Guillaume Brat, Neha Rungta, and Arnaud Venet, editors, *NASA Formal Methods, 5th International Symposium, NFM 2013, Moffett Field, CA, USA, May 14-16, 2013. Proceedings*, volume 7871 of *Lecture Notes in Computer Science*, pages 453–458. Springer, 2013. doi: 10.1007/978-3-642-38088-4_33.

[28] Alex Caithness. The Forensic Implications of SQLite's Write Ahead Log. Technical report, CCL Solutions Group, 5 2012. URL `https://web.archive.org/web/20220922074815/https://digitalinvestigation.wordpress.com/2012/05/04/the-forensic-implications-of-sqlites-write-ahead-log/`.

[29] Michael Capellmann. Die Kriminalistik im Spannungsfeld zwischen forensischer interdisziplinärer Wissenschaft und kriminalistischer Handlungslehre. *Kriminalistik*, pages 374–377, 6 2018.

[30] Brian D. Carrier. Defining Digital Forensic Examination and Analysis Tool Using Abstraction Layers. *Int. J. Digit. EVid.*, 1(4), 2003.

[31] Brian D. Carrier. *File System Forensic Analysis*. Addison-Wesley Professional, 2005. ISBN 0-32-126817-2.

[32] Brian D. Carrier. *A hypothesis-based approach to digital forensic investigations*. PhD thesis, Purdue University, 2006.

[33] Brian D. Carrier and Eugene H. Spafford. Getting Physical with the Digital Investigation Process. *Int. J. Digit. EVid.*, 2(2), 2003. URL `http://www.utica.edu/academic/institutes/ecii/publications/articles/A0AC5A7A-FB6C-325D-BF515A44FDEE7459.pdf`.

[34] Brian D. Carrier and Eugene H. Spafford. An Event-Based Digital Forensic Investigation Framework. In *Digital Forensic Research Workshop (DFRWS)*, pages 1–12, 08 2004.

[35] Andrew Case and Golden G. Richard III. Memory forensics: The path forward. *Digit. Investig.*, 20:23–33, 2017. doi: 10.1016/j.diin.2016.12.004.

[36] Eoghan Casey. *Digital Evidence and Computer Crime*. Academic Press, 2 edition, 2004. ISBN 978-0-121-63104-8.

[37] Eoghan Casey. Introduction. In Eoghan Casey, editor, *Handbook of Digital Forensics and Investigation*, pages 1–17. Elsevier, 2010.

[38] Eoghan Casey. *Digital Evidence and Computer Crime: Forensic Science, Computers, and the Internet*. Academic Press, 3 edition, 2011.

[39] Eoghan Casey. Triage in digital forensics. *Digit. Investig.*, 10(2):85–86, 2013. doi: 10.1016/j.diin.2013.08.001.

[40] Eoghan Casey and Curtis W Rose. Forensic Analysis. In Eoghan Casey, editor, *Handbook of Digital Forensics and Investigation*, pages 21–62. Elsevier, 2010.

[41] Eoghan Casey, Sean Barnum, Ryan Griffith, Jonathan Snyder, Harm M. A. van Beek, and Alex J. Nelson. Advancing coordinated cyber-investigations and tool interoperability using a community developed specification language. *Digit. Investig.*, 22:14–45, 2017. doi: 10.1016/j.diin.2017.08.002.

[42] K. Mani Chandy and Jayadev Misra. *Parallel program design - a foundation*. Addison-Wesley, 1989. ISBN 978-0-201-05866-6.

[43] Hsinchun Chen, Homa Atabakhsh, Chunju Tseng, Byron Marshall, Siddharth Kaza, Shauna Eggers, Hemanth Gowda, Ankit Shah, Tim Petersen, and Chuck Violette. Visualization in law enforcement. In Lois M. L. Delcambre and Genevieve Giuliano, editors, *Proceedings of the 2005 National Conference on Digital Government Research, DG.O 2005, Atlanta, Georgia, USA, May 15-18, 2005*, volume 89 of *ACM International Conference Proceeding Series*, pages 229–230. Digital Government Research Center, 2005.

[44] W. Jerry Chisum and Brent E. Turvey. Evidence dynamics: Locard's exchange principle & crime reconstruction. *Journal of Behavioral Profiling*, 1(1):1–15, 2000.

[45] W. Jerry Chisum and Brent E. Turvey. *Crime Reconstruction*. Academic Press, San Diego, CA, 2 edition, 6 2011.

[46] Chia Yuan Cho, Domagoj Babić, Eui Chul Richard Shin, and Dawn Song. Inference and Analysis of Formal Models of Botnet Command and Control Protocols. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, CCS '10, page 426–439, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781450302456. doi: 10.1145/1866307.1866355.

[47] Mehzeb Chowdhury. A broken system? Examining the perilous state of quality assurance in crime scene practice. *Science & Justice*, 61(5):564–572, 2021. ISSN 1355-0306. doi: 10.1016/j.scijus.2021.07.001.

[48] Harold Chun and Norman Barbosa. Ochko123 - How the Feds Caught Russian Mega-Carder Roman Seleznev, 2017. URL `https://www.youtube.com/watch?v=6Chp12sEnWk`.

[49] Alessandro Cimatti, Edmund M. Clarke, Enrico Giunchiglia, Fausto Giunchiglia, Marco Pistore, Marco Roveri, Roberto Sebastiani, and Armando Tacchella. NuSMV 2: An OpenSource Tool for Symbolic Model Checking. In Ed Brinksma and Kim Guldstrand Larsen, editors, *Computer Aided Verification, 14th International Conference, CAV 2002,Copenhagen, Denmark, July 27-31, 2002, Proceedings*, volume 2404 of *Lecture Notes in Computer Science*, pages 359–364. Springer, 2002. doi: 10.1007/3-540-45657-0_29.

[50] R. Cook, I.W. Evett, G. Jackson, P.J. Jones, and J.A. Lambert. A model for case assessment and interpretation. *Science & Justice*, 38(3):151–156, 1998. ISSN 1355-0306. doi: 10.1016/S1355-0306(98)72099-4.

[51] R. Cook, I.W. Evett, G. Jackson, P.J. Jones, and J.A. Lambert. A hierarchy of propositions: deciding which level to address in casework. *Science & Justice*, 38 (4):231–239, 1998. ISSN 1355-0306. doi: 10.1016/S1355-0306(98)72117-3. URL `https://www.sciencedirect.com/science/article/pii/S1355030698721173`.

[52] Antony K Cooper and Peter MU Schmitz. Mapping crime scenes and cellular telephone usage in South Africa. *Crime Mapping News*, 3(1):4–6, 2001.

[53] Donald R Cooper, Pamela S Schindler, and Jianmin Sun. *Business Research Methods*, volume 12. McGraw-Hill Education, 2014.

[54] Hannelore Crolly. Spur 4334 und eine Sternstunde der Kriminalistik. *Welt*, 06 2017. URL `https://www.welt.de/vermischtes/article165231684/Spur-4334-und-eine-Sternstunde-der-Kriminalistik.html`.

[55] Peter R. De Forest. Recapturing the essence of criminalistics. *Science & justice : journal of the Forensic Science Society*, 39(3):196—208, 1999. ISSN 1355-0306. doi: 10.1016/s1355-0306(99)72047-2.

