

Towards a Digital Forensic Science

Martin S Olivier
Computer Science
University of Preoria
Pretoria, South Africa
Email: ms.olivier@olivier.ms

Abstract—The scientific principles that underlie digital forensic science are still not clear. Possible foundations have been proposed by Gladyshev, Carrier, Cohen, The Scientific Working Group on Digital Evidence of the US Department of Justice and others. However, all these proposals, although valuable contributions, do not provide a satisfactory scientific basis. The current article argues that the search for a suitable foundation is hampered by the conflation of science used for investigative purposes and science used for probative purposes. We argue that some aspects of forensic science are indeed useful for both purposes, but that large portions of the application of science for use in matters of law are only applicable to one of these categories.

The paper focuses on the probative use of science for matters of law. We suggest that the explicit focus on proof (rather than investigation) leads to a clearer understanding of the notion of the claims to be proven. Based on this it is shown that such claims may be expressed as propositions that can be proven, disproven, or determined to be ambiguous given the available evidence using well-known computing concepts. It also indicates how this approach helps one to determine the accuracy (which will not always be the opposite of error rates) of any findings. Given its specific focus the paper establishes a basis for digital forensic science without asserting that it is the only possible scientific basis.

Index Terms—Digital forensic science.

I. INTRODUCTION

The scientific status of *digital forensics* is a topic of debate [1]. There are a number of points of contention, including the nature of the science that forms the basis of digital forensics, how much forensics can be performed prior to an incident occurring and what the error rates of examination procedures should be to be considered useful or to have probative value.

A premise of this paper is that one of the obstacles to achieve clarity on the scientific status is the failure to distinguish between the investigative and probative aspects of an investigation where digital evidence plays some role. This distinction is, in fact, one of the key points emphasised in the SWGDE (Scientific Working Group on Digital Evidence) statement on *Digital and Multimedia Evidence (Digital Forensics) as a Forensic Science Discipline* [1]. However, the SWGDE document posits that these two facets are, in effect, two sides of the same coin — it effectively conflates these two issues. In contrast, the current paper argues that clarity can only be achieved once these two issues are disentangled.

Some support for disentangling these two elements may be found in the (relatively new) discipline of investigative psychology. The use of the adjective *investigative* points to the

focus of this discipline. The field naturally explores the nature of an investigation [2]. An investigation is seen as a process in which the investigator is faced with a series of decisions to make (such as which leads to prioritise). Such decisions may be made based on a variety of considerations, many of which are not (and do not have to be) scientifically sound bases. Investigative psychology aims to provide a scientific basis for some of these decisions. The field evolved from criminal profiling. While such ‘classical’ profiling is no longer the essence of investigative psychology, profiling helps to illustrate the issue at the core of the current argument. A profiling claim that a given crime (with its detailed characteristics) is usually committed by, say, a male in his forties, who works in a bureaucratic context (such as government) will clearly aid an investigation — especially if the profiling is solidly grounded in science. However, it is unlikely that a court will afford much weight to evidence that a given accused fits the profile. Profiling is based on the behaviour of the ‘group’ of such criminals, while the trial deals with an individual. And, as famously stated by Sherlock Holmes “while the individual man is an insoluble puzzle, in the aggregate he becomes a mathematical certainty. You can, for example, never foretell what any one man will do, but you can say with precision what an average number will be up to. Individuals vary, but the percentages remain constant” [3, chap. X]. The investigative process and a somewhat deeper assessment of the probative value of investigative psychology will be revisited below.

Investigations, in general, are not scientific processes. In contrast to what one sees in a Sherlock Holmes novel, police detectives are not (applied) scientists. To conduct their investigations they may use tools that were developed from scientific principles; these tools range from ballpoint pens, cars and cameras to tools that detect the presence of blood at a crime scene, where such blood is not visible to the naked eye. Such detectives are users of tools (and sometimes ‘users’ of scientists) to do their detective work. As an example retired Brigadier Piet Byleveld’s ability to solve crimes is still legendary [4]. The author of a book about his cases [4] explains in the book’s opening paragraph that “Piet Byl knows about writing docketts [...] not books.” The virtues of a good detective and a good scientist may be very different. There is no need (and it may be impossible) to find a scientific basis for investigations in general.

Note that this does not deny the value of scientific procedures that are valuable for investigative purposes, such as in-

vestigative psychology. In addition some scientific procedures are useful for both investigative and probative purposes. DNA found at a crime scene is one obvious example.

