AI Reading, or Automatic Semantic Decomposition into Knowledge Graphs and Symbolic reasoning through Marker Passing

Prof. Dr. Johannes Fähndrich[1] Dr. Frank Trollmann[2]

Abstract: Marker passing algorithms have been applied to solve problems in artificial intelligence related to the semantic of written words. Such approaches could also prove to be useful in digital forensics, e.g., to reduce the effort of extracting evidence from confiscated data. We call creating semantic or even pragmatic understanding of text: AI reading. With this, we show that the aggregation of knowledge out of heterogeneous information sources can be a combination of symbolic and connectionist approaches. With that, the extraction of knowledge graphs can be automated. This approach has the benefit, when used correctly, that both the creation and the use of the knowledge graph through Marker Passing stay explainable. In this paper we describe a tool chain of a Marker Passing approach from the point of view of digital forensics and discuss challenges and opportunities arising from the application of such an approach.

Keywords: Semantic Decomposition; Digital Investigation; Marker Passing.

1 Introduction

Understanding language is a basic ability of humans, separating us from other animals. Language is developed in the first years of our life but is continually learned over the entire life [Bl98]. Understanding language is also an important part of artificial intelligence (AI). Here the first question which arises is the meaning of understanding [Gr68]. The application of methods of AI to digital forensic is not a new idea e.g. [CDO19] approaches digital forensic with answer set programming, and [TR14] analyze supervised multi class discrimination. Language understanding can also be of utility to the law enforcement community. Sighting large amounts of textual data for relevant details can require considerable human resources, especially as the amount of seized data increases [QC14]. However, using AI for this purpose also comes with significant challenges [Vi16].

In the scope of AI, we have worked on an approach to combine Demantic Decomposition and Marker Passing [Fä18] to solve tasks that require language understanding. This approach has been to domain like semantic descriptions of software agents [FWA16] and solving Winograd schemes [FWK18]. We argue that this approach will also prove useful in the scope of digital forensics.

[1] Hochschule für Polizei Baden-Württemberg, Germany, JohannesFaehndrich@hfpol-bw.de
[2] CODE University of Applied Sciences, frank.trollmann@code.berlin

In this paper we consider this approach from the point of view of digital forensics to point out possible applications and challenges for future work that arise from this domain. Among these challenges is the requirement to provide a traceable reasoning chain that can be understood and verified by a human law enforcement expert.

The paper is divided into two parts. The first part (cf. Section 2), describes the components of the approach and illustrates. The second part (cf. Section 3) analyzes this algorithm from the point of view of digital forensics and discuss strengths and weaknesses to derive future research avenues. Section 4 concludes the paper.

2 Marker Passing

This section describes the Marker Passing algorithm and its four main components. While we use examples from the domain of law enforcement as illustrations, we keep the descriptions independent of this domain. These four components will be picked up in Section 3 where we focus more closely on the requirements in digital forensics.

Marker Passing is an algorithm that is able to answer a question about an information source. As an example, the information could be all data contained on a persons hard drive and the question could be "which communication did a person have over which channels regarding a certain topic". Usually, answering the question isn't easily possible, e.g., because the data spread over multiple files with different formats. To tackle this problem, the Marker Passing algorithm employs a graph-based representation of this text, called a Semantic Graph. This representation is generated from the original information source via a process we call Semantic Decomposition. The Semantic Graph represents important entities from the original data source as well as their relationships. The relations in this Semantic Graph may extend beyond the relations captured by the original text. They may, for example, relate names that occur in multiple data sources.

Marker Passing operates on this Semantic Graph. It propagates data (so called markers) along the graph to detect connections between pieces of information. The algorithm does this by placing markers with information onto the graph and propagating these markers along its edges. The resulting placement of markers can be interpreted to answer the original question.

This algorithm can be subdivided into four components:

1. *Semantic Graph*: A graph structure that is appropriate to capture the entities and relationships relevant in the application domain.
2. *Semantic Decomposition*: An algorithm that generates a Semantic Graph out of information sources.
3. *Marker Passing*: An algorithm that propagates a marking along the edges of the Semantic Graph.

4. *Result Interpretation*: An algorithm that interprets the result of the Marker Passing to answer the original question.

These four components are discussed in more detail in the following subsections.

