

Narrative Shift Detection: A Hybrid Approach of Dynamic Topic Models and Large Language Models

Kai-Robin Lange^{1,*}, Tobias Schmidt², Matthias Reccius³, Henrik Müller², Michael Roos³ and Carsten Jentsch¹

¹Department of Statistics, TU Dortmund University, 44221 Dortmund, Germany

²Institute of Journalism, TU Dortmund University, 44221 Dortmund, Germany

³Faculty of Management and Economics, Ruhr University Bochum, 44780 Bochum, Germany

Abstract

With rapidly evolving media narratives, it has become increasingly critical to not just extract narratives from a given corpus but rather investigate, how they develop over time. While popular narrative extraction methods such as Large Language Models do well in capturing typical narrative elements or even the complex structure of a narrative, applying them to an entire corpus comes with obstacles, such as a high financial or computational cost. We propose a combination of the language understanding capabilities of Large Language Models with the large scale applicability of topic models to dynamically model narrative shifts across time using the Narrative Policy Framework. We apply a topic model and a corresponding change point detection method to find changes that concern a specific topic of interest. Using this model, we filter our corpus for documents that are particularly representative of that change and feed them into a Large Language Model that interprets the change that happened in an automated fashion and distinguishes between content and narrative shifts. We employ our pipeline on a corpus of The Wall Street Journal news paper articles from 2009 to 2023. Our findings indicate that a Large Language Model can efficiently extract a narrative shift if one exists at a given point in time, but does not perform as well when having to decide whether a shift in content or a narrative shift took place.

Keywords

change point, narrative, story, Large Language Models, Latent Dirichlet Allocation

1. Introduction

With the rise of populism in western democracies, it has become increasingly important to evaluate narratives in politics, economics and other areas of interest that shape our society. Broadly speaking, narratives are linguistic constructs that boil complex connections between events down to an explainable form. Such narratives do not have to fall in line with facts and have become increasingly important in many areas of our society as they can act as a substitute for fact-based decision making. Consequently, methods to extract and analyze such narratives from large corpora have garnered the interest of many researchers. With both constantly changing and also conflicting narratives spread by entities such as policy makers, news outlets or social media personalities, we consider it especially important to observe, how a narrative develops over time instead of just globally extracting narratives from a corpus that spans over a long time period.

Along with the development and improvement of Natural Language Processing (NLP) methods, narrative extraction methods have also improved. The latest big development in NLP has brought a push towards more reliable human-like narrative extraction: the language understanding capabilities of Large Language Models (LLMs) such as ChatGPT [1], Llama [2], and others show the ability to detect narratives with increasingly complex definitions and thus blur the lines between qualitative and quantitative analyses. Given billions of parameters in the neural networks these models are based on,

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia, M. Litvak (eds.): *Proceedings of the Text2Story'25 Workshop, Luca (Italy), 10-April-2025*

*Corresponding author.

✉ kalange@statistik.tu-dortmund.de (K. Lange); tobias3.schmidt@tu-dortmund.de (T. Schmidt);

Matthias.Reccius@ruhr-uni-bochum.de (M. Reccius); henrik.mueller@tu-dortmund.de (H. Müller);

michael.roos@ruhr-uni-bochum.de (M. Roos); jentsch@statistik.tu-dortmund.de (C. Jentsch)

🆔 0000-0003-1172-9414 (K. Lange); 0000-0001-7330-4943 (T. Schmidt); 0000-0002-0716-0432 (M. Reccius);

0009-0004-1747-3303 (H. Müller); 0000-0002-5465-9893 (M. Roos); 0000-0001-7824-1697 (C. Jentsch)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

they can answer questions about almost any input document and the narratives contained within, but this enormous size also comes with drawbacks. Similarly to how expensive it is to pay experts of a certain field of interest to annotate texts, using LLMs for large scale corpora is not feasible for many researchers due to their computational demand or the financial cost of commercial models. Additionally, their size does not allow users to fully train the models on just a single corpus from scratch. Thus, many methods that use training from scratch to model temporal developments in a corpus, such as many dynamic topic models and diachronic word embedding models [3], cannot be directly transferred to LLMs, leaving fewer options for temporal narrative change detection.

We propose a pipeline that combines the best qualities of both approaches, using dynamic topic models and LLMs. Leveraging the approach of [4], we use a bootstrap-based topical change detection on the topics resulting from the dynamic topic model “RollingLDA” [5]. While this model succeeds in providing us with change points that are based on differences in the word count vectors of individual topics, it does not give us any intricate details about the changes themselves except for the topic of the change, the time chunk in which it happened and some key words that are responsible for the detection of the change. Furthermore, when a change is detected, it is not always certain whether it actually signifies a genuine shift in the narrative or in another dimension of the discourse, such as the factual content or the contextual framework provided. We can however use this method to narrow a large corpus down to a small curated number of documents that are suspected to contain information about some discursive shift. We do this by filtering the documents within the time period in which the change occurred, given the information we are provided about the change. We then use the language understanding capabilities of an LLM by processing these documents to explain the topical change that occurred and to decide whether the change signifies a narrative shift or not. We additionally provide the LLM with information gathered by the topic model to put the change into context. To accurately guide our identification of shifts in narratives, we use the Narrative Policy Framework (NPF) [6, 7], an analytical approach from political science. Our findings indicate that the LLM performs well when explaining a narrative shift, if one exists, but hallucinates when judging whether a detected change is a narrative shift or not, claiming most content shifts to also be narrative shifts.

We evaluate our model on a corpus of news articles of The Wall Street Journal ranging from 2009 to 2023. As the documents are copyright-protected, we do not use a commercial LLM in the cloud, but rather a local instance of the open source LLM Llama 3.1 8B [2].

2. Related Work

In our research, we focus on media narratives stemming from news articles with a business and finance focus. Thus, we base our definition of narratives on the existing literature on narratives from research in the field of economics and political economics.

