Narrative Shift Detection

A Hybrid Approach of Dynamic Topic Models and Large Language Models

Kai-Robin Lange, Tobias Schmidt, **Matthias Reccius**, Henrik Müller, Michael Roos, Carsten Jentsch

TU Dortmund & Ruhr-Universität Bochum

Text2Story 2025
The 47th European Conference on Information Retrieval

April 10, 2025

Motivation

- Media **narratives** shape public discourse; understanding **narrative shifts** is crucial.
- Traditional methods (e.g. BoW-methods) are scalable but offer limited **interpretability** & language **comprehension**.
- Large Language Models (LLMs) have powerful language-understanding abilities, but are **expensive** to scale.
- We propose a **hybrid pipeline** combining **dynamic topic modeling** and **LLMs** to detect and interpret **narrative shifts**.

Defining "Narrative Shifts"

- Following the **Narrative Policy Framework (NPF)**[1]
  1. **Setting**: Where? (context or environment)
  2. **Characters**: Who? (heroes, villains, victims)
  3. **Plot**: Structure (events, causal link)
  4. **Moral**: Take-home point & value judgement
- A true **narrative** shift occurs if these elements change.
- Can we systematically detect narrative changes in large corpora?

[1] Shanahan et al. (2018), *The narrative policy framework*

Sequence-of-Methods overview

1) **Topic Model** (RollingLDA)[2]
2) Bootstrap **Change Detection** flags topical shift points[3].
3) **LLM-based annotation** (Llama 3.1 8B[4]) interprets **narrative structure** of shift points.

[2]Rieger et al. (2021), *RollingLDA: An Update Algorithm of Latent Dirichlet Allocation to Construct Consistent Time Series from Textual Data captures topics dynamically*

[3]Rieger et al. (2022), *Dynamic change detection in topics based on rolling LDAs*

[4]Dubey et al. (2024), *The Llama 3 Herd of Models*.

RollingLDA in a nutshell

- **Rolling window** approach to LDA:
  1. Foundation: Train initial LDA on first w $(= 12)$ time chunks to learn **stable** topic distributions.
  2. Rolling memory: iteratively incorporate next month, compare to past m $(= 4)$ chunks to update topics **without losing coherence**.
- Advantages:
  - **Temporal consistency** of topics
  - No forward-looking bias
  - Allows for **abrupt changes** when appropriate

Detecting Topical Changes

- For each topic-timechunk combination (k, t), compare LDAs (pooled) **word-topic vectors** to chunks $(t-1, ..., t-z)$:
- Apply **bootstrap-based tests** to cosine distances to see if distribution shifts significantly.

  $\rightarrow$ **Mixture parameter** accounts for steady evolution.

LLM-based Narrative Interpretation

- We use word-level *LOO*-impact to **identify words** driving the shift in the LDA (Rieger et al., 2021).
- After detecting a **topic change**, we:
  1. Identify the 5 **documents most affected** by *LOO*-impact words.
  2. Provide ① (plus relevant top words) as input to LLM.
  3. Prompt the model: Do **NPF structural elements** appear? ⇒ identify *narrative shifts* vs. mere *content shifts*.

Data and Experimental Setup

- Corpus: **The Wall Street Journal**, 2009–2023, ~800k articles.
- Monthly time chunks ($w = 12, m = 4$) for RollingLDA with $K = 50$ topics.
- Topical changes tested over the look-back window (z), significance level $\alpha = 0.01$.
- LLM: Local instance of Llama 3.1 (8B params) with **Temperature = 0**

Illustration of Detected Changes

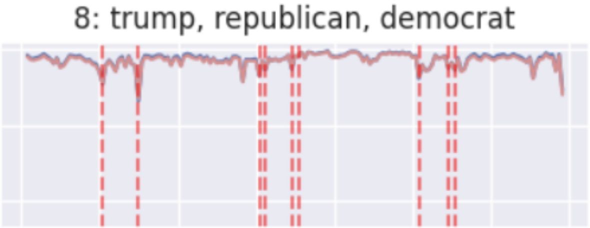

Figure 1: Topic 8 – US domestic politics

- **Red vertical lines**: significant topic shift.

Key Findings

- **68 topic changes** over 156 monthly chunks ($\sim$13 years).
- Expert coding: **37** true *narrative shifts* vs. **31** *content shifts*.
- $A_{\text{Llama}} \approx 57\%$, $\text{Recall}_{\text{Llama}} \approx 84\%$
- LLM **captures structural elements** accurately *when* a true narrative exists.
- **over-identification** of narratives (FP-rate drags performance).

Conclusion

Narrative Shift Detection **locates potential narrative shifts** within broader discourses.

- **Strengths**:
  - "Expert"-model principle – **efficient** use of computation
    1. RollingLDA handles evolving discourses.
    2. LLM captures **semantic depth** for narrative detection.
- **Limitation**:
  - LLM **hallucination** risk: RLHF induces "oversatisfaction" of user prompts
  - $\Rightarrow$ systematic **over-reporting** of narratives.
- Take care selecting hyperparameters (w, m, z, mixture rate)!

Thank you!

6 Example LLM Prompt

Prompt (Part I)

```
## You are an expert journalist. You will be asked to explain, why a
topical change in a corpus of news articles has 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.
```

Prompt (Part II)

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 specifiy 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.

Prompt (Part III)

```
## 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]
```

Prompt (Part IV)

```
## 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.
```