

# THE USE OF ARTIFICIAL INTELLIGENCE IN DIGITAL FORENSICS: AN INTRODUCTION

By Dr Faye Mitchell

**Artificial Intelligence (AI) is an important and well established area of modern computer science that can often provide a means of tackling computationally large or complex problems in a realistic time-frame. Digital forensics is an area that is becoming increasingly important in computing and often requires the intelligent analysis of large amounts of complex data. It would therefore seem that AI is an ideal approach to deal with many of the problems that currently exist in digital forensics. The purpose of this paper is to give a high level introduction to AI as it might be used in digital forensics.**

## Introduction

Duce and others<sup>1</sup> outline what might be considered, from both an academic and a practitioner point of view, to be three of the main challenges in digital forensics: (1) the exponential growth in storage capacity, in single drives (hard drives, USB sticks, optical media); (2) the growth in distributed systems and the sophisticated forms of attack that can now be launched; (3) the degree of technical sophistication employed by opponents and the apparent inability of existing tools and methodologies to keep pace.<sup>2</sup> To that a fourth challenge might now reasonably add: (4) the ubiquity of electronic storage and the range and prevalence of disparate storage systems.

The requirements needed to solve these challenges might be stated, in a greatly simplified form, as the *ability to reason and discover over a large amount of complex, potentially disparate, data in a realistic time frame*. The conventional intensive and manual

approaches currently used to search for data are not capable of dealing with the size of the problem currently found in digital forensics. It is for this reason that it is becoming apparent that a more selective and intelligent approach is needed for digital forensic analysis. There is a branch of computer science that tries to tackle this type of problem – the branch known as Artificial Intelligence (AI). This paper concentrates on considering AI generally, and how it might be applied to the challenges that face digital forensics.

## Artificial Intelligence

Defining Artificial Intelligence is not simple. There is no one clear definition of AI. Most of the definitions that do exist tend to define AI in terms of “creating a computer process that acts intelligently” (but what is intelligence?) or “creating a computer process that can mimic human behaviour” (do humans always act intelligently, what happens if a computer can normally perform better than a human being?). Other definitions refer to “rational behaviour” (but what is rational?) or “doing things that are hard for a computer to do” (does this mean that when an AI system has been developed to do the task, it is no longer AI?) and are equally unhelpful in this discussion. Therefore in order to simplify the task for the purposes of this article, a pragmatic approach is adopted, and AI is defined as “creating a computer process that acts in a manner that an ordinary person would deem intelligent”, and consideration is given to some of the various types of AI and AI technologies that might be of concern to people in the digital forensics community. Consideration will not, therefore, be given to techniques such as robotics, which at present have no direct relevance to digital

<sup>1</sup> D. A. Duce, F. R. Mitchell and P. Turner, ‘Digital Forensics: Challenges and Opportunities’, in John Haggerty and Madjid Merabti, (eds.), ACSF 2007:

Proceedings of the 2nd Conference on Advances in Computer Security and Forensics, (Liverpool John Moores University, School of Computing &

Mathematical Sciences, 2007).  
<sup>2</sup> NIST Computer Forensics Tool Testing Project at <http://www.cftt.nist.gov/>.

forensics; other material that may potentially be useful will also not be covered, including related fields such as data visualization. This article will concentrate on the conventional areas, and interested readers are recommended to consult Luger, or Russell and Norvig as a suitable starting point for more detailed discussion of AI and other possibly relevant branches of AI.<sup>3</sup>

### Representation of knowledge

The most important concept in the majority of AI systems is the representation of knowledge (referred to in AI as *knowledge representation*) and ontology. That is, how to represent the information we wish to reason about (representation of knowledge) and how we formally structure that representation of knowledge in such a way that we can reason about it (ontology). It is important to note that our representation of knowledge can be about the properties of objects in the domain (information), how those facts can be processed (knowledge about what rules and techniques to apply in a particular situation) or even how those processes are applied (strategic or meta knowledge).

