



# Fast Few Shot Self-attentive Semi-supervised Political Inclination Prediction

Souvic Chakraborty<sup>(✉)</sup>, Pawan Goyal, and Animesh Mukherjee

Indian Institute of Technology, Kharagpur, West Bengal, India  
chakra.souvic@gmail.com

**Abstract.** With the rising participation of the common mass in social media, it is increasingly common now for policymakers/journalists to create online polls on social media to understand the political leanings of people in specific locations. The caveat here is that only influential people can make such an online polling and reach out at a mass scale. Further, in such cases, the distribution of voters is not controllable and may be, in fact, biased. On the other hand, if we can interpret the publicly available data over social media to probe the political inclination of users, we will be able to have controllable insights about the survey population, keep the cost of survey low and also collect publicly available data without involving the concerned persons. Hence we introduce a self-attentive semi-supervised framework for political inclination detection to further that objective. The advantage of our model is that it neither needs huge training data nor does it need to store social network parameters. Nevertheless, it achieves an accuracy of 93.7% with no annotated data; further, with only a few annotated examples per class it achieves competitive performance. We found that the model is highly efficient even in resource-constrained settings, and insights drawn from its predictions match the manual survey outcomes when applied to diverse real-life scenarios.

## 1 Introduction

Political inclination refers to the political stance of an individual. Polling and surveying to understand the political leanings of people within a particular community, in a particular geopolitical region, or a specific context is a common approach. However, the manual polling mechanisms used today are hard to scale. Also, there is a significant chance of biased sampling as the samples are often too small in terms of the number of individuals surveyed and localized. On the other hand, if a survey or poll is conducted on online social platforms, it is impossible to control voters' distribution to calibrate it to resemble a random sample of opinions. Often the voters in these polls are limited to being the active audience of the pollsters sharing similar political inclination. Thus the result of the same poll can be completely different if introduced by a different pollster. Therefore, algorithmic labeling of people chosen from a controllable distribution is important, rather than asking for bias-prone active participation by influencers sharing particular political inclination or cost-inefficient manual polling.

Most of the existing approaches of political inclination detection (PID) on social networks focus on probabilistic models [8,9,13,14], which are in turn based on the texts generated by users. Researchers have also tried to exploit the network structure by making use of GCNs [24] (Graph Convolutional Networks). This method uses all second-degree features (neighbors of neighbors of the node/user to be classified in the graph) for a rich representation which makes the classification more accurate at the expense of speed. The data collection process involves collecting features of followers of followers of the user whose political inclination needs to be detected. As the collection of followers and all their tweets itself is a slow process limited by Twitter<sup>1</sup>, the time required for the collection of the features of the second-degree neighbours increases quadratically in terms of average unique neighbors per node.

Also, the GCN-based models need to store the huge Twitter-subnetwork involving political persons and their followers. This severely violates the users’ *right to erasure* as per the Article 17 of GDPR (See footnote 2) which reads as follows:

**“The data subject shall have the right to obtain from the controller the erasure of personal data concerning him or her without undue delay and the controller shall have the obligation to erase personal data without undue delay ...”**

Further, these models are trained on a huge number of annotated examples. This makes the approaches hard to scale for newer settings and countries. In contrast, we show that certain *easy to collect features* plugged into a novel self-attentive framework can be very accurate in predicting the political inclination even if trained on a handful of annotated examples.

Our main contributions are as follows:

- (1). Graph-based methods used previously raise many ethical questions [15, 17, 21, 23]. The users on social media platforms have the right<sup>2</sup> to deletion of their data from other storage systems which are dependent on social media as data source whenever their public profile on social media is deleted. Graph-based methods violate this by storing their information such as retweets, mentions, likes, and the follower-followee network. The time required to build and update such networks is huge as it will require everyday monitoring for (i) the existence of each connection and (ii) the arrival of new connections. So, the only way to use these features at inference time is to store them permanently in memory. We eliminate the need for storing such large relational graphs from past social media data of a huge number of users. We achieve this by using richer first-degree features that we collect directly at inference time along with their second degree neighbors which can be collected from the tweets of the concerned user/person to be classified directly (e.g., we collect the hashtags used by the retweeted user as it is readily available with the retweet, same for

<sup>1</sup> <https://developer.twitter.com/en/docs/twitter-api/v1/rate-limits>.

<sup>2</sup> <https://gdpr-info.eu/art-17-gdpr/>.

replies). Using smart augmentation of these features we beat the performance obtained by the graph-based approaches [24] at a reduced inference time.

(2). We propose a novel Fast Self-attentive Semi-supervised Political Inclination Predictor *FSSPIP* (Fig. 1). The experimental results show that even without using any gold annotation, we can achieve high accuracy of  $\sim 94\%$  using weak supervision. The model is highly scalable and free from manual intervention unlike Darwish et al. (2020) [10] which needs human supervision or cluster inspection.

(3). We bring on board multiple additional datasets to show that our model can be used in many other similar settings for political inclination detection with a handful of labeled examples (or even without it). In specific, we present several case studies on media bias and political polarization using our classifier in zero-shot settings.

