

Understanding Politics via Contextualized Discourse Processing

Rajkumar Pujari and Dan Goldwasser

Department of Computer Science

Purdue University

{rpujari, dgoldwas}@purdue.edu

Abstract

Politicians often have underlying agendas when reacting to events. Arguments in contexts of various events reflect a fairly consistent set of agendas for a given entity. In spite of recent advances in Pretrained Language Models, those text representations are not designed to capture such nuanced patterns. In this paper, we propose a Compositional Reader model consisting of encoder and composer modules, that captures and leverages such information to generate more effective representations for entities, issues, and events. These representations are contextualized by tweets, press releases, issues, news articles, and participating entities. Our model processes several documents at once and generates composed representations for multiple entities over several issues or events. Via qualitative and quantitative empirical analysis, we show that these representations are meaningful and effective.

1 Introduction

Over the last decade political discourse has moved from traditional outlets to social media. This process, starting in the '08 U.S. presidential elections, has peaked in recent years, with former-president Trump announcing the firing of top officials as well as policy decisions over Twitter. This presents a new challenge to the NLP community, *how can this massive amount of political content be used to create principled representations of politicians, their stances on issues and legislative preferences?*

This is not an easy challenge as in political texts perspective is often subtle rather than explicit (Fan et al., 2019). Choices of mentioning or omitting certain entities or attributes can reveal the author's agenda. For example, tweeting *“mass shootings are due to a huge mental health problem”* in reaction to a mass shooting is likely to be indicative of opposing gun control measures, despite the lack of an explicit stance in the text.

Recent advances in Pretrained Language Models (PLMs) in NLP (Devlin et al., 2019; Yang et al., 2019; Liu et al., 2019) have greatly improved word representations via contextualized embeddings and powerful transformer units, however such representations alone are not enough to capture nuanced biases in political discourse. Two of the key reasons are: (i) they do not directly focus on entity/issue-centric data and (ii) they only represent linguistic context rather *external* political context.

Our main insight is that effectively detecting such bias from text requires modeling the broader political context of the document. This can include understanding relevant facts related to the event addressed in the text, the ideological leanings and perspectives expressed by the author in the past, and the sentiment/attitude of the author towards the entities referenced in the text. We suggest that this holistic view can be obtained by combining information from multiple sources, which can be of varying types, such as news articles, social media posts, quotes from press releases and historical beliefs expressed by politicians.

For example, consider the following tweet in context of a school shooting: *We need to treat our teachers better! We should keep them safe.* If the author of the tweet is Kamala Harris (known to be pro-gun control), this tweet is likely to be understood as *“ban guns to avoid mass shootings in schools”*. However, if the same tweet is from Mike Pence, whose stance on guns is: *“firearms in the hands of law abiding citizens makes our communities safer”*, the tweet could mean *“arming school teachers stops active shooters”*. This example demonstrates that depending on the context, the same text could signal completely different real-world actions. Hence, we need to model the broader context of the text in order to understand its true meaning. Visualization projecting the tweet representation into a 2D space is given in figure 1, and shows how contextualization from our

model helps disambiguate this example. First, we show the BERT-base representation of the tweet (Tweet-BERT). We also show the BERT-base representations of the known stances of Pence and Harris on gun control ({Mike Pence, Kamala Harris} Stance-BERT). Finally, we apply our model, contextualizing the ambiguous tweet representation with speaker information ({Mike Pence, Kamala Harris} Tweet-Contextualized). The visualization captures how this representation can disambiguate the different interpretations of the same text, and capture their differences.

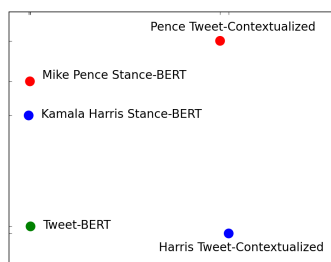


Figure 1: BERT vs. Author-Contextualized Encoder Composer Representation of an Ambiguous Tweet

A computational setting for this approach, *combining text and context analysis*, requires two necessary attributes: (i) an *input* representation that combines all the different types of information meaningfully and (ii) the ability to create a *meaningful unified representation* in one-shot, that captures the complementary strengths of the different inputs.

We address the first challenge by introducing a graph structure that ties together first-person informal (tweets) and formal discourse (press releases and perspectives), third-person current (news) and consolidated (Wikipedia) discourse. These documents are connected via their authors, the issues/events they discuss and the entities mentioned in them. As a clarifying example consider the tweet by former-President Trump “*The NRA is under siege by Cuomo*”. This tweet will be represented in our graph by connecting the text node to the author node (Trump) and the referenced entity node (NY Gov. Cuomo). These settings are shown in Fig. 2.

We propose a novel neural architecture that unifies all the information in the graph in one-shot. *Our architecture generates a distributed representation for each item in the graph that is contextualized by the representations of others*. It can dynamically respond to queries, focusing the induced representation on a specific context. In our exam-

ple, this results in a modified tweet representation helping us characterize Trump’s opinion of Cuomo *in the context of the guns issue*. Our architecture consists of an *Encoder* combining all documents related to a given node to generate an initial node representation and a *Composer*, a Graph Attention Network (GAT), composing the graph structure to generate contextualized node embeddings.

We design two self-supervised learning tasks to train the model and capture structural dependencies over the rich discourse representation, predicting *Authorship* and *Referenced Entity* links over the graph structure. Intuitively, the model is required to understand subtle language usage; *Authorship* prediction requires the model to differentiate between: (i) the language of one author from another and (ii) the language of the author in context of one issue vs another issue. *Referenced Entity* prediction requires understanding the language used by a specific author when discussing a particular entity, given the author’s past discourse.

We focus on a specific graph element—*politicians*, and evaluate their resulting discourse representation on several empirical tasks which capture their stances and preferences. Our evaluation demonstrates the importance of each component of our model and usefulness of the learning tasks. To summarise, our research contributions include:

1. A novel graphical structure connecting various types of documents, entities, issues and events.
2. An effective neural architecture, *Compositional Reader*, processing all information in one-shot, and designing two effective tasks for training it.
3. Designing & performing quantitative and qualitative evaluation showing that our graph structure, neural architecture and learned representations are meaningful and effective for representing politicians and their stances on issues.¹

2 Related Work

Due to recent advances in text representations catalysed by Peters et al. (2018), Vaswani et al. (2017) and followed by Devlin et al. (2019), Liu et al. (2019) and Yang et al. (2019), we are now able to create rich textual representations, effective for many NLP tasks. Although contextual information is captured by these models, they are not explicitly designed to capture entity/event-centric information. Hence, tasks that require such information

¹Repository: https://github.com/pujari-raj कुमार/compositional_learner

(Biessmann, 2016; Johnson and Goldwasser, 2018, 2016; Kornilova et al., 2018a; Chen et al., 2019), would benefit from more focused representations.