2.1. Defining Narratives

The study of narratives has recently gained traction in economic research, though scholars have yet to converge on a single, universally accepted definition of the concept. Shiller [8], as an early and influential example, characterizes economic narratives as “stories that offer interpretations of economic events, or morals, or hints of theories about the economy”, thus providing a rather vague definition that leaves ample room for interpretation. A more formal modeling strategy was pioneered by Eliaz and Spiegler [9], who adapt concepts from the literature on Bayesian networks [10]. The authors highlight the causal connections among events and economic variables as sufficient for shaping people’s economic and political beliefs. While analytically appealing, this conceptualization sets aside value judgments, a crucial aspect of narratives. We hold that these normative implications ground the impact of political and economic narratives by motivating individuals and groups to act on their beliefs [11, 12, 13].

To empirically capture narratives, including their ideological valence, the Narrative Policy Framework (NPF) offers an alternative lens [6, 7, 14]. The NPF distinguishes narratives by their content and form, systemizing the latter through four elements: a setting, certain characters, a plot and a moral of the

story. The setting describes the scenery in which the narrative takes effect, such as a presidential election, a military conflict or a time of high inflation. The characters of a narrative can be persons or organizations, but also other entities take actions in the narrative including even non-sentient entities such as a spreading virus. The plot establishes relationships between the characters in space and time. Lastly, the moral of the story acts as a “takeaway” that often includes implicit or explicit calls to action. With these components, the NPF highlights both the role of ideological charge and the centrality of causal attributions, particularly by emphasizing that identifying “who is to blame for the problem” is an essential part of every narrative [15].

Similarly, Müller et al. [16] propose a narrative definition that is based on six key elements. They build their framework around the theory of media frames [17], proposing that narratives provide dynamic, evolving depictions of events. According to this view, a media narrative comprises one or more media frames (all of which are built around four key elements, see [17]) combined with protagonists (e.g., individuals and institutions) and events arranged chronologically and often presented as causally linked.

For a more comprehensive overview of narrative definitions across disciplines, see the overview papers by Roos and Reccius [11] or Santana et al. [18], among others.

Building upon foundational works in narrative theory, our research focuses on detecting narrative structures and their evolution in large text corpora. To achieve this, we draw on methodologies that have proven effective in identifying thematic shifts or *change points* within textual data. These change points often coincide with shifts in journalistic focus, and we hypothesize that such shifts frequently reflect underlying narrative changes. By leveraging established methods that combine topic modeling with change point detection, we aim to capture these transitions, providing valuable insights into how narratives develop and evolve over time.

2.2. Extracting Narratives from Text

Following the identification of discursive change points, the second critical step in our approach involves the automated analysis of the documents associated with these transitions. To date, there is no widely established (language) model specifically optimized for narrative extraction in this context. Instead, a diverse range of methods has emerged, each attempting to identify narratives or their components through diverse techniques ranging from word-count-based analyses [19] to Large Language Model-based methods [20].

A detailed review of existing NLP-methods to identify systematic parts of narratives is outlined by Santana et al. [18]. The authors focus on identifying key components such as events, participants, and temporal and spatial data, and linking these components to form coherent narratives. The methodology discussed includes part-of-speech tagging, event extraction, semantic role labeling, and entity linking, among others. However, this approach using classical NLP-tools struggles with challenges such as narrative complexity and cross-document narratives.

A similar approach, focusing on political and economic narratives, is proposed by Ash et al. [21]. Their methodology *RELATIO* employs semantic role labeling to identify key narrative components such as agents, actions, and patients within sentences, which culminates in the production of interpretable narrative statements. While their approach effectively identifies simple narrative building blocks, it falls short in detecting and extracting more complex narratives that integrate causality and sense-making. Lange et al. [22] advanced this approach by enhancing the existing *RELATIO* method with additional pre- and post-processing steps. By combining multiple *RELATIO*-extracted narrative blocks, the authors were able to link related statements and extract complex narrative structures that better align with causality-based definitions of narratives. At the same time, they emphasize that the increased complexity of their pipeline can amplify error cascades, where even minor changes in a longer input sequence may lead to significantly different results.

A recent study leveraging large language models (LLMs) for narrative extraction is presented by Gueta et al. [20]. The authors explore whether LLMs can effectively capture macroeconomic narratives from social media platforms like X (formerly Twitter). However, the study falls short of providing a robust definition of narratives, focusing instead on sentiment analysis and *RELATIO*-style statements,

leaving key aspects of narrative complexity and causality underexplored.

Collectively, these studies demonstrate the evolving landscape of narrative extraction methodologies, highlighting the integration of advanced NLP techniques to unravel complex narrative structures. By employing a hybrid approach, we seek to address the limitations identified in previous studies, particularly concerning the temporal dynamics of narratives and their computational feasibility. Our proposed pipeline, which combines dynamic topic models and LLMs, aims to provide a more comprehensive understanding of narrative changes over time, thereby contributing to the broader discourse on narrative extraction and analysis.

The choice to utilize LLMs for annotating and categorizing economic narratives stems from their demonstrated ability to excel at complex natural language processing tasks. Modern state-of-the-art LLMs, such as OpenAI’s GPT-4 [1] or Anthropic’s Claude 3.5 Sonnet, have consistently outperformed traditional NLP models in natural language understanding, classification tasks, and information retrieval. These models enable efficient processing of large corpora, reducing the time needed to annotate thousands of documents from weeks or months to just hours [23, 24, 25].

Unlike traditional NLP tools, which often rely on predefined models like sentiment analysis or topic modeling, LLMs can understand nuanced, contextual relationships in text. Previous approaches to analyzing economic narratives frequently employed machine learning pipelines [21], topic modeling [26], or sentiment analysis [27]. While these methods are effective for identifying broad patterns, as discussed earlier, they lack the depth to identify and categorize predefined, backward-looking narratives, particularly when such narratives involve subtle linguistic cues or complex causal relationships. On the other hand, studies in computational social sciences have shown that LLMs can match or even surpass human coders in annotating political, social, and economic texts [28], underscoring their potential for content analysis. For example, Mellon et al. [29] reported that LLMs achieved 95% agreement with expert annotators when analyzing British election statements. Similarly, Gilardi et al. [30] demonstrated that LLMs like GPT-3.5 could classify tweet content, author stances, and narrative frames more accurately than trained crowd workers. Building on this work, we explore the ability of LLMs to identify narratives and narrative-like structures in text. Detailed prompting and the continuous involvement of humans in the loop ensure that LLM annotations align as closely as possible with human intuition.

