Analyzing Debate Dynamics in the Portuguese Parliament with Dialogue Action Flows

Patrícia Ferreira[1,2], 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
Correspondence: patriciaf@dei.uc.pt, ana@dei.uc.pt, catarina@dei.uc.pt, hroliv@dei.uc.pt

Abstract

Analyzing how large-scale multi-party dialogues shape collective behavior is a central challenge in computational linguistics. However, traditional text-based methods often overlook the complex, non-linear turn-taking dynamics defining these interactions. To address this gap, we propose a framework based on Dialogue Action Flows (DAFs) that integrates verbal utterances and non-verbal actions into a unified probabilistic representation of interactional behavior. Interactions are encoded as speaker-action states, forming a probabilistic DAF that reveals dominant behavioral trajectories and recurrent patterns. We validate this framework on a case study comprising five years of Portuguese State of the Nation debates. Analysis reveals systematic behavioral asymmetries driven by party roles: while government parties exhibit increasing alignment, opposition forces, particularly the radical wing, maintain persistently high conflict. Additionally, the rising volume of interactions across legislative years indicates a progressively heated environment. Overall, our framework provides a quantitative and interpretable approach for modeling polarization, alignment, and interactional dynamics in multi-party political discourse.

1 Introduction

The growing volume and diversity of available large-scale text collections has driven significant progress in information extraction from various types of documents. Modern Natural Language Processing (NLP) excels at processing static documents, enabling effective text summarization (Zhang et al., 2024) and information extraction (Chen et al., 2023) regardless of length. There is a significant presence of conversational data in the digital landscape, ranging from social media (Suhaim and Berri, 2021) to transcripts of spoken dialogue on TV shows (Pant et al., 2022) and political debates (Cribb and Rochford, 2018). In these contexts, the complexity lies not only in what is said, but in how the interaction unfolds.

This challenge becomes more pronounced in multi-party environments. Although tasks such as Dialogue State Tracking (DST) (Williams et al., 2016; Pais et al., 2024) employ Text Classification (Said and Ismail, 2025) and Information Extraction in goal-oriented interactions, these techniques often prove insufficient for unstructured dialogues involving multiple participants. In these dynamic scenarios, communication is shaped not only by verbal content but also by non-verbal signals (e.g., interruptions or reactions). By ignoring this structural dimension, text-only approaches fail to capture the true progression and the specific dynamics of real-world dialogues.

To address the limitations of existing approaches, we propose a framework for modeling interaction dynamics in multi-party dialogues that integrates both verbal utterances and non-verbal actions into a unified structure. Each interaction is encoded as a Dialogue Action State, defined by the speaker and the type of behavior, allowing conversational data to be transformed into a structured and interpretable representation of the dialogue. These states are interconnected according to computed probabilities, forming a directed graph that captures how interactions typically unfold in multi-speaker environments. We refer to this representation as the Dialogue Action Flow (DAF), an artifact that enables the identification of dominant transition paths, recurring interaction patterns, and significant moments in the dialogue, such as interruptions, heated exchanges, or changes in communicative roles.

Beyond the structural representation, the proposed framework incorporates quantitative metrics that allow for characterizing and comparing interaction patterns among participants over time. These metrics facilitate a deeper analysis of complex dynamics, including the emergence of alliances, patterns of support and conflict, and shifts in interac-

Proceedings of the 17th International Conference on Computational Processing of Portuguese (PROPOR 2026)
Vol. 1, pages 369–379

tion strategies. In this sense, the approach provides a quantitative and interpretative framework for examining how participants position themselves and respond to one another within a dialogue.

To demonstrate the practical utility of this approach, we apply it to a real-world case study: debates of the Portuguese Parliament between 2021 and 2025, conducted in European Portuguese. This application illustrates how combining quantitative metrics with visual flow modeling significantly enhances the interpretability of long, multi-speaker interactions, enabling a identification of shifts in the political landscape, specifically by revealing the dynamics of support and opposition among parties.

The main contributions of this work are:

- Integration of both verbal utterances and non-verbal actions into a unified DAF, enabling the modeling and analysis of complex interaction dynamics that extend beyond textual content.
- A structured and interpretable framework to identify dominant transition paths, recurring interaction patterns, and salient moments in dialogues, offering quantitative and visual insights into multi-party communication.
- A large-scale case study on Portuguese Parliament debates, demonstrating the practical utility of the approach, revealing trends in political polarization, shifts in patterns of support and opposition between parties, and the evolution of their reciprocal behaviors.

The remainder of the paper is structured as follows: Section 2 reviews related work on multi-speaker dialogue mining and the analysis of political debates; Section 3 describes the proposed pipeline, covering utterance act classification, clustering, and flow discovery; Section 4 describes the dataset and the applied pre-processing, followed by annotation; Section 5 reports and discusses the results, validating the proposed framework through its application to the parliamentary dataset and the verification of several hypotheses; Section 6 concludes the paper and outlines future work.

