FlowDisco: Interactive Exploration of Dialogue Flows

Patrícia Ferreira[1,2], **Carolina Loureiro**[2], **Ana Alves**[1,3]
Catarina Silva[1,2], **Hugo Gonçalo Oliveira**[1,2]

[1]CISUC/LASI – Centre for Informatics and Systems of the University of Coimbra, Portugal
[2]Department of Informatics Engineering, University of Coimbra, Portugal
[3]Polytechnic Institute of Coimbra, Coimbra Institute of Engineering (ISEC), Portugal
{patriciaf, ana, catarina, hroliv}@dei.uc.pt
carolina.g.loureiro@gmail.com

Abstract

Analyzing large conversational datasets is often inefficient due to the linear nature of text, which hinders the tracking of interaction evolution over time. To address this, we present FlowDisco, an interactive platform for the automatic discovery and exploration of dialogue flows. The framework uses semantic embeddings and modular clustering to transform raw text into probabilistic dialogue flows. By providing a web interface with dynamic filtering and a suite of analytical metrics, FlowDisco simplifies the visual identification and validation of conversational behaviors at scale. The platform's utility is demonstrated through real-world application scenarios, including customer support interactions and multi-party political debates, where it successfully uncovers complex patterns and sentiment shifts that traditional sequential analysis often overlooks.

1 Introduction

The increasing volume of conversational data makes manual analysis inefficient and difficult to scale (Ammar and Bennani, 2025). Traditional NLP techniques often treat utterances as isolated events, hindering the tracking of interaction patterns across thousands of dialogues (Cardoso et al., 2025). Linear reading prevents an understanding of structural dynamics and the multiple paths a conversation can take, creating a need for tools that bridge the gap between a global flow view and the original text (Ganesh et al., 2023; Bouraoui et al., 2019).

To address this, we present FlowDisco, an interactive platform for the automatic discovery and exploration of dialogue flows. It transforms raw data into probabilistic transitions through a modular pipeline combining semantic vectorization with clustering algorithms (Ferreira et al., 2025). The framework provides an intuitive web interface for real-time manipulation of interaction networks, simplifying the identification of behaviors at scale (Bouraoui et al., 2019). By utilizing probability thresholds and node inspection, users can explore complex trajectories and validate model coherence through real utterance examples.

The practical utility of FlowDisco is demonstrated through two distinct scenarios: Study-AI, a dataset with task-oriented human-machine dialogues from a Portuguese schoolbook campaign and complex multi-party debates from the 2025 Portuguese Parliament. These scenarios showcase the platform's ability to handle diverse domains and dynamics, ranging from commercial customer support to adversarial political discourse.

2 FlowDisco

FlowDisco is a full-stack framework for the automatic discovery and interactive exploration of dialogue flows[1]. Its architecture consists of a backend for modular data processing and a frontend for visual analysis, transforming conversational data into structured probabilistic models. This setup facilitates the identification of underlying patterns, following the pipeline illustrated in Figure 1.

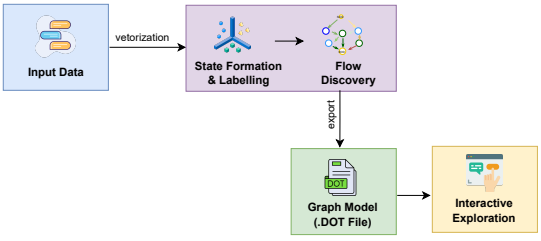

Figure 1: The FlowDisco processing pipeline.

2.1 Processing and Modeling Architecture

The backend, implemented in Python, transforms tabular datasets with three mandatory fields (dialogue_id, speaker, and utterance) into probabilistic graphs. First, utterances are

[1]Available at: https://github.com/NLP-CISUC/FlowDisco

Proceedings of the 17th International Conference on Computational Processing of Portuguese (PROPOR 2026)
Vol. 2, pages 14–17
April 13-16, 2026

vectorized using specific Sentence Transformers (e.g., a multilingual model like `paraphrase-multilingual-MiniLM-L12-v2`) to capture contextual meaning and accurately handle the specificities of the Portuguese language. These are then grouped into dialogue states ($s \in S$) through clustering algorithms like K-Means or DBSCAN, optimized by Silhouette Score or V-measure, or based on keywords and semantic categories. Once states are defined, they receive interpretable labels via strategies like LLMs, verb extraction, or KeyBERT [2]. The system then computes transition probabilities (p_{ij}) and dialogue-flow metrics for analysis. The final graph is exported as a .dot file for rendering. Additionally, FlowDisco can incorporate sentiment or metadata annotations to enrich the flow (Ferreira et al., 2024).