Of late, several works attempted to solve such tasks, such as analyzing relationships and their evolution (Iyyer et al., 2016; Han et al., 2019), analyzing political discourse on news and social media (Demszky et al., 2019; Roy and Goldwasser, 2020) and political ideology (Diermeier et al., 2012; Preoțiuc-Pietro et al., 2017; Kulkarni et al., 2018). Various political tasks such as roll call vote prediction (Clinton et al., 2003; Kornilova et al., 2018b; Patil et al., 2019; Spell et al., 2020a; Davoodi et al., 2020), entity stance detection (Mohammad et al., 2016; Fang et al., 2019), hyper-partisan/fake news detection (Li and Goldwasser, 2019; Palić et al., 2019; Baly et al., 2020) require a rich understanding of the context around the entities that are present in the text. But, the representations used are usually limited in scope to specific tasks and not rich enough to capture information that is useful across several tasks.

The Compositional Reader model, that builds upon Devlin et al. (2019) embeddings and consists of a transformer-based Graph Attention Network inspired from Veličković et al. (2017) and Müller et al. (2019), aims to address those limitations via a generic entity-issue-event-document graph, which is used to learn highly effective representations.

Representing legislative preferences is typically done by modeling the ideal point of legislators represented in a Euclidean space from roll-call records (Poole et al., 1997). Recent approaches incorporate bill text information into this representation (Gerrish and Blei, 2011; Nguyen et al., 2015; Kraft et al., 2016; Kornilova et al., 2018c). Most relevant to our work is (Spell et al., 2020b) which uses social media information. We significantly extend these approaches by contextualizing the social media content using a novel architecture.

3 Data

Data	Count	Data	Count
News Events	367	Tweets	86,409
Author Entities	455	Press Releases	62,257
Ref. Entities	10,506	Perspectives	30,446
Wikipedia	455	News Articles	8,244
Total Docs	187,811		

Table 1: Summary statistics of data

We collected US political text related to 8 broad

topics: *guns, LGBTQ rights, abortion, immigration, economic policy, taxes, middle east & environment*. The data focused on 455 members of the US Congress. We collected political text data relevant to above topics from 5 sources: press statements by political entities from ProPublica Congress API², Wikipedia articles describing political entities, tweets by political entities (Congress Tweets, Baumgartner (2019)), perspectives of the senators and congressmen regarding various political issues from ontheissues.org and news articles & background of the those political issues from allsides.com. A total of 187,811 documents were used to train our model, as shown in Tab. 1.

3.1 Event Identification

To identify news events, we use news article headlines. We find the mean (μ) and standard deviation (σ) of the number of articles published per day for each issue. If more than $\mu + \sigma$ number of articles are published on a single day for a given issue, we identify it as the beginning of an event. Then, we skip 7 days and look for a new event.

In our setting, events *within* an issue are non-overlapping. We divide events for each issue separately, hence events for different issues do overlap. These events last for 7 – 10 days on average and hence the non-overlapping assumption within an issue is a reasonable relaxation of reality. To illustrate our point: coronavirus and civil-rights are separate issues and hence have overlapping events. An example event related to coronavirus could be “First case of COVID-19 outside of China”. Similarly an event about civil-rights could be “Officer part of George Floyd killing suspended”. We inspected the events manually by random sampling. More example events are in the appendix.

3.2 Data Pre-processing

We use Stanford CoreNLP tool (Manning et al., 2014), Wikifier (Brank et al., 2017) and BERT-base-uncased implementation by Wolf et al. (2019) to preprocess data for our experiments. We tokenize the documents, apply coreference resolution and extract referenced entities from each document. The referenced entities are then wikified using Wikifier tool (Brank et al., 2017). The documents are then categorized by issues and events. News articles from allsides.com and perspectives from ontheissues.org are already classified by issues. We

²<https://projects.propublica.org/api-docs/congress-api/>

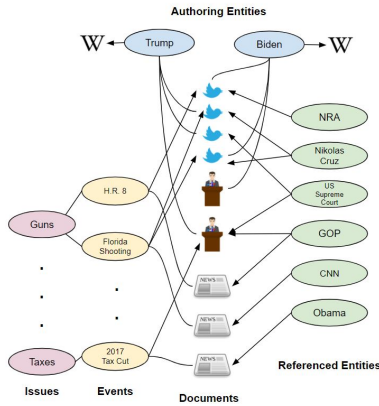


Figure 2: Example Text Graph from Graph Generator

use keyword based querying to extract issue-wise press releases from Propublica API. We use hashtag based classification for tweets. A set of gold hashtags for each issue was created and the tweets were classified accordingly³. Sentence-wise BERT-base embeddings of all documents are computed.

3.3 Query Mechanism

We implemented a query mechanism to obtain relevant subsets of data from the corpus. Each query is a triplet of *entities, issues & lists of event indices corresponding to each of the issues*. Given a query triplet, news articles related to the events for each of the issues, Wikipedia articles for each of the entities, background descriptions of the issues, perspectives of each entity regarding each of the issues and tweets & press releases by each of the entities related to the events in the query are retrieved. Referenced entities for each of the sentences in documents and sentence-wise BERT embeddings of the documents are also retrieved.

4 Compositional Reader

In this section, we describe the architecture of the proposed ‘Compositional Reader’ model in detail. It contains 3 key components: Graph Generator, Encoder and Composer. Given a query output of the query mechanism from Sec. 3.3, Graph Generator creates a directed graph with entities, issues, events and documents as nodes. Encoder is used to generate initial node embeddings for each of the nodes. Composer is a transformer-based Graph Attention Network (GAT) followed by a pooling layer. It generates the final node embeddings and a single summary embedding for the query graph. Each component is described below.

³Data collection is detailed in appendix

4.1 Graph Generator

Given the output of the query mechanism for a query, the Graph Generator creates a directed graph with 5 types of nodes: authoring entities, referenced entities, issues, events and documents. Directed edges are used by Composer to update source node representations using destination nodes. We design the topology with the main goal of capturing the representations of events, issues and referenced entities that reflect author’s opinion about them. We add edges from issues/events to author’s documents but omit the other direction as our main goal is to contextualize issues/events using author’s opinions.

Bidirectional edges from authors to their Wikipedia articles, tweets, press releases and perspectives, from issues to their background description, events and from events to news articles describing them are added. Uni-directional edges from events to tweets and press releases, from issues to author perspectives and from referenced entities to the documents that mention them are added. An example graph is shown in Fig. 2.

