AUSTRALIAN JOURNAL OF FORENSIC SCIENCES
https://doi.org/10.1080/00450618.2024.2381535

Taylor & Francis
Taylor & Francis Group

OPEN ACCESS Check for updates

Understanding and comparing digital traces

Graeme Horsman*

Cranfield Forensic Institute, Cranfield University, Bedford, UK

ABSTRACT

Digital forensic practitioners will encounter digital traces during their examinations which they must take steps to understand. This may involve trying to attribute an '*activity*' to a trace (what created it) or determine where it came from (its '*source*') – Trace-to-Activity/Source interpretation. Alternatively, they may need to determine if an activity has taken place on a system by identifying traces denoting it – Activity-to-Trace interpretation. In both instances, practitioners may need to conduct tests and/or identify research which will help them understand a trace, and compare any results of their testing/research to the traces in their casework. This work describes both the Trace-to-Activity/Source and Activity-to-Trace interpretive journeys, as well as the steps contained in both. In addition, six 'trace comparison criteria' are proposed and discussed to help those carrying out a trace comparison, notably: 'trace location', 'trace structure', 'trace examination method', 'trace metadata', 'trace content', and 'trace context'.

ARTICLE HISTORY
Received 30 May 2024
Accepted 9 June 2024

KEYWORDS
Digital forensics; digital trace; digital evidence; testing; interpretation

1. Introduction

As part of any digital forensic (DF) examination, practitioners must identify and understand any digital traces that may aid both their and their wider investigative team's apprehension of any alleged offence[1]. While the identification of pertinent traces poses a challenge in its own right, this work focuses on the task of trying to accurately comprehend them. Interpreting the meaning of a digital trace (what it is, why it is present, etc.) is a complex task[2], and one that must be distinguished from merely stating the outcome of any given tool and any label it assigns to a trace that is processed. When interpreting digital traces, often practitioners will be faced with one of two scenarios; either they must try to understand a trace they have discovered; what activity could have created it or where it has come from, coined '*Trace-to-Activity/Source*' and reference to Cook et al.'s[3] hierarchy of propositions is made. Or second, a practitioner may examine a device following the belief that a relevant activity has occurred, making it necessary to try and establish whether traces denoting this activity exist – '*Activity-to-Trace*'. Both interpretative pathways require a practitioner to evaluate any hypotheses they hold regarding a trace's potential significance, either by conducting their own tests or by identifying relevant research that describes their target trace, then comparing the results

CONTACT Graeme Horsman graeme.horsman@cranfield.ac.uk
*The author is responsible for all work in this article.

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

of their work to the trace found in their case – a 'trace comparison'. Conducting a trace comparison is a fundamental part of the digital trace interpretation process and the outcome will both influence and form the basis of any subsequent opinion offered by a practitioner as to what a trace could signify.

This work describes both the Trace-to-Activity/Source and Activity-to-Trace interpretive journeys, the steps contained in both, and the role they play in DF examinations. In addition, six 'trace comparison criteria' are proposed to help those carrying out a trace comparison, notably: *'trace location', 'trace structure', 'trace examination method', 'trace metadata', 'trace content'*, and *'trace context'*.

2. Trace-to-Activity/Source

Trace-to-Activity/Source (TTtoAS) describes a scenario where a practitioner encounters a digital trace during their examination of a device which they consider to be of potential relevance to their overarching investigation, necessitating further inquiries into what that trace could actually mean. There are two elements here that require additional detail. First, a '*potentially relevant trace*' (PRT) is one that a practitioner feels may have an impact upon the outcome of a case and/or help those involved in an investigation to understand any alleged or applicable conduct to a greater degree. Typically, PRTs maintain characteristics that a practitioner would consider as '*of initial interest*' to them within the confines of their current investigative brief when they are initially discovered; at a surface level and on face value, they appear to provide something of worth to their inquiries. Some PRTs may be more obviously potentially relevant to an inquiry than others, for example, a picture depicting child sexual abuse will always be important to an investigation, whereas traces showing communication between certain parties may require additional deliberation as to their investigative value. PRTs may be found as a result of a practitioner manually reviewing data or via automated data-screening processes (for example, a keyword 'hit' to an indicative list of criteria; a hash match; a file within a location of interest). In all cases, a PRT maintains traits that when they are first identified, there is something that compels a practitioner to learn more about it.

