On a Scientific Theory of Digital Forensics

Martin Olivier

► To cite this version:

Martin Olivier. On a Scientific Theory of Digital Forensics. 12th IFIP International Conference on Digital Forensics (DF), Jan 2016, New Delhi, India. pp.3-24, 10.1007/978-3-319-46279-0_1. hal-01758695

HAL Id: hal-01758695

<https://hal.inria.fr/hal-01758695>

Submitted on 4 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

Chapter 1

ON A SCIENTIFIC THEORY OF DIGITAL FORENSICS

Martin Olivier

Abstract A suitable theory to serve as scientific grounds for a digital forensic science is still elusive. Such a theory needs to satisfy the demands imposed by science and justify the facts derived as evidence using the theory. A number of grounding theories have been proposed. This chapter revisits three prominent theories, those of Gladyshev, Carrier and Cohen, and: (i) determines the requirements they suggest for a digital forensics theory; (ii) analyzes their primary differences; and (iii) assesses them using the norms that exist for science. This enables us to sketch the outlines of a new theory that better reflects the scientific requirements and the intended application of forensic science in a digital context.

Keywords: Forensic science, digital evidence, digital forensic science, theory

1. Introduction

Forensic science is, quite literally, a science intended for use in the forum – the place in society where matters of law are debated. The role of forensic science is to contribute to the debate by providing answers to issues relevant to the debate, where the veracity of the answers stems from the use of science.

Exactly what constitutes science is a question that has occupied (reflective) scientists and philosophers for over a century. The problem of distinguishing between the scientific and non-scientific is often referred to as the “demarcation problem.” Although much insight has been gained about the demarcation of science, the problem is far from settled. Gratzner [9], for example, recounts a number of instances where theories that were not scientific were for periods deemed to be scientific. He is not interested in cases where scientists actively attempted

to deceive others by, for example, fabricating data. He is also not interested in “tales of scientific lunacy.” He studies cases that may be referred to as “pathological science.” He is fascinated by “the way that false theories and imagined phenomena sometimes spread through the scientific community. . . . Sometimes such a perversion of the scientific method results from external . . . pressures, but at other times it is a spontaneous eruption” [9].

Where Gratzer looks at aberrations in the broad scientific enterprise, Turvey [22] uses much stronger language when he speaks about forensic fraud that thrives in the “science cultures” of forensic examiners. Examiners are often not scientists, but they are put in a position where they act as scientists. Even when examiners are scientists, Turvey shows convincingly that the culture of law enforcement is that of a “noble cause” – one that brings criminals to justice. In contrast, the scientist is one who ought to value scientific truth above all else. Yet, many (if not most) forensic examiners are employed by law enforcement agencies. This juxtaposes two cultures that differ in their primary virtues. Foucault [5], for example, claims that, in order to understand the penal system, one has to determine whether “the insertion in legal practice of a whole corpus of ‘scientific’ knowledge, is not the effect of a transformation of the way in which the body itself is invested by power relations.”

This latter question resonates in the questions posed by Harding [10] about whose science and whose knowledge is valued and pursued. In her critique, she distinguishes between “bad science” and “science-as-usual.” Her critique of bad science overlaps with Turvey’s critique and contains some elements of Gratzer’s accounts. Her critique of science-as-usual is primarily based on the power structures in which the science is situated and the impossibility of the situated knowledge to be value neutral.

The questions that have been raised about the scientificity of digital forensics as a forensic discipline are primarily about demarcation. The underlying questions are about the scientific method used in digital forensics. Some responses have dealt with characteristics of digital forensic techniques; in general, the characteristics that have been mentioned are necessary, but not sufficient to declare that digital forensics is a science. Some answers have taken the form of suggesting a foundation for digital forensics that could make digital forensics scientific.

This chapter revisits the line of thought that seeks a theory for digital forensics that may serve as a scientific foundation for the discipline. In its quest to find such a theory, it attempts to stay within the boundaries of what would be considered scientific by science at large. It recognizes that scientific facts may be useful to solve a crime that has (possibly) been committed. However, it attempts to focus on science as an arbiter of fact

(within the limits of science), rather than science used in the pursuit of a criminal. In simple terms, it explores the possibility that digital forensics can consider some claim c and express scientifically-justifiable support for or rejection of the claim — if (and only if) scientific grounds exist to reach such a conclusion. Such proven claims (c or $\bar{c}$) may, if relevant, be accepted by a court of law.

In order to justify this course of action, the next section recounts the underlying logic employed by society in its use of scientific evidence. This is followed by a discussion of prominent attempts to formulate a foundational theory for digital forensic science. The intention to canvass these earlier theories is twofold. On one hand, the theories shed light on what scholars in the field have deemed to be important in a theory. On the other hand, it is possible to critique the theories based on the role of digital evidence that is elucidated in the initial sections of this chapter. Finally, the chapter identifies the characteristics that underpin a theory of digital forensic science.

2. Rationale for Scientific Evidence in Society

Science cannot prove the superiority of scientific knowledge over other forms of knowledge. However, the superiority of scientific knowledge is deeply ingrained in the fabric of many modern societies.

