Knowledge-Enhanced Argument Mining and Reasoning for Argumentative Dialogues

Ameer Hassan Saadat-Yazdi

The University of Edinburgh

Abstract

My work seeks to bridge the gap between NLP and computational argumentation in order to reason about the validity of online arguments. In order to do so, we wish to leverage commonsense knowledge to reveal implicit lexical and logical inferences and improve reasoning in the presence of enthymemes. We do this with the vision of building an argumentative dialogue agent to help with question answering and elaboration, particularly in the domain of online reviews. This agent will be composed of two components: (i) argument understanding, to extract and build bipolar argumentation frameworks from text, and (ii) argument generation, to identify relevant and valid arguments to provide to the user based on their input.

1 Introduction

Individuals create and share content online more than ever. Processing such large amounts of data and making sense out of it is a challenging problem. Artificial Intelligence (AI) techniques are useful to address such challenges. Many example AI systems are successful in helping users to achieve their tasks. However, current AI systems are far from trustworthy AI systems that the users need. We focus on human-centered AI systems that should be: (i) transparent in providing rationales underlying their recommendations, (ii) fair in providing all possible outcomes instead of one most likely outcome, (iii) adaptive to the inputs provided by their users while capturing the contextual information needed.

Our vision is to build an AI system, shown in Figure 1, that will analyze data from heterogeneous information sources to answer a specific query of the user about a product, book, movie and so on. The system will provide various opinions by displaying supporting or attacking arguments as justifications. For this, we will be using Argument Mining techniques to extract arguments and the relations among them (support or attack) from natural language text (Argument Understanding). We will then use Computational Argumentation techniques in order to evaluate and generate relevant arguments based on user input. Such an approach will be used to evaluate the strength and validity of arguments within the discourse at large. Finally, the system will be able to provide personalized explanations to the user based on the obtained argumentation graph while giving an option to the user to choose between options (Argument Generation).

Computational argumentation has been applied in AI and multi-agent systems as a formal approach to reasoning with evidence, where evidence is represented in terms of arguments for and against conclusions. One big challenge in this domain is the automatic discovery of arguments and the relationships among them in unstructured text. Our work so far has focused on the relation classification component of our system which will be used to construct abstract argumentation frameworks [Saadat-Yazdi et al., 2022].

Identifying argumentative relationships often requires the ability to make implicit inferences, draw

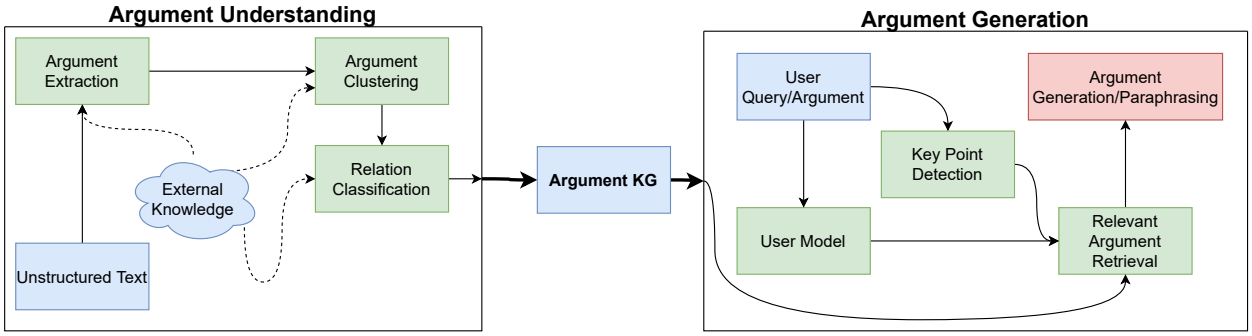

Figure 1: Our vision for the complete argumentative agent describing the elements of the understanding and generation components.

from domain knowledge and/or identify sarcasm and irony. Our work seeks to alleviate the most prevalent of these challenges, the need to make inferences in the face of incomplete arguments. Some recent works have identified the role that explicit knowledge could play in identifying these relations. [Abels et al., 2021] provides an example of this approach where a long short-term memory (LSTM) is used to compute representations of relevant paths in the Wikidata knowledge graph that links the two arguments at hand. The authors find that this significantly improves the accuracy of the models on argument mining tasks. Similarly, [Fromm et al., 2019], shows that context cues derived from Wikipedia based knowledge graphs and pretrained transformers provide useful signals for relational argument mining.