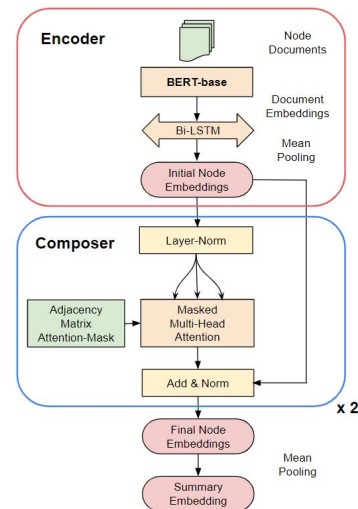


Figure 3: Encoder-Composer Architecture

4.2 Encoder

Encoder is used to compute the initial node embeddings. It consists of BERT followed by a Bi-LSTM. For each node, it takes a sequence of documents as input. The documents are ordered temporally. The output of Encoder is a single embedding of dimension d_m for each node. Given a node $\mathcal{N} = \{D_1, D_2, \dots, D_d\}$ consisting of d documents, for each document D_i , contextualized embeddings of all the tokens are computed using BERT. Token embeddings

are computed sentence-wise to avoid truncating long documents. Then, token embeddings of each document are mean-pooled to get the document embeddings $\vec{N}^{bert} = \{\vec{D}_1^{bert}, \vec{D}_2^{bert}, \dots, \vec{D}_d^{bert}\}$ where $\vec{D}_i^{bert} \in \mathbb{R}^{1 \times d_m}$, d_m is the dimension of a BERT token embedding. The sequence \vec{N}^{bert} is passed through a Bi-LSTM to obtain an output sequence $\vec{E} = \{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_d\}$, $\vec{e}_i \in \mathbb{R}^{1 \times h}$, where $h/2$ is the hidden dimension of the Bi-LSTM, we set $h = d_m$ in our model. Finally, the output of Encoder is computed by mean-pooling the sequence \vec{E} . We use BERT-base-uncased model in our experiments where $d_m = h = 768$. Initial node embeddings of all the document nodes are set to Encoder output of the documents themselves. For authoring entity nodes, their Wikipedia descriptions, tweets, press releases and perspective documents are passed through Encoder. For issue nodes, background description of the issue is used. For event nodes, all the news articles related to the event are used. For referenced entities, all documents that mention the entity are used.

4.3 Composer

Composer is a transformer-based graph attention network (GAT) followed by a pooling layer. We use the transformer encoding layer proposed by Vaswani et al. (2017), without the position-wise feed forward layer, as graph attention layer. Position-wise feed forward layer is removed as in contrast with sequence-to-sequence prediction tasks, nodes in a graph usually have no ordering relationship between them. Adjacency matrix of the graph is used as the attention mask. Self-loops are added for all nodes so that updated representation of the node also depends on its previous representation. Composer module uses $l = 2$ graph attention layers in our experiments. Composer module generates updated node embeddings $\mathbb{U} \in \mathbb{R}^{n \times d_m}$ and a summary embedding $\mathbb{S} \in \mathbb{R}^{1 \times d_m}$ as outputs. The output dimension of node embeddings is 768. Equations that describe Composer unit are:

$$\begin{aligned}
 \mathbb{E} &\in \mathbb{R}^{d_m \times n}, \mathcal{A} \in \{0, 1\}^{n \times n} \\
 \mathbb{G} &= LN(\mathbb{E}) \\
 Q &= W_q^T \mathbb{G}, K = W_k^T \mathbb{G}, V = W_v^T \mathbb{G} \\
 M &= \frac{Q^T K}{\sqrt{d_k}}, M = \text{mask}(M, \mathcal{A}) \\
 \mathbb{O} &= MV^T, \mathbb{U} = W_o^T \mathbb{O} + \mathbb{E} \\
 \mathbb{S} &= \text{mean-pool}(\mathbb{U})
 \end{aligned} \tag{1}$$

where n is number of nodes in the graph, d_m is the dimension of a BERT token embedding, d_k ,

d_v are projection dimensions, n_h is number of attention heads used and $Q \in \mathbb{R}^{n_h \times d_k \times n}$, $K \in \mathbb{R}^{n_h \times d_k \times n}$, $V \in \mathbb{R}^{n_h \times d_v \times n}$, $\mathbb{O} \in \mathbb{R}^{n_h \times d_v \times n}$, $M \in \mathbb{R}^{n_h \times n \times n}$. $W_q \in \mathbb{R}^{d_m \times n_h d_k}$, $W_k \in \mathbb{R}^{d_m \times n_h d_k}$, $W_v \in \mathbb{R}^{d_m \times n_h d_v}$ and $W_o \in \mathbb{R}^{n_h d_v \times d_m}$ are weight parameters to be learnt. $\mathbb{E} \in \mathbb{R}^{d_m \times n}$ is the output of the encoder. $\mathcal{A} \in \{0, 1\}^{n \times n}$ is the adjacency matrix. We set $n_h = 12$ and $d_k = d_v = 64$.

5 Learning Tasks

We design two learning tasks to train the Compositional Reader model: *Authorship Prediction* and *Referenced Entity Prediction*. Both the tasks are intuitively designed to train the model to learn the association between the author node representation and the language used by the particular author. These tasks are two variations of link prediction over the graph. The tasks are detailed below.

5.1 Authorship Prediction

Authorship Prediction is designed as a binary classification task. In this task, the model is given a graph generated by the graph generator in subsection 4.1, an author node and a document node. The task is to predict whether or not the document was authored by the input author.

Intuition behind this learning task is to enable our model to learn differentiating between: 1) language of an author’s first-person discourse vs. third person discourse in news articles, 2) language of an author vs. language used by other authors and 3) language of an author in context of one issue vs. in context of other issues. The model sees documents by the author in the graph and learns to decide whether or not the input document is by the same author and talking about the same issue.

Model	IS Acc	IS F1	OS Acc	OS F1
Authorship Prediction				
BERT Adap.	93.01	92.31	95.56	95.20
Comp. Reader	99.49	99.47	99.42	99.39
Reference Entity Prediction				
BERT Adap.	76.57	75.21	76.26	73.67
Comp. Reader	78.52	77.51	78.98	78.62

Table 2: Learning Tasks In-Sample & Out-Sample Results on Test Data. Acc.denotes Accuracy. F1 Score for the Positive Class is Reported.

Data Training data for the task was created as follows: for a particular author-issue pair, we obtain a data graph similar to Fig. 2 using the query mechanism in subsection 3.3. To create a positive data sample, we sample a document d_i authored by the

entity a_i and remove the edges between the nodes a_i and d_i . Negative samples were designed carefully in 3 batches to align with our above task objectives. In the first batch, we sample news article nodes from the same graph. In the second batch, we obtain tweets, press releases and perspectives of the same author but from a different issue. In the third batch, we sample documents related to the same issue but from other authors. We generate 421,284 samples in total, with 252,575 positive samples and 168,709 negative samples. We randomly split the data into training set of 272,159 samples, validation set of 73,410 samples and test set of 75,715 samples.

Architecture We concatenate the initial and final node embeddings of the author, document and also the summary embedding of the graph to obtain inputs to the fine-tuning layers for Authorship Prediction task. We add one hidden layer of dimension 384 before the classification layer.

Out-sample Evaluation We perform out-sample experiments to evaluate generalization capability to unseen author data. We train the model on training data from two-thirds of politicians and test on the test sets of others. Results are shown in Tab. 2.