As a society, we (subjectively) value scientific claims; we rely on scientific claims about the safety and efficacy of our medicines; we are willing to board aircraft under the assumption that science assures us that the planes will indeed fly; and we will allow driverless cars to share our roads once we believe their safety has been demonstrated scientifically. We know that science is not infallible, but we generally experience an increase in the knowledge and understanding of our world and our technology and are willing to increasingly rely on science. In our narrower circle, as digital forensic practitioners and researchers, we are, after all, the readers and authors of scientific literature and are unlikely to devote our lives to an enterprise in which we do not believe. It is only natural that we want to have facts at our disposal when decisions in legal matters are made. And when such facts are not obvious, we prefer scientific proof of allegations above any other means of knowing, whenever possible. Our reliance on science is not shared globally – in some legal systems, religious insights are deemed more valuable than scientific ones. To the scientific mind, the grounds of beliefs of others seem particularly unreliable when used as proof – beliefs that are often built on a metaphysical foundation that does not allow objective scrutiny.

Accepted scientific theories are expressed as concise laws, such as the iconic $E = mc^2$. Sometimes these laws are conditional – they apply in a physical vacuum or where no external force has an effect on the body of interest. Once the preconditions are met, it is possible to apply science in even more specific terms: a feather dropped in a vacuum from the top of the Leaning Tower of Pisa will reach the ground in a time that can be calculated with extreme precision. These are the types of facts that, we hope, science will put on the table in our courts of law: this specific revolver fired the round that was found in the body of the victim; the fingerprints of the suspect match those found at the crime scene; Jack is the father of Jill. In the end, the scientist has to testify about some claim c . We prefer to hear a resounding yes, c is true; or a resounding no, c is false. In reality, the scientist will say that c is true or false – to some degree of certainty. We also realize that a scientist sometimes may not be in a position to confirm or deny c – tests are not always conclusive or even always possible. However, when tests purport to be conclusive, they must be reliable.

As a society, we have formed (absolute) rules to deal with uncertainty. In the current context, the rule that a suspect is innocent until proven guilty is a mechanism that deals decisively with uncertainty.

However, much of life is cloaked in uncertainty and no fixed rules exist to deal with the uncertainty. When a crime is committed, law enforcement may have much more uncertainty than certainty at its disposal. There may be leads, suspects, hunches and many metaphysical opinions may be expressed. A few facts may be available and some of them may be scientifically provable. However, in these uncertain conditions, it is often unclear whether such facts will turn out to be relevant – and one of the defining characteristics of evidence is relevance.

It is in this nebulous world of uncertainty that a detective investigates a crime (without even being certain that a crime was committed). Facts – even scientifically-proven facts – may help the detective plot a course through the leads, facts and opinions in a process known as an investigation. The detective is a puzzle solver, but not in the sense that Kuhn [12] uses to describe the activities of the (normal) scientist. The detective formulates theories about what happened, but these theories are very different from the theories formulated in science; the detective's theories correspond to something like the general theory of relativity only in name. The detective may formulate hypotheses about what happened, but the hypotheses again only have a family resemblance to the hypotheses that are formulated (and formally tested) in a scientific endeavor. At the end of the investigation, the detective turns all the uncovered evidence over to a prosecutor, who decides whether a *bona fide* (or, more

cynically, a winnable) case can be made against certain parties. If so, the matter is referred to a court, where facts become the currency presented to the presiding officer or jury. It is up to the court to decide whether or not the case has been proven. It is in this forum where the claims of the forensic scientist are offered as facts. These same facts may have played a role during the investigation – just like other leads, claims and hunches. However, it is in this forum where the evidence is proffered as scientific facts. It is in this forum where the claims may impact the life of the suspect (or of the scientist) for years to come.

The disentanglement of investigation and examination is important, not only because they are two inherently different activities that superficially mimic one another, but also because many of the problems in current forensic practice stem from investigative methods that drifted to the forensic laboratories without ever establishing scientific grounds – a problem highlighted by the seminal U.S. National Academies’ report on forensic science [13].

From the narrative above, it is clear that a scientific theory of digital forensics needs to: (i) consider the domain of digital forensics; (ii) reflect on the nature of claims that may be offered as scientific fact; and (iii) reflect on the certainty with which the claims can be offered as facts. This is the journey that is followed in the remainder of this chapter. During the journey, it is important to also reflect on what others who previously attempted similar journeys said. From their journeys, we learn about the expectations of such journeys. We also learn about the pitfalls that prevent their theories from being the ultimate blueprint of digital forensic science. It is also prudent to be aware that the journey – like those that have gone before – is bound to be found lacking in many respects. However, we hope that one more foray into the world of digital forensic science will shine more light on a topic that is likely to remain murky and foggy for some time to come.

3. Domain of Digital Forensic Science

Digital forensics is often seen as an activity with the purpose of extracting evidence from all things digital for use in judicial proceedings. One consequence of this informal definition is that, taken to its logical conclusion, digital forensics encompasses almost every form of scientific (and much non-scientific) evidence. Almost every piece of laboratory equipment runs some software that processes values from a range of sensors and presents the results of a test or measurement in a visual form to the scientist (or technician or even novice) who interprets it. This would put digital forensics in a peculiar relationship with other forensic

Table 1. Technical specifications of a DNA analysis system [7].

