

Does Twitter know your political views? POLiTweets dataset and semi-automatic method for political leaning discovery

Joanna Baran, Michał Kajstura, Maciej Ziółkowski, Krzysztof Rajda

Wrocław University of Science and Technology

Wyspiańskiego 27, 50-370 Wrocław, Poland

{joanna.baran, krzysztof.rajda}@pwr.edu.pl, {242491, 242475}@student.pwr.edu.pl

Abstract

Every day, the world is flooded by millions of messages and statements posted on Twitter or Facebook. Social media platforms try to protect users' personal data, but there still is a real risk of misuse, including elections manipulation. Did you know, that only 10 posts addressing important or controversial topics for society are enough to predict one's political affiliation with a 0.85 F1-score? To examine this phenomenon, we created a novel universal method of semi-automated political leaning discovery. It relies on a heuristical data annotation procedure, which was evaluated to achieve 0.95 agreement with human annotators (counted as an accuracy metric). We also present POLiTweets - the first publicly open Polish dataset for political affiliation discovery in a multi-party setup, consisting of over 147k tweets from almost 10k Polish-writing users annotated heuristically and almost 40k tweets from 166 users annotated manually as a test set. We used our data to study the aspects of domain shift in the context of topics and the type of content writers - ordinary citizens vs. professional politicians.

Keywords: political affiliation, political leaning, political profiling, Twitter dataset, political dataset

1. Introduction

Digital traces such as social media posts or interactions appear to be extremely powerful in gaining knowledge about our personal lives and predicting political beliefs (Kosinski et al., 2013). Automated analysis of electoral support on social media could effectively replace traditional surveys, being more cost-effective, allowing for examining a much larger portion of the population and giving more insights into voters' profiles. Unfortunately, as history has shown with the Cambridge Analytics example (Boldyreva, 2018), it is easy to misuse personal data, even to manipulate election results. Since then, social media companies have done a lot to strengthen the security of users' personal data (Hu, 2020). Have they done enough? Is it still possible to easily access enough personal data to model users' political affiliation?

To verify these questions, we propose a universal, semi-automatic political affiliation discovery method and a POLiTweets dataset - the novel political leaning discovery dataset, consisting of over 147k tweets from almost 10k Polish-writing users. As far as we know, this is the first published Polish dataset for predicting an individual's political orientation.

The main objective of this article is to answer the following Research Questions: **(RQ1)** Is it possible to determine one's political affiliation using one's social media activity? **(RQ2)** How many social media posts are needed to accurately determine one's political leaning? **(RQ3)** Do professional politicians use the same language as their party supporters?

Our main contribution includes 3 points. Firstly, we designed a novel method of semi-automatic annotation of political leaning data. Secondly, we used so to collect the first Polish publicly open dataset for predicting

political affiliation in a multi-party setup. Finally, we performed an analysis of the current political polarization in the Polish Twitter community.

2. Related Work

Assessing political affiliation - which party a given user is a supporter - is the main challenge of constructing political affiliation detection dataset. There have been several works on labelling political orientation using Twitter or Facebook.

Manual annotation. Probably the most accurate method is to survey volunteers willing to declare their political opinion (Kosinski et al., 2013; Preotiuc-Pietro et al., 2017) or manual annotation by a group of specialists who have access to publicly available user account information (Cohen and Ruths, 2013; Samih and Darwish, 2020). Unfortunately, both approaches are very time-consuming and practically unattainable for a large amount of data. For this reason, many automatic annotation methods have been proposed.

Annotation by relations. This approach assumes that the majority of the politicians we observe on Twitter reflect our views to some extent. The most popular method is to propagate party labels through analysis of followers and following users of politician accounts. (King et al., 2016; Golbeck and Hansen, 2014; Sylwester and Purver, 2015; Barberá, 2015). However, the follower relationship can introduce some noise, given the existence of users observing opposing politicians (eg. to get a balanced opinion). An et al. (2012) provided annotation by exploiting the political bias of popular news media provider accounts in the USA and their mutual followers' group.

Annotation by keywords. The use of hashtags or keywords in tweets has been explored in some work as

an indicator of user’s affiliation, as exemplified by the Scottish Independence Referendum study of social media public opinions (Fang et al., 2015). Similarly, Tattman (2017) focused only on specific electoral slogans in users’ bios or usernames.