The works mentioned above explore the role of fact-based knowledge graphs, which contain information about entities and their relationships. While this is useful for providing contextual information to the model, it does not necessarily clarify the implicit inferences made by humans in their arguments. For example, the argument "Jack seemed tired this morning", supports the argument "he was up all-night doing homework" given the inference step "up all night" $\Rightarrow$ "Jack is tired", where $X \Rightarrow Y$ means the argument X implies the argument Y (i.e., logical implication). Yet this implication would not be found in a fact-based knowledge graph such as Wikidata or DBPedia. Even if such an implication exists in a knowledge base, it is unclear if the implication would be valid for a specific context.

2 Method

2.1 Current Work

My work has so far focused on the role of commonsense knowledge in the identification of attack and support relations between arguments. In particular my supervisory team and I have investigated the role of the commonsense knowledge graph, ATOMIC [Sap et al., 2019], in providing useful knowledge to our models. Alongside this, we designed a novel model architecture to incorporate external commonsense knowledge with a pre-trained language model for relation classification [Saadat-Yazdi et al., 2022], see Figure 2.

We developed three methods for extracting knowledge from ATOMIC:

1. Extracting knowledge directly based on the input arguments (ATOMIC),

2. Using a transformer model trained on ATOMIC, to generate knowledge based on the input (COMET),

3. Enhancing the above model with link prediction (COMET+LINK).

The use of transformers to generate commonsense knowledge is described in [Hwang et al., 2021] which introduces the COMET model used in our work. However, in contrast with existing work, ours introduces a novel approach to generate multi-hop knowledge that connects separate arguments together by a thread of reasoning. This is something we wish to explore further in future works.

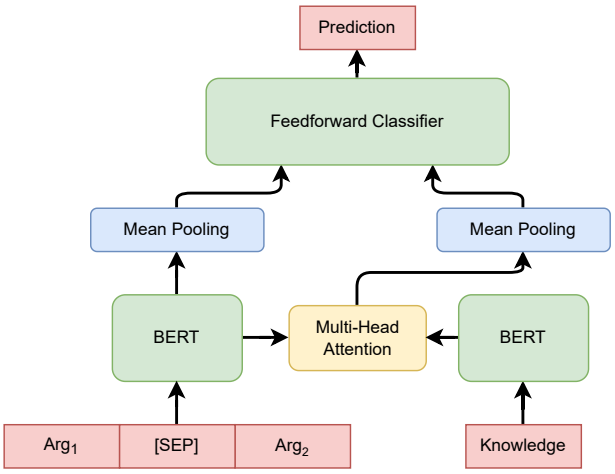

Figure 2: Architecture of the model we used for argument relation classification with external knowledge.

2.2 Computational Argumentation

Alongside further advancing the work of relation classification, I wish to begin work on more theoretical aspects of argumentation, that is the Argument Generation component as depicted in Figure 1. Computational argumentation provides a solid foundation upon which we wish to build our system, due to its ability to reason about uncertainty and explicitly model conflicts. In particular, the semantics described in [Amgoud and Ben-Naim, 2018], which provides a weighted semantics for bipolar argumentation frameworks (BAFs), and [Baier et al., 2021], which provides an analysis of probabilistic semantics of abstract argumentation frameworks and their properties, provide solutions to integrating uncertainty about arguments into our system. This is essential given that the outputs of our argument understanding components are probabilistic in nature.

The notion of augmenting argumentation frameworks with background knowledge has recently began gaining popularity. Argumentation Knowledge Graphs [Al-Khatib et al., 2020] are a proposed structure to capture the dialectical relationships of various arguments while grounding the entities mentioned in the real world. This approach could prove useful when additional background knowledge is needed to explain to users why one argument attacks/supports another. We hope that this would allow the agent to have more natural conversations with users by drawing on the external knowledge when necessary.

