Discourse Analysis for Participatory Discussions: A Topic Shift Detection Approach Using Text Mining and Network-Based Analysis

Shun Ikumi[1,] , Gregory Gay[2], and Maiko Sakamoto[1],

[1] Department of International Studies, Graduate School of Frontier Sciences, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 2778563, Japan.
[2] Department of Computer Science and Engineering, Chalmers University of Technology and University of Gothenburg, Horselgången 5, 417 56 Goteborg, Sweden
ikumi.shun.22@dois.k.u-tokyo.ac.jp, greg@greggay.com, m-sakamoto@k.u-tokyo.ac.jp

Abstract. Effective decision-making relies on well-organized discussions and seamless topic transitions. However, real-world discussions often suffer from inefficiencies such as unclear topic shifts or fragmented discussion, making it challenging to track discussion flow and participant influence. This study refines a novel Topic Change Indicator (TCI) metric for detecting topic shifts in real-time. In contrast to traditional fixed segmentation methods, TCI dynamically adjusts analytical segmentations based on discussion intensity and structural changes, enabling more accurate segmentation. This study employs network-based analysis with TCI, applying speaker networks to examine conversational dynamics and word co-occurrence networks to analyze topic evolution. This paper shows this framework's ability to identify topic transitions and interaction patterns using discussion data from "12 Angry Men." Future work in this Ph.D. research will extend its application to real-world discussions and aims to contribute to real-time discourse analysis by providing facilitators with quantitative insights into topic evolution and speaker influence.

Keywords: text mining, network-based analysis, discourse analysis, topic segmentation

1 Introduction

Effective decision-making relies on well-structured discussions; however, real-world conversations often suffer from fragmented discussion and ambiguous topic shifts. Topic modeling identifies latent themes by analyzing word frequency and co-occurrence patterns. Building on topic modeling, discourse segmentation structures conversations into meaningful units using linguistic structures such as speaker transitions and syntactic structures. Discourse structure analysis further examines coherence, interaction patterns, and conversational flow to reveal the overall organization of discussions. Integrating these methods provides a more comprehensive understanding of discussions, supporting clearer and more effective decision-making.

Efforts to enhance this integration have been made through advancements in argument mining and discourse segmentation. Argument mining identifies logical relationships between claims and supporting evidence using machine learning [1]. Discourse segmentation has evolved from fixed-length segmentation methods like TextTiling [2] to dynamic approaches such as RollingLDA [3], which adaptively capture discussion shifts. However, these methods remain largely constrained to pre-recorded or non-interactive textual datasets, limiting their applicability to interactive or real-time discussions [4].

Network-based analysis has also emerged as a critical tool for discourse analysis, offering insights into linguistic structures and interaction patterns. Word co-occurrence networks facilitate topic modeling by capturing relationships between frequently co-occurring terms, identifying topic emergence and shifts [5]. Community detection techniques, such as Louvain clustering, enable the identification of latent discourse structures and thematic clusters [6]. In addition, Social Network Analysis (SNA) techniques, including speaker interaction networks, analyze conversational dominance and networked argumentation patterns [7, 8].

However, existing approaches face challenges in discourse segmentation and structuring. First, RollingLDA and other segmentation methods struggle to detect abrupt topic shifts in real-time, reducing their effectiveness in interactive discussions. Second, network-based analysis provides structural insights but often lacks integration with linguistic topic analysis, making it difficult to correlate discourse structure with thematic transitions. Finally, most existing techniques are designed for post-hoc analysis, limiting their application in real-time decision-making environments. A key challenge, as previously pointed out [9], is how to determine analytical segmentations that reflect discussion flow, as both linguistic-based topic analysis and network-based interaction analysis are more effective when applied to meaningful discourse segments rather than to the discourse as a whole. Thus, a more adaptive and integrative approach is needed to enhance discourse segmentation and enable real-time application.

