Distributed Hypothesis Generation and Evaluation

Jordan I. Robinson

Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, UK
J.Robinson9@liverpool.ac.uk

Abstract. Inferring the potential outcome of ambiguous events using acquired intelligence is a challenging task, even for the most seasoned investigator. Analysts must use abductive and defeasible reasoning in combination with a variety of analytical techniques to assess hypotheses for situations as they materialise. Analysis is a time-consuming task and requires multidisciplinary teams of subject-matter experts to draw conclusions. This research focuses on the development of explainable Artificial Intelligence approaches which aid the task of hypothesis generation and evaluation in investigative settings. Techniques from computational argumentation will be used to help analysts forecast potential future events by identifying the sets of collectively acceptable arguments from the available information. Then, the evaluated arguments will be assessed using probability theory to quantify the relative uncertainty in all the possible hypotheses. The outcome of this research should help the intelligence, legal and business communities by providing decision-support tools which increase the accuracy and speed of decision-making.

Keywords: Computational Argumentation · Decision-making · Explainable AI · Hypothesis Generation and Evaluation.

1 Introduction

With any investigation, an analyst must use the information available to them to conclude what has or could happen. Although investigators do not know the amount of evidence they will receive during analysis, they are required to draw conclusions about available data anyway. For example, an analyst might receive a report from an agent with one piece of information in it, and that could be the only evidence available to them when conducting analyses and drawing inferences. On the other hand, an analyst may have access to a large number of documents where the pertinent information is sparsely populated throughout. In both instances, analysts must find the critical information required to make well-founded judgements during analysis. Analysts must work with data which is often: distributed – i.e., analysts can receive a variety of reports from agents who could be situated in different places; incomplete – that is to say, analysts do not know if they have all the relevant information when reporting it; conflicting – as in, an agent might report information which contradicts another person's version of events; and possibly deceptive – in other words, a government or an adversary may intentionally try to mislead analysts so that the inferences drawn from data are incorrect (Operation Fortitude is a good example of this [22]).

Identifying arguments within natural language corpora is a laborious task which requires a high level of analytical acumen, especially if the most important arguments are scattered throughout a number of different documents. After identifying and making sense of the evidence, investigators must understand the impact that each piece of evidence has on the set of mutually exclusive and comprehensively exhaustive (MECE) hypotheses. Here, the term *hypothesis* is used in a broad sense to mean a potential outcome, explanation, or conclusion made from available information. A good hypothesis is defined by Heuer and Pherson in [24] below:

- It is posed as a definite statement, not a question;
- It is based on observed evidence, facts and knowledge;
- It is testable and can be proven wrong;
- It forecasts anticipated events correctly.

Investigators in the intelligence, law enforcement and business communities use analytical techniques to help them assess situations which could be detrimental to national security, public safety, or even a business' market capitalisation [24]. Analytical techniques encourage methodical and deliberate reasoning amongst analysts in a way which eliminates many of the cognitive biases and inadequacies of intuitive judgements. Although using these techniques makes analysis less complicated, they dramatically increase the time-taken to conduct it. In critical situations, analysts may not be time rich so automating parts of the process will benefit investigators who need to make accurate decisions quickly.

As analysis progresses, new information may become available which can invalidate an investigator's prior conclusions. Acquired evidence in real-world analysis is dynamic; thus analysts are required to reason defeasibly – i.e., the conclusions made by an investigator may change in light of new evidence.

The focus of this research project is to develop explainable Artificial Intelligence (AI) techniques and algorithms which identify the relative probability of all the MECE hypotheses for ambiguous scenarios, implemented through computational argumentation and probability theory.

This paper is organised as follows. In Section 2, we describe this research project's problem setting and define the aim to address this. Section 3 provides an overview of computational argumentation, analytical techniques and probability theory used in investigative analysis. Section 4 briefly details the proposed methodology for this project and Section 5 describes the current state. Section 6 concludes with a discussion regarding the expected contribution of this research.

