Stance Detection for Macro Topics based on Multi-factor Aggregation Analysis

Anonymous ACL submission

Abstract

This paper presents an innovative exploration into stance detection, with a specific focus on subjects characterized by their inherently abstract and macroscopic nature, termed as "macro topics." Due to the intricate complexity associated with these subjects, individuals often refrain from explicitly stating their opinions, thereby introducing challenges to stance detection when the target is implicit or unmentioned in the text. To address this complexity, we propose a tailored representation model designed to effectively encapsulate the nuanced aspects of macro topics. Our model relies on a comprehensive multidimensional analysis of sub-topics within a given macro topic, employing a specially designed discourse-based Latent Dirichlet Allocation (LDA) model. Utilizing this representation, an aggregation analysis is implemented to deduce stances on the macro topic by examining the array of sub-topic stances. The analysis of stances associated with sub-topics expressed in text is achieved by leveraging the semantic analysis capability of large language models (LLMs). Our approach attains superior stance detection accuracy, as validated through extensive experiments conducted on large-scale social media and finance text datasets.

1 Introduction

Stance detection aims to automatically discern an individual's opinionated perspective or attitude embedded in text regarding a particular entity or viewpoint, commonly referred to as the target (Augenstein et al., 2016). The scope of targets in stance detection is extensive, ranging from abstract concepts to tangible entities such as products and policies (Mohammad et al., 2016). The descriptions of these targets in the text exhibit considerable variability, with some targets being implicitly conveyed or, in certain instances, entirely absent from the text (Zhou et al., 2018). Consequently, it becomes imperative to predefine the target for which a stance judgment is required (Aldayel and Magdy, 2019).

Considerable research in stance detection addresses intricate scenarios and complex targets. Approaches for handling implicit or unexpressed targets include inferring their stance based on correlation or similarity of targets (Dong et al., 2017; Sobhani et al., 2019), identifying primary target claims from conversation sequences (Li et al., 2019), and generating public opinions through a combination of micro-level predictions (Qiu et al., 2015). Despite these advancements, to the best of our knowledge, there is currently no research on the stance analysis of the macro topics discussed in this paper.

This paper pioneers the exploration of a relatively complex challenge in stance detection, termed as "**macro topic stance detection**." This challenge emerges when individuals face difficulty expressing their attitudes directly toward abstract or unfamiliar concepts, collectively referred to as macro topics. However, a more explicit stance is often evident regarding concrete sub-topics associated with these macro topics. Analyzing these specific sub-topic stances facilitates the inference of stances towards the macro topics. To illustrate this concept, consider the following examples: 1) Economic Perspective: Economists predict macroeconomic trends, but their perspectives may be confined to specific facets, such as foreign trade or consumption. Aggregation of these diverse viewpoints is required for a comprehensive evaluation of the economic perspective. 2) Political Standpoints: In countries lacking distinct political party divisions, the political inclinations of the general populace—whether left-wing or right-wing—remain ambiguous. Individuals express viewpoints on issues such as the legalization of abortion, LGBT rights, etc. A thorough assessment of their political leanings demands the synthesis of diverse particulars and specific aspects of their stances.

Addressing the aforementioned challenges, this paper introduces a novel stance detection model tailored for macro topics. Given the absence of explicit information regarding the intended target of stance detection in the text, a **discourse-based topic modeling** approach is employed to extract pertinent topics (referred to as sub-topics) associated with a given macro topic. The judgment of the stance taken on a macro topic relies on an aggregation analysis of the stances identified within its corresponding sub-topics. To handle the diversity and openness of these sub-topics, a **zero-shot stance detection framework** is designed for identifying their stances, leveraging the semantic analysis capability of **large language models (LLMs)**. The contributions of this paper are as follows:

- Pioneering research is conducted on an intricate stance detection issue related to the macro topics, utilizing a combination of statistical analysis and LLMs to perform semantic analysis on extensive textual data.
- A discourse-based Latent Dirichlet Allocation (LDA) method is designed to facilitate topic modeling of short texts, especially suited for analyzing contents on online media platforms.
- A novel approach for stance analysis via sub-topic aggregation analysis is proposed, improving the interpretability of stance detection by analyzing the key factors influencing the stance of a given topic.

2 Related Work

2.1 Techniques Applied in Stance Detection

Stance is defined as the speaker's standpoint toward a given proposition (Darwish et al., 2017a). The prevailing focus in stance detection research involves the application of Natural Language Processing (NLP) techniques, framing the task as text entailment or classification (Siddiqua et al., 2019; Sobhani et al., 2019). The primary objective of this approach is to ascertain whether a specific piece of text supports or opposes the target proposition (Dias and Becker, 2016; Igarashi et al., 2016).

The data and features underpinning stance judgment encompass various types and sources, including: 1) linguistic features such as n-gram modeling (Hosseinia et al., 2020), sentiment polarity (Raghunathan and Kandasamy, 2023), and latent semantics (GomezSuta et al., 2023); 2) individual identity (Zhu et al., 2019; Darwish et al., 2020); and 3) social activity in social media, such as social connections (Darwish et al., 2020), retweets (Darwish et al., 2018), and hashtags (Dey et al., 2017).

The techniques employed for stance detection fall into three primary categories: 1) Supervised Learning (Lai et al., 2020) entails the use of classification techniques; 2) Weakly-supervised and Transfer Learning methods are implemented based on Pre-trained Language Models or Graph Convolutional Networks to model the relationship between the target and the text (Conforti et al., 2021; Li et al., 2022); 3) Unsupervised measures, such as clustering (Rashed et al., 2021).

Large Language Models (LLMs) such as ChatGPT achieve state-of-the-art or comparable performance on widely-used stance detection datasets (Zhang et al., 2022). Notably, ChatGPT offers explanations for its predictions, a feature absent in existing models (Zhang et al., 2022). However, potential biases towards specific targets have been identified (Zhang et al., 2023).

