

Multimedia Forensics Is Not Computer Forensics

Rainer Böhme¹, Felix C. Freiling², Thomas Gloe¹, and Matthias Kirchner¹

¹ Technische Universität Dresden, Institute of Systems Architecture,
01062 Dresden, Germany

² University of Mannheim, Laboratory for Dependable Distributed Systems,
68131 Mannheim, Germany

Abstract. The recent popularity of research on topics of multimedia forensics justifies reflections on the definition of the field. This paper devises an ontology that structures forensic disciplines by their primary domain of evidence. In this sense, both multimedia forensics and computer forensics belong to the class of digital forensics, but they differ notably in the underlying observer model that defines the forensic investigator's view on (parts of) reality, which itself is not fully cognizable. Important consequences on the reliability of probative facts emerge with regard to available counter-forensic techniques: while perfect concealment of traces is possible for computer forensics, this level of certainty cannot be expected for manipulations of sensor data. We cite concrete examples and refer to established techniques to support our arguments.

1 Introduction

The advent of information and communication technology has created a digital revolution which is about to change our world fundamentally. Digital information stored in computing systems increasingly defines tangible parts of our lives and thereby becomes an ever larger part of reality. Moreover, many physical or 'real-world' social interactions are being replaced by their virtual counterparts through computer-mediated communication. As a consequence, the rule of law has to be extended to the digital sphere, including enforcement and prosecution of crimes. This raises the need to reconstruct, in a scientific and reliable way, sequences of actions performed in the digital sphere to find—or at least to approach—the truth about causal relationships. This is a prerequisite to hold potential perpetrators accountable for their actions and to deter imitators.

Endeavors to use scientific methods to gain probative facts in criminal investigations are referred to as forensic sciences (short: *forensics*). This term has its etymologic roots in the Latin word 'forum', which means 'main square', a place where public court hearings took place in ancient times. The term *computer forensics* has emerged to describe similar endeavors when computers are involved in criminal activities [1]. However, the definition of computer forensics is somewhat blurred, as computers can stand in manifold relations to crimes: they can be tools to commit crimes in the real world, or means that merely create a digital sphere in which crimes take place. In both cases, forensic investigators may strive to extract probative facts from the computers involved.




Fig. 1. Ontology of forensics, digital forensics and multimedia forensics

The situation becomes even more complex when we introduce *sensors* to the scenario. Sensors can capture parts of the reality and transform them into digital representations, such as images or audio files, which are then stored and processed in computers. Such digital representations of parts of reality can then be subject to forensic investigations, but they can only serve as probative facts if they are reliable and authentic. Realizing this goal defines the field of *multimedia forensics*.

This paper strives to clarify the definition of the various new variants of forensics, and to reflect on their underlying assumptions in comparison to classical forensic sciences known from the analog world. To do so, we employ an ontology of terms as illustrated in Fig. 1. One can subdivide all forensic sciences by their *domain of evidence*. This is the domain from which facts are extracted: classical (analog) forensics sets out to find traces of *physical evidence*, whereas *digital forensics* is limited to explore *digital evidence* [2]. While most people have a good intuition about the various forms of evidence derived from physical matters, digital evidence is intangible and therefore appears more abstract. Whenever we speak of digital evidence, we mean finite sequences of discrete and perfectly observable symbols, typically drawn from a binary alphabet, such as bit strings extracted from a computer’s memory and storage devices. So both *computer forensics* and *multimedia forensics* share their reliance on digital evidence and thus can broadly be subsumed to digital forensics. In the following, however, we argue that they differ substantially in their underlying assumptions, which justifies the distinction made in the title of this work. We would like to point out that we intentionally draw a very black-and-white picture of the addressed sub-disciplines in order to highlight their basic differences. Many practitioners from the one or the other field may be inclined to disagree with some of the assertions made. In practical investigations, of course, we will see a more grayish picture with combinations of different disciplines. We believe that such combinations in practice blur the important differences. This calls for a structured approach, to which this paper shall make a novel contribution.

In the remainder of this paper, we discuss each branch in more detail, taking classical (analog) forensics as a starting point (Sect. 2). Adhering to our terminology, we recall the principles of computer forensics in Section 3 to distinguish it from the discipline of multimedia forensics in Section 4. In Section 5 we

change the perspective to counter-forensics and discuss the main challenges for each of the two branches of digital forensics. The final Section 6 concludes with remarks on the possibility to combine the various sub-disciplines in practical investigations.