## 2 Related Work

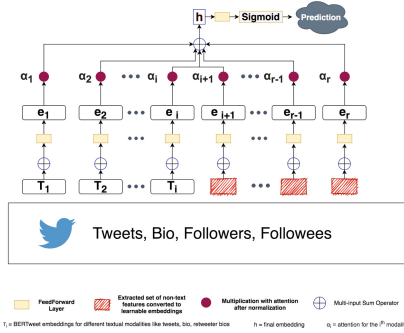
Stefanov et al. (2020) [20] and Baly et al. (2020) [2] used Wikipedia, Twitter, YouTube, and other channels of information to detect the political leanings of the media houses. This approach is not scalable in context of persons. Conover et al. (2011) [9] had used a corpus of 1000 annotated data points to test the supervised approaches based on bag of words. Iyyer et al. (2014) [13] used advanced neural techniques like RNNs on a labeled corpus of sentences taken from speeches of democratic and republican parliamentarians. Chen et al. (2017) [8] used graph-based approaches to show the efficacy of using an opinion-aware knowledge graph. However, these techniques fail to take the richer network features into account. They also completely rely on annotated data failing to take advantage of domain knowledge of the task in hand.

Aldayel et al. (2019) [1] analyzed the features responsible for higher accuracy in stance detection setups using network features, tweet texts and text derived features. Darwish et al. (2020) [10] on the other hand used clustering based unsupervised setup to detect stance of users mainly relying on three channels of features: retweeted tweets, retweeted accounts and hashtags. Xiao et al. (2020) [24] approached the same task using manual annotation and collecting a large dataset of non-politician social media users and politicians on Twitter. They relied on variants of relational GNNs coupled with multi-task learning. However, given the need for explicit storage of information in the graph structures, even after the training phase, the graph-based algorithms often violate privacy rights of a large section of users.

Therefore, in this paper, we attempt to solve political inclination detection in a resource-constrained setup with no storage of user data after model training. We use several task-dependent augmentation techniques and unsupervised learning methods which have not been used in this context earlier thus making our model robust, easily adaptable and scalable without any human help/supervision. We only use public data available at the time of inference.

### 3 Model Architecture

**The Base Architecture:** Like previous state-of-the-art approaches [24] we too use a *GCN-like framework*. However, we, in contrast, do not store the user/feature graphs nor do we need a list of politicians in the country of the users to be classified with huge set of labelled examples. We use a long list of different *feature types* derived from *follow, mention, reply, retweet, tweets* and *likes*. We hypothesize that for a good representation of the political inclination, there are many important but easy to collect features which can be retrieved from the web directly during inference time with no need of storage. We describe these features in details below:



**Fig. 1.** The FSSPIP base architecture: The Twitter profile of a user is taken as input from which 22 different feature types are extracted and processed to predict political inclination.

#### Base Features

*User Descriptions:* We collected user descriptions of users retweeted and quoted, forming two separate documents. These user descriptions/bios often contain key information like the user’s occupation, gender, religion etc.

*Hashtags:* Hashtags are important as similar hashtags are used to express opinions for/against a polarizing topic by users of different leanings.

*Mentions:* IDs mentioned in tweets are used as features.

*Media Domains:* It is no secret that users of different political leanings share different sets of news items that fit their ideological perspective. Considering their importance in our task, we collect domain names and domain + co-domain names from users’ tweets. We use them as separate features.

*Textual Content:* We use pre-trained models like BERTweet [18], and Google’s Universal Sentence Encoder [4] to convert the content of tweets of a user into embeddings. In our experiments, we found that BERTweet performs better (possibly because BERTweet is trained on text with vocabulary more similar to ours). Thus we report BERTweet numbers only.

*Features Connected to Neighborhood:*

We mentioned a total of 6 features till now. We repeat the same features for retweets and replies separately. So, the total number of features become  $6+6+6=18$ .

In addition to all these features, we also use *friend ids*, *follower ids*, *mention ids*, *ids replied to*, and *ids retweeted* as features collected at test time.

So, in total we use  $18+4=22$  features.

### Attending to Different Modalities

We use  $|R| = 22$  features in our architecture. For each feature type  $r$  and user  $i$ , we obtain an embedding  $e_{ir}$  of size  $d = 8$  as follows.

$$\begin{aligned} e_{ir} &= W_r \times BERTweet(T_{ir}), \text{ if } r \in T \\ &= W_r A_{ir} H_r, \text{ if } r \in T' \end{aligned} \quad (1)$$

where  $A_{ir} \in \{0, 1\}^{1 \times Vlen_r}$  is the feature presence-absence vector for the  $r^{\text{th}}$  feature and  $H_r \in \mathbb{R}^{Vlen_r \times d}$  is the embedding matrix containing embeddings of all the features for feature type  $r$ . While pre-processing we chose only those features in the vocabulary which appeared at least in five instances of the training data to ensure having enough training instances. The length of the vocabulary for  $r^{\text{th}}$  feature type is represented as  $Vlen_r$ .  $T'$  and  $T$  are respectively the set of non-textual and textual feature types. For each feature type  $r$ ,  $T_r$  denotes the textual content of that feature for user  $i$ ,  $W_r \in \mathbb{R}^{d \times d_{em}}$ , where  $d_{em}$  is the embedding dimension of the output of BERTweet [18]. Now, we calculate  $h_i$ , the final embedding for the  $i^{\text{th}}$  user as follows.