2.1 Semantic Graph

The first step of adapting the Marker Passing algorithm to a new domain is defining the Semantic Graph. The goal is to capture entities that are relevant to the application domain and relationships that are relevant to the types of questions that should be answered. Traditionally, a graph consists of nodes edges, where each edge connects two nodes. In the example of communication, we could choose to model topics of interest and persons as nodes and connect them via edges whenever we discover a message about a topic that is sent from one person to another. As discussed in more detail in [Fä18], a more complex graph structure is usually needed to represent AI Reading problems. For this reason, the following extensions have been made:

- **n-ary edges**: can connect more than two nodes. They can represent relations between multiple entities (e.g., a message from a sender to a recipient about a topic).
- **edges connecting edges**: can point towards other edges. This enables the representation of meta-information. E.g., we can use these edges to keep track of which file a message edge has been extracted from.
- **node attributes**: can represent nodes with additional attributes. E.g., a person can be annotated with first and last name.
- **typing mechanism**: can be used to treat nodes and edges differently based on which kind of information they represent. E.g., we may decide to treat nodes representing persons differently than nodes representing topics.
- **type inheritance**: can be used to represent similarities on a type level. E.g., we can express a hierarchy of topics such as ä sexual assault is a specific type of assault".

Figure 1 shows an example of a semantic graph modeling the offense of “sexual assaultäs 3-ary relation. It connects an “offenderäs the source, a “victimäs the target and a “weaponäs the object of the sexual assault. This is enabled by the typing concept in our Semantic Graph. The example also illustrates edges connecting edges. The edge “is a"represents the fact that a “sexual assaultïs a special type of “assault".

The application of the Marker Passing algorithm to a new domain requires the definition of a type graph. This requires a list of node and edge types and their possible attributes. This is a one-time effort: once this type graph has been defined it can be used to run the algorithm on all problems with a similar structure. It is also often possible to use and adapt existing type graphs for similar problems, as was done, e.g., in [FWK18].

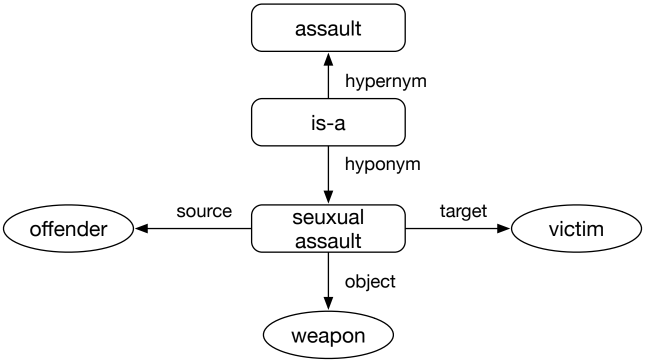

Fig. 1: Example of a typed graph.

2.2 Semantic Decomposition

Semantic Decomposition is the process of extracting the Semantic Graph from the original information source. The idea of breaking complex concepts down into less complex concepts is a well known solution strategy in explaining complex concepts.

While the algorithm for decomposition can be custom-made for specific data sources and file formats, a lot of the data available on a computer will be in natural language. For this type of data a variety of syntactic and semantic approaches already exist and can be integrated into such an algorithm. Riemer [Ri15] describes a syntactical decomposition of parts. For example, sentences can be decomposed into less complex syntactical structures. A first approach on a Semantic Decomposition algorithm has been introduced in [FAA14] and was detailed in [FWK18]. This algorithm has been improved and made available to the community[3].

Fig. 2 depicts a simplified version of the graph created from different data sources. Beyond the original data files, Semantic Decomposition can include additional online data sources, e.g., a dictionary to better understand words and their synonyms [FAA15]. With these additional data sources (here called dictionaries), it is possible to establish relations that go beyond uses of the same word in multiple files. The decomposition can take custom ontologies or domain specific databases like mobile device forensics, special information retrieval or structured data sources into account, which makes it customizable for most contexts. The domain DB in Fig. 2 is an example of typical information extracted from confiscated phones.

[3] `https://github.com/Datenverlust/SemanticDecomposition` Please contact the first author if the access is still not public on the date of publication. Access will be granted.