2 Classical (Analog) Forensics

Classical forensics refers to the endeavor to extract probative facts from physical evidence in the reality, i.e., the ‘analog’ world. It has been argued that the discipline draws on two principles: (a) *divisibility of matter* [3], and (b) *transfer* [4,3]. The first principle means that matter divides into smaller parts when sufficient force is applied. The smaller parts retain characteristics of the original matter as well as acquire characteristics generated by the separation itself.

The second principle, also known as *exchange principle*, states that whenever two entities interact in the real world, e.g., a burglar and a padlock, each entity will retain some physical matter of the other [4,5,6, among others]. Such exchanges can include for example fingerprints and footprints, hair, fibres of clothes, scratches, wounds, or oil stains. The examples show that transfer should not only be reduced to transfer on a microscopic scale. As Inman and Rudin [3] emphasize, transfer also includes the exchange of *patterns* (like footprints). So transfer means not only transfer of matter, but also transfer of traits.

If one accepts the principles as given, then it is straight to follow Kirk [7]:

“Physical evidence cannot be wrong, it cannot perjure itself, it cannot be wholly absent. Only human failure to find it, study and understand it, can diminish its value.”

This means that an *unconstrained* forensic investigator—at least theoretically—is free to analyse the scene (i.e., reality) from infinitely many perspectives (though not all at the same time). So he or she would have a non-zero chance to find even the subtlest trace. However, in modern epistemology it is accepted that human cognition of reality is in fact constrained in several ways. Most importantly, the human sense organs and the perceptual processes give us an incomplete picture of reality. While this can be considered as a filter which can be partly compensated for with technical means (e.g., microscopes increase the resolution of the human visual system), the Heisenberg uncertainty principle imposes an even greater constraint: the observer is always part of the very same reality and the sheer fact that he or she interacts with it, changes the object to be observed. This is also consistent with Inman and Rudin’s reflections on the division and transfer of physical matter [6] which do not distinguish between perpetrators and forensic investigators as entities taking part in the exchange.

What is important about this view for the argument in this paper is the reliability of probative facts derived from physical evidence. In other words, how difficult is it for a very sophisticated perpetrator to wipe out all traces, or even worse, to forge traces that can lead to false accusations? This corresponds to the attempt to modify reality in order to create a different picture of his actions.

Since both forensic investigator and perpetrator are part of the same reality and therefore subject to similar physical and cognitive constraints, even the most sophisticated perpetrator can never be sure whether his ‘modification of reality’ is fully consistent with the (imaginary) reality if the action had not taken place. So committing a ‘perfect crime’ in reality, and pretending a consistent picture of reality that hides all traces, is an incredibly difficult problem. As a result, careful investigations of physical evidence are likely to deliver either reliable probative facts or none. (Ignoring the possibility of lapses, which can never be fully excluded.)

3 Computer Forensics

Computers are physical machines that form part of our reality. Therefore, at first sight, if one accepts the divisibility and transfer principles, then they should equally hold for computer forensics. However, when people speak of computer forensics, they often make the implicit assumption that the forensic analysis is limited to the digital evidence stored in the status of the finite automata each computer represents. This implies an observer model with drastically reduced view on reality: bits alone are theoretical concepts that carry no side-information about their history. For example, the common practice to make a copy of the digital evidence stored in a computer and to base further investigations exclusively on this (read-only) copy, implements this observer model [8,9]. This observer model also implies that in computer forensics, divisibility of matter is not relevant. The transfer of traits remains as a possible basis for a theory of computer forensics.

This is not without consequences on the reliability of probative facts derived from such digital evidence. As the number of states in a closed system is finite, there is *always* a non-negligible chance that a sophisticated perpetrator leaves a computer in a state which *perfectly* erases all traces. Assuming that the entire persistent state of a computer is stored on hard disk, this can be achieved, for example, by using the computer after booting from a live-CD (and not altering anything on the disk).

Perfectly erasing all traces in practice is of course not always easy. The number of possible states that need to be controlled quickly grows intractably high. For example, nowadays standard PCs are equipped with about 100 GB of disk space. This translates to about $2^{10^{11}}$ states; for comparison, the number of atoms in the universe is estimated in the order of magnitude of 2^{10^3} . Especially in the complex modern networked systems with their many software components and hardware interfaces, perpetrators often fail to control parts of the state space. However, the perpetrator could use more technology, e.g., another system that simulates the relevant computer at the crime scene in a virtual machine. This helps to construct a valid and plausible state with reasonable time and effort, as only a negligibly small fraction of all possible states is actually relevant for finding a clean system state. This, however, implies an observer model that only sees parts of the entire state space; the observer ignores the additional technology

used to create the clean state. In practice, it is often hard for investigators to determine the borders of the system to be seized and analyzed, especially if it uses (wireless) network links.