2 Problem Setting and Research Aim

Trying to forecast the outcome of ambiguous events is an extremely challenging, detailed, and time-consuming task which requires interdisciplinary teams of experts and investigators to collate and assimilate distributed, incomplete, conflicting, and potentially deceptive data. The outcomes from analyses are then used to draw conclusions, which can ultimately affect, for example, business decisions, legal cases, or international relations.

The aim of this research project is to design, develop and evaluate explainable AI approaches for hypothesis generation and evaluation in investigative settings. The proposed approach will use techniques from computational argumentation and probability theory to automate parts of the analytical process and help analysts perform logical and defeasible reasoning when assessing complex situations. Techniques from computational models of argument will be used to find critical pieces of evidence (or arguments) in natural language corpora. The extracted arguments, and the relations between them, will then be used to evaluate the sets of collectively acceptable arguments within the debate, aiding the generation of the set of MECE hypotheses. Then, the evidence will be assessed using probability theory to quantify the relative uncertainty in all the hypotheses. The proposed implementation should be able to update prior probabilities when new information becomes available. This should result in decision-support tools which aid analysts in making well-founded judgements about available data, increasing the accuracy and speed of analyses. The techniques developed will be applicable across multiple fields which require explainable decision-making, such as in the intelligence, legal, and business communities.

3 Overview on state-of-the-art

Argumentation Theory is a well established field of research which concerns how argumentation-based reasoning is used to resolve disputes and conflicts of interest [20]. It has its roots in logic, philosophy, and law. Arguments form the basis of argumentation. An argument consists of a claim (or conclusion) which is supported by one or more reasons (or premises) to believe it [37].

In Section 1, we saw that investigations require defeasible reasoning as new conclusions can be drawn when analysts acquire new information from distributed sources. Argumentation is also defeasible which is why a wealth of propositions are asserted during a debate [18]. For example, if we think that another person's assertions are strong enough to defeat our argumentation, we will often provide more reasons to accept our initial claim (unless one does not have any further justifications for why the other party should believe their standpoint), yielding a dialogue through which arguments are exchanged.

Computational argumentation is considered a relatively new but increasingly important study within the AI community [8]. The field of computational models of argument combines Argumentation Theory and computer science to formalise and automate defeasible reasoning through modelling the generation and evaluation of arguments, counterarguments, and the addition of new information in a debate [6]. The rapid adoption of argumentation techniques within the field of AI can be attributed to their effectiveness in modelling decision-making [17]. Over recent years, there has been much research on knowledge representation, defeasible reasoning and multi-agent systems within the field of AI – see Pollock [32], Simari and Loui [34], and Dung [16] for a few examples of pioneering works. The different computational models of reasoning proposed in the literature focus on three types of argumentation: rhetorical, dialogical and monological models [11]. Rhetorical argumentation concerns the persuasive intention of a speaker to an audience. Dialogical argumentation studies how agents exchange arguments and counterarguments in a debate by simulating how they interact, change their beliefs, retract arguments, etc. Monological argumentation concerns the internal structure of arguments by modelling how a set of arguments and counterarguments might be generated from a knowledge-base, evaluated to find which arguments are acceptable, and used to draw conclusions which could inform decision-making [12]. Computational argumentation can also be loosely divided into two categories: abstract and structured argumentation.

Abstract argumentation has its foundations in a seminal paper by Dung which considers the relations between arguments as entities that have no internal structure [16]. Dung's argumentation framework (DAF) can model and evaluate the attack relations between arguments in a debate through computation of semantics, described in [7], which calculate the sets of acceptable arguments in a framework. Although DAF's are extremely powerful reasoners, much subsequent work has been completed to enrich their expressiveness. Dung's work has been extended to capture social preferences and rational disagreement in abstract argumentation frameworks (AFs) through, for example, preference-based AFs [2], value-based AFs [9], and the bipolar AF (BAF) which models the notion of support as well as attack between arguments [15]. Further research has been completed to extend BAFs to incorporate evidential [30], deductive [13] and necessary [29] support. Attempts have also been made to extend Dung's framework using probability theory by assigning strengths to arguments and their attacks [25]. One of the challenges within investigative analyses and abstract argumentation concerns dealing with dynamic information – i.e., as new information becomes available during an investigation, this may alter our evaluation of the evidence. The conditioned AF is an approach that deals with additional information well by using a division-based method to cope with dynamic information in the framework [26]. The different types of argumentation systems and semantics proposed in the literature could allow investigative analysts to conduct logical and defeasible reasoning in a consistent manner while also allowing them to evaluate information in a rich variety of ways.