Our approach combines supervised and unsupervised techniques to analyze macro topics through diverse information facets, including text and statistical features, to discern individuals' stances. In terms of interpretability, unlike ChatGPT's direct reasoning for stance judgment, our method interprets the stance on macro-level topics based on computed weights of their sub-topics.

2.2 Intractable Targets in Stance Detection

Different from sentiment analysis tasks such as Aspect-Based Sentiment Analysis, stance detection tasks often involve cases where the target of the stance is not explicitly mentioned in the text (Hardalov et al., 2022). For instance, a tweet stating "Jeb Bush is the only sane candidate in this Republican lineup" may express an 'against' stance towards the unmentioned topic "Donald Trump as President". Moreover, stance detection primarily addresses ideological topics (e.g., atheism, media bias), posing challenges for accurate identification (Alturayeif et al., 2023).

To address implicitly expressed targets in stance detection, various strategies have been proposed, including Multi-Related-Target Stance Detection (Darwish et al., 2017b), Claim-Based Stance Detection (Kochkina et al., 2017), Collaborative Filtering-based approaches (Gottipati et al., 2013), and Target-Independent Models (Alturayeif et al., 2023). Another notable category is Stance Prediction, which aims to infer social media users'

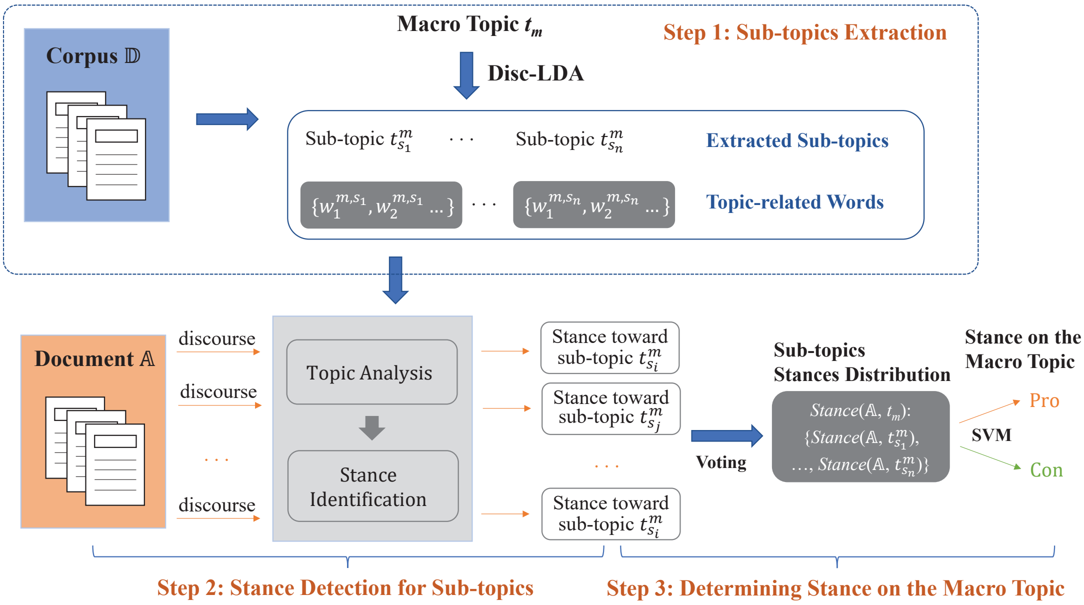

Figure 1: The architecture of STS4MTS. In Step 2, "Topic Analysis" identifies the topic for each discourse, while "Stance Identification" determines the discourse's stance on the topic. In Step 3, a voting mechanism aggregates discourses' stances to derive an overall stance for each sub-topic of t_m.

stances when not explicitly expressed in their interactions, and even forecast their viewpoints on forthcoming events (Qiu et al., 2015).

Existing stance detection methods typically focus on single or multiple sentences to determine stances. However, this article addresses the challenge of "macro topic stance detection," which involves making stance judgments on broad topics using corpora comprising tens of thousands of sentences. The input data at this scale exceeds current models' capacity. Stance judgment for subtopics, based on a limited number of sentences, is well-suited for LLMs due to the diversity and openness inherent in these sub-topics (Zhang et al., 2023). Nevertheless, selecting appropriate prompt templates for each sub-topic remains a significant challenge (Zhang et al., 2022).

3 Methodology

The problem of macro topic stance detection is formally defined as follows: Given a document $\mathbb{A}$ (e.g., a collection of articles from an individual) containing discussions on a diverse set of topics, the objective is to determine the stance regarding a macro topic t_m, which may not be explicitly mentioned within the given document.

The proposed model for macro topic stance detection is abbreviated as **STS4MTS** (Sub-topic Stances for Macro Topic Stance), comprising three key steps (illustrated in Figure 1):

1) **Sub-topics Extraction**: Acquire a set of subtopics $T^m_s = \{t^m_{s1}, ..., t^m_{sn}\}$ associated with t_m, based on a corpus $\mathbb{D}$ (e.g., a collection of articles from multiple individuals).

2) **Stance Detection for Sub-topics**: Obtain the stances expressed in the discourses within $\mathbb{A}$ regarding each sub-topic in T^m_s.

3) **Determining Stance on Macro Topic**: Determine the stance on t_m through aggregation analysis of the stances of all sub-topics in T^m_s.

3.1 Sub-topics Extraction for a Macro Topic

The task of extracting associated sub-topics for a given macro topic from a corpus resembles the effort involved in aspect/feature mining within the domain of product opinion mining and aspect-based sentiment analysis (Wang et al., 2019). Given our focus on analyzing an extensive collection of texts, we opted for the technique of corpus-level aspect mining, often facilitated by topic model-based approaches, such as LDA (Brody and Elhadad, 2010).