In the early days of AI, ontology was not considered an issue and a new representation of knowledge was created for each application. However, in the last ten years, there has been a realisation that being able to reason over multiple sources of knowledge is very important. This has focused interest in producing ontologies for domains that can be shared amongst applications and systems. For the most part, this has focused on XML, RDF<sup>4</sup> and related technologies, although other notations such as ontolingua<sup>5</sup> are also occasionally used. It is perhaps here that AI has the potential to have the most effect on digital forensics, in providing expertise to help with the standardisation of the representation of knowledge and information in the digital forensic domain. This lack of standards hinders the exchange of information for even the most basic of tasks in digital forensics, such as the exchange of image information between forensic imaging tools,<sup>6</sup> and this

unfortunately means that digital forensics is behind the accepted good practice in many other scientific domains where there has been a concerted effort to produce a standard domain ontology.

The creation of standardised international domain ontology for digital forensics would have obvious benefits in, for instance, a multi-national case covering a number of jurisdictions, in that it would provide a formal framework for the discussion of digital evidence, but it would also provide other benefits in that it would enable the creation of a large, re-usable case repository.<sup>7</sup> Such a case repository would contain known, sanitised examples of digital forensic investigations with known properties and results. This could be useful in testing the performance of experts, be they human or AI systems, and could provide a useful method for training digital forensic practitioners, and has proven extremely valuable in other areas of AI.<sup>8</sup> The use of a standardised ontology could also prove invaluable in creating a standard, reusable collection of background knowledge<sup>9</sup> that could be used by AI techniques.

### Explaining the reasoning process

An important issue for AI in the forensic arena is the ability to explain the reasoning process. That is, the ability of the AI technique or algorithm used to explain the reasoning process. AI techniques are often divided into two categories: symbolic (those that reason with discrete entities in a knowledge base) and sub symbolic (those where the knowledge is spread around the representation structure). One of the most common types of symbolic reasoning is the expert system. Expert systems follow a predefined rule base,<sup>10</sup> and normally have a limited strategy for choosing which rule to use at any particular moment in time. Expert systems can, therefore, at any point, provide an explanation of the reasoning for the conclusions obtained. This enables an outside entity to criticise the reasoning process and to highlight any flaws there might be with the reasoning

<sup>3</sup> George F. Luger, *Artificial Intelligence: Structures and Strategies for Complex Problem Solving* (6th edition, 2009 Addison-Wesley); Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (3rd edition, 2010, Prentice Hall).

<sup>4</sup> W3C (2004) *RDF Primer* at <http://www.w3.org/TR/rdfl-primer/>; W3C (2006) *Extensible Markup Language (XML) 1.1 (Second Edition)* at <http://www.w3.org/TR/xml11/>.

<sup>5</sup> Adam Farquhar, Richard Fikes and James Rice, 'The Ontolingua Server: a Tool for Collaborative Ontology Construction', *International Journal of Human-Computer Studies*, 1997, Volume 46, pp 707-727.

<sup>6</sup> Philip Turner, 'Unification of digital evidence from disparate sources (Digital Evidence Bags)', *Digital Investigation* (2005) 2(3), pp 223-228.

<sup>7</sup> D. A. Duce, F. R. Mitchell and P. Turner, 'Digital Forensics: Challenges and Opportunities', in John Haggerty and Madjid Merabti, (eds), *ACSF 2007: Proceedings of the 2nd Conference on Advances in Computer Security and Forensics*, (Liverpool John Moores University, School of Computing & Mathematical Sciences, 2007).

<sup>8</sup> The UCI Machine Learning Repository (<http://archive.ics.uci.edu/ml/>) is an example of such a case repository, and is used by the Machine Learning community to test new algorithms.

<sup>9</sup> Background knowledge is the term given to knowledge about a domain that is often common sense, and often extremely large (e.g. If I throw a ball in the air it will normally come down; this windows file is normally found in this position in the directory tree). AI systems can be set up to use this knowledge to help their reasoning processes.