2.2 Interactive Frontend

The FlowDisco frontend, illustrated in Figure 2, provides a web-based dashboard for the interactive exploration of dialogue flows.

To manage the complexity of dense networks, the interface implements navigation controls and dynamic filtering. A central feature is the threshold slider, which allows for hiding transitions below a specific probability ($p_{ij} < \theta$), revealing the main conversation paths. Users can further isolate trajectories through speaker filters, by individual or group, or locate specific states via keyword node search. Navigation is supported by zoom, pan, and a selection-based box-zoom tool. For improved readability, nodes can be individually repositioned, and the original view can be restored at any time. Additionally, the upload panel allows for loading new .dot files, while the screenshot function saves the current visualization for external use.

The framework utilizes a visual encoding strategy where edge thickness is proportional to transition probability and a color gradient, from red to green, maps the sentiment flow. A legend facilitates the interpretation of these visual elements by identifying node types as either utterances or actions, marking special states such as the start and end of dialogues, and aiding in the identification of participants, which is particularly useful in multi-speaker datasets. To ensure model transparency, FlowDisco enables node inspection through contextual tooltips. Upon hovering, the system highlights incoming and outgoing paths and displays up to five real utterance examples, featuring a “Show More” button that indicates the remaining phrases available in the cluster. This allows for the validation of semantic coherence within groups without leaving the visualization.

A sidebar provides real-time statistical support through a metrics panel. This panel organizes information into Dialogue Metrics, including total statements, actions, and dialogues; Grouping Metrics, covering the number of dialogue and action states; and Flow Metrics, which track transition counts, sentiment data, and coverage metrics. These indicators are dynamically updated; high threshold values reduce visible nodes to focus on likely flows, while lower values increase graph coverage, allowing for the correlation of visual abstraction with the statistical precision of the model.

Overall, by combining automated modeling with interactive visualization, FlowDisco provides a comprehensive framework for understanding and analyzing conversational dynamics.

3 Study Cases

To validate the versatility of FlowDisco, we explored two cases with opposing characteristics to demonstrate how the interface enables extracting conclusions in different domains.

The first case, Study-AI, involves a customer support system with 105,376 utterances. The probability threshold slider was critical to hide rare transitions and focus on main paths. Sentiment colors allowed us to distinguish success flows (in green) from frustation zones (in red). Through node inspection, we found that red areas corresponded to topics where the system failed to solve problems or understand the user, creating negative repetitions that are difficult to detect through linear reading.

The second case, the Parliament, analyzes the 2025 State of the Nation debate, characterized by multi-speaker dynamics and long interactions. Speaker filters were fundamental to reducing visual noise and separating the polarities of each party. The tool also differentiates action nodes (such as Applause or Laughter) from speech nodes, identifying behavioral patterns like self-applause by specific party benches or systematic protests (e.g., consecutive interruption nodes like "False!" or "Shame!" clustered around specific political topics) that transcripts do not clearly evidence.