Second, once a PRT has been found, in order to establish the actual potential value it has to their casework, a practitioner must begin to determine what they think that trace could mean; in essence, they must try to understand it; what it is, and why it is there. In the DF field practitioners will often form an investigative opinion of a PRT which will be derived from 'hypotheses (theories) that can account for their scientific observations'[4], made from within their current examination scenario and the data it contains. Whilst a practitioner may have to consider what any PRT is (in terms of type), possibly of more value to an investigation is the need to try and address how it got there (the 'activity' creating it) and/or where it came from (the 'source' of it)[3]. Answering these questions helps any wider investigative team to understand any behaviours occurring on the system and operations undertaken by a system-user, and ultimately what this will mean for their investigation. It also allows a practitioner to begin to determine a trace's investigative worth, and whether (and how) such a trace should be reported to their client.

Practitioners will routinely encounter TtoAS scenarios as part of their casework as in many instances they will be unaware of the contents of any device or dataset they are tasked to examine, requiring them to interpret the meaning of any PRTs as and when they encounter them. This interpretative process can be complex if

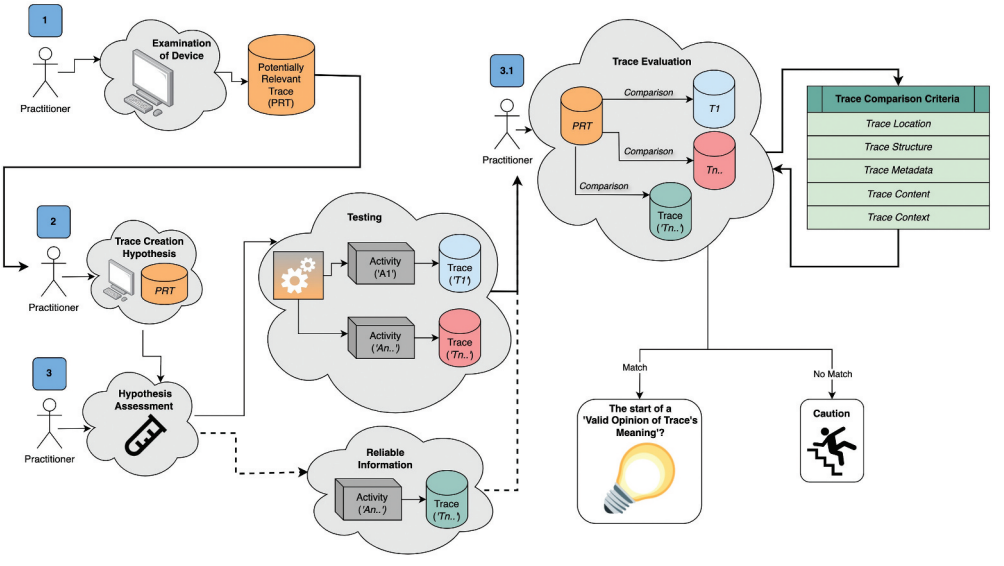

Figure 1. Stages of TtoAS interpretation.

a PRT is of a type that a practitioner does not understand at the time in which it is found (as opposed to a trace that is of a known type or one that is well understood within the forensic community)[5], requiring them to conduct additional interpretative work in order to ascertain their meaning. This is a process that this work suggests has four steps - *'Device Examination'*, *'Trace Creation Hypothesis'*, *'Hypothesis Assessment'* and '*Trace Evaluation*', shown in Figure 1, and discussed below.

2.1. Device examination

Prior to commencing any examination of a device, practitioners should be provided with information that appropriately outlines the need to undertake any DF work in the first instance, and the scope and remit of it[6]. This allows an examination strategy to be created which should define the processes and procedures required to address the needs of the investigation and to allow any investigative hypotheses that are held to be evaluated[7]. Practitioners may do this by deploying automated techniques that are designed to sort and sift data efficiently based on the use of specific set criteria which, if met, matching data is flagged for review. An example here would include the use of keyword searching techniques for terms that, if found to be in any files or data, such traces would be considered '*of interest*'. Alternatively, practitioners may manually and systematically traverse a given data set, reviewing data as they go. At any given point during an examination of case data, a practitioner may encounter a PRT, and when/if one is found, they must then proceed to try and understand what it is and validate any initial thoughts concerning its relevance to their work. In all DF examinations, multiple PRTs may be discovered.