3. Methodology

To extract change points in our documents from our corpus, we use the Topical Changes method [4, 31], which is based on the models RollingLDA [5] and LDAPrototype [32], that improve the reliability of the classical Latent Dirichlet Allocation [33] and allows us to apply it to temporal data. After extracting the change points, we further analyze our documents using the LLM Llama 3.1 8B [2]. In this section, we detail the models’ functionalities and advantages for the task at hand.

3.1. LDAPrototype

Because the modeling of LDAs is inherently non-deterministic due to its sampling and initialization, there is no way of telling if a single given run does well to represent the corpus or whether it creates “bad” topics by chance. To prevent relying on randomness, we use LDAPrototype.

In terms of language modeling, LDAPrototype follows an LDA [33], but instead of training just one LDA, N LDAs are trained from which one model is chosen as a prototypical LDA. To “represent” the N models, the model with the highest average pairwise similarity to every other model is chosen as the prototype. To do this, Rieger et al. [32] proposed a similarity measure to compare LDAs with. All pair-wise combinations of N LDA models are compared by clustering their topics based on the cosine similarity of the topic’s top words into clusters of size two. The topic model similarity between models A_1 and A_2 is then given by

$$\frac{\text{\#topics of model } A_1 \text{ that are matched with a topic of model } A_2}{K}$$

with K as the number of topics in both models A_1 and A_2 . This similarity can thus be interpreted as the percentage of topics that is “more similar” to a topic of the different model rather than to a topic of the same model, so a form of topic overlap between two models. The chosen prototype will thus be one of the models with the highest average topic overlap and, conversely, the least “unique” topics that have not been generated by other LDAs and can therefore be considered the most stable out of the N generated models.

3.2. RollingLDA

RollingLDA [5] is a dynamic topic model based on LDA and can, by extension, also be used with LDAPrototype as a back-end. The model uses a rolling window approach that does not train an LDA from all time chunks at once, but rather begins to model the first w time chunks and then proceeds to model the remaining time chunks based on the information and topic assignments of the last m time chunks.

As the training of RollingLDA is thus based on the initially trained model, we use w instead of just one time chunk as a warm-up to ensure the model that initializes RollingLDA is properly trained. This parameter should thus be chosen sufficiently large that a proper topic model that covers most important repeating topics. For instance, choosing $w = 12$ when using monthly chunks, enables the model to initialize on the first 12 months of the data without a temporal component, ensuring its initial model has been trained on observing topics that return in yearly trends.

The memory parameter m creates the rolling window effect of the model. In each time chunk, we provide our model with previous topic assignments from the last m time chunks, which are considered while estimating the document-topic and word-topic distributions during the Gibbs-sampling step. This enables RollingLDA to efficiently model trends in temporal copropra, if the memory parameter is tuned accordingly. For instance, when using monthly chunks, $m = 4$ or $m = 3$ can be chosen to generate topics “remembering” quarterly trends, while “forgetting” older information from years prior. Only providing the model with a bit of past information allows it to change flexibly while still forcing it to keep coherent topics over time. This model is specifically designed to model abrupt changes in rapidly changing news media, which sets it apart from other dynamic topic models, such as the original dynamic topic model [34] and other early iteration of the idea [35, 36, 37].

3.3. Topical Changes

The topical changes model [4, 31] detects change points within topics put out by RollingLDA by comparing the development of word-topic vectors over time. For this, we observe both the current and previous count vectors of word assignments and compare the resulting similarity scores. We then perform a Bootstrap-based monitoring procedure that tests whether a change occurred.

To construct these word count vectors, we count the number of occurrences of each word in that topic over the last z time chunks (or since the last detected change, if it occurred less than z time chunks ago). Thus, z is the maximum number of time chunks the change detection can “look back” to, enabling a rolling window based change detection similar to RollingLDA’s rolling window based topic modeling. z can be tuned to, for instance, stop the detection from capturing repeatedly appearing trend effects by setting it as the assumed length of the trend. As the model focuses on finding abrupt changes rather than the slow natural development of language over time, a “mixture” parameter controls, how much language change the user expects from one time chunk to the next, which alters the look-back word topic vectors by mixing them with the current word vector to a certain degree.

The change detection is then performed by sequentially performing a bootstrap-based test with a significance level α for each time chunk in each topic. In this test, the cosine similarities between the word topic distribution in the current time chunks and B Bootstrap samples of the look-back word topic vector are compared. If the cosine distance of the observed look-back word topic vector to the current word topic vector is larger than $B * (1 - \alpha)$ of the bootstrap samples, a change is detected.

This detection does however not give indication about which tokens have actually caused the change.

Rieger et al. [4] propose to use words with high leave-one-out word impacts, that is words for which the cosine distance is reduced the most when leaving the word out of both word-topic vectors that are compared during the detection step. These words can thus be interpreted as the main causes for the detected change, as they had the highest impact on the drop of similarity.

3.4. Llama as a change interpreter

With the topical changes model, we are able to identify not only points in time where a change point in the topics of our corpus occurred, but also which topics are affected by this change. However, as topics are abstract constructs consisting of word distributions, the model can only give us abstract information about what changed. To gather more information, we further analyze our documents using a large language instruction model with great language understanding and summarization abilities. As we are handling copyrighted texts, we will not use a commercial model such as GPT [1], but rather an open source model that runs on a local machine. We use Llama 3.1 8B model, a instruction model designed by Meta with 8 billion parameters [2].

We first narrow down the number of potential documents to give our LLM as an input. Given a change point in topic k between the time chunks t and $t - 1$, we employ a filtering strategy that looks for change-related documents to feed into our LLM.

For this, we use the leave-one-out word impacts native to the Topical Changes model as a foundation. We count the number of occurrences of the words found to be significant due to the leave-one-out measure in each document in time chunk t . We then select those 5 documents with the highest count of these words and feed them into our model, letting it compare them to the topic of the previous time chunk, judging from 10 top words. We do not include documents from time chunk $t - 1$ here, as preliminary experiments showed that this “confused” the model. Leveraging from the Narrative Policy Framework, we give our LLM the following instruction, alongside the 5 chosen documents:

You are an expert journalist. You will be asked to explain, why a topical change in a corpus of news articles has been found and what the change consists of. To fulfill this task, you will be provided information from other text analysis models such as parts of the output of a RollingLDA topic model.