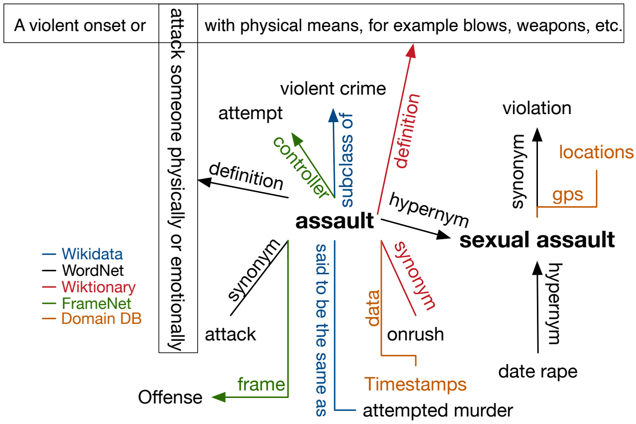

Fig. 2: Example of a Semantic Graph created by a decomposition.

2.3 Marker Passing Algorithm

The Marker Passing algorithm has been published in [FWA16] and has been detailed in [Fä18]. It generalizes the idea of activation spreading, like it is used in neuronal networks, to more structured information. Marker Passing has been used in NLP taks as a model of thought e.g. [Me19] With that we want to achieve an amalgamation of connectionist and symbolic reasoning like formulated in the vision of Charniak [Ch86a; Ch86b; DF; He89]. An implementation has been made available to the community [4].

The information (e.g., a number, a location, a Person) is represented by so-called “markers"that are associated to the nodes of the graph. They denote the attention of the Marker Passing algorithm and annotate additional information. This marking is not static. It changes over time according to the rules of the Marker Passing algorithm.

The initial placement of markers and the Marker Passing behavior are problem dependent. The initial placement of markers is based on the question asked. E.g., if our question is to establish who knew about a sexual assault, we could place markers on nodes representing this topic.

The Marker Passing algorithm moves these markers through the Semantic Graph. The rules according to which markers move are problem specific. This has, e.g., been applied to implement a semantic similarity measure [FWA16]. In our conversation example, the algorithm may spread markers from the topic of sexual assault to all persons that knew about it according to chat conversations to identify the offender or possible victims.

[4] `https://github.com/Datenverlust/MarkerPassingAlgorithm`

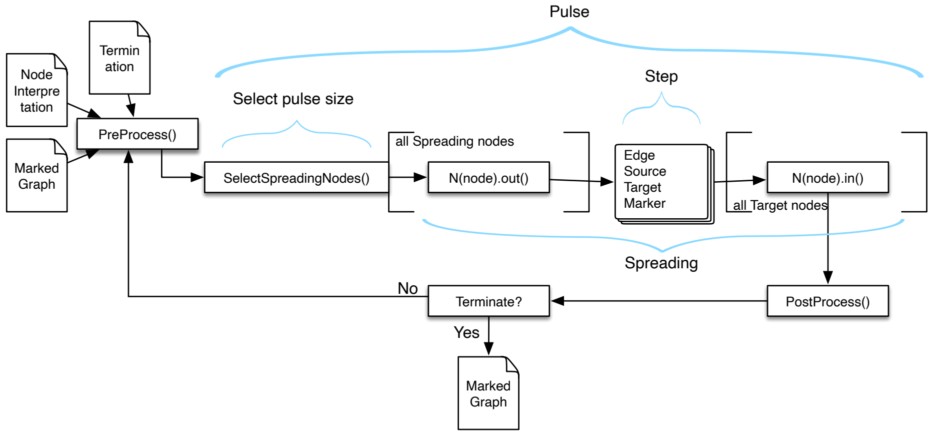

Fig. 3: Abstract description of our Marker Passing algorithm.

Figure 3 shows an abstract representation of the algorithm. The input of the algorithm consists of a marked graph, a node interpretation and a termination condition. These serve the following purposes:

Marked Graph This is a Semantic Graph as described in Section 2.1. It is annotated with markers that represent the initial marking.

Node Interpretation The Node Interpretation describes the Marker Passing behavior of specific node types. It consists of a set of concept interpretations (representing the note types). Each node type defines how it processes markers passed to it, when it is activated, and how it passes markers to other concepts. These functions allow each concept type to react to different markers differently by specifying.

Termination condition This condition defines the stopping point of the spreading. If no condition is given, the activation spreads until no concept is activated. This only accomplishes the goal of Marker Passing in some cases [Be09].

The output of the Marker Passing algorithm is a marked graph. The structure of the Semantic Graph does not change during the Marker Passing, thus the output graph differs from the input graph only in the location and value of markers.

