Tracking the Evolution of Foresight Signals in News Data: The Case of the European Electric Vehicle Market

Karine Navasartian
University of Bremen, Germany
karine@uni-bremen.de

Abstract

To prepare for an uncertain future, organizations must continuously monitor emerging trends and early signals of change. The increasing availability of web-based textual data has boosted natural language processing (NLP) methods in strategic foresight, particularly in the scanning phase. While prior studies have extensively focused on the identification of signals in such data, considerably less attention has been paid to how these signals evolve over time and gain relevance as they become more visible. This study addresses this gap by examining whether tracking the temporal dynamics of signals can improve their assessment for strategic decision-making. Demonstrated on the use case of the European electric vehicle market, we find three dominant signal trajectories and show that burst dynamics tend to surface signal consolidation rather than the early detection of weak signals. The results indicate that foresight research should move beyond static, one-off analyses toward a dynamic temporal perspective capable of identifying signals at earlier stages of emergence.

1 Introduction

The early identification of emerging changes is crucial for organizations and industries seeking to plan strategically, anticipate risks, and maintain a competitive advantage. In an increasingly interconnected and complex world, the growing speed and variety of information, particularly from web-based sources, has led to the use of big data and machine learning to support environmental scanning and improve foresight processes (Muraro and Salles-Filho, 2024). Consequently, a growing body of literature has emerged on the detection of foresight signals in textual data streams.

Common NLP approaches largely build on the work of Yoon (2012), who propose the degree of visibility and degree of diffusion of signals in documents as indicators of emerging issues. With the advent of large language models (LLMs), more semantically informed approaches have become possible, enabling richer representations of emerging topics and signals (Ebadi et al., 2024b; Boutaleb et al., 2024). Despite these advances, weak signal analysis often remains a largely static exercise and rarely captures the continuous evolution of signals into stronger patterns or broader trends over time (Mühlroth et al., 2023; Rousseau et al., 2021).

One early attempt to incorporate temporal dynamics is provided by Ebadi et al. (2024b), who identify topics across the entire observation period and calculates the topic emergence map across different time frames to make signal development over time more visible. Nevertheless, the approach remains essentially static and is applied to scientific literature, which is generally less vivid and dynamic than web-based sources such as news media.

Other approaches more explicitly aim to track newly emerging signals over time. For example, Poumay and Ittoo (2022) show that contextual embeddings and similarity measures can improve the tracking of identified weak signals. However, their study still relies on keyword-based representations and does not support long-term signal tracking. A more recent contribution to neural topic tracking is presented by Boutaleb et al. (2024), who link topics over time based on the cosine similarity of cluster representations and dynamically classify topics as noise, weak signals, or strong signals according to popularity trends composed of topic size, growth, and temporal decay. Nonetheless, they do not provide systematic validation or a comparative evaluation of the resulting trajectories. This would be particularly relevant because weak signals are often sparse and fragmented in their early stages, increasing the risk of being overlooked (Mühlroth et al., 2023).

This paper addresses this gap by introducing a novel perspective on weak signal tracking that

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop), pages 352–362

explicitly accounts for trajectory dynamics. We combine neural topic modeling with a single-pass clustering algorithm to model the temporal evolution of signals and bursty time series analysis to analyze the resulting trajectories of signals. The proposed approach is evaluated on a multilingual corpus of European electric vehicle (EV) news, reflecting the diversity and heterogeneity of real-world web data. To validate the results, we examine three empirically relevant cases: concerns related to Chinese espionage, employment implications associated with the ban on combustion engines, and geopolitical dependence arising from shortages of critical minerals. We assess how this approach can inform strategic relevance by detecting how weak signals intensify and potentially mature into strong signals or broader trends. Hence, our research contributes to foresight by enabling the anticipation of emerging developments and allowing organizations to act early, prepare for external risks, and potentially secure first-mover advantages.

2 Literature Review

Strategic foresight contains the practice of anticipating early signs that may drive future change in order to inform strategic responses and sustain competitiveness in dynamic and uncertain environments (Brandtner and Mates, 2021). Within this context, environmental scanning represents a central framework for identifying and monitoring signals of change beyond an organization's immediate environment (Kohler, 2021).

Signals identified through scanning are commonly classified according to their level of maturity, ranging from weak signals over strong signals to established (mega-)trends (Ansoff, 1975). Hereby, signals are understood as general events that can inform future-oriented decision-making, whereas trends are conceptualized as streams of events with a clear directional development, reflecting a sustained increase or decrease in a given metric (Kohler, 2021). Accordingly, while weak signals are initially characterized by low visibility and uncertain interpretation, their strategic relevance may increase over time as they gain momentum (Kohler, 2021; Mühlroth and Grottke, 2018).

