BACK TO THE FUTURE: PREDICTING CAUSAL RELATIONSHIPS INFLUENCING OIL PRICES

Jože M. Rožanec, Beno Šircelj *
Department for Artificial Intelligence
Jožef Stefan Institute
Jamova cesta 39, 1000 Ljubljana, Slovenia
{joze.rozanec,beno.sircelj}@ijs.si

Michael Cochez
Computer Science department
Vrije Universiteit Amsterdam
Amsterdam, Netherlands
{m.cochez}@vu.nl

Gregor Leban
Event Registry d.o.o.
{gregor}@eventregistry.com

ABSTRACT

Strategic foresight has been identified as a key tool to enhance policymaking and guide decision-making related to environmental protection. Artificial intelligence (AI) has been scarcely applied to this domain. We present a novel approach to predict causal relationships between entities leveraging graph neural networks (GNN) for link prediction. We test our approach on a real-world dataset, achieving a precision of 0.4905 when predicting causal relationships over 2000 pairs of entities corresponding to events reported between 2022 and 2023.

1 INTRODUCTION

Background Strategic Foresight offers a systematic method for collecting information on potential future scenarios to prepare for change effectively. It delivers expert perspectives on evolving trends and emerging challenges that can be factored into strategic planning and policy development. Strategic foresight is essential for policymaking; especially in the context of a VUCA (Volatile, Uncertain, Complex, and Ambiguous) world order and TUNA (Turbulence, Unpredictability, Uncertainty, Novelty, and Ambiguity) conditions Wilkinson (2017); Kaivo-oja & Lauraeus (2018).

Motivation A study conducted in the USA identified scenario planning as one of the most frequently used strategic foresight techniques (Greenblott et al., 2019). Moreover, strategic foresight has been recognized as a key tool to guide environmental decision-making and achieve conservation objectives (Cook et al., 2014; Ednie et al., 2023). While artificial intelligence (AI) has been identified as a key technology that could enhance strategic foresight capabilities, research invested in this domain has been scarce (Reez, 2020).

State of the Art and research gap Research applying artificial intelligence to strategic foresight has mostly focused on automating repetitive and time-consuming tasks. Parrish et al. (2019) and Brandtner & Mates (2021) leveraged AI to automate information scanning and data analysis. Geurts et al. (2022) used text mining to perform topic modeling to analyze trends and identify weak signals, which domain experts consider to identify plausible futures and their impact. Recently, researchers have explored the use of LLMs for causal reasoning, but their findings are contradictory: while some researchers embrace them, others suggest LLMs are not causal (Kıcıman et al., 2023; Zečević et al., 2023). Nevertheless, little research has been devoted to modeling causal relationships in the context of strategic foresight and predicting possible future scenarios grounded on facts using graph-based approaches. LLMs could be used to create scenarios grounded on fact-based causal predictions.

Key insights This paper introduces a novel approach to AI-assisted strategic foresight. We leverage LLMs to analyze media events and extract causal relationships reported in media news. We map

*Corresponding author: Jože M. Rožanec (joze.rozanec@ijs.si)

strings to concepts, build a causality graph, and train a GNN for link prediction. By building the causality graph for some period of interest, we can predict causal relationships for new and future events.

2 Experimental setup

Relevant period For our research, we extracted media events (cluster of media news reporting about the same event occurring in the real world) reported by EventRegistry Leban et al. (2014), querying media events related to the "*Price of Oil*" concept. Based on geopolitical events highlighted by the U.S. Energy Information Administration [1] and the Russo-Ukrainian War, we considered media events that took place in the first quarter of 2015, 2020, 2022, and 2023. A total of 5486 media events were extracted.

Causal relationships extraction ChatGPT 3.5 was prompted to extract entities and their causal relationships as reported in media news events. When extracting the causal relationship, we disregarded the causal relationship type to (i) increase the critical mass of edges modeled in the causality graph and (ii) ease later forecasting. We envision a system with a domain expert in the loop where this person would easily identify the causal relationship type. The strings identifying entities were processed, removing non-alphanumeric characters and converting them to lowercase. The dates of the media news reporting the causal relationship were kept to understand how the causal graph evolved. Matching was performed to turn strings into concepts, leaving a causality graph with a mixture of strings and Wikidata entities (Vrandečić, 2012). Further details on matching are provided in Appendix B. As alternatives, matching against other broad domain knowledge graphs (DBpedia and KBPedia) was tried. However, the best results were obtained with Wikidata. We weighted the edges (causal relationships) based on how frequently they were reported in the media news events, providing an empirical estimate of the probability of occurrence. If a causal relationship was reported multiple times within the same media event, we counted it only once.