Given the prevalent use of **brief textual forms**, such as tweets, for expressing stances on specific topics, insights are drawn from the concept of Sentence Latent Dirichlet Allocation (SLDA) (Balikas et al., 2016; Büschken and Allenby, 2016) and an approach termed Discourse-based LDA (**Disc-LDA**) is proposed in this paper to ascertain the topic distribution of **short articles or discourses**, thereby facilitating the generation of a sub-topic list from the corpus. A discourse can take various forms, such as a tweet or a paragraph in an article. **Disc-LDA** incorporates the assumption that all words within a discourse originate from a singular topic, aligning well with empirical data patterns, especially in online social networks.

Algm. 1: Text Generation with **Disc-LDA**.

☐ **Topic Plate**:
for *every topic* $t \in [1, T]$ **do**
  Draw a word distribution for t: $\phi_t \sim Dirichlet(\beta)$
end
☐ **Document Plate**:
for *every document* $d \in [1, D]$ **do**
  Draw a topic distribution for d: $\theta_d \sim Dirichlet(\alpha)$;
  Sample a discourse number for d: $R_d \sim Poisson(\xi_{doc})$;
  ☐ **Discourse Plate**:
  for *every discourse* $r \in [1, R_d]$ **do**
    Sample a word number for r: $W_{d,r} \sim Poisson(\xi_{disc})$;
    Sample a topic for r: $t_{d,r} \sim Multinomial(\theta_d)$;
    ☐ **Word Plate**:
    for $w \in [1, W_{d,r}]$ **do**
      Sample a word for w: $word \sim Multinomial(\phi_{t_{d,r}})$;
    end
  end
end

Suppose the corpus $\mathbb{D} = \{d_1, d_2, ..., d_D\}$, from which the sub-topics of t_m are extracted. Each d_i, $i \in [1, D]$, is a document that represents a collection of discourses from an individual. The generative process with Disc-LDA proceeds through the steps demonstrated in Algorithm 1. The meanings of some notations are elucidated in Table 1.

In contrast to traditional LDAs (Blei et al., 2003), Disc-LDA incorporates an additional component called "Discourse Plate", dedicated to generating a unified topic for each discourse. We assume symmertric Dirichlet priors, i.e., the values of α and β remain consistent across all documents and topics, and they are model hyper-parameters. ξ_{doc} and ξ_{disc} are statistical values that can be acquired from $\mathbb{D}$. Different values of T can be tested experimentally to determine the optimal settings, and finally

Nota.	Meaning
D	Number of documents in the corpus $\mathbb{D}$
T	Number of topics within $\mathbb{D}$
R_d	Number of discourses in document d
V	Size of the word vocabulary
ϕ	Distribution over words for a topic
θ	Distribution over topics for a document
α	Dirichlet prior for θ
β	Dirichlet prior for ϕ
ξ_{doc}	Parameter of Poisson distribution for the number of discourses in a document
ξ_{disc}	Parameter of Poisson distribution for the number of words in a discourse
t_r	The topic of discourse r
$\vec{t}_{\neg r}$	The topic assignments for all discourses except discourse r
$\vec{w}_d$	The word sequence of document d
$C_k^{\neg r}$	Number of discourses assigned topic k in document d excluding r
$F_{v,k}^{\neg r}$	Frequency of the word v assigned the topic k in discourses excluding r
N_r	Total number of words in discourse r
$N_{r,v}$	Number of occurrences of word v in r

Table 1: Meaning of Notations in Disc-LDA.

T sub-topics are generated for t_m, forming the set $T_s^m = \{t_{s1}^m, t_{s2}^m, ..., t_{sT}^m\}$, where each sub-topic $t_{s_i}^m$ is associated with a set of **topic-related words** $TW(t_{s_i}^m) = \{w_1^{m,s_i}, w_2^{m,s_i}, ...\}$.

For inference, the topic distribution for a document d is calculated using collapsed Gibbs sampling and approximated by the full conditional $p(\vec{t_r}|\vec{t}_{\neg r}, \vec{w})$ (Heinrich, 2005) as follows:

$$p(t_r = k|\vec{t}_{\neg r}, \vec{w}_d) \propto \frac{C_k^{\neg r} + \alpha}{\sum_{t'=1}^{T}(C_{t'}^{\neg r} + \alpha)} \times$$

$$\frac{\Gamma(\sum_{v=1}^{V}(F_{v,k}^{\neg r} + \beta))}{\Gamma(N_r + \sum_{v=1}^{V}(F_{v,k}^{\neg r} + \beta))} \times \prod_{v=1}^{V} \frac{\Gamma(F_{v,k}^{\neg r} + \beta + N_{r,v})}{\Gamma(F_{v,k}^{\neg r} + \beta)} \quad (1)$$

$p(t_r = k|\vec{t}_{\neg r}, \vec{w}_d)$ is the conditional probability that the topic of r is k, given the complete set of words and topics for all discourses except r. The initial derivation of Equation (1) is detailed in (Balikas et al., 2016) and Section 5.5 of (Heinrich, 2005). Our primary contribution is the extension of sentence-based topic sampling (Balikas et al., 2016) to discourse-based topic sampling.

3.2 Stance Detection for Sub-topics

The problem of stance detection for sub-topics can be formally described as follows: given a document $\mathbb{A} = \{r_1, ..., r_R\}$, where r_i, $i \in [1, R]$, is a discourse, the stance expressed in $\mathbb{A}$ towards each sub-topic in T_s^m is determined based on stance analysis of discourses within $\mathbb{A}$. The specific steps involved in the process are:

1) **Topic Identification for Discourse**: Identify the topic t_r, $t_r \in T_s^m$, for each discourse r within $\mathbb{A}$;

2) **Stance Determination for Discourse**: Determine the stance for t_r expressed in r;

3) **Stance Detection for Document**: Identify the stance expressed in $\mathbb{A}$ for each topic in T_s^m.