2 Related Work

Automated analysis of conversational data has become increasingly relevant, driven by the growing availability of large-scale dialogue collections (Ammar and Bennani, 2025; Vakulenko et al., 2021). These resources are often used to support the development of chatbots (Lin et al., 2023; Kim et al., 2022) or to analyze communication trends (Ferreira et al., 2024a; Carvalho et al., 2024; Holstrup et al., 2020). While most research focuses on dyadic, task-oriented dialogues (Carvalho et al., 2023; Budzianowski et al., 2018), multi-party environments remain less explored due to challenges like thread entanglement (Ganesh et al., 2023). Nevertheless, resources such as the MRDA corpus (Shriberg et al., 2004) and EmotionLines (Hsu et al., 2018) are pivotal for capturing social dynamics in group settings.

Within this landscape, political debates represent a complex multi-party scenario. Initiatives like PTPARL-D (Almeida et al., 2021; Sousa and Cardoso, 2024) and ParlaMint-PT (Aires et al., 2024) provide crucial corpora for Portuguese debates. Research utilizing these resources typically focuses on textual content to extract political insights. For instance, Fournier-Tombs and Di Marzo Serugendo (2020) used machine learning to measure discourse quality and polarization. In the Portuguese context, Cardoso et al. (2025) identified far-right ideological shifts via topic modeling, while Costa et al. (2021) evaluated the coherence between parties' Twitter sentiment and official programs. However, these approaches often treat debates as collections of isolated texts, overlooking the structural flow of interaction and non-verbal signals, which this work aims to address.

To capture such structural dynamics, we turn to Dialogue Flow Discovery. Beyond its traditional use in chatbot development (Bouraoui et al., 2019; Sastre Martinez and Nugent, 2022), this field has been increasingly applied to the analysis of human-human conversations to uncover behavioral patterns and communication trends (Bouraoui et al., 2019). For instance, discovered flows have been used to identify frequent sequences of dialogue acts (Ritter et al., 2010), analyze the evolution of sentiment during interactions (Ferreira et al., 2024a), or characterize user engagement (Holstrup et al., 2020). These approaches abstract conversations into directed graphs, where nodes represent dialogue states and edges model the transitions between them (Ritter et al., 2010; Ferreira et al., 2024b). While a common strategy involves using unsupervised clustering algorithms (e.g., K-means (Ferreira et al., 2023), DBSCAN (Martinez and Nugent, 2022)) to discover latent states from text, other approaches structure flows on Dialogue

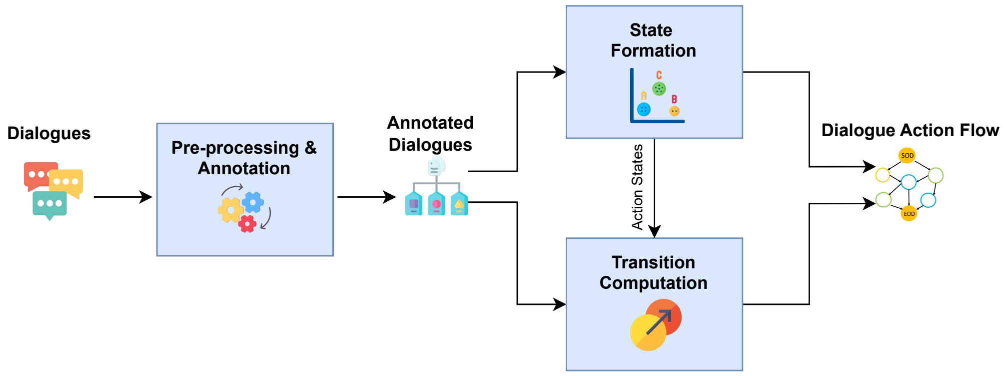

Figure 1: Overview of the proposed approach, comprising three main steps: Pre-processing and annotation, State Formation, and Transition Computation.

Acts (Ribeiro et al., 2019). We adopt this latter perspective to capture the adversarial dynamics of the debate rather than just semantic topics.

In summary, existing research focuses largely on dyadic interactions or isolated text analysis, often overlooking the complex structure of multi-party debates. In contrast, our framework combines DAF with quantitative metrics, integrating both speech and non-verbal actions. This approach allows us not only to visualize interaction dynamics but also to measure polarization and alignment patterns that traditional methods cannot capture.

3 Proposed Approach

To enable the analysis of multi-party dialogues, we propose the pipeline illustrated in Figure 1, which represents the dynamics in a graph of speech acts and actions, dubbed Dialogue Action Flows (DAF). The pipeline comprises three main stages:

Pre-processing & Annotation: raw dialogue data is normalized and converted to a structured format, where interactions and their authors are clearly identified in a list of events. Specifically, non-verbal interactions (e.g. reactions, interruptions) are identified, and the speech act of each utterance is annotated. The latter will depend on the adopted taxonomy of speech acts.

State Formation: interactions are grouped in states ($s \in S$), each representing a type of interaction (i.e., speech act or action) and its author. This step significantly reduces dimensionality and enables generalization, by transforming varied interactions into a set of representative nodes.

Transition Computation: Transitions between action states are extracted from the sequence of actions, resulting in a directed graph $G(S, E)$. Nodes $s \in S$ represent the action states, and edges $e(s_i, s_j, p_{ij}) \in E$ represent the transitions between states. Each edge is weighted by the transition probability p_{ij}, i.e., the proportion of instances in which the state s_i is immediately followed by s_j. This constitutes the DAF. Two additional nodes are added for representing the start (SOD) and the end of dialogue (EOD). To ensure readability and focus on dominant patterns, a threshold θ can be applied to prune low-probability transitions (i.e., with $p_{ij} < \theta$).