Structured argumentation differs from above as it concerns the internal structure of arguments. The goal of structured argumentation is to extract the components of arguments by identifying stereotypical patterns of reasoning [36]. A structured argument can be split into three parts: a set of premises, a conclusion, and an inference from the premises to the conclusion [35]. *Argumentation Schemes* are an example of one of the methods used in structured argumentation [38]. Although these schemes are typically used in monological models, their critical questions can be employed to further dialogical models by identifying attack relations between instantiated arguments; thus aiding in the abstraction of arguments to abstract AFs for evaluating the acceptability of arguments [3,4].

Turning to hypothetical reasoning, this is a common practice in investigative settings. Attempts have been made to model hypothetical reasoning using structured argumentation in [10] and [21]. In intelligence settings, various analytical techniques are employed to force analysts to consider the following question: *"What would the evidence have to look like for a particular hypothesis to be true?"* Foresight techniques use a similar type of hypothetical reasoning to make predictions about what could happen in the future – Pherson and Donner provide a good example of this in [31]. To extend hypothetical reasoning to abstract argumentation, we could consider choosing the set of collectively acceptable arguments in a framework before evaluation, and use this to predict the required attack relations between arguments for the acceptable set to hold – i.e., we could try to model how each piece of evidence would have to relate for a particular hypothesis to be true.

Now consider argumentation mining (AM) which is a task that involves the automated identification and extraction of arguments from natural language corpora. The process can be viewed as a pipeline. An AM system will take a set of documents as input, and output the argumentative components and their support and attack relations. AM is a rapidly-developing field within AI and natural language processing – refer to [27] and [14] for a more detailed overview.

Within the intelligence, law enforcement and business communities, many Structured Analytical Techniques are used to assess information so that informed decisions can be made [24]. The evaluative techniques used in these analyses emanate from Karl Popper's philosophy of science, where he argues that science should attempt to disprove theory, rather than attempt to continually support theoretical hypotheses [33]. A commonly used technique for hypothesis evaluation is the Analysis of Competing Hypotheses (ACH) which uses strict falsification to test hypotheses [23]. It is interesting to note the similarity between ACH and abstract argumentation. For example, an argument is said to attack another in an AF if that argument rebuts or undercuts another argument within the framework. During ACH, a piece of evidence attacks a hypothesis if that evidence is inconsistent with that explanation of events.

In 2015, Murukannaiah et al. published the only study to date which measures the benefit of combining structured argumentation with investigative analytical techniques. They combined argumentation schemes with ACH, calling the approach Arg-ACH [28]. Their experiment found that participants performed better analysis using Arg-ACH compared to participants that used ACH only. The study provides an initial example of the benefits when using structured argumentation to evaluate complex hypotheses. It's worth noting that each participant in the experiment did not have any previous experience of argumentation or analytical techniques before it began. One would expect better results when extending this study to include experts who are familiar with the techniques being used.

Very few attempts have been made to use probability theory to quantify the relative uncertainty in hypotheses or assess the impact of new information when forecasting the outcome of events. In [19] and [5], a Bayesian approach to predictive analyses is adopted where the likelihoods of all the MECE hypotheses are updated in light of new evidence. Investigators may only have access to a few pieces of evidence, especially when analysing unprecedented situations such as terrorist attacks, making this approach a powerful one when dealing with novel situations which have a limited amount of data.