[56] Peter R. De Forest, Robert E. Gaensslen, and Henry C. Lee. *Forensic science: an introduction to criminalistics*. McGraw-Hill, New York, 1983.

[57] Daan de Graaf, Ahmed F. Shosha, and Pavel Gladyshev. BREDOLAB: Shopping in the Cybercrime Underworld. In Marcus K. Rogers and Kathryn C. Seigfried-Spellar, editors, *Digital Forensics and Cyber Crime - 4th International Conference, ICDF2C 2012, Lafayette, IN, USA, October 25-26, 2012, Revised Selected Papers*, volume 114 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 302–313. Springer, 2012. doi: 10.1007/978-3-642-39891-9_19.

[58] Hinrich de Vries. Ist die Kriminalistik eine Wissenschaft? *Kriminalistik*, 62:213–217, 4 2018. ISSN 0023-4699.

[59] Debian. *Manual systemd-tmpfiles*. Debian, 2022. URL `https://dyn.manpages.debian.org/testing/systemd/systemd-tmpfiles.8.en.html`.

[60] Waldo Delport and Martin S. Olivier. Isolating Instances in Cloud Forensics. In Gilbert L. Peterson and Sujeet Shenoi, editors, *Advances in Digital Forensics VIII - 8th IFIP WG 11.9 International Conference on Digital Forensics, Pretoria, South Africa, January 3-5, 2012, Revised Selected Papers*, volume 383 of *IFIP Advances in Information and Communication Technology*, pages 187–200. Springer, 2012. doi: 10.1007/978-3-642-33962-2_13.

[61] Department of Justice. Russian Cyber-Criminal Convicted of 38 Counts Related to Hacking Businesses and Stealing More Than Two Million Credit Card Numbers, 08 2016. URL `https://www.justice.gov/opa/pr/russian-cyber-criminal-convicted-38-counts-related-hacking-businesses-and-stealing-more-two`.

[62] Dominic Deuber, Jan Gruber, Merlin Humml, Viktoria Ronge, and Nicole Scheler. Argumentation Schemes for Blockchain Deanonymisation. *FinTech*, 3:236–248, 2024. doi: 10.3390/fintech3020014.

[63] Andreas Dewald. *Formalisierung digitaler Spuren und ihre Einbettung in die Forensische Informatik*. PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), 2012.

[64] Andreas Dewald. Characteristic evidence, counter evidence and reconstruction problems in forensic computing. *it Inf. Technol.*, 57(6):339–346, 2015.

[65] Andreas Dewald and Felix Freiling. From Computer Forensics to Forensic Computing: Investigators Investigate, Scientists Associate. Technical Report CS-2014-04, Friedrich-Alexander-University Erlangen-Nuremberg (FAU), 5 2014. URL `https://opus4.kobv.de/opus4-fau/files/4750/computer_forensics_is_not_forensic_science.pdf`.

[66] Oxford English Dictionary. evidence, n. OED Online. URL `https://www.oed.com/oed2/00079136`.

[67] Die Welt. Phantommörderin gibt es nachweislich nicht. *Die Welt*, 03 2009. URL `https://www.welt.de/vermischtes/article3457025/Phantommoerderin-gibt-es-nachweislich-nicht.html`.

[68] Edsger W. Dijkstra. Guarded Commands, Nondeterminacy and Formal Derivation of Programs. *Commun. ACM*, 18(8):453–457, 1975. doi: 10.1145/360933.360975.

[69] Carsten Dominik. *The Org Mode 9.1 Reference Manual*. 12th Media Services, 2018.

[70] Harper Douglas. hypothesis (n.). Online Etymology Dictionary. URL `https://www.etymonline.com/word/hypothesis#etymonline_v_16140`.

[71] Valentina Dragos. Developing a core ontology to improve military intelligence analysis. *Int. J. Knowl. Based Intell. Eng. Syst.*, 17(1):29–36, 2013. doi: 10.3233/KES-130253.

[72] United Nations Office On Drugs and Crime. *Criminal Intelligence: Manual for Analysts*. UNODC Criminal Intelligence Manual. United Nations Office on Drugs and Crime (UNODC), Austria, 4 2011.

[73] DSTL. Hue, Smartwatches and Nintendo Switch. Technical Report Digital Forensics Bulletin, Edition 13, DIIS: DSTL/PUB125049, Defence Science and Technology Laboratory, 8 2020. URL `https://web.archive.org/web/20230119073315/https://us5.campaign-archive.com/?u=a5a2a1131e612711f02b96e2c&id=34cb5884cb`.

[74] Aric Dutelle. *An introduction to crime scene investigation*. Jones & Bartlett Learning, Sudbury, Mass, 2014. ISBN 978-1449645427.

[75] John Ellson, Emden Gansner, Lefteris Koutsofios, Stephen C North, and Gordon Woodhull. Graphviz – open source graph drawing tools. In *International Symposium on Graph Drawing*, pages 483–484. Springer, 2001.

[76] Satu Elo and Helvi Kyngäs. The Qualitative Content Analysis Process. *Journal of Advanced Nursing*, 62(1):107–115, 2008.

[77] Europol. World's most dangerous malware EMOTET disrupted through global action, 1 2021. URL `https://www.europol.europa.eu/newsroom/news/world%E2%80%99s-most-dangerous-malware-emotet-disrupted-through-global-action`.

[78] Jason Farina, Mark Scanlon, and M. Tahar Kechadi. BitTorrent Sync: First Impressions and Digital Forensic Implications. *Digit. Investig.*, 11(Supplement 1):S77–S86, 2014. doi: 10.1016/j.diin.2014.03.010.

[79] Fabian Faust, Aurélien Thierry, Tilo Müller, and Felix C. Freiling. Selective Imaging of File System Data on Live Systems. *Digit. Investig.*, 36 Supplement:301115, 2021. doi: 10.1016/j.fsidi.2021.301115.

[80] Federal Criminal Police Office Germany. Bundeslagebild Cybercrime 2021, 2022. URL `https://www.bka.de/DE/AktuelleInformationen/StatistikenLagebilder/Lagebilder/Cybercrime/cybercrime_node.html`.

[81] Jens Finke and Olivier Sessink. Thumbnail Managing Standard. Technical report, freedesktop.org, 2001. URL `https://specifications.freedesktop.org/thumbnail-spec/thumbnail-spec-latest.html`.

[82] Ane Elida Fonneløp, Helen Johannessen, Thore Egeland, and Peter Gill. Contamination during criminal investigation: Detecting police contamination and secondary DNA transfer from evidence bags. *Forensic Science International: Genetics*, 23:121–129, 7 2016. ISSN 1872-4973. doi: 10.1016/j.fsigen.2016.04.003.

[83] Frank Fransen and Richard Kerkdijk. Cyber threat intelligence sharing through national and sector-oriented communities. In Florian Skopik, editor, *Collaborative Cyber Threat Intelligence: Detecting and Responding to Advanced Cyber Attacks at the National Level*, pages 187–224. CRC Press, 2017.