The aforementioned aspects of our system are purely symbolic, and would in fact be data independent. This would allow us to separate the data-intensive NLP aspects of our work from the symbolic manipulations. The benefits of this are twofold: firstly, the symbolic manipulations allow for the interpretation of our system's inner reasoning mechanisms, and second, these components can be transferred across domains and even languages, requiring only the retraining of NLP components.

Once relevant arguments have been identified by aggregating relevance measures with argumentative semantics, the agent will be required to paraphrase the text to fit within the context of the conversation being had with the user. This kind of conditional paraphrasing is understudied in NLP domain and so we would like to carry out a feasibility study to determine whether or not a working solution can be designed given existing technologies.

The evaluation of our system will pose a challenge given the lack of data and ambiguity of what are considered valid, persuasive and appropriate arguments. While there is some work in this regard for automatic argument evaluation as well as automatic measures of dialogue quality, user studies will be necessary to evaluate the performance of our system in real-world settings. We aim to carry out our user studies in the context of online reviews, where the user can ask about various aspects of a product and be provided with arguments in support or against the product based on existing reviews. Our system would also allow users to explore why a certain response was

given by exploring portions of the the Argumentation Knowledge Graph. We will measure the quality, usefulness and transparency of our system in helping users make decisions about online purchases.

3 Discussion

In our paper [Saadat-Yazdi et al., 2022], we compared [Paul et al., 2020]. The previous work uses LSTM models to integrate commonsense knowledge extracted from ConceptNet into a relation classification model. We show that BERT with additional knowledge performs exceptionally well, benefiting greatly from the introduction of ATOMIC knowledge. Our approach was found to give significant improvements to the state-of-the-art on the two datasets, Debatepedia [Paul et al., 2020] and Student Essay [Opitz and Frank, 2019]. We found that the addition of the transformer generated knowledge into our model improved significantly upon a BERT baseline.

Our analysis also emphasised that while knowledge can be useful, all the knowledge extraction methods we used are prone to making invalid or counterfactual logical connections between arguments. In order to build useful explanations of implicit links between arguments, in future work, I wish to explore more robust methods that rely on natural language inference.

The dialogical aspects of our system will build off of [Chalaguine and Hunter, 2021] and [Hadoux et al., 2021] where an agent is created to choose suitable arguments from a bipolar argumentation framework to respond to a user's queries; these responses are conditioned on models of the user's beliefs and preferences in order to offer more persuasive answers. The BAFs used in these works are constructed by crowdsourcing arguments from a large number of participants. A similar work, [Kökciyan et al., 2021], applies metalevel argumentation frameworks to identify possible treatment plans for patients based on their preferences and needs in a text-based format. A key difference between our desired system and these works is that the use of NLP to construct the argumentation frameworks automatically. The latter work provides an initial step towards this by using a rule-based system to mine arguments from government health websites and to develop an argumentation-based chatbot to communicate with the user in a semi-structured manner. By combining these ideas with our existing work on argument mining, we hope to implement argumentative chatbots that draw knowledge from automatically constructed argumentative knowledge graphs in order to help users make decisions about product reviews, medical treatment, business decisions etc. . . .

4 Conclusion

In this paper I have described our vision for a argumentative agent which can construct bipolar argumentation frameworks from the web and use this as a source of knowledge to engage in a conversation with users. Beginning with the relation classification component, I have outlined our findings on improving commonsense reasoning for relational argument mining using generated knowledge and described some of the key limitations of this approach. In addition to this, I have discussed next steps on improving this component as well as theoretical research questions we wish to address once we begin constructing bipolar argumentation frameworks.

My work up until this point has focused exclusively on arguments that appear in the same context; i.e., within a single essay or debate. However, it is often useful to identify argumentative relations from different contexts, e.g. tweets that do not belong in the same conversation. I also have not addressed the classification of neutral relations between arguments in the same context. This is a necessary aspect in argument extraction, where the argumentative sentences are not known in advance. These are problems which I wish to address in future work.