Even with an observer model that captures the entire system, it follows from the limitation of the analysis to digital evidence that we can never ignore the possibility that a sophisticated perpetrator has covered all digital traces perfectly. In practice, such sophisticated perpetrators may be rare, but some skepticism is appropriate with regard to the residual probability of error whenever digital evidence is used in court to judge about capital crimes.

So does the principle of transfer apply to computer forensics? Many practitioners today will be inclined to agree, because from their experience every perpetrator makes mistakes and will leave patterns of criminal activity on the evidence. However, the digital nature of evidence makes it possible to cover traces perfectly. Furthermore, unlike practical limitations of the observer in classical (analog) forensics, the perpetrator knows all about this ‘blind spot’ of the investigator in advance and thus can adapt his action and pretend false or misleading facts. The advantages of inexpensive (due to automation) and convenient¹ computer forensics—most of the work can be carried out from the forensic investigator’s office—come at the cost of lower probative force. As social interactions move into the digital sphere, state-funded investigation offices have to make delicate decisions on the allocation of resources between exploitation of physical and digital evidence.

A completely different situation emerges if computer forensics is understood in a broader sense that comprises both physical and digital evidence (unlike in this paper). Such additional physical evidence, although sometimes costly and cumbersome to obtain, can be very indicative side-information. Features such as wear and tear, recordings of electromagnetic emanations [10], temperature [11], as well as all kinds of analog traces on storage devices [12] might reveal information about previous states of a target computer and thus thwart efforts to conceal traces. Even digital (or digitized) evidence stored in other devices (e. g., computers connected over a network link) can form such additional information *if* their integrity is secured, e. g., by means of secure logging [13]. For example, US agent Oliver L. North was convicted in the Iran-Contra affair in 1986 because he had overlooked evidence that was stored on backup tapes.

4 Multimedia Forensics

An important class of digital data which is often found (and analyzed) on seized mass storage devices is digital multimedia data. While digital and digitized media nowadays affect (and mostly enrich) our everyday life in innumerable ways, critics have expressed concerns that it has never been so easy to manipulate media data. Sophisticated editing software enables even unexperienced users to

¹ Note that the largest inconvenience in modern digital investigations results from the enormous amounts of data on seized computers.

substantially alter digital media with only small effort and at high output quality. As a result, questions regarding media authenticity are of growing relevance and of particular interest in court, where consequential decisions might be based on evidence in the form of digital media.

Over the past couple of years, the relatively young field of multimedia forensics has grown dynamically and now brings together researcher from different communities, such as multimedia security, computer forensics, imaging, and signal processing. Although multimedia forensics, like computer forensics, is based on digital evidence, the fact that symbols are captured with a sensor makes a difference which has implications on the reliability of probative facts. In the following, we will briefly define the field and then discuss the relation to other forensic sciences.

4.1 A Short Introduction to Multimedia Forensics

Scholars in multimedia forensics aim at restoring some of the lost trustworthiness of digital media by developing tools to unveil conspicuous traces of previous manipulations, or to infer knowledge about the source device. We call these two basic branches of multimedia forensics *manipulation detection scenario* and *identification scenario*, respectively. Note that multimedia forensics in this sense is not about analyzing the semantics of digital or digitized media objects. Techniques from multimedia forensics merely provide a way to test for the authenticity and source of digital sensor data.² This is a prerequisite for further analysis: probative facts derived from the content of multimedia data (for instance speaker identification from a microphone recording or license plate identification from a CCTV video) are only useful when the underlying data is reliable and authentic.

In multimedia forensics, it is generally assumed that the forensic investigator does not have any knowledge of a presumed original. Such methods are called ‘blind’ [15] and typically exploit two main sources of digital traces:

- ▷ Characteristics of the acquisition device can be checked for their very presence (identification scenario) or consistence (manipulation detection scenario).
- ▷ Artifacts of previous processing operations can be detected in the manipulation detection scenario.

The first class of traces is inseparably connected with the process of capturing digital media [16]. Since different sensors systematically vary in the way they transform (parts of) reality into a discrete representation, each capturing device is believed to leave characteristic features in its output data. The level of variation determines whether the corresponding traces can be used to distinguish the class [17,18], model [19,20,21,22,23] or specific device [24,25,26] of an acquisition device. Today’s multimedia forensic techniques mostly focus on the analysis of digital images. Here, one of the probably most-studied device characteristics is the CCD/CMOS sensor noise, which occurs in practically all

² A related discipline, which could be even framed into the general concept of multimedia forensics, is steganalysis. The link becomes obvious whenever we think of embedding a secret message as a manipulation of genuine sensor data [14].