The focus of this paper is on the application of science to prove facts that are useful for matters of law. More specifically, the focus is on using science to prove facts in the digital realm that may be useful for matters of law. To be considered *evidence* a claim needs to meet a number of legal constraints, including the requirement that it should be relevant to the matter being heard, that an appropriate chain of custody should have been maintained for any artefacts on which the claim is based, that the bases of the claim should have been accessed or obtained lawfully and any other requirements that may be prescribed in a given jurisdiction. There should also be grounds to believe that the claim is correct. Such bases may include the direct observation by a witness (such as an eyewitness) or the experience of an expert. The particular concern of this paper are those cases where the grounds for accepting the veracity of a claim is science (rather than observation, experience or other ways of ‘knowing’).

The term *forensic science* is still typically defined as the application of science to answer questions that are of interest to matters of law: *forensic* indicate that this field is used to argue matters in a forum — typically in the modern setting in a court of law or similar context; *science* refers to both the nature of the discourse in the discipline and the grounds for accepting claims made. It should be noted that the phrase has been diluted in practice to refer to almost any activity where law and science meets. However in this paper we return to the original meaning where forum, scientific method and epistemology are the essential features of forensic science.

The intention of this paper is therefore to explore the nature of a branch of science that (also) uses science to ground claims to be made in a court of law. Phrased differently, it explores the possible nature of a digital forensic science. We assume that anything that can be proven can be formulated as an assertion, proposition, claim or hypothesis — which we deem to be equivalent for the purposes of this paper. Therefore, the problem to be explored is the nature of provable assertions about the digital realm that may be offered as testimony.

The first item on the agenda of the remainder of the paper is to reflect on investigations to contrast it to provable scientific claims, which we posited as the essence of digital forensic science. Section III then uses examples to illustrate the nature of (potential) provable scientific claims — albeit not from the digital realm, in general. Next, Section IV explores (and critiques) some of the previously proposed theories for a digital forensic science given the examples from other fields and the requirements for a digital forensic science. Section V determines how the earlier theories need to adjust to meet these requirements. Section VI casts the findings of the earlier sections as ‘provable scientific propositions’ to answer the primary question posed in the paper. Section VII concludes.

II. ON INVESTIGATIONS

From the new field of investigative psychology we learn that “the work of investigators is essentially a decision-making process” [2, p.10]. Detectives are regularly faced with alternatives to pursue, and information to serve as a basis for a decision is often sparse. In the end the investigation proceeds as a cycle: suspects and theories are eliminated, forcing detectives to return to earlier decision points. Better guidance at each decision point may lead to better initial choices and more effective investigations. Investigative psychology helps in this process by providing psychological insights that assist when such decisions are made.

When the decision can be bolstered by scientific information, forensics may play a similar guiding role. Similarly, hearsay, hunches and eyewitness reports may help in this regard.

In fact, this is one of the processes that the National Academies Report [5] comments on: How investigative techniques have found their way into courtrooms as forensic science without any scientific basis (and no consideration of the technique’s error rate).

One of the key distinctions that need to be made is the difference between work done by a scientist and work done by a technician. This distinction is not intended to diminish the value of technicians in forensic, research and other functions. The role of technicians is also not limited to only support of the scientist or scholar. The role and contributions of technicians will be discussed in more details below. However, the distinction between science and technique remains — and the mere fact that technique has proven its value does not make technical work a science. In the investigative process science and technology both contribute. However, in the probative context just the justifiable true belief of the scientist may be used to arrive at trustworthy conclusions; only the scientific theory provides a basis to predict. For Turvey [6] when technology masquerades as science it forms one of the forms of forensic fraud [6].

III. ON PROBATIVE TESTING

Requests for experts (or expertise) are occasionally distributed via a mailing list to members of the Digital & Multimedia Sciences (DMS) Section of the American Academy of Forensic Sciences (AAFS). This section contains requests for expertise based on the requests distributed via the DMS AAFS mailing list over the past 18 months. Since at least some of the requests were marked as confidential and many of the others not intended for public distribution, the requests have been edited to remove or modify information that could tie it to a specific person or case; the requests below have also been paraphrased to convey their essence as succinctly as possible.

It should be noted that the DMS section includes members with multimedia expertise who specialise in voice and image recognition and that many of the requests received over this period requested multimedia expertise (rather than digital evidence expertise).

We contend that the bulk of requests demonstrate a need to prove (or disprove) a claim in the manner highlighted earlier; there is already some hypothesis or interpretation available — the need is not to interpret, but to validate (or refute) an interpretation. We find support for this metaclaim in the fact that it emerges from the requests. A comprehensive list is therefore presented without discussing the individual cases. The list is presented in chronological order. Given the intention of this section — to establish a pattern, rather than make its case in the usual narrative style — it should be noted that it is not necessary to read all the examples to follow the remainder of the paper.