[84] Felix Freiling and Konstantin Sack. Zur Authentizität und Integrität bei (digitalen) Beweismitteln. In Klaus Vieweg, editor, *Festgabe Institut für Recht und Technik: Erlanger Festveranstaltungen 2011 und 2016*, volume 111 of *Recht - Technik - Wirtschaft*, pages 319–337. Carl Heymanns Verlag, Köln, 2017. ISBN 978-3-452-28832-5. URL `https://faui1-files.cs.fau.de/public/publications/22-Freiling-Sack_RTW-111.pdf`.

[85] Felix C. Freiling and Michael Gruhn. What is Essential Data in Digital Forensic Analysis? In Jana Dittmann and Holger Morgenstern, editors, *Ninth International Conference on IT Security Incident Management & IT Forensics, IMF 2015, Magdeburg, Germany, May 18-20, 2015*, pages 40–48. IEEE Computer Society, 2015. doi: 10.1109/IMF.2015.20.

[86] Felix C. Freiling and Leonhard Hösch. Controlled experiments in digital evidence tampering. *Digit. Investig.*, 24 Supplement:S83–S92, 2018. doi: 10.1016/j.diin.2018.01.011.

[87] Felix C. Freiling, Thomas Glanzmann, and Hans P. Reiser. Characterizing loss of digital evidence due to abstraction layers. *Digit. Investig.*, 20 Supplement:S107–S115, 2017. doi: 10.1016/j.diin.2017.01.012.

[88] Aya Fukami, Radina Stoykova, and Zeno J. M. H. Geradts. A new model for forensic data extraction from encrypted mobile devices. *Digit. Investig.*, 38:301169, 2021. doi: 10.1016/j.fsidi.2021.301169.

[89] Carlos Hernandez Gañán, Michael Ciere, and Michel van Eeten. Beyond the pretty penny: the Economic Impact of Cybercrime. In *Proceedings of the 2017 New Security Paradigms Workshop, NSPW 2017, Santa Cruz, CA, USA, October 01-04, 2017*, pages 35–45. ACM, 2017. doi: 10.1145/3171533.3171535.

[90] Rod Gehl and Darryl Plecas. *Introduction to criminal investigation: processes, practices and thinking*. Justice Institute of British Columbia, 2017.

[91] Peter Gill, Tacha Hicks, John M. Butler, Ed Connolly, Leonor Gusmão, Bas Kokshoorn, Niels Morling, Roland A.H. van Oorschot, Walther Parson, Mechthild Prinz, Peter M. Schneider, Titia Sijen, and Duncan Taylor. DNA commission of the International society for forensic genetics: Assessing the value of forensic biological evidence - Guidelines highlighting the importance of propositions. Part II: Evaluation of biological traces considering activity level propositions. *Forensic Science International: Genetics*, 44:102186, 2020. ISSN 1872-4973. doi: 10.1016/j.fsigen.2019.102186.

[92] Simone Gittelson, Charles E.H. Berger, Graham Jackson, Ian W. Evett, Christophe Champod, Bernard Robertson, James M. Curran, Duncan Taylor, Bruce S. Weir, Michael D. Coble, and John S. Buckleton. A response to "Likelihood ratio as weight of evidence: A closer look" by Lund and Iyer. *Forensic Science International*, 288: e15–e19, 2018. ISSN 0379-0738. doi: 10.1016/j.forsciint.2018.05.025.

[93] Pavel Gladyshev. *Formalising event reconstruction in digital investigations*. PhD thesis, University College Dublin, 2004.

[94] Pavel Gladyshev. Finite State Machine Analysis of a Blackmail Investigation. *Int. J. Digit. EVid.*, 4(1), 2005. URL `http://www.utica.edu/academic/institutes/ecii/publications/articles/B4A163F6-0AE6-1C11-14CBB310F7E43F08.pdf`.

[95] Pavel Gladyshev and Ahmed Patel. Finite state machine approach to digital event reconstruction. *Digit. Investig.*, 1(2):130–149, 2004. doi: 10.1016/j.diin.2004.03.001.

[96] Roland Grassberger. Pioneers in Criminology XIII–Hans Gross (1847-1915). *J. Crim. L. Criminology & Police Sci.*, 47(4):397–405, 1956.

[97] Ivan Griffin and Ita Richardson. Using LATEX for qualitative data analysis. *The PracTEX Journal*, 1, 2010.

[98] Hans Gross and Friedrich Geerds. *Handbuch der Kriminalistik: Wissenschaft und Praxis des Verbrechensbekämpfung*. M. Pawlak, 1977.

[99] Jan Gruber. Identifizierung von Malware-Infrastruktur mittels verteilter Spamtrap-Systeme. In Albrecht Ude, editor, *Sicherheit in vernetzten Systemen: 30. DFN-Konferenz*, pages A1–A27. BoD–Books on Demand, Hamburg, 02 2023. ISBN 3756881393.

[100] Jan Gruber and Felix Freiling. Fighting Evasive Malware. *Datenschutz und Datensicherheit - DuD*, 46(5):284–290, 5 2022. doi: 10.1007/s11623-022-1604-9.

[101] Jan Gruber and Merlin Humml. A Formal Treatment of Expressiveness and Relevance of Digital Evidence. *Digital Threats*, 7 2023. ISSN 2692-1626. doi: 10.1145/3608485.

[102] Jan Gruber, Dominik Brodowski, and Felix C. Freiling. Die polizeiliche Aufgabe und Pflicht zur digitalen Gefahrenabwehr. *Zeitschrift für das gesamte Sicherheitsrecht (GSZ)*, 5:171–176, 2022. ISSN 2567-3823.

[103] Jan Gruber, Lena L. Voigt, Zinaida Benenson, and Felix C. Freiling. Foundations of cybercriminalistics: From general process models to case-specific concretizations in cybercrime investigations. *Forensic Sci. Int. Digit. Investig.*, 43(Supplement):301438, 2022. ISSN 2666-2817. doi: 10.1016/J.FSIDI.2022.301438.

[104] Jan Gruber, Christopher J. Hargreaves, and Felix C. Freiling. Contamination of digital evidence: Understanding an underexposed risk. *Forensic Sci. Int. Digit. Investig.*, 44 (Supplement):301501, 2023. ISSN 2666-2817. doi: 10.1016/j.fsidi.2023.301501.

[105] Jan Gruber, Merlin Humml, Lutz Schröder, and Felix C. Freiling. Formal Verification of Necessary and Sufficient Evidence in Forensic Event Reconstruction. In Edita Bajramovic and Ricardo J. Rodríguez, editors, *Proceedings of the Digital Forensics Research Conference Europe* (*DFRWS EU*), pages 1–11, Bonn, 3 2023. dfrws.org.

[106] Jan Gruber, Lena L. Voigt, and Felix C. Freiling. Faktoren erfolgreicher Cybercrime-Ermittlungen. *Kriminalistik*, 77:266–271, 5 2023. ISSN 0023-4699.

[107] Jan Grübler. 'Kontamination'. In Ingo Wirth, editor, *Kriminalistik-Lexikon*, pages 362–363. C.F. Müller GmbH, Heidelberg, 5 edition, 2021.

[108] Mayank R. Gupta, Michael D. Hoeschele, and Marcus K. Rogers. Hidden Disk Areas: HPA and DCO. *Int. J. Digit. EVid.*, 5(1), 2006. URL `http://www.utica.edu/academic/institutes/ecii/publications/articles/EFE36584-D13F-2962-67BEB146864A2671.pdf`.