<sup>10</sup> A rule base is essentially an ordered collection of rules where each rule is in the form IF  $\langle \dots \text{antecedent set} \dots \rangle$  THEN  $\langle \dots \text{consequents} \dots \rangle$ . The rule base can use certainty factors, probability and fuzzy sets where the domain has to deal with vagueness and uncertainty.

used. However, systems that exhibit the property of the ability to explain the reasoning process often have two major drawbacks.

The first of these drawbacks is that they operate in a closed world.<sup>11</sup> That is, if it is not in the rule base, then it does not exist or get taken into consideration. This can be a serious issue in an area such as computing, where the technology changes at an extremely rapid rate. Rebuilding a rule base is known to be a time consuming task, and adding additional rules (a processes known as rule base repair) can damage the original performance and result in rules that would have previously worked, but no longer function.<sup>12</sup> The second drawback is that expert systems do not cope well with large quantities of data. This is a particularly major disadvantage for the direct use of expert systems in digital forensic investigations where the amount of data investigated is becoming larger and larger, and is increasing at an almost exponential rate.

Where techniques such as expert systems might prove to be useful, however, is in higher order situations such as guiding an investigator on what to try next, or to advise on what the policy of the organisation is in a given situation.

## Cases

Case Based Reasoners (CBRs) are a type of (normally) symbolic AI that are an attempt to avoid some of the problems associated with symbolic rule based systems such as expert systems. CBRs are based on well understood notions from psychology on how domain experts themselves represent information.<sup>13</sup> Most domain experts rely heavily on their past experiences, and when faced with a problem, will attempt to match the problem to one they have experienced before. Only when an expert has exhausted all possible similar cases in their experience do they use first principles to attempt to find a solution to the problem.

A CBR system works in a similar fashion, in that a large collection of cases (and in digital forensics, the resultant actions) is obtained, and a metric<sup>14</sup> is used to

match the current situation with one found in the case base. If a perfect match is found, then the action carried out in the initial case is applied to the existing situation. If no perfect match is found, but a match is found that is deemed to be close enough, then the system may attempt to adapt the action of the matched case to the current situation using what are called 'repair' rules.<sup>15</sup>

CBR systems have the advantage of approaching a problem in a way that is familiar to the expert, can cope with large amounts of data, and can deal with situations that have not previously been encountered. They address in part the ability to explain the reasoning process, because the reasoning can be inspected (this case was most like X, and in X you did Y). This, however, means that the user might rely very heavily on the quality of the cases in the case base, together with a good coverage of the possible scenarios. CBRs are also limited, in that although they can help guide the process of the investigation, they are perhaps ill suited to helping to automate the lower level activities (such as "find all pictures with naked people in them").

## Pattern recognition

Identifying specific types or clusters of data in an investigation is best handled by a type of AI known as pattern recognition. The type of pattern recognition that people are most familiar with is perhaps image recognition, where the software attempts to identify parts of a picture. Other forms of pattern and image recognition also exist, such as detecting a pattern in an e-mail message which indicates SPAM, or a pattern in a disk image that might indicate it is part of a sound file. Many of the techniques used rely very heavily on statistics or probabilistic reasoning or both. The more complex and accurate forms of image recognition that might be used to locate certain types of picture, rely on an understanding of how the human perceptual system works. However, at present these have a high rate of false positives or false negatives (depending on where the thresholds are set) as well as being very computationally intensive.

<sup>11</sup> This is also an issue with respect to the data available. If the expert system is not provided with all the necessary information available, the output may not be reliable.

<sup>12</sup> There is a technique known as 'Knowledge Base Refinement' which can help automate the rule base changes, but even that can still result in breaking the rule base unless steps are taken. This is discussed further under 'Knowledge Refinement', below.

