Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)

TEC: A Time Evolving Contextual Graph Model for Speaker State Analysis in Political Debates

Ramit Sawhney[1*], **Shivam Agarwal**[2*], **Arnav Wadhwa**[1] and **Rajiv Ratn Shah**[1]
[1]IIIT Delhi
[2]Manipal Institute of Technology
{ramits, rajivratn}@iiitd.ac.in, shivamag99@gmail.com

Abstract

Political discourses provide a forum for representatives to express their opinions and contribute towards policy making. Analyzing these discussions is crucial for recognizing possible delegates and making better voting choices in an independent nation. A politician's vote on a proposition is usually associated with their past discourses and impacted by cohesion forces in political parties. We focus on predicting a speaker's vote on a bill by augmenting linguistic models with temporal and cohesion contexts. We propose TEC, a time evolving graph based model that jointly employs links between motions, speakers, and temporal politician states. TEC outperforms competitive models, illustrating the benefit of temporal and contextual signals for predicting a politician's stance.

1 Introduction

Parliamentary debates discuss affairs affecting the future development of a country, such as budget revisions and policy reforms. The outcome of such debates entail wide-ranging consequences; for instance, analysts show that withdrawal of the UK from the European Union without a free trade agreement can cost UK's GDP £140 billion within 10 years of the agreement [Ries *et al.*, 2017]. Recent studies [Davoodi *et al.*, 2020] show that analyzing the records of such debates provide information regarding viewpoints of politicians over critical societal factors. These records also aid assessing political candidates and basing voting decisions [Utych, 2019]. However, the esoteric language and opaque parliamentary jargon form a barrier to ordinary citizen's insights into standings of politicians. Recent advances in computational social sciences [Abercrombie, 2018a] have made progress in assessing political stance by using language features of voluminous parliamentary debates. However, analyzing the linguistic traits of a speech is not sufficient to accurately predict a politicians stance on a motion [Bhavan *et al.*, 2019]. Political debates involves multiple contextual elements beyond language, such as party affiliations, topic of debates [Van Dijk, 1977] and historic debates [van Dijk, 2004]. Consider Figure 1, where we

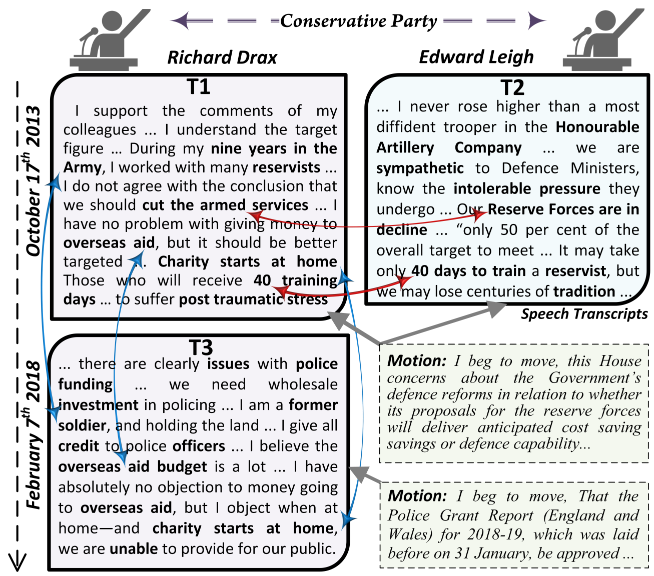

Figure 1: We present Hansard debate transcripts of two speakers from the UK Parliament's House of Commons, over two different motions. We study a speakers association with their party members and their historic speeches for predicting their political stance.

present speech transcripts of two speakers who belong to the Conservative party at two different times. We observe that in the first two transcripts, T_1 and T_2, the speakers support each other and express similar outlooks on the motion. Such partisan forces [Owens, 2003] within political parties is indicative of *intra-party contexts*. Next, amongst transcripts T_1 and T_3, we observe remarkable similarity between two different speeches from the same speaker, debating over two distinct motions, almost four years apart. Such temporal similarities [Curato, 2012] reflect unique characteristics of a speaker, which is reflected across their speeches and plays an essential role for modeling the *temporal state* of a speaker. We also note that the debate's motion governs the overall content of the speech, indicating the presence of *motion-level* contexts in debates. Identifying such similarities in the political ecosystem unfolds the possibility to learn latent patterns among speakers, their political affiliations, and how they target various motions.

[*]Equal contribution.