—o(1)o—

We require a forensic analysis of two sets of photographs. One set depicts a known subject, while the other are of an unknown subject. We need an opinion whether the photographic evidence establishes that the known subject and unknown subject are the same individual or two different individuals.

—o(2)o—

I have two, good quality head shots that are allegedly of the same person about 20 years apart. Because of the dissimilarities in the photos, I am not convinced the photos are of the same person. It is very important that I am accurate in my representations about whether they are the same man.

—o(3)o—

We seek the possibility of having an animated reconstruction of a fatal motor vehicle collision.

—o(4)o—

I have a corrupted video that will not play. The file that was created is about 1.6 GB in size, but will not load in any players. The agent believes that the problem stems from the fact that the stop recording button was not pressed to end the recording, but rather, the recording device was simply powered down. Is there any way to fix the file, or salvage what was recorded/captured?

—o(5)o—

We need a biometrics expert to assess another experts facial recognition analysis of several photographs. He used local binary pattern, three-patch, and four-patch methods. This expert may be required to testify in federal court about the analyses.

—o(6)o—

Two sets of pictures of injuries formed part of a case where a man was convicted for rape. We believe that the photographs and testimony do not match. Questions include (1) Who is in the photographs? Is it the same person in both sets? and (2) What size are the injuries depicted in the photographs?

—o(7)o—

We are looking for an “age progression expert” to assist with an investigation.

—o(8)o—

A person was convicted for ‘bombing’ another man’s vehicle as well as for illegal possession of some explosive devices (grenades) and weapons. The conviction was based on circumstantial evidence, rather than forensic evidence. Photographs were on two occasions taken of a hole in which grenades were found (near the convicted man’s house). In order to appeal the case an expert is needed who can testify about apparent differences in the hole that seem apparent in the two sets of photographs, because this may raise doubts about the manner in which the grenades was found in the hole (or even whether the grenades were really found in the hole), based on the size of the box that contained the found grenades.

—o(9)o—

My late husband served in the military in the Vietnam War. While he was there, there was a picture taken of him and he was given a Xerox copy of the photograph while he was in Vietnam. More than 20 years later this photograph was published in a book. After his death I became aware of a larger problem with his military records. There are two Vietnamese people that are claiming to be the two people I identify as my husband and a (now deceased) colleague in the Vietnam picture. I have a few other photographs of my husband from this time period. Perhaps there might be different levels of authentication that could be performed? My ultimate goal is to have my late husband’s service recognized.

Arguably, the fact that most requests deal with so-called multimedia issues is to be expected. The need to authenticate photographs, voice recordings and other artefacts represented on or using some medium predates digital technology by many decades. During its history questions about authenticity naturally arose, and the question about the accuracy with which matches can be made followed as an inevitable consequence. This arguable educated the law enforcement community about the value of expertise; to request expertise when faced with such questions became second nature.

In contrast, the issues about which digital evidence experts can make definitive claims is not yet clear. On the one hand, recovery of information from a digital medium is a skill that is appreciated and requested; however, this is often recovery in a technical, rather than scientific sense. And recovery tools improve continuously, with the implication that digital evidence expertise that is sought is often limited to expertise to use a tool.

What is arguably required is an understanding of claims that a ‘digital scientist’ can support or refute, rather than instrumentalist requirements that merely expect the digital information to be retrieved for analysis by the photographic,

sound, medical, or other science that can deal with the extracted content.

IV. EXISTING DIGITAL FORENSIC SCIENCE PROPOSALS

The literature on the scientific foundations of digital forensic science is still surprisingly sparse.

The work by Gladyshev [7] is one prominent attempt to establish such a scientific basis. It views the computations Turing machine calculations and posits that running the Turing machine in reverse will reproduce the original state (or, the intermediate state of interest for the investigation. Cohen [8] critiques this approach by pointing out that reverse execution is non-deterministic and therefore impractical to obtain information of interest.