Annotations by interactions. These methods take advantage of users’ interactions with the content they view. Rajadesingan and Liu (2014) introduced a semi-supervised retweet-based label propagation algorithm, based on the belief that retweeting a tweet may indicate endorsement of its content. Nevertheless, one can easily find tweets shared with a negative comment, retweeted to show disapproval. A more reliable premise is a user’s ”like” (Papakyriakopoulos et al., 2018). Analysis of the likes distribution towards political parties’ posts has proven to be very effective in this task and was verified by a manual survey in Kristensen et al. (2017). Like is a very strong signal of support towards the post creator and, in our opinion, the best choice to develop heuristics for assigning political labels to users’ social media profiles.

Finally, we need to point out that most of the published collections have focused on the two-party system, as is present in the USA (Conover et al., 2011; Yan et al., 2017) (binary classification setup). The multi-party prediction problem is far more challenging and yet still underestimated in research works, even though the majority of democratic countries are characterized by a diverse political arena. This highlights the importance of creating multi-party datasets for a more complex study of social media users.

3. Political Afliation Discovery Method

In this section, we describe the proposed semi-automatic method of political leaning discovery and test its quality.

3.1. Method Steps

Our approach to detect one’s political affiliation goes as follows (Figure 1):

- (1) Manually create the list of prominent political figures and their Twitter accounts in examined country/culture. Assign them a party label according to the politician’s affiliation.
- (2) Aquire tweets posted from those accounts (we used an official Twitter API).
- (3) Save a list of social media users that liked collected politicians’ tweets.
- (4) Filter out individuals with less than 10 likes to reduce the noise in acquired labels.
- (5) Count and group each users’ likes by a political party. Choose the most frequent one as a label. Exclude inconclusive users with an equal number of likes for different parties for training purposes. For testing purposes, they could be manually annotated and included in the test set.
- (6) Collect selected users’ tweets. Each scraped entry had to contain at least one predefined hashtag or keyword related to controversial topics - a post addressing

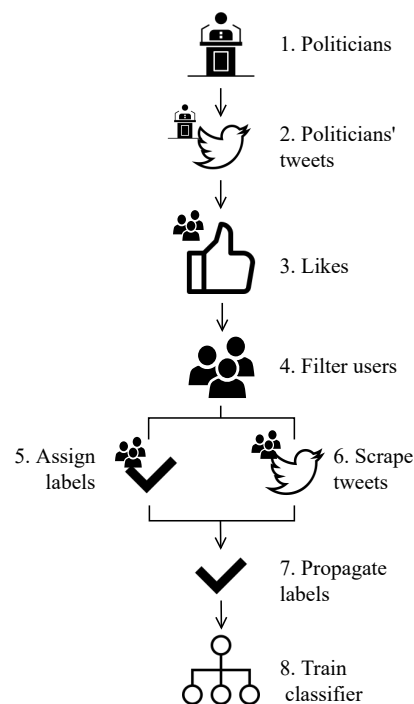


Figure 1: An illustration of the data acquisition pipeline

this kind of politically loaded and divisive subject is more likely to reflect the political views of a person writing it Bail et al. (2018).

(7) Propagate obtained users’ labels to all users’ tweets, resulting in a heuristically annotated text classification dataset.

(8) Train a text classifier, which can be then used for analysing the political affiliation of any provided text. It can be also applied to a user-level classification, aggregating individual posts’ labels for each user, further increasing the accuracy of the method.

3.2. Heuristic Quality Evaluation

As the proposed method relies on a heuristic approach to data annotation, we ensured its performance by comparing it with human annotators. All users present in the test set were manually verified and labelled by three independent annotators using all available public data, such as posts, bio, followers network and uploaded images. When a discrepancy occurred, a majority vote was taken to select a party name. Determining a person’s political preferences based on social media posts is a highly subjective task, but annotators scored a high inter-annotator agreement of 0.74 Krippendorff’s alpha coefficient. In total, we had 29 960 posts from 133 users for which we had human annotations.

Labels obtained using the proposed automatic labelling scheme matched the manual annotations with a **0.95 accuracy**, which confirms heuristics reliability.

4. POLiTweets Dataset Summary

We used our proposed political leaning discovery method to collect a dataset consisting of **186 868**

tweets in Polish language, written by 9 837 Twitter users. Description and full hashtag list of controversial topics selected in method Step 6 are presented in Section 9.2. Data were posted between March 2021 and January 2022. As labels set, we have chosen five main political groups with the greatest impact on the current Polish political scene - our selection was based on Parliament representation and election polls, more details in Section 9.1.




Figure 2: POLiTweets party label distribution

The distribution of assigned labels was unbalanced (Figure 2), which is quite consistent with the Polish political reality.