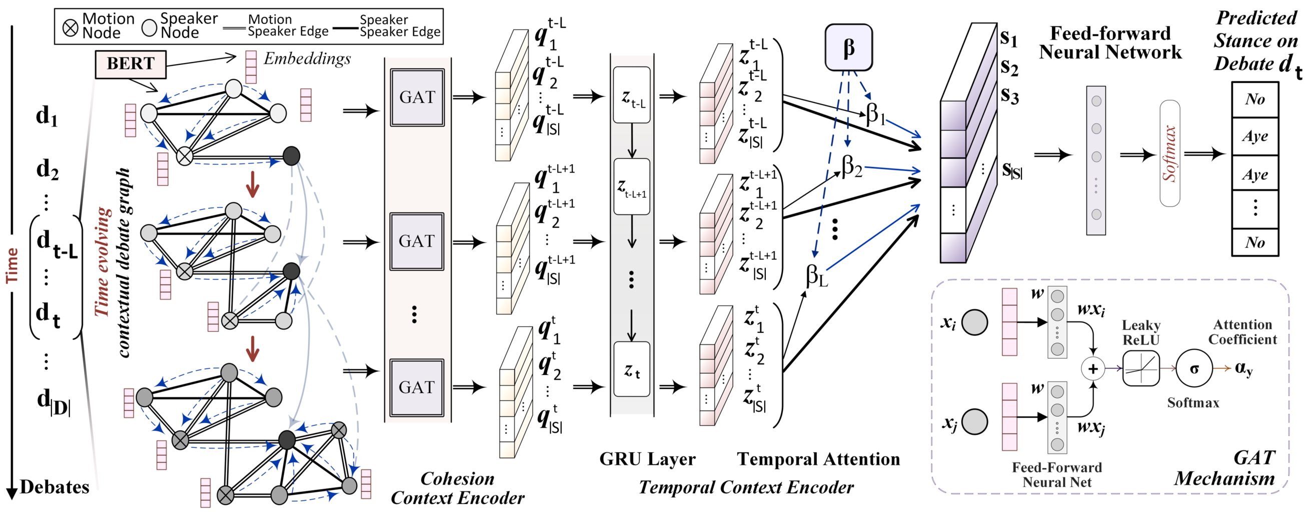

Figure 2: An overview of the context graph, temporal speaker state modelling, TEC model and its components.

Building on existing work [Sawhney *et al.*, 2020] and political theories [van Dijk, 2004; Van Dijk, 1977], we propose **TEC**: **T**ime **E**volving **C**ontextual Graph model, the first dynamic-graph based model for analysing parliamentary debates. Our model enhances prior solutions for political stance analysis [Bhavan *et al.*, 2020; Abercrombie, 2020] by extending the political cohesion based model [Sawhney *et al.*, 2020] to incorporate speaker's temporal states (§3.2,§3.4). We propose a dynamic graph attention networks to jointly learn from political cohesion (§3.3) and the temporal states of a speaker (§3.5). We demonstrate practical applicability through qualitative analysis (§5.5,§5.6). Our contributions are:

1. We identify three types of debate contexts: intra-party, motion-level, and temporal contexts in parliamentary discourse and represent these contexts via dynamic graphs.

2. We propose the first dynamic-graph attention network for jointly learning from language, cohesion context and temporal speaker states in parliamentary debates.

3. Through experiments on more than 30K parliamentary speeches of the UK House of Commons spanning over 20 years, we demonstrate TEC's applicability to provide non-expert citizens insights into political ideologies.

2 Related Work

Politics and Linguistics Analyzing political data acts as a knowledge source that provides insights into cohesion within political parties, stances of politicians towards critical motions for both the general public and across domains including humanities and linguistics [Slembrouck, 1992]. Research at the intersection of Politics and Linguistics spans agreement detection [Duthie and Budzynska, 2018], topic-opinion analysis [Abercrombie, 2018b] and, debate stance classification [Bhavan *et al.*, 2020]. Existing work focuses on these tasks via legislative speeches from the US Congress [Chen *et al.*, 2017], and EU Parliament [Frid-Nielsen, 2018].

Political Stance and Sentiment Analysis NLP has seen a growth in analyzing and mining opinions from political discourse [Abercrombie, 2018a]. Recent approaches [Abercrombie, 2020] show the ability of transformers such as BERT in capturing domain-specific jargon better for debate feature extraction. A promising new direction at the intersection of Politics and NLP is temporal-speaker state modeling and inclusion of context such as engagement in social circles. Recent works [Bhavan *et al.*, 2020] have shown the presence of herd mentality in political stances through shallow graph embeddings by identifying the linguistic similarities between members of the same party. [Sawhney *et al.*, 2020] uses a static graph to encode the cohesion contexts in parties for stance prediction. Despite their success, a common limitation is that they do not consider the dynamic nature of debates and the temporal evolution of speaker states.

3 Time Evolving Contextual Graph Model

Problem Formulation For a debate $d_t \in D = \{d_1, d_2, \ldots, d_{|D|}\}$ at time t, we denote the transcript as $r_{ij} \in T$ corresponding to the speech made by a MP (speaker) $s_i \in S = \{s_1, s_2, \ldots, s_{|S|}\}$ on one specific motion $m_j \in M = \{m_1, m_2, \ldots, m_{|M|}\}$. Where, T is a set of transcripts, and $|\cdot|$ denotes the cardinality of the set. Each speaker s_i is affiliated to only one political party $p \in P = \{p_1, p_2, \ldots, p_{|P|}\}$. Given parliamentary debates for a historical lookback window of L (i.e., $[t-L, t-1]$), the task is to classify the stance $Y \in \{$*'Aye'*, *'No'*$\}$ of the MP s_i on the motion m_j based on the transcript r_{ij} on day t. *'Aye'* and *'No'* denote positive and negative stance, respectively.