The discovered DAF is combined with quantitative metrics to elucidate interaction dynamics and behavioral trajectories, offering a visualization of how the conversational structure evolves. This approach not only measures the interaction structure but also characterizes the social climate and reveals the development of group alignments or divergences, making it applicable to diverse multi-party contexts ranging from political debates to social interactions.

4 Data Preparation

This section starts by presenting the source dataset, followed by the steps to normalize and structure the raw transcripts. Finally, it describes the annotation method, specifying the speech act taxonomy and the LLM configuration for enriching the data.

4.1 Dataset

Experimentation was conducted on a dataset of official transcripts of the State of the Nation debates of the Portuguese Parliament, spanning a five-year period (2021–2025)[1]. These debates were

[1] https://tinyurl.com/ywpwb4e2

Entity	Utterance	Reaction
President	Tem a palavra, também para um pedido de esclarecimentos, o Sr. Deputado André Ventura. *(You have the floor, also for a request for clarification, Mr. Deputy André Ventura.)*	—
CH	Sr. Presidente, pedi para fazer esta intervenção [...] para dizer que não devem temer nenhuma "assombração" *(Mr. President, I asked to intervene [...] to say you should not fear any "haunting")*	—
PS	Assombração! Assombração! *(Haunting! Haunting!)*	—
CH	nem do Governo, nem da oposição. Também não devem temer que este seja o Governo-sombra do Chega *(neither from the Government, nor the opposition. Also, do not fear this is Chega's shadow-Gov*	—
CH	—	Applause
PSD	—	Protests
CDS-PP	Vê-se! *(We can see that!)*	—
CH	Dez vezes! Dez vezes! *(Ten times! Ten times!)*	—
CH	Por exemplo, com um governo do Chega, a esta hora não haveria um incendiário que não estivesse algemado *(For example, with a Chega government, every arsonist would be in handcuffs)*	—

Table 1: Excerpt from a dialogue from the 2025 State of the Nation debate, illustrating the multi-party dynamics.

selected as a representative example of a complex, high-intensity, multi-speaker environment, allowing for the observation of behavioral patterns and group dynamics. The proposed framework is not restricted to parliamentary debates and can be applied to other contexts. To illustrate the nature of these interactions, a representative excerpt from the 2025 debate is presented in Table 1.

The transcription includes both verbal and non-verbal interactions. The former consists of utterances by the participants, while non-verbal interactions encompass reactions such as "Applause", "Protests", "Laughter", as well as more descriptive events, e.g. "The speaker displayed the cited document" or "The speaker's microphone was automatically muted for exceeding the allocated time". Including these actions allows for a richer representation of the dialogue dynamics beyond textual content. For analysis purposes, each annual session is segmented into distinct dialogues, delimited by significant pauses (e.g. recesses or interruptions).

Table 2 presents descriptive statistics evidencing an intensification of parliamentary activity: the volume of utterances ($|U|$) and actions ($|A|$) increased significantly along the years, contrasting with the stability of the total duration (ΔT), which reveals a higher debate density. Furthermore, we observe increasing structural fragmentation, evidenced by the rise in the number of dialogues ($|D|$), which lowers the average duration per dialogue and results in high standard deviations, indicating that the debate alternates between long periods of discussion and very brief interactions.

Debates involve three entities: the **President of the Assembly (P)**, acting as the moderator; the **Government (GOV)**, represented by the Ministers; and party representatives. The parliamentary spectrum during this period includes, in ideological order, roughly Left to Right: Bloco de Esquerda (**BE**), Partido Comunista Português (**PCP**), Partido Ecologista "Os Verdes" (**PEV**), Livre (**L**), Partido Socialista (**PS**), Pessoas-Animais-Natureza (**PAN**), Juntos pelo Povo (**JPP**), Partido Social Democrático (**PSD**), Iniciativa Liberal (**IL**), Partido Popular (**CDS-PP**), and Chega (**CH**). The previous order allows to visualize the evolution of political polarization, alignments, and the structural isolation of specific forces in the DAFs.

Year	|U|	|A|	|D|	|U|/|D|	|A|/|D|	ΔT
'21	900	197	1	900.00 ± 000.0	197.00 ± 00.0	4h21
'22	1184	320	3	394.67 ± 305.4	106.67 ± 66.3	4h48
'23	1377	339	6	229.50 ± 236.7	56.50 ± 55.2	4h22
'24	1466	525	11	133.27 ± 192.9	47.73 ± 65.9	5h16
'25	1598	506	12	133.17 ± 151.7	42.17 ± 45.0	5h01

Table 2: Descriptive statistics of the parliamentary datasets. $|U|$, $|A|$, and $|D|$ represent the total count of utterances, actions, and dialogues, respectively. $|U|/|D|$ and $|A|/|D|$ correspond to the average number of utterances and actions per dialogue. ΔT denotes the total dialogue duration between SOD and EOD.

The composition of the Government changed during the analyzed period. Between 2021 and 2023, the government was led by the Center-Left PS. Following the 2024 elections, a Center-Right government formed by a coalition between PSD and CDS-PP, called the Aliança Democrática (AD) took office. This shift allows us to analyze how the transition from Government to Opposition roles affects the discourse and behavior of the parties.