[109] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52 (5):91–98, 2009. doi: 10.1145/1506409.1506429.

[110] Allen Hall. 'Woman Without A Face' leaves German police in the dark. *The Age*, 11 2008. URL `https://web.archive.org/web/20090211004359/http://www.theage.com.au/world/woman-without-a-face-leaves-german-police-in-the-dark-20081116-683p.html?page=-1`.

[111] Christopher James Hargreaves. *Assessing the reliability of digital evidence from live investigations involving encryption.* PhD thesis, Cranfield University, UK, 2009.

[112] Christopher James Hargreaves and Howard Chivers. Recovery of Encryption Keys from Memory Using a Linear Scan. In *Proceedings of the The Third International Conference on Availability, Reliability and Security, ARES 2008, March 4-7, 2008, Technical University of Catalonia, Barcelona , Spain*, pages 1369–1376. IEEE Computer Society, 2008. doi: 10.1109/ARES.2008.109.

[113] Ryan Harris. Arriving at an anti-forensics consensus: Examining how to define and control the anti-forensics problem. *Digit. Investig.*, 3(Supplement):44–49, 2006. doi: 10.1016/j.diin.2006.06.005.

[114] Durdica Hazard. The Relevant Physical Trace in Criminal Investigation. *Journal of Forensic Science and Medicine*, 2(4):208–212, 10 2016. doi: 10.4103/2349-5014.164662.

[115] Durdica Hazard. Discussion Between Forensic and Evidence Law Practitioners About the Relevancy Concept. In *Archibald Reiss Days 2017, Belgrade, Serbia, November 7–9, 2017*, pages 3–10. Academy of Criminalistic and Police Studies Belgrade, 2017.

[116] Himmelreich, Claudia. Germany's Phantom Serial Killer: A DNA Blunder. *Time Magazine Online*, 03 2009. URL `https://content.time.com/time/world/article/0,8599,1888126,00.html`.

[117] Robert R. Hoffman, Nigel R. Shadbolt, A. Mike Burton, and Gary Klein. Eliciting Knowledge From Experts: A Methodological Analysis. *Organizational Behavior and Human Decision Processes*, 62(2):129–158, 1995.

[118] Thomas J. Holt. *Subcultural Theories of Crime*, pages 513–526. Springer International Publishing, Cham, 2020. ISBN 978-3-319-78440-3. doi: 10.1007/978-3-319-78440-3_19.

[119] Edith Huber. *Cybercrime: Eine Einführung*. Springer, 2019.

[120] Ewa Huebner, Derek Bem, Frans A. Henskens, and Mark Wallis. Persistent systems techniques in forensic acquisition of memory. *Digit. Investig.*, 4(3-4):129–137, 2007. doi: 10.1016/j.diin.2008.02.001.

[121] Paul Hunton. The growing phenomenon of crime and the internet: A cybercrime execution and analysis model. *Comput. Law Secur. Rev.*, 25(6):528–535, 2009. doi: 10.1016/j.clsr.2009.09.005.

[122] Paul Hunton. Cyber Crime and Security: A New Model of Law Enforcement Investigation. *Policing: a journal of policy and practice*, 4(4):385–395, 8 2010. ISSN 1752-4520. doi: 10.1093/police/paq038.

[123] Paul Hunton. The stages of cybercrime investigations: Bridging the gap between technology examination and law enforcement investigation. *Comput. Law Secur. Rev.*, 27(1):61–67, 2011. doi: 10.1016/j.clsr.2010.11.001.

[124] Michael Huth and Mark Dermot Ryan. *Logic in computer science - modelling and reasoning about systems (2. ed.)*. Cambridge University Press, 2004.

[125] Keith Inman and Norah Rudin. *Principles and Practice of Criminalistics: The Profession of Forensic Science*. CRC Press, 8 2000. ISBN 9780429247620. doi: 10.1201/9781420036930.

[126] Keith Inman and Norah Rudin. The origin of evidence. *Forensic Science International*, 126(1):11–16, 2002. ISSN 0379-0738. doi: 10.1016/S0379-0738(02)00031-2.

[127] ISO/IEC JTC 1/SC 27. Guidelines for identification, collection, acquisition and preservation of digital evidence. Standard, International Organization for Standardization, Geneva, CH, 10 2012.

[128] ISO/TC 272. Minimizing the risk of human DNA contamination in products used to collect, store and analyze biological material for forensic purposes – Requirements. Standard, International Organization for Standardization, Geneva, CH, 2 2016.

[129] ISO/TC 272. Forensic sciences – Part 1: Terms and definitions. Standard, International Organization for Standardization, Geneva, CH, 8 2018.

[130] Joshua James, Pavel Gladyshev, Mohd Taufik Abdullah, and Yuandong Zhu. Analysis of Evidence Using Formal Event Reconstruction. In Sanjay Goel, editor, *Digital Forensics and Cyber Crime - First International ICST Conference, ICDF2C 2009, Albany, NY, USA, September 30-October 2, 2009, Revised Selected Papers*, volume 31 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 85–98. Springer, 2009. doi: 10.1007/978-3-642-11534-9_9.

[131] David-Olivier Jaquet-Chiffelle and Eoghan Casey. A formalized model of the Trace. *Forensic Science International*, 327:110941, 2021. ISSN 0379-0738. doi: 10.1016/j.forsciint.2021.110941.

[132] Mark A. Jobling and Peter Gill. Encoded evidence: DNA in forensic analysis. *Nature Reviews Genetics*, 5(10):739–751, 2004.

[133] Christopher S. Johnson, Mark Lee Badger, David A. Waltermire, Julie Snyder, and Clem Skorupka. Guide to Cyber Threat Information Sharing. Technical Report NIST Special Publication 800-150, National Institute of Standards and Technology, 10 2016.

[134] Sven Kälber, Andreas Dewald, and Felix C. Freiling. Forensic Application-Fingerprinting Based on File System Metadata. In Holger Morgenstern, Ralf Ehlert, Felix C. Freiling, Sandra Frings, Oliver Göbel, Detlef Günther, Stefan Kiltz, Jens Nedon, and Dirk Schadt, editors, *Seventh International Conference on IT Security Incident Management and IT Forensics, IMF 2013, Nuremberg, Germany, March 12-14, 2013*, pages 98–112. IEEE Computer Society, 2013. doi: 10.1109/IMF.2013.20.

[135] Sven Kälber, Andreas Dewald, and Steffen Idler. Forensic Zero-Knowledge Event Reconstruction on Filesystem Metadata. In Stefan Katzenbeisser, Volkmar Lotz, and Edgar R. Weippl, editors, *Sicherheit 2014: Sicherheit, Schutz und Zuverlässigkeit, Beiträge der 7. Jahrestagung des Fachbereichs Sicherheit der Gesellschaft für Informatik e.V. (GI), 19.-21. März 2014, Wien, Österreich*, volume P-228 of *LNI*, pages 331–343. GI, 2014.