<sup>13</sup> F. R. Mitchell, 'An introduction to Knowledge Acquisition', School of Computing and Mathematical Sciences, Oxford Brookes University Technical Report (1998, CMS-TR-98-06; also

published as F. Mitchell, 'An introduction to Knowledge Acquisition', (1998, AUCS/TR9804, Department of Computing Science, University of Aberdeen Technical Report; although not within the context of AI, others have discussed the same issues relating to experts: Peter M. Bednar, Vasilios Katos and Cheryl Hennel, 'On the complexity of collaborative cyber crime investigations', Digital Evidence and Electronic Signature Law Review, 6 (2009) pp 241-219.

<sup>14</sup> In AI, the term 'metric' is used to mean any system of measurement by which items may be compared. For instance a 'similarity metric' measures how similar two items are or a distance metric measures

the distance (under some notion of distance) between two items.

<sup>15</sup> These repair rules tell the system in what ways and in what order a rule can be changed. For instance, if the original case specified a hard disc, but the particular instance was about a USB stick, the CBR system might reason that they both contain writeable file systems, so the rule could be used in this situation. However, if the particular instance was about a CD-R, then it might indicate that because this system is not writeable to, this rule can never be made to apply.

Pattern recognition systems are essentially classifiers – that is they answer the question: is this piece of data a member of the class X, where X is the type of data the user is interested in. In order to work successfully, pattern recognition techniques have therefore to try to match against all possible pieces of data (or as near as is computationally feasible) which can involve a large amount of matches, and the patterns have to have sufficient generality to match all positive matches but sufficient specificity to not match any of the negative examples. In practice, this is often very hard to achieve, although Machine Learning techniques (for which see below) can help with the generality or specificity problem by allowing patterns to adapt, and in the case of certain systems, such as Artificial Neural Nets or decisions trees, can be used to learn the initial patterns.

### Knowledge discovery

Another field of AI that might have benefit in the forensic arena is Data Mining and Knowledge Discovery in Databases (Datasets). Although these terms technically refer to different things, the two terms are colloquially used interchangeably to refer to process of finding useful information in a large collection (normally sparse) of data. Data Mining/Knowledge Discovery in Databases (DM/KDD) is not a single technique but is a mixture of AI, statistical analysis and probabilistic techniques used together in an integrated manner to analyse large collections of data. It can be viewed as a form of pattern recognition, but with a few significant differences.

First, the sheer size of the data (in some cases petabytes) means that more computationally intensive techniques cannot be effectively used, therefore any AI technique involving the use of a complex knowledge representation is unlikely to be used for DM/KDD. Similarly, background knowledge about the domain may also not be used, or may only be used in a limited fashion.

Secondly, DM/KDD is often directed by the user. Technically this process is a form of Exploratory Data Analysis (EDA) where the user asks the system to, for instance, highlight files with characteristic X, and the system uses Data Visualisation (DV) to highlight

information and potential relationships to the user. This is particularly useful, because the human perceptual system has the ability to distinguish patterns in extremely complex data – even in data with a large number of attributes<sup>16</sup> – if the data can be represented properly. Care does, however, have to be taken, because the human perceptual system can find patterns that do not, in reality, exist.

Thirdly, DM/KDD has the concept of an interestingness measure (often called a J measure) that helps to decide whether there are any meaningful patterns in the data. This helps avoid the situation where a DM/KDD system ‘discovers’ the extremely obvious, but extremely unhelpful fact, such as that you only ever find female patients in the maternity ward of a hospital.

It is extremely likely then that, given the increase in quantities of data, the forensics community will have to rely on DM/KDD techniques to help with the initial assessment. To date they are the best AI method for dealing with large quantities of data, but they are also potentially the most likely to miss relevant pieces of information, because the reasoning processes do not normally use the background knowledge or complex reasoning of more complex AI approaches.