4.2 Pre-processing

Raw debate transcripts were converted from PDF to a structured tabular format using a custom Python pipeline. The extraction process relied on regular expressions to detect speaker transitions, which follow a strict syntactic pattern in the official transcripts. This approach allowed for the automatic extraction of metadata, including the speaker's name and political affiliation, directly from the turn headers. Non-dialogue elements, such as page numbers and timestamps, were filtered out. The resulting schema includes sequential identifiers (dialogue_id, turn_id) to preserve the debate flow, entity metadata (participant), and content fields (utterance, action) to distinguish between verbal and non-verbal interactions.

Speeches were segmented by paragraph to create shorter utterances, facilitating speech act annotation. Actions were standardized to focus on the primary signals of parliamentary alignment. Specifically, "Counter-protests" were recoded as "Applause" to consolidate supportive reactions, while any interactions falling outside the categories of "Applause", "Protests", or "Laughter" were categorized as "Other". Finally, although "Laughter" forms a distinct category, it was excluded from the polarization metrics due to its ambiguity (signaling either irony or genuine amusement), ensuring that the analysis focuses exclusively on explicit indicators of support or dissent.

Speeches or reactions involving multiple parties (e.g., "Voices from PSD and CDS-PP" or "Applause from PSD and CDS-PP") were decomposed into distinct events, ensuring that each party is associated with its own interaction. Finally, pauses in the debate were used as delimiters to mark the EOD, preventing each annual session from being treated as a single continuous dialogue.

4.3 Annotation

The speech act of every utterance was annotated according to a detailed schema adapted for political discourse analysis (Reinig et al., 2024), which provides clear criteria for categorizing utterances into two top-level classes: Cooperation and Conflict.

Cooperation represents constructive contributions to the dialogue (e.g. "*A Nação realiza-se todos os dias, com cada decisão tomada, com cada obstáculo vencido, com cada português que acredita que é possível fazer diferente e melhor.*"[2]). **Conflict** represents opposing or challenging contributions (e.g. "*Que vergonha! Que vergonha um Primeiro-Ministro que perde 16 segundos a falar de saúde no debate sobre o estado da Nação, quando há portugueses que passam anos e anos até conseguirem ter uma consulta ou um ato cirúrgico.*"[3]).

Annotation was performed automatically by Llama 3.3-70b, deployed within the Ollama framework[4] in Q4_K_M quantization, running on a system equipped with an NVIDIA RTX A6000 GPU (48 GB VRAM). To ensure maximum consistency and deterministic results, temperature was set to 0.0.

To validate the automated annotations, a human annotator classified a random sample of 100 utterances. We computed Cohen's Kappa (κ) and obtained a value of $\kappa = 0.88$, indicating strong agreement between human and LLM annotations regarding the binary classification of Cooperation versus Conflict. These results support the reliability of the automatic labelling process for political discourse analysis. The prompt template used for classification, in Portuguese, is in Figure 2.

> És um assistente especializado em análise linguística política. A tua tarefa é classificar a seguinte fala numa das 2 categorias seguintes: Cooperação ou Conflito. Responde APENAS com o nome da categoria (Cooperação ou Conflito).
> 1. Cooperação: Inclui atos para moderar o debate, transmitir informação relevante, expressar apoio e promover a coesão entre os participantes como MACRO, EXPRESSIVO, RELATO, PERGUNTA, PEDIDO e APOIO.
> 2. Conflito: Inclui atos marcados pela ausência de consenso, envolvendo autopromoção, crítica a outros intervenientes ou confronto direto como AUTO-REPRESENTAÇÃO, PROMESSA, ACUSAÇÃO, REJEIÇÃO, EXIGÊNCIA e PERGUNTA RETÓRICA.
> FALA A CLASSIFICAR: `{utterance}`

Figure 2: Utterance act classification prompt, where `{utterance}` represents the input text.

5 Results and Analysis

In this section, we analyze trends identified through the application of the proposed framework to the State of the Nation debates. First, we establish a set of expected observations based on political theory and parliamentary conventions; second, we validate these hypotheses using the quantitative metrics and the generated DAF.

5.1 Expected Observations and Hypothesis

Parliamentary debates are governed by specific institutional rules and political strategies. Based on the context of the Portuguese political landscape, which experienced a shift in government in 2024

[2]The Nation is built every day, with every decision made, every obstacle overcome, every Portuguese citizen who believes that it is possible to do things differently and better.

[3]What a shame! A Prime Minister that wastes 16 seconds talking about health in the State of the Nation debate, when there are Portuguese people who wait years and years to get a consultation or a surgical procedure.

[4]https://ollama.com/

and the rise of new political forces, we formulate the following hypotheses (H) to guide our analysis:

H1. The President of the Assembly, acting as the moderator, shows the lowest conflict ratios, maintaining a strictly cooperative stance to manage the debate regardless of the political climate.

H2. The volume of utterances ($|U|$) and actions ($|A|$) is structurally concentrated in the three main political forces (PS, PSD, CH).