Learning causal relationships We trained a GraphSAGE model (Hamilton et al., 2017) to learn causal relationships between entities, which leverages node information to generalize to unseen data. Following a similar rationale to Daza et al. (2021), nodes were characterized with the average word2vec embedding (Church, 2017) for the words describing a particular entity. Doing so helps keep the embeddings for entities with similar semantic meanings close in the embedded space. The GNN was trained on 10% of the causal relationships (about 2000 pairs of entities) extracted from 2015 and 2020. The model's predictive power was tested on causal relationships in the first quarter of 2022 and 2023.

Experimental scenarios We were interested in two experimental scenarios: (i) given a random pair of entities, how accurately can the GNN model predict a causal relationship between them? and (ii) given a stream of events, is the GNN model able to predict a chain of causal relationships?

For the first experiment, we predicted causal relationships on about 2000 pairs of entities from the first quarters of 2022 and 2023. For the second experiment, we randomly sampled an entity for the first month (January 2022) and streamed the rest of the months, one month at a time. The link prediction model was applied from the original randomly sampled entities to all newly added ones. To select the final node, we started with the link prediction ranking. We removed all predictions with a probability of causal relationship below 0.5 and then kept ten top candidates at most. Among these, we sampled randomly with a probability proportional to the score. The same procedure was applied until the pool of six test months was exhausted.

3 Results, conclusion, and future work

Results and analysis Our experiment results show that our setup with GraphSAGE accurately predicted causal relationships: we observed a precision of 0.4905 when predicting over 2000 pairs of entities from the test set. Regarding the causality chains, we observed that most were partially

[1]The geopolitical events were highlighted in the following report, last accessed on August 25th 2023: `https://www.eia.gov/finance/markets/crudeoil/spot_prices.php`.

accurate, and few provided deep insights. However, generating causality chains remains a challenge. Some examples we observed: (i) environmental opposition → OPEC → agreement → futures → EV sector; (ii) OPEC → stock market index → currency → fuel pricing → Russo-Ukrainian war → energy industry; and (iii) armed conflict → OPEC → fossil fuel → currency → fuel pricing.

The results validate our research direction toward creating a fully AI-driven strategic foresight tool. Future work will focus on enhancing the causal graph quality, generating better causality chains, and ranking them based on multiple criteria. LLMs may be used to turn causality chains into fully-fledged scenarios.

Acknowledgements

Funding acknowledgments are hidden due to the double-blind review. To be updated on acceptance.

URM Statement

The authors acknowledge that at least one key author of this work meets the URM criteria of the ICLR 2024 Tiny Papers Track.

References

Patrick Brandtner and Marius Mates. Artificial intelligence in strategic foresight–current practices and future application potentials: Current practices and future application potentials. In *The 2021 12th International Conference on E-business, Management and Economics*, pp. 75–81, 2021.

Kenneth Ward Church. Word2vec. *Natural Language Engineering*, 23(1):155–162, 2017.

Carly N Cook, Sohail Inayatullah, Mark A Burgman, William J Sutherland, and Brendan A Wintle. Strategic foresight: how planning for the unpredictable can improve environmental decision-making. *Trends in ecology & evolution*, 29(9):531–541, 2014.

Daniel Daza, Michael Cochez, and Paul Groth. Inductive entity representations from text via link prediction. In *Proceedings of the Web Conference 2021*, pp. 798–808, 2021.

Gabrielle Ednie, Tyreen Kapoor, Olga Koppel, Morgan L Piczak, Jessica L Reid, Alyssa D Murdoch, Carly N Cook, William J Sutherland, and Steven J Cooke. Foresight science in conservation: Tools, barriers, and mainstreaming opportunities. *Ambio*, 52(2):411–424, 2023.

Amber Geurts, Ralph Gutknecht, Philine Warnke, Arjen Goetheer, Elna Schirrmeister, Babette Bakker, and Svetlana Meissner. New perspectives for data-supported foresight: The hybrid ai-expert approach. *Futures & Foresight Science*, 4(1):e99, 2022.

