Digital Forensic Science: A Manifesto

Martin S Olivier

Department of Computer Science, University of Pretoria

ABSTRACT

Forensic examination of evidence holds the promise of making claims about the truth of certain propositions with the inherent accuracy and reliability that characterises scientific endeavours. The propositions may relate to the artefacts examined or related artefacts. The nature of propositions about which claims can be made depend on the extent to which given propositions fall within the ambit of scientific knowledge and on the extent to which the examined evidence is suitable for the application of established science. A continuing series of incidents illustrate that in many forensic disciplines that promise is not met — often because some branch of forensic science happen to not being scientific at all. In fact, serious assessments of forensic science have shown that many (if not most) branches of forensic science are not scientifically valid.

Digital forensic science is one of the newest members of the family of forensic sciences. A number of reasons for concern exist that it is following in the footsteps of its more established footsteps and repeating many of the mistakes of those other branches of forensic science.

This viewpoint is written in the form of a manifesto that is situated in the current discourse about digital forensic science and practice. If challenges the current developments in digital forensic science by positing a number of demands that digital forensic science have to meet to be deemed scientific. The demands are posited as necessary, but not sufficient to ensure that digital forensic science uses science to contribute to justice. Appropriate responses to the manifesto is a change in digital forensic developments or an informed debate about the issues raised in the manifesto.

KEYWORDS: Digital forensic science; Foundational science

1 INTRODUCTION

Most members of the public probably had a rather vague notion of forensic science until various TV shows — starting with CSI — carried an image of a forensic utopia into our living rooms on a weekly basis. In general we were impressed — to the extent that jurisdictions where juries are used had to deal with the so-called new *CSI-effect*: Juries wanted the detailed and authoritative evidence they got used to in their favourite shows in order to make what should have been simple decisions during their deliberations. Unfortunately, the reality did not match these expectations.

While many reports suggested that much of the forensic science on these shows was in principle realistic (apart from the speed at which tests results became available) the computer scientists (and some technically inclined computer users) amongst us were usually not impressed when digital evidence needed to be recovered. The ability to type a command or two to trace the exact physical location from which some message was received was often beyond what we could accept as science fiction.

However, the disillusion was not limited to the situations where jury members (or victims of crimes) learned that forensic science could often not provide the answers. A more significant problem emerged. In reality, various groups of people knew about these problems for many years, but a wider audience became aware of them as the popular media carried reports about the problem more frequently. The problem was that forensic science was not always as reliable as touted. Actually, it was worse: Much of what was used as forensic science had no scientific basis. These knowledge soon enough made news headlines. Stories of innocent people who were wrongfully convicted based on flawed forensic conclusions and spent much of their lives in prison (or were executed) before being exonerated are indeed stories of human tragedy, and deserve to be told. Such stories also sell newspapers. Unfortunately, such stories cast a shadow of doubt over forensic science in general. The greatest tragedy of all occurs when forensic practice in general deserves such a blanket of distrust.

Comparative bullet-lead analysis and microscopic hair analysis are just two examples of 'forensic sciences' that were discredited. Currently bite-mark analysis is clinging to straws to regain some of its former reputation. Are finger prints unique? Do we really interpret blood spatter correctly? The simple answer, using science as a yardstick, is *no*. In a September 2016 report (cited below) the US President's Council of Advisors on Science and Technology reconfirmed what is obvious to so many: They found that of the seven forensic disciplines they assessed, almost none could be deemed to be founded on science. (Their study only considered pattern comparison methods.)

And yet, we use those methods to send people to jail and (in some jurisdictions) to justify capital punishment.

Much research has been done in the field of digital forensics over the past decade. However, where so few of the forensic disciplines — despite their practitioners' best efforts — are not scientific, self-reflection is indicated for all disciplines.

The document below describes the context and then highlights some course adjustments that should be made

if digital forensic science want to be a 'real' science. As we move more of our daily lives into the digital realm, and realise that we have not yet (as of September 2016) worked out how to do proper forensic science in the physical world (with some notable exceptions), the need to think about digital forensic science has become imperative.

