Deciphering Conversational Networks: Stance Detection via Hypergraphs and LLMs

Daniele De Vinco
Università degli Studi di Salerno
Fisciano, Italy
ddevinco@unisa.it

Alessia Antelmi
Università degli Studi di Torino
Torino, Italy
alessia.antelmi@unito.it

Carmine Spagnuolo
Università degli Studi di Salerno
Fisciano, Italy
cspagnuolo@unisa.it

Luca Maria Aiello
IT University of Copenhagen
Copenaghen, Denmark
luai@itu.dk

ABSTRACT

Understanding the structural and linguistic properties of conversational data in social media is crucial for extracting meaningful insights to understand opinion dynamics, (mis-)information spreading, and the evolution of harmful behavior. Current state-of-the-art mathematical frameworks, such as hypergraphs and linguistic tools, such as large language models (LLMs), offer robust methodologies for modeling high-order group interactions and unprecedented capabilities for dealing with natural language-related tasks. In this study, we propose an innovative approach that blends these worlds by abstracting conversational networks via hypergraphs and analyzing their dynamics through LLMs. Our aim is to enhance the stance detection task by incorporating the high-order interactions naturally embedded within a conversation, thereby enriching the contextual understanding of LLMs regarding the intricate human dynamics underlying social media data.

CCS CONCEPTS

• **Mathematics of computing → Hypergraphs**; • **Computing methodologies → Natural language processing**.

KEYWORDS

Conversational networks, hypergraphs, stance detection, LLMs

ACM Reference Format:
Daniele De Vinco, Alessia Antelmi, Carmine Spagnuolo, and Luca Maria Aiello. 2024. Deciphering Conversational Networks: Stance Detection via Hypergraphs and LLMs. In *16th ACM Web Science Conference (Websci Companion '24), May 21–24, 2024, Stuttgart, Germany*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3630744.3658418

1 INTRODUCTION

Social interactions through online media are increasingly pervasive [12], offering opportunities to study a wide range of social human behavior, including the mechanisms behind collective action, the spread of (mis)information, and the evolution of harmful behaviors. The widespread adoption of Large Language Models (LLMs) like ChatGPT or Gemini are changing the way in which people produce and consume content in online social media. LLMs are influencing how people create, share, and consume content from textual descriptions to images to videos [5]. Moreover, from an academic perspective, a growing body of work demonstrates how LLMs outperform state-of-the-art models in various tasks [6, 11, 13]. Traditionally, online interactions between social media users have been successfully studied through graphs abstracting the underlying relations with nodes and edges connecting pairs of interacting components. Yet, many online actions are characterized by group interactions that cannot be described simply in terms of dyads (e.g., all users commenting on the same discussion thread). Hypergraphs are the perfect candidates to model group-wise interactions, as these structures are a generalization of graphs where a (hyper)edge allows the connection of an arbitrary number of nodes.

In this line of work, we focus on analyzing social media conversational data and, specifically, on the task of stance detection. In this preliminary study, we introduce an innovative approach to address this task, which combines the use of hypergraphs for abstracting conversational networks with the analytical power of LLMs to examine their dynamics. We opted to use LLMs based on compelling evidence from the literature demonstrating their rising dominance over graph neural networks as the predominant tool for graph analysis [7–9]. Additionally, LLMs have proven their efficacy in stance detection tasks, demonstrating versatility across datasets with varying prompting schemes and even surpassing supervised models in performance while demanding fewer resources [4]. While previous studies have explored the integration of LLMs and graphs for stance detection [3], they have overlooked group-wise interactions, which are recognized as pivotal in understanding many real-world phenomena [1]. Building upon these foundations, our objective is to leverage LLMs to address the stance detection task within online social media conversations represented as hypergraphs. We envision that this modeling approach has the potential to not only advance the current state-of-the-art but also deepen our comprehension of social media dynamics.

Websci Companion '24, May 21–24, 2024, Stuttgart, Germany

ACM ISBN 979-8-4007-0453-6/24/05
https://doi.org/10.1145/3630744.3658418

2 METHODOLOGY

Problem definition. In more formal terms, our objective is to address the following task: given input represented as a conversational hypergraph and a target pair, our goal is to classify each comment within the conversation based on the author's stance toward the target. In our case, the stance is categorized into one of three labels: favor, against, or neutral.

Based on previous literature [2], we model a conversational hypernetwork as a Text-Attributed Hypergraph (TAH). A TAH is formally defined as a tuple $H = (V, E, T)$, where V is the node set representing the users of a conversation, E is the hyperedge set denoting the group-wise relationship of the users in the conversation and T is the comment set. Specifically, each node $v \in V$ is associated with a textual comment $t \in T$.

Research questions (RQs). In our work, we aim to answer the following RQs.

· RQ_1: Could the performance of a model benefit from integrating high-order interactions capturing the nuanced local context of conversations when addressing a stance detection task?

· RQ_2: Given the scarcity of labeled conversational datasets, how feasible is it to use LLMs as generators of ground truth?

