Empowering Supply chains Resilience: LLMs-Powered BN for Proactive Supply Chain Risk Identification

Maryam **Shahsavari**[1,*], Omar Khadeer **Hussain**[1], Morteza **Saberi**[2] and Pankaj **Sharma**[1]

[1]*Univeriity of New South Wales, Canberra, Australia*

[2]*University of Technology Sydney,*

Abstract

The dynamic and unpredictable nature of today's global risk landscape renders supply chains (SCs) susceptible to vulnerabilities, potentially leading to significant business disruptions if left unaddressed. This paper endeavors to construct a proactive risk identification model aimed at enhancing SC resilience. Our approach incorporates agent models, capable of continuous monitoring and early warning recommendations. To imbue these agents with intelligence, we harness the capabilities of Large Language Models (LLMs) to facilitate text comprehension. Specifically, we employ a Bayesian network (BN) as an agent, utilizing news feeds as its primary information source. We introduce a novel methodology, leveraging the expertise of risk managers and LLMs, to determine the relevance of detected events to the targeted SC risks. This research not only strives to equip businesses with the foresight to anticipate potential risk events but also emphasizes the identification and analysis of contributing events. These contributing events are systematically evaluated to understand their potential to precipitate primary risk events, thereby providing a more nuanced insight into the causative chains that lead to SC disruptions. Our methodology enables the proactive quantification of risk likelihood, enhancing predictive capabilities in SC management.

Keywords

supple chain risk management, Large Language Model, Bayesian Network, risk identification, risk assessment

1. Introduction

Supply chains are complex networks involving suppliers, manufacturers, distributors, and retailers, all coordinated to deliver products to consumers [1]. The operational success of these chains is pivotal, influencing product availability, cost, and quality. With globalization, supply chains have become more complex and exposed to lots of risks, necessitating sophisticated management strategies [2]. The landscape of supply chain risks is diverse, originating from geopolitical tensions, economic instabilities, environmental catastrophes, and health crises, such as the COVID-19 pandemic [3]. These events underscore the fragility of global supply chains, demonstrating the need for robust risk management practices to mitigate disruptions

and ensure continuity [4].

In our previous work [5], we stated that a critical aspect of understanding and managing supply chain risks involves recognizing the cause-and-effect relationship inherent in these risks. Each risk in a supply chain, which is called risk event in this research, does not occur in isolation but is often the result of multiple preceding events. By identifying these causative events and their causal relationship, businesses can adopt a more proactive approach to risk management. This proactive stance is central to the proposed framework, which utilizes the concept of cause and effect to systematically identify contributing events to a risk. Through this approach, it becomes possible to not only identify but also quantify or assess the likelihood of the main risk event occurring, enabling more effective mitigation strategies. Incorporating advanced technologies, such as LLMs [6], into this framework enhances its capability [7]. LLMs can analyze extensive data sets to detect potential causative events and evaluate their impact on the supply chain, providing a novel and powerful tool for proactive risk management. This methodology not only aids in navigating the complexities of modern supply chains but also contributes to building resilience against unforeseen disruptions.

In our research project, we have leveraged the potential of AI specifically Natural Language Processing (NLP) and LLMs in developing Contributing Event-based Risk Identification and Assessment (CERIA) [5], a novel framework capable of analyzing past news to find the causal links between the events that if they occur, they can cause a risk event to happen. CERIA stands out by its ability to continuously scan daily news, identifying occurrences of Contributing Events (CEs). Utilizing this intelligence, the framework quantifies the likelihood of the risk event's occurrence. In CERIA we have used Bayesian Networks, which serves to model these causal connections and estimate the probability of the main risk event based on the occurrence of CEs. In this paper, our focus is on a specific module of CERIA dedicated to assigning probabilities to the events in the causal network of events. Subsequently, these probabilities are utilized in the next module to assess the likelihood of the primary risk event. By "contributing event" we mean any event that, if it happens, can cause a risk event occurrence. This definition is pivotal to understanding the CERIA framework, as it directly influences the creation of the Bayesian Network and the assessment of risk probabilities. Through this approach, CERIA provides a systematic and dynamic method for early warning and risk management by anticipating and quantifying the impact of potential future disruptions based on current events.

2. Methodology

The CERIA framework [5] is structured with six interconnected modules, each playing a crucial role in the analysis and forecasting of risk events through the the news articles. (See Fig. 1). A summary explanation of each module is provided in the following section.