TEC consists of two integral components: Cohesion Context Encoder and Temporal Context Encoder to model intra-party and motion contexts across a time-series of Parliamentary debates shown in Figure 2.

3.1 Encoding Parliamentary Texts via BERT

A debate comprises of *Motions*, i.e., expressions over policy positions taken by the government, MPs, etc. The mo-

Dataset Statistics	
No. of Speakers	1,346
No. of Motions	1,995
No. of Transcripts	33,461
Maximum Node Degree	149
Average Node Degree	9 ± 16.4

Table 1: Graph construction statistics

tion is followed by responses from other MPs in the form of *Speeches*. We adopt Bidirectional Encoder Representations from Transformers (BERT) [Devlin *et al.*, 2019] to encode debate transcripts r and motion descriptions m. The domain-specific nature of political speeches and political jargon motivate us to *fine-tune* pre-trained BERT [Sawhney *et al.*, 2020].

We fine-tune BERT on the ParlVote dataset [Abercrombie, 2020], a corpus of debates from the UK Parliament's House of Commons. For each transcript r and, motion $m \in M$, we obtain feature vectors $h_r, h_m \in \mathbb{R}^F$ where, $F = 768$ obtained as the output of the $[CLS]$ token from the final BERT layer.

3.2 Cohesion Contexts in Debates

Decision-makers are subtly influenced by the environment around them [Bode *et al.*, 2014]. Following [Sawhney *et al.*, 2020] we identify two types of cohesion contexts in parliamentary debates— *intra-party context*, and *motion context*. The first, *Intra-Party Context*, captures the influence of the same political affiliations and fellow party members over a speaker's speeches. During debates, some speakers tend to express their *raw* individual opinions, while some may exhibit *homophily*: the likeliness of associated individuals to adopt similar viewpoints [Boucek, 2002]. Next, we present a *Motion Context*, that captures the relationship between a speaker's speech and the motion of the debate.

For each debate d, we represent the relations between speakers and motions in the form of a graph $G = (V, E)$. V and E represent the nodes and edges in the graph. The nodes V consists of the set of speakers S and motions M in the debate d. The edges E are of two types: Speaker-Speaker edges E_{ss} based on intra-party context, and Speaker-Motion edges E_{sm} based on motion context. G is a heterogeneous graph as it has different types of nodes and edges. We now describe the two relations that capture different contexts.

Intra-party Context models the relationship between a MP and the party they belong to. We build on the hypothesis that speakers are influenced by other party members, and there exists a partisan mentality like political cohesion within parties [Owens, 2003]. We represent intra-party context by a Speaker-Speaker edge E_{ss} between two MPs $s_i, s_j \in S$ if both s_i and s_j are affiliated to the same political party $p \in P$.

Motion-level Context encodes the relation between a speaker's speech and the motion. In debates, a speech is based on the current motion of discussion. We represent motion context via Speaker-Motion edge E_{sm} between a speaker s and a motion m corresponding to that debate transcript.
We show the statistics of the graph G in Table 1.

3.3 Cohesion Context Encoder

To capture the cohesion context among speakers and the motions they speak on, we use the cohesion context encoder over the constructed graph G for each debate d. The cohesion context encoder is a graph neural network that propagates information between different speakers and motions. Following [Sawhney *et al.*, 2020] we use we use Graph Attention Networks (GATs) [Velickovic *et al.*, 2018], as each context has a different degree of influence on a speaker's speech.

Here, we define a GAT layer that is applied to one debate d. The input to this GAT is a set of node features $h = \{x_1, x_2, \ldots, x_{|V|}\}$ For each speaker, s and motion m, we set the node features to h_r and h_m extracted from the speech and motion transcript using BERT, respectively. The node features are transformed to context dependent features, $h' = [q_1, q_2, \ldots q_{|V|}]$ based on the influence by its neighbors. We first apply a shared linear transform parameterized by W to all the nodes. Then, we apply a shared self-attention mechanism to each node i in its neighborhood $\mathcal{N}_i$. For each node $j \in \mathcal{N}_i$, we compute attention coefficients α_{ij} which shows the importance of context between nodes i and j. Formally, the attention weights α_{ij} is given as:

$$\alpha_{ij} = \frac{\exp\left(\text{LeakyReLU}(a_w^T[Wx_i \oplus Wx_j])\right)}{\sum\limits_{k \in \mathcal{N}_i} \exp\left(\text{LeakyReLU}(a_w^T[Wx_i \oplus Wx_k])\right)} \quad (1)$$

where, $.^T, \oplus$ represent transposition and concatenation, respectively and a_w is the parameter matrix of a fully connected layer. We use multi-head attention to stabilize training. Formally, U independent executors apply the above attention mechanism. Their outputs are concatenated to yield:

$$q_i = \bigoplus_{k=1}^{U} \text{LeakyReLU}\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^k W^k x_j\right) \quad (2)$$

where α_{ij}^k and W^k denote the attention coefficients and linear transform weight matrix computed by the k^{th} attention head.