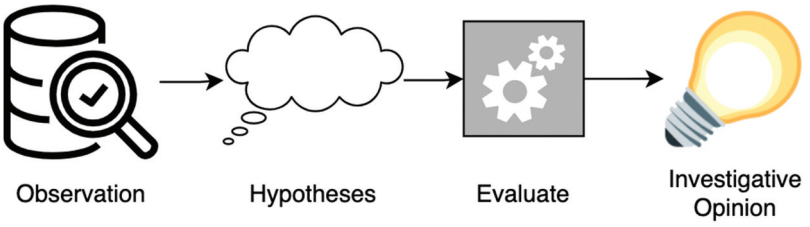

Figure 2. Stages of investigative opinion formation.

2.2. Trace Creation Hypothesis

When a PRT is identified, a practitioner should begin to assess why they have initially considered it to be potentially relevant and want to know more about it. To do this, they may generate one or more hypotheses with regard to what they think the trace is (or could be) and what its presence within their examination circumstances may mean, based on their observations. Any hypotheses must then be evaluated in order to determine whether they can form the basis of a valid opinion (Figure 2).

In many cases when a practitioner is engaged in a TtoAS interpretative task, the trace they have found and are seeking to understand is the result of past behaviour on a device. The focus therefore remains on understanding what actions *could* have led to the trace being present (the activity responsible for its creation), or determining where it came from (its source). This will typically involve practitioners defining a list of all of the activities that they reasonably consider could have been responsible for the presence of the PRT, or any potential sources of it. This list may be based upon a practitioner's knowledge, experience, research or from observations that have been made within the case itself. Circumstances commonly considered may include determining if a trace is the result of activities on a certain application or software package, for example, an artefact that has been created when a certain functionality has been engaged. The validity of any hypotheses generated about a trace must then be assessed.

2.3. Hypothesis assessment

All hypotheses generated regarding a PRT's potential meaning must be assessed in order to determine their legitimacy. Any activity that forms the basis of a hypothesis must be examined in order to determine the type of trace that it is capable of creating. Typically, a practitioner will begin to assess their hypothesis in one of two ways; either by trying to recreate the actions/sources they believe to be responsible for a PRT and evaluating any results generated (testing), or by accessing and evaluating information (existing research) which describes the processes they believe to have been responsible for the creation of a PRT.

(a) *Testing*: A practitioner may opt to test each of their hypothesized activities/sources (i.e. execute each activity within an appropriate environment) and observe and capture the traces that are generated as a result[8]. Any traces generated during testing can be considered '*synthetic traces*', and a practitioner must then proceed to determine if such synthetic traces are comparable to any PRT identified on a system (discussed below in 'trace evaluation'). This work will not explore the conditions

required for conducting robust and reliable testing within the DF domain, as this has already been covered in detail by Horsman[9] and Oliveira Jr et al.[10].

(b) *Reliable Information*: A practitioner may have access to information that documents how their hypothesized activity performs and the traces that it generates[11]. Such information must be reliable if a practitioner is going to base their interpretation of a PRT in their case upon it, requiring them to conduct an evaluation of it. When assessing the reliability of any information, a practitioner should consider both where the information has come from and who has created it. In addition, the level of evaluation and scrutiny that any information has been subject to (i.e. has it been subject to peer review) should be assessed, including whether there is general acceptance of the information within an appropriate community[12].

Any information source must also describe both the activity tested and the trace generated in sufficient detail for a subsequent evaluation of these actions against any PRT within the practitioner's case to take place. It is hoped that when reliable information sources are identified, any findings documented would have been as a result of appropriate testing being carried out by the information's author, and the findings are simply being made available for review.

Whichever route is taken, a practitioner must try to get to a position where they are able to make reliable observations and/or determinations as to how any hypothesized activities behave. A practitioner must have obtained sufficient detail to be able to conduct a trace evaluation. This entails performing a comparison between any synthetic traces generated from testing or details of traces documented by reliable information, and the PRT identified in their case.