[136] Vladimir Katalov. iOS Device Acquisition with checkra1n Jailbreak. Technical report, ElcomSoft Co.Ltd., 2019. URL `https://web.archive.org/web/20220914142952/https://blog.elcomsoft.com/2019/11/ios-device-acquisition-with-checkra1n-jailbreak/`.

[137] Christoph Keller. Einführung in die Kriminalistik. In Christoph Keller, editor, *Basislehrbuch Kriminalistik: Strategien und Techniken der Verbrechensaufklärung und -bekämpfung*, pages 57–67,. Verlag Deutsche Polizeiliteratur, Hilden, 11 2019.

[138] Paul L. Kirk. Journal of Criminal Law and Criminology. *Journal of Criminal Law and Criminology*, 38:165, 1947.

[139] Paul L. Kirk. Criminalistics. *Science*, 140(3565):367–370, 1963.

[140] Paul L. Kirk. The ontogeny of criminalistics. *J. Crim. L. Criminology & Police Sci.*, 54: 235, 1963.

[141] Paul L. Kirk. *Crime investigation*. John Wiley & Sons, Nashville, TN, 2 edition, 7 1974.

[142] Janet L. Kolodner. An introduction to case-based reasoning. *Artificial intelligence review*, 6(1):3–34, 1992.

[143] Jihène Krichène, Mohamed Hamdi, and Noureddine Boudriga. Collective computer incident response using cognitive maps. In *Proceedings of the IEEE International Conference on Systems, Man & Cybernetics: The Hague, Netherlands, 10-13 October 2004*, pages 1080–1085. IEEE, 2004. doi: 10.1109/ICSMC.2004.1398448.

[144] Michael Y. K. Kwan, Kam-Pui Chow, Frank Y. W. Law, and Pierre K. Y. Lai. Reasoning About Evidence Using Bayesian Networks. In Indrajit Ray and Sujeet Shenoi, editors, *Advances in Digital Forensics IV, Fourth Annual IFIP WG 11.9 Conference on Digital Forensics, Kyoto University, Kyoto, Japan, January 28-30, 2008*, volume 285 of *IFIP*, pages 275–289. Springer, 2008. doi: 10.1007/978-0-387-84927-0_22.

[145] Leslie Lamport. The Temporal Logic of Actions. *ACM Trans. Program. Lang. Syst.*, 16 (3):872–923, 1994. doi: 10.1145/177492.177726.

[146] Karl Landsteiner. Forensic application of serologic individuality tests. *Journal of the American Medical Association*, 103(14):1041–1044, 1934.

[147] Tobias Latzo and Felix C. Freiling. Characterizing the Limitations of Forensic Event Reconstruction Based on Log Files. In *18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications / 13th IEEE International Conference On Big Data Science And Engineering, TrustCom/BigDataSE 2019, Rotorua, New Zealand, August 5-8, 2019*, pages 466–475. IEEE, 2019. doi: 10.1109/TrustCom/BigDataSE.2019.00069.

[148] Henry C. Lee, Timothy Palmbach, and Marilyn T. Miller. *Henry Lee's crime scene handbook*. Academic Press, 2001.

[149] Ryan Leigland and Axel W. Krings. A Formalization of Digital Forensics. *Int. J. Digit. Evid.*, 3(2), 2004. URL `http://www.utica.edu/academic/institutes/ecii/publications/articles/A0B8472C-D1D2-8F98-8F7597844CF74DF8.pdf`.

[150] Nena Lim and Anne Khoo. Forensics of computers and handheld devices: identical or fraternal twins? *Commun. ACM*, 52(6):132–135, 2009. doi: 10.1145/1516046.1516080.

[151] Yao Liu, Ming Xu, Jian Xu, Ning Zheng, and Xiaodong Lin. SQLite Forensic Analysis Based on WAL. In Robert H. Deng, Jian Weng, Kui Ren, and Vinod Yegneswaran, editors, *Security and Privacy in Communication Networks - 12th International Conference, SecureComm 2016, Guangzhou, China, October 10-12, 2016, Proceedings*, volume 198 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 557–574. Springer, 2016. doi: 10.1007/978-3-319-59608-2_31.

[152] Edmund Locard. *L'enquête criminelle et les méthódes scientifiques*. Bibliothèque de philosophie scientifique. E. Flammarion, 1920.

[153] Jonathan Lusthaus. *Industry of anonymity*. Harvard University Press, 2018.

[154] James R. Lyle. A strategy for testing hardware write block devices. *Digit. Investig.*, 3 (Supplement):3–9, 2006. doi: 10.1016/j.diin.2006.06.001.

[155] Gabriele Margiotta, Giorgia Tasselli, Federica Tommolini, Massimo Lancia, Susanna Massetti, and Eugenia Carnevali. Risk of DNA transfer by gloves in forensic casework. *Forensic Science International: Genetics Supplement Series*, 5:e527–e529, 12 2015. ISSN 1875-1768. doi: 10.1016/j.fsigss.2015.09.208.

[156] Pierre Margot. Forensic science on trial-What is the law of the land? *Australian Journal of Forensic Sciences*, 43(2-3):89–103, 2011.

[157] Pierre Margot. Traçologie: la trace, vecteur fondamental de la police scientifique. *Revue internationale de criminologie et de police technique et scientifique*, 67(1):72–97, 2014.

[158] Pierre Margot. Traceology, the bedrock of forensic science and its associated semantics. In Quentin Rossy, David Décary-Hétu, Olivier Delémont, and Massimiliano Mulone, editors, *The Routledge international handbook of forensic intelligence and criminology*, pages 30–39. Routledge, 2017.

[159] Moxie Marlinspike. Exploiting vulnerabilities in Cellebrite UFED and Physical Analyzer from an app's perspective. Technical report, Signal, 4 2021. URL `https://web.archive.org/web/20220909054459/https://signal.org/blog/cellebrite-vulnerabilities/`.

[160] Rodney McKemmish. *What is forensic computing?* Australian Institute of Criminology Canberra, 1999.

[161] Rodney McKemmish. When is Digital Evidence Forensically Sound? In Indrajit Ray and Sujeet Shenoi, editors, *Advances in Digital Forensics IV, Fourth Annual IFIP WG 11.9 Conference on Digital Forensics, Kyoto University, Kyoto, Japan, January 28-30, 2008*, volume 285 of *IFIP*, pages 3–15. Springer, 2008. doi: 10.1007/978-0-387-84927-0_1.

[162] Georgina Meakin and Allan Jamieson. DNA transfer: Review and implications for casework. *Forensic Science International: Genetics*, 7(4):434–443, 7 2013. ISSN 1872-4973. doi: 10.1016/j.fsigen.2013.03.013.

[163] Ronny Merkel, Claus Vielhauer, Jana Dittmann, Robert Fischer, Mario Hildebrandt, and Christian Arndt. From classical forensics to digitized crime scene analysis. In *2015 IEEE International Conference on Multimedia and Expo, ICME 2015, Turin, Italy, June 29 - July 3, 2015*, pages 1–6. IEEE Computer Society, 2015. doi: 10.1109/ICME.2015.7397313.

[164] Michigan Legal Publishing Ltd. *Federal rules of evidence; 2023 edition*. Michigan Legal Publishing, 2023 edition, 11 2022.

[165] Clifford Miller. Electronic evidence – can you prove the transaction took place. *Computer Lawyer*, 9(5):21–33, 1992. Contains a seminal definition of digital evidence.