3.4 Temporal Speaker State Modelling

Parliamentary debates occur over time, and historic debates often set a stage for upcoming discourses [Curato, 2012]. Furthermore, the sentiment of some MPs towards a motion may evolve over time while, some MPs tend to reflect similar ideologies and linguistic styles all the time [Hampsher-Monk, 1984]. The inherent temporal nature of parliamentary debates suggests the presence of contextual information across a time series of debates [van Dijk, 2004]. We identify a speaker temporal context in Political debates.

Dynamic Graph Construction We model the speaker temporal context by creating a sequence of graphs $\mathcal{G} = \{G_1, G_2, \cdots G_L\}$ corresponding to each time-step in the lookback L. For two consecutive debates d_i and d_{i+1} occurring on time τ_i and τ_{i+1}, respectively with $\tau_{i+1} > \tau_i$. We construct heterogeneous graphs $G_i(V_i, E_i)$ and $G_{i+1}(V_{i+1}, E_{i+1})$ corresponding to the debates d_i and d_{i+1}. The nodes V_i of the graph $G_i(V_i, E_i)$ contains a set of speaker nodes S_i and motion nodes M_i, and the edge set E_i contains

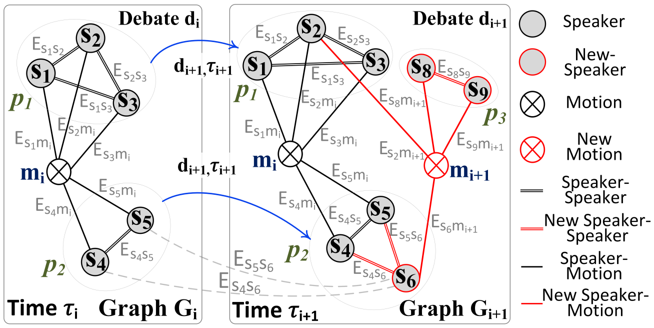

Figure 3: Dynamic Graph Construction

speaker-speaker edges E_{ss_i} and speaker-motion edges E_{sm_i}. As shown in Figure 3, we construct graph $G_{i+1}(V_{i+1}, E_{i+1})$ by updating the graph G_i. The set of nodes V_{i+1} consists of new speakers S_{i+1} that join the debate d_{i+1} along with speaker nodes S_i of debate d_i. A single motion node m_{i+1} (since there is only one motion per debate) is added to the motion nodes M_i of the debate d_i. Similarly, the set of edges E_{i+1} consists of speaker-speaker edges $E_{ss_{i+1}}$ along with the speaker-speaker edges E_{ss_i} of the debate d_i and speaker-motion edges $E_{sm_{i+1}}$ along with the speaker-motion edges E_{sm_i} of the debate d_i. We obtain graph G_{i+1} as,

$$G_{i+1}(V_{i+1}, E_{i+1}) = G_i(V_i \cup V_{i+1}, E_i \cup E_{i+1}) \tag{3}$$

where, $\cup$ denotes union operator. When $i = 1$ i.e., for the first debate d_1 in the lookback, we create the graph $G_1(V_1, E_1)$ with nodes and edges of the current debate d_1 only.

3.5 Temporal Context Encoder

To capture the cohesion as well as speaker temporal contexts in Parliamentary debates, we use the temporal context encoder over the dynamic graph $\mathcal{G}$. We integrate the GAT into a gated recurrent unit which can encode cohesion contexts as well as model temporal speaker states. The temporal context encoder first applies the GAT (§3.3) on each time step of the lookback. As shown in Figure 2, at time τ we obtain the input node features for the GAT $h^\tau = \{x_1^\tau, \ldots, x_{|V_\tau|}^\tau\}$. We update these node features to cohesion context dependent features, $h'^\tau = \{q_1^\tau, \ldots, q_{|V_\tau|}^\tau\}$ via Equation 2 given by,

$$q_i^\tau = \bigoplus_{k=1}^{U} \text{LeakyReLU}\left(\sum_{j \in N_i^\tau} \alpha_{ij}^{k^\tau} W^k x_j^\tau\right) \tag{4}$$

where, N_i^τ denotes the neighborhood of node i on time τ and $\alpha_{ij}^{k^\tau}$ is the attention weight between nodes i and j on time τ.

We now use a GRU to capture the temporal dependencies in the node features q_i^τ. We feed the features q_i^τ of each node i in the graph to obtain the hidden states z_i^τ for day τ as:

$$z_i^\tau = \text{GRU}(q_i^\tau, z_i^{\tau-1}) \quad t - L \leq \tau \leq t \tag{5}$$

Temporal Attention Political studies have shown that some debates have more impact on the upcoming discussions [Box-Steffensmeier and Jones, 1997]. To this end, we apply a temporal attention mechanism which learns to weigh more critical debates. The temporal attention mechanism aggregates the hidden states $\overline{z}_i = [z_i^{t-L}, \ldots, z_i^t]$ from different days into an overall representation o_i using learnt attention weights β_i^τ for each day τ. We formulate this mechanism as:

$$o_i = \sum_\tau \beta_i^\tau z_i^\tau, \;\; \beta_i^\tau = \frac{\exp\left(\overline{z}_i^{\tau\top} W z_i^t\right)}{\sum_{j=t-L}^{t} \exp\left(\overline{z}_i^{j\top} W z_i^t\right)} \tag{6}$$

where, W is a learned linear transform. We now feed the outputs o_i into a fully connected layer followed by softmax activation which outputs the stance of each MP $y_i \in$ 'Aye'/'No'. We optimize TEC using the cross-entropy loss function.