2.1. CERIA Framework

To develop a proactive model capable of identifying potential risks to SCs, the following research questions should be addressed:
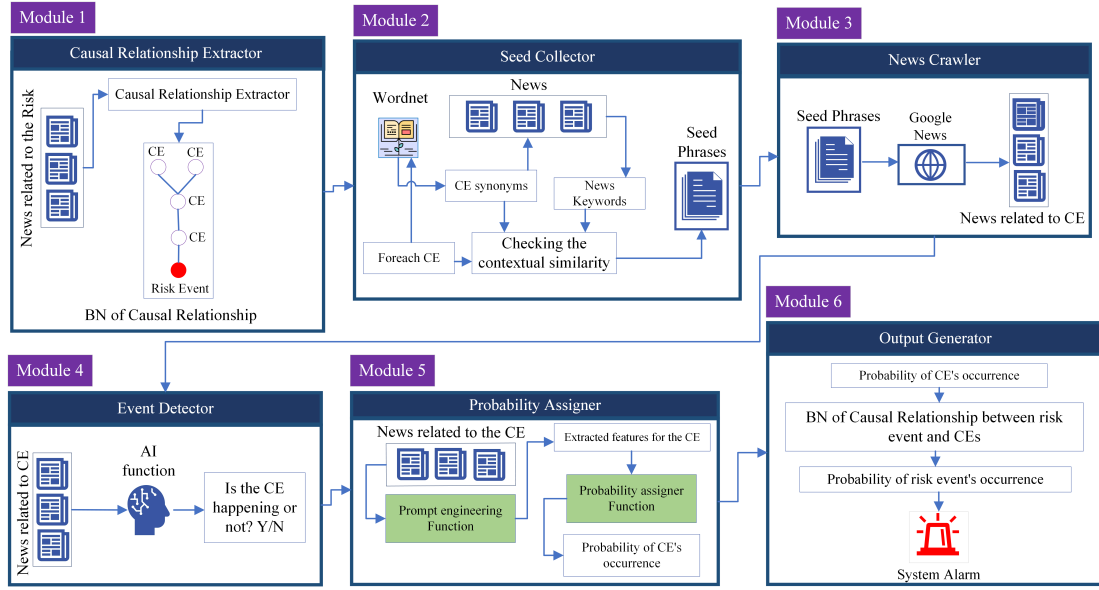
Figure 1: An overview of CERIA framework [5]

- **BN construction**: What are the causal relationships between events that ultimately pose or contribute to SC risks?
- **Event detection (direct inference)**: How can risk events be detected from textual sources? What are the relevant phrases?
- **Probability of risk event occurrence**: How can the chance of risk event occurrence be quantified?
- **CEs detection**: How can CEs be detected from textual sources? What are the relevant phrases?
- **CEs Impact assessment**: Do the detected CEs have an impact on the SC?
- **Probability of CEs occurrence**: How can the chance of CEs occurrence be quantified?

To answer these questions, the CERIA framework is structured into six distinct modules (See Fig. 1). The **Causal Relationship Extractor Module** undertakes a thorough analysis of historical news data to identify events that precipitated a risk event, which are called CEs. It identifies the causal connections between CEs, representing these relationships within a Bayesian Network (BN). The **Seed Collector Module** is crucial for gathering news related to CEs or risk event in the phase of risk identification. It identifies and aggregates the most relevant keywords for each risk event/CE into a collection known as seed-phrases. The **News Crawler Module** leverages the seed-phrases to searches for news articles related to each CE. Then the gathered news are passed to the next module. The **Event Detector Module** employs AI algorithms based on LLMs to evaluate each piece of news extracted by the previous module to ascertain its relevance to the targeted CE. The **Probability Assigner Module** is responsible for conducting a deeper analysis to assign a probability to the occurrence of the events, enhancing the predictive capability of the framework. Finally the **Output Generator Module** propagates

the occurrences of different CEs through the BN to forecast the probability of the occurrence of the main risk event and alarm the risk manager about the potential risk.

This paper focuses on the **Probability Assigner Module** which is preceded by the **Event Detector Module** that searches and scores the news articles. In the next sections of the paper, we briefly describe the **Event Detector Module** before investigating the details of the **Probability Assigner Module**

2.2. Event Detector Module

In detailing the analytical process, we begin by searching for seed phrases on the Google News search engine. The framework then analyzes the content of each news article based on its context. To do this, the body of each article is converted into a vector using the BERT LLM. Similarly, the event of interest is also translated into a vector form. Then, Cosine similarity is used to determine the contextual similarity between the event and the news article's content, with scores ranging from -1 to 1. A score closer to -1 signifies lower contextual similarity, while a score closer to 1 indicates higher contextual similarity.