The Marker Passing algorithm is divided into a sequence of **pulses**. Each pulse activates some concepts and spreads their activation to neighbors. A common visualization is to imagine the activated concepts to light up and create a pulsing light which “wanders"through the graph. This visualization gives the Marker Passing pulse its name.

The pulse has two intermediary steps: the **Pre-** and **Post-Processing**. They can be used to integrate general tasks such as cleaning up or normalizing the activation.

For each pulse, the pulsing concepts are selected during **Pulse Size Selection**. This step selects some activated concepts to spread their activation. While the intuition of the reader may be to pulse all activated concepts at once, this turns out to be impractical in reality. The ability to select a smaller pulse size enables us to avoid situations in which the termination condition becomes fulfilled by activation of one concept and then negated by activation of another concept. Furthermore, it enables more fine-grained control in cases where the order of activation influences the result.

After selecting these concepts, the **spreading** of activation takes place. This step consists of three stages:

Out-function All pulsing concepts determine which of the connected relations to spread information over and which markers they should receive by using their node interpretation. These markers are then forwarded to the respective relations.

Spreading Step The relations take these incoming markers and process them. They also select which of the connected concepts receive which markers. The markers are then forwarded to these concepts.

In-function The receiving concepts get their respective markers via their in-function. This function determines how to integrate the new markers with the already placed markers. Again, the node interpretation has an influence on this step.

Confiscated data rarely contains relevant facts directly. With the large amount of data collected, the data triage needs to be efficient but reduce false negatives. In our running example, Marker Passing could extract a Semantic Graph on which information has been known to whom (e.g., by spreading over email and messaging protocols) and when (e.g., via the associated timestamps). Since the graph will include the same concepts if they appear in different contexts, the marker can reveal connections between information sources.

The interpretation of the result Result is still a manual task. The output of the Marker Passing algorithm is the final position of the propagated markers in the Semantic Graph. The goal of the result interpretation is to lift this representation back to the level of the original question.

Since the Semantic Graph and the algorithm have been defined with this type of question in mind, this step can be trivial - e.g., it may involve reading the result directly from the markers. However, more complex methods - e.g., making statements about the distribution of markers or their movement during the algorithm, are also possible and are highly dependent on the question to be answered.

3 Applications in Digital Forensics

In this section we discuss possible uses of Marker Passing in the field of digital forensics and highlight challenges and extensions derived from these uses. We start with an example of using AI reading to process mass data. With that we want to highlight design decisions in the parameters of the Marker Passing. Some of these parameters can also be learned experimentally and are subject to fine-tuning.

The application to digital forensics starts with the creation of a format for the Semantic Graph and an algorithm for the Semantic Decomposition. Previous work on Semantic Decomposition as been based on public OSINT dictionaries. This work needs to be extended to include data found in private police databases. This customization can go further: The case file and the collected data in one case make a good starting point. Depending on the investigated crime, specific data sources can be integrated e.g. car registration, password or username lists, social networks, cellphone GIS data or traffic information. The reasoning capabilities of the Marker Passing are heavily influenced by the information included in this stage.

The Marker Passing needs to be configured to suit the reasoning use case. While this parameter selection is context dependent, we argue for a few rules to keep the marker passing deterministic and explainable.

Marked Graph We argue to keep the graph structure constant while the Marker Passing is running. This assures that the graph structure is only influenced by the initial data sources and can be traced back to them. [5]

Node Interpretation Node interpretation gives the Marker Passing the ability to change passing behavior depending on the node type. E.g. secured facts can have another passing behavior then non-secured facts. Another node dependent interpretation might be dependent on the data source the node came from[6].

Termination condition The termination condition should be deterministic. Here an optimization needs to be found, when the markers should stop passing. This stop on the activation prevents markers to move away from relevant concepts.

Out-function The out-function should not be probabilistic. It should always be reproducible which markers it passes, given a specific marking of the concept.

SelectSpreadingNotes The selection of the spreading nodes can be modeled in many ways. Some of these options (e.g. inspired of natural observation, like in neuronal networks) are modelled with e.g. inhibition or decay over time. Theses are dependent on system

[5] If relations are found that where not in the graph, one might be tempted to include these new relations into the graph: e.g. a ownership relation of a piece of evidence.