The increasing availability of large-scale textual data, particularly from web-based sources such as online news, has significantly expanded the role of NLP in environmental scanning (Mühlroth and Grottke, 2018; Muraro and Salles-Filho, 2024; Kohler, 2021). Consequently, a broad range of NLP-based methods has been proposed to automatically identify weak signals and emerging trends from textual corpora. Early methods for NLP-based signal detection largely build on Keyword Portfolio Maps, first introduced by Yoon (2012) and grounded in the theoretical framework proposed by Hiltunen (2008). In this framework, future-oriented signs are assessed along three dimensions: signal, issue, and interpretation. These dimensions capture the broader visibility and diffusion of a phenomenon as well as the clarity with which its implications can be understood. Building on this conceptual structure, Yoon (2012) operationalize the signal and issue dimensions by evaluating changes in keyword frequency and growth over time. While this approach provides a structured mechanism for identifying emerging keywords, its reliance on isolated keyword-based representations limits its ability to capture semantic context, which in turn makes signal interpretation more difficult. More recent research has therefore shifted toward clustering and topic modeling techniques that represent signals as groups of semantically related documents or terms (Mühlroth et al., 2023; El Akrouchi et al., 2021; Ebadi et al., 2024a).

For instance, Mühlroth et al. (2023) apply clustering methods to news articles in order to identify emerging signals across heterogeneous sources and aggregate them into semantically related clusters. Similarly, El Akrouchi et al. (2021) employ Latent Dirichlet Allocation (LDA) to extract topics from news data and introduce a weakness function to identify signals that are semantically distant from mainstream discourse. Temporal autocorrelation is then used to assess signal directionality, indicating whether a signal is gaining or losing prominence over time. However, both approaches focus on snapshots of short observation windows. While this is well suited to detecting novel developments in news streams, it lacks insight into historical patterns and the longer-term evolution of signals.

A further shift from lexical to contextual representations is proposed by Ebadi et al. (2024b), who apply neural topic modeling to extract semantically coherent and interpretable signals from scientific literature. Their topic emergence map extends traditional keyword-based maps by analyzing whole topics rather than individual keywords as proxies for signals. By dividing the observation period into discrete time windows and analyzing growth metrics separately within each window, the approach aims to assess whether signals persist and strengthen over time. Importantly, their findings underscore the need for continuous monitoring of signals and for a more systematic analysis of their temporal dynamics and evolutionary patterns. In particular, how weak signals can be continuously monitored and analyzed by dynamic metrics represents an open research question (Mühlroth et al., 2023; Ebadi et al., 2024a).

To address this challenge, topic tracking has been proposed as an approach for linking signals identified in different time periods. Early work by Xu et al. (2019) demonstrates that topics extracted using Latent Dirichlet Allocation (LDA) can be tracked over time through measuring the Jensen–Shannon divergence of their keyword distribution, enabling the analysis of topic change and evolution. However, because topics are represented as word distributions, this approach remains largely lexical, limiting its ability to capture semantic shifts in dynamically evolving and multilingual news corpora. The work of Poumay and Ittoo (2022) further shows that contextual embeddings and similarity measures can improve topic tracking. However, their study still relies on keyword-based representations and does not support long-term signal tracking. A more recent contribution toward neural topic tracking is presented by Boutaleb et al. (2024), who link topics over time based on the cosine similarity of their cluster representations and dynamically classify topics as noise, weak signals, or strong signals based on popularity trends, combining topic size and growth with temporal decay. However, validation of the results and a comparative evaluation of the detected trajectories is not provided.

Building on this line of research, the present study combines neural topic modeling with a single-pass clustering algorithm to detect and track the temporal evolution of signal trajectories and to assess their relevance based on recurrence and trajectory patterns. In doing so, it contributes to the literature on weak signal analysis in foresight by examining the added value of incorporating historical information into the environmental scanning process.

3 Data and Methodology

This section introduces the data and analytical methods used to detect and capture the emergence and temporal dynamics of weak signals.

3.1 Data

In this study, we focus exclusively on news articles and select the European electric vehicle (EV) market as our primary use case. To construct a suitable corpus, we collected data via the LexisNexis API (LexisNexis Developer Portal, 2025), which provides access to licensed and credible news content from established publishers. To ensure a European focus, we restricted the search to European news providers and queried EV-related keywords in the five most widely spoken European languages: English, German, Spanish, French, and Italian. In addition, LexisNexis industry and source-rank filters were applied to retain automotive-related content from high-quality and relevant publishers. This mitigates the risk of insufficient rigor, given that web-based data can be less reliable for environmental scanning when the credibility of publishing sources is limited (Harris and Brooker, 2025). The complete search string used for data retrieval is provided in Section A.

Following data collection, duplicate articles were removed based on title and URL. We further excluded irrelevant sources (e.g., stock-market-only pages) and filtered out articles in the lower and upper 2.5% of the article length distribution that did not contain meaningful news content (Adämmer and Schüssler, 2020). The resulting dataset comprises 56,994 articles published between October 2018 and September 2025. The distribution of the articles over time is depicted in Figure 1.