2 THE CHALLENGE OF A DIGITAL FORENSIC SCIENCE

Forensics is the application of science to determine facts that contribute to reaching a just deciding in a legal case. While the focus is on the law, these insights are also required in some other situations, such as when the root cause of, say, an aviation accident needs to be determined with sufficient certainty to prevent similar accidents in future whenever possible. As a society we often rely on science to make informed decisions about important matters. In the safety and efficacy of medicine, the prediction of severe weather conditions, the safety of new technologies and the determination of the root cause of disastrous accidents scientific answers are preferred over other forms of knowledge; in fact legislation often requires scientific proof of, for example, the safety and efficacy of a new medication before the medicine can be registered and offered for sale.

Forensic science, in principle, enables one to make similar informed decisions in courtrooms and elsewhere where the law is to be applied. However, if forensic science is not scientific but a pretence of science, trust in the endeavour is misplaced. Note that even where the word *forensics* (rather that the phrase *forensic science*) is used, the notion of science is implied; almost every academic paper on, for example, digital *forensics* contain some definition that invokes science as an inherent foundation of such forensics.

Mistakes (or even deception) occur in all forms of testimony in legal matters. However, mistakes in the context of forensics introduce a systematic bias with far-reaching effects on the justness of the justice system. It is not hard to find extensive lists of examples where forensic evidence was wrong and possibly lead to an incorrect determination that an accused was guilty or innocent. The Innocence Project[1] is a good starting point to find such examples; it should be noted that they make extensive use forensic science — in particular of DNA evidence — to exonerate the wrongfully convicted. Arguably the most thorough critique of forensic science (including recommendations about reforming the discipline) is the US National Academies of Science report [1] on the state of forensic science. In its assessment of various forensic disciplines it repeatedly finds that the specific discipline is not grounded in science. A more recent report by the US President's Council of Advisors on Science and Technology [2] finds that almost none of the (selected) forensics disciplines it examined meets the requirements of scientific foundational validity. They highlight, in particular, that "an expert's expression of *confidence* based on personal professional experience or expressions of *consensus* among practitioners about the accuracy of their field is no substitute for error rates estimated from relevant studies" [2, p.6].

The discourse in digital forensics has only seen limited self-reflection about the use of science (or scientific methods) in its activities [3]. While some notable exceptions exist, the few published claims that digital forensics is indeed scientific are often based on a limited understanding of science.

Interactions in the world in which we live increasingly occur in the digital realm; hence one would expect that criminal activities (and civil disputes) will increasingly rely on evidence obtained from the digital domain. The purpose of this manifesto is to — in a rather informal manner — reflect on the inherent qualities that a discipline needs to meet to be viewed as a forensic *science*.

For the sake of brevity we simply posit that large parts of work done under the digital forensics label ought not to be considered forensic science. Often a debate about such a statement reveals that different parties in the debate view the notion of science differently. However, while some differences in opinion about the nature of science will always exist, an activity cannot simply be 'designated' as scientific based on some notion of science. Both reports referred to above emphatically reject various disciplines claims to be scientific — despite the strong belief in some of those communities that their work is indeed scientific.

Below a number points pertinent to a digital forensic science are raised as a basis for reflection. They are not intended to form a comprehensive argument about the nature of digital forensic science, but are a reaction to some common themes in current research in the discipline; some points provide basic background information; others are introduced to either support or oppose some prevalent lines of thought in the research literature. The manifesto provided in the final section of this paper should similarly be seen as a document situated in the current state of digital forensics and the current discourse of its 'scientificness'. It is hoped that the manifesto will have some impact on the future course of digital forensic science — if not by correcting inappropriate lines of inquiry, then by a deeper reflection of how the discipline ought to proceed.

3 MUSINGS ABOUT DIGITAL FORENSIC SCIENCE: TRUTH, SCIENTIFIC TRUTH AND LEGAL TRUTH

3.1 Some remarks on when and where the 'science' is performed

If forensic science is the use of science to help answer disputes in legal and related matters a question that arises is when this science is actually performed. Consider, as a comparative example, the amount of science that underlies the operation of a modern motorcar. However, this does not make the average driver a scientist. Most mechanics won't be deemed scientists. In fact, very few car service facilities or repair shops would employ any scientists. This does not imply that such places and people do not possess extensive expertise in their specific domain. In fact, it will not raise many eyebrows if such a person is called as an expert witness in some case. However, the person will not

[1] `http://www.innocenceproject.org/`

be able to testify as a scientist.