The previously manually annotated users and their tweets were taken as a test set (see Section 3.2). It contained only profiles with at least 15 posts per each. The rest of the data (over 147k tweets) was randomly split in a 9:1 ratio creating a training and validation split. We ensured that each entry has a minimum of 5 words that are not a hashtag, user mention or URL. The dataset is available in CSV files containing TweetID instead of explicit text, according to Twitter’s guidelines about redistributing their content for scientific purposes. Test collection file appears twice due to different annotation sources - from heuristics method (*heuristics-test*) and human annotators (*manual-test*). We provide also a file with posts of the Polish parties’ official Twitter accounts and their active politicians’ list. Additionally, we prepared an *ambiguous test set* of manually annotated tweets of 33 users for whom we skipped filtering by the ratio of likes in Step 5 of our method (Section 3.1). We will refer to them as *ambiguous* because discovering their political views is a much more difficult task even for a human annotator due to a more uniform distribution of those users’ likes across political parties. All files are publicly available in the GitHub repository ¹.

5. Experiments and Results

The final step of our method was to train a classifier on the acquired tweets’ textual data, without using any additional knowledge. We used it for analysis of current political polarization among the Polish Twitter community.

Experimental setup. As a classifier, we finetuned the base version of HerBERT - a BERT-based language

model dedicated to Polish language (Mroczkowski et al., 2021) - with a sequence classification head. The choice of this architecture was justified by its numerous effective applications in the NLP field. We considered it a good starting point for obtaining preliminary results on our dataset. The training took maximally 50 epochs with an early stopping patience parameter equal to 15 epochs. The batch size was set to 32 and the model was optimized using AdamW with a learning rate of 1e-5 along with a warmup linear scheduler. The pre-processing stage included the removal of hashtags, user mentions and URLs from posts to prevent data leakage. Due to the large imbalance of classes in the dataset, we applied a weighted sampler.

Model performance. The text classifier achieved 0.64 micro F1-score on the *manual-test* split. More detailed results are presented in Table 1.

Party name	Precision	Recall	F1-score
PiS	0.78	0.75	0.76
PO	0.64	0.67	0.65
Konfederacja	0.35	0.36	0.36
Lewica	0.18	0.26	0.21
PL2050	0.31	0.07	0.12
Total	-	-	0.64

Table 1: Scores for particular classes

Classifier errors as Polish political scene descriptor.

The confusion matrix, presented in Figure 3, reveals differences between political parties in Poland. Rows and columns are ordered by the place in political spectrum, *Lewica* being the furthest to the left and *Konfederacja* on the rightmost side (Kosowska-Gastol, 2021). The text classifier makes most errors between political groups with similar social and economic views. The lowest score is achieved for the centrist party - *Pol-ska2050*, which can be explained mostly by a strong class imbalance and hard to distinguish political opinions.

	Lewica	PL2050	PO	PiS	Konfederacja
Lewica	0.26	0.0	0.37	0.16	0.21
PL2050	0.13	0.072	0.45	0.23	0.12
PO	0.08	0.0079	0.67	0.16	0.081
PiS	0.02	0.0012	0.14	0.75	0.087
Konfederacja	0.082	0.0026	0.16	0.39	0.36

Figure 3: F1-score confusion matrix

Effects of Domain Shift. We investigated the effect of domain shift on our classifier by considering it from two perspectives. To examine writers’ shift, the clas-

¹<https://github.com/Joanna065/politweets>

sifier trained previously on tweets from regular Twitter users (ordinary citizens) was used to predict party labels on professional politicians’ posts, which were obtained in Step 2 of our method (Section 3). To study domain shift among topics, the training was carried out using the same experimental setup as stated above, but with a training set containing posts from 3 topics, leaving the last topic’s posts as a test set.

Domain-out	F1-score
Politicians	0.35
Topic - abortion	0.40
Topic - lexTVN	0.45
Topic - EU & CJEU	0.48
Topic - The Polish Order	0.46

Table 2: Results for domain shift study. Each row states test set used

Results shown in the Table 2 prove that the model’s performance under domain shift is severely affected. The model tested on politicians’ tweets achieved only 0.35 of micro averaged profile-level F1-score. A similar situation occurred for the experiment on topics - the highest result was 0.48 F1-score, compared to 0.65 without the domain shift scenario (Table 1).

6. Discussion

RQ1: Is it possible to determine one’s political affiliation using one social media activity? Yes, but we have to take into account the quality and balance of the input data. The research showed that the underrepresented parties were less recognizable by the classifier - model performance was the highest for more frequent political labels (like *PiS* or *PO* in the Polish case).