$$h_i = \sum_r \alpha_{ir} \times \frac{e_{ir}}{|e_{ir}|} \quad (2)$$

FSSPIP uses a dynamic dot-product self-attention mechanism to calculate the weights for each of the feature types to finally get a weighted sum of the normalized embeddings of each feature type. We use learnable parameters  $p, q, k \in [0, 1]$  to allow some flexibility in attention calculation. Learnable parameters  $q_r$  and  $k_r \in \mathbb{R}^d$  are queries and keys, respectively, for each feature type  $r$  (Here a feature type is specific social media attribute, so a collection of hashtags coming from tweets is a feature type different from the collection of hashtags coming from the retweets/replies. Please refer to the list of features mentioned at the start of the section for a broader understanding). So,

$$\alpha_{ir} = p * \frac{e^{q_{ir} \times k_{ir}}}{\sum_r e^{q_{ir} \times k_{ir}}} + (1 - p) * |e_{ir}| \quad (3)$$

$$q_{ir} = q \times e_{ir} + (1 - q) \times q_r \quad (4)$$

$$k_{ir} = k \times e_{ir} + (1 - k) \times k_r \quad (5)$$

An illustration of this base architecture is presented in Fig. 1. It shows how input from each feature type goes through different transformation functions (BERTweet in case of textual data, trainable embeddings in case of follower

ids etc.) to transform into embeddings which are then weighted by attention values calculated through three different architectural scheme as mentioned. The weighted summation of the embeddings (vector size 768) denote representation of the node/person to be classified. This embedding is further multiplied with a vector of size  $768 \times 1$  and passed through a sigmoid function to obtain probability of a person being a republican in a binary classification setup. We use binary cross entropy as the loss function for supervision.

## 4 Augmented Semi-supervision for Superior Representation Learning

To make our architecture ready for few-shot learning, we make the model robust using regularization and multi task learning. We also use weak supervision producing high accuracy without any labelled example. Specifically, we use three different categories of techniques which are described below.

### Dynamic Augmentation

*Mixup*: Mixup [25] is a technique that enforces linear change in output given linear change in input by training a neural network on convex combinations of pairs of examples and their annotated labels for a particular task. We adopt the method to our network data by mixing two random users for each channel (e.g. : hashtags, domains, retweeters from both users are present for the augmented user) increasing the diversity of data points and regularizing the model for unseen data points.

*Sampling*: Twitter users can be imagined as generative agents who generate tweets on selective issues and follow/reply/mention/interact with other users following some implicit probability distribution. So, if some of the points from the distribution are sampled out uniformly, the distribution will not change. So, we uniformly(chosen from a random uniform distribution for each feature type) sample out features from labeled examples for augmentation masking out 0–15% of the features randomly during training.

*Feature Channel Dropout*: While some feature types may influence the result more than others, it is important to learn to predict from the cues available if one influential feature type (e.g. *hashtags*, *followers*, *retweets* etc.) is absent. So, we randomly drop random feature types while training for better performance through adversarial training.

**Weak Supervision**: We hypothesize that the followers/retweeters of a particular political party often share the bias of having that particular political inclination. So, they are statistically more likely to follow the leaning of that particular political party which they are following on social media than any other. This provides some silver labels in the Twitter space for weak supervision. We crawled the Twitter handles of each political party (i.e., the official Twitter handles of *The Democratic & The Republican* party in case of US and *AAP, Congress & BJP* party in case of India) to collect the last 75,000 (set heuristically to contain enough examples) followers and the last 75,000 retweeters for each of these

parties. We randomly selected 2,500 from each of them to get a sample representative of the timeline (as the most recent followers appear first, so collecting a big pool and resampling may help) and collected their relevant data for training. Users following both parties were removed.

**Self-supervision:** Self-supervision is a semi-supervised learning technique that trains the model in a new dummy task predicting part of input data using the other part of the data [18]. While in case of textual data masked language modelling and next sentence predictions [11, 18] are the most frequently used pre-training technique, graph neural nets use predicting masked edges between nodes as the pre-training task. Following these methods, we pretrain our model with the task of prediction of the non-textual features which are masked in the dynamic augmentation phase during sampling the input features. We use self-supervision as pre-training method while performing few-shot learning and fine-tune later on the annotated data points. Hyperparameter details and loss function of the pretraining phase has been put in the Appendix available at <https://tinyurl.com/icadlappendix>).

## 5 Data Preparation

**Dataset for the Main Task:** As provided by Xiao et al. (2020) [24], we have 2,976 labeled data points (labelled *republican* or *democratic*) along with 583 politicians’ data in the US setting. For a nuanced analysis, we retain the partition of the data points, used in the dataset – **PureP**<sup>3</sup>, **P50**<sup>4</sup>, **P20~50**<sup>5</sup> and **P+all**<sup>6</sup>.

**Table 1.** Descriptive statistics of the labeled dataset.

	PureP	P50	P20-50	P+all
#User	579	730	946	1,125
#Retweetees	26,034	46,734	67,239	87,802
#Replies	1,738	16,261	12,665	29,713
#Mentions	30,363	53,125	62,274	84,962

We collected the Twitter ids and labels provided by Xiao et al. (2020) [24]. We crawled the last 3,200 tweets (some tweets got deleted, some tweets were retweets, quotes, and replies), follower ids and friend ids of each labeled id in

<sup>3</sup> This dataset contains only the politicians.

<sup>4</sup> This dataset contains people highly interested in politics being followed or following at least 50 politicians, including the politicians themselves.

<sup>5</sup> This dataset contains people moderately interested in politics being followed or following anywhere between 20 to 50 politicians and the politicians themselves.

<sup>6</sup> This dataset contains members of **PureP**, **P20-50** along with many outliers who are following or being followed by maximum five politicians.

November, 2020 using the Twitter API<sup>7</sup>. We also collected the user objects (containing bios) for each id. So, after pre-processing, we have data for each feature type described in the previous section. We extracted the domain and co-domain names from the URLs shared using the tldextract<sup>8</sup> library. Out of 2,976 labeled users, 2,665 users were available on Twitter at the time of crawling of the tweets (November, 2020). We report our results on this dataset. A major point to be noted here is that we do not store this data once the training is over, nor do we need to collect neighborhood data at inference time making the inference faster and memory efficient. A detailed statistics of this dataset with the count of unique features for some feature types is provided in Table 1.