In the realm of forensic science the following scenario is common: After extensive research a test is developed to detect the presence of some substance in, say, blood. Then a device is developed to execute the test. In a particular case a phlebotomist will typically draw blood from an individual, put it in the device and obtain a reading (or printout) from the device. A phlebotomist is not a scientist (and, in particular, not a forensic scientist): In the UK there is no formal qualification required to become a phlebotomist. Very few states in the US require phlebotomists to hold any particular qualification. In a court case they can attest to the fact that they labelled the blood samples correctly and operated the test device according to standard operating procedures. However, they are not qualified to offer any conclusions to the court based on the results reported by the device. Science occurs during the development of the test. Evidence on the interpretation of the result (as well as on the accuracy of the results) will have to be given by a scientist, who knows and understands the operation of the underlying science. In the case of forensic laboratories the report of the test will (officially) be prepared and signed by such a qualified scientist.

Note that the example in the previous paragraph does not imply that the actual 'testing' in the laboratory never requires a test to be performed by a scientist. The point is merely that many people involved in a forensic science process on a daily basis are not (and need not be) scientists. The process itself (and hence is development) needs to be scientifically sound; the scientific laws that underlie the process needs to be understood (and justified) by the developers of such a process.

Checking authenticity of data using a hash function is an example where 'science in a box' may be used by a technician in a digital forensic science laboratory. However, if the similarity is disputed — say due to new results about hash collisions — the active involvement of the scientist may be required. A hash function, after all, maps an infinite number of inputs to a limited number of outputs — and hence there will be hashes corresponding to an infinite number of inputs. Hence, using a hash to claim uniqueness is far from obvious and needs an underlying scientific basis before conclusions may be drawn that a given match is unique.

3.2 On the truths

Science is a quest for truth. The law, when considering disputes, often need to determine facts. Facts are claims that are true. It was inevitable that at some point the paths of science and law would meet.

As the first step in a project to reassess scientific truth for application in law it is necessary to recognise that two different notions of truth are involved: Scientific truth is truth that helps to explain our world. As we learn more about the world, the truth often needs to be adjusted. But these adjustments are not arbitrary. As an example, science explains how aerodynamic forces impact on a body that moves through air. Such knowledge can be used to design wings that cause sufficient lift to keep aeroplanes in the air. If it is at some point determined that the scientific theories were not perfect when that plane was built, that plane will not suddenly stop flying. The old truth was 'good enough'. The new truth is (hopefully) better. (Formally: it has more explanatory power.)

Legal truth, on the other hand, is whatever the court decides. Such a truth typically remains a legal truth unless it is changed through some judicial process (such as an appeal to a higher court or a change in legislation). Legal truth is often deemed as absolute (unless changed in such an explicit manner). Case law (that may date back to Roman times) in common law jurisdictions are deemed law until it is changed by a party authorised to do so.

When a scientific truth changes, the use of an older scientific truth in legal proceedings is often still 'good enough'. Problems arise where the older theory was not 'good enough', but where it was wrong. The question to ask when science is used in legal proceedings is not whether that science is perfect, but whether it is sufficiently reliable. Newtonian physics, for example, is adequate to consider the trajectory of a bullet even though such physics is often deemed to have been replaced by relativity theory. In a dispute about the height of some building or the boundaries of some property Euclidean geometry or traditional trigonometry may be used, even when it (in principle) proceeds from the assumption that the earth is flat. On the other hand, to use an old example, the discovery of the (non-existent) planet Vulcan was a mistake (and the theory that 'predicted' its existence was corrected by relativity theory).

While science does not claim to be infallible; it is not hard to find examples of scientific theories that were incorrect. However, many of the failed forensic disciplines were not based in science. When a 'non-scientific' forensic science discipline fails it does not represent a failure of science. It does represent a failure of the legal system which never explored the grounds on which such a discipline made claims that purported to be based on forensic science.

To testify, is to express propositions that one believes to be true. The belief that a certain proposition is true may, in the case of an eye-witness, be based on the fact that the witness observed what is being testified to. However, courts limit the nature of the grounds on which beliefs may be based to accept the belief as testimony. Most courts will not accept a belief based on divine revelation as evidence. Similarly, what one has heard from someone else is normally not admissible as evidence and typically rejected as *hearsay*. In some situations knowledge of those who have experience of actions, context or other related matters may share their expertise with the court as expert testimony, which may assist the court to better understand the matter at hand.

Forensic science, in contrast, bases its belief in the truth of a given proposition on science.