The more important critique from the perspective of the current paper is that fact that the approach proposed by Gladyshev asks a “hat happened?” or “What was the initial states?” or even “What inputs were provided to the process?” These are all investigative questions. As noted earlier, we do not dispute the potential value of answers to such questions during an investigation. However, given our focus, the relevant questions in the context of program execution are of the form: “Prove that X happened/did not happen (if possible).” Proving (or, at least, verifying) such an assertion X is, in principle, much easier, because it may be possible to execute the computation in its normal (forward) direction: In its simplest form the assertion to be tested will be of the form $y=f(x)$ where x is the posited input value(s), f is the computation being considered and y is the posited result (or the state found at the crime scene). In this simple case just repeating the computation verifies or refutes the assertion. Note that this reformulation is not entirely new because it largely corresponds to Cohen’s reconstruction phase of his evidence examination process [8] (but it also differs from it in some important aspects to be discussed below).

Carrier’s work [9] on a hypotheses-based approach to digital forensics is another classic contribution to the science of digital forensics. As the title indicates, it stresses the importance of hypotheses in the forensic process to justify the scientific epithet. However, as has been shown, his treatment of hypotheses differs from their use in typical scientific research [10], and hence do not per se make the process scientific. However, his work is a stark reminder that artefacts may not be what they claim to be. A log file, for example, is hypothetically a log file — until its status is proven. There is a fine line between accepting what is clearly a log file without further testing and accepting what seems to be a log file — again without further testing. And how far does this process have to continue, because those artefacts that create the apparent log file or the entries in the log file may also not be what they purport to be. Cohen [8] approaches the same problem from the opposite point of departure. For him the (potential) digital evidence is a bag of bits — and work is required to infer any meaning from such bits (such as those bits constituting a log file). He suggests that this should be achieved by illustrating that facets of the bit sequences are consistent with what one

would expect to find in a log file (including actual entries) and showing that no inconsistencies exist (including occurrences that ought to have been logged, but were not, and occurrences that were logged, but where sufficient evidence exists that those logged incidents never happened).

In contrast to the work mentioned above, Cohen [8] posits the need for a new science, an information science, to form the basis of digital forensic science. The starting point of such a science is a number of self-evident truths, including the fact that the smallest unit of digital information is a bit, that effect has to *follow* cause and that digital space converges. (In contrast, units in analogue or physical space are continually split into smaller parts and physical space diverges; obviously effect in an analogue space also has the *follow* cause.) This leads to a number of very useful results. However, according to Cohen, the digital examiner starts with a bag of bits and assigns meaning to it through a process based on consistencies and inconsistencies. This assignment of meaning suggests an investigative mode of thinking, rather than a pure probative testing of potential evidence. In a legal context, if the opposing parties stipulate to the fact that certain files are log files and the examiner is requested to test some hypothesis about the log files, the starting point for the examination will be log files, rather than a bag of bits. This is not to say that an examination will never start with a bag of bits: if an examiner is one or more disc image and requested to express an opinion on the correctness of a certain part of a log file, the starting point will arguably be a sequence of bits to which meaning has to be assigned to discover artefacts in the image(s) that corroborate or refute the log entries. However, it is not clear that such a new science is required. In fact, Cohen later embraced the field of diplomatics as a better foundation for digital forensic science. As was the case for Gladyshev’s and Carrier’s work, Cohen also makes extensive use of formal models of computation including Turing machines and algorithmic complexity.

All of the approaches above fail to address the question of accuracy in terms of quantifiable error rates in any meaningful manner.

Olivier [11] tentatively suggests that algorithmics may be a suitable foundation for digital forensic science. There are many problems in computing that are intractable, but may be answered using probabilistic algorithms. The question whether a certain (large) number p is prime is one example. If the forensic examiner is faced with a question that can only be answered probabilistically the question about error rates can be answered precisely. For problems that are computationally tractable the error rate is then obviously zero. However, such theoretical results have to be converted to tools and the correctness of tool implementation becomes an issue. It should be noted that error rates of forensic tests assume randomness in the errors. If errors are biased the mere disclosure of the error rates is insufficient; the nature of the bias should also be disclosed (because bias implies a different error rate when the test is applied to a certain population). Tool implementation errors often provide incorrect answers for certain categories

of inputs and hence are not random. Hence, the zero error rate of a tractable problem cannot simply be replaced by the error rate of the implementation. One option is to assume that the tool has been calibrated, which may mean that it has been sufficiently tested or proven correct to accept that it implements the given problem accurately — and hence the zero error rate, as unlikely as it may seem — has to be accepted as correct. We will reflect on this issue in more detail below.