[166] Virginie Redouté Minzière, Anne-Laure Gassner, Matteo Gallidabino, Claude Roux, and Céline Weyermann. The relevance of gunshot residues in forensic science. *WIREs Forensic Science*, 5(1):467–470, 8 2022. doi: 10.1002/wfs2.1472.

[167] Steven Morris. Rape accused was victim of forensics error, regulator finds. *The Guardian*, 10 2012. URL `https://www.theguardian.com/world/2017/mar/12/netherlands-will-pay-the-price-for-blocking-turkish-visit-erdogan`.

[168] Patrick Mullan, Christian Riess, and Felix C. Freiling. Forensic source identification using JPEG image headers: The case of smartphones. *Digit. Investig.*, 28 Supplement: S68–S76, 2019. doi: 10.1016/j.diin.2019.01.016.

[169] Wilmuth Müller, Dirk Mühlenberg, Dirk Pallmer, Uwe Zeltmann, Christian Ellmauer, and Konstantinos P. Demestichas. Knowledge Engineering and Ontology for Crime Investigation. In Ilias Maglogiannis, Lazaros Iliadis, John MacIntyre, and Paulo Cortez, editors, *Artificial Intelligence Applications and Innovations - 18th IFIP WG 12.5 International Conference, AIAI 2022, Hersonissos, Crete, Greece, June 17-20, 2022, Proceedings, Part I*, volume 646 of *IFIP Advances in Information and Communication Technology*, pages 483–494. Springer, 2022. doi: 10.1007/978-3-031-08333-4_39.

[170] Jennifer C. Noble. *White-Collar and Financial Crimes: A Casebook of Fraudsters, Scam Artists, and Corporate Thieves*. University of California Press, 2021.

[171] Joseph D. Novak and Alberto J. Cañas. The Origins of the Concept Mapping Tool and the Continuing Evolution of the Tool. *Information visualization*, 5(3):175–184, 2006. ISSN 1473-8724. doi: 10.1057/palgrave.ivs.9500126.

[172] Richard E. Overill and Kam-Pui Chow. Measuring Evidential Weight in Digital Forensic Investigations. In Gilbert L. Peterson and Sujeet Shenoi, editors, *Advances in Digital Forensics XIV - 14th IFIP WG 11.9 International Conference, New Delhi, India, January 3-5, 2018, Revised Selected Papers*, volume 532 of *IFIP Advances in Information and Communication Technology*, pages 3–10. Springer, 2018. doi: 10.1007/978-3-319-99277-8_1.

[173] Richard E. Overill and Jantje A. M. Silomon. Uncertainty Bounds for Digital Forensic Evidence and Hypotheses. In *Seventh International Conference on Availability, Reliability and Security, Prague, ARES 2012, Czech Republic, August 20-24, 2012*, pages 590–595. IEEE Computer Society, 2012. doi: 10.1109/ARES.2012.17.

[174] Richard E. Overill and Jantje AM Silomon. Digital meta-forensics: quantifying the investigation. In *Proc. 4th International Conference on Cybercrime Forensics Education & Training (CFET 2010), Canterbury, UK (September 2010)*, 2010.

[175] Richard E. Overill, Jantje A. M. Silomon, and Kam-Pui Chow. A Complexity Based Model for Quantifying Forensic Evidential Probabilities. In *ARES 2010, Fifth International Conference on Availability, Reliability and Security, 15-18 February 2010, Krakow, Poland*, pages 671–676. IEEE Computer Society, 2010. doi: 10.1109/ARES.2010.42.

[176] Richard E. Overill, Jantje A. M. Silomon, Kam-Pui Chow, and Hayson Tse. Quantification of digital forensic hypotheses using probability theory. In *Eighth International Workshop on Systematic Approaches to Digital Forensic Engineering, SADFE 2013, Hong Kong, China, November 21-22, 2013*, pages 1–5. IEEE, 2013. doi: 10.1109/SADFE.2013.6911547.

[177] Uygar Özesmi and Stacy L. Özesmi. Ecological models based on people's knowledge: a multi-step fuzzy cognitive mapping approach. *Ecological modelling*, 176(1-2):43–64, 2004. doi: 10.1016/j.ecolmodel.2003.10.027.

[178] Fabio Pagani, Oleksii Fedorov, and Davide Balzarotti. Introducing the Temporal Dimension to Memory Forensics. *ACM Trans. Priv. Secur.*, 22(2):9:1–9:21, 2019. doi: 10.1145/3310355.

[179] Helen Page, Graeme Horsman, Anna Sarna, and Julienne Foster. A review of quality procedures in the UK forensic sciences: What can the field of digital forensics learn? *Science & Justice*, 59(1):83–92, 2019. ISSN 1355-0306. doi: 10.1016/j.scijus.2018.09.006.

[180] Gary Palmer. A road map for digital forensic research. Technical Report DTR-T001-01 Final, Air Force Research Laboratory, Rome, New York, 2001.

[181] Mark Phythian. *Understanding the intelligence cycle*. Routledge, 2013.

[182] Ines Pickrahn, Gabriele Kreindl, Eva Müller, Bettina Dunkelmann, Waltraud Zahrer, Jan Cemper-Kiesslich, and Franz Neuhuber. Contamination when collecting trace evidence—An issue more relevant than ever? *Forensic Science International: Genetics Supplement Series*, 5:e603–e604, 12 2015. ISSN 1875-1768. doi: 10.1016/j.fsigss.2015.09.238.

[183] Ines Pickrahn, Gabriele Kreindl, Eva Müller, Bettina Dunkelmann, Waltraud Zahrer, Jan Cemper-Kiesslich, and Franz Neuhuber. Contamination incidents in the pre-analytical phase of forensic DNA analysis in Austria—Statistics of 17 years. *Forensic Science International: Genetics*, 31:12–18, 11 2017. ISSN 1872-4973. doi: 10.1016/j.fsigen.2017.07.012.

[184] Amir Pnueli. The Temporal Logic of Programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977. doi: 10.1109/SFCS.1977.32.

[185] Mark Pollitt. An Ad Hoc Review of Digital Forensic Models. In Ming-Yuh Huang and Deborah A. Frincke, editors, *Second International Workshop on Systematic Approaches to Digital Forensic Engineering, SADFE 2007, Seattle, Washington, USA, April 10-12, 2007*, pages 43–54. IEEE Computer Society, 2007. doi: 10.1109/SADFE.2007.3.

[186] Karl Raimund Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. London, England: Routledge, 1962.

[187] Karl Raimund Popper. *The logic of scientific discovery*. Routledge Classics (Hardcover). Routledge, London, England, 2 2002.

[188] A. Poy and R.A.H. van Oorschot. Beware; gloves and equipment used during the examination of exhibits are potential vectors for transfer of DNA-containing material. *International Congress Series*, 1288:556–558, 4 2006. ISSN 0531-5131. doi: 10.1016/j.ics.2005.09.126.

[189] Hans H. Rath and Steve Pepper. Topic maps: introduction and allegro. In *Proceedings of the Markup Technologies*, volume 99, pages 7–9, 1999.