4 Proposed Methodology

To successfully complete this research, we will investigate explainable AI approaches which emulate parts of the analytical processes used during investigation.

The first part of the investigative pipeline requires analysts to gather information and make sense of it. This is usually the most laborious part of analysis as analysts must find the most important information from potentially copious amounts of data. If an analyst has a lapse in concentration at this stage and fails to notice a critical piece of information, then the outcome of their analysis will be wrong. This part of the pipeline could be automated using AM which would help analysts by extracting all the arguments from the reports received from different agents.

After this, analysts need to generate the set of MECE hypotheses in light of the evidence at hand. This is also an important stage in the analytical process. If a critical hypothesis is overlooked here, then it will not factor into evaluation. There are many different ways to do this. One consideration might be to instantiate an abstract AF using the predicted support and attack relations generated from the AM task. Then, this framework could be evaluated using preferred extension semantics. This would provide analysts with the maximal sets of acceptable arguments within the debate and help them when generating the set of MECE hypotheses.

Finally, we need to evaluate the set of hypotheses and infer the likelihoods of every potential event. In the intelligence community, ACH is the tool of choice. However, ACH only provides the least inconsistent hypotheses which require further attention and does not provide the relative likelihoods of each hypothesis. Attempts have been made using Bayes' theorem to quantify the change in likelihood of a hypothesis using the available information. Attempts have also been made to reason hypothetically using structured argumentation. For this research project to be a success, there needs to be investigation into methods which combine computational argumentation with probability theory to quantify the relative uncertainties in all of the MECE hypotheses for a situation.

4.1 Measuring the Utility of Argumentation in Hypothesis Generation and Evaluation

A technique or tool is only as useful as how it compares to a trained professional. Measuring the benefit of argumentation in investigative settings is necessary for the success of this project. This motivates the following question to be considered by this research: *To what extent does the use of argumentation benefit MECE hypothesis generation and evaluation?* This question can sensibly be answered by conducting an experiment that measures this benefit. Participants in the studies will be familiar with argumentation and analytical techniques prior to the study commencing. Participants taking part in the study will be divided into five groups based on the techniques they're familiar with. They will be divided as follows:

1. Seasoned argument annotators;
2. Participants familiar with investigative analytical techniques;
3. Participants familiar with the techniques mentioned in (1) and (2);
4. Participants unfamiliar with the techniques mentioned in (1) and (2);
5. Experienced investigators.

By splitting participants into the aforementioned five groups, we will establish whether the analysis conducted by analysts with a background in argumentation, but no prior knowledge of analytical techniques, can emulate that of a seasoned investigator.

A further study comparing the output of experienced investigators with fully automated approaches, using a combination of the techniques discussed in Section 3, will be conducted when the initial study concludes.

5 Current Work

A manual AM analysis of two documents has been conducted for a specific case study: an explosion which occurred on the Kettle Valley railway on 29th October, 1924 [1]. This case study was conducted to gain an understanding of how to automate the process (Table 1). No investigative analytical techniques were used to generate or evaluate MECE hypotheses in this analysis. The situation was chosen because it is sufficiently ambiguous such that the cause of the explosion is unknown; similar to many real-world investigations. There are two commonly recognised hypotheses as to the cause of the explosion on the train. The first hypothesis was that the explosion was accidental and caused by gas pipes which had malfunctioned. The second was that the explosion was caused by detonation of an explosive.

Table 1. Manual and automated analyses to be performed on the archives relating to the Kettle Valley train explosion.

Task	**Manual**	**Automated**
Text Segmentation	✓	✓
Argument / Non-Argument Classification	✓	In progress
Simple Structure	✓	In progress
Refined Structure	✓	In Progress
AF Instantiation and Evaluation	✓	In Progress

6 Expected Contribution

The main contribution of this research will be the demonstration of the feasibility of generating and evaluating the set of MECE hypotheses using arguments modelled within unknown, investigative, or inquiry-based settings. The hypotheses will be evaluated using computational argumentation and probability theory to provide analysts with the relative uncertainty in each hypothesis. Any decision-support tools created should increase the speed and accuracy of analyses, and their utility will be tested accordingly. The implemented techniques are envisaged to provide benefits to the intelligence, legal and business communities.