4 Experimental Setup

4.1 Dataset and Preprocessing

We evaluate TEC on the ParlVote dataset consisting of $33,461$ debate transcripts. On average, a speech in ParlVote has 760.2 ± 901.3 tokens. Following [Abercrombie, 2020], we remove non-speech elements, tokenize motions and transcripts and, preserve the texts' original casing. The dataset is fairly balanced with 53.57/46.43% *Aye/No* labels. The transcripts are labeled based on a speaker's vote to their speech, with votes for 'Aye' and 'No' representing positive and negative stance. We split the dataset temporally and obtain 1396 debates from 7/5/1997 and 16/1/2012 for training, 200 debates from 16/1/2012 to 6/1/2014 for validation, and 398 debates with from 6/1/2014 to 5/11/2019 for test.[1]

4.2 Training Setup and Evaluation

We perform all experiments on Tesla T4 GPU. We use grid search for hyperparameter selection for all models and select optimal values based on validation accuracy. We fine-tune BERT using learning rate = $5e-5$, batch size 8, with the AdamW optimizer, for 7 epochs. The GAT layer's output size is set to 16, and number of attention heads to 4. The output space of the GRU is 16 and that of the last fully connected layer to 2 in order to classify the stance into two classes. We use Adam optimizer, which is set with learning rate of $1e-4$.

We compare TEC with the following baselines on accuracy (Acc) & Matthews correlation coefficient (MCC): [2]

Majority class Does not use any features. The majority class in the training set as the predictions for test set.

Multi-layer Perceptron (MLP) A BoW model that utilises only unigram textual features from transcripts as input with TF-IDF selection [Abercrombie, 2018a].

BERT+MLP BERT embeddings fed into a MLP. It is a textual model without additional contexts [Abercrombie, 2020].

Deepwalk Concatenates intra-party graph embeddings with TF-IDF based BoW text features [Bhavan *et al.*, 2020].

GPolS BERT embeddings are fed into the GAT to encode cohesion context. However, they do not use an MPs past speeches for making predictions [Sawhney *et al.*, 2020]

[1]Code at: https://github.com/midas-research/tec-ijcai

[2]All baselines are recreated via the creator's settings.

Model name	Acc↑	MCC↑
Majority class	50.01	0.01
MLP [Abercrombie, 2018a]	65.34	0.36
BERT+MLP [Abercrombie, 2020]	67.31	0.40
Deepwalk [Bhavan *et al.*, 2020]	$70.12^{\dagger}$	$0.42^{\dagger}$
GPolS [Sawhney *et al.*, 2020]	$73.69^{\dagger}$	$0.45^{\dagger}$
TEC (Ours)	$\mathbf{75.07}^{*\dagger}$	$\mathbf{0.47}^{*\dagger}$

Table 2: Performance comparison with baselines (mean of 10 runs). $*$ & † imply improvement over GPolS and BERT+MLP, respectively, is significant ($p < 0.01$), under Wilcoxon's Signed Rank test.

5 Results and Analysis

5.1 Performance Comparison with Baselines

We compare TEC with baseline methods in Table 2. We note that BERT+MLP significantly ($p < 0.01$) outperforms Majority class and BoW based approaches (MLP). We postulate this improvement to fine-tuning BERT to obtain rich embeddings that better represent the language used in the political realm. We observe that graph-based models (GPolS, TEC, Deepwalk) outperform text-only models (MLP, BERT-MLP) likely because graph-based learning captures intra-party and motion level contexts in parliamentary debates. Additionally, TEC and GPolS perform better than featureless embedding method, Deepwalk, potentially because the GAT component augments language features of different nodes for propagating contextual information across nodes. Our work is in line with political studies [van Dijk, 2004] who find that a contextual approach benefits discourse analysis where MPs take up different roles, such as communicative, interactional, and social roles, and contexts are identified when MPs account for group memberships or speak for/against motions. Finally, TEC beats GPolS likely because TEC uses a time-evolving graph rather than a static graph to represent the correlation between an MP's speech across time and enriches contextual node features via temporal speaker state modelling [AJ Willingham, 2017]. We now further probe into each of TEC's components through ablative studies.