Before proceeding with the above steps, it is necessary to provide an explicit representation for each sub-topic in T_s^m, as detailed in Section 3.2.1.

3.2.1 Explicit Representation of Sub-topics

The sub-topics derived from Disc-LDA, represented by sets of topic-related words $TW(t_{s_i}^m) = \{w_1^{m,s_i}, w_2^{m,s_i}, ...\}$, are hidden variables (Balikas et al., 2016). Providing explicit representations of these sub-topics is crucial for subsequent stance analysis. Previous work has involved representing emerging sub-topics based on existing topics, but due to the diversity and openness of explored sub-topics, existing training data is often insufficient (Allaway and McKeown, 2020). To address this limitation, a prompt learning-based approach is proposed to generate explicit representations for sub-topics based on topic-related words lists, leveraging the capabilities of generative LLMs in summarization and expression. The main part of the prompt template for instructing LLMs to generate the explicit representation of $t_{s_i}^m$ is as follows:

Create a contentious statement based on the words in the set: $\{w_1^{m,s_i}, w_2^{m,s_i}, ...\}$. These words are ordered by diminishing importance and originate from a variety of documents that all relate to a certain controversial topic... Present the result in the format: {Statement}...

In the above template, $\{w_1^{m,s_i}, w_2^{m,s_i}, ...\}$ is the list of topic-related words of $t_{s_i}^m$, which should be instantiated based on different topics in practice. The complete template and examples of the generated explicit representation are provided in **Appendix A**. Given the potentially large number of sub-topics, employing LLMs for automatic summarization of representations of sub-topics is necessary for efficiency.

The representation generated for $t_{s_i}^m$ is referred to as the title of $t_{s_i}^m$, denoted as **Title**($t_{s_i}^m$). Title($t_{s_i}^m$) provides a clear expression of a viewpoint or assertion (e.g., "Equal pay for equal work"), as emphasized in the complete template (refer to **Appendix A**). However, due to the inherent randomness in LLMs-generated results, slight modifications may be necessary to obtain the final value of Title($t_{s_i}^m$). A formal description is as follows:

$$Title(t_{s_i}^m) = ER\_LLM(w_1^{m,s_i}, w_2^{m,s_i}...) \quad (2)$$

ER_LLM is a function of LLMs-based explicit representation generation. The specific LLMs model in ER_LLM can be GPT (Brown et al., 2020) and PaLM (Chowdhery et al., 2023), etc.

3.2.2 Stance Detection for Discourse

This section presents the steps of **Topic Identification for Discourse** and **Stance Determination for Discourse**.

The explored sub-topics in T_s^m often lack sufficient training data for stance detection. Existing strategies, such as zero-shot stance detection (Liang et al., 2022), aim to utilize existing training data for stance detection of emerging topics. However, the diversity and openness of explored sub-topics pose challenges for stance detection with limited training data. This section presents an LLMs-based approach for determining discourse stance toward specific sub-topics.

Firstly, the **top λ candidate topics** for a discourse r, denoted as $\{t_{r1}, ..., t_{r\lambda}\}$, are obtained based on Equation (1) and ranked in descending order of scores. λ is a model hyperparameter. LLMs is employed to re-evaluate the candidate topics and select the optimal one. The prompt template for stance judgment on r for t_{ri} is as follows:

Assess the relevance of the given statement to the topic {Title(t_{ri})}. If not relevant, output 'NA'. Otherwise, determine its stance on {Title(t_{ri})}, choosing from {1. Support, 2. Oppose, 3. Neutral}. The statement is: {r}.

The above process can be formalized as follows:

$$Stance(r, t_{ri}) = SD\_LLM(r, t_{ri}) \quad (3)$$

SD_LLM is a function of LLMs-based stance detection. $Stance(r, t_{ri})$ can be 'NA' (if t_{ri} not discussed in r), 'Support', 'Oppose', or 'Neutral'. The process of determining the final topic t_r, and the stance towards t_r, for a discourse r, is shown in Algorithm 2.

Algorithm 2: Discourse Stance Detection

Input : Discourse r, Candidate Topics $\{t_{r1}, ..., t_{r\lambda}\}$
Output : Topic of r: t_r, Stance toward t_r: $Stance(r, t_r)$

```
t_r ← 'None';
Stance(r, t_r) ← 'NA';
for topic t_ri in {t_r1, ..., t_rλ} do
    Stance(r, t_ri) ← SD_LLM(r, t_ri);
    if Stance(r, t_ri) is not 'NA' then
        t_r ← t_ri;
        Stance(r, t_r) ← Stance(r, t_ri);
        Terminate;
    end
end
```

According to Algorithm 2, topics in the set $\{t_{r1}, ..., t_{r\lambda}\}$ are sequentially assessed, and the first topic that yields a non-'NA' result is set to the topic of r. If none of the topics are discussed in r, t_r is set to 'None'. An alternative approach is to have LLMs evaluate all candidate topics at once and then select the topic from the results. In this case, the principle of topic selection is consistent with Algorithm 2.

3.2.3 Stance Detection for Document

For each discourse r_j in $\mathbb{A} = \{r_1, ..., r_R\}$, t_{r_j} and $Stance(r_j, t_{r_j})$ are determined based on Algorithm 2. As different discourses may express varying stances toward the same topic, the stance of $\mathbb{A}$ toward t_i is derived using a voting mechanism to select the majority value (excluding 'NA').

$$Stance(\mathbb{A}, t_i) = Majority_{\neg NA}\{Stance'(r_1, t_i), ..., Stance'(r_R, t_i)\} \quad (4)$$

$Stance'(r_j, t_i)$=$Stance(r_j, t_{r_j})$ if $t_i = t_{r_j}$, and $Stance'(r_j, t_i)$ is 'NA' otherwise. $Majority_{\neg NA}$ is a function that identifies the value with the most occurrences in a set, excluding instances of 'NA'.