3.2 Signal Extraction

To extract signals at each time step, we employ the neural topic modeling framework BERTopic (Grootendorst, 2022). BERTopic uses contextual embeddings from transformer-based language models to cluster semantically similar documents. Unlike traditional keyword-based approaches such as Latent Dirichlet Allocation (LDA), BERTopic captures semantic relationships and contextual meaning. This results in more coherent and interpretable topics and has contributed to its increasing adoption in recent literature (Son and Park, 2025; Egger and Yu, 2022).

The BERTopic pipeline comprises four main steps: document embedding, dimensionality reduction, clustering, and cluster representation. We use the open-source, pre-trained multilingual embedding model *ibm-granite/granite-embedding-107m-multilingual* from Hugging Face to generate doc-

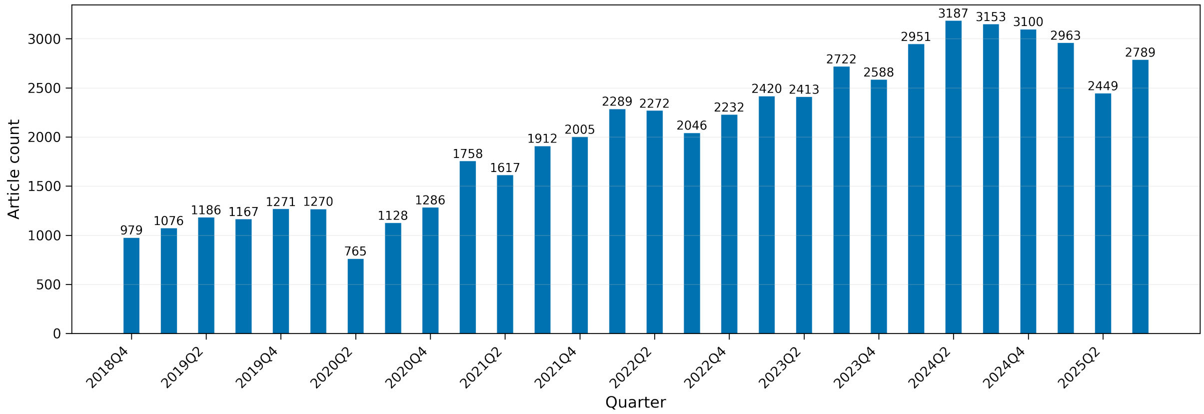

Figure 1: Distribution of news articles per quarter over the observation period.

ument representations. The model is based on the RoBERTa transformer architecture and was selected for its ability to embed all required languages into a shared semantic space while demonstrating strong clustering performance on the MTEB benchmark (Leaderboard, 2025). All hyperparameters were kept at their default values except for the critical parameter *min_cluster_size*, which was set to 5 (instead of the default value of 15) to allow for very small groups of articles to be detected and treated as potential weak signals.

For interpretability, cluster representations were generated using the generative language model Claude Sonnet 4.5, which produces concise, human-readable labels based on representative articles within each cluster. The final embedding representation of a topic is defined as the mean vector of all document embeddings assigned to that cluster.

3.3 Signal Tracking

To track signals over time, we apply the single-pass clustering algorithm commonly used in topic tracking (Xu et al., 2019; Boutaleb et al., 2024). This online clustering approach processes each signal exactly once, assigning it either to an existing cluster or using it to form a new cluster based on a predefined threshold. Once assigned, signals are not revisited or reassigned, enabling low computational complexity and scalability for temporally evolving corpora. Topic similarity across time steps is computed using cosine similarity between cluster embeddings.

The key parameter governing this process is the similarity threshold: values that are too low lead to excessive topic merging, while values that are too high prevent meaningful temporal linking. To determine an appropriate similarity threshold, we performed a parameter grid search over a range of threshold values and examined the resulting number of global topics (Figure 5 in Appendix B), following a procedure similar to that of Boutaleb et al. (2024). The results reveal a clear change in growth behavior around a threshold of approximately 0.95. However, this transition does not indicate an optimal threshold choice. To resolve this ambiguity, we constructed a small validation set consisting of temporally aligned clusters that should be merged based on semantic similarity. Clustering performance was then evaluated across the threshold range using this validation set. Based on this evaluation, a similarity threshold of 0.9475 achieved the highest F1 score and was therefore selected for all subsequent experiments. Additional details on the grid search are provided in Appendix B.

By merging similar clusters across consecutive time steps, the final output can be interpreted as a set of temporal network clusters k, each active only during specific months t. From these clusters, two time-series representations can be derived: (i) a binary appearance series indicating whether a signal is active in a given month, and (ii) a count-based series reflecting signal intensity. These representations enable downstream analysis of cluster trajectories to detect transitions from weak to strong signals.