RQ2: How many social media posts are needed to accurately determine one’s political leaning? To answer that question, we tested the classifier incrementally by increasing a subset of users’ posts (from 1 to 15), choosing the most common label for the profile. We conducted evaluations on all 3 available test sets. The experiment was repeated 30 times with different tweet selection orders. Basing final predictions on more tweets drastically improves the performance of our method, which is presented in Figure 4. Accumulated predictions for well-defined users, for whom the number of collected likes was significantly higher for one party than the others, allowed to achieve a micro-averaged F1-score of 0.85 when used on 10 posts. For the more difficult user cases, the classifier scored significantly lower, but still above 0.7 F1-score.

RQ3: Do professional politicians use the same language as their party supporters? Unfortunately, no. The classifier we chose proved to be sensitive to any changes in test data compared to training examples. It occurs not only in writer shift but also in the aspect of topic shift. This leads to the conclusion that despite the impressive results on in-domain texts, the deployment of such models should be carried out with special care.




Figure 4: Impact of the number of posts on model performance

Method advantages. The main advantage of our method is resource and time savings, compared to manual dataset annotation or traditional political affiliation surveying. Such surveys typically last for days to acquire a representative population, while our method can perform in only a few hours. It is also quite general - the data collection scheme can be applied to other social media platforms and any country or culture, due manual selection of controversial topics and seed profiles. Moreover, models trained on such data can also generalize to any kind of in-domain text, even those from users not actively liking any content produced by politicians.

Method limitations. We assessed political leaning in a semi-automatic way. It requires choosing an initial set of politicians’ accounts to scrape tweets from. It also needs a selection of some controversial, politically-loaded topics to filter out the discussion. Such topics depend mostly on country culture and current political debate and need to be chosen with care.

7. Conclusions and Future Works

In this work, we proposed a semi-automatic, universal schema for political affiliation discovery, based on the heuristic method of data acquisition and annotation. Our approach proved to be highly effective, achieving a 0.95 user-level accuracy agreement when compared with the manually annotated dataset. We also introduce POLiTweets - the first publicly available dataset for political leaning analysis in Polish, with almost 187k tweets from nearly 10k users. We proved that using tweets from popular and controversial topics, it is possible to associate Twitter users with political parties they support with sufficient confidence. Such knowledge may be exploited for microtargeting purposes, similar to how it was in the Cambridge Analytica scandal (Boldyreva, 2018). Publishing even a few personal opinions on such topics may uncover users’ political views - as our experiments showed, 10 posts are enough to classify political affiliation with a 0.85 F1-score. Being aware of the high dependency on the data domain - topics and writers’ type - we’d like to apply the domain adaptation techniques (eg. Ma et al. (2019)) to support classifier stability. Also, more model architectures are needed to examine. We find it a good start for the follow-up work we tend to perform.