Name	Description
Data Output Files	.bmp, .fsa and .cmf formats
...	...
Internal Memory	80 GB solid-state drive
External Connections	USB 2.0 GPS (USB 2.0; L1 frequency reception; sensitivity > -150 dBm) Wi-Fi 802.11 Ethernet (RJ45 10/100/1,000 Mb data rates) SVGA, DVI
Security	Multiple encryption systems for stored data WPA2 encryption Strong passwords; secure logging of all accesses to local database
System Clock	Clock synchronizes with GPS signal

sciences. As an example, consider the relationship between a forensic biologist and a digital forensic specialist. The brochure of a DNA analysis system states [7]:

“The instrument’s Expert System software analyzes the data after run completion and provides real-time feedback on the usability of the STR profiles for database searching.”

Table 1 lists some of the specifications of the device. The quotation and selected specifications leave no doubt that the instrument is a digital computer. If the definition of digital forensics as provided in the opening sentence of this section is accurate, it follows that evidence about the output of this device should be in the domain of digital forensics. In this manner, digital forensics becomes the universal forensics because very few instruments do not operate in a digital fashion at their core.

As noted in Table 1, the DNA test instrument produces files in the .bmp, .fsa and .cmf formats. Identification and classification are core forensic competencies [11]. The average digital forensic practitioner will probably be willing to identify and/or classify a .bmp file. However, it is quite possible that the average digital forensic practitioner has never encountered the other two file types and the question then becomes: Which forensic branch is able to express an opinion when the tampering of such a file is alleged?

However, we should be cautious when introducing the notion of tampering: digital forensics has, for a long time, been preoccupied with identifying tampering and information hiding and, in general, trying to outsmart the really clever hacker. The average computer user is ar-

guably equally able to store secret information in, say, the slack space of an ext4 volume as he/she is to modify the code in the DNA analysis instrument. This does not mean that either is impossible. This also does not mean that a capability to detect such interference with normal system operation is unimportant. Instead, the consequence is that most digital information is an accurate representation of what it purports to be (subject to the well-known volatile nature of data where a file date is, for example, inadvertently modified by someone who opens the file to read it and then saves it when closing it).

Hence, we contend that digital forensics should, in the first place, focus on digital evidence that has not (maliciously) been tampered with (beyond what average users are able to do, such as deleting files). However, this assumption of correct data does not diminish the need to identify inaccurate data – it just means that there is a gap in our work on the scientific extraction of evidence from normal data.

As noted elsewhere, the forensic scientist should be able to testify to facts, which are usually formulated as claims or propositions. When initially faced with a claim, the forensic scientist considers it as a hypothesis (in a general sense), tests it, decides whether to reject or accept it (with due regard to the inherent problem of accepting a statistical hypothesis) and then testifies to the finding based on the result of the test.

The question then is: What facts can the digital forensic scientist testify to? The assumption that the facts may relate to all things digital is clearly problematic. We suggest that postconditions of computations may form this foundation.

This chapter explores a system that is simple, but powerful enough to form a foundation for digital forensics. Note that the intention is not to describe forensic methods or a real-world approach to conducting digital forensics.

3.1 Foundational Concepts

At the core, the digital forensic scientist examines digital artifacts that are acted upon by computational processes.

Definition 1. *A digital artifact is a sequence of bits that has (or represents) meaning. The meaning is often (but not always) determined by context.*

As an example of contextual meaning, consider the first few bytes of a file. They often indicate the type of the file. These bytes are the file signature or magic number. In other contexts, the same bytes may have no meaning (or a different meaning). In our definition, a sequence of bytes that function as a signature is a digital artifact.

A file is arguably a more obvious example of a digital artifact. The file technically is a file because it exists on a digital medium where metadata links its various blocks together, names it and details its other attributes.

Corollary 1. *A digital artifact may contain or consist of other digital artifacts.*

Corollary 2. *A combination of digital artifacts may constitute a digital artifact.*

One interpretation of this corollary is that the sum of the parts may be more than the parts on their own; however, no grounds for this interpretation are provided at this time.

Corollary 3. *The component parts of a digital artifact often provide details that help an examiner to classify (or even individualize) the artifact.*

The definition and corollaries enable one to denote (or name) artifacts, as well as to refer to the attributes of artifacts. A file f may, for example, have a name. The claim that

$$\text{name}(f) = \text{example.txt}$$

may be confirmed or refuted through examination.

Definition 2. *Processing (or computing) may create, modify or destroy digital artifacts. Stated differently, a computational action α may have one or more effects. It may be useful to indicate that such an action pertains to a specific artifact, although this will not always be practical.*

As an example, consider the editing of file f . An example of an action that causes a change to f is saving (or writing) the file:

$$w(f) \rightsquigarrow f'$$

In addition, writing may cause a backup file to be created (or updated):

$$w(f) \rightsquigarrow \exists f^\sim \text{ where } f^\sim = f$$

Other consequences of the write operation are that the file date is updated (or that the modification date and time of the “new” file f' is set to the time and date of the write operation). The update also impacts the containing disk:

$$w(f) \rightsquigarrow d' \text{ where } f \in d$$

where E is used to indicate an element of the structure – in this case, the disk d . Such a structure may sometimes be equivalent to a set; however,

a disk allows multiple instances of the same file and, hence, more closely resembles a bag. Note that E is used to indicate bag membership.

This chapter does not intend to introduce new notation. Many specification languages, such as Z, already have a rich notation in which postconditions can be specified.

