POLVIZX: A Multiplatform Prototype for Stance Tree Construction and Exploration

Tobías Afonso[1], Gabriela A. Diaz[1,2], Carlos Chesñevar[1,2], and Ana G. Maguitman[1,2]

[1] Departamento de Cs. e Ing. de la Computación, Universidad Nacional del Sur, Argentina
[2] Instituto de Cs. e Ing. de la Computación - ICIC (CONICET - UNS), Argentina
ttobiasafonso@hotmail.com, {gabriela.diaz,cic,agm}@cs.uns.edu.ar

Abstract. Opinion polarization in social media has become a major challenge for understanding public discourse and supporting informed decision-making. In previous work, POLVIZ was introduced as a simple web-based prototype designed to analyze polarized opinions by integrating stance detection, topic modeling, argument synthesis, and interactive visual analytics. In this paper, we present POLVIZX, a cross-platform version of POLVIZ that implements the framework for flexible application on any dataset of opinions, offering both a desktop web application and a mobile application for accessible exploration of polarized opinion spaces across devices. The system organizes opinions into hierarchical stance trees that group arguments by topic and stance (pro/anti), enabling users to navigate complex debates and identify polarization patterns. The pipeline supports partial execution and reuse of intermediate results, facilitating iterative analysis. Each application introduces interaction and visualization strategies tailored to its respective environment. POLVIZX further supports user authentication and cross-device synchronization via Firebase, enabling stance tree sharing and collaborative dataset exploration. This cross-platform extension broadens access to opinion polarization tools and supports emerging e-deliberation scenarios.

Keywords: stance trees, opinion polarization, cross-platform prototype, topic modeling, large language models

POLVIZX: Un Prototipo Multiplataforma para la Construcción y la Exploración de Árboles de Posturas

Resumen La polarización de opiniones en redes sociales se ha convertido en un desafío mayor para comprender el discurso público y apoyar la toma de decisiones informada. En trabajos previos, POLVIZ fue presentado como un prototipo web simple diseñado para analizar opiniones polarizadas integrando detección de posturas, modelado de tópicos, síntesis de argumentos e interfaces visuales interactivas. En este artículo presentamos POLVIZX, una versión multiplataforma de POLVIZ que implementa el marco para su aplicación flexible en cualquier dataset de opiniones, ofreciendo tanto una aplicación web de escritorio como una aplicación móvil para la exploración accesible de espacios de opiniones polarizadas en múltiples dispositivos. El sistema organiza opiniones en árboles de posturas jerárquicos que agrupan argumentos por tópico y postura (pro/anti), permitiendo a los usuarios navegar debates complejos e identificar patrones de polarización. El pipeline soporta ejecución parcial y reutilización de resultados intermedios, facilitando el análisis iterativo. Cada aplicación introduce estrategias de interacción y visualización adaptadas a su entorno respectivo. POLVIZX además soporta autenticación de usuarios y sincronización entre dispositivos vía Firebase, habilitando la posibilidad de compartir árboles de posturas y la exploración colaborativa de datasets. Esta extensión multiplataforma amplía el acceso a herramientas de análisis de polarización y soporta escenarios emergentes de e-deliberación.

Palabras clave: árboles de posturas, polarización de opiniones, modelado de tópicos, grandes modelos de lenguaje

1 Introduction

The highly polarized landscape of social media presents a critical challenge to understanding public discourse and supporting decision-making. Online platforms enable large-scale participation in civic debates, but they also produce massive, noisy, and highly unstructured streams of opinions that are difficult to interpret. In this context, polarization is not simply a matter of how many people support or oppose a topic, but also of how opinions cluster into distinct stance groups and how the reasoning behind those groups is organized. To make this structure visible, earlier work introduced POLVIZ [4], a prototype that combined visual analytics and computational argumentation to explore polarized opinions through stance trees. A stance tree is a hierarchical visualization where nodes represent topics at different granularity levels, from broad themes near the root to specific opinion clusters at the leaves, with colors encoding stance polarization patterns. POLVIZ showed that this representation can reveal topic-based conflicts, support human-in-the-loop analysis, and integrate multiple coordinated views for exploring polarized discourse [3].

Building on that foundation, we present **POLVIZX**, a cross-platform system that turns the POLVIZ framework into a practical and deployable environment for stance-tree analysis. POLVIZX adds end-to-end processing of user-provided datasets, Firebase-based authentication and persistence, cross-device synchronization, support for personal and public stance trees, and optimized visualization for both desktop and mobile environments. It also enables partial execution and reuse of intermediate results, making the analysis pipeline more suitable for iterative use in real scenarios. The system combines stance detection, topic modeling through BERTopic [6], argument synthesis through Large Language Models (LLMs) [3], and interactive visual analytics to organize opinions by topic and stance, enabling users to explore polarization patterns at different levels of granularity. This paper presents POLVIZX's components that make cross-platform exploration of polarized discourse feasible in practice.