Graph Trimming We perform graph trimming to make the computation tractable on a single GPU. We randomly drop 80% of the news articles, tweets and press releases that are not related to the event to which d_i belongs. We use graphs with 200-500 nodes and batch size of 1.

5.2 Referenced Entity Prediction

This is also a binary classification task. Given a data graph, a document node with a masked entity and a referenced entity node the graph, the task is to predict whether the referenced entity is same as the masked entity. Intuition behind this learning task is to enable our model to learn the correlation between the language of the author in the document and the masked entity. For example, in context of recent Donald Trump’s impeachment hearing, consider the sentence ‘X needs to face the consequences of their actions’. Depending upon the author, X could either be ‘Donald Trump’ or ‘Democrats’. Learning to understand such correlations by looking at other documents from the same author is effective in capturing meaningful author representations.

Data To create training data, we sample a document from the data graph. We mask the most

frequent entity in the document with a generic <ENT> token. We remove the link between the masked entity and the document in the data graph. We sample another referenced entity from the graph to generate a negative example. We generated 252,578 samples for this task, half of them positive. They were split into 180,578 training samples, validation and test sets of 36,400 samples each.

Architecture We use fine-tuning architecture similar to Authorship Prediction on top of Compositional Reader for this task as well. We keep separate fine-tuning parameters for each task as they are fundamentally different prediction problems. Compositional Reader is shared. We apply graph trimming for this task as well. We also perform out-sample evaluation for this learning task.

Results Performance of the BERT Adaptation baseline and Compositional Reader model are shown in Tab 2. On Authorship Prediction, out-sample performance doesn’t drop for either model. This shows the usefulness of our graph formulation which allows the models to learn linguistic nuances. On Referenced Entity Prediction, F1 score for our model improves from 77.51 from in-sample to 78.62 on out-sample while BERT adaptation baseline’s F1 drops slightly from 75.21 to 73.67

6 Evaluation

We evaluate our model and pre-training tasks in a systematic manner using several quantitative tasks and qualitative analysis. Quantitative evaluation includes *Grade Paraphrase* task, *Grade Prediction* on *National Rifle Association (NRA)* and *League of Conservation Voters (LCV)* grades data followed by *Roll Call Vote Prediction* task. Qualitative evaluation includes entity-stance visualization for issues and Opinion Descriptor Generation. We compare our model’s performance to BERT representations, the BERT adaptation baseline and representations from the Encoder module. Baselines and the evaluation tasks are detailed below. Further evaluation tasks are in the appendix.

6.1 Baselines

BERT: We compute the results obtained by using pooled BERT representations of relevant documents for each of the quantitative tasks. Details of the chosen documents and the pooling procedure is described in the relevant task subsections. We chose BERT-base over BERT-large due to the complexity of running the learning tasks on embedding

Model	Paraphrase All Grades	Paraphrase A/F Grades	NRA Val Acc	NRA Test Acc	LCV Val Acc	LCV Test Acc
BERT	41.55%	38.52%	55.93 ± 0.72	54.83 ± 1.79	54.28 ± 0.31	52.63 ± 1.21
BERT Adap.	37.54%	42.62%	71.23 ± 3.93	69.95 ± 3.33	60.58 ± 1.56	59.09 ± 1.77
Encoder	56.16%	48.36%	83.95 ± 1.24	81.34 ± 0.86	65.10 ± 0.46	63.42 ± 0.35
Comp. Reader	63.32%	63.93%	84.19 ± 0.98	81.62 ± 1.23	65.55 ± 1.33	62.24 ± 0.56

Table 3: Results of *Grade Paraphrase* and *Prediction* tasks. Acc denotes Accuracy, NRA and LCV denote Grade Prediction tasks. Mean ± Std. Dev for 5 random seeds for Grade Prediction showing statistical significance.

Session	Majority Class (%)	Accuracy (%)		Precision (%)		Recall (%)		F1 (%)	
		NW-GL	CR	NW-GL	CR	NW-GL	CR	NW-GL	CR
106	83.23	85.04	85.65	91.89	91.67	90.22	91.27	91.05	91.47
107	85.78	87.62	88.30	90.12	89.48	95.37	97.17	92.67	93.16
108	87.02	92.03	92.27	93.46	93.52	97.59	97.83	95.48	95.32
109	83.57	85.42	87.23	88.38	88.39	93.84	97.33	91.49	92.65
Average	84.90	87.53	88.36	90.96	90.77	94.26	95.90	92.67	93.15

Table 4: Roll Call Prediction Results. NW-GL represents the best performing model of Patil et al. (2019) as replicated by us using their official implementation. CR represents Compositional Reader results. The improvements are statistically significant as per McNemar’s test.

dimension 768 vs 1024. A bigger embedding dimension results in lesser context (lesser number of nodes in the graph).

Encoder Representations: We compare the performance of our model to the results obtained by using initial node embeddings generated from the Encoder for each of the quantitative tasks.

BERT Adaptation Model: We design a BERT adaptation baseline for the learning tasks. BERT adaptation architecture is same as the Encoder of the Compositional Reader model. While Encoder’s parameters are trained via back-propagation through the Composer, BERT adaptation model is directly trained on learning tasks. In BERT adaptation, once we generate the data graph, we pass the mean-pooled sentence-wise BERT embeddings of the node documents through a Bi-LSTM. We mean-pool the output of Bi-LSTM to get node embeddings. We use fine-tuning layers on top of thus obtained node embeddings for both the learning tasks. BERT Adaptation baseline allows us to showcase the importance of our proposed training tasks via comparison with BERT-base representations. It also demonstrates the usefulness of Composer.

6.2 Grade Paraphrase Task

National Rifle Association (NRA) assigns letter grades (A+, A, . . . , F) to politicians based on candidate questionnaire and their gun-related voting. We evaluate our representations on their ability to predict these grades. We collected the historical data of politicians’ NRA grades from everytown.org.

In *Grade Paraphrase* task, we evaluate our rep-

resentations directly *without* training on the NRA data. Grades are divided into two classes: grades including and above *B+* are in positive class and grades from *C+* to *F* are clustered into negative. We formulate representative sentences for them:

- POSITIVE: *I strongly support the NRA*
- NEGATIVE: *I vehemently oppose the NRA*

For each politician, we obtain data graph for the issue *guns*. We input the data graph to Compositional Reader model and use the node embeddings of the author politician (\vec{n}_{auth}), issue *guns* (\vec{n}_{guns}) and referenced entity *NRA* (\vec{n}_{NRA}). For some politicians, \vec{n}_{NRA} is not available as they have not referenced *NRA* in their discourse. We just use \vec{n}_{auth} and \vec{n}_{guns} for them. We compute BERT-base embeddings for the representative sentences to obtain $\vec{p\vec{o}s}_{NRA}$ and $\vec{n\vec{e}g}_{NRA}$. We mean-pool the three embeddings \vec{n}_{auth} , \vec{n}_{guns} and \vec{n}_{NRA} to obtain \vec{n}_{stance} . We compute cosine similarity of \vec{n}_{stance} with $\vec{p\vec{o}s}_{NRA}$ & $\vec{n\vec{e}g}_{NRA}$. Politician is assigned the higher similarity class.