To quantify such transitions, we draw on measures from temporal network analysis and conceptualize signal relevance as the re-bursting of previously weak clusters. Burstiness captures the extent to which events are temporally clustered rather than uniformly or randomly distributed. Following Kleinberg (2002), burstiness is well-suited for discrete event sequences with irregular temporal patterns. For each cluster, we construct the corresponding event sequence and compute inter-event times (τ). The coefficient of variation r quantifies how unevenly the inter-event times are distributed around their average. Higher values indicate that events tend to cluster closely together, followed by longer inactive periods, which is characteristic of bursty dynamics. However, the classical burstiness formulation is known to be biased for finite event sequences and systematically underestimates burst intensity when the number of events n is small (Kim and Jo, 2016). To address this limitation, we adopt the finite-size corrected burstiness measure $A_n(r)$ proposed by Kim and Jo (2016), which analytically accounts for finite length and preserves comparability across empirical settings. It maps the theoretical reference cases regular, random, and extremely bursty, to -1, 0, 1, respectively, thereby removing the finite-size ceiling effect. We compute $A_n(r)$ for each signal of a cluster, and interpret $A_n(r) > 0$ as evidence of bursting activity. The full definition is provided in Equation 1.

$$
\begin{aligned}
\tau_i &= t_{i+1} - t_i, \\
\mu &= \tfrac{1}{n-1} \sum_{i=1}^{n-1} \tau_i, \\
\sigma &= \sqrt{\tfrac{1}{n-1} \sum_{i=1}^{n-1} \tau_i^2 - \mu^2}, \\
r &= \sigma / \mu, \\
A_n(r) &= \frac{\sqrt{n+1}\, r - \sqrt{n-1}}{(\sqrt{n+1} - 2)\, r + \sqrt{n-1}}.
\end{aligned}
\tag{1}
$$

In the following we present the empirical results of the signal tracking analysis.

4 Results

This section first presents quantitative results on the data and the identified cluster structure, followed by three use cases to validate the findings. Given the inherent challenges of establishing ground truth for weak signal analysis, the evaluation adopts a qualitative approach, as is common in this research area, based on retrospective analysis of known outcomes and comparison with existing literature.

4.1 General Results

We first present the quantitative characteristics of the extracted signal trajectories for the use case of the European electric vehicle market. Over the observation period from October 2018 to September 2025, the volume of EV-related news articles increased substantially, more than doubling over time. As expected, the number of detected signals per month follows a similar upward trend. The temporal evolution of article volume and detected signals is shown in Figure 2. Under the selected parameter configuration, newly emerging signals—defined as clusters that cannot be linked to any previously observed cluster—account on average for 42% of all detected signals. This indicates a persistent inflow of novel themes and topics rather than a purely incremental evolution of existing discussions. Analyzing the resulting signal trajectories, we identify three distinct cluster types.

The first type, *singleton clusters*, consists of a single signal in a given month without any subsequent merges across the observation period. Since these signals are semantically isolated from all other clusters, they may encode potentially important but weak early information. In contrast, *mega clusters* emerge from repeated merges and represent broad, stable supertopics that dominate large portions of the discourse. While they can be decomposed further to analyze subtle shifts and changes in the discussion over time, they primarily reflect well-known and widely-covered themes. Table 1 provides selected examples of both cluster types. The third and most relevant cluster type of signal trajectories is the *sparse temporal network*. These networks are built of signals that recur only sporadically and irregularly over time, resulting in varying structures depending on the specific cluster. Some exhibit short activity over successive months, while others appear with substantial temporal gaps, or display gradual expansion. By analyzing these trajectories for bursty behavior we aim to detect those weak signals which possibly transition into stronger ones.

Using the metrics introduced in Section 3.3, we quantify the temporal dynamics of these sparse networks to detect re-occurrence and early bursts that may indicate increasing signal strength. In the following subsection, we present representative trajectories to validate this process.

4.2 Trajectory Analysis

For the analysis of signal trajectories, we excluded networks with fewer than 3 nodes, as they do not exhibit temporal structure and cannot be quantified by the burst metric. Across the 124 identified trajectories, 53 show at least one burst event during the observation period. Since there is no valida-

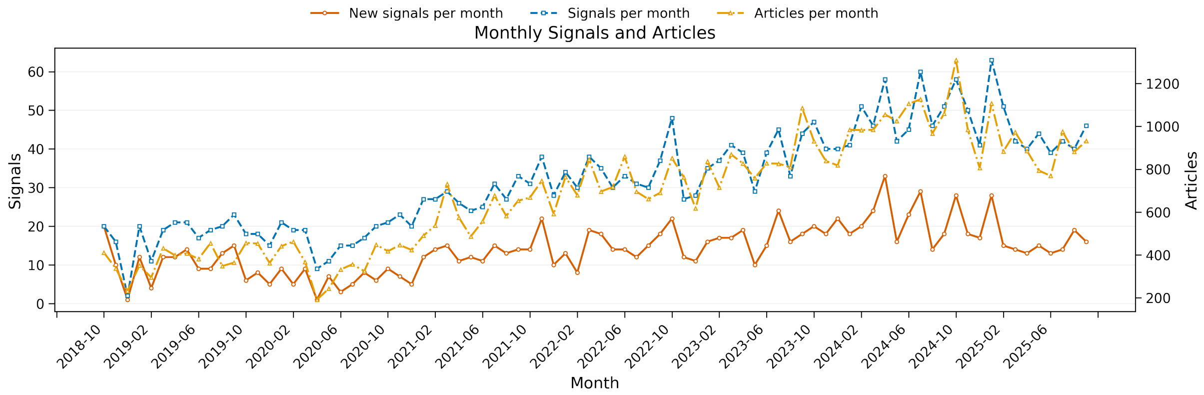

Figure 2: Monthly counts of articles, detected signals, and newly emerging signals over the observation period.