To illustrate, the ballistic trajectory of a projectile will follow after launch is known in physics. It can be calculated based on characteristics of the projectile, the speed and direction at which it is launched, the impact of gravity and a number of other variables. The accuracy with which the trajectory is calculated does not depend on the experience of the person who performs the calculation. It uses theories formulated by people who may have died

centuries ago and could therefore be seen as hearsay evidence.[2] The calculation of such ballistic tables (also known as range tables) was a routine component of artillery used in battle. In fact, one of the prime reasons the development of computers became important at the time of the Second World War was to automate such calculations.

While such a calculation is relatively simple to perform for application in the artillery context, it may be much harder to perform in the forensic context, where reconstruction may be the aim. In the forensic context muzzle elevation angle, muzzle velocity, barometric pressure, wind speed and even the original size and shape of the projectile may be known to a reasonable degree of accuracy. However, to reconstruct the trajectory (or, at least, determine the area from which it was fired), the same theories — momentum, gravity, drag, drift and other factors — are used. However, compared to those who fire the projectile, the forensic examiner cannot readily determine the values of all these variables at the time of firing, which makes the computation significantly more complex and introduces errors, which the forensic examiner will (hopefully) be able to quantify.

It is worth pointing out that software behaves in a very similar manner. It is typically easy to predict what a program will do; however, reconstructing what a program did is much harder. A number of pertinent questions should be asked:

1. Are there scientific theories (akin to those use in ballistic trajectory calculation) that are useful to understand (or predict what will happen in) the digital realm?
2. To what extent (if at all) are such theories used in the reconstruction of events in a computing context in the current discourse on digital forensics?

The latter question may be reformulated as follows: What are the scientific theories that a digital forensic scientist can use to justify that his or her testimony is true? Do these theories meet the requirements of foundational scientific validity?

3.3 On the origins of forensic science

An alternative route to explore the nature of forensic science is an exploration of the roots of forensic science. This section explores two aspects of these roots: it explores the semantics of the phrase and the original recognised use of science in a court case.

Prediction observes a phenomenon (the 'cause') and predicts an outcome (the 'effect'). Therefore, if A (predictably) causes B, and A is the only cause of B, then if A and B happened, one can infer that A caused B. Stated differently, A now *explains* B. This is exactly how forensic science uses laws to explain phenomena; forensic science is often defined as a scientific analysis performed to determine the root cause of one or more events.

[2]The first use of scientific evidence in English Law occurred in *Folkes v. Chadd and Others (1782)* (often referred to as the Wells Harbour case). In summary it determined that "In an action of trespass for cutting a bank, where the question is, whether the bank, which had been erected for the purpose of preventing the overflowing of the sea, had caused the choking up of a harbour, *the opinions of scientific men* [sic]*, as to the effect of such an embankment upon the harbour, are admissible evidence…*" (emphasis added) [4, p.157].

Locard, by many seen as the father of forensic science, formulated what has become known as *Locard's exchange principle*; in his 1934 book *La Police et les Méthodes Scientifiques* he formulates it as "Any action of a human … cannot unfold without leaving some mark" [5, p.7][3] It has been formulated in a number of ways — often in the short form: every contact leaves a trace. While this principle is not a scientific law, it works remarkably well, and in many ways seem even more valuable in the digital realm. If we know from science that contact between X and Y leaves some trace T, observing X, Y and T may enable us to explain T (assuming the usual caveats about determining causes form what are deemed to be effects).

For such an explanation to be accepted as testimony, law is required to make two concessions: (a) It needs to recognise some notion of scientific truth, that may be conveyed from one scholar to another in a 'hearsay' fashion and that science has checks and balances in place to ensure 'truth'; these checks and balances override the need to hear (and cross examine) the original scholar to determine ('legal') truth; and (b) the scientist is allowed to *conclude* that the presence of A explains the occurrence of B. This latter concession is important because the law prefers to hear the 'facts' and then reach its own conclusions. Now some conclusion, based on science rather than the law, effectively becomes a fact in the legal process.