Note that a computation often has a wide array of effects (manifested as postconditions) on the system. The algorithm used by software transforms its input into output, with (often) some logical relation between the input and the output. In some cases, details of the implementation may be left to the developer and, thus, specific software may leave its fingerprints in the output; different software suites may, for example, encode pictures slightly differently. Depending on the software, the input, output or both may remain in the system as artifacts. Software may leave other traces of its operation such as entries in a log file. All these characteristics are potentially postconditions of the computation.

It is not always useful (or even possible) to consider all the postconditions of a computation. The use of postconditions with the greatest discriminatory power and highest resilience are (obviously) the ones to focus on, but certainty also derives from redundancy. A sufficient number of postconditions need to be verified before a conclusion can be reached. It is not obvious what would constitute a sufficient number; this problem is part of ongoing work [18] and is not considered further in this discussion. It is worth noting that DNA evidence faces a similar problem, albeit in the realm of natural phenomena. The human genome is incredibly complex but, in general, of little forensic value. The major variations that tie genetic material to individuals occur in a small portion of genetic material. DNA evidence achieves its success by focusing on a very small set of loci that show enough variation in a given population to discriminate between individuals, if enough of the loci are used for comparison.

4. Achieving Scientific Status

We now turn our attention to the manner in which a truth claim can be justified in a scientific manner. This is the question that lies at the core of the seminal work on the scientific bases of digital forensic science. The best known answers are arguably the following (risking some oversimplification for the sake of brevity):

- **Gladyshev [8]:** Find a path through a finite state automaton that fits all known facts and terminates in the current state of the finite state automaton; such a path is known as an explanation.

- **Carrier** [3]: Formulate and test a hypothesis or a sequence of hypotheses.
- **Cohen** [4]: Establish the physics of information as a new science and use the facts in this science to find consistencies or inconsistencies to support or refute a claim.

In addition to these proposals, one has to consider digital forensic practice as, for example, embodied in the best known forensic tools. These tools typically extract data from devices and enable forensic examiners to inspect or observe data. In addition, the relationships between data can be rendered in a variety of ways (for example, using a visual representation). This enables the examiner to observe relationships, patterns and other potentially interesting characteristics of the data being examined.

In fact, observation is one of the key elements of all three theories, as well as of digital forensic practice. However, the perspectives of the theories with regard to observation differ and, therefore, require deeper exploration.

The three theories use automata theory as a foundational concept, albeit in different roles. The fundamentally different uses of automata theory also require exploration.

The theories (and practice) disagree about what they deem to be the primary artifacts of a digital examination. Cohen starts with a “bag of bits” and identifies the larger units of meaning in the bits. Carrier distinguishes between primitive and complex artifacts; complex artifacts are of primary interest, but one examines them with a hypothesis (or assumption) that the primitive artifacts that constitute them are sound (or can be shown to be sound – or not).

4.1 Observation

An important difference between the three theories is the relationship between (scientific) theory and observation.

Carrier and Cohen say that digital evidence is latent – that it cannot be observed by the naked eye. This is arguably consistent with the view of vendors of digital forensic tools and does not contradict Gladyshev’s statements. Carrier, for example, states that “[t]he investigator cannot directly observe the state of a hard disk sector or bytes in memory. He can only directly observe the state of output devices.” He immediately follows this remark about the nature of digital evidence by postulating that, “[t]herefore, all statements about digital states and events are hypotheses that must be tested to some degree.”

The requirement to observe, inspect or measure a feature of a (digital) object is a common activity in most scientific enterprises. In forensic science, in particular, these activities commonly occur in forensic laboratories. Examples include measuring the alcohol level of a blood sample, observing the presence of Y chromosomes in genetic material or inspecting blood spatter patterns. When evidence is latent, some instrument is, by definition, required for observation or measurement.

The question whether mere observations can be the foundation of a science is an old one. The logical positivists were already deeply divided about the role of observation (or protocol sentences [14]) in science, with the non-verifiability of personal observations at the core of the controversy. Much of the controversy can be traced back to Wittgenstein's *Tractatus Logico-Philosophicus* [23].

The outcome of this controversy was an almost unanimous rejection of observation as a building block in science. Bunge [1], in his comprehensive review of the philosophy of science, mentions parts of this controversy in passing; the role of observation is not a point of debate in his exposition. French [6] uses an apt analogy to illustrate the inability to generalize from observation to theory by reflecting on what and how a botanist would observe in a forest:

"Is she just going to parachute into the jungle, unbiased and without presuppositions and simply start observing, left, right and center? And observing what? All the plants, all the trees, all the strange animals and insects? No, of course not. She will know what she is looking for, what counts and what doesn't, what conditions are relevant, etc.; she may even have some theory in mind, some set of hypotheses that is being tested. Of course, serendipitous observations happen, new plants or animals are discovered, for example, but a botanist observing without bias in the field would be overwhelmed."

Carrier attempts to avoid direct observation (and measurement); the formulation of hypotheses and the subsequent testing of the hypotheses form the core of his claim that his work contributes to the discourse of the science of digital forensics. However, a large number of his hypotheses deal with situations where direct observations or measurements appear to be natural and the use of hypotheses artificial. An example is a scientist who observes the disk capacity on the label of the disk. The scientist then formulates a hypothesis that the disk capacity is indicated by the label (possibly after finding the relevant documentation of the device and determining the specifications from the manual). A tool may then be used to confirm the capacity. The debate about observation arguably explains this approach. However, simply measuring the capacity of the device appears to be a much simpler approach to achieve the same result.