[190] Slim Rekhis and Noureddine Boudriga. A formal logic-based language and an automated verification tool for computer forensic investigation. In Hisham Haddad, Lorie M. Liebrock, Andrea Omicini, and Roger L. Wainwright, editors, *Proceedings of the 2005 ACM Symposium on Applied Computing (SAC), Santa Fe, New Mexico, USA, March 13-17, 2005*, pages 287–291. ACM, 2005. doi: 10.1145/1066677.1066745.

[191] Slim Rekhis and Noureddine Boudriga. A Temporal Logic-Based Model for Forensic Investigation in Networked System Security. In Vladimir Gorodetsky, Igor V. Kotenko, and Victor A. Skormin, editors, *Computer Network Security, Third International Workshop on Mathematical Methods, Models, and Architectures for Computer Network Security, MMM-ACNS 2005, St. Petersburg, Russia, September 25-27, 2005, Proceedings*, volume 3685 of *Lecture Notes in Computer Science*, pages 325–338. Springer, 2005. doi: 10.1007/11560326_25.

[192] Slim Rekhis, Jihène Krichène, and Noureddine Boudriga. Cognitive-Maps Based Investigation of Digital Security Incidents. In *Third International Workshop on Systematic Approaches to Digital Forensic Engineering, SADFE 2008, Berkeley, California, USA, May 22, 2008*, pages 25–40. IEEE Computer Society, 2008. doi: 10.1109/SADFE.2008.20.

[193] Olivier Ribaux and Stefano Caneppele. Forensic intelligence. In Quentin Rossy, David Décary-Hétu, Olivier Delémont, and Massimiliano Mulone, editors, *The Routledge international handbook of forensic intelligence and criminology*, pages 136–148. Routledge, 2017.

[194] Ralph R. Ristenbatt III, Jack Hietpas, Peter R. De Forest, and Pierre A. Margot. Traceology, criminalistics, and forensic science. *Journal of Forensic Sciences*, 67(1): 28–32, 2022.

[195] Bernard Robertson, George A. Vignaux, and Charles E.H. Berger. *Interpreting evidence: evaluating forensic science in the courtroom*. John Wiley & Sons, 2 edition, 2016.

[196] Holger Roll. 'Version'. In Ingo Wirth, editor, *Kriminalistik-Lexikon*, pages 382–383. C.F. Müller GmbH, Heidelberg, 5 edition, 2021.

[197] Quentin Rossy and Carlo Morselli. The contribution of forensic science to the analysis of crime networks. In Quentin Rossy, David Décary-Hétu, Olivier Delémont, and Massimiliano Mulone, editors, *The Routledge international handbook of forensic intelligence and criminology*, pages 191–204. Routledge, 2017.

[198] Claude Roux, Rebecca Bucht, Frank Crispino, Peter De Forest, Chris Lennard, Pierre Margot, Michelle D. Miranda, Niamh NicDaeid, Olivier Ribaux, Alastair Ross, and Sheila Willis. The Sydney declaration – Revisiting the essence of forensic science through its fundamental principles. *Forensic Science International*, 332:111182, 2022. ISSN 0379-0738. doi: 10.1016/j.forsciint.2022.111182.

[199] Konstantin Sack. *Selektion in der digitalen Forensik*. PhD thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), 2017.

[200] Johnny Saldaña. *The coding manual for qualitative researchers*. sage, 2021.

[201] Janine Schneider, Julian Wolf, and Felix C. Freiling. Tampering with Digital Evidence is Hard: The Case of Main Memory Images. *Digit. Investig.*, 32 Supplement:300924, 2020. doi: 10.1016/j.fsidi.2020.300924.

[202] Eric Schulte, Dan Davison, Thomas Dye, and Carsten Dominik. A multi-language computing environment for literate programming and reproducible research. *Journal of Statistical Software*, 46:1–24, 2012.

[203] Giuliana Schwendener, Sébastien Moret, Karen Cavanagh-Steer, and Claude Roux. Can "contamination" occur in body bags? The example of background fibres in body bags used in Australia. *Forensic Science International*, 266:517–526, 9 2016. ISSN 0379-0738. doi: 10.1016/j.forsciint.2016.07.012.

[204] Siti Rahayu Selamat, Robiah Yusof, and Shahrin Sahib. Mapping process of digital forensic investigation framework. *International Journal of Computer Science and Network Security*, 8(10):163–169, 2008.

[205] Kishore Sengupta, Tarek K. Abdel-Hamid, and Luk N. Van Wassenhove. The experience trap. *Harvard Business Review*, 86(2):94–101, 2008.

[206] Michael Soiné. 'Beweis'. In Ingo Wirth, editor, *Kriminalistik-Lexikon*, pages 86–86. C.F. Müller GmbH, Heidelberg, 5 edition, 2021.

[207] Somayeh Soltani and Seyed-Amin Hosseini-Seno. A formal model for event reconstruction in digital forensic investigation. *Digit. Investig.*, 30:148–160, 2019. doi: 10.1016/j.diin.2019.07.006.

[208] Peter Sommer. Downloads, logs and captures: Evidence from cyberspace. *Journal of Financial Crime*, (5):138–151, 1997.

[209] Peter Sommer. Digital Footprints: Assessing Computer Evidence. *Criminal Law Review*, pages 61–78, 12 1998.

[210] Peter Sommer. Intrusion detection systems as evidence. *Comput. Networks*, 31(23-24): 2477–2487, 1999. doi: 10.1016/S1389-1286(99)00113-9.

[211] Christopher D. Steele and David J. Balding. Statistical Evaluation of Forensic DNA Profile Evidence. *Annual Review of Statistics and Its Application*, 1(1):361–384, 2014. doi: 10.1146/annurev-statistics-022513-115602.

[212] Timo Steffens. *Attribution of Advanced Persistent Threats - How to Identify the Actors Behind Cyber-Espionage*. Springer, 2020. ISBN 978-3-662-61312-2. doi: 10.1007/978-3-662-61313-9.

[213] Peter Stelfox. *Criminal investigation: an introduction to principles and practice*. Willan, Cullompton, 2009.

[214] Brett Stone-Gross, Thorsten Holz, Gianluca Stringhini, and Giovanni Vigna. The Underground Economy of Spam: A Botmaster's Perspective of Coordinating Large-Scale Spam Campaigns. In Christopher Kruegel, editor, *4th USENIX Workshop on Large-Scale Exploits and Emergent Threats, LEET '11, Boston, MA, USA, March 29, 2011*. USENIX Association, 2011. URL `https://www.usenix.org/conference/leet11/underground-economy-spam-botmasters-perspective-coordinating-large-scale-spam`.

[215] Johannes Stricker. *Tatortarbeit: Spurensuche und-sicherung bei verschiedenen Tat-und Einsatzorten*. Richard Boorberg Verlag, 2018.

[216] Südkurier. Festnahme im Mordfall Endingen - Lkw-Maut bringt Polizei auf die Spur. *Südkurier*, 06 2017. URL `https://www.suedkurier.de/baden-wuerttemberg/Festnahme-im-Mordfall-Endingen-Lkw-Maut-bringt-Polizei-auf-die-Spur;art417930,9278929`.