Whenever you are asked to analyze a “narrative”, assume the definition of a narrative that is laid out in the paper “The Narrative Policy Framework: A Traveler’s Guide to Policy Stories”. Specifically, respect and apply the following definitory aspects of a narrative: “The NPF posits that while the content of narratives may vary across contexts, structural elements are generalizable. For example, the content of a story about fracking told by a Scottish environmentalist is certainly different from the story told by a right-wing populist who attacks a public agency in Switzerland. However, these stories share common structural elements: They take place in a setting, contain characters, have a plot, and often champion a moral.” Keep in mind that a moral must feature a value judgement. When asked to specify a moral of a narratives, you must refer to this value judgement or note that there is no moral and thus no narrative! A narrative change must satisfy the four structural criteria, while a content change can simply be caused by an event that shifts the focus of the topic without a clear narrative. Your goal is to determine if a narrative change occurred or if it was a mere content change.

Please explain an apparent change within a RollingLDA topic that has occurred in [date]

The following topic top words might give you an idea of what the topic was about before the change: [10 top words of the topic in chunk $t - 1$]

The following topic top words might give you an idea of what the topic was about after the change: [10 top words of the topic in chunk t]

The following words were found to be significant to the detected change: [leave-one-out word impacts]

The following are those articles from the period that make the most use of the words

found to be significant to the detected change: [Filtered articles]

```
## Provide your output in a strict JSON format. First, summarize each article in one sentence:
{"summaries": [{"article_1": "...", {"article_2": "...", ...}]. Then formulate what the topic was
about before and after the change based on the topic top words, emphasizing the changes
induced to the topic, judged by the articles and the change words: "topic_change": ...
Explain how this change in topic indicates a shift in narrative. How did the narrative shift?
"narrative_before": "Before the change, the narrative centered around ...", "narrative_after":
"After the change, the narrative centers around ...". Finally, walk through the four structural
criteria that true narratives must satisfy according to the Narrative Policy Framework and
confirm or disconfirm their existence in the narrative after the break by briefly naming
what they are in the texts provided {"narrative_criteria": [{"setting": "...", {"characters": "...",
{"plot": "...", {"moral": "...}]]. Make sure to specify the exact source of the moral judgement
that you may have found. Lastly, make a final judgement if there is a narrative shift to be
found with {"true_narrative": True/False}. Do not answer in anything but JSON.
```

This filtering strategy enables us to specifically observe documents that are found by the model to be significant to the change. It does, however, not capture a larger picture of the topic itself and might lead the model to focus too much on a few significant words. This can happen if more than a few words are significant to the change with similar intensity, not just those that were captured by the leave-one-out word impacts. This might indicate an abrupt and broad topic change that shook the entire word topic distribution.

We also tested out different filtering strategies, such as providing the LLM with documents that are particularly representative of the topic k (i.e. documents with the highest topic share of k) in both time frames $t - 1$ and t , but they generally yielded worse results. We aim to further improve our prompting to further optimize our pipeline in the future.

4. Evaluation

We use Python 3.9 to run our scripts. We will publish the code corresponding to this paper as a part of the [Tools for Temporal Text Analysis \(tta\)](#) Python package[38].

We evaluate our model on a The Wall Street Journal data set containing 795,800 articles dating from 01/01/2009 to 12/31/2023. This high count of documents allows us to use small time chunks for our RollingLDA analysis. We choose monthly time chunks to enable a fine-grained analysis while ensuring that the number of detected changes remains manageable for annotators. We also choose our memory to check the last four months $m = 4$, thus enabling RollingLDA to remember quarterly trends in the data. To provide a stable initial LDA model and to incorporate all yearly trends into that initial model, we choose a warm-up period of $w = 12$ months. As a trade-off between computational efficiency and model reliability, we conduct 10 LDA runs to determine an LDAPrototype. We generated the initial model multiple times with $\{20, 30, \dots, 100\}$ topics and decided to choose $K = 50$ topics. Our Topical Changes model is then performed with a look-back-window of 4 months (also to remember for quarterly trends). We use a mixture parameter of 95% and evaluate the similarities to a significance level of $\alpha = 0.01$ to control the severity of changes that are observed to be major changes. We choose $B = 500$ Bootstrap samples to generate the bootstrap percentiles. For our LLM we set the temperature parameter to 0 in the hopes of minimizing hallucinations.

Our Topical Changes model found a total of 68 changes across 156 time chunks. In Figure 1 each topic is represented with its overall top words across the entire corpus in its title for interpretability. Each line plot includes a blue line, representing the word vector similarity, going from one chunk to the next while the orange line signalizes the dynamic threshold calculated using bootstrap samples. A change is detected when the blue line falls below the orange line, resulting in a vertical red line. At these change points, we filter the documents for candidate documents to feed into our LLM to check for a narrative shift.