One possible approach to addressing this challenge is the use of the Topic Change Indicator (TCI), a quantitative metric for detecting topic shifts in discussions [10]. Unlike traditional fixed segmentation adjustment strategies, TCI dynamically adjusts segmentation windows, allowing finer segmentation adjustment during rapid shifts while maintaining broader segmentations in stable phases. Similar to RollingLDA, which models topic transitions over time, TCI provides adaptive segmentation. However, while RollingLDA relies on probabilistic topic distributions, TCI directly detects real-time shifts in discussion flow, enabling more precise segmentation in interactive settings. To further enhance segmentation accuracy, it is crucial to examine the calculation unit for TCI, such as utterance-based, sentence-based, time-based, word-count-based methods, or potentially a combination of these approaches, which affects how TCI is computed and influences segmentation granularity.

This Ph.D. research adopts the Design Science Research (DSR) methodology [11] to systematically develop and evaluate a framework for topic shift detection and discussion analysis. In this paper, the framework is validated with the script of the film "12 Angry Men," where twelve jurors engage in a deliberation process to determine a verdict, focusing on refining TCI and examining its alignment with network-based discourse analysis. Future work in this Ph.D. research will extend its application to real-world discussions to assess generalizability.

2 Research Aims and Questions

2.1 Research Aims

This Ph.D. research aims to propose a novel framework integrating linguistic-based topic shift detection with network-based interaction analysis to enhance the accuracy of dynamic discourse segmentation and structuring.

To achieve this, the Ph.D. research iteratively explores and implements methods for adjusting topic segmentation and interaction analysis within the framework by conducting empirical validation using comparative experiments and real-world discussions. By integrating topic segmentation with network-based analysis, this framework enhances the ability to model discussion structures by providing adaptive insights into discussion flow and topic shifts. Through the identification of key transitions, argumentation density, and speaker influence patterns, the framework aims to assist facilitators and decision-makers in managing discussions effectively.

2.2 Research Questions

This Ph.D. research is guided by four key research questions (RQs):

RQ 1. How can TCI be refined through calculation unit adjustments to enhance the accuracy of topic shift detection?

RQ 2. How can TCI-driven segmentation adjustment enhance network-based analysis of discussion structures?

RQ 3. How do topic shifts influence the progression of discussions and their impact on decision-making processes?

RQ 4. How can network-based approach integrating TCI be applied to real-world discussions to support facilitators and improve discussion structuring?

RQ1 focuses on refining the TCI calculation strategy to improve topic shift detection accuracy, ensuring adaptive discussion structuring. Here, the calculation unit refers to the fundamental unit used to analyze topic shifts within discussions. Since different units affect segmentation granularity and sensitivity, this Ph.D. research evaluates their impact on TCI performance to identify the most effective approach.

RQ2 examines whether TCI-driven segmentation adjustment enables more effective network-based analysis of discussion structures. It investigates whether dynamically adjusted analytical segmentation improves the identification of speaker influence patterns and topic evolution. Using network metrics such as centrality, connectivity, and modularity, this Ph.D. research evaluates how TCI-driven segmentation adjustments contribute to a more accurate representation of discussion dynamics.

RQ3 examines how topic shifts impact discussion structures and decision-making. It assesses whether network-based analysis with TCI can identify key structural patterns, such as shifts in conversational focus, argumentation density, and engagement.

Additionally, it evaluates how different topic shifts affect decision-making effectiveness, distinguishing between productive and ineffective discussions.

RQ4 focuses on real-world applications of network-based analysis with TCI in meetings, workshops, and policy deliberations. It evaluates their effectiveness in capturing discussion flow, detecting key transitions, and supporting facilitators. This Ph.D. research also assesses whether real-time feedback improves discussion management and decision-making, ensuring adaptability across diverse environments.

3 Theoretical Foundations

3.1 Topic Change Indicator (TCI) Definition

In our previous study [10], we introduced the Topic Change Indicator (TCI) to quantify shifts in discussion content over time. The detection of topic shifts is based on tracking variations in word usage patterns. A discussion can be represented as a sequence of utterances, where each utterance introduces new words or alters the frequency of previously used words. When a previously unused word appears or when the frequency distribution of existing words significantly shifts, it suggests that the discussion topic has changed.