3.3 Determining Stance on the Macro Topic

After determining the stance toward each sub-topic, an SVM model is employed to identify the overall stance for the macro topic. Following AlDayel and Magdy, 2021, where neutrality is not considered, macro topic stance detection is treated as a binary classification task, yielding 'Pro' or 'Con' stances. The SVM model is chosen for its suitability in small sample learning, a significant advantage given the challenges in acquiring extensive training data. Moreover, the SVM's reliance on only a few support vectors is beneficial for identifying essential sub-topics while discarding less relevant ones, contributing to interpretability of macro topic stance detection. Formally, stance detection for the macro topic t_m is described as:

$$Stance(\mathbb{A}, t_m) = Stance(\mathbb{A}, T_s^m) = SVM\_CLF\{Stance(\mathbb{A}, t_{s1}^m), ..., Stance(\mathbb{A}, t_{sT}^m)\} \quad (5)$$

SVM_CLF is an SVM-based stance classifier. Training data for the classifier should be collected in advance, as detailed in Section 4.1.

4 Experimental Setup

The section introduces the experimental setup for macro topic stance analysis on the political left-right division problem ("Political Leaning").

4.1 Experimental Datasets

The experimental data comprises tweets from individuals with distinct left/right inclinations, sourced from voteview[1], including 1,178 Twitter accounts of political figures supporting either the Republican or Democratic party. A dataset of 5.41 million tweets from these accounts was collected following the method outlined in AlDayel and Magdy, 2021, and it was used for sub-topic analysis, referred to as "**Corpus** $\mathbb{D}$". All tweets from an individual account form a document (referred to as "**Document** $\mathbb{A}$"), and each tweet serves as a discourse for Disc-LDA.

A subset of 842 Twitter accounts and their associated 4.12 million tweets serves as the training and test data for **SVM_CLF**. These accounts were selected as they encompass at least 60% of the extracted sub-topics for "Political Leaning" and are evenly distributed between Republican and Democratic support. 75% of these accounts were assigned as training data, and the remaining 25% as test samples.

[1] https://github.com/voteview/tag_twitter

4.2 Comparison Models

The comparison methods are introduced below:

- **BERT-CLS**, a BERT-based classifier, utilizes the same experiment data as SVM_CLF. Due to BERT's limited input capacity, the model input for an individual account comprises randomly selected tweets with diverse sub-topics.

- **LLM-Prompt**, a designed zero-shot prompt for directly querying LLMs for stance results, follows the same experimental data and model input selection as BERT-CLS.

Ablation studies were conducted as follows:

- **Subtopic-manual**, a variant of STS4MTS excluding the step "Sub-topic Extraction", manually selects sub-topics for the macro topic.

- **LDA-based**, a variant of STS4MTS, applies traditional LDA instead of Disc-LDA.

- **ERST-less**, a variant of STS4MTS, utilizes the topic-related words for sub-topic representation instead of 'ER_LLM'.

- **STS-ZSSD**, a variant of STS4MTS, applies the zero-shot stance detection model (Allaway and McKeown, 2020) for sub-topic stance detection instead of 'SD_LLM'.

In addition, different values of λ were tested, where $\lambda = 1$ indicates the exclusion of re-evaluating candidate topics using LLMs.

4.3 Implementation Details

In Disc-LDA, experimentation with different numbers of sub-topics (i.e., the value of T) revealed that utilizing 112 sub-topics for "Political Leaning" resulted in optimal average topic coherence across all topics. The values of α and β were set to 0.1 and 0.001, respectively. Disc-LDA based "Sub-topic Extraction" ran for 2 hours on a 24GB NVIDIA GeForce RTX 3090 GPU server, iterated 50 times to optimize sub-topics selection.

The specific LLMs used in all experiments was GPT 4.0 (Achiam et al., 2023), accessed through APIs. The BERT model employed was the pre-trained uncased BERT-base (Devlin et al., 2019). The learning rate was set to 3e-5, and the Adam optimizer was used with a mini-batch size of 16.

For SVM_CLF, a radial basis function (RBF) kernel was selected due to the moderate sample size and relatively small feature dimension (i.e., the number of sub-topics). Hyperparameter optimization was performed using grid search to select regularization and penalty coefficients within a predefined range, enabling the SVM model to autonomously identify optimal hyperparameters.

Following Allaway and McKeown, 2020, the Macro-averaged F1 metric for each label was utilized to evaluate the performance of each model.

Model		Pro	Con	All
BERT-CLS		52.5%	55.3%	53.9%
LLM-Prompt		47.9%	49.4%	48.6%
Subtopic-manual		47.4%	46.9%	47.2%
LDA-based		53.5%	55.1%	54.3%
ERST-less		55.7%	56.4%	56.6%
STS-ZSSD		54.4%	53.8%	54.1%
STS4MTS	λ=1	59.1%	60.7%	59.9%
	λ=3	63.2%	63.8%	63.5%
	λ=5	63.5%*	64.3%*	63.9%*

Table 2: Performance of stance detection for macro topic "Political Leaning" of different models.

4.4 Main Results

Table 2 presents the outcomes of different models for stance detection towards "Political Leaning". Examples of extracted sub-topics and corresponding topic-related words are provided in **Appendix A**. Based on analysis of the experimental results, the following conclusions can be drawn:

1) **BERT-CLS** and **LLM-Prompt** exhibited sub-optimal performance, mainly attributed to their limited inputs and the abstract and implicit nature of the targets, while **STS4MTS** maximizes the utilization of comprehensive information from various sub-topics.