3.4 On the digital

Science, truth and reality in a sense form a triad: Science helps us to discover the truths about the reality in which we live. Conversely, if science make correct claims (in particular, correct predictions) about the reality in which we live, science has uncovered truth. (Postmodernists will arguably disagree here, but it is not clear that postmodern forensic science is possible…)

Now enter the digital realm. This is an environment that seems to be human-made. Many of its prominent concepts, such as *cyberspace* derived from science fiction. Reality in this context may be virtual — that is, reality may be 'unreal'. Yet, despite these idiosyncrasies, we have moved into this world lock, stock and barrel. Whether one views this 'cyberspace' as an alternative world, or just use the Internet for shopping, banking and talking, does not matter. The digital is integrated in our lives (or vice versa). If things in life go wrong we often need to prove claims we make. And, in this integrated world, many relevant events may have happened on the digital side.

Hence, it is no surprise that a branch of forensics — digital forensics — developed to find truths about what happened in the digital sphere. But, unlike the physical world, there seem to be very few rules that constrain the digital world. In the physical world, the rules of physics enable us to predict what will happen (or explain what happened). Is there a basis on which we can make such claims about the digital space?

This is the purpose of this text: to explore *which* questions about the digital space can be answered in a scientific

[3]This is a direct translation as he formulates it in the cited book: "Toute action de l'homme … ne peut pas se drouler sans laisser quelque marque."

manner, so that we can demonstrate a scientific truth for our claims — in particular claims that may be useful as evidence in a legal context.

3.5 Digital forensic science

From the preceding it is rational to question whether a digital forensic science can ever exist. Most (if not all other) forensic sciences deal with natural phenomena, natural substances or human nature (which is also natural in some senses at least). Even human-made tools are made of natural materials that will, when it interacts with any other natural material, behave in a predictable manner. Note that the word "natural" is used loosely here: plastics and other synthetic materials exhibit "natural" characteristics — that is, in contact with other materials, they will (to a lesser or greater degree) react in a predictable manner — this artificial material possesses (an artificial) "nature".

In contrast, computing and the various artefacts produced by it are as close to alchemy that humanity has ever come. It is trivial to program a computer to, for any inputs x_i, generate any desired outputs y_j; it is almost equally simple to modify 'trusted' software to produce arbitrary outputs for given inputs — unless security mechanisms are in place that will ensure that the software cannot be modified. Stated differently, it seems one needs to build systems that are so secure (and correct) that they perform as reliably and as consistently as a law of nature does. However, I think very few experts would be willing to stake their reputations on such an assumption that a piece of software in infallible (or even, say, 99.999% reliable).

The shift from digital forensic science to computing in the previous paragraph may not seem logical. However, digital evidence is — as Fred Cohen [6] so aptly states — a bag of bits[4] out of which the examiner has to extract some 'evidence'. Evidence (or, at least, meaning) may be inferred from one of only two processes: (1) If the bits through some justifiable process can be arranged to form some meaningful artefact,[5] then meaning has obviously been found. Alternatively, (2) if the bag of bits is the result of some computational process it may sometimes be possible to make claims about the inputs to that process and/or the process itself.

Conjecture 1 *Digital forensic science claims can only assume one or both of two forms, namely*

1. *That the digital data examined is an example of a specific class of artefact; and/or*
2. *That the digital data examined proves or disproves a claim that the data was the result of specific data transformed by a specific computational process.*

More formally these two claims may be stated as follows:

1. *For some 'recognised category' C (to be elaborated on later) and some sequence of bits s, the digital forensic scientist can* conclude *that $s \in C$; and/or*
2. *That, given some computational process P, some inputs x and some output s, the digital forensic scientist may* conclude *that, depending on the specific values, $P(x)$ did or could have produced s.*

Both of these claims, for the sake of simplicity, have been stated in a somewhat more limited form than intended. This will be addressed below. In fact, it will be shown that in this limited form the conjecture has much wider application than what it may seem initially.

To 'prove' this conjecture in one direction (namely that scientific forensic claims can indeed take one — or both — of these two forms) examples will suffice. However, conjectures are conjectures because they cannot (yet) be proven; to convince the reader that the examples to be provided are indeed correct, we need conjecture 2 to be introduced below. 'Proving' the conjecture in the other direction is harder and may indeed be shown to be false.

Note that conjecture 1 refers to digital forensic science specifically; it excludes branches of forensic science that may deal with digital artefacts, but where the claims made are not in the digital realm. Thus the intention here is *not* to deal with, for example, voice recognition or authorship attribution of a recording or a document, respectively, that happen to be in a digital format. Those branches include 'natural' properties (such as the properties of human speech or the vocabulary and style used to compose a document), and hence do not face the same challenges as 'pure' digital forensic science.