[6] e.g. the information in a DMV might be more trustworthy than some testimony out of a case file.

time, thus depend on the underlying machine and its performance, which makes the results only reproducible in a controlled environment. To be explainable, such mechanisms should be avoided.

In-function The in-function should be deterministic, meaning the resulting marking of the concept should only be dependent on the previous marking and the newly passed markers.

We ignored the pre- and post-processing function, because they are thought for clean up and special (out of the algorithm) modification of the graph and the markers. However, these functions should be written in a way that they don't manipulate the relevant information in the graph.

One important concern in law enforcement is traceability - the ability of humans to understand the results of the algorithm and to trace it back to the original data sources. An algorithm that does not fulfill this requirement would not be usable as it hurdles police and court proceedings to provide conclusive evidence.

In our Marker Passing algorithm, this requirement for traceability means that we need to be able to trace the answer back to the original data sources and how they are processed by the algorithm. This means, all components of the algorithm need to support traceability. In specific (working backward from the answer) the following requirements need to be fulfilled by the respective components:

Result Interpretation: the process of creating the answer from the final marking of the Semantic Graph needs to be traceable. It needs to be clear, which of the markers present in the network were involved in creating the answer, and how so.

Marker Passing: The path of the markers through the Semantic Graph needs to be traceable. Even if the algorithm itself is deterministic, it is not generally possible to understand the path of tokens from the final marking. Traceability of Marker Passing requires a recording of the marking history, either by having tokens keep track of their path or by recording the states of the activation spreading and enabling a replay. Depending on the use case, more specialized methods could be implemented that only keep track of information relevant for the end result. In addition to keeping track of the history, visualization techniques may play an important role in enabling an investigator to understand the evolution of the marking over time.

Semantic Decomposition: Any concept in the Semantic Graph needs to be traceable to the original data sources. This can be guaranteed by a Semantic Decomposition that keeps track of meta-data regarding the origin of its concepts (e.g., pointer to file and line number). This needs to be possible both for concepts and relations.

Semantic Graph: The Semantic Graph is static and thus there is no need for traceability. However, it needs to provide the appropriate edges, concepts and data types for keeping track of the above-mentioned traceability information.

If these requirements are fulfilled, it is possible to trace back the answer to the final marking of the graph, understand which concepts / relations have been involved in producing it and find the respective information in their original files. This makes it possible to understand and validate the reasoning taken by the algorithm.

Other properties that should be discussed are correctness and completeness to assess the reliability of the approach if it is to be used in law enforcement. Correctness means the answer of the algorithm is correct (e.g., a social network connection pointed out by the algorithm is indeed a provable connection). Completeness means the algorithm detects all answers that could be given (e.g., that it finds all connections contained in an email dump). This can be measured in false positives and false negatives. Both correctness and completeness are dependent on the specific implementation of the algorithm, and their importance may also depend on the specific use case.

4 Conclusion

In conclusion, the definition of a Semantic Graph enables us to describe knowledge representations in a connectionist way. The Semantic Decomposition algorithm can be used to create such graphs from text sources like Wikipedia or WordNet. Additional information sources can be integrated by defining new dictionaries to extract concepts and relations. The application of the Semantic Decomposition can be manyfold.

Based on such a graph, the Marker Passing algorithm allows a symbolic reasoning on such a graph. With the interpretation of the resulting markers, many graph based algorithms can be simulated (PetriNets, Artificial Neuronal Networks, or other state transition systems). The idea of this combination of connectionist graphs and symbolic reasoning could be applied to many challenging tasks. The mechanism realizes the fuzzy understanding of text while reading, thus the title of this paper. Utilizing text to extract concepts and relations, and then perform reasoning upon those extracted concepts and relations, could be seen as reading.

To become explainable, two things need to be parametrized in this approach:

- Decomposition with source specific concepts and relations
- Marker passing with history on the nodes or the markers.

By specifying where concepts and relations come from, the text, which has been used to create the Semantic Graph, can be traced back and analyzed. With the history, the concepts and relations that have participated in the algorithm can be reconstructed and the reasoning of the algorithm can be traced. The result of this approach thereby is fully explainable, if certain conditions hold. E.g. the graph should not be modified by the Marker Passing algorithm. Or the decomposition is done before the Marker Passing is started.