2.4. Trace evaluation

Trace evaluation is a sub-task within trace assessment. Once a practitioner is able to ascertain and describe the traces created by their list of hypothesized activities, they must compare them to the PRT found in their casework. Trace evaluation is a crucial component of establishing the meaning of a trace and is a process that intends to determine whether any trace found during experimental work or research is comparable to a PRT in casework, allowing the practitioner to opine about what the PRT could mean. If any trace generated by a tested or researched activity is comparable to a PRT, a practitioner may proceed to offer an appropriate explanation of it. For example, if activity 'A1' is capable of producing a trace that is comparable to the PRT observed within their case, they may proceed to explain that their PRT could be a record of a system conducting activity A1 (see Figure 3). A real world example may include a trace being left by a web browser following visits to a webpage. If such a comparable trace is found within casework, a practitioner may opine that this trace is a record of web pages being visited on the device. A comparison of traces is therefore a crucial step of the trace interpretation process, and whilst comparing traces may seem straightforward, it is argued here that it is not. There are no documented methods or steps for conducting trace comparisons where it is suggested that this task must be structured around a series of metrics that allow a practitioner to accurately evaluate the similarity of both a synthetic or researched trace, and a PRT.

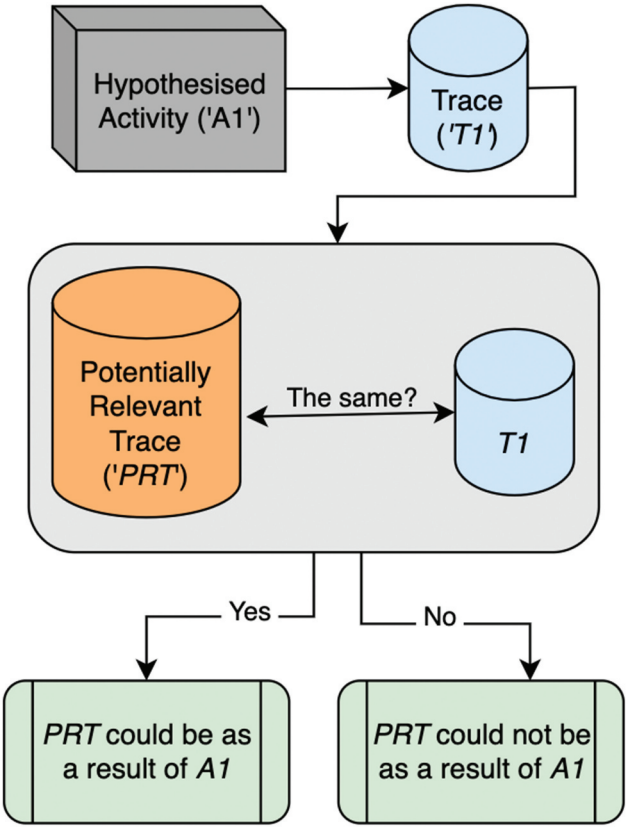

Figure 3. Trace evaluation steps.

Currently there is no formalisation of such metrics where this work proposes six – 'trace comparison criteria'- 'trace location', 'trace structure', 'trace examination method', 'trace metadata', 'trace content', and 'trace context'. Given the importance of this task, a discussion of the proposed trace comparison criteria is reserved for Section 2.2.

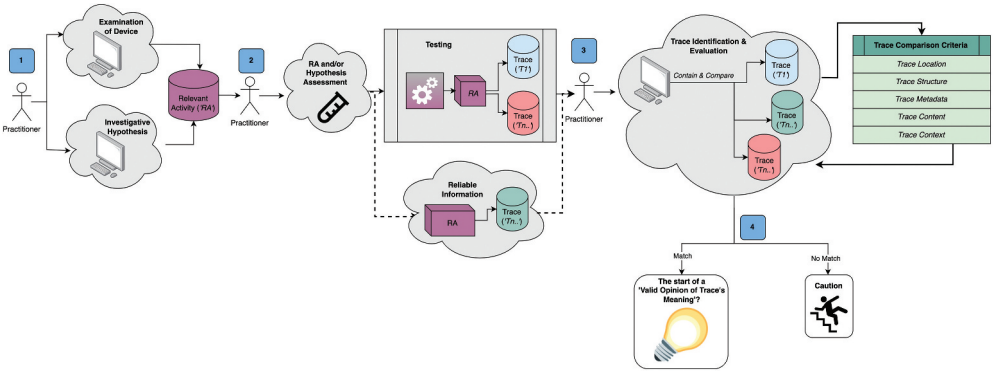

Figure 4. Stages of AtoT interpretation.