5.2 Ablation Study Over Model Components

Table 3 shows how TEC's performance benefits from each of its components. We observe that modeling the speaker temporal context by feeding the text representation of their historical debates to a GRU (T-BERT) leads to significant ($p < 0.01$) gains likely because MPs tend to reflect similar opinions across time [Hampsher-Monk, 1984]. Next, we note large improvements on adding the temporal attention mechanism. We postulate that this improvement stems from the ability of TEC to identify more salient trends in MPs speech across time, similar to observations in [Fort, 2019] who show that there are certain historic speeches that MPs keep referring to. We note a significant improvement ($p < 0.01$) on augmenting speaker-temporal context with cohesion contexts to encode correlations between speakers via GCNs, likely because of its ability to capture the herd mentality in political parties. We also note that adding graph attention mechanism leads to gains, likely because the GAT weighs more critical relations between MPs. Our observation is in line with other

Temporal Context	Cohesion Context	Temporal Attention	Ablation Component	Acc↑	MCC↑
✗	✗	✗	BERT+MLP	67.31	0.40
✓	✗	✗	T-BERT	69.60^{*}	0.42^{*}
✓	✗	✓	T-BERT	70.23^{*}	0.43^{*}
✓	✓	✗	GCN+GRU	71.90^{*}	0.44^{*}
✓	✓	✓	GCN+GRU	72.91^{*}	0.45^{*}
✓	✓	✗	GAT+GRU	$74.45^{*\dagger}$	$0.46^{*\dagger}$
✓	✓	✓	**TEC**	$\mathbf{75.07}^{*\dagger}$	$\mathbf{0.47}^{*\dagger}$

Table 3: Ablation over TEC (mean of 10 runs). * and † imply improvement over BERT+MLP and GPolS, respectively, is statistically significant ($p < 0.01$), under Wilcoxon's Signed Rank test.

Intra Party Context	Motion Context	Acc↑	MCC↑	p-value
✗	✗	70.23	0.43	-
✗	✓	73.56^{*}	0.45^{*}	2e-4
✓	✗	74.20^{*}	0.46^{*}	4e-6
✓	✓	75.07^{*}	0.47^{*}	3e-5

Table 4: Ablation over cohesion contexts (mean of 10 runs). p-values on the right are shown with respect to the row above it. $*$ indicates improvement over T-BERT with temporal attention is statistically significant ($p < 0.01$) under Wilcoxon's Signed Rank test.

studies [Bawn *et al.*, 2012] who find that there exists variance in relationship strength between MPs due to the presence of interest groups and key members of a party. Finally, we observe that the GAT and speaker-temporal context encoder complement each other and capture the dynamic structure [van Dijk, 2004] of political debates.

5.3 Ablation Study Over Cohesion Contexts

We perform an ablation study over the different kinds of cohesion contexts that TEC models in Table 4. We remove the contexts one by one and find that adding intra-party and motion context to speaker-temporal context leads to significant improvements ($p < 0.01$) individually. Interestingly, we note that the performance drop on removing the intra-party context is significantly ($p < 0.01$) larger compared to the removal of motion context, suggesting that while the motion context models the global similarities between speeches on the same motion, the intra-party context is more important to capture the actual decision taken by the MP. We postulate this observation to the existence of herd mentality in political parties, wherein MPs support their party via their votes to make one collective decision [Sawhney *et al.*, 2020].

5.4 Sensitivity to Lookback Period

Figure 5 shows TEC's performance with varying lengths of lookback periods. First, we observe that TEC using temporal attention outperforms GAT+GRU model due to its ability to capture critical historic debates that influence the speaker's current state. Next, we note that using shorter lookbacks leads to poorer performance, likely because of fewer speeches for a coherent speaker state modeling [AJ Willingham, 2017]. As we increase the length of lookback, we note deterioration potentially due to inclusion of old speeches which may not contribute to the speaker's current state. We note that TEC using