Acknowledgements

I would like to thank Prof. Katie Atkinson, Prof. Simon Maskell, Prof. Chris Reed, and Todd Robinson for their contributions to this research project thus far. I would also like to thank the Defence Science and Technology Laboratory and the UK Engineering and Physical Sciences Research Council (Grant number: EP/S023445/1) for financially supporting this project.

References

1. Explosion on the kettle valley line: The death of peter verigin. `https://www.canadianmysteries.ca/sites/verigin/home/indexen.html`, accessed: 15th March 2021
2. Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. Annals of Mathematics and Artificial Intelligence **34**(1–3), 197–215 (2002). https://doi.org/10.1023/A:1014490210693
3. Atkinson, K., Bench-Capon, T.: Abstract argumentation scheme frameworks. In: Dochev, D., Pistore, M., Traverso, P. (eds.) Artificial Intelligence: Methodology, Systems, and Applications. pp. 220–234. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
4. Atkinson, K., Bench-Capon, T., McBurney, P.: Computational representation of practical argument. Synthese **152**, 157–206 (10 2006). https://doi.org/10.1007/s11229-005-3488-2
5. Barbieri, D.: Bayesian intelligence analysis (10 2013)

6. Baroni, P., Gabbay, D., Giacomin, M.: Handbook of Formal Argumentation. College Publications (2018)
7. Baroni, P., Giacomin, M.: Semantics of Abstract Argument Systems, pp. 25–44. Springer US, Boston, MA (2009)
8. Bench-Capon, T., Dunne, P.E.: Argumentation in artificial intelligence. Artificial Intelligence **171**(10), 619–641 (2007). https://doi.org/https://doi.org/10.1016/j.artint.2007.05.001, `https://www.sciencedirect.com/science/article/pii/S0004370207000793`, argumentation in Artificial Intelligence
9. Bench-Capon, T.J.M.: Persuasion in Practical Argument Using Value-based Argumentation Frameworks. Journal of Logic and Computation **13**(3), 429–448 (06 2003). https://doi.org/10.1093/logcom/13.3.429, `https://doi.org/10.1093/logcom/13.3.429`
10. Bench-Capon, T.J.M., Prakken, H.: Using argument schemes for hypothetical reasoning in law. Artificial Intelligence and Law **18**, 153–174 (2010)
11. Bentahar, J., Moulin, B., Belanger, M.: A taxonomy of argumentation models used for knowledge representation. Artif. Intell. Rev. **33**, 211–259 (03 2010). https://doi.org/10.1007/s10462-010-9154-1
12. Besnard, P., Garcia, A., Hunter, A., Modgil, S., Prakken, H., Simari, G., Toni, F.: Introduction to structured argumentation. Argument & Computation **5** (02 2014). https://doi.org/10.1080/19462166.2013.869764
13. Boella, G., Gabbay, D., van der Torre, L., Villata, S.: Support in abstract argumentation. vol. 216, pp. 111–122 (01 2010). https://doi.org/10.3233/978-1-60750-618-8-111
14. Cabrio, E., Villata, S.: Five years of argument mining: a data-driven analysis. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18. pp. 5427–5433. International Joint Conferences on Artificial Intelligence Organization (7 2018). https://doi.org/10.24963/ijcai.2018/766, `https://doi.org/10.24963/ijcai.2018/766`
15. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolar abstract argumentation systems, pp. 65–84. Springer US, Boston, MA (2009), `https://doi.org/10.1007/978-0-387-98197-0_4`
16. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–357 (1995). https://doi.org/https://doi.org/10.1016/0004-3702(94)00041-X, `https://www.sciencedirect.com/science/article/pii/000437029400041X`
17. Dung, P., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. Artificial Intelligence **171**(10), 642–674 (2007). https://doi.org/https://doi.org/10.1016/j.artint.2007.05.003, `https://www.sciencedirect.com/science/article/pii/S000437020700080X`, argumentation in Artificial Intelligence
18. Eemeren, F.H.v., Garssen, B., Verheij, B., Krabbe, E.C.W., Snoeck Henkemans, A.F., Wagemans, J.H.M.: Handbook of Argumentation Theory. Springer Netherlands (2014)
19. Fisk, C.E.: The sino-soviet border dispute: A comparison of the conventional and bayesian methods for intelligence warning. Studies in Intelligence **6**(2), 53–62 (1972)
20. Freeley, A., Steinberg, D.: Argumentation and Debate. Cengage Learning (2013), `https://books.google.co.uk/books?id=1YOWAAAAQBAJ`