Carrier, in fact, points out that “[m]ost forensic science disciplines have theories that are published, generally accepted, and testable.” However, while Carrier’s work emulates this process, it does not lead to theories that are tested, generally accepted and published. His process leads to contextualized theories – that are, in general, not published and do not become theories that are generally known – and, therefore, are not theories that would be generally accepted.

Philosophers who study measurement encounter problems similar to those who reflect on observation [20]: measurement is not theory-free, yet measurement (for example when testing hypotheses) helps create theories. This mutual influence of measurement and theory questions the very foundations of scientific theories. However, Carrier’s approach is not one that is intended to shape scientific theories. When Carrier formulates a hypothesis about disk capacity, his intention is to determine the capacity of the disk. Measurement would have had the same value as a hypothesis-based approach.

Cohen reminds us that digital evidence, by its very nature, is latent: all observation occurs via instrumentation. Observation in this context is very similar to measurement. Observation is, as noted, not theory independent; however, because observation only provides information about the case at hand – rather than impact a general theory – the problem of circularity is avoided. The claims of the proponents of protocol sentences apply; the critiques of opponents are avoided. An attempt to introduce hypothesis testing where observation or measurement suffice does not improve the scientific validity of digital forensics.

It is also worth noting that, unlike Carrier, both Gladyshev and Cohen appear to accept observation (through instrumentation).

4.2 Automata Theory

All three theories discussed here use automata theory as one of the tenets on which claims of being scientific are based. From this one may infer that justifiable facts about computational processes are deemed important by all three theories. This extends the domain beyond the examination of digital artifacts – which seems to be the primary focus of commercial digital forensic tools, where artifacts such as pictures and other documents may yield partial indications of the events that occurred during processing.

All three theories use finite state machines. Given that Turing machines define computability, the use of much less powerful finite state machines is somewhat surprising.

Cohen initially discusses finite state machines, but later correctly notes that a finite state machine with an unlimited tape added to it becomes a Turing machine. This happens before the core of his theory is presented. He notes that complexity and computability are based on Turing machines and derives a number of insights that are key to his theory of the physics of information. While both finite state machines and Turing machines are important concepts in Cohen's theory (and complexity, an important part in the theory itself), automata theory *per se* is not a core element of the theory.

Carrier employs an interesting variation of finite state machines. Instead of using the finite number of states that characterize a finite state machine, he suggests that the states between transitions are variable – that they play the role of memory that can change over time. In addition, the finite state machines may change from transition to transition, given the fact that a computer system changes as disks are mounted and unmounted. If an infinite number of finite state machines are available to substitute one another, the resulting device may well be Turing-complete. However, finite state machines play a minor role in Carrier's eventual theory and we, therefore, do not explore the computing power of these machines.

In contrast to the other two theories, finite state machines form the basis of Gladyshev's theory. The finite state machines do not directly model computations in the sense that a finite state machine is a representation of a program being examined. Instead, the states of the finite state machine are key indicators of important parts of the system state. While it may, in theory, be possible to consider the entire state of a system, the (theoretical) costs are prohibitive. A system with n bits of storage has 2^n possible states, ignoring the fact that additional storage may be added. In addition, the transitions between the states would be hard to determine; although, it is clearly not Gladyshev's intention to follow such an approach. Unfortunately, even when Gladyshev uses simple examples to illustrate his theory, the number of states quickly becomes unmanageable and raises the question if such an approach is practical.

Carrier remarks that his theory is not intended to be used in all its details for investigations in practice; just like a Turing machine is a very useful model of a real computer, but programming a Turing machine is unwieldy at best. The question then is whether an impractical theory of digital forensic science sheds useful insights on such science; if it does, it clearly has merit. Unfortunately, Gladyshev does not explore the nature of such (abstract) insights that may be gained from his model – and state explosion limits its practical use. Some questions may also be

raised about how one would determine the parts of the global machine state that should be included in the states modeled in the finite state machine – in particular, where multiple programs may have an impact on the state. However, this is an interesting problem for future research rather than a critique of Gladyshev’s model.

Gladyshev’s model presents an intriguing mechanism to test hypotheses. His model assumes that the current state of the finite state machine is known and that some intermediate states may be known. A trace of a path through the finite state machine states that covers all the known intermediate states and terminates at the current state explains the observations (or knowledge of intermediate and current states). Clearly, at least one explanation has to exist, else the observed evidence (or model of the states) cannot be correct. To test a hypothesis about whether an intermediate state of interest was reached, it is inserted as evidence and a search for explanations is initiated. The hypothesis is refuted when no explanations are found.

Unfortunately, this method of hypothesis testing also raises concerns about practical use of the theory. The time frame in which evidence was observed and the time frame in which the hypothesized state occurred has to be specified (with a finite degree of deviation allowed). This makes it impossible to ask and answer whether a given state may ever have been reached. Again, the practical implications of this aspect of the model are not quite clear. It is possible that this choice was required to prevent the complexity of computing the explanations from reaching intractable levels.

4.3 Bits, Bytes or Files?

A third area raised by the three theories is the notion of layers. Cohen asserts that the examiner starts with “a bag of bits” and infers the meaning at a higher layer (e.g., these bits are a JPEG file) using criteria he details.