2) **BERT-CLS** outperforms **LLM-Prompt** as it acquires knowledge of the macro topic from training samples, whereas **LLM-Prompt** lacks such understanding. This highlights the rationale behind **STS4MTS** for mining sub-topics of macro topics.

4.5 Ablation Analysis

Analysis of the model components is as follows:

1) Impact of sub-topics extraction methods. The result of **Subtopic-manual** indicates that manual sub-topic selection inadequately captures the breadth of a macro topic due to its large

number of associated sub-topics. In contrast, **STS4MTS** employs statistical analysis for automated sub-topic selection, improving accuracy and efficiency. The result of **LDA-based** demonstrates inferior performance compared to the Disc-LDA-based approach, attributed to less precise extraction of subject words, consistent with findings in (Jo and Oh, 2011).

2) Impact of explicit representation of sub-topics. The result of **ERST-less** indicates that using topic-related words to represent sub-topics may hinder accurate comprehension due to noise interference. In contrast, **STS4MTS** creates an explicit representation that enhances sub-topic understanding and utilization.

3) Impact of sub-topic stance detection methods. The result of **STS-ZSSD** suggests that representing new topics based on existing ones cannot provide accurate representation of sub-topics due to limited training data. LLMs demonstrate significant advantages in the task of zero-shot stance detection.

4) Impact of λ. Expanding the pool of candidate topics and conducting re-evaluation using LLMs can effectively mitigate computational errors of Equation (1).

4.6 Unsupervised Application of STS4MTS

The training data limitation may constrain the application of STS4MTS, while the unsupervised nature of sub-topic extraction and stance detection highlights their inherent value. This section illustrates an application of sub-topic stance detection using the macro topic "Economic Expectation" as an example. From an online forum[2], 10,120 economic review articles were obtained and analyzed, and sub-topics related to "Economic Expectation" were extracted and the stance for each sub-topic was detected. The average positive stance per month constitutes the Economist Confidence Index, which was subjected to correlation analysis with the official Economist Confidence Index.

Table 3 demonstrates that the computed confidence index exhibits a notably high correlation with the official index. Sub-topic 'Output' demonstrates a stronger correlation compared to Sub-topic 'Risk', suggesting avenues for macroeconomic analysis and revealing factors closely linked to economic indicators and growth trends.

[2]http://www.chinacef.cn/index.php/index/articlemore

	IndexOffi	CurOffi	ExpOffi
STS4MTS	0.64	0.67	0.37
Price	0.4	0.56	0.12
Output	0.69*	0.75*	0.39
Finance	0.42	0.55	0.15
Risk	0.3	0.44	0.07
International	0.51	0.53	0.29
Policy	0.6	0.52	0.42*

Table 3: Correlation analysis of "Economic Expectation" confidence index variables. 'IndexOffi' represents the official quarterly index, 'CurOffi' represents the official current prosperity index, 'ExpOffi' represents the official expected prosperity index. The last six items (starts from 'Price') represent the confidence indices of the sub-topics of "Economic Expectation". The numbers in the table represent Pearson correlation coefficients.

4.7 Interpretability Analysis

STS4MTS provides explanatory insights into macro topic stance judgments by examining the magnitudes of coefficients, which indicate the significance of sub-topics within the macro topic. SVM coefficients allow for the identification of primary classification features, aiding in the removal of less crucial elements with lower variance.

Appendix A presents five influential and representative sub-topics and their topic-related words for the "Political Leaning" macro topic. These pivotal topics significantly shape individuals' political inclinations and facilitate a comprehensive analysis of relevant subjects within the macro topic.

5 Conclusions

This paper proposes a strategy for analyzing the stance of macro-level topics, often implicit in texts but conveyed through relevant sub-topics. By extracting sub-topics and identifying stances towards them in text, a more accurate determination of macro topic stances is achieved. Moreover, interpreting macro topic stances can be enhanced by analyzing the influence of sub-topics. Leveraging stance identification on sub-topics extends stance analysis to macro topics using unsupervised methods, broadening its applicability.

Future research should prioritize the construction of higher-quality, larger-scale evaluation datasets, and the development of more effective evaluation methodologies for macro topic stance detection. Moreover, exploring additional categories and domains of macro topics holds promise for future investigations.

Limitations

Firstly, our proposed method relies on a substantial dataset for conducting sub-topic analysis of macro topics. Given the diversity of macro topics, there are currently no standardized criteria for the requisite sub-topic analysis data. In STS4MTS, following the SVM_CLF analysis, essential sub-topics for analyzing macro topic stances can be identified. Therefore, in future research, we aim to employ an iterative analysis approach: initially analyzing a limited dataset to eliminate less relevant sub-topics, followed by a gradual refinement process to remove increasingly irrelevant sub-topics. This iterative method aims to optimize data collection and computation.

Secondly, when evaluating an individual's stance on macro topics, their stances on various key sub-topics are necessary but often unavailable in practice. In such cases, an alternative approach involves conducting a questionnaire survey to determine the individual's stances on the relevant sub-topics.

Ethical Statement

The data utilized in our experiments, primarily consisting of tweets and review articles, was legally acquired from public websites and strictly employed for research purposes only. Additionally, we took measures to anonymize the data by removing any identifying information. ChatGPT was employed to aid in refining the rhetoric of certain sentences in the paper; however, it did not contribute to the generation of any content.

One potential risk associated with our approach is the inadvertent disclosure of an individual's stance on a macro topic, which they may prefer not to explicitly express. However, in practical applications, our methodology primarily involves analyzing the evolution of public stances regarding macro topics, where personal information is concealed. Therefore, there is minimal risk of exposing personal viewpoints.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Abeer Aldayel and Walid Magdy. 2019. Your stance is exposed! analysing possible factors for stance detection on social media. *Proceedings of the ACM on Human-Computer Interaction*, 3:205:1–205:20.