### Additional Datasets for Lateral Verification

We collect several other datasets to demonstrate the usefulness of FSSPIP in zero/few-shot setting. The statistics of these datasets are detailed in Table 2.

**Table 2.** Descriptive statistics of the collected datasets. MB: MediaBias; C: Community; MP: Multiparty; S: Statewise TPC: Topicwise; HTU: HashTagUsers (4 hashtags subset as mentioned in Fig. 3b , details in Appendix).

	MB	C	MP	S	TPC	HTU
#User	806	400	1000	2900	4000	4000
#Domains	5,629	3,829	6,651	17,285	25,679	27,287
#Hashtags	83,245	18,109	28,891	69,341	71,189	68,423
#Mentions	52,178	96,123	67,478	89,651	89,765	81,319

**The Media Bias Dataset:** Following Stefanov et al. (2020) [20], we use the crowdsourced labels<sup>9</sup> for media bias prediction. There are 806 labeled instances in the dataset with labels *left*, *center-left*, *least biased*, *center-right* and *right*. In order to binarize the label space (to fit in our classification model which is a binary classifier), we first discard the instances with label *least biased*; next, we merge *left*, *center-left* to a single label *left* and *center-right*, *right* to a single label *right*. We collect the friend ids, follower ids, and the last 3,200 tweets of these media houses to employ the FSSPIP classifier for prediction.

**The Ethnic Community Dataset:** Many post-poll surveys establish how different communities vote differently. We try to use our model to identify such divisions. We first sample recent tweets using the Twitter API (See footnote 7) mentioning names of any of the communities. Among the users tweeting, we

<sup>7</sup> <https://developer.twitter.com/en/docs/twitter-api>.

<sup>8</sup> <https://pypi.org/project/tldextract/>.

<sup>9</sup> <https://mediabiasfactcheck.com/>.



select only those who mention their community as one (or more) of the communities/ethnicities being probed ('black', 'white', 'hispanic/latino', 'asian'), in their bio. We put a user to a particular community if that community is mentioned in their bio.

**Multi-party Leaning Dataset:** In order to collect a set of users residing in a multi-party democratic system, we filter the latest 10,000 tweets (and tweeters) containing the term 'Delhi Election' using Twitter API on 17<sup>th</sup> March, 2021. We annotate random 1000 Twitter users from this list into followers of three political parties: *AAP: Aam Aadmi Party*, *BJP: Bhartiya Janta Party* and *Congress/INC: Indian National Congress* (AAP: 203 users, INC: 435 users, BJP: 362 users) to form the multi-party inclination dataset. We check the residency of the users and confirm it to be India through the self-declared location tag in twitter while annotating each user.

**Statewise Inclination Dataset:** Here, we use the Twitter API to collect tweets against a politically neutral query 'election' (all datapoints are collected before 17.05.2021). If the user has a state's name mentioned in the *Twitter's location tag*, we categorise that user to that particular state. We collected 100 users for each state in India for representative sample collection.

**Hashtags User Data:** In order to find out the inclination distribution behind each hashtag, we collect the tweets containing some trending hashtags (on or before 17.05.2021). We collect 1000 ( $\times 30$ ) tweets, excluding the retweets and replies containing each trending hashtags using the Twitter API. For a manual verification, we annotate 30 hashtags with tags Congress, BJP, and Neutral (In the date of collection, we could not find hashtags which can be attributed to be inclined towards AAP. Moreover, politically unmotivated hashtags are termed as 'neutral'). This annotation was done by a PhD student, expert in Indian Politics by reading the tweets with the hashtags.

## 6 Main Task: Experiments and Analysis

**Baselines:** We use the best performing models provided by Aldayel et al. (2019) [1] and Darwish et al. (2020) [10] (UMAP+DBSCAN with tweets containing chosen hashtags included in the Appendix). **NTF** [1] uses network/graph and textual features together in its model just like our model without attention. **UUS** [10] on the other hand uses weak supervision (a quite different method compared to ours) through dimensionality reduction and clustering, manual inspection (which also makes the algorithm less scalable) and labelling of the clusters with only three features (retweeted tweets, retweeted accounts and hashtags). We also added a modified version of the UUS algorithm for a fair comparison with our fine-tuned model as the UUS algorithm is completely unsupervised and incapable of using any supervisory signal for few-shot learning. We took the unsupervised UUS model and fine-tuned it using annotated data points, terming it **UUS+**.

We also add non-neural baselines like SVM, Logistic regression (LR) and Random Forest (RF) as we are interested to show how simple algorithms with smaller inference time tally with our methods. Here, we use the concatenation of tweets for each user as input. We used TIMME-hierarchical [24] and its other two variants as the other baselines using self-supervision on graphs with higher inference time due to second order data collection on large graph. However, we only report TIMME-hier results as it was the best-performing variant (hyperparameter stats and details on other TIMME variants in Appendix). A qualitative comparison of the baselines is added in Table 3.

**Table 3.** A pointwise comparison of the models used as baselines. {NNeur : Non Neural baselines}.

Models	FSSPIP	TIMME	NTF	UUS+	UUS	NNeur
Uses only neighbors?	✓		✓	✓	✓	✓
Uses textual features?	✓	✓	✓	✓	✓	✓
Uses unsupervised pretraining also while finetuning on labelled examples?	✓	✓		✓		
Performs well even without manual annotation?	✓			✓	✓	
Performs well even without any human supervision ?	✓					

**Results:** In Table 4, we show that our best performing model *FSSPIP*<sup>10</sup> fairly beats other baselines for all datasets. We see that we gain most compared to other models when very few training datapoints (50) are present<sup>11</sup>. In the case of the non-politician datasets, i.e., **P50**, **P20–50** and **P+all**, the performance obtained by our model is significantly higher than other baselines, even with only 50 training data points. This may be because the non-politician datasets do not purely contain political features unlike the PureP dataset making the feature learning task less straight-forward needing finer features like domain names a user is interested in or the tweets from the retweetees.