Fig. 2. Typical image manipulation and detection with multimedia forensics. Presumably original photograph of Iranian missile test with one non-functioning missile (left, source: online service of the Iranian daily Jamejam today) which was replaced by a copy and paste forgery (middle, source: Iranian Revolutionary Guards). The detector output marks regions which were copied with high probability (right) [32].

digital cameras [25] or scanners [27]. Estimates of the so-called photo response non-uniformity (PRNU)—a noise source that reflects small but systematic deviations in the light sensitivity of single sensor elements—serve as ‘digital fingerprints’ that allow to identify individual acquisition devices. A useful analogy is the analysis of bullet scratches in classical forensics, which assign projectiles to weapons [28].

Besides the usefulness of such device-specific traces in the identification scenario, they also found wide application in the detection of manipulations [29,30,31,25]. By testing for the existence of consistent device characteristics in the whole digital media object, deviations from the genuine sensor output can be detected. For example, a block-by-block analysis that signals the absence of the expected PRNU in certain image regions can be seen as indication for possible (local) post-processing.

There are more ways to uncover manipulations of digital media. Traces of the applied post-processing itself can also be very indicative [32,33,34,35,36]. Forensic methods that exploit this type of traces approach the problem of manipulation detection from the opposite direction than techniques based on device characteristics. Not the absence, but the very presence of particular features is used as a probative fact for possible post-processing. Typical traces of manipulation include periodic inter-pixel correlations after geometric transformations, like scaling or rotation of digital images [33], or the identification of (almost-)duplicate regions after copy and paste operations [32]. An recent example for the latter type of manipulations is depicted in Fig. 2, which shows a forged image of an Iranian missile test that was analyzed with a copy and paste detector.

4.2 Relation to Computer Forensics

Even though both computer forensics and multimedia forensics explore digital evidence, we believe that they form two distinct sub-categories of digital forensics. This may seem counter-intuitive at first sight, since in any case, the domain of evidence is limited to the set of discrete symbols found on a particular device. In multimedia forensics, however, it is assumed that these discrete symbols were captured with some type of a sensor and therefore the symbols are a digital

representation of an incognizable reality. The existence of a sensor that transforms natural phenomena to discrete projections, which are then subject to investigation, implies that multimedia forensics has to be seen as *empirical* science. This resembles the epistemological argument brought forward in the context of steganography in digitized covers [37]. Literally, a forensic investigator can never gain ultimate knowledge about whether a piece of digital media reflects reality or not. Neither can a sophisticated perpetrator be sure whether his manipulation really has not left any detectable traces. Unlike computer forensics, digital evidence in multimedia forensics is linked to the outside world and cannot be reproduced with machines. Thus, while the principle of transfer does not necessarily apply to computer forensics, it does have a place in multimedia forensics.

To make complex matters like a ‘projection of reality to discrete symbols’ more tractable with formal methods, multimedia forensics employs *models* of reality (though rarely stated explicitly). PRNU-based camera identification, for instance, assumes that the sensor noise follows some probability distribution, which can be reasonably approximated with a Gaussian distribution. This way, the problem can be formulated as a hypothesis testing problem with an optimal detector for the applied model [25]. Another sort of models is implied in methods that try to detect copy and paste operations. The assumption here is that connected regions of identical, but not constant, pixel values are very unlikely to occur in original images [32]. The two examples stress that typical models function as yet another dimensionality reduction within the domain of digital evidence. (A first reduction is the projection of physical evidence to digital evidence, see Sect. 1.) So the models provide a very simplistic view of reality.

Obviously, the quality of probative facts resulting from multimedia forensic methods depends on the quality of the model. The better an underlying model can explain and predict (details of) reality, the more confident we can base decisions on it. A model of PRNU which incorporates different image orientations is definitely preferable to one that does not. It can help to decrease the probability of missed detection. False alarms can be reduced by removing so-called non-unique artifacts like traces of demosaicing from the PRNU estimates [25].