Carrier distinguishes between (low level) primitive categories and (high level) complex categories. He suggests that investigations are usually performed at the complex level, with the primitive level serving as a theoretical backdrop on which the complex level rests. If necessary, remarks about complex categories can be mapped to primitive categories and examined at the primitive level. However, the important thing is the knowledge that this is possible, rather than an operational practice that is sometimes (or even ever) done.

Gladyshev notes that programs can be mapped to finite state machines (as discussed above), and proceeds to work at the lower level, without exploring the relationships between layers any further.

All three theories acknowledge the existence of layers. It may be argued that two layers do not adequately allow the creation of abstractions in the digital domain. A database built on a filesystem is a new reality, with the filesystem providing primitive support [15]. A dedicated application on a computer may use a database as a primitive, but leave no clue to the user of the system that a database is running in the background.

Other branches of forensic science are also confronted by such layers. A DNA forensic scientist and a forensic pathologist both work with human tissue, but at very different layers. Both know that tissue consists of atoms that, in turn, consist of sub-atomic particles. However, the existence of these sub-atomic particles does not directly impact the work of either scientist. Cohen observes that the division of digital artifacts cannot proceed beyond the bit-level – in contrast to the physical example above. From this, Cohen derives valuable insights for his physics of information. However, this does not detract from the fact that physical and digital artifacts may both be examined at different levels.

In this regard, Carrier’s position is most closely aligned with other forensic disciplines: while particles of an artifact may shed some light on its working (and may be important to understand the operation of the higher layer artifact), the artifacts that confront the forensic examiner should be the starting point of an examination. The question whether a file f was digitally signed by some user u is best addressed using notions of files, signatures, public/private keys and the algorithm(s) that could have been used given the known keys. Determining the truth of such a claim involves processing bytes (or even bits), but the terms listed are those that are used to formulate and answer the question.

Stated differently, the question posed to a digital forensic scientist determines the level at which an examination occurs. Analysis at a deeper level is only needed if a specific reason for the deeper analysis exists. The reasons may include anomalies found during processing or a separate request to express an opinion on the integrity of an artifact.

To defend the claim made in the previous paragraph, consider that an artifact may exist independently or exist embedded in some context. Copying a file from one medium to another does not change the content of the file. (The usual chain of integrity may be confirmed by computing cryptographic hashes of the original and copy, and comparing the hashes – a standard practice in digital forensics.) However, copying the file from one medium to the other may significantly change the order of the sectors that the file occupies on the medium, may change the number of sectors

per cluster and, hence, the total number of sectors occupied by the file, may change the content of the slack space in the final cluster, and so on. Whether or not these changes are important depends on the nature of the question to be answered. If there is no reason to doubt that the artifact to be examined is, indeed, a file, it may be examined as a file instead of as a sequence of clusters of sectors existing on some medium.

4.4 Investigations, Examinations and Analyses

The distinction between investigative and probative work was highlighted earlier in this chapter. This distinction is also present in the three theories being discussed.

Carrier sees the terms investigation and forensics as equivalent: “Digital investigations, or digital forensics, are conducted ...” Carrier also states that some definitions of digital forensics “consider only data that can be entered into a legal system while others consider all data that may be useful during an investigation, even if they are not court admissible.” He continues by pointing out that laws of evidence differ from one jurisdiction to another and concludes that “the set of objects with legal value is a subset of the objects with investigative value.”

We have argued elsewhere [17] that facts are useful during an investigation; facts are useful as evidence; however, facts are only one of the elements of an investigation – experience, hunches, tip-offs and the gut feel of the experienced detective are often just as useful, if not more useful – to investigating an incident than the mere facts. The detective follows leads and identifies possible sources of further evidence, possibly identifying and excluding suspects. When a suspect, for example, has a convincing alibi, the attention shifts to other suspects. While one may talk about the detective’s theory or hypothesis, the meanings of these words are very different from what a scientist means when using the same words (see, e.g., the critique of Carrier’s use of the term hypothesis [21]).

In the context of investigative psychology, Canter and Youngs [2] describe an investigation as a decision making problem, where the investigator is continually confronted by choices between alternatives to explore deeper next. Suspect profiling is one way in which investigative psychology helps in this regard: if it is known that the act being investigated is most often committed by someone who meets certain personal criteria (age, gender, employment status, relationship status, ethnic origin, and so on), then the investigation can be expedited by focusing on the suspects who match the profile. If the science underlying such profiling is sound, it will, indeed, expedite most investigations because it will more

often than not point the investigator in the right direction. However, such profiling does not constitute evidence.

When investigating a case based on a profile, an investigator may uncover facts that do constitute evidence. For example, the investigator may prioritize obtaining a DNA sample from a suspect who matches the predicted profile. If the sample from the suspect matches the DNA found at a crime scene, a fact has been established that not only guides the subsequent actions of the investigator, but that may also have probative value in a legal context. Note that both claims (the profiling and the DNA match) are factual and grounded in science. However, only one of them would be considered evidence, even in a colloquial sense.

Carrier attempts to avoid the issue of admissibility to work with “a general theory of evidence [...] which can be restricted to satisfy the rules of evidence in a specific legal system.” This statement is made after the observation that the definitions of digital forensics differ in their emphasis on an investigative process and obtaining evidence for use in legal proceedings. It is clear that evidence is a subset of the useful inputs to an investigation; it is not clear that, if all the elements of an investigation without probative value are removed, then one would be left with a non-empty set of elements with probative value. Phrased in simpler terms, if a theory describes an investigative process, it does not follow that the theory inherently says something about evidence. This is elucidated below using an example taken from Carrier.