21. Grabmair, M., Ashley, K.: Argumentation with value judgments - an example of hypothetical reasoning. vol. 223, pp. 67–76 (01 2010). https://doi.org/10.3233/978-1-60750-681-2-67
22. Hesketh, R.: Fortitude: The D-Day Deception Campaign. The Overlook Press, 1 edn. (2000)
23. Heuer, R.J.: Psychology of intelligence analysis. Center for the Study of Intelligence, Central Intelligence Agency (1999)
24. Heuer Jr., R., Pherson, R.: Structured Analytic Techniques for Intelligence Analysis. CQ Press (7 2014)
25. Li, H.: Probabilistic argumentation. Ph.D. thesis, Citeseer (2015)
26. Liao, B., Jin, L., Koons, R.C.: Dynamics of argumentation systems: A division-based method. Artificial Intelligence **175**(11), 1790–1814 (2011). https://doi.org/https://doi.org/10.1016/j.artint.2011.03.006, `https://www.sciencedirect.com/science/article/pii/S0004370211000518`
27. Lippi, M., Torroni, P.: Argumentation mining: State of the art and emerging trends. ACM Trans. Internet Technol. **16**(2) (Mar 2016). https://doi.org/10.1145/2850417, `https://doi.org/10.1145/2850417`
28. Murukannaiah, P.K., Kalia, A.K., Telangy, P.R., Singh, M.P.: Resolving goal conflicts via argumentation-based analysis of competing hypotheses. 2015 IEEE 23rd International Requirements Engineering Conference (RE) (2015). https://doi.org/10.1109/re.2015.7320418
29. Nouioua, F., Risch, V.: Argumentation frameworks with necessities. In: Benferhat, S., Grant, J. (eds.) Scalable Uncertainty Management. pp. 163–176. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
30. Oren, N., Norman, T.: Semantics for evidence-based argumentation. In: Besnard, P., Doutre, S., Hunter, A. (eds.) Computational Models of Argument. pp. 276–284. Frontiers in artificial intelligence and applications, IOS Press (2008)
31. Pherson, R., Donner, O.: The case of kabul: Foresight and the hindsight bias trap (2021), `http://www.globalytica.com/foresight-and-hindsight/`
32. Pollock, J.L.: Defeasible reasoning. Cogn. Sci. **11**, 481–518 (1987)
33. Popper, K.R.: The logic of scientific discovery. New York: Harper and Row (1968)
34. Simari, G.R., Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation. Artificial Intelligence **53**(2), 125–157 (1992). https://doi.org/https://doi.org/10.1016/0004-3702(92)90069-A, `https://www.sciencedirect.com/science/article/pii/000437029290069A`
35. Walton, D.: Argumentation theory: A very short introduction. In: Argumentation in Artificial Intelligence (2009)
36. Walton, D.: Argument mining by applying argumentation schemes. Studies in Logic **4** (04 2012)
37. Walton, D., Reed, C., Macagno, F.: Argumentation Schemes. Cambridge University Press (2008). https://doi.org/10.1017/CBO9780511802034
38. Walton, D.N.: Argumentation Schemes for Presumptive Reasoning. L. Erlbaum Associates (1996)