It is important to note that the uncertainty about the generally incognizable reality is not the only fundamental difference between multimedia forensics and computer forensics. The transformation from analog world to discrete symbols itself adds further degrees of freedom on the sensor level. Especially the extent of *quantization* is a very influential factor for all multimedia forensic techniques, but in general every sort of post-processing inside the sensor has to be taken into account for a thorough analysis of digital media data. By definition, quantization causes information loss and thus introduces uncertainty in the forensic analysis. Here, quantization not only refers to lossy compression schemes like JPEG, but for instance also to the resolution of the output data.³ When reasoning about what constitutes a *sensor* in a wider sense, i. e., including possible attached

³ Note that JPEG compression is by far the most relevant source of uncertainty in practical applications: virtually all known forensic methods are more or less vulnerable to strong JPEG compression.

data compression mechanisms, sooner or later one stumbles over the question of ‘legitimate’ post-processing. For example, scans of printed and dithered images in newspapers can result only in a very coarse digital representation of reality, but traces of inconsistent lighting may still be detectable [36]. Generally, it appears that the quality of sensor output necessary for sound forensic analysis heavily depends on the applied techniques—yet another aspect which has no counterpart in computer forensics (in the narrow sense), where digital symbols are not linked to the world outside the closed and deterministic system.

Following our previous comments on computer forensics in a broader sense (cf. Sect. 3), we have to point out the general similarity of multimedia forensics with additional side-information about previous states. Strictly speaking, a computer becomes a sensor whenever it records signals of its environment, i. e., the reality. Recordings can happen for various reasons, for example key stroke pattern are used to seed pseudo-random generators. Such pattern also convey information about the typist, who no doubt belongs to the reality [38].

5 Counter-Forensics⁴

Theory and practices to reconstruct crime scenes and, hence, to identify evidence are not reserved to the special group of forensic investigators. Most state-of-the-art methods are published in publicly available conference proceedings or journal articles. In general, transparency is a welcome security principle [39], but at the same time it makes it a bit easier for potential perpetrators to refine their strategies and to develop counter-forensic methods, which reduce the formation or availability of probative facts to the forensics process [40].

The horizontal order of sub-disciplines in Fig. 1 has been chosen intentionally to reflect gradual differences in the reliability of probative facts. This is emphasized by the two scales on the bottom. The measure *degrees of freedom* captures the amount of possibilities through which an investigator—at least theoretically—is able to collect evidence from the scene. Obviously, it is highest for analog forensics (physical evidence) and lowest for computer forensics due to the restricted observer model. The more restricted and the better predictable the observer model is, the easier it becomes for sophisticated perpetrators to manipulate the facts undetectably. This can be expressed in a measure of *forgeability*, which forms a kind of mirror image to the degrees of freedom and directly relates to counter-forensics. So in the following, we will explain the order of sub-disciplines with respect to forgeability in more detail.

In classical forensics, and adhering to the principles of division and transfer, counter-forensic methods that completely avoid the formation of probative facts cannot exist. Even the most sophisticated perpetrator can merely compete with forensic investigators to find the best abstraction of the real world in order to include as many as possible traces in their actions. For example, a perpetrator

⁴ The terms ‘counter-forensics’ and ‘anti-forensics’ appear synonymously in the literature. We prefer the former because it better reflects the *reaction to* forensics, as opposed to *disapproval of* forensics.

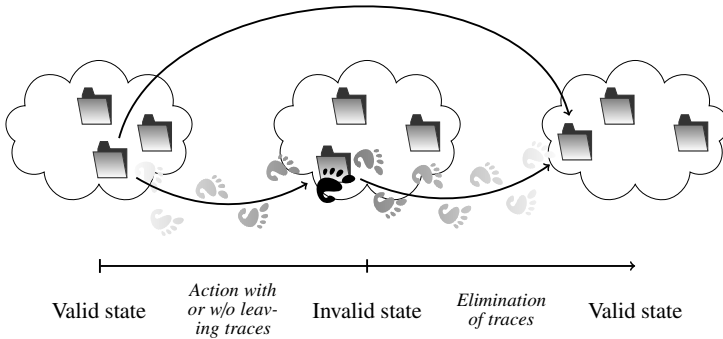


Fig. 3. Malicious actions cause traces by setting the machine to a suspiciously ‘invalid’ state that is detectable by forensics. Counter-forensics either eliminate all occurring traces subsequently (bottom path), or avoid traces preemptively (top path).

can remove traces like fingerprints by cleaning all touched surfaces. But indeed, cleaning surfaces with a cleanser and a cloth introduces new evidence. So the efforts to hide evidence transitively will most likely end up in an infinite recursion. Consequently, only human failure to find probative facts enables a perpetrator to be successful in the analog world (cf. Sect. 2).