To quantify topic shifts, TCI relies on descending rankings derived from the term frequency-inverse document frequency (TF-IDF). This approach ensures that the significance of words within a discussion segment is measured dynamically. The key idea is to compare the ranking of words in consecutive discussion segments rather than simply tracking their frequency changes. This approach enhances robustness against repetitive word use, which does not necessarily indicate a topic shift.

Formally, S_t represents all sentences spoken up to break point[1] t. Furthermore, let W_t be the set of all words that appear in S_t. The frequency of a word w in each segment S_t, denoted as f_{w,S_t}, represents the number of times w appears in S_t. Term frequency is computed as follows:

$$\mathrm{tf}(w, S_t) \coloneqq \frac{f_{w,S_t}}{\sum_{w' \in W_t} f_{w',S_t}}. \quad \textbf{(1)}$$

To account for the importance of words across multiple segments, inverse document frequency is defined as follows:

$$\mathrm{idf}(w, S) \coloneqq \log \frac{\mathrm{card}(S)}{\mathrm{card}(\{S_t \in S \mid w \in W_t\})} + 1\,, \quad \textbf{(2)}$$

[1] A point in audio-based content where no voice input is detected for a certain period, or a logical segmentation of the text in text-based content.

where $\mathrm{card}(*)$ is the cardinality of the set *. TF-IDF, represented as $\text{tf-idf}(w, S_t, S)$, can be calculated as the product of $\mathrm{tf}(w, S_t)$ and $\mathrm{idf}(w, S)$:

$$\text{tf-idf}(w, S_t, S) \coloneqq \mathrm{tf}(w, S_t) \cdot \mathrm{idf}(w, S)\,. \tag{3}$$

$\text{TF-IDF}(S_i, S_j)$ is the set of TF-IDF values of the words contained in W_i and W_j:

$$\text{TF-IDF}(S_i, S_j) \coloneqq \left\{ \text{tf-idf}\left(w, S_i, \{S_i, S_j\}\right) \mid w \in W_i \cup W_j \right\}, \tag{4}$$

A is a subset of real numbers and a is an element of A. Then, $\mathrm{rank}(a, A)$ denotes the descending rank of a among all the elements in A, defined as follows:

$$\mathrm{rank}(a, A) \coloneqq \begin{cases} \{the\ descending\ rank\ of\ a\ in\ set\ A\}, & a \in A \\ 0, & a \notin A \end{cases}. \tag{5}$$

Finally, the Topic Change Indicator (TCI) is defined as follows:

$$\begin{aligned} \mathrm{TCI}(t) \coloneqq \frac{1}{\mathrm{card}(W_t)} \sum_{w \in W_t} & \Big| \mathrm{rank}\big(\text{tf-idf}(w, S_t, \{S_{t-1}, S_t\}), \text{TF-IDF}(S_t, S_{t-1})\big) \\ & - \mathrm{rank}\big(\text{tf-idf}(w, S_{t-1}, \{S_{t-1}, S_t\}), \text{TF-IDF}(S_{t-1}, S_t)\big) \Big|\,. \end{aligned} \tag{6}$$

$\mathrm{TCI}(t)$ is the sum of the change in TF-IDF descending rankings for all the words used between $t-1$ and t, divided by the number of words.

To determine when a topic shift has occurred, a dynamic threshold is applied. Instead of using a fixed threshold, the mean and standard deviation of recent TCI values are used to establish an adaptive detection criterion. Let $m_{l,t}$ and $s_{l,t}$ represent the mean and standard deviation of the last l TCI values up to time t. A topic shift is identified if the following condition is satisfied:

$$\mathrm{TCI}(t) > m_{l,t} + k s_{l,t}\,, \tag{7}$$

where k is a sensitivity parameter. This method is based on the idea of the "three-sigma rule," which considers values more or less than three times of standard deviation from the mean as outliers [12].