### Adaptation

A system that has a fixed knowledge source is unlikely to be able to cope well with the change in pace in computer technology, and therefore it is likely that some measure of adaptability will be required for any long term forensic system. The branch of AI that deals with the ability of the software of a system to adapt is called Machine Learning (ML). From the point of view of interest in the applicability to digital forensics, ML techniques can be divided into two: ones that use ML as a method of trying to refine the knowledge source<sup>17</sup> to keep it current (the refiners), and those who use ML to gather the initial knowledge (the learners). There are, of course, techniques which combine both approaches, but they can be thought of as a subset of the learners. Each type can in turn be divided into supervised (a human, called an oracle, gives the correct answer) and unsupervised (the system is left to find out its own

<sup>16</sup> In AI an ‘attribute’ is a dimension of the ‘problem space’. Although the terms ‘attribute’ and ‘dimension’ are sometimes used interchangeably in AI, the preferred usage is to refer to the dimensions of a problem, rather than the attribute of an object. However, for those readers that might not be familiar with the terms used in AI, the more familiar term ‘attribute’ is used here rather than the more accurate AI term ‘dimension’. Contrary

to popular belief, a standard two dimensional computer screen can display considerably more than two attributes in a comprehensible manner. In some cases, through the use of complex shapes, the amount of attributes displayed can easily reach double figures. Interested readers are referred to Vitaly Friedman, ‘Data Visualization: Modern Approaches’ Smashing Magazine, 2 August 2007, at

<http://www.smashingmagazine.com/2007/08/02/data-visualization-modern-approaches/> for some examples of modern visualisation techniques.

<sup>17</sup> The term ‘knowledge source’ is used to mean any place that the system can obtain knowledge from. This may be the knowledge base in a symbolic system or the network of weights in a sub symbolic system.

answers).

### Refinement of knowledge

The ‘refiners’ are often symbolic systems where an oracle has created the original knowledge source which has become outdated. Rather than try and manually patch the knowledge source (which in a complex knowledge base can be very error prone), it is possible to direct the software to automatically refine the knowledge base, so that repairs can be made which are consistent with the existing knowledge base and are guaranteed not to break ‘chestnut’ cases.<sup>18</sup> The problem with this approach is that it is not possible to be sure that the knowledge base is rendered useless by the refinement.

Also, many of these techniques do not consider performance, so repeated refinements can result in a knowledge base that does not work effectively. When dealing with even simple knowledge bases, this becomes an issue when the data to be reasoned about is as large as it is in a digital forensics case.

### Machine Learning

Often it is possible to say that object X is an example of type Y, but not necessarily why, or at least not in a way that can be used easily to assign meaningful attributes to a symbolic concept (e.g. you might be able to tell by taste whether or not a whisky was a malt whisky, but not be able to say exactly what about the taste made it a malt whisky). In such a situation, it is still however possible to use an AI system to learn what the concept is by using a learning system. Such systems normally rely on the use of training sets which contain pre-classified examples which, along with the algorithm, form the basis of the learning system. The success or failure of the learning will depend on the power and suitability of the learning algorithm and the quality of the data set used.

### Sub symbolic learners

Sub symbolic systems are perhaps one of the more successful forms of learning that can deal with previously unseen data. In sub symbolic systems, the information about a concept is not stored in a particular

part of the system but is instead spread around the system as a network of nodes and links connecting those nodes. The most common type of sub symbolic system is an Artificial Neural Network (ANN) or a variant thereof. ANNs take the concept of the brain as a collection of synapses where the information is stored in the links between the synapses (this is of course a great over simplification, but will do for the purposes of the discussion) and represent it in a simplified computer model. A set of inputs (example cases) is repeatedly given to the ANN, and the links between the parts of the model are slowly adjusted until the actual output is the one that is desired. This means that ANNs are capable of learning. Unfortunately this usefulness comes at a price. ANNs are fast when they have learnt what something is, but can be very slow to learn, and what it learnt is not in normally in a form that can be inspected or used elsewhere in any meaningful manner. This means that although ANNs might be useful to help identify items in an investigation which can have their relevance or importance verified by a human expert (such as initial investigation of a disk), they are perhaps not best suited to situations such as guiding the course of an investigation where the investigator might be required to explain their reasoning. There are also hybrid techniques such as ANNs combined with decision trees. However, to date none really address the ability to explain the reasoning process without losing the learning power of the ANN. A related problem is that it is necessary to be careful that they learn what the investigator wants them to learn, and not some unrelated feature.<sup>19</sup> There exists other forms of sub symbolic learner such as genetic algorithms or emergent behaviour systems – however they all exhibit similar properties of being good at learning concepts which are not easily definable, and not having the ability to explain the reasoning process.