This is totally different for computer forensics, where the discrete and finite nature of computer systems allows farsighted perpetrators, first, to determine valid states, and second, to reset a system from an invalid (i. e., suspicious) state back to a previously recorded valid state (see Fig. 3). For example, to cover data theft from a standard desktop PC, a perpetrator stores the initial state of the source device, transfers all portions of the demanded data, and finally, resets the source device to the initial state to remove occurring traces. While the success of this type of counter-forensic methods typically is limited to scenarios without any system inspection before the traces are eliminated, other counter-forensic methods can avoid detectable traces preemptively. A simple practical example is to boot an operating system from CD-ROM and mount the local hard drive in read-only mode. Figure 3 illustrates both approaches of counter-forensic methods.

Regarding multimedia forensics, we can distinguish two possible goals a perpetrator could strive for:

- ▷ altering the result of identification schemes by either suppressing the true source or counterfeiting a different one, and
- ▷ hiding post-processing by synthesis of authentic characteristics of the acquisition device or suppression of post-processing artifacts.

Practical examples for the first goal include attempts to suppress the device-specific noise pattern in a given image and replace it with the pattern of another camera [41,42]. Techniques to realize the second goal include a method for undetectable resampling [14] or the attempt to synthesize authentic demosaicing

artifacts in arbitrary, possibly manipulated, images [43]. The two counter-forensic approaches sketched in Fig. 3 also apply to multimedia-forensics [14]: While generating a valid demosaicing pattern after a manipulation clearly aims at removing suspicious traces, employing a undetectable resampling approach is intended to never leave any traces at all.

However, eliminating or avoiding traces by setting a valid state is not so simple in the case of multimedia data. This is so because the discrete symbols, via the sensor, depend on the scene that is part of reality. Although the number of possible states is finite, unlike in the analog world, it is too large to determine valid states with reasonable effort. And in contrast to computer forensics in deterministic machines, it is impossible to escape this problem by simply ‘virtualizing’ reality in a larger system. Consequently, sophisticated perpetrators and forensic investigators compete for the best abstractions to model relations between digital data and real world scenes. Their goal is either to hide or to counterfeit digital evidence (perpetrator’s point of view), or to detect even the subtlest modifications in media content (investigator’s point of view).

The reason why practical counter-forensics happen to work in laboratory settings is that current forensic techniques base their decisions on very low-dimensional criteria. In other words, they rely on a very simplistic models of reality. It is unlikely that these counter-forensic methods will still be successful against a combination of a handful of forensic techniques, so that the dimensionality is somewhat higher. And it is an open research question whether models can be found good enough to fool such combinations with novel counter-forensic techniques. In the meantime, an alternative could be to discourage forensic analysis by increasing the uncertainty through lossy, but inconspicuous post-processing (information loss through lossy compression or size reduction). The focal point here is the question which post-processing will be perceived as inconspicuous. This seems to be an inverse problem to the question for legitimate post-processing in Sect. 4.2. Both answers ultimately depend on established habits and conventions, which themselves are conditional to context information and may change over time.

6 Concluding Remarks

In this paper, we have devised an ontology to structure the various kinds of forensic disciplines by their primary domain of evidence. We deem such a distinction appropriate to clarify the assumptions and the logic of inference behind the different sub-disciplines. In particular, it became evident that the fact whether digital evidence is collected from the real world with a sensor, or merely represents an internal state of a closed and deterministic system, makes a difference with respect to the reliability of the extracted probative facts: it is harder to forge media data undetectably than to manipulate other digital evidence. Further, the notion of an observer model helps to distinguish the two extremes—computer forensics and classical (analog) forensics.

One may rejoin that this distinction is fairly artificial, as our conceptual borders are quite blurred in practice. For example, a police search could result in

a hard disk image, on which digital photographs are to be found with computer forensic methods. Then, multimedia forensics is applied to assign these photographs to a particular digital camera, which has been seized elsewhere. Fingerprints on this camera ultimately lead to the identity of the perpetrator via a police database. In this example, all kinds of forensic disciplines interact, spanning both digital and physical evidence, and jointly form a complete chain of evidence. While this holistic approach hopefully helps to convict the right person, such combinations in practice could hide the subtle differences between the various methods involved, and thus complicate the exercise to study each of them separately.

We see the contribution of this paper in a modest attempt to structure the field and to reflect on the (often implied) assumptions and models more explicitly and critically. Our proposal of an ontology and its accompanying terminology are understood as a starting point to stimulate fruitful discussion. Further refinements are envisaged for future research, along with an attempt to replace the informal arguments with more formal rigor. This implies that the deterministic view in this paper has to be replaced by the probabilistic theory of hypothesis testing.