Cluster Type	Title	#Articles	#Signals	Date Range
Mega Cluster	EV Charging Infrastructure Development	5605	301	10/18–09/25
	EV Market Adoption and Sales Dynamics	2994	176	11/18–08/25
	Strategic Partnerships, Investments and Expansions	2166	145	10/18–09/25
Singleton	Silicon Carbide Chips Revolutionizing Electric Vehicle Efficiency	8	1	12/21
	Volvo the first Automaker to introduce EV Battery Passports	6	1	06/24
	Chinese EV Makers Face Major Safety Recalls Over Driver-Assistance Defects	8	1	09/25

Table 1: Examples of mega clusters and singleton signals identified in the European electric vehicle news corpus.

tion set of originating events and their transitions, weak-signal evaluation still remains largely qualitative in practice (Mühlroth et al., 2023; Boutaleb et al., 2024). We validate the practical applicability of the results by examining representative trajectories in which a burst follows an initial weak-signal phase and compare the observed patterns for the use cases to stages described in scientific and industry literature. While we identified different signal categories, such as new market entrants, strategic corporate moves, policy or investment announcements, we focus on three risk-related signals for detailed interpretation. The corresponding temporal trajectories are shown in Figure 3.

The first example is a relatively new signal that first appears in July 2023 and concerns the UK's national security risks associated with Chinese electric vehicles, specifically the potential for large-scale mobility data collection and remote control. While concerns in Western markets regarding China's rapidly growing auto exports had already circulated earlier, mainly around economic disruption due to lower prices, job losses, and competitive pressure for established brands, the security dimension only starts to emerge in 2024 and follows earlier technology-related narratives around Huawei and TikTok regarding data access and espionage (Wang and Xie, 2025). Hence, the first occurence of this signal in the news corpus is relatively early compared to general awareness. Media attention increases in early 2024 in the context of tariff debates and regulatory responses and continues to grow through 2025, as Chinese EV market shares rise in the United Kingdom. Using our approach, the third signal observed in March 2024 is already classified as a burst in discussion, representing the earliest point at which the cluster can be robustly identified. After the burst, first discussions about possible defense options for Western governments appear, mainly involving restrictions on Chinese EV imports and requirements that smart vehicle technologies come from trusted suppliers.

This development ultimately leads to a proposed ban in April 2025, indicating that the signal grew over time and allowed for interpretation and action, hence a transition from a weak to a strong signal could be determined.

The second now well-known risk relates to the availability of lithium and other critical battery materials for the electrification of the automotive industry and the resulting dependency of Europe on Chinese supply chains. Although this issue has been discussed in academic contexts in early 2012 (Beißwenger, 2012), it first appeared in our news corpus in 2019 due to historical data limitations. Around this time, the signal had reached some early maturity, as it was analyzed and assessed for future implications in literature (Vaalma et al., 2018). This becomes visible in the news discourse as well, since there is a notable gap until late 2020, when discussions resume around new mining investments, domestic extraction, and growing concerns that emerging battery alternatives could quickly devalue lithium-related assets. A burst is detected only at the end of 2021, at a stage when collaborations, industrial expansions, and strategic initiatives are already underway, and their consequences are being actively discussed in the media. Consequently, the burst emerges too late to enable anticipatory action, effectively capturing an already well-established strong signal rather than an early-stage development.

The third example concerns potential unemployment and social disruption arising from the transition to electric mobility. This issue, too, was previously addressed in the academic literature (Yoann Le Petit, 2017). Within the news corpus, the signal first emerges in mid-2021 in connection with the forthcoming ban on internal combustion engines in the European Union. The discourse is characterized by strong interventions from industry representatives, who increasingly position themselves as critics of the 2035 ban, due to societal and economic concerns. A burst occurs toward the end of 2022, at a time when the ban is finally signed and media discourse expands to assessments of its social consequences (Doll and Wietschel, 2022). After the burst, reporting shifts toward regulatory uncertainty, and arguments emerge that potential technological breakthroughs could make large-scale battery gigafactory investments obsolete if policy changes again. In this example, the burst mainly indicates a revival of debate in response to regulatory changes rather than sustained growth of the signal itself. The trajectory still allows for the comparison of different strategies within the EV transition, such as Ferrari's continued combustion-engine strategy, illustrating divergent industry responses. Hence, the continued monitoring of such trajectories remains important, even if the signal reaches a certain maturity.

5 Discussion and Conclusion

This study set out to explore whether leveraging burst metrics on the temporal evolution of signals extracted from large-scale news discourse can help detect progressions into strong signals for strategic foresight.