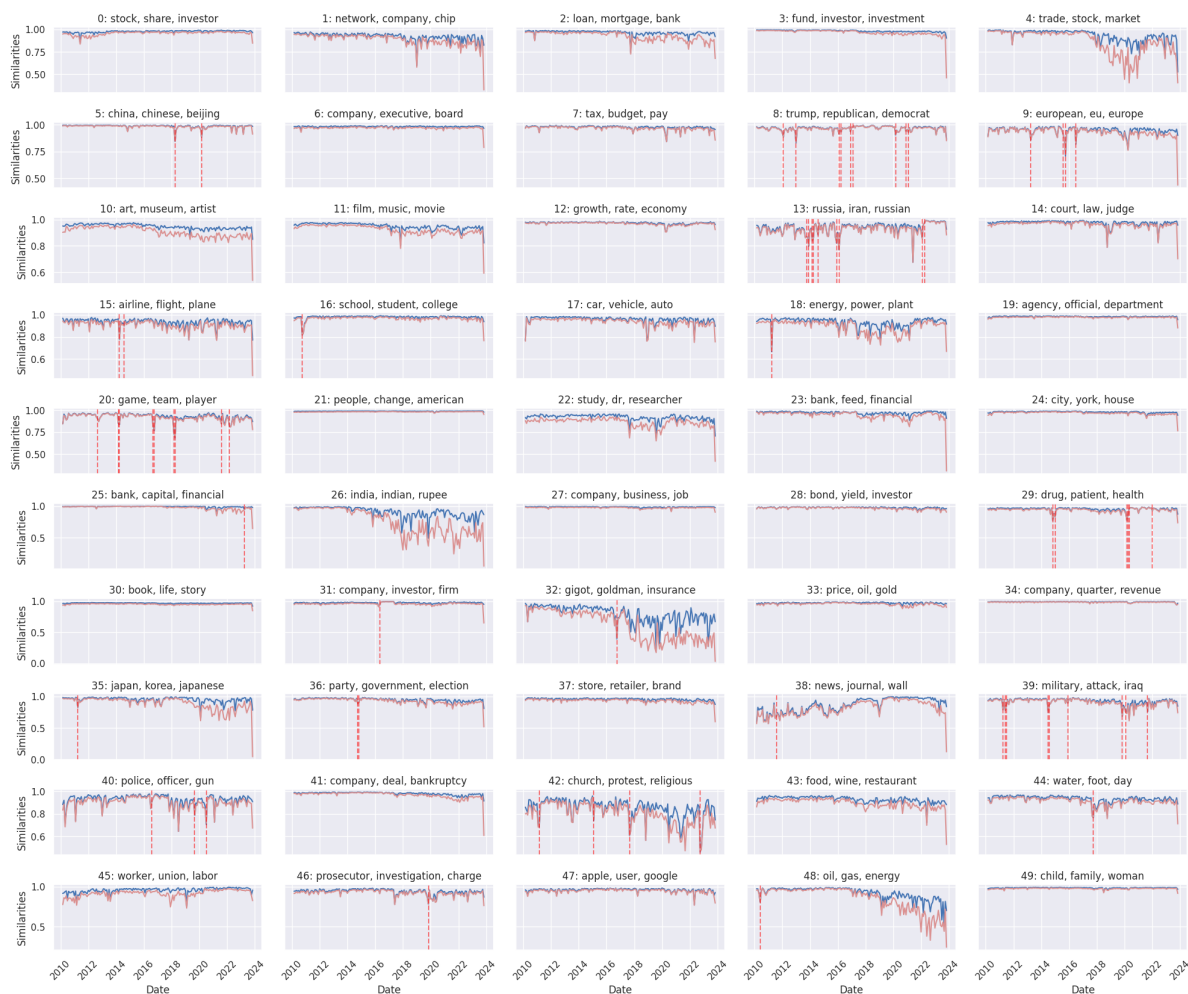


Figure 1: Results of the topical changes model with 50 topics on The Wall Street Journal corpus from 2010 to 2023. Vertical red lines indicate a detected change in the topic at that time chunk.

Three expert annotators discussed each change and annotated them according to the Narrative Policy Framework and afterwards classified the answers of the Large Language Model. They found 37 of them to contain narrative rather than mere content shifts. To exemplify Llama’s performance, we show three samples of detected changes as well as their narrative shift evaluation of the LLM in the appendix. We will release the list of annotated narratives upon publication of this paper in a [GitHub repository](#).

We split our evaluation into two parts: We see the binary classification between a narrative shift and a content shift as a first step. In the second step, we check if Llama correctly categorized aspects of the existing narratives according to the NPF. The results show that the LLM does not perform particularly well in classifying the changes into content shifts and narrative shifts, as it finds a narrative in 60, so all but eight cases. This results in a accuracy score of 57.35%, and an f1 score of 0.7010. Since the task involves Llama reproducing all 4 aspects of narratives laid out by NPF, this misclassification is likely caused by hallucinatory behavior, which is a well-known tendency of LLMs [39]. The LLM hallucinates trying to give a satisfactory output – ie. provide the user with aspects of narratives according to the NPF – at all times, resulting in a large false positive rate.

The model does, however, perform better when explaining an existing narrative shift. In those cases the model accurately defines the narrative in 31 out of 37, so 83,78%, of the cases.

Overall, the LLM performs well when a narrative shift exists for a given change point, accurately applying the NPF definition to capture the narrative. However, it does not perform as well when a mere content change occurs, stretching the definitional aspects of the NPF, thus hallucinating narratives. An improved prompt or an additional filtering step could help to solve this issue in future research.

5. Summary

When interpreting political or economic events, people align the corresponding information with their internal world view, combining the two into a narrative. Media narratives have become a big research topic in recent years due to the rise of spreaders of “simple”, often populist narratives. While recent advances in Natural Language Processing, namely the emergence of LLMs, have resulted in improvements in the task of narrative extraction from texts due to their language understanding capabilities, these models are resource intensive. Thus, using them to label narrative in large corpora is often not always feasible.

We introduce a novel pipeline that combines the scalability of the dynamic topic model RollingLDA and its extension Topical Changes with the language understanding of the LLM Llama 3.1 8B. The topic model is used to detect changes within a corpus of The Wall Street Journal dating from 2009 to 2023 in 50 topics. We use two filtering strategies to identify documents that contain information about the nature of the detected change or the before-and-after of the topic. These texts are then fed to Llama to analyze the change according to the Narrative Policy Framework and detect whether a mere shift in content or in the narrative took place.

We processed the articles in monthly time chunks and detected 68 changes within 13 years, as one year served as a warm up period. After manually labeling the changes, we find that 37 of those 68 changes show signs of narrative shifts. While our LLM managed to distinguish content shifts from narrative shifts only 57.35% of the time due to hallucinatory behavior, it correctly explained a narrative shift, if one exists, 83.78% of the times. The incorrect narrative shift detections stem from the model showing a preference to report a narrative shift for each input rather than considering the option of a change that is not caused by a narrative shift. While this result might improve with more careful prompting, it indicates that the language understanding capabilities of an LLM are functional enough to properly fit an existing narrative shift into a given complex definition, yet may fail not because of insufficient language understanding, but rather due to an inclination to oversatisfy the user’s prompt.