Abeer AlDayel and Walid Magdy. 2021. Stance detection on social media: State of the art and trends. *Inf. Process. Manag.*, 58(4):102597.

Emily Allaway and Kathleen McKeown. 2020. Zero-Shot Stance Detection: A Dataset and Model using Generalized Topic Representations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8913–8931, Online. Association for Computational Linguistics.

Nora Saleh Alturayeif, Hamzah Luqman, and Moataz A. Ahmed. 2023. A systematic review of machine learning techniques for stance detection and its applications. *Neural Comput. Appl.*, 35(7):5113–5144.

Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. Stance detection with bidirectional conditional encoding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 876–885, Austin, Texas. Association for Computational Linguistics.

Georgios Balikas, Massih-Reza Amini, and Marianne Clausel. 2016. On a topic model for sentences. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*, pages 921–924. Association for Computing Machinery.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.

Samuel Brody and Noemie Elhadad. 2010. An unsupervised aspect-sentiment model for online reviews. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 804–812, Los Angeles, California. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Joachim Büschken and Greg M. Allenby. 2016. Sentence-based text analysis for customer reviews. *Mark. Sci.*, 35(6):953–975.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.

Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd, and Nigel Collier. 2021. Synthetic examples improve cross-target generalization: A study on stance detection on a Twitter corpus. In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 181–187, Online. Association for Computational Linguistics.

Kareem Darwish, Walid Magdy, Afshin Rahimi, Timothy Baldwin, and Norah Abokhodair. 2018. Predicting online islamophobic behavior after #parisattacks. *The Journal of Web Science*, 4:34–52.

Kareem Darwish, Walid Magdy, and Tahar Zanouda. 2017a. Improved stance prediction in a user similarity feature space. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 145–148, Sydney, Australia. Association for Computing Machinery.

Kareem Darwish, Walid Magdy, and Tahar Zanouda. 2017b. Trump vs. hillary: What went viral during the 2016 US presidential election. In *Social Informatics - 9th International Conference, SocInfo 2017, Oxford, UK, September 13-15, 2017, Proceedings, Part I*, volume 10539 of *Lecture Notes in Computer Science*, pages 143–161. Springer.

Kareem Darwish, Peter Stefanov, Michaël J. Aupetit, and Preslav Nakov. 2020. Unsupervised user stance detection on twitter. In *Proceedings of the Fourteenth International AAAI Conference on Web and Social Media (ICWSM 2020)*, pages 141–152, Atlanta, Georgia. AAAI Press.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Kuntal Dey, Ritvik Shrivastava, Saroj Kaushik, and Vaibhav Mathur. 2017. Assessing the effects of social familiarity and stance similarity in interaction dynamics. In *Proceedings of Complex Networks 2017 (The Sixth International Conference on Complex Networks and Their Applications) (COMPLEX NETWORKS 2017)*, pages 843–855, Lyon, France. Springer.

Marcelo Dias and Karin Becker. 2016. INF-UFRGS-OPINION-MINING at SemEval-2016 task 6: Automatic generation of a training corpus for unsupervised identification of stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 378–383, San Diego, California. Association for Computational Linguistics.

Rui Dong, Yizhou Sun, Lu Wang, Yupeng Gu, and Yuan Zhong. 2017. Weakly-guided user stance prediction via joint modeling of content and social interaction. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM 2017)*, pages 1249–1258, Singapore. Association for Computing Machinery.

Manuela GomezSuta, Julian David Echeverry Correa, and Jose A. Soto Mejia. 2023. Stance detection in tweets: A topic modeling approach supporting explainability. *Expert Systems with Applications*, 214:119046.

Swapna Gottipati, Minghui Qiu, Liu Yang, Feida Zhu, and Jing Jiang. 2013. Predicting user's political party using ideological stances. In *Proceedings of the 5th International Conference on Social Informatics*, pages 177—-191, Kyoto, Japan. Springer.

Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2022. A survey on stance detection for mis- and disinformation identification. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1259–1277, Seattle, United States. Association for Computational Linguistics.

Gregor Heinrich. 2005. Parameter estimation for text analysis. Technical report, Citeseer.

Marjan Hosseinia, Eduard Dragut, and Arjun Mukherjee. 2020. Stance prediction for contemporary issues: Data and experiments. In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 32–40, Online. Association for Computational Linguistics.

Yuki Igarashi, Hiroya Komatsu, Sosuke Kobayashi, Naoaki Okazaki, and Kentaro Inui. 2016. Tohoku at SemEval-2016 task 6: Feature-based model versus convolutional neural network for stance detection. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 401–407, San Diego, California. Association for Computational Linguistics.

Yohan Jo and Alice Oh. 2011. Aspect and sentiment unification model for online review analysis. In *Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM 2011, Hong Kong, China, February 9-12, 2011*, pages 815–824. Association for Computing Machinery.

Elena Kochkina, Maria Liakata, and Isabelle Augenstein. 2017. Turing at SemEval-2017 task 8: Sequential approach to rumour stance classification with branch-LSTM. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 475–480, Vancouver, Canada. Association for Computational Linguistics.

Mirko Lai, Alessandra Teresa Cignarella, Delia Irazú Hernández Farías, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. Multilingual stance detection in social media political debates. *Computer Speech and Language*, 63:1–43.

Chen Li, Hao Peng, Jianxin Li, Lichao Sun, Lingjuan Lyu, Lihong Wang, Philip S. Yu, and Lifang He. 2022. Joint stance and rumor detection in hierarchical heterogeneous graph. *IEEE Transactions on Neural Networks and Learning Systems*, 33:2530–2542.