3.2 Network-Based Analysis for Discourse Analysis

Network-based analysis provides a structural perspective on discussions by capturing participant interactions and topic evolution. Unlike purely linguistic approaches, network-based analysis examines relational dynamics between speakers and semantic connections between words, offering a complementary discourse analysis method. This Ph.D. research constructs and analyzes both speaker networks and word networks to investigate discussion structures and topic shifts.

A speaker network is a directed graph where nodes represent participants and edges denote turn-taking interactions. If speaker A speaks after speaker B, a directed edge is drawn from B to A, mapping conversational flow and participant influence. Degree centrality measures engagement, betweenness centrality identifies intermediaries guiding topic transitions, and clustering coefficients reveal discussion subgroups. By correlating these metrics with TCI fluctuations, this Ph.D. research examines whether topic shifts align with changing speaker roles, such as rising influence from central participants or emerging conversational leaders.

A word co-occurrence network captures structural relationships between co-occurring words, identifying thematic clusters and topic transitions. Louvain community detection maps subtopics. Core-periphery analysis [13] differentiates stable from emerging topics and refines word co-occurrence networks by distinguishing between central and peripheral terms, allowing for a clearer representation of topic structures across segmentations. Unlike conventional word networks that include all words as nodes, which may lead to excessively complex and noisy representations, this method focuses on identifying the most structurally relevant terms. By tracking how core words evolve across segmentations, we can observe whether key discussion terms remain stable or shift in prominence, offering insights into topic progression. These techniques track how topic transitions occur and whether discussion fragmentation signals topic shifts.

Integrating TCI with both speaker and word co-occurrence networks offers a multi-dimensional framework for discourse analysis, linking linguistic topic shifts with interactional structures. In this framework, TCI captures real-time linguistic topic shifts while speaker networks, modeled as directed graphs of turn-taking interactions, offer insights into conversational dynamics and participant influence. Meanwhile, refined word co-occurrence networks, utilizing community detection and core-periphery analysis, reveal thematic clusters and transitions. This synthesis of text mining and network-based analysis establishes a theoretical foundation for understanding discussion structure and dynamic topic evolution.

4 Research Approach

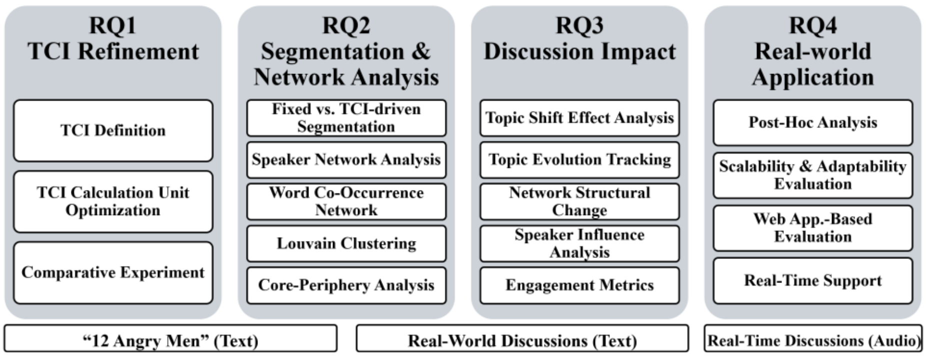

Fig. 1 Research Approach.

Fig. 1 illustrates the Ph.D. research overview, which integrates TCI-based topic shift detection, network-based discourse analysis, and segmentation adaptation to enhance discussion structuring.

Addressing RQ1, this Ph.D. research evaluates the TCI calculation unit strategy to enhance detection precision. Optimizing the TCI calculation unit strategy improves detection accuracy and ensures more adaptive discussion structuring.

For RQ2, discourse analysis with TCI-driven segmentation examines speaker and word co-occurrence networks. The comparison between fixed and TCI-driven segmentation assesses its role in discussion structure identification, such as topic transitions, interaction patterns, and argumentation flow. Speaker networks analyze conversational dynamics by measuring how participants interact and influence discussions, while word networks capture thematic structures by identifying clusters of co-occurring words, allowing for the tracking of topic evolution over time.