### Symbolic learners

Sub symbolic learners are not the only type of learner, and there are two common types of symbolic learner that have the potential to help with a forensic investigation: decision trees and observational learners.

A decision tree at its most basic can be thought of as learning the rules for pattern classification or possibly

<sup>18</sup> Chestnut cases were a concept developed in the early days of knowledge refinement (Susan Crow and D. Sleeman, ‘Automating the Refinement of Knowledge-Based Systems’, in: L. C. Aiello (ed), *Proceedings of the Ninth European Conference on Artificial Intelligence (ECAI90)* (Pitman, Stockholm, Sweden, 1990), pp 167-172.) and are a set of cases

that the user has decided the knowledge base must be able to solve correctly.

<sup>19</sup> This is best illustrated in a (possibly apocryphal) story about the US military, who tried to train an ANN to recognise tanks hiding in trees. To this end they took pictures of forests with no tanks, pictures of forests with tanks and showed them to

the ANN. The pictures without tanks were taken on a cloudy day, and the pictures with tanks were taken on a sunny day, so the ANN learnt how to tell if it was sunny or not. Because an ANN has no ability to explain the reasoning process, this fact was not found out until much later in the testing process.

for use in an expert system. Decision trees use information theoretic and other similar measures to decide on the minimum set of decisions to classify an entity as belonging to one of the given sets. In order to do this, they must be given a set of entities and told what the attributes of a particular entity are, and to what class that entity belongs. As such it is potentially useful for many aspects of the digital forensic analysis from the high level, such as indicating what type of investigative process should be undertaken, to locating examples of a particular type of file. They also, because of their rule based nature, are very good at explaining the reasoning process. However, a serious limitation is that the information must be categorized and labelled first, often in some detail by a human expert, and their ability to generalise the rules to deal with similar unseen situations is often not as good as can be found in sub symbolic approaches.

One type of learner that has the potential to help with guiding the course of an investigation is the observational or apprentice learner systems (ALS). An ALS is based on the idea of human apprenticeships, in that the ALS is assigned a human expert. The ALS observes what the human does, and when the software is designed to conclude when there is enough information to form a rule based on observations of the human action. The software then forms the rule, and poses the question "I think when X happens you do Y". The expert then agrees or criticises the rule, and the knowledge base is then updated accordingly. If the ALS records the human doing something that contradicts the rules it has learnt, the software will also ask the human for guidance. Although this can become annoying for the human expert, it can prove a quick way of automating basic tasks, encoding best practice and advising new practitioners on what company practice is, for instance.

## Conclusion

This paper has provided an introduction to some of the basic AI techniques. This is by no means an exhaustive summary, and there are many other possible AI techniques that might be applied to the domain of digital forensics that there is insufficient space to

discuss here (for instance, Support Vector Machines might be used to reduce the amount of attributes of the problem that you need to consider, agent based systems could be used to utilize distributed resources better, conceptual clustering could be used to help identify important areas in the evidence under investigation, and so on). Unfortunately, most of the work on applying AI to digital forensics is still at a very early stage and can be split into two areas – (1) where AI is used to help automate an individual part of the forensic process (e.g. to look for a particular file type) and (2) where AI is used to guide the expert in their task.