In both cases, the metrics panel validated vi-

[2] Maarten Grootendorst. KeyBERT: Minimal keyword extraction with BERT., 2020

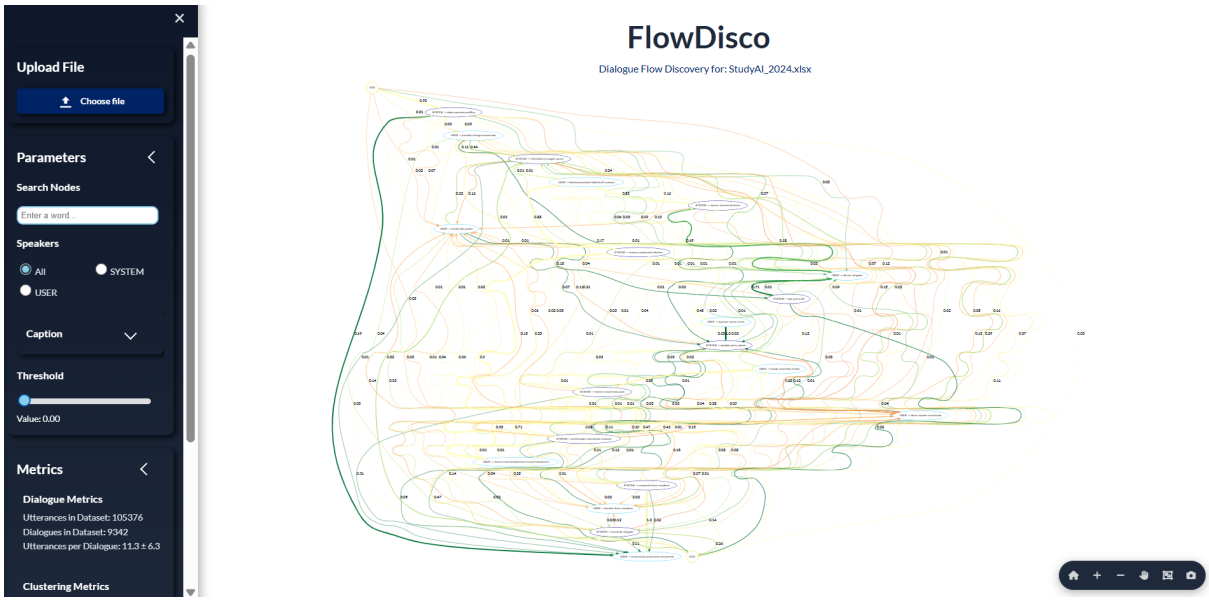

Figure 2: The FlowDisco interface, showcasing the sidebar with parameters and the main dialogue flow visualization.

sual observations. In Study-AI, a flow density of 0.33 confirmed a logical and closed structure despite the volume. In contrast, the low density in Parliament (0.12) corroborated the unpredictable nature of the debate. Monitoring node and transition counts allowed us to balance visual detail with statistical representativeness as we moved the threshold slider.

4 Conclusion

This work presented FlowDisco, an interactive tool designed to make the analysis of conversational data easier and more efficient. By transforming raw text into visual dialogue flows, the platform allows users to explore patterns that are usually hidden when reading text line by line. Through semantic clustering, probabilistic modeling, and dynamic filtering, FlowDisco provides a clear view of how conversations evolve, even in very large or complex datasets.

The application to two different datasets demonstrated the tool's flexibility. In the Study-AI case, the interface allowed for the rapid identification of points of dissatisfaction within thousands of short messages. In the Parliament case, the use of speaker filters and action nodes successfully organized the complex interactions of a multi-speaker debate. These results show that FlowDisco is a valuable framework for anyone needing to understand large-scale dialogue structures. Future work will focus on making the tool work with live data, allowing users to monitor and analyze conversations as they occur.

Acknowledgments

This work was financed by the Portuguese Recovery and Resilience Plan (PRR), through project C645008882-00000055 – Center for Responsible AI.

This work was also supported by FCT – Foundation for Science and Technology, I.P., within the scope of the research unit UID/00326 - Centre for Informatics and Systems of the University of Coimbra. Patrícia Ferreira was supported by FCT – Foundation for Science and Technology, I.P. through the PhD scholarship with reference 2024.01240.BD.

References

Mohamed Achref Ben Ammar and Mohamed Taha Bennani. 2025. A computational approach to modeling conversational systems: Analyzing large-scale quasi-patterned dialogue flows. In *IEEE EUROCON 2025-21st International Conference on Smart Technologies*, pages 1–6. IEEE.

Jean Léon Bouraoui, Sonia Le Meitour, Romain Carbou, Lina M Rojas Barahona, and Vincent Lemaire. 2019. Graph2bots, unsupervised assistance for designing chatbots. In *20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 114–117. ACL.

Mauro Cardoso, Eugénio Ribeiro, and Fernando Batista. 2025. Portuguese far-right discourse on social media: Insights from topic modeling. In *14th Symposium on Languages, Applications and Technolo-*

gies (SLATE 2025), pages 12–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

Patrícia Ferreira, Isabel Carvalho, Ana Alves, Catarina Silva, and Hugo Gonçalo Oliveira. 2024. Sentiment-aware dialogue flow discovery for interpreting communication trends. In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 274–288, Kyoto, Japan. Association for Computational Linguistics.

Patrícia Ferreira, Daniel Martins, Ana Alves, Catarina Silva, and Hugo Gonçalo Oliveira. 2025. Unsupervised flow discovery from task-oriented dialogues. In *Hybrid Intelligent Systems*, pages 336–347, Cham. Springer Nature Switzerland.

Ananya Ganesh, Martha Palmer, and Katharina von der Wense. 2023. A survey of challenges and methods in the computational modeling of multi-party dialog. In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 140–154.