Methodology. To address the RQs outlined above, we devised the pipeline shown in Figure 1, and opted to focus on the specific topic of climate change. Our initial step involved selecting a dataset of conversational threads. We began by choosing the SPINOS dataset [10], which features human annotations of Reddit comments across six distinct topics, including climate change. To enrich this dataset, we collected all messages from the three climate-related subreddits (`climate`, `climatechange`, `climateskeptics`) spanning from January 4th, 2019, to July 31st, 2020. After constructing the dataset (|messages| = 230122, |authors| = |V| = 21496, |submissions| = |E| = 18771), the next step involved assigning a stance to each unlabeled message. For this task, we employed Gemma 1.1 7B [1], using the same prompt as the human annotators who compiled the SPINOS dataset, by obtaining an accuracy of 68%. The second step involved building the TAH associated with the so-built dataset and feeding it to *HyperStance*, our proposed stance detection model. This model will then evaluate the stance attributed to a user by simultaneously learning the semantics of the textual description and the hypergraph structure, which encapsulates the high-order local context of a conversation. In the final stage of validation, we will assess our model by comparing it with the current state-of-the-art approaches, which either leverage the conversational structure or solely focus on individual pieces of text. The model, experiments, and data are freely available on GitHub[2].

Challenges. Our primary challenge lies in the absence of ground truth data against which we can assess the developed model. Current labeled datasets include only isolated text fragments, yet we require access to complete conversation threads to capture the high-order dynamics of online platforms. Leveraging pre-trained LLMs to construct ground truths for this purpose could bridge this gap, albeit potentially introducing bias. Moreover, handling (potentially large-scale) hypergraphs introduces an additional layer of computational complexity, which significantly challenges the scalability and efficiency of the developed pipeline.

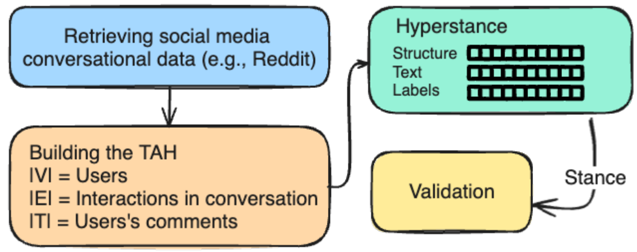

Figure 1: Overview of the pipeline.

3 CONCLUSION AND CURRENT WORK

This paper presented a novel approach to leverage LLMs to address the stance detection task within online social media conversations modeled as hypergraphs. Currently, we are working on enhancing the training dataset and implementing HyperStance.

Acknowledgement. This work has been partially supported by the spoke "FutureHPC & BigData" of the ICSC – Centro Nazionale di Ricerca in High-Performance Computing, Big Data and Quantum Computing funded by European Union – NextGenerationEU.

REFERENCES

[1] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri. 2020. Networks beyond pairwise interactions: Structure and dynamics. *Physics Reports* 874 (2020), 1–92. https://doi.org/10.1016/j.physrep.2020.05.004 Networks beyond pairwise interactions: Structure and dynamics.

[2] A. Bazaga, P. Liò, and G. Micklem. 2024. HyperBERT: Mixing Hypergraph-Aware Layers with Language Models for Node Classification on Text-Attributed Hypergraphs. arXiv:2402.07309 [cs.LG]

[3] R. Cao, X. Luo, Y. Xi, and Y. Qiao. 2022. Stance detection for online public opinion awareness: An overview. *International Journal of Intelligent Systems* 37, 12 (2022), 11944–11965.

[4] I. Cruickshank and L. Ng. 2023. Use of large language models for stance classification. *arXiv preprint arXiv:2309.13734* (2023).

[5] L. Da Silva, J. Samhi, and F. Khomh. 2024. ChatGPT vs LLaMA: Impact, Reliability, and Challenges in Stack Overflow Discussions. *arXiv preprint arXiv:2402.08801* (2024).

[6] X. He, X. Bresson, T. Laurent, A. Perold, Y. LeCun, and B. Hooi. 2023. Harnessing explanations: Llm-to-lm interpreter for enhanced text-attributed graph representation learning. In *The Twelfth International Conference on Learning Representations*.

[7] Y. Hu, Z. Zhang, and L. Zhao. 2023. Beyond Text: A Deep Dive into Large Language Models' Ability on Understanding Graph Data. *arXiv preprint arXiv:2310.04944* (2023).

[8] J. Huang, X. Zhang, Q. Mei, and J. Ma. 2023. Can llms effectively leverage graph structural information: when and why. *arXiv preprint arXiv:2309.16595* (2023).

[9] Y. Li, Z. Li, P. Wang, J. Li, X. Sun, H. Cheng, and J.X. Yu. 2023. A survey of graph meets large language model: Progress and future directions. *arXiv preprint arXiv:2311.12399* (2023).

[10] F. Sakketou, A. Lahnala, L. Vogel, and L. Flek. 2022. Investigating User Radicalization: A Novel Dataset for Identifying Fine-Grained Temporal Shifts in Opinion. arXiv:2204.10190 [cs.CL]

[11] M. Tsimpoukelli, J. L. Menick, S. Cabi, S.M. Eslami, O. Vinyals, and F. Hill. 2021. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems* 34 (2021), 200–212.

[12] P. Verduyn, N. Gugushvili, K. Massar, K. Täht, and E. Kross. 2020. Social comparison on social networking sites. *Current opinion in psychology* 36 (2020), 32–37.

[13] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* (2022).

[1] https://ai.google.dev/gemma
[2] https://github.com/ddevin96/HyperStance