An alternative approach that moves away from the formal modelling of computation is one that uses the social science case study research method as proposed by Yin [12] as its basis; the details of this foundation are being explored by Oyelami and Olivier [13]. A number of parallels exist between this approach and the well-known certainty scale proposed by Casey [14]. However, where Casey's certainty scale is intended as an indication of accuracy, the work by Oyelami and Olivier is intended to provide a scientific basis for digital forensic science (adapted from a social science approach). In addition, this research method suggests a method to conduct forensic examinations based on a long history (and much reflection) of conducting "case" research in the social sciences. It is worth noting that case study research can prove causality, which seems promising when using it to conduct a digital forensic examination. However, as a qualitative method it will not enable the examiner to quantify error rates. This recalls the age old battle between qualitative and quantitative research, where a large body of work exists that indicates the value of qualitative research. The accuracy of a case study approach will arguably have to be expressed qualitatively. Note that Casey's certainty scale, despite its use of numeric confidence levels, is, in fact, a qualitative "scale". Given that qualitative certainty is frequently sufficient to add to the body of scientific knowledge, it is certainly worth considering the use of qualitative certainty in forensic evidence. Note that quantifiability is not required in all forensic disciplines; the forensic pathologist is not expected to say that "I am 99.9% certain that the cause of death of the victim was the gun shot to the head." The accuracy is established qualitatively, by exploring (and finding evidence to reject) alternative explanations — such that the victim was already dead when the gun shot was inflicted to the head.

V. TAKING IT FORWARD

Science and technology occupy separate realms, but often exist in a symbiotic relationship. As a simple example that applies to law enforcement consider the various mechanisms to calculate the speed of a vehicle. One simply has to measure the time it takes for the vehicle to travel from some point a to some point b ; the distance between a and b divided by the time it took determines the average speed of the vehicle between points a and b . Given the accuracy with which the distance between a and b was measured and as well as the accuracy of the clock it is possible to determine the accuracy of the speed measured. Suppose, for the sake of argument, that points a and b are $1m$ apart, then a car travelling at $60kmh^{-1}$ will traverse that distance in $0.06s$. If we know that

the maximum time measurement error for such periods is, say, 2% too fast or too slow, the speed measured may range from $58.8kmh^{-1}$ to $61.2kmh^{-1}$. However, if we also know that the points a and b may each shift up to $1mm$ in either direction (due to vehicle movement, changes in ambient temperature or other human or environmental factors, this has to be taken into account, and the range of speeds that may be reported for any actual speed correspondingly adjusted. If, on the other hand timing errors are normally distributed with, at most, a 1% error in 99% of cases and the court wants to be at least 99.99% accuracy, the calculation becomes slightly more complex. Clock accuracy may be determined theoretically, based on the known characteristics of the crystal or other mechanism that 'drives' the clock. Similarly, various methods may be used to determine possible variance of the distance between the two points. If points a and b are placed kilometres apart (say, at toll gates) the relevant factors change dramatically, but the concerns remain similar. Note that speed detection by radar uses a fixed time and then measures the distance travelled over that time. Now concerns about the frequency with which the two distance measurements measure the distance to the same object (and the same point on that object) become an issue. The questions raised in this paragraph (such as clock accuracy and error distribution) are complex issues that can only be answered authoritatively through science (which may even determine that the issues are not significant after all).

Clearly, determining the speed at which vehicles travel is an activity used for law enforcement. The results of such measurements are often offered in courts as testimony. And it is based on scientific principles that may become rather complex if all factors that may impact on the accuracy of measurement are to be included. Yet, few people would consider 'speedcops' (or the metro police or other authority responsible for enforcement of traffic related laws) as forensic scientists (or, as scientists at all). Neither would it seem that they are 'doing' science in some form, and presumably few of the practitioners understand the science that underlies their practice. Instead, they are required to operate the instruments according to very specific guidelines — and, it seems one of the simpler ways of escaping punishment when caught speeding, is to demonstrate that the officers in control deviated from the prescribed protocol to operate the equipment in some (minor) manner. This is often what is expected from the technician: to operate

We note that such speed enforcement meets many definitions of forensics, but postpone reflection on whether it is a form of forensics until later.

In other cases the technicians play a more direct role during research. Consider, for example, the first successful human heart transplant [15]. From a research perspective this procedure was one that tested the hypothesis if the donor heart is attached to the recipient's cardiovascular system in a specific manner it will continue to function. Prior to this a number of heart transplant attempts elsewhere in the world failed miserably. The hypothesis had merit because this first heart transplant team already demonstrated that the technique

enabled successful transplants in dogs.

In order to conduct this experiment a number of professionals were required: Dr Chris Barnard who tested his hypothesis, his bother, Dr Marius Barnard, who removed the heart of the (dead) donor, and 28 other people [16]. In his autobiography Dr Marius Barnard mention “the ‘pump’ technicians, who were responsible for managing the heart-lung machine that kept the patient alive while we operated. This is an extremely important function and the technicians never really received the recognition they deserved” [15].