Acknowledgements

I thank my supervisors, Nadin Kokciyan and Jeff Z. Pan, for their continued support of my research. I also wish to acknowledge the Edinburgh-Huawei Research Lab for funding my work.

References

[Abels et al., 2021] Abels, P., Ahmadi, Z., Burkhardt, S., Schiller, B., Gurevych, I., and Kramer, S. (2021). Focusing Knowledge-based Graph Argument Mining via Topic Modeling. *arXiv:2102.02086 [cs]*. arXiv: 2102.02086.

[Al-Khatib et al., 2020] Al-Khatib, K., Hou, Y., Wachsmuth, H., Jochim, C., Bonin, F., and Stein, B. (2020). End-to-End Argumentation Knowledge Graph Construction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7367–7374.

[Amgoud and Ben-Naim, 2018] Amgoud, L. and Ben-Naim, J. (2018). Weighted Bipolar Argumentation Graphs: Axioms and Semantics. *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 5194–5198.

[Baier et al., 2021] Baier, C., Diller, M., Dubslaff, C., Gaggl, S. A., Hermanns, H., and Käfer, N. (2021). Admissibility in Probabilistic Argumentation. In *Proceedings of the Eighteenth International Conference on Principles of Knowledge Representation and Reasoning*, pages 87–98, Hanoii, Vietnam. International Joint Conferences on Artificial Intelligence Organization.

[Chalaguine and Hunter, 2021] Chalaguine, L. and Hunter, A. (2021). Addressing Popular Concerns Regarding COVID-19 Vaccination with Natural Language Argumentation Dialogues. In Vejnarová, J. and Wilson, N., editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, Lecture Notes in Computer Science, pages 59–73, Cham. Springer International Publishing.

[Fromm et al., 2019] Fromm, M., Faerman, E., and Seidl, T. (2019). TACAM: Topic And Context Aware Argument Mining. *IEEE/WIC/ACM International Conference on Web Intelligence*, pages 99–106. arXiv: 1906.00923.

[Hadoux et al., 2021] Hadoux, E., Hunter, A., and Polberg, S. (2021). Strategic argumentation dialogues for persuasion: Framework and experiments based on modelling the beliefs and concerns of the persuadee. *CoRR*, abs/2101.11870.

[Hwang et al., 2021] Hwang, J. D., Bhagavatula, C., Bras, R. L., Da, J., Sakaguchi, K., Bosselut, A., and Choi, Y. (2021). COMET-ATOMIC 2020: On Symbolic and Neural Commonsense Knowledge Graphs. *arXiv:2010.05953 [cs]*. arXiv: 2010.05953.

[Kökciyan et al., 2021] Kökciyan, N., Sassoon, I., Sklar, E., Modgil, S., and Parsons, S. (2021). Applying Metalevel Argumentation Frameworks to Support Medical Decision Making. *IEEE Intelligent Systems*, 36(2):64–71. Conference Name: IEEE Intelligent Systems.

[Opitz and Frank, 2019] Opitz, J. and Frank, A. (2019). Dissecting content and context in argumentative relation analysis. In *Proceedings of the 6th Workshop on Argument Mining*, pages 25–34, Florence, Italy. Association for Computational Linguistics.

[Paul et al., 2020] Paul, D., Opitz, J., Becker, M., Kobbe, J., Hirst, G., and Frank, A. (2020). Argumentative Relation Classification with Background Knowledge. *Computational Models of Argument*, pages 319–330. Publisher: IOS Press.

[Saadat-Yazdi et al., 2022] Saadat-Yazdi, A., Kokciyan, N., and Pan, J. Z. (2022). Situational reasoning with transformers for argument mining. Under review.

[Sap et al., 2019] Sap, M., LeBras, R., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N. A., and Choi, Y. (2019). ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning. *arXiv:1811.00146 [cs]*. arXiv: 1811.00146.