2.4.1. Activity to trace

In contrast to TtoAS, Activity-to-Trace (AtoT) interpretation (shown in Figure 4) occurs when a practitioner encounters one of the following scenarios.

(1) During the examination of a device/data, a practitioner finds information which indicates that an activity that is relevant to their investigation (a '*relevant activity*') may have taken place on a device. This invokes a requirement to establish if there are traces present in their examination that describe if this activity has actually been undertaken by a device-user (or another process) and any details that describe the extent of it. Consider the scenario where a file encryption application is discovered on a device, leading to a practitioner needing to determine if it has actually been used to encrypt data, and how. In these cases, a practitioner must first assess how any relevant activity performs and establish any traces that are left following the activity being carried out, then determine whether comparable '*traces of usage*' exist within their case.
(2) A practitioner may have or is supplied with a relevant investigative hypothesis regarding any alleged conduct which they must assess – i.e. it is believed that device '*D*' has been used to carry out relevant activity '*A*'. From the outset of an examination, a practitioner may be given instructions that indicate that a device has been used to do a certain task, resulting in a need to establish if such tasks have taken place.

AtoT scenarios are the converse of TtoAS, where instead of encountering PRTs which must be interpreted by working back to an activity, here, practitioners must interpret how any relevant activities function, and then establish whether any traces of these activities being performed exist within their case.

When engaged with AtoT, a practitioner must still assess any investigative hypotheses that they are given or have or relevant activities identified during their examination. When conducting this assessment practitioners have the same two options to do this; either via testing or through the use of reliable information that documents the functionalities of their relevant activity and the traces generated by it. Finally, a practitioner must determine whether any of the traces they have observed as a result of their testing or research regarding a relevant activity are present in their current case. This, as in TtoAS, requires trace evaluation and the comparison of synthetic/researched traces and any traces identified with a case.

2.3.2. 'Trace comparison criteria'

Regardless of whether practitioners are engaged in a TtoAS or AtoT interpretative task, both require a comparison to be made between any traces found within the data they are examining in their case and the traces that a practitioner has identified and assessed during their tests or research. The purpose of a trace comparison is to determine whether a practitioner can reliably recreate a trace that they have observed within their case or determine whether any case-trace can be deemed comparable to one that has been documented in research by someone else. If they can establish that any traces from these categories are comparable, they may be in a position to posit an explanation as to what any trace observed within their case may mean, and a practitioner's ability to offer an opinion of a trace's meaning hinges on this comparative task.

When simplified, a practitioner must try to determine if two given traces are fundamentally the same; however, when reference is made to the term '*same*', we do not mean that both are required to be a binary-level match; rather, they maintain a set of characteristics which only traces of that type should maintain. To conduct a trace comparison, there is a need for a set of criteria which a practitioner can utilize when trying to determine if any two traces are comparable and this work proposes six criteria.

(1) *Trace Location*: The location of a trace found within an examination and a synthetic/researched trace should be the same. For example, if, following testing, activity '*A*' creates a trace '*T*' at location '*L*', any trace found during a forensic examination that is the subject of a comparison should exist within an equivalent location to '*T*' if a practitioner is to begin to consider them comparable. If a discrepancy in location between both trace types is revealed, caution must be exercised as to whether the activity being tested/researched could have been responsible for the trace observed within an examination. When reference is made to a trace's location, this means logical file path, log-location, or in some cases, physical location on a disk. Whilst the terms logical file path and physical location are well understood in the DF domain, the term log-location refers to a trace embedded within an artefact where file path information alone would be insufficient for locating it. For example, consider a trace within a database; log-location would include comparing the table name where any trace was stored and information about any record/field it was contained within.