2 Related work

Tools for polarization analysis range from aggregate dashboards to interactive visual frameworks, though most remain research prototypes [9]. AllSides [1] and Ground News [7] provide real-time polarization indicators through media bias ratings and source diversity metrics. OpinionFlow [8] visualizes opinion evolution through animated flows, while OpMap [2] maps opinions onto argumentation frameworks using probabilistic justification. Other studies combine stance classification with topic modeling to reveal the thematic organization of polarized debates [5]. The POLVIZ prototype [4] introduces stance trees with coordinated multiple views (stance hierarchy, nested topics, and geographic maps) to reveal polarization patterns interactively. However, it remained a web-based demonstration limited to a single climate change dataset without user accounts, cross-device support, or custom dataset processing.

POLVIZX differs from prior work by providing a complete system implementation that bridges the gap between research prototypes and deployable analysis environments. Unlike static visualization tools or aggregate dashboards, POLVIZX supports the full pipeline from raw CSV upload to interactive stance-tree exploration across desktop and mobile platforms.

3 System overview

PolvizX is a cross-platform system that transforms a user-provided opinion corpus into an interactive stance tree for exploration on desktop and mobile devices. Figure 1 summarizes the five main stages of the workflow. First, **Data collection** gathers relevant social media posts using keyword and hashtag filtering targeted at specific issues, with pre-existing research datasets also supported for efficiency. Second, **Stance detection** classifies each opinion as pro, anti, or neutral using either fine-tuned BERT models trained on labeled data or LLM prompting strategies (0-shot, 1-shot, few-shot), with neutral opinions excluded to focus on polarized discourse. Third, **Topic modeling** analysis applies BERTopic to create hierarchical topic dendrograms, assigning polarization valence to each topic (green for pro, red for anti, yellow for balanced, with intermediate shades like lime and orange). Fourth, **Argument synthesis** transforms leaf-node opinion clusters into coherent, stance-aligned arguments via LLM prompt engineering that abstracts informal language into structured reasoning. Fifth, **Stance tree construction** organizes synthesized arguments and topics into a hierarchical tree structure where leaf nodes represent arguments and internal nodes represent topics synthesized via prompt engineering for global opinion synthesis (combining child node summaries/arguments), with colors encoding polarization valence to enable comprehensive visual exploration of polarized discourse organization.

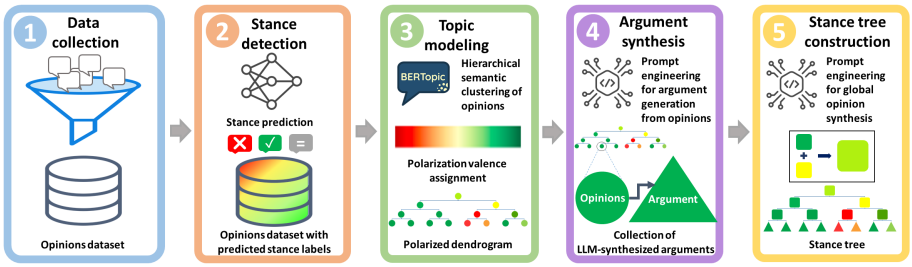

Fig. 1. PolvizX workflow, from data collection to stance tree construction.

The PolvizX system follows a client-server architecture. Users interact with the front-end through either a desktop web application or a mobile application, while the computational pipeline runs on a Linux server. This server-side design is essential because the processing stage integrates models and libraries for LLM-based synthesis and BERTopic clustering that benefit substantially from GPU execution. After a dataset is uploaded as a CSV file, the server executes the analysis pipeline and stores the resulting stance tree in a user-specific directory identified by the Firebase user ID. Each tree is then stored in a dedicated subfolder with a unique tree name, which supports persistent organization, retrieval, and cross-device access.

The visualizer is optimized to reduce client-side computation. Server-side processing precomputes node rendering information (position, color, and parent) into JSON, minimizing client-side layout computation. The tree is drawn on the mobile side using the Canvas API in Jetpack Compose, which leverages GPU acceleration and makes the

rendering fast even on older devices. In practice, the result behaves similarly to displaying a static image, while still preserving interactive behavior. The visualization is redrawn when interaction events occur, such as selecting a node. In those cases, the chosen node is highlighted while the rest of the tree is visually dimmed. Figure 2 shows an example of the stance tree visualization.