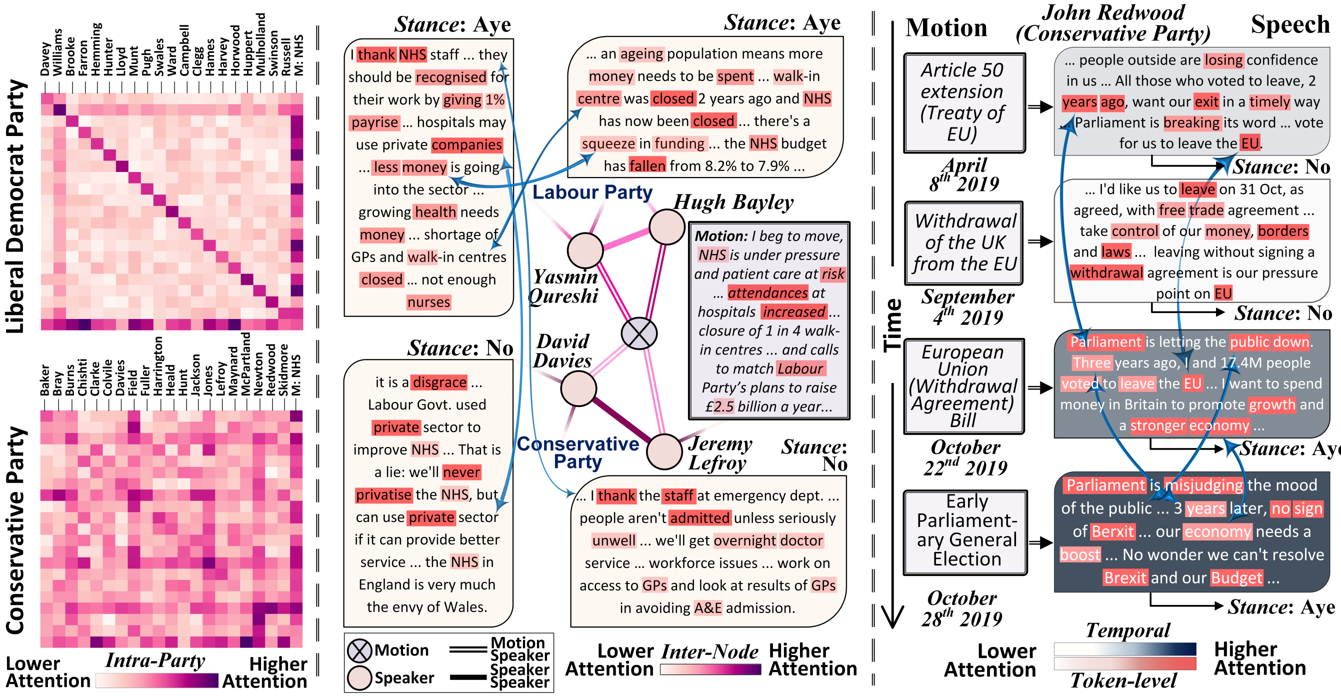

Figure 4: **Left**: Graph attention weights between MPs. **Middle**: Debate sample between Conservative and Labour party with graph and token level attention. **Right**: Historic speeches of an MP with temporal and token level attention. Arrows represent similar terms across transcripts.

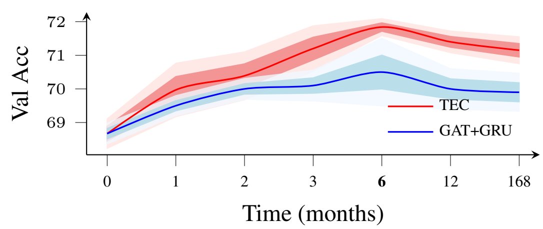

Figure 5: Sensitivity to parameter L

temporal attention can filter crucial debates from large windows to an extent and works best with mid-sized windows.

5.5 Effectiveness of Cohesion Context

We first calculate the attention scores amongst MPs of two parties in Figure 4 (left) and observe high political cohesion in the Conservative party, which is further consolidated by the old age and large size of the party, as such parties show better party discipline [Hayton, 2012]. On the other hand, the Liberal-Democrats, a relatively newer and smaller party, show high attention scores along the diagonal, indicating a large self-dependency of speakers. Next, we perform a qualitative analysis in Figure 4 (middle), and observe that the Conservatives strongly oppose the Labour party's plans to raise £2.5 billion for the national health service. We note higher attention weights between speakers that support each other than those who present diverging arguments, suggesting that the GAT accounts for different degrees of interactions between speakers. However, we note an overall unity since the MPs have a uniform stance on the motion despite their diverging speeches, indicating the presence of partisan forces within these parties. Interestingly, we also note *inter-party* contexts when Lefroy, Jeremy of the conservative party, acknowledges the hard work of health workers similar to Qureshi, Yasmin of the Labour Party. This observation ties up with [Hix *et al.*, 2005] who show that MPs engage in negotiations of inter-party boundaries to create realities that fit their political ends.

5.6 Effectiveness of Temporal Context

We now analyze how an MP's speech evolves over time via a qualitative study in Figure 4 (right). We take an example of John Redwood a member of the Conservative party and a long-term critic of the European Union [BBC, 2016][3]. We observe that John consistently supports the Brexit referendum in his speeches and often makes references to his previous debates, suggesting the presence of temporal patterns in his speeches. Through temporal speaker state modeling TEC is able to capture some of these cues by lending higher attention to debates that influence his current state.

6 Conclusion

Motivated by political studies [van Dijk, 2004] we present TEC, that enriches linguistic features with cohesion contexts and temporal speaker state modeling for accurate political stance estimation. In the future, we plan to adopt such a contextual approach to other political tasks such as roll-call prediction. We also wish to add speaker-specific features such as religion, the constituency they represent, etc. We would also explore other data sources, such as news and social media.

[3]John Redwood: https://en.wikipedia.org/wiki/John_Redwood

References

[Abercrombie, 2018a] Riza Abercrombie, Gavin Batista-Navarro. 'aye' or 'no'? speech-level sentiment analysis of hansard UK parliamentary debate transcripts. In *LREC*. European Language Resources Association, 2018.

[Abercrombie, 2018b] Riza Theresa Abercrombie, Gavin Batista-Navarro. Identifying opinion-topics and polarity of parliamentary debate motions. In *Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 2018.

[Abercrombie, 2020] Riza Abercrombie, Gavin Batista-Navarro. ParlVote: A corpus for sentiment analysis of political debates. In *LREC*, 2020.