We compare our model’s results to BERT-base, BERT adaptation and Encoder embeddings. For BERT-base, we compute \vec{n}_{stance} by mean-pooling the sentence-wise BERT embeddings of tweets, press releases and perspectives of the author on all events related to the issue *guns*. Results are shown in Tab. 3. Compositional Reader achieves 63.32% accuracy. Encoder embeddings get 56.16%. Mean-pooled BERT-base embeddings get 41.55%. Using node embeddings from BERT adaptation model yields 37.54%. When we evaluate using only ‘A’/‘F’ grades, we obtain 63.93% accuracy for Com-

positional Reader, 48.36% for Encoder, 42.62% for BERT adaptation and 38.52% for BERT-base.

6.3 Grade Prediction Task

NRA Grades This is designed as a 5-class classification task for grades $\{A, B, C, D \& F\}$. We train a simple feed-forward network with one hidden layer. The network is given 2 inputs \vec{n}_{auth} & \vec{n}_{guns} . When \vec{n}_{NRA} is available for an author, we set $\vec{n}_{guns} = \text{mean}(\vec{n}_{NRA}, \vec{n}_{guns})$. The output is a binary prediction.

We perform $k = 10$ -fold cross-validation for this task. We repeat the entire process for 5 random seeds and report the results with confidence intervals. We perform this evaluation for BERT-base, BERT adaptation, Encoder and Compositional Reader. To compute \vec{n}_{auth} for BERT-base, we mean-pool the sentence-wise embeddings of all author documents on *guns*. For \vec{n}_{guns} , we use the background description document of issue *guns*. Results on the test set are in Tab. 3.

LCV Grades This is similar to NRA Grade Prediction task. This is a 4-way classification task. *League of Conservation Voters* (LCV) assigns a score ranging between 0-100 to each politician depending upon their environmental voting activity. We segregate politicians into 4 classes (0 – 25, 25 – 50, 50 – 75, 75 – 100). We obtain input to the prediction model by concatenating \vec{n}_{auth} and $\vec{n}_{environment}$. We use same fine-tuning architecture as NRA Grade Prediction task.

Results of Grade Prediction task are shown in Tab. 3. On *NRA Grade Prediction*, which is a 5-way classification task, our model achieves an accuracy of 81.62 ± 1.23 on the test set. Our model outperforms BERT representations by 26.79 ± 3.02 absolute points on the test set. On *LCV Grade Prediction* task which is a 4-way classification, our model achieves 9.61 ± 1.77 point improvement over BERT representations.

6.4 Roll Call Vote Prediction Task

This task was proposed in Patil et al. (2019). We skip the finer details of the task for brevity. The task aims to predict the voting behaviour of US politicians on roll call votes. Given the bill texts and voting history of the politicians, the aim is to predict future voting patterns of the politicians. We inject our politician author embeddings from Compositional Reader model to improve the performance on the task. We input all the politician first-person discourse from our data to compute politician author

embeddings using Compositional Reader model. We use these embeddings to initialize the legislator embeddings in their news-augmented glove model, which is their best performing model. We use the data splits provided in their official implementation. We use their code to reproduce their results. Results are shown in Tab. 4.

6.5 Qualitative Evaluation

Politician Visualization We perform Principle Component Analysis (PCA) on issue embeddings (\vec{n}_{issue}) of politicians obtained using the same method as in NRA Grade prediction. We show one such interesting visualization in Fig. 4. Sen. McConnell, a Republican who expressed right-wing views on both *environment* and *guns*. Sen. Sanders, a Democrat that expressed left-wing views on both. Rep. Rooney, a Republican who expressed right-wing views on *guns* but left-wing views on *environment*. Fig. 4 demonstrates that this information is captured by our representations. Additional such visualizations are included in the appendix.

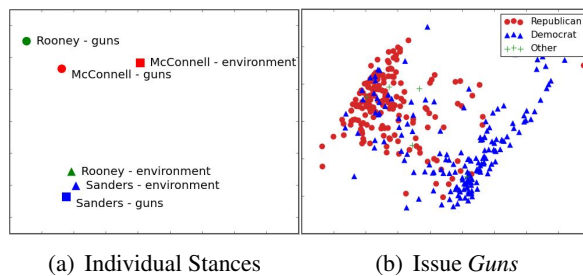


Figure 4: PCA Visualizations of Politician Embeddings

Issue Visualization We present visualization of politicians on the issue *guns* in Fig. 4. We observe that *guns* tends to be a polarizing issue. This shows that our representations are able to effectively capture relative stances of politicians. We observe that issues that have traditionally had clear conservative vs liberal boundaries such as *guns* & *abortion* are more polarized compared to issues that evolve with time such as *middle-east* & *economic-policy*. Visualization for issue *abortion* is in the appendix.

6.6 Opinion Descriptor Generation

This task demonstrates a simple way to interpret our contextualized representations as natural language descriptors. It is an unsupervised qualitative evaluation task. We generate opinion descriptors for authoring entities for specific issues. We use the final node embedding of the issue node (\vec{n}_{issue})

Issue	Opinion Descriptors	Issue	Opinion Descriptors
Mitch McConnell	Republican	Nancy Pelosi	Democrat
<i>abortion</i>	fundamental, hard, eligible, embryonic, unborn	<i>abortion</i>	future, recent, scientific, technological, low
<i>environment</i>	achievable, more, unobjectionable, favorable, federal	<i>environment</i>	forest, critical, endangered, large, clear
<i>guns</i>	substantive, meaningful, outdone, foreign, several	<i>guns</i>	constitutional, ironclad, deductible, unlawful, fair
<i>immigration</i>	federal, sanctuary, imminent, address, comprehensive	<i>immigration</i>	immigrant, skilled, modest, overall, enhanced
Donald Trump	Republican	Joe Biden	Democrat
<i>guns</i>	terrorist, public, ineffective, huge, inevitable, dangerous	<i>guns</i>	banning, prohibiting, ban, maintaining, sold
<i>immigration</i>	early, dumb, birthright, legal, difficult	<i>taxes</i>	progressive, economic, across-the-board, annual, top

Table 5: Opinion Descriptor Labels for Politicians. They show the most representative adjectives used by the politicians in context of each issue.

Model	Accuracy	Model	Accuracy
Comp.Reader	63.32%		
w/o Tweets	63.32%	Only Tweets	40.11%
w/o Press	63.04%	Only Press	55.87%
w/o Persp.	59.31%	Only Persp.	60.74%

Table 6: Ablation Study on *Grade Paraphrase* task for various types of documents

for each politician to generate opinion descriptors.