A defence of the conjecture that these are the only two valid forms of scientific forensic claims will be attempted later. However, to proceed in the former direction (that there are indeed two forms) another conjecture is required, which will be colloquially formulated as follows.

Conjecture 2 *In digital forensic science the notion of 'intelligent design' will often be sufficient to correctly classify an artefact. The degree of certainty with which this can be done depends on the nature of the class.*

To illustrate this conjecture, suppose that an investigator obtains a set of bytes for which some reasonable grounds exist to infer that a subset of the bytes are intended to be interpreted in a given fashion. To make this concrete, suppose one obtains a sequence of bytes from a system that are purportedly a JPG file. The claim (or hypothesis) that it is a JPG file may come from the file extension (if the file name is available), the initial bytes of the sequence and/or a variety of other clues. Conjecture 2 claims that we are able to determine whether the sequence of bytes indeed a JPG file or not and make that claim with a specified degree of certainty.

Of course, if the sequence of bytes conforms to all syntactic and semantic requirements for a JPG file, it opens in a variety of JPG viewers and (possibly) yields an identifiable picture, the sequence is a JPG structure, without any doubt. The only source of uncertainty is whether it existed on the medium from which it was retrieved as a JPG file. This is where the level of certainty needs to be determined. It is extremely unlikely that a random sequence of bytes from a medium will form a JPG image. On the other

[4] Perhaps the phrase "bag of of bit sequences" would have been more apt.

[5] The term *artefact* is meant to refer to something digital produced for later use; it forms the traces available to the digital forensic examiner. In a number of forensic science branches the term *artefact* refers to something artificially introduced into a photograph or recording that was not part of what was originally recorded; in those branches artefacts are ignored as artificial additions to recorded observations.

hand, in the unlikely case that one tries all permutations of subsets of bytes on a medium, the likelihood that one of the permutations will conform to the JPG specifications increases dramatically.

Another example may be useful: if one recovers an 8-bit value from a medium it obviously can be a member of the class of 8-bit unsigned binary numbers. The question whether it existed (that is, was used) as an 8-bit binary number on the system in question can only be answered after much more context has been studied.

Clearly, the likelihood of error does depend on the complexity of the artefact being examined: A JPG file has a header which not only has a standard format, but also has fields that impact the interpretation of the remainder of the file. As noted, when opening the file in an image viewer one would normally expect to see an intelligible image. If this is true of a file confidence grows that we indeed have a JPG file. A series of additional checks may be desired, such as the EXIF metadata to increase confidence — if required. In contrast, other formats may have much less inherent structure (effectively, much less redundancy/meaning) and it may be much harder (or even impossible) to confirm whether they exhibit "intelligent design traits."

It is now time to return to conjecture 1 to fulfil the promise that examples of the two conjectured claims would be provided. For the first form the example alluded to will suffice: If the investigator obtains a file that claims to be a JPG image (say, through its extension), determines that it conforms to the rules and specifications of a valid JPG image and, when opened with an image viewer displays what is clearly a picture, the conclusion that the file is indeed a JPG file is obvious. The contents of that file can then be reproduced in a form that will enable the court (or some suitably qualified expert) to make its findings. In some instances the digital forensic scientist will be qualified to do this, given the second form of conjecture 1.

As an example of the second form of conjecture 1 consider the case where the computational process P is the calculation of a known reliable hash function and the input x is a sequence of bytes. Then the digital forensic scientist may conclude that s is (or is not) the hash of x. To say that $P(x) = s$ is straightforward; however, the intention is also to conclude that $P^{-1} : s \mapsto x$, which needs to be qualified by the confidence (or error rate) of such a claim, because this is inherently a probabilistic claim. However, note that this example does not suggest that P should be a standard, well-studied computation: P may, for example, also be a piece of malware never encountered before.

Given the fact that these two forms of claims are used on a daily basis in digital forensics no further elaboration is required to substantiate their utility. What needs attention is their sufficiency and (eventually) a stronger justification that there is a scientific basis for (some) such claims.

4 THE MANIFESTO

The manifesto that follows represents the insights that may be gained from the discussion in the preceding sections. Not every point contained in the manifesto can be deduced in full from the preceding discussion, though.