[217] Nina Sunde and Itiel E. Dror. A hierarchy of expert performance (HEP) applied to digital forensics: Reliability and biasability in digital forensics decision making. *Digit. Investig.*, 37:301175, 2021. doi: 10.1016/j.fsidi.2021.301175.

[218] Zareen Syed, Ankur Padia, Tim Finin, M. Lisa Mathews, and Anupam Joshi. UCO: A Unified Cybersecurity Ontology. In David R. Martinez, William W. Streilein, Kevin M. Carter, and Arunesh Sinha, editors, *Artificial Intelligence for Cyber Security, Papers from the 2016 AAAI Workshop, Phoenix, Arizona, USA, February 12, 2016*, volume WS-16-03 of *AAAI Technical Report*. AAAI Press, 2016. URL `http://www.aaai.org/ocs/index.php/WS/AAAIW16/paper/view/12574`.

[219] Andrew S. Tanenbaum and David Wetherall. *Computer networks, 5th Edition*. Pearson, 2011. ISBN 0132553171. URL `https://www.worldcat.org/oclc/698581231`.

[220] April Tanner and David A. Dampier. Concept Mapping for Digital Forensic Investigations. In Gilbert L. Peterson and Sujeet Shenoi, editors, *Advances in Digital Forensics V - Fifth IFIP WG 11.9 International Conference on Digital Forensics, Orlando, Florida, USA, January 26-28, 2009, Revised Selected Papers*, volume 306 of *IFIP Advances in Information and Communication Technology*, pages 291–300. Springer, 2009. doi: 10.1007/978-3-642-04155-6_22.

[221] The European Commission. Communication from the Commission on the EU Security Union Strategy, 2020.

[222] The White House. Executive Order on Improving the Nation's Cybersecurity, 5 2021.

[223] Kurt Thomas, Damon McCoy, Chris Grier, Alek Kolcz, and Vern Paxson. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and Abuse. In Samuel T. King, editor, *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, pages 195–210. USENIX Association, 2013. URL `https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/thomas`.

[224] Heli Tiirmaa-Klaar, Jan Gassen, Elmar Gerhards-Padilla, and Peter Martini. *Botnets*. Springer Briefs in Cybersecurity. Springer, 2013. ISBN 978-1-4471-5215-6. doi: 10.1007/978-1-4471-5216-3.

[225] Hayson Tse, Kam-Pui Chow, and Michael Y. K. Kwan. Reasoning about Evidence using Bayesian Networks. In Gilbert L. Peterson and Sujeet Shenoi, editors, *Advances in Digital Forensics VIII - 8th IFIP WG 11.9 International Conference on Digital Forensics, Pretoria, South Africa, January 3-5, 2012, Revised Selected Papers*, volume 383 of *IFIP Advances in Information and Communication Technology*, pages 99–113. Springer, 2012. doi: 10.1007/978-3-642-33962-2_7.

[226] UK Forensic Science Regulator Guidance. The Control and Avoidance of Contamination in Scene Examination involving DNA Evidence Recovery. Standard, The Forensic Science Regulator, Birmingham, UK, 07 2016. URL `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/915221/FSR_G-206_Issue_2_Final.pdf`.

[227] University of Chicago Press. *The Chicago manual of style*. University of Chicago Press, Chicago, IL, 16 edition, 8 2010.

[228] U.S. District Court Nebraska. USA v. Maksim V. Yakubets – Criminal Complaint, Case No. 4:19MJ3142, 11 2019. URL `https://www.justice.gov/opa/press-release/file/1223591/download`.

[229] U.S. District Court West. Dist. Of. Pennsylvania. USA v. Maksim V. Yakubets and Igor Turashev, Criminal No. 19 342, 11 2019. URL `https://www.justice.gov/opa/press-release/file/1223586/download`.

[230] Frits Vaandrager. Model Learning. *Commun. ACM*, 60(2):86–95, 01 2017. ISSN 0001-0782. doi: 10.1145/2967606.

[231] Harm M. A. van Beek. A forensic visual aid: Traces versus knowledge. *Science & Justice*, 58(6):425–432, 2018. ISSN 1355-0306. doi: 10.1016/j.scijus.2018.08.006.

[232] Harm M. A. van Beek, E. J. van Eijk, Ruud B. van Baar, M. Ugen, J. N. C. Bodde, and A. J. Siemelink. Digital forensics as a service: Game on. *Digit. Investig.*, 15:20–38, 2015. doi: 10.1016/j.diin.2015.07.004.

[233] Roland A. H. van Oorschot, Sally Treadwell, James Beaurepaire, Nicole L. Holding, and Robert J. Mitchell. Beware of the possibility of fingerprinting techniques transferring DNA. *Journal of forensic sciences*, 50(6):1417—1422, November 2005. ISSN 0022-1198.

[234] Jan Peter van Zandwijk and Abdul Boztas. The iPhone Health App from a forensic perspective: can steps and distances registered during walking and running be used as digital evidence? *Digit. Investig.*, 28 Supplement:S126–S133, 2019. doi: 10.1016/j.diin.2019.01.021.

[235] Marielle Vennemann, Claus Oppelt, Stefanie Grethe, Katja Anslinger, Rolf Fimmers, Harald Schneider, Carsten Hohoff, Martin Eckert, Thomas Rothämel, Peter M. Schneider, and der gemeinsamen Kommission rechtsmedizinischer und kriminaltechnischer Institute als Mitglieder der Spurenkommission. Möglichkeiten und Grenzen der forensischen DNA-Analyse unter dem Gesichtspunkt verschiedener Szenarien zur Spurenentstehung. *Rechtsmedizin*, 31(5):395–404, 10 2021. ISSN 1434-5196. doi: 10.1007/s00194-021-00508-2.

[236] Jacobus Venter. Process Flow Diagrams for Training and Operations. In Martin S. Olivier and Sujeet Shenoi, editors, *Advances in Digital Forensics II - IFIP International Conference on Digital Forensics, National Centre for Forensic Science, Orlando, Florida, USA, January 29 - February 1, 2006*, volume 222 of *IFIP Advances in Information and Communication Technology*, pages 331–342. Springer, 2006. URL `http://dl.ifip.org/db/conf/ifip11-9/df2006/Venter06.pdf`.

[237] VERBI Software GmbH. *MaxQDA*, 2023. `https://www.maxqda.com/` (accessed 13 Dec. 2023).

[238] Lena Voigt. Collection and Systematization of Phenomenon-Specific Cybercrime Knowledge at the Example of Botnet Crime. Master's thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2022.

[239] Stefan Vömel and Felix C. Freiling. Correctness, atomicity, and integrity: Defining criteria for forensically-sound memory acquisition. *Digit. Investig.*, 9(2):125–137, 2012. doi: 10.1016/j.diin.2012.04.005.

[240] Hans Walder and Thomas Hansjakob. *Kriminalistisches Denken*. Kriminalistik, C.F. Müller, 10 edition, 2016.

[241] David A. Wheeler and Gregory N. Larsen. Techniques for Cyber Attack Attribution. Technical report, Institute for Defense Analyses, Alexandria VA, 10 2003.

[242] Jennifer Jie Xu and Hsinchun Chen. Criminal network analysis and visualization. *Commun. ACM*, 48(6):100–107, 2005. doi: 10.1145/1064830.1064834.

Index