Acknowledgements

The authors would like to thank the anonymous reviewers for their detailed and helpful comments and their colleague Stefan Köpsell for challenging our view of the world. Matthias Kirchner gratefully receives a doctorate scholarship from Deutsche Telekom Stiftung.

References

1. Kruse, W., Heiser, J.: *Computer Forensics: Incident Response Essentials*. Addison Wesley, Reading (2001)
2. Carrier, B., Spafford, E.H.: Getting physical with the digital investigation process. *International Journal of Digital Evidence* 2(2) (2003)
3. Inman, K., Rudin, N.: The origin of evidence. *Forensic Science International* 126, 11–16 (2002)
4. Locard, E.: *L'Enquête criminelle et les Methodes scientifiques*, Flammarion, Paris (1920)
5. Saferstein, R.: *Criminalistics: An Introduction to Forensic Science*, 7th edn. Prentice Hall, Englewood Cliffs (2000)
6. Inman, K., Rudin, N.: *Principles and Practices of Criminalistics*. CRC Press, Boca Raton (2000)
7. Kirk, P.L.: *Crime Investigation*. John Wiley & Sons Inc, Chichester (1974)
8. Casey, E.: *Digital evidence and computer crime*, 2nd edn. Academic Press, London (2004)
9. The Common Digital Evidence Storage Format Working Group: Standardizing digital evidence storage. *Communications of the ACM* 49(2), 67–68 (2006)
10. Kuhn, M.G.: *Compromising emanations: eavesdropping risks of computer displays*. PhD thesis, University of Cambridge Computer Laboratory (2003)

11. Zander, S., Murdoch, S.J.: An improved clock-skew measurement technique for revealing hidden services. In: *SSYM 2008: Proceedings of the 17th USENIX Security Symposium*. USENIX Association, Berkeley (2008)
12. Wright, C., Kleiman, D., Sundhar, S.: Overwriting hard drive data: The great wiping controversy. In: Sekar, R., Pujari, A.K. (eds.) *ICISS 2008*. LNCS, vol. 5352, pp. 243–257. Springer, Heidelberg (2008)
13. Schneier, B., Kelsey, J.: Cryptographic support for secure logs on untrusted machines. In: *SSYM 1998: Proceedings of the 7th USENIX Security Symposium*. USENIX Association, Berkeley (1998)
14. Kirchner, M., Böhme, R.: Hiding traces of resampling in digital images. *IEEE Transactions on Information Forensics and Security* 3(4), 582–592 (2008)
15. Ng, T.T., Chang, S.F., Lin, C.Y., Sun, Q.: Passive-blind image forensics. In: Zeng, W., Yu, H., Lin, C.Y. (eds.) *Multimedia Security Technologies for Digital Rights*, pp. 383–412. Academic Press, London (2006)
16. Khanna, N., Mikkilineni, A.K., Martone, A.F., Ali, G.N., Chiu, G.T.C., Allebach, J.P., Delp, E.J.: A survey of forensic characterization methods for physical devices. *Digital Investigation* 3(suppl. 1), 17–28 (2006)
17. Khanna, N., Chiu, G.T.C., Allebach, J.P., Delp, E.J.: Forensic techniques for classifying scanner, computer generated and digital camera images. In: *Proceedings of the 2008 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2008)*, pp. 1653–1656 (2008)
18. McKay, C., Swaminathan, A., Gou, H., Wu, M.: Image acquisition forensics: Forensic analysis to identify imaging source. In: *Proceedings of the 2008 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2008)*, pp. 1657–1660 (2008)
19. Kharrazi, M., Sencar, H.T., Memon, N.: Blind source camera identification. In: *Proceedings of the 2004 IEEE International Conference on Image Processing (ICIP 2004)*, 709–712 (2004)
20. Böhme, R., Westfeld, A.: Feature-based encoder classification of compressed audio streams. *Multimedia Systems Journal* 11(2), 108–120 (2005)
21. Bayram, S., Sencar, H.T., Memon, N.: Classification of digital camera-models based on demosaicing artifacts. *Digital Investigation* 5, 46–59 (2008)
22. Farid, H.: Digital image ballistics from JPEG quantization: A followup study. Technical Report TR2008-638, Department of Computer Science, Dartmouth College, Hanover, NH, USA (2008)
23. Gloe, T., Borowka, K., Winkler, A.: Feature-based camera model identification works in practice: Results of a comprehensive evaluation study. In: *Accepted for Information Hiding 2009, Darmstadt, Germany, June 7–10*. LNCS (to appear, 2009)
24. Geradts, Z.J., Bijhold, J., Kieft, M., Kurosawa, K., Kuroki, K., Saitoh, N.: Methods for identification of images acquired with digital cameras. In: Bramble, S.K., Carapezza, E.M., Rudin, L.I. (eds.) *Proceedings of SPIE: Enabling Technologies for Law Enforcement and Security*, vol. 4232, pp. 505–512 (2001)
25. Chen, M., Fridrich, J., Goljan, M., Lukáš, J.: Determining image origin and integrity using sensor noise. *IEEE Transactions on Information Forensics and Security* 3(1), 74–90 (2008)
26. Dirik, A.E., Sencar, H.T., Memon, N.D.: Digital single lens reflex camera identification from traces of sensor dust. *IEEE Transactions on Information Forensics and Security* 3(3), 539–552 (2008)
27. Gloe, T., Franz, E., Winkler, A.: Forensics for flatbed scanners. In: Delp, E.J., Wong, P.W. (eds.) *Proceedings of SPIE: Security and Watermarking of Multimedia Content IX*, vol. 6505, p. 65051I (2007)