H3. The stance of conflict or cooperation depends on whether a party is in power. Following the change of government, a reciprocal exchange of behaviors is expected: from 2024 onward, the former Government (PS) becomes conflictual, while the former Opposition (AD) becomes cooperative.

H4. Based on ideological proximity, with the shift to a right-wing government (AD) in 2024, the right-wing spectrum (PSD, CDS-PP, IL, CH) should form a cohesive support bloc (high level of applause/cooperation), contrasting with opposition from the left.

H5. The proportion of protests is expected to remain stable over the years.

H6. Reactions are expected to align with speech sentiment: conflict triggers protests and cooperation elicits applause, driven by the content rather than the speaker's identity.

5.2 Validation

To validate the previous hypotheses, we analyze the evolution of parliamentary activity, the behavioral shifts in stance, and the structural patterns revealed by the DAF. Table 2 and Figure 3 illustrate the progressive intensification of parliamentary activity over the analyzed period.

We observe a marked increase in debate intensity, with the total volume of utterances ($|U|$) rising by 77% between 2021 and 2025. This growth is largely driven by the exponential rise of CH, which evolved from minimal expression to become the party with the highest number of utterances ($|U| > 300$) in 2025. Data confirms that activity is concentrated within the three main political forces (PS, PSD, CH), while smaller parties maintain a smaller number of interventions, validating H2. Despite the significant surge in volume, the relationship between utterances and actions remained stable over the years, indicating a constant level of interactivity and tension in the debate.

Year	#Coop	#Conf	Conf Ratio	#App	#Prot	Prot Ratio
'21	418	482	0.54	163	20	0.11
'22	565	618	0.52	209	95	0.31
'23	622	755	0.55	187	100	0.35
'24	610	856	0.58	300	169	0.36
'25	767	831	0.52	283	157	0.36

Table 3: Distribution of speech acts (Cooperation (Coop) vs. Conflict (Conf)) and reactions (Applause (App) vs. Protests (Prot)) across the years.

Year	BE	PCP	PEV	L	PS	PAN	JPP	GOV	IL	PSD	CDS	CH	P
'21	0.8	0.8	0.8	–	0.5	0.7	–	0.3	1.0	0.8	0.8	0.8	0.0
'22	0.8	0.7	–	0.4	0.4	0.8	–	0.4	0.8	0.8	–	0.7	0.0
'23	0.8	0.7	–	0.5	0.5	0.9	–	0.4	0.7	0.7	–	0.7	0.0
'24	0.7	0.8	–	0.8	0.8	0.9	–	0.6	0.8	0.5	0.5	0.8	0.1
'25	0.6	0.8	–	0.8	0.7	0.6	0.7	0.4	0.8	0.5	0.3	0.7	0.1

Table 4: Conflict ratios by political entity. The color scale indicates the prevalence of conflict, ranging from green (cooperation) to red (conflict). The symbol '–' denotes absence of participation in the debate.

Regarding institutional roles, the data confirm the neutrality of the President of the Assembly (H1). Throughout the five years, the President consistently exhibits the lowest conflict ratio, acting as a cooperative hub in the dialogue flow. However, his $|U|$ increased in parallel with the general intensification, a symptom of the growing need for moderation in a conflict-prone environment.

Table 3 shows stable conflict ratios but a marked rise in protests, refuting H5. This signals intensifying tension, explained by a thematic shift: 2021's consensus topics ("Pandemic", "Health") gave way to polarized issues like "Housing", "Immigration", and "Justice".

To analyze the impact of institutional roles, we examine Table 4 and Figure 4. Data highlights a shift in 2024: PS's moderate conflict spiked upon moving to Opposition, whereas PSD's high levels dropped after taking office.

Role reversal is clearly visible in the 2025 Reaction Matrix (Figure 4), a heatmap that shows sharp polarization: the governing coalition (PSD, CDS-PP) consistently supports the Government (green), while PS, now in Opposition, consistently protests (red). These findings validate H3, confirming that the stance is dictated by the institutional role. This partisan volatility contrasts with the structural stability of the President (H1), who remains the only entity with consistently low conflict ratios in Table 4, regardless of the political cycle.

To analyze the structural alignment between parties, we examine the Dialogue Action Flow generated for the 2025 debate (Figure 5). This graph

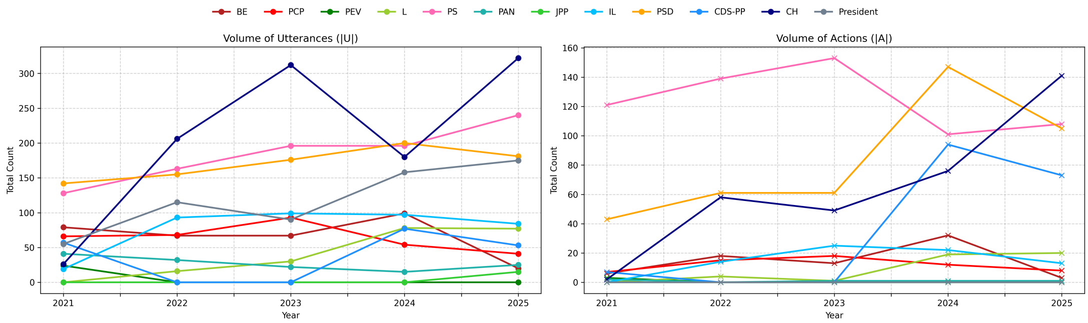

Figure 3: Annual evolution of parliamentary activity by political entity (2021–2025). The left chart displays the volume of utterances ($|U|$), while the right chart shows the volume of actions ($|A|$) for each entity.