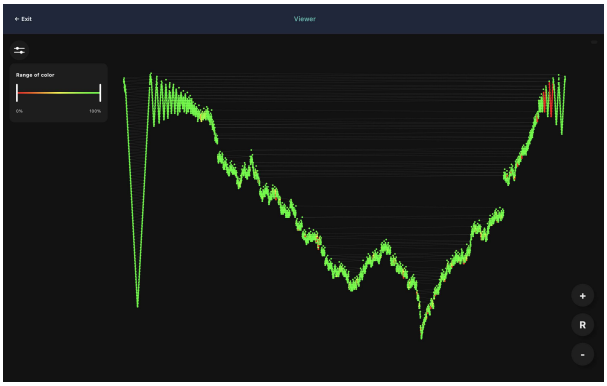

Fig. 2. Stance tree (∼6K nodes) from ∼1M climate change tweets, showing hierarchical organization and polarization patterns. Green dominance over red shows believers outnumbering deniers.

Selecting a node opens a more detailed view of that subtree, showing a navigable local structure centered on the selected position. This secondary view presents only the nearby levels of the hierarchy, which helps users inspect a focused portion of the stance tree without losing orientation in the larger structure. The node detail panel also exposes information that depends on node type. For leaf nodes, users can inspect the most representative opinions associated with the selected cluster, together with the argument synthesized by an LLM from those opinions. For internal nodes, the panel provides access to parent navigation and other structural or descriptive information associated with the cluster. Figure 3 shows the node details visualizer.

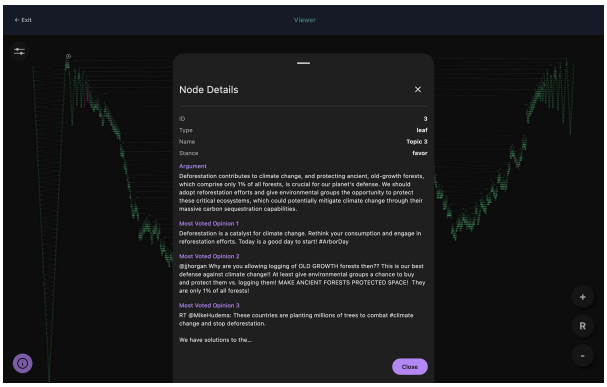

Fig. 3. Leaf node details: representative opinions and LLM-synthesized argument.

This interaction model supports exploratory analysis at multiple levels of granularity. Users can first scan the full tree to identify broad polarization patterns and then navigate into specific topics or stance clusters to inspect the underlying opinions. The result is a system that combines overview, zoom, and detail-on-demand in a compact interface tailored to opinion analysis. The supplementary material[3] provides additional details on the POLVIZX system operation.

4 Conclusion and future work

POLVIZX transforms stance-tree analysis into a practical cross-platform system with Firebase authentication, server-side processing, GPU-accelerated rendering, and synchronized personal/public tree management across devices. Its customizable Canvas-based rendering enables seamless integration of additional CSV data without architectural changes. Future work includes encoding node prominence through node size variations to highlight larger tweet clusters, adding a temporal heatmap with slider control to visualize changing tree regions over time, incorporating vertical x-axis labels summarizing regional topics, and developing animations of progressive tree construction. These enhancements will strengthen POLVIZX's utility for researchers, policymakers, and e-deliberation platforms.

Acknowledgments

This research is supported by a Google Academic Research Award (GARA), the Gemma 3 Academic Program, ANPCyT (PICT 2019-03944), and UNS (PGI 24/N060).

References

1. AllSides Technologies, Inc. Allsides media bias ratings, 2026. Accessed 2026-04-12.
2. Gregor Betz, Michael Hamann, Tamara Mchedlidze, and Sophie Von Schmettow. Applying argumentation to structure and visualize multi-dimensional opinion spaces. *Argument & Computation*, 10(1):23–40, 2018.
3. Gabriela Diaz, Carlos Chesñevar, Roma Patel, and Ana Maguitman. Modeling polarization in public opinion through llm-synthesized arguments and stance trees. *PeerJ Computer Science*, 2026. In press.
4. Gabriela Andrea Diaz, Dana Urribarri, Maria Luján Ganuza, Carlos Chesñevar, Elsa Estevez, and Ana Maguitman. Polviz: assessing opinion polarization in social media through visual analytics and argumentation. In *Proceedings of the 17th International Conference on Theory and Practice of Electronic Governance*, pages 337–347, 2024.
5. Manuela Gómez-Suta, Julián Echeverry-Correa, and José A Soto-Mejía. Stance detection in tweets: A topic modeling approach supporting explainability. *Expert Systems with Applications*, 214:119046, 2023.
6. Maarten Grootendorst. BERTopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*, 2022.
7. Ground News. Rating system, 2026. Accessed 2026-04-12.
8. Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. Opinionflow: Visual analysis of opinion diffusion on social media. *IEEE transactions on visualization and computer graphics*, 20(12):1763–1772, 2014.
9. Yunfei Xing, Xiwei Wang, Chengcheng Qiu, Yueqi Li, and Wu He. Research on opinion polarization by big data analytics capabilities in online social networks. *Technology in Society*, 68:101902, 2022.

[3] Available at https://osf.io/6mrgh/overview?view_only=fb8379c15c21402a978d1c23102c5335