Threshold determination: Upon reviewing and assessing the scores of around 2,500 news articles, our observations include:

- News articles with scores above 0.5 indicate that the CE is happening.
- The highest score observed for relevant news articles is 0.7.

The latter finding can be explained by our method of calculating similarity scores between a long document, which is the news article, and a short sentence as the CE, like "AFL matches will be happening." Achieving a score of 1, which denotes perfect similarity, is highly unlikely because it's rare for a large document and a short sentence to be exactly alike. Hence, a score above 0.5 is deemed relevant and is further processed in the next module, the Probability Assigner, where each news article has a score between 0.5 and 0.7 based on its similarity to the CE [5].

2.3. Probability Assigner Module

As the previous **Event Detector** module establishes a contextual similarity between the CE and the news article, **Probability Assigner** module of the CERIA is used to find the relevance of the event to a specific Supply Chain (SC) and the probability that the event will occur. More specifically, this module answers the following two questions.
1- Is the detected event impacting the SC?
2- What is the probability of occurrence of the event?
In the subsequent section, we discuss the specifics of the module, outlining how it addresses the aforementioned questions.

2.3.1. Is the detected event impacting the SC?

In order to find out if a specific event is impacting our SC or not, it is necessary to define key factors (features) for each event within our BN. This approach enables a structured analysis of

the event's impact on the SC.

Case study: We're considering a scenario of a delivery sector of a supply chain. Thus, we have defined following features for CEs.

1. **Location**: If the geographical area of the CE of interest is close to the SC of interest the identified event may impact the SC. By the supply chain nodes we mean areas such as manufacturing facilities, warehouses, or major transportation hubs.

2. **Time and Duration**: The timing and duration of an event (CE) are important factors which can tell us whether an identified event impacts our SC of interest or not. This timing factor includes the specific date and time it occurs and how long it lasts. For example, events during peak production or shipping periods may have more severe impacts. If the system identifies the truck drivers' strike, it is crucial to determine whether the strike is occurring for a single day, a week, or some other duration.

Prompt engineering approach: In order to extract the features of each CE, we utilized "GPT-3.5-Turbo", an optimized version of the GPT-3.5 language model developed by OpenAI [8, 9], to analyze the content of news articles. Subsequently, we formulated the following prompt to extract the result that we want related to the x^{th} detected CE:

Prompt= "within the text which is in triple backticks, answer these questions about CE_x:

1- where is the location of the CE_x?

2- What is the date of the CE_x.

3- For how long will the CE_x happen?

"' news_body"' "

Then for each CE of interest, the name of the CE would be passed to the prompt as CE_x and the body of the news with the highest similarity score would be passed as news_body in to the prompt.

2.3.2. What is the probability of the occurrence of the event?

As explained earlier, our system aggregates relevant news articles for each event. Each article receives a relevance score determined by its contextual similarity to the event, with scores ranging between -1 and 1. If an article is relevant, its score is 0.5 or more. If it's not relevant, its score is less than 0.5. This scoring mechanism is important for assessing the likelihood of an event's occurrence; On the significance of the score, it's important to note that, a higher contextual similarity score not only indicates a closer contextual match to the event of interest but also gives us more confidence about the event's occurrence. This added confidence comes from our specific approach to defining events using active verbs rather than only nouns. For instance, rather than identifying "rain" or "flood" as events, we use phrases like "flood is coming" or "it's raining" Consequently, the contextual similarity score reflects the degree to which news articles confirm that an event is happening, thus making our judgement of the event happening more reliable.

To refine our analysis, we break down the score range from 0.5 to 0.7 into three equal parts, reflecting different chances of the event taking place:

- `range 1, From 0.5 to 0.57` : indicating a low chance of occurrence.

- range 2, From 0.57 to 0.64 : indicating a medium chance of occurrence.
- range 3, From 0.64 to 0.7 : indicating a high chance of occurrence.