[AJ Willingham, 2017] Aaron Kessler Brandon Griggs AJ Willingham. The two-word phrase president trump relies on most. https://edition.cnn.com/2017/04/21/politics/donald-trump-president-speeches-favorite-phrases-trnd/index.html, 2017. Accessed: 2020-12-23.

[Bawn *et al.*, 2012] Kathleen Bawn, Martin Cohen, David Karol, Seth Masket, Hans Noel, and John Zaller. A theory of political parties: Groups, policy demands and nominations in american politics. *Perspectives on Politics*, 2012.

[BBC, 2016] News BBC. Eu exit 'would banish uk austerity', says john redwood. https://www.bbc.com/news/uk-politics-eu-referendum-35873996, 2016. Accessed: 2020-12-23.

[Bhavan *et al.*, 2019] Anjali Bhavan, Rohan Mishra, Pradyumna Prakhar Sinha, Ramit Sawhney, and Rajiv Ratn Shah. Investigating political herd mentality: A community sentiment based approach. In *Proceedings of the 57th Association for Computational Linguistics*, 2019.

[Bhavan *et al.*, 2020] Anjali Bhavan, Mohit Sharma, Ramit Sawhney, and Rajiv Ratn Shah. Analysis of parliamentary debate transcripts using community-based graphical approaches. In *Conference on Artificial Intelligence*, 2020.

[Bode *et al.*, 2014] Stefan Bode, Carsten Murawski, Chun Siong Soon, Philipp Bode, Jutta Stahl, and Philip L. Smith. Demystifying "free will": The role of contextual information and evidence accumulation for predictive brain activity. *Neuroscience & Biobehavioral*, 2014.

[Boucek, 2002] Francoise Boucek. The structure and dynamics of intra-party politics in europe. *Perspectives on European Politics and Society*, 3(3):453–493, 2002.

[Box-Steffensmeier and Jones, 1997] Janet M. Box-Steffensmeier and Bradford S. Jones. Time is of the essence: Event history models in political science. *American Journal of Political Science*, 1997.

[Chen *et al.*, 2017] Wei Chen, Xiao Zhang, Tengjiao Wang, Bishan Yang, and Yi Li. Opinion-aware knowledge graph for political ideology detection. In *IJCAI*, 2017.

[Curato, 2012] Nicole Curato. A sequential analysis of democratic deliberation. *Acta Politica*, 2012.

[Davoodi *et al.*, 2020] Maryam Davoodi, Eric Waltenburg, and Dan Goldwasser. Understanding the language of political agreement and disagreement in legislative texts. In *Association for Computational Linguistics*, 2020.

[Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.

[Duthie and Budzynska, 2018] Rory Duthie and Katarzyna Budzynska. A deep modular rnn approach for ethos mining. In *IJCAI*, 2018.

[Fort, 2019] Hugh Fort. Sir john redwood mp slams parliament for 'letting public down' over brexit. https://www.getreading.co.uk/news/reading-berkshire-news/john-redwood-mp-brexit-vote-17133130, 2019. Accessed: 2020-12-23.

[Frid-Nielsen, 2018] Snorre Sylvester Frid-Nielsen. Human rights or security? positions on asylum in european parliament speeches. *European Union Politics*, 19, 2018.

[Hampsher-Monk, 1984] Iain Hampsher-Monk. Political languages in time - the work of j. g. a. pocock. *British Journal of Political Science*, 14(1):89–116, 1984.

[Hayton, 2012] Richard Hayton. *Reconstructing Conservatism?: The Conservative party in opposition*. 2012.

[Hix *et al.*, 2005] Simon Hix, Abdul Noury, and Gérard Roland. Power to the parties: cohesion and competition in the european parliament 1979-2001. *British Journal of Political Science*, 2005.

[Owens, 2003] John E. Owens. Part 1: Cohesion. *The Journal of Legislative Studies*, 9(4):12–40, December 2003.

[Ries *et al.*, 2017] Charles P. Ries, Marco Hafner, Troy D. Smith, Frances G. Burwell, Daniel Egel, Eugeniu Han, Martin Stepanek, and Howard J. Shatz. *After Brexit: Alternate forms of Brexit and their implications for the UK, the EU and the US*. RAND Corporation, 2017.

[Sawhney *et al.*, 2020] Ramit Sawhney, Arnav Wadhwa, Shivam Agarwal, and Rajiv Ratn Shah. GPolS: A contextual graph-based language model for analyzing parliamentary debates and political cohesion. In *the 28th International Conference on Computational Linguistics*, 2020.

[Slembrouck, 1992] Stef Slembrouck. The parliamentary hansard 'verbatim'report: the written construction of spoken discourse. *Language and literature*, 1992.

[Utych, 2019] Stephen M. Utych. Speaking style and candidate evaluations. *Politics, Groups, and Identities*, 2019.

[Van Dijk, 1977] Teun Adrianus Van Dijk. *Text and context: Explorations in the semantics and pragmatics of discourse*. Longman London, 1977.

[van Dijk, 2004] Teun A. van Dijk. Text and context of parliamentary debates. In *Cross-Cultural Perspectives on Parliamentary Discourse*. 2004.

[Velickovic *et al.*, 2018] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.