For our primary use case of the European electric vehicle news corpus, the analysis identified three characteristic trajectory types: singletons without temporal continuity, large and persistent mega clusters reflecting established discourse, and sparse, recurrent trajectories marked by irregular gaps and bursty behavior. Although the majority of signals fall into the first category and therefore do not permit temporal trajectory analysis, they remain highly relevant from a foresight perspective. In line with the research of Mühlroth et al. (2023), such signals can still serve as valuable inputs for early sense-making by deriving potential implications through expert judgment.

Moreover, the analysis of sparse recurrent trajectories indicates that signals detected early in the corpus often reappear as more consolidated and prominent discourses over longer time horizons. These extended delays between initial emergence and subsequent consolidation suggest that many weak signals are likely to be overlooked by globally aggregated modeling approaches, such as those employed in earlier work by Ebadi et al. (2024a). Although previous studies have proposed methods for linking semantically related signals across time (Xu et al., 2019; Poumay and Ittoo, 2022), to our knowledge, this study offers the first systematic comparison of long-term trajectories with analytical metrics.

By comparing the outcomes of three different empirical examples, the results highlight important limitations of burst-based relevance metrics. While burst detection reliably captures phases of rapid signal recurrence, they often do so only once a phenomenon has already reached a relatively advanced stage. As illustrated by cases such as lithium supply risks or concerns around Chinese in-

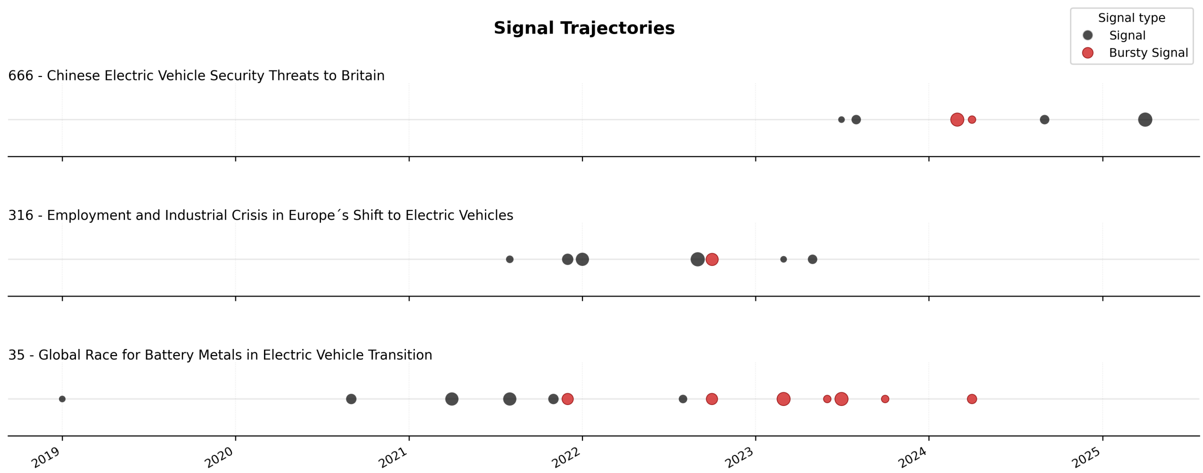

Figure 3: Signal trajectories over the observation period for the three selected use cases presented in the results section. Dot sizes are proportional to the article count. Red markers indicate periods in which a burst (burstiness score > 0) was detected.

dustrial espionage, burst detection tends to confirm the emergence of strong signals rather than functioning as an effective early-warning mechanism. Hence, this method is best understood as an indicator of signal maturation rather than anticipatory in nature. Similar limitations occur for related temporal metrics, such as those employed by Boutaleb et al. (2024), which also exhibit difficulties in early-stage signal detection. In their study, zero-shot topic modeling was applied using expert-defined fields of interest, implying that detection becomes more tractable when the target phenomenon is already known. Consequently, weak-signal detection and early-stage foresight require complementary approaches, including alternative temporal metrics, semi-automated analysis with AI-based methods and expert judgment (Geurts et al., 2022).

Nonetheless, the results highlight the potential of temporal signal tracking for strategic foresight and policy monitoring, particularly in contexts where emerging technological risks, regulatory shifts, or market developments unfold slowly and unevenly. By demonstrating that signals often reoccur at irregular intervals and precede observable strategic actions, the findings underscore the value of systematically integrating temporal signal analysis into environmental scanning processes. From a managerial perspective, the study suggests that such approaches can strengthen competitive positioning by enabling organizations to proactively adjust their strategies in response to emerging competitive actions and, in some cases, sustain first-mover advantages.

Building on these insights, the study points to several directions for future research. First, from a theoretical perspective, there is a clear need for the development and systematic evaluation of new metrics capable of capturing consistent weak-signal patterns at earlier stages. Second, linking observed signal trajectories to downstream market, technological, or regulatory outcomes would enable stronger validation of their predictive relevance. Finally, future work should explore how temporal signal-tracking approaches can be operationalized in real-world foresight workflows. In particular, the design of hybrid and iterative workflows that combine data-driven insights with expert interpretation, strategic reflection, and feedback from decision outcomes may more effectively support strategic decision-making under uncertainty.