Inspired from Han et al. (2019), we define our candidate space for descriptors as the set of adjectives used by the entity in their tweets, press releases and perspectives related to an issue. Although Han et al. (2019) uses verbs as relationship descriptor candidates, we opine that adjectives describe opinions better. We compute the representative embedding for each descriptor by mean-pooling the contextualized embeddings of that descriptor from all its occurrences in the politician’s discourse. This is the one of the key differences with prior descriptor generation works such as Han et al. (2019) and Iyyer et al. (2016). They work in a static word embedding space. But, our embeddings are contextualized and also reside in a higher dimensional space. In an unsupervised setting, this makes it more challenging to translate from distributional space to natural language tokens. Hence, we restrict the candidate descriptor space more than Han et al. (2019) and Iyyer et al. (2016). We rank all the candidate descriptors according to cosine similarity of its representative embedding with the vector \vec{n}_{issue} .

We present some of the results in Tab. 5. In contrast to Iyyer et al. (2016) and Han et al. (2019), our model doesn’t need the presence of both the entities in text to generate opinion descriptors. This is often the case in first person discourse. Results are shown in table 5.

6.7 Ablation Study

Further, we investigate the importance of various components. We perform ablation study over various types of documents on the *NRA Grades Paraphrase* task. the results are shown in Tab. 6. They indicate that *perspectives* are most useful while *tweets* are the least useful documents for the task. As *perspectives* are summarized ideological leanings of politicians, it is intuitive that they are more effective for this task. Tweets are informal discourse and tend to be very specific to a current event, hence they are not as useful for this task.

7 Conclusion

We tackle the problem of *understanding politics*, i.e., creating unified representations of political figures capturing their views and legislative preferences, directly from raw political discourse data originating from multiple sources. We propose the Compositional Reader model that composes multiple documents in one shot to form a unified political entity representation, while capturing the real-world context needed for representing the interactions between these documents.

We evaluate our model on several qualitative and quantitative tasks. We outperform BERT-base model on both types of tasks. Our qualitative evaluation demonstrate that our representations effectively capture nuanced political information.

Acknowledgements

We thank Shamik Roy, Nikhil Mehta and the anonymous reviewers for their insightful comments. This work was partially supported by an NSF CAREER award IIS-2048001.

References

Ramy Baly, Georgi Karadzhov, Jisun An, Haewoon Kwak, Yoan Dinkov, Ahmed Ali, James Glass, and Preslav Nakov. 2020. [What was written vs. who](#)

- read it: News media profiling using text analysis and social media context. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3364–3374, Online. Association for Computational Linguistics.
- Jason Baumgartner. 2019. *Twitter Tweets for Donald J. Trump (@realdonaldtrump)*. *Harvard Dataverse*.
- Felix Biessmann. 2016. Automating political bias prediction. *arXiv preprint arXiv:1608.02195*.
- Janez Brank, Gregor Leban, and Marko Grobelnik. 2017. Annotating documents with relevant wikipedia concepts. In *Proceedings of Slovenian KDD Conference on Data Mining and Data Warehouses (SiKDD)*.
- Sihao Chen, Daniel Khashabi, Wenpeng Yin, Chris Callison-Burch, and Dan Roth. 2019. *Seeing Things from a Different Angle: Discovering Diverse Perspectives about Claims*. In *Proc. of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Joshua Clinton, Simon Jackman, and Douglas Rivers. 2003. *The statistical analysis of roll call data*. *American Political Science Review - AMER POLIT SCI REV*, 98.
- Maryam Davoodi, Eric Waltenburg, and Dan Goldwasser. 2020. *Understanding the language of political agreement and disagreement in legislative texts*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5358–5368, Online. Association for Computational Linguistics.
- Dorottya Demszky, Nikhil Garg, Rob Voigt, James Zou, Matthew Gentzkow, Jesse Shapiro, and Dan Jurafsky. 2019. Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. *arXiv preprint arXiv:1904.01596*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Daniel Diermeier, Jean-François Godbout, Bei Yu, and Stefan Kaufmann. 2012. *Language and ideology in congress*. *British Journal of Political Science*, 42(1):31–55.
- Lisa Fan, Marshall White, Eva Sharma, Ruisi Su, Prafulla Kumar Choubey, Ruihong Huang, and Lu Wang. 2019. *In plain sight: Media bias through the lens of factual reporting*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6343–6349, Hong Kong, China. Association for Computational Linguistics.
- Wei Fang, Moin Nadeem, Mitra Mohtarami, and James Glass. 2019. *Neural multi-task learning for stance prediction*. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 13–19, Hong Kong, China. Association for Computational Linguistics.
- Sean M Gerrish and David M Blei. 2011. Predicting legislative roll calls from text. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256.
- Xiaochuang Han, Eunsol Choi, and Chenhao Tan. 2019. *No permanent Friends or enemies: Tracking relationships between nations from news*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1660–1676, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. *Feuding families and former Friends: Unsupervised learning for dynamic fictional relationships*. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544, San Diego, California. Association for Computational Linguistics.
- Kristen Johnson and Dan Goldwasser. 2016. *Identifying stance by analyzing political discourse on Twitter*. In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 66–75, Austin, Texas. Association for Computational Linguistics.
- Kristen Johnson and Dan Goldwasser. 2018. *Classification of moral foundations in microblog political discourse*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 720–730, Melbourne, Australia. Association for Computational Linguistics.
- Anastassia Kornilova, Daniel Argyle, and Vlad Eidelman. 2018a. *Party matters: Enhancing legislative embeddings with author attributes for vote prediction*. *CoRR*, abs/1805.08182.
- Anastassia Kornilova, Daniel Argyle, and Vladimir Eidelman. 2018b. *Party matters: Enhancing legislative embeddings with author attributes for vote prediction*. In *Proceedings of the 56th Annual Meet-*