Our model performs better than other models in terms of time required to predict for a single user. Compared to the networks using second order relational data (TIMME) we are at least  $\sim 10x$  faster as shown in Table 4.

<sup>10</sup> † represents p value less than 0.05 in student’s t-test while comparing FSSPIP’s result with the best performing baseline.

<sup>11</sup> Training on the whole data, we got comparable accuracy with other state-of-the-art architectures (result in Appendix).

**Table 4.** Results of few-shot learning {NTF: Model proposed by [1]; UUS: Model proposed by [10]; UUS+: Model proposed by [10] fine-tuned on annotated data points; TIMME: TIMME-hier (other TIMME variants’ result in Appendix); **FSSPIP: FSSPIP base architecture with the few-shot learning framework**; #T: Number of training datapoints; TTI: Time Taken for Inference per datapoint with Twitter ids as inputs. For each framework, it also includes the time taken to collect the data}.

Dataset	#T	LR	RF	SVM	NTF	UUS	UUS+	TIMME	FSSPIP
PureP	50	81.8	90.1	79.2	89.6	94.1	95.2	95.4	<b>97.1</b> <sup>†</sup>
	250	85.1	93.2	83.8	93.2	94.1	97.5	95.8	<b>98.9</b> <sup>†</sup>
	500	89.2	96.9	88.5	97.9	94.1	98.1	99.2	<b>99.2</b>
P50	50	78.2	81.7	64.2	90.9	92.2	95.5	84.6	<b>97.3</b> <sup>†</sup>
	250	79.2	83.8	65.4	92.3	92.2	96.8	95.1	<b>98.1</b> <sup>†</sup>
	500	80.4	84.1	69.3	97.8	92.2	97.5	95.2	<b>98.9</b> <sup>†</sup>
P20-50	50	86.1	86.6	81.2	90.2	92.5	94.9	79.1	<b>96.9</b> <sup>†</sup>
	250	89.6	87.3	84.8	91.8	92.5	95.8	96.9	<b>97.8</b> <sup>†</sup>
	500	92.1	89.9	86.1	97.7	92.5	97.1	97.1	<b>98.9</b>
P+all	50	84.2	77.1	76.2	91.5	89.1	92.8	69.5	<b>95.2</b> <sup>†</sup>
	250	86.5	79.5	77.9	93.9	89.1	94.5	93.8	<b>97.1</b> <sup>†</sup>
	500	87.6	83.2	80.8	96.7	89.1	97.4	95.2	<b>98.1</b> <sup>†</sup>
TTI (secs)		28.0	28.9	28.1	65.3	63.1	63.2	973.2	62.1

**Table 5.** Ablation study of different model variants {F1: FSSPIP-fixedattn; F2: FSSPIP-auto; FSSPIP- - -: FSSPIP base architecture; FSSPIP- -: FSSPIP without weak supervision and self supervision; FSSPIP-: FSSPIP without self supervision.}

Dataset	#T	F1	F2	FSSPIP- - -	FSSPIP- -	FSSPIP-	FSSPIP
PureP	50	90.4	91.3	91.6	92.3	94.2	97.1
	250	94.1	94.7	94.8	95.1	98.1	98.9
	500	98.9	98.7	99.0 <sup>†</sup>	99.1	99.1	99.2
P50	50	87.2	89.4	90.9 <sup>†</sup>	92.1	94.9	97.3
	250	92.1	91.5	92.7 <sup>†</sup>	93.5	97.3	98.1
	500	95.6	96.8	98.4 <sup>†</sup>	98.9	98.9	98.9
P20-50	50	87.1	87.6	90.2 <sup>†</sup>	92.8	93.8	96.9
	250	91.6	92.4	92.8	95.2	97.2	97.8
	500	97.8	98.2	98.4	98.5	98.5	98.9
P+all	50	88.3	88.5	89.2 <sup>†</sup>	92.5	93.0	95.2
	250	92.1	93.8	94.3 <sup>†</sup>	95.8	96.9	97.1
	500	95.9	96.3	97.8 <sup>†</sup>	97.8	97.9	98.1

Also, our model performs better than NTF [1] and UUS [10] by a significant margin using weak supervision<sup>12</sup> with better augmentation while utilizing carefully extracted network features like NTF inputs [1].

## Ablation Study

*Model Variants* - In order to ablate our attention mechanism we employ two other varieties of attention in place of ours in FSSPIP base architecture which are as follows: We recall Eq. 2 here to understand the two new attention mechanisms: *FSSPIP-fixedattn* (F1): FSSPIP-fixedattn uses fixed learnable attention to calculate a weighted sum of embeddings of each feature type. Thus here the Eq. 2  $\alpha_r$  values are learnable parameters and  $\alpha_{ir} = \alpha_r, \forall i$ . *FSSPIP-auto* (F2): FSSPIP-auto simply sums up each of the normalized embeddings of each feature type, assuming equal attention to all the feature types while computing the final embedding vector. So, here we assume  $\alpha_{ir} = 1 \forall i, r$ .