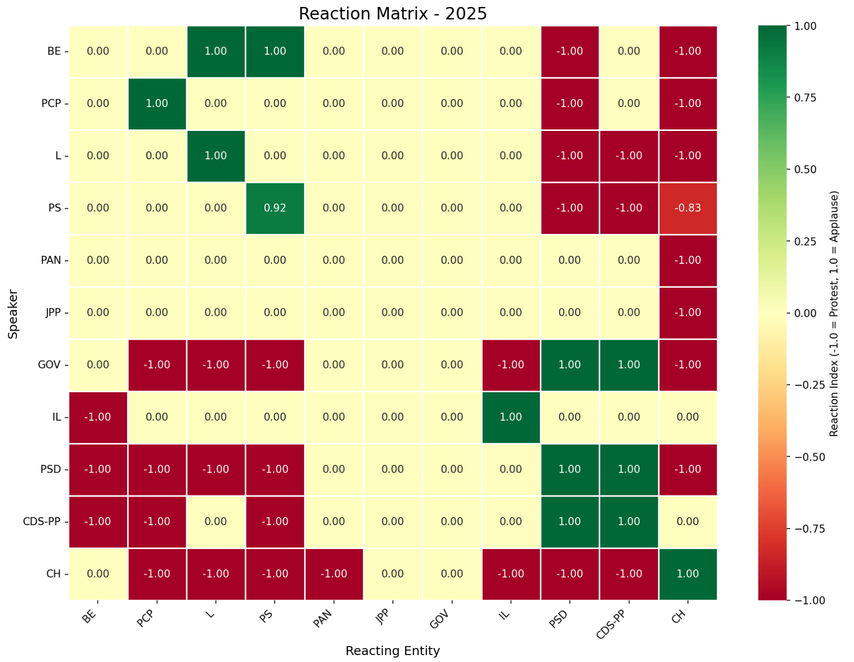

Figure 4: Reaction Matrix for the 2025 debate. Rows represent speakers, and columns reactors. Colors indicate interaction valence: green (Applause, 1.0), red (Protests, -1.0), and yellow (Neutral/No reaction, 0.0).

was generated with $\theta = 0.07$ to show only the most relevant transitions. Numbers in parentheses within each node represent the count of utterances or actions for that specific state, except for SOD and EOD nodes, which instead contain the start and end times of the debate, respectively.

Figure 5 highlights the structural prominence of the dominant entities (PS, PSD, CH, GOV), characterized by a high $|U|$ and $|A|$. In contrast, smaller parties (e.g., PAN, Livre, IL) have minimal counts or disappear entirely due to the threshold θ, demonstrating that they do not play a central role in the main flow of the debate, which confirms H2.

The expectation of a unified right-wing bloc (H4) is only partially verified. The graph demonstrates strong cohesion within the government coalition (AD): there is a very high probability ($p = 0.78$) that CDS-PP applause will come right after PSD applause. However, the remaining right-wing parties do not reflect this alignment. IL begins to protest against the new government. CH remains isolated. This is evidenced in the DAF by the absence of support transitions from other parties and by self-congratulatory loops, where CH's conflictual discourse transitions into internal applause. This pattern corroborates the analysis of the 2025 heatmap (Figure 4), where CH's interaction with the Government remains consistently red (protest), regardless of the change in government.

Table 5 shows that the probability of receiving applause is almost the same, whether the speech expresses support or attack. This finding refutes H6, demonstrating that reactions depend on who is speaking rather than on what is being