While we observed narratives as purely binary cases in this paper, we plan to perform a more nuanced evaluation of narrative extraction techniques in the future, considering a wider array of subjectivity at this complex task.

Limitations

While we demonstrate an exemplary usage of our pipeline here, parameter choices such as the size of time chunks for RollingLDA and the mixture parameter for the Topical Changes can impact the number and granularity of changes detected, which could also alter the results we observed in our study. Furthermore, we didn’t observe cases in which the Topical Changes model incorrectly detected a change, we therefore cannot tell, how the LLM would react to such a case.

Acknowledgments

This study is part of a project of the Dortmund Center for data-based Media Analysis (DoCMA) at TU Dortmund University. The work was also part of the Narrative Economic Alliance Ruhr (NEAR) project, supported by the Mercator Research Center Ruhr (MERCUR) with project number Ko-2022-0015.

References

- [1] OpenAI, J. Achiam, S. Adler, S. Agarwal, ..., B. Zoph, GPT-4 Technical Report, 2024. ArXiv:2303.08774 [cs].
- [2] A. Dubey, A. Jauhri, A. Pandey, ..., Z. Zhao, The Llama 3 Herd of Models, 2024.
- [3] A. Kutuzov, L. Øvrelid, T. Szymanski, E. Velldal, Diachronic word embeddings and semantic shifts: a survey, in: Proceedings of the 27th International Conference on Computational Linguistics,

- Association for Computational Linguistics, Santa Fe, New Mexico, USA, 2018, pp. 1384–1397. URL: <https://aclanthology.org/C18-1117>.
- [4] J. Rieger, K.-R. Lange, J. Flossdorf, C. Jentsch, Dynamic change detection in topics based on rolling LDAs, *Proceedings of the Text2Story’22 Workshop CEUR-WS 3117* (2022) 5–13.
 - [5] J. Rieger, C. Jentsch, J. Rahnenführer, RollingLDA: An Update Algorithm of Latent Dirichlet Allocation to Construct Consistent Time Series from Textual Data, in: *Findings of the Association for Computational Linguistics: EMNLP 2021*, Association for Computational Linguistics, Punta Cana, Dominican Republic, 2021, pp. 2337–2347. doi:10.18653/v1/2021.findings-emnlp.201.
 - [6] M. D. Jones, M. K. McBeth, A narrative policy framework: Clear enough to be wrong?, *Policy studies journal* 38 (2010) 329–353.
 - [7] E. A. Shanahan, M. D. Jones, M. K. McBeth, C. M. Radaelli, The narrative policy framework, in: *Theories of the policy process*, Routledge, 2018, pp. 173–213.
 - [8] R. J. Shiller, Narrative economics, *American economic review* 107 (2017) 967–1004.
 - [9] K. Eliaz, R. Spiegler, A model of competing narratives, *American Economic Review* 110 (2020) 3786–3816.
 - [10] J. Pearl, *Causality*, Cambridge University Press, 2009. URL: https://www.ebook.de/de/product/20478918/judea_pearl_causality.html.
 - [11] M. Roos, M. Reccius, Narratives in economics, *Journal of Economic Surveys* 38 (2024) 303–341.
 - [12] S. R. Shenhav, Political narratives and political reality, *International Political Science Review* 27 (2006) 245–262. doi:10.1177/0192512106064474.
 - [13] K. Eliaz, S. Galperti, R. Spiegler, False narratives and political mobilization (2024). doi:10.1093/jeea/jvae047.
 - [14] C. Schlauffer, J. Kuenzler, M. D. Jones, E. A. Shanahan, The narrative policy framework: a traveler’s guide to policy stories, *Politische Vierteljahresschrift* 63 (2022) 249–273.
 - [15] D. Crow, M. Jones, Narratives as tools for influencing policy change, *Policy & Politics* 46 (2018) 217–234.
 - [16] H. Müller, G. von Nordheim, K. Boczek, L. Koppers, J. Rahnenführer, Der wert der worte—wie digitale methoden helfen, kommunikations-und wirtschaftswissenschaft zu verknüpfen, *Publizistik* 63 (2018) 557–582.
 - [17] R. M. Entman, Framing: Toward clarification of a fractured paradigm, *Journal of communication* 43 (1993) 51–58.
 - [18] B. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, S. Nunes, A survey on narrative extraction from textual data, *Artificial Intelligence Review* 56 (2023) 8393–8435.
 - [19] N. Benner, K.-R. Lange, C. Jentsch, Named Entity Narratives, *Working Paper 962*, Ruhr Economic Papers, 2022. doi:10.4419/96973126.
 - [20] A. Gueta, A. Feder, Z. Gekhman, A. Goldstein, R. Reichart, Can llms learn macroeconomic narratives from social media?, *arXiv preprint arXiv:2406.12109* (2024).
 - [21] E. Ash, G. Gauthier, P. Widmer, Relatio: Text semantics capture political and economic narratives, *Political Analysis* 32 (2024) 115–132.
 - [22] K.-R. Lange, M. Reccius, T. Schmidt, Towards Extracting Collective Economic Narratives from Texts, *Ruhr Economic Papers* (2022). doi:<https://doi.org/10.4419/96973127>.
 - [23] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al., A survey on evaluation of large language models, *ACM Transactions on Intelligent Systems and Technology* (2023).
 - [24] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, Y. Du, C. Yang, Y. Chen, Z. Chen, J. Jiang, R. Ren, Y. Li, X. Tang, Z. Liu, P. Liu, J.-Y. Nie, J.-R. Wen, A survey of large language models, *arXiv preprint 2303.18223* (2023).
 - [25] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. T. Ribeiro, Y. Zhang, Sparks of artificial general intelligence: Early experiments with GPT-4, *arXiv preprint 2303.12712* (2023).
 - [26] A. Macaulay, W. Song, News media, inflation, and sentiment, in: *AEA Papers and Proceedings*, volume 113, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2023,