Joseph M Greenblott, Thomas O'Farrell, Robert Olson, and Beth Burchard. Strategic foresight in the federal government: A survey of methods, resources, and institutional arrangements. *World futures review*, 11(3):245–266, 2019.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Jari Roy Lee Kaivo-oja and Iris Theresa Lauraeus. The vuca approach as a solution concept to corporate foresight challenges and global technological disruption. *foresight*, 2018.

Emre Kıcıman, Robert Ness, Amit Sharma, and Chenhao Tan. Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050*, 2023.

Gregor Leban, Blaz Fortuna, Janez Brank, and Marko Grobelnik. Event registry: learning about world events from news. In *Proceedings of the 23rd International Conference on World Wide Web*, pp. 107–110, 2014.

Nathan H Parrish, Anna L Buczak, Jared T Zook, James P Howard, Brian J Ellison, and Benjamin D Baugher. Crystal cube: Multidisciplinary approach to disruptive events prediction. In *Advances in Human Factors, Business Management and Society: Proceedings of the AHFE 2018 International Conference on Human Factors, Business Management and Society, July 21-25, 2018, Loews Sapphire Falls Resort at Universal Studios, Orlando, Florida, USA 9*, pp. 571–581. Springer, 2019.

Norbert Reez. Foresight-based leadership. decision-making in a growing ai environment. In *International Security Management: New Solutions to Complexity*, pp. 323–341. Springer, 2020.

Denny Vrandečić. Wikidata: A new platform for collaborative data collection. In *Proceedings of the 21st international conference on world wide web*, pp. 1063–1064, 2012.

Angela Wilkinson. Strategic foresight primer. *European Political Strategy Centre*, 2017.

Matej Zečević, Moritz Willig, Devendra Singh Dhami, and Kristian Kersting. Causal parrots: Large language models may talk causality but are not causal. *arXiv preprint arXiv:2308.13067*, 2023.

A Appendix: details on ChatGPT prompting

We used the following prompting instructions to retrieve causal relationships from news samples:

> *Given the title and summary of a text, extract the causal relationships. Each causal relationship consists of cause and effect.*
>
> *Cause contains an entity which is an item, individual or company that an event happened to. Event is an action, development, happening or state of the entity that is causing the relationship. Locations are a list of locations where the event in the cause took place.*
>
> *Effect contains an entity which is an item, individual or company who undertook changes because of the event in the cause. Event is an action, development, happening or state of the entity that was affected because of the cause. Locations are a list of locations where the event of the affected entity took place.*
>
> *Output should be in a json format.*

We trained ChatGPT on a media news example, providing also the expected output. Both, the media news example and expected output, have been provided in the supplementary materials. Furthermore, in the supplementary material we also included a sample of nearly 2000 media news events that we used to extract causalities.

B Appendix: matching strings to concepts

A series of preprocessing steps were undertaken to establish successful associations between strings and semantic concepts. Initially, non-letter symbols and stopwords were eliminated, followed by the stemming of individual words. A match was deemed valid if at least one identical string was identified in the original string and the entity of interest. It's important to note that not all semantic concepts identified by the Wikifier were considered. Specifically, criteria included: (a) a requirement for concepts to possess a PageRank higher than 0.0001; (b) for location data, only concepts categorized as "place" were included; and (c) when replacing the original entity with the associated semantic concept, preference was given to the concept with the highest cosine similarity with the corresponding Wikipedia page of the article. We created a dictionary with nearly 250 entries to further cleanse the entities, proposing new semantic terms for frequently observed strings. Such a dictionary will be later structured into a custom ontology. While the current causality graph does not include rich data about the entities, we expect to enrich it by later integrating it with knowledge graphs.

C Appendix: implementation details

To create the Graph Neural Network models, we used the STELLAR Graph library, and trained a GraphSAGE model for link prediction, considering two layers with 32 hidden feature dimensions, with a bias vector learned on each layer and considering a dropout of 0.3. Furthermore, a ReLU output activation function was used, and the model was trained with an Adam optimizer with a learning rate of 0.001 and a cross-entropy loss. The model achieved a precision of 0.4905, recall of 0.6875, and F1 equal to 0.5725.