VI. FORENSIC PROPOSITIONS

Above the assumption that the details of some computation may have probative value in a legal case. It is not hard to substantiate this assertion by simply providing an example. Consider, for example, whether a given digital document was signed using a given user’s private key. Proving the authenticity of the signature may determine how the case proceeds in terms of contract law. Ultimately we need to make claims about the authenticity of a signature, the similarity of two files or some other pertinent fact about an abstract concept, similar to the signature and file metaphors (or abstractions) mentioned. Quite a number of relatively detailed examples of claims were provided in Section III. In the digital realm Pollitt [17] similarly states that hypotheses (as narratives) about the evidence be made about the evidence *prior* to examination of the evidence; more specifically, the narrative should be tested against the evidence, rather than be derived from the evidence.

Where we differ from Pollitt is the greater prominence of computation in the formulation and testing of claims about digital artefacts and processes. Therefore, before proceeding to consider claims about such computational artefacts and processes it is necessary to briefly consider the underlying notion where some value is computed.

A. Computation

The notation used above to abstractly indicate a computation was $y = f(x)$. When initially introduced, it was indicated that both x and y may be vectors, which may make $\bar{y} = f(\bar{x})$ a more appropriate notation. In fact, the intention was that x and y may be the complete machine state before and after some computation f .

The range of meanings intended by the simple $y = f(x)$ notation is therefore wide, and more specific notations will arguably be required. However, we proceed by exploring the possible intended notions based on a believe that notation can best be selected once we have achieved a clearer understanding of what needs to be (de-) notated.

Let us start our exploration by what is, at least, marginally speaking, a scientific proposition. If, in $y = f(x)$, the function f is the *md5* function, x is a string of bytes and y is a value, the claim that $y = md5(x)$ is relatively easy to test. However, if *md5sum* is a program on some computer the claim that $y = md5sum(x)$ may be somewhat harder to test — a claim that will be revisited below.

Thus far the equality test tacitly assumed a single, deterministic outcome to computing f . Suppose r is a function that calculates a pseudo-random number between zero and the input provided to it. Then, if the claim $1.234 = r(5)$ is seen as the question of such a deterministic process, it would be absurd — a random number generator cannot always return the same value. However, it is possible to consider the case where the investigator found a case where r seemed to have returned this value — and the probative question may be whether r can ever return this value. Although we noted that our intention is not to introduce notation, the distinction between these two cases is important enough to introduce some tentative notation for this: $1.234? = r(5)$.

In terms of Turing-computability the difference between the two formulations may be formulated as follows. If x represents some specific part of the tape and y some other part of a Turing machine f . Then the first form would claim that running f yields y . The second form claims that y may result when running f with x on the tape. It is reasonable to assume that software on a modern computer can be represented as a deterministic Turing machine. Then, the only reason why x may in some cases yield y and in other cases not is that y depends on values on the tape that do not form part of x .

More specifically multiple possible results may result from values that are not part of x and/or values that are part of x , but known to be unknown.

There are two practical cases to consider. If f is a forensic tool then it ought to be assized to produce the correct *and only correct* answers — it cannot depend on any unknown inputs. The truth about a claim about what values other software would compute for some given input need to be supported using appropriate evidence — in particular that no unknown inputs impact on the result.

B. Assertions

After thorough deliberation, Inman and Rudin [18, p. 115] “address that aspect of forensic science that virtually defines it, the process of comparing an evidence and reference sample and forming and forming a conclusion about their relationship.” The questions that are amenable to answers through forensic science deal with identification, classification and individualisation (with the realisation that such questions often overlap). They continue (p.116): “Both classification and individualization attempt to answer the question of source” with possibly subtle but important differences regarding the source. In a related paper, Inman and Rudin [19, p.11] mention “the processes of identification, classification or individualization, association, and reconstruction describe the practice of forensic science starting with the recognition of an item as evidence.” This latter point — the recognition as something that may have probative value. Identification asks (or proves) “what is it?”; classification asks questions about a common origin of artefacts (that may already be implicit given identification; individualisation asks “which one is it?” or “whose is it?”).

More specifically [18, p.115], “identification, classification and individualization all depend on understanding the fundamental nature of matter or, for our purposes, the *nature of evidence*” (emphasis in original). In the digital realm this suggests that forensics depend on our understanding of the fundamental nature of ‘digital matter’.¹ The source of ‘digital matter’ is computational processes. With reference to the previous section, in $y = f(x)$ the input x and the process f are the source of y . This corresponds with an earlier claim [11] that algorithmics should form the basis of digital forensic science.