- ing of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 510–515, Melbourne, Australia. Association for Computational Linguistics.
- Anastassia Kornilova, Daniel Argyle, and Vladimir Eidelman. 2018c. Party matters: Enhancing legislative embeddings with author attributes for vote prediction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 510–515.
- Peter Kraft, Hirsh Jain, and Alexander M Rush. 2016. An embedding model for predicting roll-call votes. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 2066–2070.
- Vivek Kulkarni, Junting Ye, Steven Skiena, and William Yang Wang. 2018. [Multi-view models for political ideology detection of news articles](#). *CoRR*, abs/1809.03485.
- Chang Li and Dan Goldwasser. 2019. [Encoding social information with graph convolutional networks for political perspective detection in news media](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2594–2604, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. [A dataset for detecting stance in tweets](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3945–3952, Portorož, Slovenia. European Language Resources Association (ELRA).
- Thomas Müller, Francesco Piccinno, Massimo Nicosia, Peter Shaw, and Yasemin Altun. 2019. Answering conversational questions on structured data without logical forms. *arXiv preprint arXiv:1908.11787*.
- Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1438–1448.
- Niko Palić, Juraj Vladika, Dominik Čubelić, Ivan Lovrenčić, Maja Buljan, and Jan Šnajder. 2019. [TakeLab at SemEval-2019 task 4: Hyperpartisan news detection](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 995–998, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Pallavi Patil, Kriti Myer, Ronak Zala, Arpit Singh, Sheshera Mysore, Andrew McCallum, Adrian Benton, and Amanda Stent. 2019. [Roll call vote prediction with knowledge augmented models](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 574–581, Hong Kong, China. Association for Computational Linguistics.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Keith T Poole, Howard Rosenthal, et al. 1997. *Congress: A Political-economic History of Roll Call Voting*. Oxford University Press on Demand.
- Daniel Preoțiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. [Beyond binary labels: Political ideology prediction of twitter users](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 729–740, Vancouver, Canada. Association for Computational Linguistics.
- Shamik Roy and Dan Goldwasser. 2020. [Weakly supervised learning of nuanced frames for analyzing polarization in news media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7698–7716, Online. Association for Computational Linguistics.
- Gregory Spell, Brian Guay, Sunshine Hillygus, and Lawrence Carin. 2020a. [An Embedding Model for Estimating Legislative Preferences from the Frequency and Sentiment of Tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 627–641, Online. Association for Computational Linguistics.
- Gregory Spell, Brian Guay, Sunshine Hillygus, and Lawrence Carin. 2020b. An embedding model for estimating legislative preferences from the frequency and sentiment of tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 627–641.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.

Appendices

A Event Examples

In this section, we provide examples of events that were identified by our event identification heuristic. For each automatically extracted event, we observe that the news headlines within the cluster usually describe the same real world event. The span of each event is 10 days at most. Hence, the assumption that the events within each issue are non-overlapping is a reasonable relaxation of reality. We made event segregated document data available for future research along with our code. Examples are shown in Tab 7.

B Reproducibility

We use seeds (set to 4056 for both tasks) for both random example generation and training neural networks. For fine-tuning layers of learning tasks we initialize the models using Xavier uniform (Glorot and Bengio, 2010) initialization with gain=1.0. We optimize the parameters using Stochastic Gradient Descent with an initial learning rate=0.0075 and momentum=0.4. We used 4 Nvidia GeForce GTX 1080 Ti GPUs with 12 GB memory and linux servers with 64 GB RAM for our experiments. CPU RAM and GPU memory are the main bottlenecks for training the model. It takes 80 hours to train authorship prediction for 5 epochs and 14 hours to train referenced entity prediction task for the same. Generating test results for both tasks together takes 3 hours. We use a batch size of 1 for both training and evaluation. For NRA Grade Prediction task we use 5 random seeds: {5, 7, 11, 13, 17} and report mean and standard deviation. The encoder-composer architecture is made up of 8.26M parameters, encoder consisting of 3.54M and composer 4.72M. Due to long

training time, the only hyper-parameter we experimented with is the graph size. We retained as many nodes as possible without exceeding GPU memory (500 nodes).

We divide the 3,640 queries into 151 batches of 24 queries each (3 politicians \times 8 issues) and 1 batch of 16 queries (2 politicians \times 8 issues). Train, val and test data examples are generated for each query batch. For Authorship Prediction and Referenced Entity Prediction tasks, Compositional Reader model is trained on one batch for 5 epochs, the best parameters are chosen according to the validation performance of that batch and we proceed to training on future batches. Politicians are ordered randomly when generating queries.

B.1 Data Collection

We collected data from 5 sources: Wikipedia, Twitter, ontheissues.org, allsides.com and ProPublica Congress API. We scraped articles from Wikipedia related to all the politicians in focus. We collected tweets from [Congress Tweets](#) and [Baumgartner \(2019\)](#). We used a set of hand build gold hashtags to separate them by issues. They are shown at the end of this document. We collected all news articles related to the 8 issues in focus from allsides.com. We collected press releases from Propublica API using key word search. We use issue names as keywords. We only maintain pointers to processed tweet and text data in data releases. All social media text analyzed is by public political figures, not private citizens.

C Evaluation Tasks

C.1 Grade Prediction Additional Ablation

For Grade Prediction task, we perform experiments by training the model on a fraction of the data. We monitor the validation and test performances with change in training data percentage. We observe that, in general, the gap between Compositional Reader model and the BERT baseline widens with increase in training data. It hints that our representation likely captures more relevant information for this task. Results are shown in figure 13.

C.2 Opinion Descriptor vs. RMN and LARN

[Han et al. \(2019\)](#) and [Iyyer et al. \(2016\)](#) both take a set of documents and entity pairs as inputs and generate relationship descriptors for the entity pairs in an unsupervised setting. They are both trained in an encoder-decoder style training process in an un-

Issue - Economic Policy	
Event #1: Donald Trump's Tax Proposal Release	Event #2: Obama's Economy Speech
Donald Trump to Propose Tax Breaks on 'Pocketbook' Issues in Economic Plan Trump's economic plan aims to please both corporations and working families Donald Trump Looks to Steady His Campaign With New Economic Speech Trump to outline economic plan in Detroit Clinton to dismiss Trump's economic plan as a 'friends and family discount' OPINION: Trump agenda looks like more of the same Clinton to dismiss Trump's economic plan as a 'friends and family discount' Trump tries to right his campaign, talking of tax cuts	Obama Economy Speech: Why The President's Plan Won't Get Past Republicans U.S. Is 'Through The Worst Of Yesterday's Winds,' Obama Says Obama Blames Five Years of a Bad Economy on "Phony Scandals" and "Distractions" Obama tries to offset current scandals by recycling talking points on economy Obama at Knox College: 'Washington has taken its eye off the ball' Obama Says Private Capital Should Take Lead Mortgage Role Why Obama might tap Summers for Fed despite harsh criticism from left Obama: Growing income inequality 'defining challenge' of this generation