Limitations

While demonstrating the initial potential of dynamic signal scanning for foresight, this study still has some limitations.

- First, systematic validation remains an open challenge. The study does not yet provide an external ground truth against which detected signal trajectories can be quantitatively assessed. Future research should therefore incorporate comparative evaluations against alternative methods and leverage expert-based validation to assess the relevance of identified signals and bursts.

- Second, the empirical analysis is limited to a single domain, namely the European electric vehicle market, which represents a relatively mature and well-institutionalized technological field. Applying the proposed approach to emerging or less stable technological domains may yield different trajectory patterns, raising questions regarding the generalizability of the observed signal dynamics across industries and stages of technological development.

- Third, the analysis focuses on a single temporal metric, which is effective for identifying signal consolidation but limited in its ability to detect relevant clusters at an early stage. This reinforces the need for complementary temporal indicators and AI-enhanced pattern recognition techniques to support early sense-making and anticipation in foresight applications.

- Finally, the selection and parametrization of the methods may be subject to legitimate debate. While the data preprocessing and topic modeling pipeline was carefully developed and grounded in prior research, certain modeling decisions were optimized for the specific characteristics of the present use case and may not transfer seamlessly to other empirical settings.

References

Philipp Adämmer and Rainer A Schüssler. 2020. Forecasting the equity premium: mind the news! *Review of Finance*, 24(6):1313–1355.

H Igor Ansoff. 1975. Managing strategic surprise by response to weak signals. *California management review*, 18(2):21–33.

Julia Beißwenger. 2012. Rohstoffmangel könnte den vormarsch der elektroautos bremsen. Deutschlandfunk.

Allaa Boutaleb, Jérôme Picault, and Guillaume Grosjean. 2024. Bertrend: neural topic modeling for emerging trends detection. *arXiv preprint arXiv:2411.05930*.

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

Claus Doll and Martin Wietschel. 2022. Aus für verbrennungsmotoren bis 2035: Statement zum beschluss des eu-parlaments zum ausstieg aus dem verbrennungsmotor für neue pkw und leichte nutzfahrzeuge bis 2035.

Ashkan Ebadi, Alain Auger, and Yvan Gauthier. 2024a. Harnessing machine learning and stock market techniques for signal detection in underwater sensing technologies. In *2024 International Conference on Machine Learning and Applications (ICMLA)*, pages 1747–1752. IEEE.

Ashkan Ebadi, Alain Auger, and Yvan Gauthier. 2024b. Wisdom: An ai-powered framework for emerging research detection using weak signal analysis and advanced topic modeling. *arXiv preprint arXiv:2409.15340*.

Roman Egger and Joanne Yu. 2022. A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts. *Frontiers in sociology*, 7:886498.

Manal El Akrouchi, Houda Benbrahim, and Ismail Kassou. 2021. End-to-end lda-based automatic weak signal detection in web news. *Knowledge-Based Systems*, 212:106650.

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

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

Benjamin Harris and Jennifer Brooker. 2025. Environmental scanning: A look to the future. *New Directions for Evaluation*, 2025(185-186):33–41.

Elina Hiltunen. 2008. The future sign and its three dimensions. *Futures*, 40(3):247–260.

Eun-Kyeong Kim and Hang-Hyun Jo. 2016. Measuring burstiness for finite event sequences. *Physical Review E*, 94(3):032311.

Jon Kleinberg. 2002. Bursty and hierarchical structure in streams. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 91–101.

Kevin Kohler. 2021. Strategic foresight: Knowledge, tools, and methods for the future. *CSS Risk and Resilience Reports*.

MTEB Leaderboard. 2025. `https://huggingface.co/spaces/mteb/leaderboard`. Accessed: 2025-03-12.

LexisNexis Developer Portal. 2025. `https://dev.lexisnexis.com/home`. Accessed: 2025-04-02.

Christian Mühlroth and Michael Grottke. 2018. A systematic literature review of mining weak signals and trends for corporate foresight. *Journal of Business Economics*, 88(5):643–687.

Christian Mühlroth, Laura Kölbl, and Michael Grottke. 2023. Innovation signals: leveraging machine learning to separate noise from news. *Scientometrics*, 128(5):2649–2676.

Vinicius Muraro and Sergio Salles-Filho. 2024. Big data, machine learning and uncertainty in foresight studies. *foresight*, 26(3):436–452.

Judicael Poumay and Ashwin Ittoo. 2022. Using meaning instead of words to track topics. In *International Conference on Applications of Natural Language to Information Systems*, pages 461–468. Springer.

Pauline Rousseau, Daniel Camara, and Dimitris Kotzinos. 2021. Weak signal detection and identification in large data sets: a review of methods and applications.

Hyunsang Son and Young Eun Park. 2025. Agenda-setting effects for covid-19 vaccination: Insights from 10 million textual data from social media and news articles using bertopic. *International Journal of Information Management*, 83:102907.