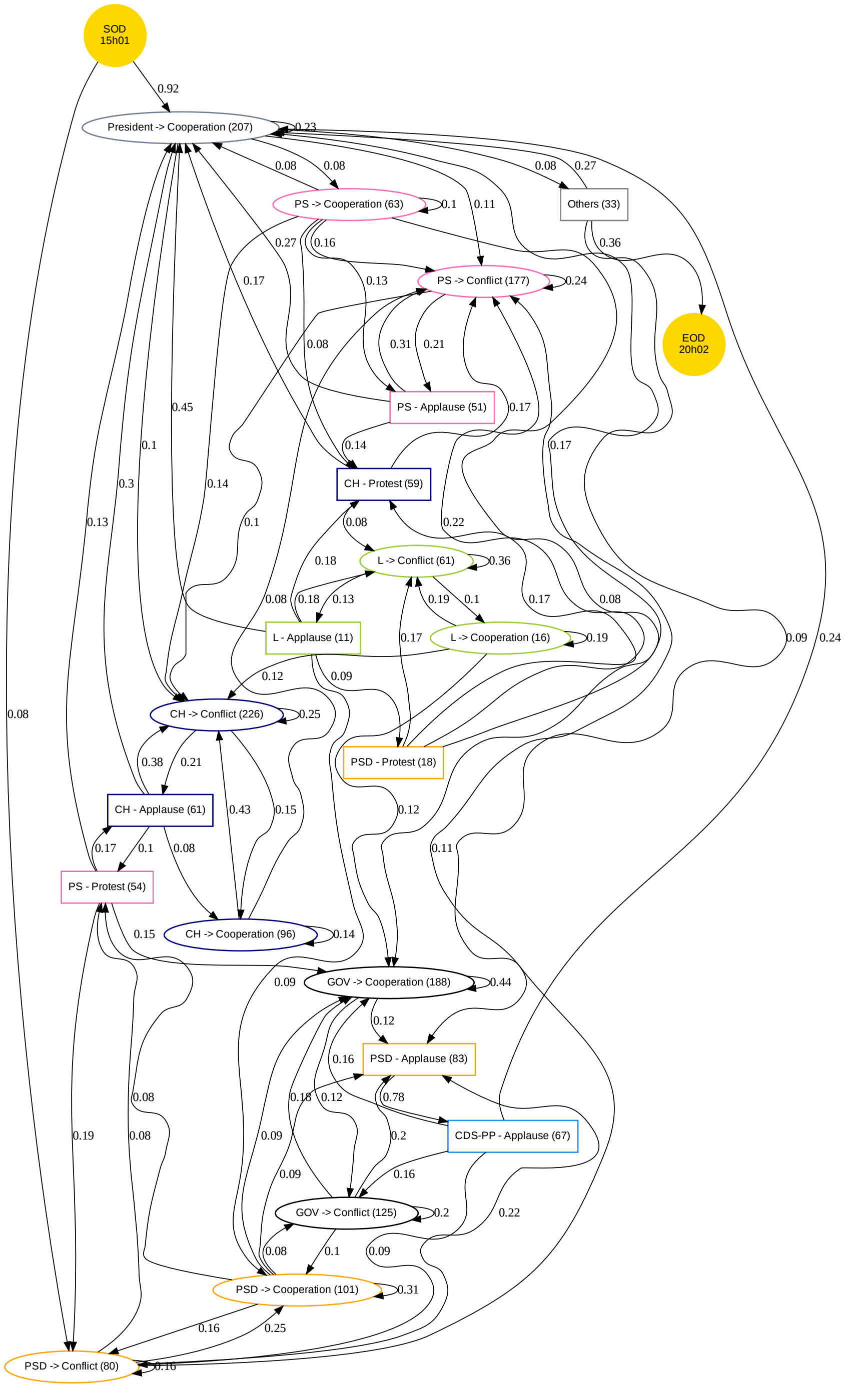

Figure 5: Dialogue Action Flow for 2025 debate.

said. This is also clear in the DAF: when CH attacks ($CH \rightarrow Conflict$), the most common reaction is applause from its own parliamentary bench ($CH \rightarrow Applause$), rather than protests from other parties.

Speech Act	Applause Prob.	Protest Prob.
Conflict	0.68	0.32
Cooperation	0.67	0.33

Table 5: Distribution of reactions for each speech act.

6 Conclusion and Future Work

We presented a framework for analyzing multi-party interaction dynamics by integrating speech and non-verbal signals into the Dialogue Action Flow (DAF). Unlike text-only methods, our approach models the conversation as a directed graph of transitions connecting utterances and actions. By combining unsupervised flow discovery with quantitative metrics, this allows for visualizing the structural dynamics of the debate, capturing patterns of alignment, polarization, and isolation that are often overlooked by isolated text analysis.

Applying this framework to Portuguese State of the Nation debates confirms that institutional roles determine conflict: the President remains neutral, while major parties shift their stance based on whether they are in Government or Opposition, rather than on intrinsic ideological traits. Structurally, despite the dominance of major parties, the right-wing proved fragmented: radical parties remained isolated rather than joining a unified coalition. Finally, non-verbal reactions are strictly partisan, driven by speaker identity over content, and have intensified, signaling a degrading atmosphere.

Future work will focus on validating these patterns along different parliamentary contexts, such as the European Parliament, and incorporating additional metrics to further refine the analysis of interaction dynamics. Additionally, we plan to refine the conflict representation by distinguishing between constructive refutation (directed at propositions) and personal attacks, to provide deeper behavioral insights. Finally, we aim to extend this framework to other multi-party dialogues, such as social media platforms and corporate meetings.

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

José Aires, Aida Cardoso, Rui Pereira, and Amália Mendes. 2024. Compiling and exploring a portuguese parliamentary corpus: Parlamint-pt. In *Proceedings of the IV Workshop on Creating, Analysing, and Increasing Accessibility of Parliamentary Corpora (ParlaCLARIN)@ LREC-COLING 2024*, pages 12–20.

Paulo Almeida, Manuel Marques-Pita, and Joana Gonçalves-Sá. 2021. PTPARL-D: an annotated corpus of forty-four years of portuguese parliamentary debates. *Corpora*, 16(3):337–348.

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 *Proceedings 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 114–117. ACL.

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. MultiWOZ – a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.

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

Isabel Carvalho, Patrícia Ferreira, Ana Alves, Catarina Silva, and Hugo Gonçalo Oliveira. 2024. Analysing customer-support trends in social networks through dialogue flow discovery. In *Proceedings of IberSPEECH 2024*, pages 56–60, Aveiro, Portugal. ISCA.