It is, however, important to briefly reflect on the importance of this issue before proceeding. During the establishment of the NIST Forensic Science Standards Board, digital forensics was not included as a forensic science because doubts were expressed about the status of digital forensics as a science. One of the key documents in this debate was “Digital and Multimedia Evidence (Digital Forensics) as a Forensic Science Discipline” published by the Scientific Working Group on Digital Evidence [19] “to assist the reader in understanding that digital forensics is a forensic science and to address confusion about the dual nature of the application of digital forensic techniques as both a forensic science and as an investigatory tool.” The document suggests that digital forensics is a science when (or because) it uses scientific methods. When used in a digital investigation, “identification and recovery” are the primary activities that (as may be seen from its final paragraph) do not necessarily have to be scientific – this is in line with our discussion thus far. Whether or not the document convinces the reader that digital forensics is a science is not important for our purposes at this stage. The distinction between forensic science, on the one hand, and investigations, on the other, is indeed important.

An additional claim in the document – that forensic science may be useful for investigations – is (as noted already) obvious (assuming that this is what is meant by the claim that “the output of digital forensics can easily be used as direct input in digital investigations”). However, the claim that “[i]nformation can easily flow from digital investigations into digital forensics, [if it is] subjected to the rigorous process demanded by the scientific method and rules of evidence” is, at best, confusing. If forensic facts are proven during an investigation, they do not have to “flow into digital forensics” – they are then already proven forensic facts. Any investigation may lead to questions that can be answered by scientific tests; a digital investigation may lead to questions about the digital realm that may be answered in a scientific manner. If the notion that information may “flow from digital investigations into digital forensics” indicates that an investigation may raise questions that may be answered by scientific methods, then the nature of these questions arguably defines digital forensic science. We contend that this is not easy and that the word flow is only justifiable if claims such as information from murder investigations can easily flow into forensic pathology or into DNA forensics can be made in a meaningful manner.

In a telling example, Carrier illustrates his use of hypotheses where contraband in the form of pictures is suspected. His initial hypothesis is that the contraband is in the form of JPG images – an act of prioritization based on prevalence; in other words, a form of profiling that in most cases increases the likelihood of successfully locating evidence early in the investigation. However, when this hypothesis has to be rejected if no incriminating JPG files are found, then the outcome of the investigation is not affected. The search simply moves to less frequently used options to search for evidence; ultimately, the search involves carving files from the disk. In the end, whether the hypotheses are accepted or rejected has no impact on the nature of the evidence. It is only after the search space is exhausted that the conclusion is that no evidence was found. This matches many aspects of investigations from the investigative psychology perspective: profiling usually accelerates evidence discovery (and may even help solve cases that would otherwise not be solved). However, the fact that the evidence matched a more or less likely profile is of little evidentiary use in itself.

5. Discussion

The preceding discussion suggests that a grounding theory for digital forensic science needs to address four areas of concern:

- 1. Observation and Measurement:** Observation is already theory laden. Neither observation nor measurement can form the essence of a grounding theory. A theory is a prerequisite for observation and measurement. The discussion suggests that the effects of computations may form a suitable theory that could guide observation and measurement. Such effects are well known in computer science, where they are typically encountered in the form of postconditions.
- 2. Automata Theory:** While automata theory appears to be a natural fit for digital forensic science applications, its use in digital forensic science theories has not resulted in the successful development of a digital forensic science. It seems too far removed from computational processes to model computations in real or ideal examinations. The insight by Cohen that computational complexity based on automata theory can help determine if examinations should be conducted deserves to be developed further. Other work [16], which indicates that probabilistic algorithms may enable the examiner to quantify error rates, also needs to be developed further. Hence, it is suggested that computability theory may be more useful than automata theory as an ingredient of an overarching digital forensic science theory.
- 3. Artifacts:** The insight (in Carrier's words) that complex objects exist out of primitive objects is valid and useful. While other theories place different emphases on the relationship between complex and primitive objects, their focus clearly increases during the examination of complex objects as the expositions of the theories progress. However, the division of digital artifacts into only two classes is an oversimplification of the digital realm. The role of layers of metadata in, for example, databases [15] better reflect the nature of this realm in general. We suggest that an ideal theory would provide a natural layer of relevance for the forensic examiner and that the examiner should start the examination at this layer of abstraction; of course, the examiner should be able to delve deeper when necessary. As a simple example, a claim about a file would, in general, be independent of the filesystem (and hence, the sectors, clusters and other system structures on which the file is stored). One should be able to copy a signed file from a filesystem to another without affecting the properties of the signature. However, when an examination involves system attributes, such as the date of modification or the use of slack space, a lower level examination is clearly indicated.

4. **Science:** A distinction should be made between a science that may be trusted to produce scientifically-justifiable evidence and a science that may help expedite an investigation.

We further suggest that computations, rather than artifacts, should form the primary focus of a theory. A lack of space precludes a full exposition, so only a brief outline can be provided here. Computations usually manifest themselves by creating digital artifacts; these artifacts are the traditional objects of inquiry in digital forensics, but they often hint at the processes that created them. When computations are the units of an examination, the known postconditions of the computations indicate whether or not a given process could have been active. Moreover, it is possible to study combinations of computations by developing an algebra that determines the expected intermediate conditions and postconditions of the composite computation. Such a change in focus will facilitate the examination of real-world processes based on the algebra. Unfortunately, space does not permit the illustration of these claims. Note that a change to the examination of computations does not preclude the current focus on the study of artifacts. An artifact is typically the result of a computation; examining an artifact is, therefore, a special case of examining computations.