Quanzhi Li, Qiong Zhang, and Luo Si. 2019. eventAI at SemEval-2019 task 7: Rumor detection on social media by exploiting content, user credibility and propagation information. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 855–859, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Bin Liang, Zixiao Chen, Lin Gui, Yulan He, Min Yang, and Ruifeng Xu. 2022. Zero-shot stance detection via contrastive learning. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, pages 2738–2747. Association for Computing Machinery.

Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.

Minghui Qiu, Yanchuan Sim, Noah A. Smith, and Jing Jiang. 2015. Modeling user arguments, interactions, and attributes for stance prediction in online debate forums. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 855–863, Vancouver, BC, Canada. SIAM.

Nilaa Raghunathan and Saravanakumar Kandasamy. 2023. Challenges and Issues in Sentiment Analysis: A Comprehensive Survey. *IEEE Access*, 11:69626–69642.

Ammar Rashed, Mücahid Kutlu, Kareem Darwish, Tamer Elsayed, and Cansin Bayrak. 2021. Embeddings-based clustering for target specific stances: The case of a polarized turkey. In *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media (ICWSM 2021)*, pages 537–548, online. AAAI Press.

Umme Aymun Siddiqua, Abu Nowshed Chy, and Masaki Aono. 2019. Tweet stance detection using an attention based neural ensemble model. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1868–1873, Minneapolis, Minnesota. Association for Computational Linguistics.

Parinaz Sobhani, Diana Inkpen, and Xiaodan Zhu. 2019. Exploring deep neural networks for multi-target stance detection. *Computational Intelligence*, 35:82–97.

Rui Wang, Deyu Zhou, Mingmin Jiang, Jiasheng Si, and Yang Yang. 2019. A survey on opinion mining: From stance to product aspect. *IEEE Access*, 7:41101–41124.

Bowen Zhang, Daijun Ding, and Liwen Jing. 2022. How would stance detection techniques evolve after the launch of chatgpt? *CoRR*, abs/2212.14548.

Bowen Zhang, Xianghua Fu, Daijun Ding, Hu Huang, Yangyang Li, and Liwen Jing. 2023. Investigating chain-of-thought with chatgpt for stance detection on social media. *CoRR*, abs/2304.03087.

Deyu Zhou, Yang Yang, and Yulan He. 2018. Relevant emotion ranking from text constrained with emotion relationships. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 561–571, New Orleans, Louisiana. Association for Computational Linguistics.

Lixing Zhu, Yulan He, and Deyu Zhou. 2019. Hierarchical viewpoint discovery from tweets using bayesian modelling. *Expert Systems with Applications*, 116:430–438.

A Examples of Generated Explicit Representation of Sub-topics.

Appendix A provides: 1) examples of extracted sub-topics for "Political Leaning" using Disc-LDA, presented as topic-related word lists; 2) the complete prompt template for generating explicit representations of sub-topics based on word lists; and 3) the corresponding explicit representations ('Title') for each example sub-topic, as shown in Table 4.

Top 25 Topic-related Words of a Sub-topic	Prompt Template for Explicit Representation Generation	Generated Explicit Representation ('Title')
bill, gun, reform, violence, law, senate, congress, mass, combat, bipartisan, legislation, authority, gun_violence, protection, ban, accountability, domestic, crime, reduce, safety, criminal, abuse, amendment, vulnerable, epidemic	Create a contentious statement based on the words in the set: $\{w_1^{m,s_i}, w_2^{m,s_i}, ...\}$. These words are ordered by diminishing importance and originate from a variety of documents that all relate to a certain controversial topic. Craft the statement in no more than 6 words, expressing a clear viewpoint or assertion, enabling people to clearly state their support or opposition. Present the result in the format: {Statement}. Provide a brief explanation of the statement in 30 words.	Title: "Ban guns to curb domestic abuse". Explanation: This statement suggests that prohibiting firearms would decrease incidents of domestic abuse and violence, assuming gun access contributes to the prevalence of such incidents.
right, woman, fight, freedom, justice, human, voice, abortion, equality, constitution, liberty, discrimination, voting_right, human_right, reproductive, lgbtq, diversity, birth, gender, body, injustice, equity, deserved, protection_act, religious		Title: "Abortion bans violate women's reproductive rights". Explanation: This statement encapsulates the belief that women should have autonomy over their reproductive choices, including the right to have an abortion.
health, care, plan, child, patient, obamacare, healthcare, coverage, medicare, quality, treatment, hospital, medicaid, low_income, option, public_health, income, poverty, tax_credit, high_quality, affordable_care, clinic, us_fda, child_care, health_coverage		Title: "Obamacare fails to ensure high-quality healthcare." Explanation: The statement critiques the Affordable Care Act (Obamacare), suggesting it does not succeed in providing high-quality healthcare to all, especially those in need.
debt, college, force, student, problem, term, civil, loan, university, student_loan, long_term, graduating, overdue, pledge, affair, long_overdue, obligation, graduate, empower, degree, student_debt, short_term, borrower, grade, talented		Title: "Forgive student loans, empower graduates now." Explanation: This statement asserts that canceling student debt would enable graduates to achieve their potential without the burden of financial obligations hanging over them.
energy, infrastructure, change, climate, water, research, clean, investment, environment, green, climate_change, awareness, environmental, natural, electric, affect, air, clean_energy, pollution, 21st_century, future_generation, funding, climate_crisis, scientist, infrastructure_investment		Title: "Clean energy investment mitigates climate crisis." Explanation: This statement posits that investing in clean energy is a crucial solution for addressing the climate crisis, implying a need for action and funding priorities.

Table 4: Five influential and representative sub-topics for the "Political Leaning" macro topic. The first column lists the top 25 associated topic-related words for each sub-topic generated by Disc-LDA. The second column shows the complete prompt template used for generating the explicit representation (i.e., 'title') of these sub-topics. The third column displays the generated explicit representation ('title') of the sub-topics, along with the corresponding explanations provided by the GPT model.