To examine RQ3, the analysis focuses on how topic shifts influence engagement, argumentation density, and discussion structure. TCI-driven segmentation examines speaker interactions to identify speaker influence patterns and discussion flow, helping distinguish productive from ineffective topic transitions. The impact of segmentation adjustments on discussion coherence and decision-making effectiveness is assessed through network structural changes and engagement metrics.

For RQ1 and RQ2, the framework is first validated using a segment of the script of the film "12 Angry Men." To address RQ2, RQ3, and RQ4, the framework is applied to real-world discussion transcripts from workshops, meetings, and policy deliberations to assess generalizability.

To evaluate RQ4, we follow three phases. First, post-hoc analysis retrospectively applies TCI and network segmentation to past transcripts. Next, a web application implementing the framework is tested in controlled discussions where participants engage in real-time debates with segmentation and network-based insights. Finally, the framework is deployed in live discussions, providing real-time segmentation feedback and network-based insights to enhance structured decision-making. By integrating network-based analysis with TCI, this Ph.D. research contributes to the development of a scalable, data-driven framework for discourse analysis. The iterative refinement ensures through empirical validation, bridging the gap between theoretical advancements with real-world applicability.

5 Preliminary Results

We show preliminary results validating the framework using "12 Angry Men," examining TCI's alignment with network-based discourse analysis. Fig. 2 illustrates the relationship between TCI shifts and structural transitions in discussion networks, as identified through Louvain clustering applied to word co-occurrence networks. The upper plot represents TCI variations, marking detected topic shifts ($k = 2,\ l = 4$), while the lower plot maps the three most prominent cohesive word clusters, where marker size represents the relative occurrence of words in each cluster. This approach enables tracking of subtopic shifts by analyzing how word clusters shift over time, complementing TCI-based topic shift detection. Results indicate that TCI shifts often align with structural transitions in discussion networks. Frequent co-occurrences of TCI spikes and community shifts highlight their role as a quantitative indicator of topic shift.

Time-based unit size significantly affects detection granularity—smaller units capture subtle shifts, while larger units smooth short-term fluctuations, emphasizing long-term trends. Dynamic TCI calculation unit adjustment enhances adaptability, making the framework suitable for structured and spontaneous discussions.

Additionally, core-periphery analysis and Louvain clustering depend on parameters like *k*-core values and edge weight thresholds, which impact topic identification accuracy. Since manual parameter tuning is impractical, future work will focus on automated parameter optimization to ensure scalable and accurate segmentation.

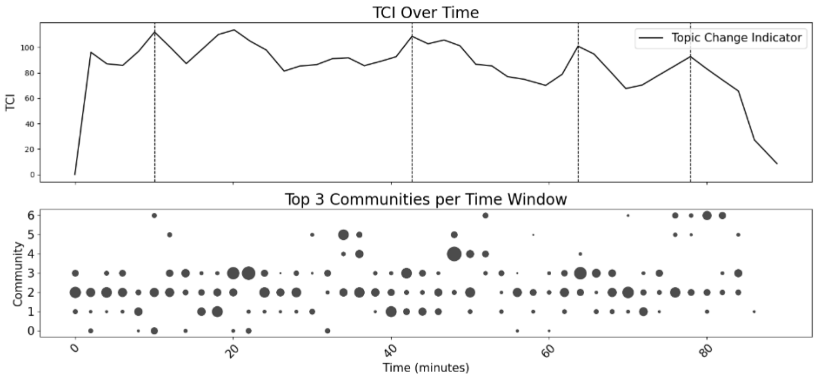

Fig. 2. TCI and Community Transitions in "12 Angry Men."