One example of the first area is a recent system called MADIK. Using a combination of expert system and agent based CBR system, Hoelz and others<sup>20</sup> have produced a collection of co-operating programs that can advise on whether or not a particular file should be investigated by the expert. However, the capabilities of the agents (small autonomous programs) are limited to a few simple tasks and the guidance of the investigation is only at a basic level. Garfinkel<sup>21</sup> also suggests a system for helping to automate parts of the forensic process, the discovery of related information across multiple hard disks using data mining techniques. The techniques used by Garfinkel are simple correlation and lexographic techniques, with none of the more powerful clustering and pattern recognition techniques that are used in other areas of data mining, nor is there any use of a J measure. However the results produced are promising.

With respect to our second area, that of guiding the expert, most of the work focuses on expert systems and the need to create a suitable ontology to describe the forensic process. Blackwell<sup>22</sup> uses an expert system for his proposed framework on evidence management, but the suggested framework is also a very simple one and there is no indication of how the rule base could be created. Turner<sup>23</sup> discusses the use of Digital Evidence Bags that can record the process of an investigation. This could be used to start to create a partial ontology for digital evidence, which could be used in wider digital forensics ontology, but no work has been done so far. Similarly Duce and others<sup>24</sup> suggest ideas on how an

<sup>20</sup> Bruno W. P. Hoelz, Célia G. Ralha and Rajiv Geeverghese, *Artificial intelligence applied to computer forensics in Proceedings of the 2009 ACM symposium on Applied Computing, Honolulu, Hawaii, (ACM, 2009)*, pp 883-888.

<sup>21</sup> Simon L. Garfinkel, 'Forensic feature extraction and cross-drive analysis', *Digital Investigation*, 3S (2006), pp 571-81

<sup>22</sup> Clive Blackwell, 'Managing evidence with an expert system', 2nd Workshop of the AI in forensics SIG, Cybersecurity KTN (London 2009), available at [http://www.ktn.qinetiq-tim.net/resources.php?page=rs\\_ktnpublications](http://www.ktn.qinetiq-tim.net/resources.php?page=rs_ktnpublications).

<sup>23</sup> Philip Turner, 'Unification of digital evidence from disparate sources (Digital Evidence Bags)', *Digital Investigation* (2005) 2(3), pp 223-228.

<sup>24</sup> D. A. Duce, F. R. Mitchell and P. Turner, 'Digital Forensics: Challenges and Opportunities', in John Haggerty and Madjid Merabti, (eds.), *ACSF 2007: Proceedings of the 2nd Conference on Advances in Computer Security and Forensics, (Liverpool John Moores University, School of Computing & Mathematical Sciences, 2007)*.

ontology used in cultural heritage might be used in digital forensics, but does not take the idea any further forward.

As this brief survey illustrates, the use of AI in digital forensics is still at a very early stage, but it does have a lot to offer the digital forensics community. In the short term it is likely that it can be immediately effective by the use of more complex pattern recognition and data mining techniques, as has discussed by Garfinkel. However, for digital forensics to take full advantage of what AI has to offer, more work is necessary. First, a suitable ontology must be produced for digital forensics, so that it is easy to record, reason about and exchange information about the evidence and processes used. This will help, both in terms of automating the digital forensic investigation and in terms of helping to record best practice in digital forensics. Secondly, this ontology needs to be used to annotate suitable cases that can be shared with both digital forensics experts

and AI experts. This collection of cases can provide help in bench marking both people and computer systems as well as opening up the opportunities for mainstream AI experts to help advance the use of AI in digital forensics.

© Dr Faye Mitchell, 2010

Dr Faye Mitchell is a senior lecturer at Oxford Brookes University. Her research focuses on applying machine learning, knowledge discovery and knowledge acquisition to real world problems, particularly in the area of computer security and digital forensics.

<http://tech.brookes.ac.uk/staff/atoz/data?username=frmitchell>

[frmitchell@brookes.ac.uk](mailto:frmitchell@brookes.ac.uk)