Identification of digital artefacts has already received much attention in the digital forensics literature; however, when the emphasis is moved from artefacts (files, databases, network packets, etc) to computing, new forms of evidence emerge — including, for example, values that were provided as input to some process. Work on identification therefore needs to be extended.

In addition, not much work has been done to prove that identification is performed correctly. For example, an image may be identified based on its file extension or the file magic number. The fact that such a file is viewable in an image viewer is an intuitive proof that it is indeed an image. However, a partial image (where some part of the file has, for example, been overwritten) may still be identifiable through its magic number (or even its extension). However, such a file typically cannot be opened in an image viewer. How certain can one be that it is (or was) an image? The file format of an image is specified (to a greater or lesser extent) as a grammar. Grammars define formal languages. The damaged image seems to be a partial string from such a language. If the manner in which the file was recovered indicates that parts are missing, the question whether the recovered string is a substring of a string that occurs in the language generated by the language, the substring may indeed be a partial image. If not, the claim that it was a partial image has been refuted. If it is a valid substring a measure such as the Hamming distance between it and the closest string may be an indication of ‘how much of an image’ we have. However, a Hamming distance does not seem like an ideal measure in this case; hopefully it inspires someone to find a better measure.

The previous paragraph dealt with images that use a highly structured format. Other artefacts have a much more flexible structure. A browser will do its utmost to display an HTML file presented to it, so the fact that a browser renders a file does not prove that the recovered file is indeed an HTML file. Validators exist that can determine whether a purported HTML file is indeed a valid HTML file. Unfortunately, a very large percentage of HTML files on the Internet are not valid HTML files. The question therefore needs to be asked what minimum

¹We acknowledge that there is a leap of faith in moving from physical matter to digital matter; Inman and Rudin [19] derive their principles from the divisible nature of matter, while Cohen [8] convincingly argues that digital ‘matter’ can only be divided up to the bit-level. This does have an impact on our adoption of Inman and Rudin’s conclusions in the digital realm. These details will be the topic of subsequent research.

criteria must be met for a file to be identified as an HTML file. The current paper does not explore this further. We do, however, note, that tentative identification may enable it to be classified or individualised — in which case it may become possible to prove that it has been identified correctly.

Classification is based on shared class characteristics. The best-known example of class characteristics may be from ballistics, where the rifling patterns of various firearms typically differ. By inspecting the striation marks on a spent bullet the make (and possibly model) of the firearm that shot it may be determined.

Classification of digital artefacts may be possible where the artefact is created by some tool — which is arguably true for almost all current digital content. Executable files are, for example, produced by compilers from source code; using the notation from earlier ($y = f(x)$) x is the source code, y the specific compiler and y the executable code (or object code, to be linked to other code). The executable (or object) file is again a string that must match some grammar. However, the target grammar often (not only for executable files) allows some flexibility; the choices made by the processor are the tool marks left on all artefacts produced by it. Studies need to be performed to see to what extent such tool marks identify a common producer.

Individualisation refers to random traces left by a tool. The best-known example is arguably again from ballistics, where small metal particles inevitably become embedded in a firearm’s barrel. However, unlike rifling, these particles are irregular — and assumed to be random. Like rifling, they will leave striation marks on a bullet discharged through the barrel. However, if the particles are indeed randomly positioned, the striations on the fired round will be unique to that barrel. Photocopiers often leave similar identifying marks on all copies they make; such marks are known as trash marks and are used for forensic purposes. Source cameras may be identified in a similar manner [20].

In many cases tools are configured to fit a given role it plays. The fully qualified domain name of an organisation, or the songs in a music collection may find their way into such configuration files and, eventually, into artefacts produced by that tool. Individualisation is not always possible; where the configuration of the tool leaves enough choices (or incorporate enough unique attributes, such as an IP address) that leaves marks on the artefacts it produces (or handles), these markers may be used for individualisation. The example of a digitally signed message mentioned earlier, is another example — if the public or private key of the user is known.

To illustrate these concepts further, the hashes of known contraband files (such as child pornography images) are often distributed to investigators working on such cases. The investigator determines the hash of all image files on a suspect’s disc. If any matches are found, it is important to recognise that nothing has been proven yet. The identified file(s) can be retrieved from the disc and rendered. The determination of whether such an identified file is indeed contraband falls outside the expertise of the digital forensic scientist. The

forensic scientist can offer proof that the file existed on the suspect's computer and offer proof of some other attributes of the file (such as whether the date of the file can be proven to be correct).