**Table 6.** Important features and feature types for the predictions.

Feature type	Example features	% drop for one feature type
Hashtags	#taxreform, #maga, #medicareforall, #healthcare, #txlege, #tcot, #coleg, #gopwomen	2.7%
Domains	senate, clk, fairandsimple, house, theguardian, congress	0.9%
Followees	JimLangevin, txstdems, COHouseDem, Sam1963, MustafaTameez, pdamerica	0.3%
Retweetees	RepJoseSerrano, PuestoLoco, MustafaTameez, SenatorMenendez, ProjectLincoln, JoeBiden	0.5%
Replies	replblumenauer, RepJoseSerrano, SenatorMenendez, Sam1963	0.3%
Mentions	kylegriffin1, MustafaTameez, WhiteHouse45, texasdemocrats, JoeBiden, Mike.Pence	0.5%

To test the few-shot learning framework we used incrementally powerful models in Table 5, where FSSPIP- - - is the base architecture without the few-shot learning framework and then each component of the framework is added sequentially to the base model (terming those intermediate models FSSPIP- -, FSSPIP-, and finally FSSPIP).

We find that the dynamic attention mechanism produces significantly<sup>13</sup> higher gains compared to the other two attention variants. We find that the

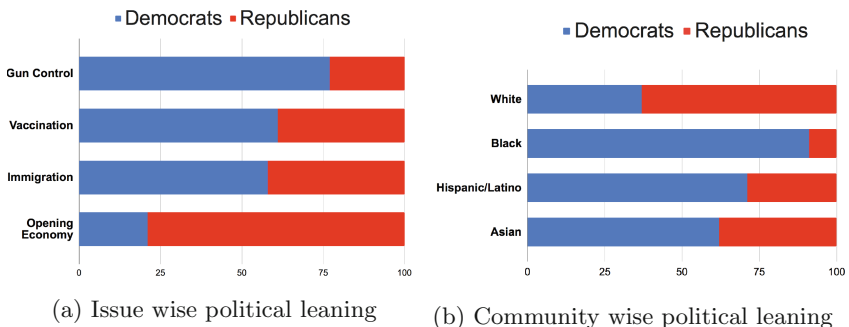
<sup>12</sup> We do not use any human intervention, unlike other approaches which cluster the datapoints and identify the cluster’s political affiliation by manually sampling users and annotating them. This may also be subjected to randomness in the clustering process and dependent on the characteristics of the specific subsets of the social network.

<sup>13</sup> We ran a significance test comparing the results of other two variants with the main (dynamic) variant and found the  $p$ -value to be less than 0.05 in all the cases (compared to both F1 and F2) as marked by † in Table 5.

gains produced are higher when fewer data points are used and weak supervision has a higher impact than adding dynamic augmentation further to the weakly supervised model. This can be explained as weak supervision already trains the model with a large number of real data points which makes the model regularized enough. However, dynamic augmentation helps in regularizing the model, specially in fewshot settings to avoid over-fitting. Similarly, self-supervision also seems more useful in case of fewer training data points. Moreover, we can see that the attention variants of the model perform very close to the original model but falls short with low number of data points.

*Most Important Feature Types* - To determine the most important feature types, we drop each feature channel and measure the information loss by calculating the deviation in performance of the classifier (FSSPIP) trained on the combined dataset (train:test:validation datapoints = 80:10:10). The results are reported in Table 6. The highest drop is witnessed when the relevant hashtags are dropped.

**Zero-Shot Gain:** Inspired by the significant improvement by weak supervision as shown in Table 5, we trained our model *FSSPIP* on the weak supervision dataset *only*, which is collected without any manual annotation. We then use the whole annotated dataset for testing this model. We obtain a **zero-shot accuracy** of 93.7% (TIMME models are based on list of politicians of each party, and thus cannot be zero-shot. UUS, which is not easily scalable due to its clustering, purity checking by experts and soft labelling methodology, performed the best among other baselines at 91.9%). This tells that the social media followers of a political party are indeed, most of the time, followers of the party in real life also. So, training a model to classify a social media user to be a follower of one party over the other on social media also trains the model for the similar task of classifying the user to be a follower of one political party over the other in real life. We verify this conclusion again in a multi-party scenario for a diverse non-English speaking democracy like India in the next section.



**Fig. 2.** Distribution of political inclinations in the USA by topic/racial demographics.

## 7 Additional Task: Experiments and Analysis

We use the additionally collected datasets to show the efficacy of the zero-shot classifier. The research questions selected for this section are easy to test but important for social scientists. They had been mostly analyzed through manual surveys till now.

**Media Bias Prediction:** We use the trained FSSPIP classifier on the media bias dataset collected by us, taking each of the Twitter handles of the media houses as the node to be classified. We obtain an accuracy of 72.6% on the task, while we do not explicitly train for this task<sup>14</sup> and rely on the assumption that {democrat  $\equiv$  left} and {republic  $\equiv$  right}.

**Topical Polarization - Bone(s) of Contention:** In order to poll users for specific contexts and issues, we collect some hashtags (see Appendix) supporting each issue mentioned in Fig. 2a. We then use the model to classify each user and plot the % of users for each leaning in the US setting, i.e., *The Democrats* & *The Republicans*.

**Multi-party Inclination Prediction:** US political system is binary consisting of only two political parties: *The Democrats* & *The Republicans*. In principle, our system can work for other countries and other kinds of political systems as well. In this section, we test *zero-shot* property classification of our model on the diverse multi-party democracy like India. We take Twitter handles of three national parties in India, namely, *Aam Aadmi Party (AAP)*, *Indian National Congress (INC)* and *Bharatiya Janata Party (BJP)*. We use the weak supervision method to train our model with sampling, mixup & feature channel dropout strategy as discussed earlier. On a random sample of 1000 Twitter accounts (*AAP*: 203, *INC*: 435, *BJP*: 362), we obtain an accuracy of 81.9%. The highest confusion scores between classes (see Appendix) were between *AAP* & *INC*. This is fairly intuitive since both these parties are left-leaning and in opposition, while *BJP* is known to be subscribing to a right-wing leaning and is currently in power.