- pp. 172–176.
- [27] S. Tilly, G. Livan, Predicting market inflation expectations with news topics and sentiment, 2021. ArXiv:2107.07155 [cs].
 - [28] C. Ziems, W. Held, O. Shaikh, J. Chen, Z. Zhang, D. Yang, Can large language models transform computational social science?, *Computational Linguistics* 50 (2024) 237–291.
 - [29] J. Mellon, J. Bailey, R. Scott, J. Breckwoldt, M. Miori, P. Schmedeman, Do ais know what the most important issue is? using language models to code open-text social survey responses at scale, *Research & Politics* 11 (2024) 20531680241231468.
 - [30] F. Gilardi, M. Alizadeh, M. Kubli, Chatgpt outperforms crowd workers for text-annotation tasks, *Proceedings of the National Academy of Sciences* 120 (2023). doi:10.1073/pnas.2305016120.
 - [31] K.-R. Lange, J. Rieger, N. Benner, C. Jentsch, Zeitenwenden: Detecting changes in the German political discourse, *Proceedings of the 2nd Workshop on Computational Linguistics for Political Text Analysis@KONVENS2022* (2022). URL: <https://old.gscl.org/media/pages/arbeitskreise/cpss/cpss-2022/workshop-proceedings-2022/254133848-1662996909/cpss-2022-proceedings.pdf>.
 - [32] J. Rieger, C. Jentsch, J. Rahnenführer, LDAPrototype: A model selection algorithm to improve reliability of latent Dirichlet allocation, *PeerJ Computer Science* 10 (2024) e2279. doi:10.7717/peerj-cs.2279.
 - [33] D. M. Blei, A. Y. Ng, M. I. Jordan, Latent dirichlet allocation, *J. Mach. Learn. Res.* 3 (2003) 993–1022.
 - [34] D. M. Blei, J. D. Lafferty, Dynamic topic models, in: *Proceedings of the 23rd international conference on Machine learning, ICML '06*, Association for Computing Machinery, New York, NY, USA, 2006, pp. 113–120. doi:10.1145/1143844.1143859.
 - [35] X. Song, C.-Y. Lin, B. L. Tseng, M.-T. Sun, Modeling and predicting personal information dissemination behavior, in: *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining, KDD '05*, Association for Computing Machinery, New York, NY, USA, 2005, pp. 479–488. doi:10.1145/1081870.1081925.
 - [36] X. Wang, A. McCallum, Topics over time: a non-Markov continuous-time model of topical trends, in: *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '06*, Association for Computing Machinery, New York, NY, USA, 2006, pp. 424–433. doi:10.1145/1150402.1150450.
 - [37] C. Wang, D. Blei, D. Heckerman, Continuous time dynamic topic models, in: *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI'08*, AUAI Press, Arlington, Virginia, USA, 2008, pp. 579–586.
 - [38] K.-R. Lange, N. Benner, L. Grönberg, A. Hachcham, I. Kolli, J. Rieger, C. Jentsch, tta: Tools for temporal text analysis, 2025. URL: <https://arxiv.org/abs/2503.02625>. arXiv:2503.02625.
 - [39] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, P. Fung, Survey of hallucination in natural language generation, *ACM Computing Surveys* 55 (2023) 1–38.

Appendix

Sample 1

Topic: 29 (healthcare)

Date: 10/2014

Leave-one-out word impacts: ebola, drug, africa, worker, dallas

Narrative shift?: Yes

Hand-coded narrative shift: After a small ebola outbreak in the U.S. in the month prior and an appropriate response that condemned that outbreak, the focus of the topic shifts. Instead of the fear of an outbreak in the U.S., the articles focus on drug tests that give hope and also shift the focus back on the African continent.

Response of the LLM:

{

summaries: [

{article_1: "Pharmaceutical companies and health officials are racing to develop and test treatments for the Ebola virus, with several experimental drugs showing promise in early trials."},

{article_2: "The US response to the Ebola outbreak has improved significantly since the first case in Dallas, with officials learning from past mistakes and taking steps to contain the virus."},

{article_3: "The number of people potentially exposed to Ebola in the US has dropped to 50 from 100, and officials are working to reassure the public that the virus can be contained."},

{article_4: "A nurse in Texas has tested positive for Ebola, highlighting the need for hospitals to be more vigilant in their infection-control procedures and raising concerns about the preparedness of US hospitals for the disease."},

{article_5: "The Ebola outbreak in West Africa is being hindered by a lack of qualified staff, with many health workers in Liberia and other affected countries lacking proper training and equipment to combat the virus."}],

topic_change: "The topic shifted from a focus on the global response to the Ebola outbreak, including the development of treatments and the US response, to a focus on the challenges of containing the outbreak in West Africa, particularly the lack of qualified staff and the need for improved infection-control procedures.",

narrative_before: "Before the change, the narrative centered around the global response to the Ebola outbreak, with a focus on the development of treatments and the US response, emphasizing the sense of urgency and the need for action.",

narrative_after: "After the change, the narrative centers around the challenges of containing the outbreak in West Africa, highlighting the difficulties of providing adequate care and the need for improved infection-control procedures, with a focus on the human cost of the outbreak and the need for more effective solutions.",

narrative_criteria: [

{**setting:** "The setting of the narrative is the Ebola outbreak in West Africa, particularly in Liberia, where the lack of qualified staff and inadequate infection-control procedures are exacerbating the crisis."},

{**characters:** "The characters in the narrative include health workers, patients, and officials, who are struggling to contain the outbreak and provide adequate care."},

{**plot:** "The plot of the narrative revolves around the challenges of containing the outbreak, including the lack of qualified staff, inadequate infection-control procedures, and the need for improved solutions."},

{**moral:** "The moral of the narrative is that the Ebola outbreak highlights the need for improved global health infrastructure and the importance of prioritizing the health and safety of healthcare workers, as emphasized by Dr. Frieden, the CDC Director, who stated that 'preventing transmission requires scrupulous attention to infection control'."}

],

true_narrative: true }

LLM evaluation: LLM detected the fact that a narrative shift correctly. It correctly identified that the focus shifted back to the African continent and condemning the disease there. It did not focus on the emotional aspect of "hope vs. fear", which might be the result of not having been provided from the month prior to the change, resulting in a lack of context.