VII. CONCLUSION

This paper suggested that legal processes where digital evidence plays a role may benefit from a distinction between the use of science for probative and investigative purposes. The probative aspect, by definition, requires proof of facts, rather than discovery of facts. Testing of claims or hypotheses provides more information, than what an open question does. This makes it unnecessary to trace programs in reverse as suggested by some researchers. In fact, the forward execution enabled by it is apparently straightforward.

Problems arise when incomplete information is available; in many cases one does not need all relevant information to prove a case — at least, as far as an investigator or judge is concerned, many facts are immaterial. As an example, it may be necessary to show that a suspect possessed contraband and actively obtained it to make a case against the suspect. The 'full' picture for an examiner may include the source from which the contraband was obtained. We think that symbolic computing may be useful when information (in a hypothesis) is incomplete and cannot readily be determined from available material.

The forward computing also becomes more complex when partial 'output' information is available (such as an incomplete file). To determine whether such incomplete information is consistent with possible true full information is one example where we foresee that the computational complexity may increase. Work is required to determine whether complexity and certainty can be balanced to still yield probative results.

REFERENCES

- [1] SWGDE, "Digital forensics as a forensic science discipline, version: 1.0," Scientific Working Group on Digital Evidence, Tech. Rep., Feb. 2014.
- [2] D. Canter and D. Youngs, *Investigative Psychology: Offender Profiling and the Analysis of Criminal Action*. Wiley, 2009.

- [3] A. C. Doyle, *The Sign of Four*. Spencer Blackett, 1890.
- [4] H. Retief, *Byleveld: dossier of a serial sleuth*. Umuzi, 2011.
- [5] Committee on Identifying the Needs of the Forensic Science Community, Committee on Science, Technology, and Law Policy and Global Affairs, and Committee on Applied and Theoretical Statistics, Division on Engineering and Physical Sciences, "Strengthening forensic science in the united states: A path forward," National Academy of Sciences, Tech. Rep., 2009.
- [6] B. E. Turvey, *Forensic Fraud: Evaluating Law Enforcement and Forensic Science Cultures in the Context of Examiner Misconduct*. Academic Press, 2013.
- [7] P. Gladyshev, "Formalising event reconstruction in digital investigations," Ph.D. dissertation, University College Dublin, 2004.
- [8] F. Cohen, *Digital Forensic Evidence Examination*, 3rd ed. Fred Cohen & Associates, 2012.
- [9] B. D. Carrier, "A hypothesis-based approach to digital forensic investigations," Ph.D. dissertation, Purdue University, 2006.
- [10] S. Tewelde, M. S. Olivier, and S. Gruner, "Notions of "hypothesis" in digital forensics," in *Advances in Digital Forensics XI*, ser. IFIP Advances in Information and Communication Technology, G. Peterson and S. Sheno, Eds., vol. 462. Springer, 2015, pp. 29–43.
- [11] M. S. Olivier, "On complex crimes and digital forensics," *Hasso-Plattner-Institut für Software-Systemtechnik an der Universität Potsdam, Technische Berichte* 63, 2013.
- [12] R. K. Yin, *Case Study Research: Design and Methods*, 5th ed. SAGE, 2013.
- [13] O. Oyelami and M. S. Olivier, "Using Yin's approach to case studies as a paradigm for conducting examinations," in *Advances in Digital Forensics XI*, ser. IFIP Advances in Information and Communication Technology, G. Peterson and S. Sheno, Eds., vol. 462. Springer, 2015, pp. 45–59.
- [14] E. Casey, *Digital Evidence and Computer Crime*, 2nd ed. Academic Press, 2004.
- [15] M. Barnard, *Defining Moments*. Zebra Press, 2011.
- [16] D. McRae, *Every Second Counts: The Race to Transplant the First Human Heart*. Penguin, 2006.
- [17] M. Pollitt, "It's not about the evidence... narrative forensic hypotheses in the age of social media," in *Advances in Digital Forensics X*, G. Peterson and S. Sheno, Eds. Springer, 2015, in press.
- [18] K. Inman and N. Rudin, *Principles and Practice of Criminalistics*. CRC, 2000.
- [19] —, "The origin of evidence," *Forensic Science International*, vol. 126, no. 1, pp. 11–16, 2002.
- [20] M. S. Olivier, "Using sensor dirt for toolmark analysis of digital photographs," in *Advances in Digital Forensics IV*, I. Ray and S. Sheno, Eds. Springer, 2008, pp. 193–206.