**Statewise Leaning:** In Fig. 3a, we plot the relative distribution of political leanings for each state of India (on a scale of 0–1, signifying the percentage of users in a state leaning toward *BJP*. We get the average of political leanings for each person in the state’s data predicted using the aforementioned classifier). This correlates quite well (Pearson’s corr coef: 0.52 with high significance and low p-value,  $p < 0.01$ ) with the vote percentage received by *BJP* in each state in the 2019 general election.<sup>15</sup>

**A Leopard Cannot Change Its Spots:** In order to check if the political inclination changes with time, we reuse the same dataset described in the last paragraph with a temporal filtering strategy. We only use the tweets and tweet

<sup>14</sup> If we train for the task explicitly using 70:30 split for train-test data, we achieve an accuracy of 84.1%.

<sup>15</sup> <https://eci.gov.in/>.

derived features for this experiment which means the bio is always left as blank and same is done for followers/retweeters. We collect all the 3,200 tweets (limit set by the twitter API) of each user ID, directly available from Twitter. For reliable prediction, we filter the users who have tweeted at least 100 times before 2017 and at least 100 times after 2018<sup>16</sup> This leaves us with 2,893 users. We then predict the inclination of these Twitter users twice. Once we use the features collected from tweets before 2017 and once we use the tweets after 2018. We observe that the predictions match for 91% of the cases, which tells that political leanings are temporally (almost) invariant.

**Hidden Agenda - Inclination Behind Promoted Hashtags:** To find out the inclination behind each hashtag, we obtain the political leanings of the users in the collected hashtag-specific dataset using the zero-shot classifier trained on followers of Congress and BJP. We plot the percentage of users leaning toward each party for each hashtag. We correctly predicted the leaning in 25 of the 30 cases using the classifier (considering a percentage distribution of 40–60% as the neutral/apolitical zone). We plot the leanings on four different India-specific issue – *#WeAreWithYouPmModiJi*, *#BengalBurning*, *#CycloneTaukte*, *#JusticeForAsif* in Fig. 3b. While we see the disaster hashtag (*#CycloneTaukte*) is non-polarizing, other trending hashtags are evidently promoted by people of particular ideologies. We include the list of other hashtags in Appendix.

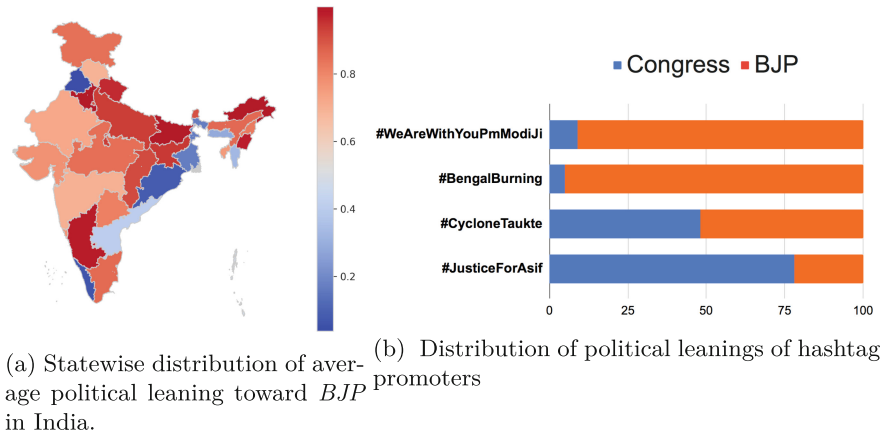


Fig. 3. Inclination distribution in India.

## 8 Limitations and Future Work

Our work is limited by the availability of social media data. If a country does not have enough political participation in the social media, then training a model

<sup>16</sup> We do not consider the follower-followee network as the past snapshots are not retrievable.

will not be possible. Moreover, if the profile of a person is kept private, the classifier will not be able to assign any label. We have discussed the related ethical implications of our work separately in Appendix.

Lastly, here we only evaluated our method on a dataset of users with high degree of political connection to very low degree of political connection. Collection of a dataset of users with *no political links* online but inclination toward a particular political party is a challenging task. In fact, Twitter matched voter registration data [3] also shows high partisanship evident in tweets and political connections. Research toward implicit(not explicitly tweeted/mentioned) political inclination detection (like implicit hate speech detection [5, 12] or implicit aspect specific sentiment detection [6, 16]) is an interesting future research direction.

## 9 Conclusions

We present an efficient, fast, and scalable few-shot learning framework for Twitter profiles for political inclination detection (FSSPIP). We showed that our model is explainable and learns features that humans find meaningful. Moreover, our model does not store any personal data of users unlike graph based models. It is also shown to be faster than graph-based methods. With the scalable representation learning framework, we achieve state-of-the-art accuracy, gaining significantly in unlabelled or few-shot learning setups on non-politician users. Enabling *zero-shot political inclination detection* with high fidelity, we provide a method to easily re-target this work to new countries and languages without any manual intervention/supervision unlike previous methods. We believe this will make a large-scale analysis of the political landscape throughout the globe easier and more accurate.