For each range, we calculate the proportion of news articles within that score range relative to the total number of news scored between 0.5 and 0.7. This proportion serves as an empirical estimate of the event's likelihood, categorized into low, medium, or high probability, as expressed in the formula:

$$P(\text{CE in range } x) = \frac{N_x}{N_{total}} \tag{1}$$

Where P(CE in range x) is the probability of the CE happening within the specified range x. N_x presents the number of articles with scores in range x. N_{total} represents the total number of articles scored between 0.5 and 0.7. This number is calculated for all three ranges of numbers and then used for further analysis, which tells us the percentage of which we are confident the probability of the occurrence of the event is low, medium and high.

3. Results and Discussion

To test the CERIA framework's performance, we utilized historical news data related to the transportation industry in Australia. Although these events have occurred in the past, for the purpose of our test, we assumed that these events were about to happen, as if we were testing right before they occurred.

Case Study: Focusing on the transportation sector of a SC in Victoria, Australia, and identifying **delay in product delivery** as a risk within this SC, our goal was to construct a BN of CEs that could lead to this risk.

3.1. BN Construction

In September 2021, the Victorian Government of Australia, announced mandates on vaccinations for construction workers as a response to the rising cases of COVID-19 linked to construction sites. The announcement sparked a significant public backlash, leading to protests and also issues in delivery of some of the products. By analyzing the past news related to the risk event of interest, the system found the following CEs as contributors to delay in delivery (output of Module 1 of CERIA):

1- There is an increase in Covid cases
2- Government mandates on vaccinations for construction workers
3- Construction workers hold strike
4- Blockade of the West Gate Bridge
5- Risk event: Delay in delivery of products

Fig 2 shows the structure of the BN, including nodes 1 to 4 as CEs and node 5 as the risk event.

Figure 2: Structure of BN containting nodes 1 to 4 as CEs and node 5 as the risk event of interest

3.2. Results

Given this chain of CEs, we assessed the framework's capability to identify any of these events by analyzing the news. For each of these CEs, seed phrases were extracted and utilized to search Google News for relevant news articles (outputs of modules 2 and 3 of CERIA). The relevance of these articles was determined by a scoring threshold (Output of module 4 of CERIA); articles with a score greater than 0.5 were classified as relevant and the corresponding CE was recognised as "occurring". Subsequently, for each detected occurring CE, CERIA compiled a dataset of news articles whose similarity scores fell within 0.5 to 0.7, as explicated in Section 2.2. Upon this foundation, the framework detected the occurrence of CEs 1 to 4 from the news articles. Then module 5 of CERIA extracted each CE's features and determined whether the detected event is related to our SC and if it has any impact on it. To do so, the text of the most relevant news article is passed into the OpenAI LLM to identify the CE's features. The framework's output for each detected CE, with the link for the most relevant news article is as follows:

1. **Event:** There is an increase in Covid cases [1]
   Extracted features: Location: Victoria, Australia - Time: September 13, 2021 - Duration: Not specified
2. **Event:** Government mandates on vaccinations for construction workers [2]
   Extracted features: Location: Victoria, Australia - Time: Announced on September 17, 2021, effective until September 23, 2021, at 11:59 pm. - Duration: Indefinitely, starting from the specified date
3. **Event:** Construction workers hold strike [3]
   Extracted features: Location: Melbourne, Australia - Time: September 21, 2021 - Duration: Not specified

---

[1] ABC News: Victoria records 473 new cases of COVID-19
[2] Important COVID-19 update: Mandatory vaccination for construction workers
[3] Protesters against vaccine mandate in Melbourne clash with police

4. **Event:** Blockade of the West Gate Bridge [4]
   Extracted features: Location: Melbourne, Victoria, West Gate Bridge - Time: September 21, 2021 - Duration: Not specified

As mentioned earlier, we assumed that this analysis is taking place at the time that these events were about to happen. Focusing on the transportation sector of a supply chain as our area of interest, the framework successfully identified all these CEs as events affecting our supply chain in Victoria, Australia, during September 2021. This implies the triggering of these CEs in our BN (Fig 2) which then leads to delay in delivery, as the risk event of interest. Consequently, the framework could detect the risk event's occurrence by identifying CEs 1 to 4. Even with the detection of just one among CEs 1 to 4 as occurring, the framework was capable of identifying the risk event by leveraging the BN's capacity to forecast an event from its antecedent events. The framework calculated the probability associated with each CE, employing the formula (see Formula 1) presented earlier (second output of module 5). The derived probabilities for each CE are documented in Table 1.