    The need for a comparable location between traces works on the assumption that computational activities often perform with a degree of consistency when observed under fixed conditions. This means that when a task is performed on a device, provided any operational parameters remain consistent, the traces generated by it often occur in the same way. It would therefore be expected, for example, that the same version of a web browser would log its activity in the same way and place on comparable systems, and any traces in these locations found during the examination of casework would have the potential to be as a result of this activity type. Of course there are always caveats, and system users may be able to manipulate the location of data (for example, change the default location of logs to be stored by a given application). In these cases, trace-comparison becomes a greater challenge, and emphasis upon the other trace comparison criteria stated here is increased.

(2) *Trace Structure*: Any two traces should be comparable in structure; their composition must be the same in terms of how any data is formatted and stored. As a basic example, any two traces may be stored as BPlists, or string data may be stored in a bespoke way or encoding-type. Discrepancies in structure would indicate that any two traces are not of the same type.

    Practitioners must pay attention to any activities being tested/researched, and particularly any variables that may have an impact on any traces generated, such as differing software versions. Where a mismatch of version types exists between any tested/researched software and the software installed in a case, the traces produced by both may differ in structure, despite the substance of both traces being comparable.

(3) *Trace Examination Method*: Both a case-trace and test/researched trace should be accessed and viewed using the same examination method. Simply put, the

practitioner must have to do the same thing to both traces in order to be able to view their data. If traces require the deployment of different techniques in order to access them, caution should be exercised as to whether they are comparable.

(4) *Trace Metadata*: Both a case-trace and a test/researched trace should maintain the same types of metadata. This may include examples such as both traces maintaining timestamps of a certain type/format, or naming conventions/formats. Both traces should maintain information that describes it, in the same way.

(5) *Trace Content*: Both a case-trace and a test/researched trace should contain the same type of data. For example, this may include specific text-types or data types that are comparable. At the most basic level, an example includes both traces containing data that shows a picture.

(6) *Trace Context*: Both a case-trace and a test/researched trace should exist upon a system/device within the same context. Here, reference is made to the files and data surrounding the presence of both traces. Often traces are bounded by additional data, tags, delimiters or are surrounded by other files. For example, a web browser may commonly store its history log file in a folder which also maintains multiple other log-types describing other parts of its functionality. Therefore, when a history log is found, a practitioner would also expect to see the additional log-types within the same folder. We consider this to be the context in which the trace exists.

2.3.2.1. Using the 'trace comparison criteria'. The trace comparison criteria described in Section 2.2 are proposed in order to try and support those conducting trace comparisons. Trace comparison is a task that almost all practitioners will undertake at some point in their roles, and many will be required to do this on a frequent basis. The outcome of a comparison directly impacts how a practitioner develops their understanding of a PRT in their case, and therefore it is of fundamental importance that they are able to do this effectively, as any failure may lead to a misinterpretation of it. Many digital traces are the product of complex activity within a device/system, and the task of not only trying to recreate the trace but also evaluating whether any recreated trace is comparable to what has been observed within a case-examination is difficult. It is argued that the introduction of the six trace comparison criteria would offer support to the practitioner when trying to determine if they can correlate their PRT to any test or research results.

When assessing whether any traces are comparable, it is suggested that any evidence of assessment should be aligned to the six criteria noted above. If a practitioner can show that their PRT is comparable to a tested/researched trace across all six of the criteria, then it is argued that they are in a strong position to speculate that their tested/researched activity could have been responsible for its creation. A practitioner also has evidence of a comparative assessment of traces being conducted in a structured and logical way, which offers objective evidence showing equivalence between traces. This is a core part of forming a robust investigative opinion of a trace. If practitioners are unable to observe conformity between traces across all six criteria, then caution should be exercised, and a practitioner must ask themselves why – casting doubt on whether their current information is sufficient for understanding a trace's meaning. The criteria themselves are as much intended to establish a match between traces as they are for determining and evidencing when and why two traces are not comparable.

Traces that show parity on some but not all of the criteria should not be immediately discarded, but clearly a practitioner must proceed tentatively when trying to consider if both traces mean the same thing at a fundamental level. It is noted above in Section 2.2 that some variables could differ between traces, despite the nature of the trace potentially being the same. We stress again with another example, consider the applications 'web browser v1' and 'web browser v1.1'. Both may create a trace for the same action-type which occurs at the same location on a system, have the same content, metadata, and context, but the updated version of the browser may generate a trace that is structured differently, requiring a different examination technique to access the information. Both traces document the same 'action'; however the traces cannot be said to be technically comparable, despite the traces being fundamentally comparable in the sense that both maintain a record of the same 'action' type. Such scenarios present an interpretative issue, and in all likelihood can be traced back to weaknesses in any deployed testing methodology or research by failing to examine a trace within a comparable environment to the one being observed within the practitioner's case itself.