Isabel Carvalho, Hugo Goncalo Oliveira, and Catarina Silva. 2023. The importance of context for sentiment analysis in dialogues. *IEEE Access*, 11:86088–86103.

Linjie Chen, Min Sun, Jiarong Liu, Pengyong Ding, Yuliang Ma, Li Li, Lili Qian, and Yang Yang. 2023. A method for extracting information from long documents that combines large language models with natural language understanding techniques. In *2023 4th International Conference on Computer, Big Data and Artificial Intelligence (ICCBD+ AI)*, pages 43–47. IEEE.

Carlos Costa, Manuela Aparicio, and Joao Aparicio. 2021. Sentiment analysis of portuguese political parties communication. In *Proceedings of the 39th ACM international conference on design of communication*, pages 63–69.

Michael Cribb and Shivani Rochford. 2018. The transcription and representation of spoken political discourse in the UK House of Commons. *International Journal of English Linguistics*, 8(2):1–14.

Patrícia Ferreira, Isabel Carvalho, Ana Alves, Catarina Silva, and Hugo Gonçalo Oliveira. 2024a. 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. 2023. Unsupervised flow discovery from task-oriented dialogues. In *International Conference on Hybrid Intelligent Systems*, pages 336–347. Springer.

Patrícia Ferreira, Daniel Martins, Ana Alves, Catarina Silva, and Hugo Gonçalo Oliveira. 2024b. Unsupervised flow discovery from task-oriented dialogues. *arXiv preprint arXiv:2405.01403*.

Eleonore Fournier-Tombs and Giovanna Di Marzo Serugendo. 2020. Delibanalysis: Understanding the quality of online political discourse with machine learning. *Journal of Information Science*, 46(6):810–822.

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.

Alexander Holstrup, Lasse Starklit, and Andrea Burattin. 2020. Analysis of information-seeking conversations with process mining. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Chao-Chun Hsu, Sheng-Yeh Chen, Chuan-Chun Kuo, Ting-Hao Huang, and Lun-Wei Ku. 2018. EmotionLines: An emotion corpus of multi-party conversations. In *Proceedings of the eleventh international conference on language resources and evaluation (LREC 2018)*.

Minju Kim, Chaehyeong Kim, Yong Ho Song, Seungwon Hwang, and Jinyoung Yeo. 2022. Botstalk: Machine-sourced framework for automatic curation of large-scale multi-skill dialogue datasets. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5149–5170.

Chien-Chang Lin, Anna YQ Huang, and Stephen JH Yang. 2023. A review of AI-driven conversational chatbots implementation methodologies and challenges (1999–2022). *Sustainability*, 15(5):4012.

Javier Miguel Sastre Martinez and Aisling Nugent. 2022. Inferring ranked dialog flows from human-to-human conversations. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 312–324.

Francisco Pais, Patrícia Ferreira, Catarina Silva, Ana Alves, and Hugo Gonçalo Oliveira. 2024. Question answering for dialogue state tracking in Portuguese. In *Proceedings of the 16th International Conference on Computational Processing of Portuguese-Vol. 1*, pages 461–471.

Sudarshan Pant, Eunchae Lim, Hyung-Jeong Yang, Guee-Sang Lee, Soo-Hyung Kim, Young-Shin Kang, and Hyerim Jang. 2022. Korean drama scene transcript dataset for emotion recognition in conversations. *IEEE Access*, 10:119221–119231.

Ines Reinig, Ines Rehbein, and Simone Paolo Ponzetto. 2024. How to do politics with words: Investigating speech acts in parliamentary debates. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8287–8300.

Eugénio Ribeiro, Ricardo Ribeiro, and David Martins De Matos. 2019. Deep dialog act recognition using multiple token, segment, and context information representations. *Journal of Artificial Intelligence Research*, 66:861–899.

Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, pages 172–180, USA. ACL.

A Said and A Ismail. 2025. Trends in natural language processing for text classification: A comprehensive survey. *International Journal of Science and Research Archive*, 14:1540–1547.

Javier Miguel Sastre Martinez and Aisling Nugent. 2022. Inferring ranked dialog flows from human-to-human conversations. In *Proceedings 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 312–324, Edinburgh, UK. ACL.

Elizabeth Shriberg, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The icsi meeting recorder dialog act (mrda) corpus.

Afonso Sousa and Henrique Lopes Cardoso. 2024. PTPARL-V: Portuguese parliamentary debates for voting behaviour study. In *Proceedings of the IV Workshop on Creating, Analysing, and Increasing Accessibility of Parliamentary Corpora (ParlaCLARIN)@ LREC-COLING 2024*, pages 38–42.

Areej Bin Suhaim and Jawad Berri. 2021. Context-aware recommender systems for social networks: review, challenges and opportunities. *IEEE Access*, 9:57440–57463.

Svitlana Vakulenko, Evangelos Kanoulas, and Maarten De Rijke. 2021. A large-scale analysis of mixed initiative in information-seeking dialogues for conversational search. *ACM Trans. Inf. Syst.*

Jason D Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 7(3):4–33.

Yang Zhang, Hanlei Jin, Dan Meng, Jun Wang, and Jinghua Tan. 2024. A comprehensive survey on process-oriented automatic text summarization with exploration of LLM-based methods. *arXiv preprint arXiv:2403.02901*.