28. Lukáš, J., Fridrich, J., Goljan, M.: Digital “bullet scratches” for images. In: Proceedings of the 2005 IEEE International Conference on Image Processing (ICIP 2005), vol. 3, pp. 65–68 (2005)
29. Popescu, A.C., Farid, H.: Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing* 53(10), 3948–3959 (2005)
30. Johnson, M.K., Farid, H.: Exposing digital forgeries through chromatic aberration. In: *MM&Sec 2006, Proceedings of the Multimedia and Security Workshop 2006*, September 26–27, pp. 48–55. ACM Press, New York (2006)
31. Mondaini, N., Caldelli, R., Piva, A., Barni, M., Cappellini, V.: Detection of malevolent changes in digital video for forensic applications. In: Delp, E.J., Wong, P.W. (eds.) *Proceedings of SPIE: Security and Watermarking of Multimedia Content IX*, vol. 6505, p. 65050T (2007)
32. Popescu, A.C., Farid, H.: Exposing digital forgeries by detecting duplicated image regions. Technical Report TR2004-515, Department of Computer Science, Dartmouth College, Hanover, NH, USA (2004)
33. Popescu, A.C., Farid, H.: Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on Signal Processing* 53(2), 758–767 (2005)
34. Kirchner, M.: Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue. In: *MM&Sec 2008, Proceedings of the Multimedia and Security Workshop 2008*, September 22–23, 2008, pp. 11–20. ACM Press, New York (2008)
35. Wang, W., Farid, H.: Exposing digital forgeries in video by detecting duplication. In: *MM&Sec 2007, Proceedings of the Multimedia and Security Workshop 2007*, Dallas, TX, USA, September 20–21, pp. 35–42 (2007)
36. Johnson, M.K., Farid, H.: Exposing digital forgeries in complex lighting environments. *IEEE Transactions on Information Forensics and Security* 2(3), 450–461 (2007)
37. Böhme, R.: An epistemological approach to steganography. In: accepted for *Information Hiding 2009*, Darmstadt, Germany, June 7–10. LNCS (to appear, 2009)
38. Joyce, R., Gupta, G.: Identity authentication based on keystroke latencies. *Communications of the ACM* 33, 168–176 (1990)
39. Kerckhoffs, A.: *La cryptographie militaire*. *Journal des sciences militaires IX*, 5–38, 161–191 (1883)
40. Harris, R.: Arriving at an anti-forensics consensus: Examining how to define and control the anti-forensics problem. *Digital Investigation* 3(suppl. 1), 44–49 (2006)
41. Lukáš, J., Fridrich, J., Goljan, M.: Digital camera identification from sensor noise. *IEEE Transactions on Information Forensics and Security* 1(2), 205–214 (2006)
42. Gloe, T., Kirchner, M., Winkler, A., Böhme, R.: Can we trust digital image forensics? In: *MULTIMEDIA 2007: Proceedings of the 15th international conference on Multimedia*, September 24–29, 2007, pp. 78–86. ACM Press, New York (2007)
43. Kirchner, M., Böhme, R.: Synthesis of color filter array pattern in digital images. In: Delp, E.J., Dittmann, J., Memon, N.D., Wong, P.W. (eds.) *Proceedings of SPIE-IS&T Electronic Imaging: Media Forensics and Security XI*, vol. 7254, p. 72542I (2009)