The trace comparison criteria should be considered a tool to support practitioners when comparing traces as part of determining their meaning. By evidencing a structured approach to trace comparison, it is argued that any overall opinion of a trace's meaning is strengthened.

3. Conclusions

This work has discussed the two interpretive journeys that DF practitioners will typically embark upon, TTtoAS and AtoT. The steps required of both have been described in order to demonstrate how a practitioner will embark upon trying to establish the meaning of any PRTs they encounter in their casework. A fundamental part of both TTtoAS and AtoT is the need to compare traces; any PRT found and any trace generated via testing or established via research. This comparison is an important part of a practitioner's process of trying to determine if they can recreate the trace they seek to understand in their casework. To support this process, six trace comparison criteria have been proposed – 'trace location', 'trace structure', 'trace examination method', 'trace metadata', 'trace content', and 'trace context', and these have been described. In totality, this work is intended to support DF practitioners by helping to bolster their trace interpretative processes.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Graeme Horsman http://orcid.org/0000-0002-0685-0650

References

1. Servida F, Casey E. 2019. IoT forensic challenges and opportunities for digital traces. Digital Invest. 28:S22–S29. doi:10.1016/j.diin.2019.01.012.

2. Casey E. 2020. Standardization of forming and expressing preliminary evaluative opinions on digital evidence. Forensic Sci Int. 32:200888. doi:10.1016/j.fsidi.2019.200888.
3. Cook R, Evett IW, Jackson G, Jones PJ, Lambert JA. 1998. A hierarchy of propositions: deciding which level to address in casework. Sci; Justice. 38(4):231–239. doi:10.1016/S1355-0306(98)72117-3.
4. Forensic Science Regulator. 2021. Forensic science regulator codes of practice and conduct'. Available from: https://assets.publishing.service.gov.uk/media/602407728fa8f5146f0769d9/FSR-C-118_Interpretation_Appendix_Issue_1__002_.pdf
5. Horsman G. 2024. Interpreting digital traces: -8 foundational pillars to support the formation of opinion in digital forensics. Sci; Justice. 64(1):38–42. doi:10.1016/j.scijus.2023.11.007.
6. Horsman G, Sunde N. 2022. Unboxing the digital forensic investigation process. Sci; Justice. 62(2):171–180. doi:10.1016/j.scijus.2022.01.002.
7. Horsman G. 2023. Digital evidence strategies for digital forensic science examinations. Sci; Justice. 63(1):116–126. doi:10.1016/j.scijus.2022.11.004.
8. Hughes N, Karabiyik U. 2020. Towards reliable digital forensics investigations through measurement science. Wiley Interdiscip Rev. 2(4):e1367. doi:10.1002/wfs2.1367.
9. Horsman G. 2018. Framework for reliable experimental design (FRED): a research framework to ensure the dependable interpretation of digital data for digital forensics. Comput Secur. 73:294–306. doi:10.1016/j.cose.2017.11.009.
10. OliveiraJr E, Zorzo AF, Neu CV. 2020. Towards a conceptual model for promoting digital forensics experiments. Forensic Sci Int. 35:301014. doi:10.1016/j.fsidi.2020.301014.
11. Horsman G. 2019. Formalising investigative decision making in digital forensics: proposing the digital evidence reporting and decision support (DERDS) framework. Digital Invest. 28:146–151. doi:10.1016/j.diin.2019.01.007.
12. Arshad H, Jantan AB, Abiodun OI. 2018. Digital forensics: review of issues in scientific validation of digital evidence. J Inf Process Syst. 14(2):346–376.

https://dspace.org/

Understanding and comparing digital traces

Horsman, Graeme

2025-08
Taylor & Francis
https://doi.org/10.1080/00450618.2024.2381535
Horsman G. (2025) Understanding and comparing digital traces. Australian Journal of Forensic Sciences, Volume 57, Issue 4, August 2025, pp. 481-491

Downloaded from a DSpace repository. https://dspace.org/