5 Research Schedule

This Ph.D. research follows a three-phase approach under DSR: theoretical validation, experimental validation, and real-world application. In theoretical validation, the study utilizes discussion data from "12 Angry Men" to assess whether TCI can accurately detect topic transitions and determine the optimal calculation unit. In experimental validation, transcripts from formal discussions, such as those from the Environmental and Social Considerations Advisory Committee [14], are analyzed to examine the interaction between TCI and network-based analysis and to evaluate their effectiveness in understanding discussion structures. In real-world application, a controlled experiment will compare discussions with and without support, alongside real-time facilitation in actual decision-making settings to evaluate its effectiveness.

Additionally, multimodal analysis, incorporating facial expression recognition and speech sentiment analysis [15, 16], will be introduced to assess the impact of non-verbal cues on discussion flow and engagement. However, real-time application of network-based analysis remains largely unexplored, presenting a key challenge in this research. By conducting stepwise validation under the DSR framework, this study aims to enhance the feasibility of real-world implementation.

References

1. Lawrence, J., Reed, C.: Argument Mining: A Survey. Computational Linguistics. 45, 765–818 (2019).
2. Hearst, M.A.: TextTiling: segmenting text into multi-paragraph subtopic passages. Comput. Linguist. 23, 33–64 (1997).
3. Rieger, J., Lange, K.-R., Flossdorf, J., Jentsch, C.: Dynamic change detection in topics based on rolling LDAs. Proc. of Text2Story Workshop. (2022).
4. Irani, A., Faloutsos, M., Esterling, K.: ArguSense: Argument-Centric Analysis of Online Discourse. Proceedings of the International AAAI Conference on Web and Social Media. 18, 663–675 (2024).
5. Pacella, D., Marocco, D.: Understanding Negotiation: A Text-Mining and NLP Approach to Virtual Interactions in a Simulation Game. Applied Sciences. 12, 5243 (2022).
6. Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. J. Stat. Mech. 2008, P10008 (2008).
7. Wu, Y.-L., Rumble, J.N., Ruth, T.K., Lamm, A.J., Ellis, J.D.: An Application of Social Network Analysis to Focus Group Discussions: Unobserved Interaction between Participants and Discussion Topics. Journal of Agricultural Education. 62, 184–195 (2021).
8. Dambanemuya, H.K., Romero, D., Horvát, E.-Á.: Emergent Influence Networks in Good-Faith Online Discussions, http://arxiv.org/abs/2306.13250, (2023).
9. Sakamoto M., Chosokabe M., Salewicz K.: Web Application Development for Dialogue Evaluation Support. SOCIOTECHNICA. 15, 83–94 (2018).
10. Ikumi, S., Maemura, Y.O., Sakamoto, M.: Development of a Real-time Discussion Support Web Application with Text-Mining and Facial Expression Recognition. In: Proceedings of the 23rd International Conference on Group Decision and Negotiation. pp. 326–332. , Tokyo Japan (2023).
11. Runeson, P., Engström, E., Storey, M.-A.: The Design Science Paradigm as a Frame for Empirical Software Engineering. In: Felderer, M. and Travassos, G.H. (eds.) Contemporary Empirical Methods in Software Engineering. pp. 127–147. Springer International Publishing, Cham (2020).
12. Shewhart, Walter.A.: Economic Control Of Quality Of Manufactured Product. D.van Nostrand Company, London (1923).
13. Borgatti, S.P., Everett, M.G.: Models of core/periphery structures. Social Networks. 21, 375–395 (2000).
14. JICA: Environmental and Social Considerations Advisory Committee, https://www.jica.go.jp/about/policy/environment/advice/index.html, last accessed 2025/03/14.
15. Ferracci, S., Giuliani, F., Brancucci, A., Pietroni, D.: Shall I Show My Emotions? The Effects of Facial Expressions in the Ultimatum Game. Behav Sci (Basel). 12, 8 (2021).
16. Kikuchi M., Ichikawa D., Imai T.: A Study on the Possibility of Using Voice, Facial Expression and Text Emotion Analysis Technology to Support Remote Counseling. The Japanese Society for Artificial Intelligence. 2021, 10–4 (2021).