Forensic science

1. The term *forensics* refers to forensic *science*. Any notion of a non-scientific forensics contradicts a generally held understanding in the academic literature and by the general public of forensics; such a notion would inherently cause confusion.
2. The utility of science in forensic science is the ability of science to explain phenomena. The explanatory ability of science is inherently related to its ability to (correctly) predict.
3. The reliability (or accuracy) of a forensic science (or forensic discipline) is limited by the accuracy with which the underlying science can predict.
4. The term *science* is contested and the problem of demarcating science remains critical. Philosophy of science provides many useful insights. In addition, standard scientific practices, such as peer review, provide a practical basis for demarcation. Both the appropriate nature and appropriate practice are necessary elements to denote an activity as scientific.
5. Forensic science ultimately has to explain why an event is seen as the root cause of other events. Forensic science therefore needs to be a science (or an application of a science or based on a science) that (a) can justifiably claim to be a science, and (b) has explanatory — and hence — predictive abilities.

The digital realm

6. Computing is used in many branches of forensic science, such as matching exemplar fingerprints with those stored in an extensive database or visualising physical phenomena in various ways. The fact that computers (and, hence, digital representations of phenomena) are used does not imply that digital forensics is being used. In these cases computing is used to support some forensic test. If a category descriptor is required for such computing the phrase *forensic computing* accurately reflects the activity.
7. The phrase *digital forensics* is commonly used to describe an examination of digital artefacts that *exist as digital artefacts* (rather than physical artefacts that have been converted to digital). To emphasise the point, fingerprints that have been transferred to paper are not examined as paper forensics; similarly fingerprints that have been converted to a digital representation do not form part of digital forensics. One possible characterisation of digital forensics is that it examines events (or traces of events) that happened in the digital realm; the purpose of digital forensics then is to determine the root cause of, or to reconstruct events that happened in cyberspace.

Examinations and investigations

8. Forensic examinations punctuate investigations. An investigation (such as a police investigation or a criminal investigation) typically includes many activities that are not scientific (and that cannot be scientific). Investigators often follow leads that are wrong or based on unreliable evidence. Decisions about which

leads to follow and when to abandon a specific line of investigation are often not based on objective criteria. The investigators' experience, intuition, the legal requirements of obtaining search warrants and other permissions, the behaviour of those implicated by a case and many other factors determine the course of the investigation. It is expected that the investigation will uncover relevant facts. The role of forensic science is to test hypotheses (or theories) that arise for (scientific) factuality. Investigators work with leads that range from unlikely to proven facts. The bar for considering something a lead is low, but leads may differ in strength; a strong lead may need much stronger grounds to be considered a strong lead. Do note that a proven fact may be a weak lead — if it, for example, turns out to be irrelevant.

9. The phrase *forensic investigation* is usually a misnomer. This phrase is often used when disasters (such as airplane accidents) are investigated. The phrase may allude to the fact that forensic science often fulfils a major function in such investigations. However, such investigations include many non-scientific aspects (such as interviews with eyewitnesses and survivors). Forensic examinations are typically conducted by laboratories best equipped for the specific test to be performed; in the case of major disasters, many forensic facilities in many countries may be involved, where each focusses on one component, one category of residue or some other specific facet of the investigation. The investigators work on the investigation; the forensic laboratories conduct their specific analysis limited to the question raised by the investigation team. In order to minimise confusion it is best to explicitly distinguish between forensic examinations (or forensic analyses) within the context of a (non-forensic) investigation.
10. With forensics being inherently scientific, care should be taken to not refer to non-scientific procedures as forensic activities. Many forensic processes may be useful during an investigation, but the converse is not always true. Forensic results will, in general, be admissible in court as evidence; leads and investigation results will not be admissible as facts unless sufficiently corroborated. Hence the investigator should clearly understand the difference between the two.
11. The word *evidence* should similarly be used with care. On the one hand, evidence may be whatever is collected in relation to a crime — whether it will have probative value or not. On the other hand *evidence* may be something that proves a claim; this is the type of evidence handed up in a court of law. In a forensic context, evidence ought to have the second meaning — (forensic) evidence will, in particular, be evidence about which forensic truth claims are made.