Table 7: Examples of extracted events

Guns: #endgunviolence, #guncontrol, #gunviolence, #nra, #gunsafety, #assaultweaponsban, #gunsense, #marchforourlives, #parkland, #hr3435, #nationalwalkoutday, #disarmhate, #guncontrolnow, #backgroundchecks, #nationalschoolwalkout, #lasvegas, #elpaso, #keepamericassafe, #gunrights, #erpoact, #lasvegasshooting, #gunreform, #hr1112, #parklandstrong, #elpasostrong, #massshootings, #parklandstudentspeak, #hr8
Taxes: #GOPTaxScam, #TaxReform, #TaxAndJobsAct, #taxreform, #goptaxscam, #taxcutsandjobsact, #taxday, #taxcuts, #smallbusinessweek, #economy, #maga, #billionairesfirst, #gopbudget, #goptaxplan, #goptaxbill, #tax, #taxscam, #trumptax
Immigration: #FamiliesBelongTogether, #Immigration, #MuslimBan, #daca, #familiesbelongtogether, #dreamers, #immigration, #protectdreamers, #dreamactnow, #muslimban, #heretostay, #keepfamiliestogether, #protectthedream, #defenddaca, #immigrants, #familyseparation, #nomuslimbanever, #immigrant, #nobanowall, #borderwall, #refugeeswelcome, #endfamilydetention, #protectfamilies, #DACA, #refugees
Abortion: #ProChoice, #ProLife, #Abortion, #prolife, #abortion, #marchforlife, #prochoice, #theyfeelpain, #bornaliveact, #paincapable, #hr36, #roevwade, #unplanned, #defundpp, #life, #standwithnurses, #endinfanticide, #righttolife, #infanticide, #ppsellsbabyparts
LGBTQ Rights: #LGBTQ, #LGBT, #Homophobia, #lgbtq, #lgbt, #equalityact, #pridemonth, #hr5, #nationalcomingoutday, #lgbtqequalityday, #loveislove, #lgbthistorymonth, #transgender, #letkidslearn, #trans, #comingoutday, #marriageequality, #protecttranstroops, #lgbtqhistorymonth, #defundconversiontherapy, #transban, #otd, #prideinprogress, #nycpride, #protecttranskids, #transrightsarehumanrights, #transdayofremembrance, #loveisthelaw, #rfa, #bathroombill
Middle-East: #MiddleEast, #Iran, #Israel, #iran, #israel, #syria, #middleeast, #iraq, #russia, #northkorea, #irandea, #jordan, #hezbollah, #gaza, #isis, #hamas, #terror, #jihad, #violence, #barbarism, #palestinians, #jewish, #antisemitism, #saudiarabia, #iranian, #lebanon, #turkey, #jerusalem, #iranprotests, #israeli, #freeiran, #sanctions, #supportisrael, #egypt, #terrorism
Environment: #ActOnClimate, #ClimateChange, #GreenNewDeal, #climatechange, #actonclimate, #greennewdeal, #climateactionnow, #parisagreement, #climatecrisis, #earthday, #climatefriday, #climate, #climatestrike, #climateaction, #cleanenergy, #climatechangeisreal, #environment, #oceanclimateaction, #cleanair, #climatechangeimpactsme, #cleanwater, #globalwarming, #renewableenergy, #worldenvironmentday, #climateemergency, #peopleoverpolluters, #greenjobs, #climatejustice, #solar, #environmentaljustice, #cleanpowerplan, #todayclimatefact, #sealevelrise, #bigoil, #climatecrisiscountdown, #stopextinction, #cleancars, #climatecosts, #cutmethane, #chamberofcarbon, #climatesolutions, #amazonrainforest, #hurricanemaria, #climatesecurityisnationalsecurity, #protectcleanwater, #renewables, #offossilfuels, #columbiaenergyexchange, #climatesolutionscaucus

Table 8: Gold Hashtag Set used to collect Politicians' Tweets for Issues

supervised manner. Given new text with an entity pair, they generate d descriptor embeddings that are used to rank candidate descriptors. Iyyer et al. (2016) uses entire vocabulary space while Han et al. (2019) uses 500 most frequent verbs.

In contrast, our model doesn't need the presence of both the entities in text to generate opinion descriptors. This often tends to be the case in tweets and press releases as they are generated directly by the author (first-person discourse). Our model is also capable of summarizing over multiple documents and generating descriptors for several referenced entities and issues at once while they deal with one entity-pair at a time.

D Additional Visualizations




Figure 5: Reps vs Dems on issue *immigration*




Figure 6: Reps vs Dems on issue *taxes*




Figure 7: Reps vs Dems on issue *middle-east*




Figure 8: Reps vs Dems on issue *abortion*




Figure 9: Reps vs Dems on issue *economic-policy*




Figure 10: Reps vs Dems on issue *environment*




Figure 11: Reps vs Dems on issue *LGBTQ Rights*




Figure 12: Normalized Agreement Score of Politicians with statement: *Believes in common-sense approach to reforming gun control.*

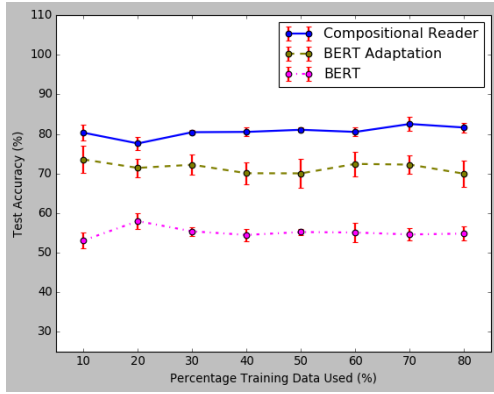


Figure 13: NRA Grade Prediction: Data % vs Test Acc




Figure 14: NRA Grade Prediction: Data % vs Val Acc

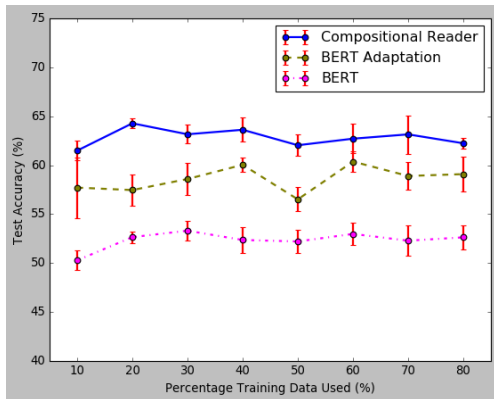


Figure 15: LCV Grade Prediction: Data % vs Test Acc




Figure 16: LCV Grade Prediction: Data % vs Val Acc




Figure 17: Comparison of Politician Stances on Issues

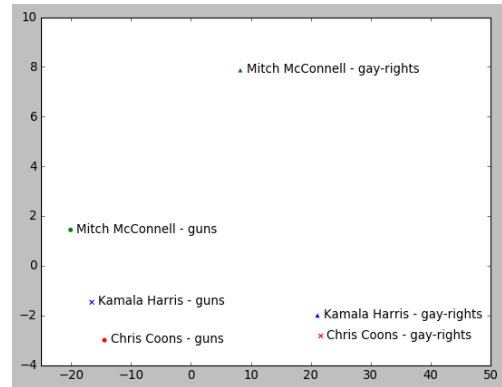


Figure 18: Comparison of Politician Stances on Issues

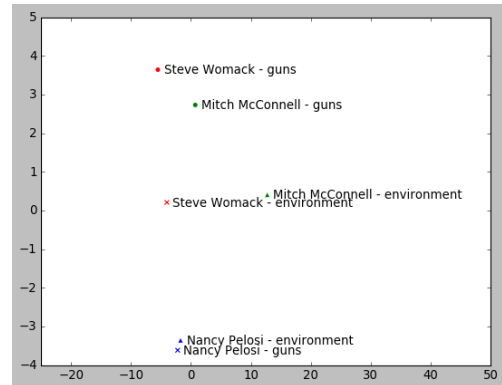


Figure 19: Comparison of Politician Stances on Issues