Sample 2

Topic: 36 (Elections and Polls)

Date: 10/2014

Leave-one-out word impacts: scotland, scottish, independence, neves, scot

Narrative shift?: No

Hand-coded narrative shift: The protagonists (and settings) in this topic have hardly anything in common. Documents deal with Catalonia's independence movement, Brazil's presidential election, and Scotland's decision to remain in the U.K., among others. There is no such thing as a common narrative here.

Response of the LLM:

{ summaries: [

{article_1: "Catalonia's independence movement remains strong despite a court injunction, with activists planning to hold a nonbinding referendum on November 9."},

{article_2: "Aécio Neves, the opposition candidate in Brazil's presidential election, is gaining momentum against incumbent President Dilma Rousseff, with polls showing him within striking distance of a runoff."},

{article_3: "Marina Silva, the third-place finisher in Brazil's presidential election, is considering endorsing Aécio Neves, the opposition candidate, in a bid to defeat President Dilma Rousseff in a runoff."},

{article_4: "The Brazilian presidential election is too close to call, with Aécio Neves and Dilma Rousseff neck and neck in the polls, with Neves focusing on the economy and Rousseff maintaining a loyal base of support."},

{article_5: "The Bank of England had planned to flood the financial system with cash and hold emergency auctions to secure banks' funding if Scotland had voted for independence in its referendum."}],

topic_change: "The topic shifted from a discussion of Scottish independence to a discussion of Brazilian politics and the presidential election, with a focus on the opposition candidate Aécio Neves and his chances of defeating incumbent President Dilma Rousseff",

narrative_before: "Before the change, the narrative centered around the Scottish independence movement, with a focus on the potential consequences of a 'yes' vote and the plans of the Bank of England to maintain financial stability",

narrative_after: "After the change, the narrative centers around the Brazilian presidential election, with a focus on the opposition candidate Aécio Neves and his chances of defeating incumbent President Dilma Rousseff",

narrative_criteria: [

{**setting:** "The setting is the countries of Scotland and Brazil, with a focus on the political and economic systems of each country."},

{**characters:** "The main characters are the politicians and activists involved in the Scottish independence movement and the Brazilian presidential election, including Aécio Neves, Dilma Rousseff, and Marina Silva."},

{**plot:** "The plot revolves around the potential consequences of a 'yes' vote in the Scottish referendum and the outcome of the Brazilian presidential election, with a focus on the opposition candidate Aécio Neves and his chances of defeating incumbent President Dilma Rousseff."},

{**moral:** "There is no clear moral judgment in the narrative, but the narrative does suggest that the opposition candidate Aécio Neves is a more viable option for Brazil's economy and that the incumbent President Dilma Rousseff's policies have been ineffective in addressing the country's economic problems."}]

true_narrative: true }

LLM evaluation: The model correctly detected some of the main characters. However, it did not come to the conclusion that, due to the different settings, the provided documents are more or less incomparable. The suggested moral may fit to a single document, but does certainly not consider all provided texts adequately.

Sample 3

Topic: 8 (US Elections)

Date: 01/2016

Leave-one-out word impacts: trump, house, iowa, cruz, sander

Narrative shift?: Yes

Hand-coded narrative shift: Just before the start of the 2016 Republican primaries, Donald Trump's campaign has gained momentum, overtaking his more mainstream competitors through his charisma and aggressive attacks on competitors, such as Ted Cruz and Marco Rubio, and despite questions surrounding his character and previous business dealings.

Response of the LLM:

{ summaries: [

{article_1: "The article discusses Donald Trump's financial struggles in the 1990s, including his debt restructuring and bankruptcy, and how he managed to come out of it by taking advantage of his creditors and using his charisma to negotiate deals."},

{article_2: "The article discusses the 2016 presidential election and how Donald Trump's campaign is gaining momentum, despite his high disapproval ratings and the objections of some voters, and how he may be able to overcome these obstacles by confronting his opponents and using his charisma to win over voters."},

{article_3: "The article discusses the 2016 presidential election and how Ted Cruz and Donald Trump are battling for the lead in Iowa, with Cruz's campaign accusing Trump of being too egotistical and Trump questioning Cruz's eligibility to be president."},

{article_4: "The article discusses the 2016 presidential election and how Donald Trump is targeting Ted Cruz, questioning his eligibility to be president and his evangelical faith, and how Cruz is handling the attacks with humor and poise."},

{article_5: "The article discusses the final debate before the Iowa caucuses, where the candidates, minus Donald Trump, sparred over policy differences and personal attacks, with Ted Cruz and Marco Rubio engaging in a heated exchange over immigration policy."}],

topic_change: "The topic shifted from discussing Donald Trump's financial struggles and personal characteristics to focusing on his presidential campaign and the 2016 election, with a particular emphasis on his rivalry with Ted Cruz and the other Republican candidates.",

narrative_before: "Before the change, the narrative centered around Donald Trump's personal characteristics, such as his financial struggles, his ability to negotiate deals, and his charisma, and how these traits might make him a successful businessman and president.",

narrative_after: "After the change, the narrative centers around Trump's presidential campaign and his rivalry with other candidates, particularly Ted Cruz, and how he is using his charisma and tactics to win over voters and gain momentum in the election.",

narrative_criteria: [

{**setting:** "The setting of the narrative is the 2016 presidential election, with the context of the Republican primary and the Iowa caucuses."},

{**characters:** "The main characters in the narrative are Donald Trump, Ted Cruz, and the other Republican candidates, who are portrayed as rivals and opponents in the election."},

{**plot:** "The plot of the narrative is the struggle for the Republican nomination, with Trump and Cruz engaging in a heated rivalry and the other candidates trying to position themselves as alternatives to the two front-runners."},

{**moral:** "There is no clear moral judgment in the narrative, but the story implies that Trump's ability to win over voters and gain momentum in the election is a testament to his charisma and tactics, rather than any inherent qualities or values."}]

true_narrative: true }

LLM evaluation: The model correctly identified the main characters, the setting and the plot of Donald Trump taking the leading role in the primary field. While the model somewhat overstated the importance of the "moral" implications of this development, all aspects of the NPF were correctly recognized and summarized.