8. Bibliographical References

- An, J., Cha, M., Gummadi, K. P., Crowcroft, J., and Quercia, D. (2012). Visualizing media bias through twitter. In *ICWSM 2012*.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.
- Barberá, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Analysis*, 23:76 – 91.
- Boldyreva, E. (2018). Cambridge analytica: Ethics and online manipulation with decision-making process. pages 91–102, 12.
- Cohen, R. and Ruths, D. (2013). Classifying political orientation on twitter: It’s not easy! *Proceedings of the 7th International Conference on Weblogs and Social Media, ICWSM 2013*, pages 91–99, 01.
- Conover, M. D., Goncalves, B., Ratkiewicz, J., Flammini, A., and Menczer, F. (2011). 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*, pages 192–199.
- Fang et al. (2015). Topic-centric classification of twitter user’s political orientation. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’15*, page 791–794, New York, NY, USA. Association for Computing Machinery.
- Golbeck, J. and Hansen, D. (2014). A method for computing political preference among twitter followers. *Social Networks*, 36:177–184. Special Issue on Political Networks.
- Hu, M. (2020). Cambridge analytica’s black box. *Big Data & Society*, 7(2):2053951720938091.
- King, A. S., Orlando, F. J., and Sparks, D. B. (2016). Ideological extremity and success in primary elections: Drawing inferences from the twitter network. *Social Science Computer Review*, 34(4):395–415.
- Kosinski, M., Stillwell, D., and Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 110, 03.
- Kosowska-Gastoł, B. (2021). Te same, ale czy takie same? : analiza oblicza ideowo-programowego polskich partii politycznych - różne perspektywy metodologiczne. In Aleksandra Kruk, editor, *Postulaty polityczne i wyborcze partii politycznych*, pages 17–35. Oficyna Wydawnicza UZ, Zielona Góra.
- Kristensen, J., Albrechtsen, T., Dahlgaard, E., Jensen, M., Pedersen, M. S., and Bornakke, T. (2017). Parsimonious data: How a single facebook like predicts voting behaviour in multiparty systems. *PLOS ONE*, 12, 04.
- Ma, X., Xu, P., Wang, Z., Nallapati, R., and Xiang, B. (2019). Domain adaptation with BERT-based domain classification and data selection. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 76–83, Hong Kong, China, November. Association for Computational Linguistics.
- Mroczkowski, R., Rybak, P., Wróblewska, A., and Gawlik, I. (2021). Herbert: Efficiently pre-trained transformer-based language model for polish. *CoRR*, abs/2105.01735.
- Papakyriakopoulos, O., Hegelich, S., Shahrezaye, M., and Medina Serrano, J. C. (2018). Social media and microtargeting: Political data processing and the consequences for germany. *Big Data Society*, 5:205395171881184, 07.
- Preotiuc-Pietro, D., Liu, Y., Hopkins, D. J., and Ungar, L. H. (2017). Beyond binary labels: Political ideology prediction of twitter users. In *ACL*.
- Rajadesingan, A. and Liu, H. (2014). Identifying users with opposing opinions in twitter debates. 02.
- Samih, Y. and Darwish, K. (2020). A few topical tweets are enough for effective user-level stance detection. *ArXiv*, abs/2004.03485.
- Sylwester, K. and Purver, M. (2015). Twitter language use reflects psychological differences between democrats and republicans. *PLOS ONE*, 10(9):1–18, 09.
- Tatman, R. (2017). #maga or #theresistance: Classifying twitter users’ political affiliation without looking at their words or friends.
- WNP.PL. (2021). Rośnie poparcie dla pis, ko i polski 2050.
- Yan, H., Lavoie, A., and Das, S. (2017). The perils of classifying political orientation from text. In *LINK-DEM@IJCAI*.

9. Appendix

9.1. Used electoral polls

Our selection of political parties to label POLiTweets dataset was based on the results of the September 24–27, 2021 CAWI survey conducted on a nationwide and representative group of Polen (WNP.PL, 2021), which are as follows: Prawo i Sprawiedliwość (PiS, *eng. Law and Justice*) - 38%, Platforma Obywatelska (PO, *eng. Civic Platform*) - 26%, Polska 2050 (PL2050, *eng. Poland 2050*) - 14%, Konfederacja (*eng. Confederation*) - 9%, Lewica (*eng. The Left*) - 8%.

9.2. Controversial topic selection

The period of data acquisition and conduction of our research was in 2021 and early 2022. At that time, the most divisive topics in Polish society with political overtones were:

1. Abortion (*pol. aborcja*) - due to the tightening of the abortion law and the resulting numerous strikes in the country,

Topic	Keywords list	No. tweets
Abortion	aborcja, strajkkobiet, godek, czarnyprotest, AniJednejWiecej, prolife, BabiesLivesMatter, LegalnaAborcja, AborcjaJestOk	40 116
EU & CJEU	tsue, turów, polexit, konstytucja, zostajeMYwUE, MyZostajemy, NieWygasiciePolski, TrybunałKonstytucyjny, trybunał, UniaToMy, ZostajewUnii, PolexitNow, ZostajemyWEuropie	73 733
LexTVN	lextvn, tvn	63 609
The Polish Order	PolskiŁad, Polski Ład, PolskiLad, Polski Lad	5 277

Table 3: Topics with descriptive keywords along with their number in the dataset

2. European Union, EU & Court of Justice of the European Union, CJEU (*pol. Unia Europejska, UE & Trybunał Sprawiedliwości Unii Europejskiej, TSUE*) - in connection with the penalties imposed on Poland resulting from the extension of active mining in Turów and the expansion of anti-EU sentiments by the ruling party PiS,
3. LexTVN - an amendment to the Broadcasting Act concerning the granting of broadcasting licenses to foreign entities,
4. The Polish Order (*pol. Polski Ład*) - the plan for the recovery of the Polish economy after the COVID-19 pandemic proposed from 2022, including new changes in tax law.

Topic distribution. It is worth mentioning that in the POLiTweets dataset we obtained there were posts concerning several topics at the same time. We present detailed statistics in Table 3 along with a list of keywords we used to collect tweets via the API. Most of the data were acquired on *EU & CJEU* and *LexTVN* - these were the main highlights of the 2021 year in Poland. The least amount of data comes from *The Polish Order* topic, as the law came into effect in early 2022 when the collection of the dataset has ended.