6. Conclusions

The chapter has explored the forces that impact a foundational theory for digital forensic science. It has considered the scientific imperatives, the intended application and the needs of digital forensic science as embodied in current digital forensic science theories. Based on these requirements, the chapter has sketched the outlines of a possible new theory – one where scientifically-justifiable claims can be made about computations based on the discovered artifacts and the known or knowable postconditions of computations. The approach suggests that intermediate conditions and postconditions of computations can be used to compare hypothesized processes against the discovered artifacts. However, it remains to be shown that the process is indeed useful and practical. In particular, the expected simplicity with which compound computations can be characterized and examined has to be proven.

Acknowledgement

The author wishes to express his thanks to his colleague and friend, Professor Stefan Gruner, for many fruitful discussions about several issues raised in this chapter.

References

- [1] M. Bunge, *Philosophy of Science: From Problem to Theory, Volume One*, Transaction Publishers, New Brunswick, New Jersey, 1998.
- [2] D. Canter and D. Youngs, *Investigative Psychology: Offender Profiling and the Analysis of Criminal Action*, John Wiley and Sons, Chichester, United Kingdom, 2009.
- [3] B. Carrier, A Hypothesis-Based Approach to Digital Forensic Investigations, CERIAS Technical Report 2006-06, Center for Education and Research in Information Assurance and Security, Purdue University, West Lafayette, Indiana, 2006.
- [4] F. Cohen, *Digital Forensic Evidence Examination*, Fred Cohen and Associates, Livermore, California, 2013.
- [5] M. Foucault, *Discipline and Punish – The Birth of the Prison*, Penguin, London, United Kingdom, 1991.
- [6] S. French, *Science: Key Concepts in Philosophy*, Continuum, London, United Kingdom, 2007.
- [7] GE Healthcare Life Science, DNAscan Rapid DNA Analysis System, Data File 29-0327-18 AB, Pittsburgh, Pennsylvania, 2014.
- [8] P. Gladyshev, Formalizing Event Reconstruction in Digital Investigations, Doctoral Dissertation, Department of Computer Science, University College Dublin, Dublin, Ireland, 2004.
- [9] W. Gratzer, *The Undergrowth of Science – Delusion, Self-Deception and Human Frailty*, Oxford University Press, Oxford, United Kingdom, 2000.
- [10] S. Harding, *Whose Science? Whose Knowledge? Thinking from Women’s Lives*, Cornell University Press, Ithaca, New York, 1991.
- [11] K. Inman and N. Rudin, *Principles and Practice of Criminalistics: The Profession of Forensic Science*, CRC Press, Boca Raton, Florida, 2001.
- [12] T. Kuhn, *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, Illinois, 1996.
- [13] National Research Council, *Strengthening Forensic Science in the United States: A Path Forward*, National Academies Press, Washington, DC, 2009.
- [14] T. Oberdan, Moritz Schlick, in *The Stanford Encyclopedia of Philosophy*, E. Zalta (Ed.), The Metaphysics Lab, Center for the Study of Language and Information, Stanford University, Stanford, California (plato.stanford.edu/entries/schlick), 2013.

- [15] M. Olivier, On metadata context in database forensics, *Digital Investigation*, vol. 5(3-4), pp. 115–123, 2009.
- [16] M. Olivier, On complex crimes and digital forensics, in *Information Security in Diverse Computing Environments*, A. Kayem and C. Meinel (Eds.), IGI Global, Hershey, Pennsylvania, pp. 230–244, 2014.
- [17] M. Olivier, Towards a digital forensic science, in *Information Security for South Africa*, H. Venter, M. Looock, M. Coetzee, M. Elooff and S. Flowerday (Eds.), IEEE Press, Danvers, Massachusetts, 2015.
- [18] O. Oyelami and M. Olivier, Using Yin’s approach to case studies as a paradigm for conducting examinations, in *Advances in Digital Forensics XI*, G. Peterson and S. Shenoï (Eds.), Springer, Heidelberg, Germany, pp. 45–59, 2015.
- [19] Scientific Working Group on Digital Evidence, Digital and Multimedia Evidence (Digital Forensics) as a Forensic Science Discipline, Version 2.0, 2014.
- [20] E. Tal, Measurement in science, in *The Stanford Encyclopedia of Philosophy*, E. Zalta (Ed.), The Metaphysics Lab, Center for the Study of Language and Information, Stanford University, Stanford, California (plato.stanford.edu/archives/sum2015/entries/measurement-science), 2015.
- [21] S. Tewelde, S. Gruner and M. Olivier, Notions of hypothesis in digital forensics, in *Advances in Digital Forensics XI*, G. Peterson and S. Shenoï (Eds.), Springer, Heidelberg, Germany, pp. 29–43, 2015.
- [22] B. Turvey, *Forensic Fraud: Evaluating Law Enforcement and Forensic Science Cultures in the Context of Examiner Misconduct*, Academic Press, Waltham, Massachusetts, 2013.
- [23] L. Wittgenstein, *Tractatus Logico-Philosophicus*, Routledge, Abingdon, United Kingdom, 2001.