The challenge of bringer methods of AI to digital forensics is that the approach needs to be validated on real world date, before it can be applied to any real data. Additionally, the inner working, the selected parameters for the marker passing and the used data sources need to be discussed with judges and prosecutors so that the resulting graph including the markers can be interpreted.

Our future work will include a modeling of "focus"to investigate in grate text data like confiscated files, which are, because of their size, no longer processable by humans. The data is used to create graphs connecting concepts with relations, and the Marker Passing algorithm needs to be parametrized to "focusön relevant data. Additionally we have to validate the approach on real world data, with help of investigators to select the parameters accordingly.

References

[Be09] Berthold, M. R.; Brandes, U.; Kötter, T.; Mader, M.; Nagel, U.; Thiel, K.: Pure spreading activation is pointless. In: Proceedings of the 18th ACM conference on Information and knowledge management. Pp. 1915–1918, 2009.

[Bl98] Bloom, L.: Language acquisition in its developmental context./, 1998.

[CDO19] Costantini, S.; De Gasperis, G.; Olivieri, R.: Digital forensics and investigations meet artificial intelligence. Annals of Mathematics and Artificial Intelligence 86/1, pp. 193–229, 2019.

[Ch86a] Charniak, E.: A Neat Theory of Marker Passing. Association for the Advancement of Artificial Intelligence/, pp. 584–588, 1986.

[Ch86b] Charniak, E.: A Neat Theory of Marker Passing. In: AAAI. Pp. 584–588, 1986.

[DF] Dyer, M. G.; Flowers, M.: Integrating Marker Passing and Connectionism for Handling Conceptual and Structural Ambiguities./.

[Fä18] Fähndrich, J.: Semantic decomposition and marker passing in an artificial representation of meaning, Doctoral Thesis, Berlin: Technische Universität Berlin, 2018.

[FAA14] Fähndrich, J.; Ahrndt, S.; Albayrak, S.: Formal Language Decomposition into Semantic Primes. ADCAIJ: Advances in distributed Computing and AI Journal/, 2014.

[FAA15] Fähndrich, J.; Ahrndt, S.; Albayrak, S.: Self-Explanation through Semantic Annotation and (automated) Ontology Creation: A Survey. In: 10th International Symposium Advances in Artificial Intelligence and Applications. 2015.

[FWA16] Fähndrich, J.; Weber, S.; Ahrndt, S.: Design and Use of a Semantic Similarity Measure for Interoperability Among Agents. In (Klusch, M.; Unland, R.; Shehory, O.; Pokahr, A.; Ahrndt, S., eds.): Multiagent System Technologies: 14th German Conference, MATES 2016, Klagenfurt, Österreich, September 27-30, 2016. Proceedings. Springer International Publishing, Cham, pp. 41–57, 2016, ISBN: 978-3-319-45889-2.

[FWK18] Fähndrich, J.; Weber, S.; Kanthak, H.: A Marker Passing Approach to Winograd Schemas. In (Ichise, R.; Lecue, F.; Kawamura, T.; Zhao, D.; Muggleton, S.; Kozaki, K., eds.): Semantic Technology. Springer International Publishing, Cham, pp. 165–181, 2018, ISBN: 978-3-030-04284-4.

[Gr68] Grice, H. P.: Utterans meaning, sentence-meaning, and word-meaning. In: Philosophy, Language, and Artificial Intelligence. Springer, pp. 49–66, 1968.

[He89] Hendler, J. A.: Marker-passing over microfeatures: Towards a hybrid symbolic/-connectionist model. Cognitive Science 13/1, pp. 79–106, 1989.

[Me19] Meyer, R.: Computerlinguistische Ansätze zur Repräsentation und Verarbeitung von Wortbedeutungen. De Gruyter, 2019.

[QC14] Quick, D.; Choo, K.-K. R.: Impacts of increasing volume of digital forensic data: A survey and future research challenges. Digital Investigation 11/4, pp. 273–294, 2014.

[Ri15] Riemer, N.: The Routledge Handbook of Semantics. Routledge, 2015, ISBN: 1317412451.

[TR14] Tallón-Ballesteros, A. J.; Riquelme, J. C.: Data mining methods applied to a digital forensics task for supervised machine learning. In: Computational Intelligence in Digital Forensics: Forensic Investigation and Applications. Springer, pp. 413–428, 2014.

[Vi16] Vincze, E. A.: Challenges in digital forensics. Police Practice and Research 17/2, pp. 183–194, 2016.