## References

1. Aldayel, A., Magdy, W.: Your stance is exposed! analyzing possible factors for stance detection on social media. Proc. ACM Hum.-Comput. Interact. **3**(CSCW) (2019). <https://doi.org/10.1145/3359307>
2. Baly, R., et al.: 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, pp. 3364–3374. Association for Computational Linguistics, July 2020. <https://doi.org/10.18653/v1/2020.acl-main.308>
3. Barberá, P.: Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. Polit. Anal. **23**(1), 76–91 (2015)
4. Cer, D., et al.: Universal sentence encoder. arXiv preprint [arXiv:1803.11175](https://arxiv.org/abs/1803.11175) (2018)
5. Chakraborty, S., Dutta, P., Roychowdhury, S., Mukherjee, A.: CRUSH: contextually regularized and user anchored self-supervised hate speech detection. In: Findings of the Association for Computational Linguistics: NAACL 2022, pp. 1874–1886. Association for Computational Linguistics, Seattle, United States, July 2022. <https://doi.org/10.18653/v1/2022.findings-naacl.144>, <https://aclanthology.org/2022.findings-naacl.144>



6. Chakraborty, S., Goyal, P., Mukherjee, A.: Aspect-based sentiment analysis of scientific reviews. In: Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020, JCDL 2020, pp. 207–216. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3383583.3398541>
7. Chakraborty, S., Goyal, P., Mukherjee, A.: (Im) balance in the representation of news? an extensive study on a decade long dataset from India. In: International Conference on Social Informatics, SocInfo (2022), [arxiv.org/abs/2110.14183](https://arxiv.org/abs/2110.14183)
8. Chen, W., Zhang, X., Wang, T., Yang, B., Li, Y.: Opinion-aware knowledge graph for political ideology detection. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, pp. 3647–3653 (2017). <https://doi.org/10.24963/ijcai.2017/510>, <https://doi.org/10.24963/ijcai.2017/510>
9. Conover, M.D., Goncalves, B., Ratkiewicz, J., Flammini, A., Menczer, F.: Predicting the political alignment of twitter users. In: 2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing, pp. 192–199 (2011). <https://doi.org/10.1109/PASSAT/SocialCom.2011.34>
10. Darwish, K., Stefanov, P., Aupetit, M., Nakov, P.: Unsupervised user stance detection on twitter. Proc. Int. AAAI Conf. Web Soc. Media **14**(1), 141–152 (2020). <https://ojs.aaai.org/index.php/ICWSM/article/view/7286>
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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, Vol.1 (Long and Short Papers), pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota, June 2019. <https://doi.org/10.18653/v1/N19-1423>, <https://aclanthology.org/N19-1423>
12. ElSherief, M., et al.: Latent hatred: A benchmark for understanding implicit hate speech. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 345–363. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, November 2021. <https://doi.org/10.18653/v1/2021.emnlp-main.29>, <https://aclanthology.org/2021.emnlp-main.29>
13. Iyyer, M., Enns, P., Boyd-Graber, J., Resnik, P.: Political ideology detection using recursive neural networks. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1113–1122. Association for Computational Linguistics, Baltimore, Maryland, June 2014. <https://doi.org/10.3115/v1/P14-1105>, <https://www.aclweb.org/anthology/P14-1105>
14. Kannangara, S.: Mining twitter for fine-grained political opinion polarity classification, ideology detection and sarcasm detection. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, pp. 751–752. Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3159652.3170461>, <https://doi.org/10.1145/3159652.3170461>
15. Kooops: Forgetting footprints, shunning shadows: a critical analysis of the right to be forgotten in big data practice. SCRIPTed **8**, 229 (2011)

16. Li, Z., Zou, Y., Zhang, C., Zhang, Q., Wei, Z.: Learning implicit sentiment in aspect-based sentiment analysis with supervised contrastive pre-training. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 246–256. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, November 2021. <https://doi.org/10.18653/v1/2021.emnlp-main.22>, <https://aclanthology.org/2021.emnlp-main.22>
17. Meeks, L.: Tweeted, deleted: theoretical, methodological, and ethical considerations for examining politicians’ deleted tweets. *Inf. Commun. Soc.* **21**(1), 1–13 (2018). <https://doi.org/10.1080/1369118X.2016.1257041>. <https://doi.org/10.1080/1369118X.2016.1257041>
18. Nguyen, D.Q., Vu, T., Nguyen, A.T.: BERTweet: a pre-trained language model for English Tweets. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 9–14 (2020)
19. Sood, G., Laohaprapanon, S.: Predicting race and ethnicity from the sequence of characters in a name (2018)
20. Stefanov, P., Darwish, K., Atanasov, A., Nakov, P.: Predicting the topical stance and political leaning of media using tweets. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 527–537. Association for Computational Linguistics, July 2020. <https://doi.org/10.18653/v1/2020.acl-main.50>, <https://www.aclweb.org/anthology/2020.acl-main.50>
21. Thompson, A., Stringfellow, L., Maclean, M., Nazzari, A.: Ethical considerations and challenges for using digital ethnography to research vulnerable populations. *J. Bus. Res.* **124**, 676–683 (2021)
22. Vahini, M., Bantupalli, J., Chakraborty, S., Mukherjee, A.: Decoding demographic un-fairness from indian names. In: International Conference on Social Informatics, SocInfo (2022), [arxiv.org/abs/2209.03089](https://arxiv.org/abs/2209.03089)
23. Williams, M.L., Burnap, P., Sloan, L.: Towards an ethical framework for publishing twitter data in social research: Taking into account users’ views, online context and algorithmic estimation. *Sociology* **51**(6), 1149–1168 (2017). <https://doi.org/10.1177/0038038517708140>, <https://doi.org/10.1177/0038038517708140>, PMID: 29276313
24. Xiao, Z., Song, W., Xu, H., Ren, Z., Sun, Y.: Timme: Twitter ideology-detection via multi-task multi-relational embedding. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2020, pp. 2258–2268. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3394486.3403275>, <https://doi.org/10.1145/3394486.3403275>
25. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=r1Ddp1-Rb>