3.3. Discussion

Utilizing of LLMs to analyze text from news articles enables the identification of detailed event features—location, time and duration—critical for matching the events into the BN. This process allows for a sophisticated understanding of an event's impact on the supply chain. For example, if an "increase in Covid cases" is happening in New South Wales, but the supply chain operations are primarily located in Victoria, the system can intelligently ignore the event, recognizing it as not immediately relevant. This contextual analysis ensures risk management efforts are concentrated on directly relevant threats.

As it's shown in Table 1 the probabilities of occurrence for four key CEs within a simulated supply chain disruption scenario, as integrated into a BN, underscores the CERIA framework's advanced capability to quantify the occurrence of events. With a high probability assigned to the "increase in COVID-19 cases" and "government mandates on vaccinations", the framework adeptly identifies the high chance of increase in covid cases and mandatory vaccination rules by government. Meanwhile, the probabilities assigned to construction workers striking and the

---

[4]Melbourne protesters swarm West Gate Freeway, blocking traffic in both directions

Table 1
Probability of occurrence of four CEs

CE	Low chance	Medium chance	High chance
There is an increase in Covid cases	3.13%	10.94%	85.94%
Government mandates on vaccinations for construction workers	13.33%	28.89	57.78%
Construction workers hold strike	12.33%	35.71%	51.94%
Blockade of the West Gate Bridge	10%	33.25%	56.75%

blockade of the West Gate Bridge showcase the system's ability to detect the high probability of the occurrence of these two events. By establishing a causal chain of events (BN), the framework is equipped to activate the BN with any identified occurring events and forecast the subsequent risk event. This capability underscores the framework's advanced analytical power in navigating the complex dynamics of supply chain disruptions.

4. Conclusion

The CERIA framework represents a novel approach to supply chain risk management, combining the analytical power of LLMs with a structured Bayesian Network to predict and quantify events/risk probabilities. This approach not only aids in the early identification of potential risks but also enables businesses to prepare and mitigate these risks proactively. The successful application of the framework to a real-world scenario underscores its potential to revolutionize supply chain risk management, especially in industries susceptible to rapid changes and disruptions. Future work will focus on refining the model's predictive accuracy and exploring its application across different sectors to further validate its effectiveness in diverse supply chain environments.

References

[1] M. Schleifenheimer, D. Ivanov, Pharmaceutical retail supply chain responses to the covid-19 pandemic, Annals of Operations Research (2024) 1–26.

[2] X. Yue, D. Mu, C. Wang, H. Ren, R. Peng, J. Du, Critical risks in global supply networks: A static structure and dynamic propagation perspective, Reliability Engineering System Safety 242 (2024) 109728. URL: https://www.sciencedirect.com/science/article/pii/S0951832023006427. doi:https://doi.org/10.1016/j.ress.2023.109728.

[3] B. A. Odulaja, T. T. Oke, T. Eleogu, A. A. Abdul, H. O. Daraojimba, Resilience in the face of uncertainty: A review on the impact of supply chain volatility amid ongoing geopolitical disruptions, International Journal of Applied Research in Social Sciences 5 (2023) 463–486.

[4] G. A. Zsidisin, B. Gaudenzi, R. Pellegrino, Strategic Sourcing as an Enabler of Supply Chain Risk Management, Springer Nature Switzerland, Cham, 2024, pp. 1–10. URL: https://doi.org/10.1007/978-3-031-52592-6_1. doi:10.1007/978-3-031-52592-6_1.

[5] M. Shahsavari, O. Hussain, M. Saberi, P. Sharma, A lightweight and unsupervised approach for identifying risk events in news articles, 2023, pp. 37–43. doi:10.1109/ICDMW60847.2023.00014.

[6] L. Fan, L. Li, Z. Ma, S. Lee, H. Yu, L. Hemphill, A bibliometric review of large language models research from 2017 to 2023, ArXiv abs/2304.02020 (2023). URL: https://api.semanticscholar.org/CorpusID:257952516.

[7] S. K. Srivastava, S. Routray, S. Bag, S. Gupta, J. Z. Zhang, Exploring the potential of large language models in supply chain management: A study using big data, Journal of Global Information Management (JGIM) 32 (2024) 1–29.

[8] OpenAI, Gpt-3.5-turbo, https://openai.com, 2023.

[9] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam,

G. Sastry, A. Askell, et al., Language models are few-shot learners, Advances in Neural Information Processing Systems 33 (2020) 1877–1901.