Christoph Vaalma, Daniel Buchholz, Marcel Weil, and Stefano Passerini. 2018. A cost and resource analysis of sodium-ion batteries. *Nature reviews materials*, 3(4):1–11.

Yongqi Wang and Yujie Xie. 2025. Legitimation contestation of the “chinese car threat” narrative in the digitized public sphere: a critical discourse analysis. *Humanities and Social Sciences Communications*, 12(1):1923.

Guixian Xu, Yueting Meng, Zhan Chen, Xiaoyu Qiu, Changzhi Wang, and Haishen Yao. 2019. Research on topic detection and tracking for online news texts. *IEEE access*, 7:58407–58418.

Yoann Le Petit. 2017. How will electric vehicle transition impact eu jobs? Briefing, September 2017.

Janghyeok Yoon. 2012. Detecting weak signals for long-term business opportunities using text mining of web news. *Expert Systems with Applications*, 39(16):12543–12550.

A Search string for LexisNexis

The following search string, depicted in Figure 4, was used to extract news articles for our use case of the European electric vehicle market.

B Grid Search for Cluster-Merging Threshold

Similar to the approach of Boutaleb et al. (2024), we conducted a parameter grid search to identify an appropriate similarity threshold for cluster merging. The objective was to select a threshold at which the number of resulting global topics remains relatively stable, avoiding both excessive aggregation and fragmentation of clusters. Figure 5 illustrates the relationship between the cluster-merging threshold and the resulting number of global clusters. Higher similarity thresholds lead to rapid topic fragmentation, whereas lower thresholds result in stronger aggregation and fewer distinct clusters. As no outstanding candidate could be identified from this analysis alone, we adopted a supervised threshold selection strategy. For this purpose, a ground truth reference was constructed manually by the authors based on all clusters identified for the year 2025. Clusters exhibiting substantively similar content were labeled as expected merge candidates. Candidate thresholds were evaluated over the interval

```
sourceRank:(1 2 3)
AND industry:Automotive
AND sourceRegion:Europe
AND language:("English" "Spanish" "Italian" "German" "French")
AND title:(
"Electric Vehicle" "Electric Vehicles" "EV" "EVs"
    "Electric Car" "Electric Cars" "E-vehicle" "E-vehicles"
"Elektrofahrzeug" "Elektrofahrzeuge" "E-Fahrzeug" "E-Fahrzeuge"
    "Elektroauto" "Elektroautos" "E-Auto" "E-Autos"
"Vehículo Eléctrico" "Vehículos Eléctricos" "VE" "VEs"
    "Coche Eléctrico" "Auto Eléctrico" "Autos Eléctricos"
    "Coches Eléctricos"
"Véhicule Électrique" "Véhicules Électriques"
    "Voiture Électrique" "Voitures Électriques"
"Véhicule E" "Véhicules E" "Veicolo Elettrico"
    "Veicoli Elettrici" "Auto Elettrica" "Auto Elettriche"
)
```

Figure 4: Search string used in LexisNexis API for retrieving European electric vehicle industry news articles.

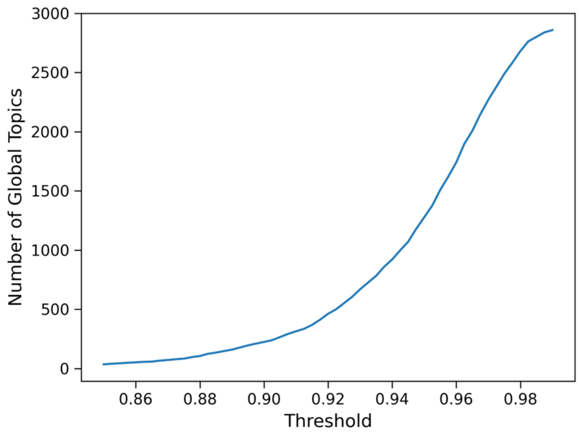

Figure 5: Relationship between the cluster-merging threshold and the resulting number of global clusters. Higher similarity thresholds lead to rapid topic fragmentation, whereas lower thresholds result in stronger aggregation and fewer distinct clusters.

$[0.90, 0.99]$. Precision and recall were used to assess the ability of each threshold to correctly merge these clusters. Based on this evaluation, the threshold achieving the highest F1 score (0.9475) was selected. The ten best-performing thresholds are reported in Table 2.

Threshold	**Precision**	**Recall**	**F1**
0.9425	0.576	0.556	0.566
0.9450	0.596	0.528	0.560
0.9400	0.532	0.587	0.558
0.9375	0.498	0.623	0.554
0.9475	0.634	0.488	0.552
0.9350	0.462	0.651	0.540
0.9500	0.651	0.444	0.528
0.9325	0.425	0.671	0.520
0.9300	0.392	0.702	0.503
0.9275	0.375	0.746	0.499

Table 2: Results of the similarity threshold grid search. The table depicts the top ten best performing thresholds on our validation set based on F1 score.