Independence

12. The use of forensic science is not limited to the criminal justice system; many forensic disciplines (with digital forensics a prime example) ought to be of use in civil matters, internal hearings and other contexts where such evidence may contribute to justice. This ought to be reflected in the language used to report research results; words such as crimes, guilt and innocence should be used judiciously in research on the topic since they ultimately affect the research agenda and application of forensic research. (This neither means that these words should be avoided at all costs, nor that work that is of particular use in a given context — criminal, civil, or other — should be discouraged; however, many forensic procedures will be applicable to more than one context, and its exposition and development should not be limited by unnecessary suggestions of context through examples, terminology or other potential biases.)
13. Forensic science is ultimately in the service of justice, rather than specific users of forensic science. Too often forensic science research focuses on its use in law enforcement (as the most prominent example). Digital forensic science will often be more useful in corporate contexts than most other forensic disciplines. Care should therefore be taken that the digital forensic research agenda is balanced and serves the interests both of those with and without access to resources. One way of gaining neutrality is for any work that produces a mechanism to prove some proposition p to also reflect on how to prove $\overline{p}$ or to prove some proposition q that could serve as a rebuttal of the claim that p happened.
14. Neither law enforcement, nor the corporate sector is, in general, equipped to do scientific research (with some significant exceptions). Hence the 'natural home' for forensic science research are the traditional research institutes, such as universities. Funding from industry or law enforcement should be recognised as possible sources of bias (if not in the research itself, then in the research agenda). Hence any such sponsorship should be explicitly declared as potential conflicts of interest.[6] In the ideal world digital forensic research will be funded by government or other bodies for whom a slogan such as a better life for all is inherently a call for justice, rather than the simplifying the task of law enforcement or big business. They now need to put their money where their mouths are.

5 CONCLUSION

Manifestos are often written by authors who deem it necessary to assume and express a position at a time when they perceive a danger that, if such an option is not expressed an opportunity will be lost to impact the direction of some discourse. The use of the term *manifesto* indicates strong convictions of the author that (a) an adjustment of the course is necessary and (b) the issues that are raised in the manifesto are those that ought to be high on the agenda for reflection. In this sense a manifesto is contextual; it speaks to the current discourse, rather than provide a conclusive, comprehensive perspective on the topic at hand.

In this manner, this manifesto does not attempt to define digital forensic science. Its intention is to highlight

[6] This is standard practice in most medical research and already strictly enforced by the American Academy of Forensic Sciences.

necessary (but not sufficient) aspects of a digital forensic science.

A manifesto is, by nature, a conviction set forth by its author(s). A conviction does not claim to be absolutely correct, but issues a challenge to others participating in the discourse to engage in further discussion on the points raised in the manifesto. As a conviction, it is a call for change in the current course of events. The weaknesses of forensic science have made newspaper headlines over many years, but the news media are often easily dismissed as being more interested in sensationalism, rather than facts. However, when organisations, such as the US National Academies of Science and the United States President's Council of Advisors on Science and Technology raises serious concerns about the absence or lack of (foundational) science in most of the forensic science disciplines that they considered, it is a loud and clear signal that introspection is required. The two reports issued by these organisations that were cited above say very little about digital forensic science, it does not absolve the digital forensic community from introspection and a well-considered response. Hopefully the manifesto above posits claims that will indeed lead to reconsideration of the 'old' answers to the issues raised (where) such answers exist, and reflection on the 'new' issues raised.

ACKNOWLEDGEMENTS

I would like to thank a number of colleagues, friends and family members who read, commented on (and critiqued) the manifesto. They will remain anonymous to protect them from the mob of digital forensic practitioners who may be upset by the publication of this manifesto. You know who you are. I appreciate your assistance.

REFERENCES

[1] 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". Tech. rep., National Academy of Sciences, 2009.

[2] President's Council of Advisors on Science and Technology. "Forensic science in criminal courts: Ensuring scientific validity of feature-comparison methods". Report to the President, Executive Office of the President, Sep. 2016.

[3] M. S. Olivier. "On a scientific theory of digital forensics". In G. Peterson and S. Shenoi (editors), "Advances in Digital Forensics XII"pp. 3–24. Springer, 2016. doi:10.1007/978-3-319-46279-0_1.

[4] H. Roscoe. *Reports of Cases Argued and Determined in the Court of King's Bench, in the Twenty-Second, Twenty-Third, Twenty-Fourth, and Twenty-Fifth Years of the Reign of George III*, vol. III. S. Sweet, Stevens and Sons and R. Milliken and Son, 1831.

[5] E. Locard. *La Police et les Méthodes Scientifiques*. Les Éditions Rieder, 1934.

[6] F. Cohen. *Digital Forensic Evidence Examination*. Fred Cohen & Associates, 3rd ed., 2012.