

DADIT: A Dataset for Demographic Classification of Italian Twitter Users and a Comparison of Prediction Methods

Lorenzo Lupo, Paul Bose, Mahyar Habibi, Dirk Hovy, Carlo Schwarz

All authors are based at Bocconi University

Via Roberto Sarfatti 25, 20136 Milano, Italy

{paul.bose, mahyar.habibi, lorenzo.lupo, dirk.hovy, carlo.schwarz}@unibocconi.it

Abstract

Social scientists increasingly use demographically stratified social media data to study the attitudes, beliefs, and behavior of the general public. To facilitate such analyses, we construct, validate, and [release publicly](#) the representative DADIT dataset of 30M tweets of 20k Italian Twitter users, along with their bios and profile pictures. We enrich the user data with high-quality labels for gender, age, and location. DADIT enables us to train and compare the performance of various state-of-the-art models for the prediction of the gender and age of social media users. In particular, we investigate if tweets contain valuable information for the task, since popular classifiers like M3 don't leverage them. Our best XLM-based classifier improves upon the commonly used competitor M3 by up to 53% F1. Especially for age prediction, classifiers profit from including tweets as features. We also confirm these findings on a German test set.

Keywords: Twitter data, demographic prediction, language models, multimodal classification

1. Introduction

Social media data are a valuable support for economic and social science research, and they have been leveraged extensively in the literature (see [Steinert-Threlkeld, 2018](#), for a review). For many use cases, social media users need to be profiled along their demographics to study differences within and between groups (e.g., [Holmberg and Hellsten, 2015](#); [Usher et al., 2018](#); [Allen et al., 2022](#)). For example, studying social and economic inequalities in a country requires stratifying users across gender, age, and geographic location.

Social media users' demographics are typically inferred from the user's profile information, such as the username ([Lockhart et al., 2023](#)), their posted pictures ([You et al., 2014](#)), or their network ([Al Zamil et al., 2012](#); [Barberá, 2015](#)). One of the most popular tools, the M3 classifier ([Wang et al., 2019](#)), predicts age and gender multimodally, leveraging usernames, users' bios, and pictures. However, these tools often depend on the availability of biographical information or profile images that contain identifiable pictures of the actual users. Whenever this information is absent, the accuracy of these approaches drops significantly, or they become unusable altogether. For example, almost 27% of US adults do not provide any text in their Twitter profile, and 14% do not provide any photo (e.g., rear shots or nature photos; see [Pew \(2022\)](#)); a lot of information is also missing on Italian profiles (Table 1). Both in the absence and the presence of these features, users' posts provide valuable information for user characteristics, but surprisingly, little attention has been paid to leveraging this information.

We bridge this gap by collecting and releasing

DADIT: a GDPR-compliant database of around 30M tweets from 20K representative Italian Twitter users, along with their profile pictures and bios. DADIT contains high-quality, manually evaluated labels for gender and age, along with location labels for 43% of the users. Given the recent challenges of gathering Twitter data for academic research, our collection supports future research.

Leveraging this data, we systematically compare gender and age classification methods. We are particularly interested if and by how much the inclusion of users' posts (tweets) improves the prediction of gender and age by comparing modern text classification models ([Brown et al., 2020](#); [Barbieri et al., 2021](#); [Chung et al., 2022](#)) to traditional users' classifiers: M3 ([Lockhart et al., 2023](#)) and another popular computer vision (CV) classifier ([Levi and Hassner, 2015](#)). We find that models leveraging bios and tweets often outperform models based on bios and pictures. In particular, a text classifier finetuned on our Italian dataset improves upon the state-of-the-art M3 by a significant margin. This finding underscores the usefulness of text-abundant, demographically annotated datasets like DADIT. Finally, we demonstrate that ensembling multiple models' predictions further improves classification performance. To strengthen the generalizability of our results, we also experiment on a German test set, corroborating the Italian results.

Our contributions are:

1. We release DADIT, an extensive database of around 30M tweets from 20K Italian Twitter users, including profile photos, bios, and tweets, enriched with high-quality demographic information (age, gender) and geographic information (Italian regions and, in

some cases, municipalities). The data will be released in GDPR and Twitter-compliant anonymized format (see also Section 5).

2. We systematically compare user attribute prediction methods in Italian and German. We demonstrate that including tweet texts brings significant improvements over profile-based classification methods like M3, especially when training data like DADIT are available.

2. The Dataset

2.1. Collection

We collected the data through the Twitter API in early 2023. We used the Twitter accounts of 802 Italian politicians and parties as a source list for user IDs.¹ For each account, we collected the user IDs of followers. In total, we obtained 3.2M unique user IDs. We then collected each user’s profile and tweets as a second step. Note that we did not obtain tweets for all users for one of three reasons: 1) some users do not tweet, 2) some accounts were deleted before we were able to scrape them, or 3) the Twitter APIs were shut down in early 2023 as a result of Elon Musk’s takeover of Twitter which prevented further data collection. Nonetheless, we obtained around 830.1M tweets for 2.1M users.

2.2. Annotation

After collection, we annotated user characteristics. More specifically, we created reference labels for gender, age, and location. Here, we briefly outline the annotation of these characteristics. First, we coded the (binary) gender of Italian Twitter users based on the full name field of each Twitter user. This process was facilitated by the fact that Italian names are gender-specific in nearly all cases.² Since many users share their full names on Twitter, we were able to assign reference labels to 1.3M users out of 2.1M. Second, we annotated the age of users, constructing a set of regex patterns matching users’ statements of birthyear or current age in either their bio or any of their tweets. Differently from the gender annotation, we were able to assign reference labels for age only to a minor fraction of the users: 33K users. This imbalance highlights an important aspect of the demographic classification task on Twitter: there is a strong need for age classifiers for the vast majority of users, while gender can be classified for large shares of the

¹See Appendix A for an overview of the number of accounts and the total number of followers by party.

²We used the most common Italian names and their associated gender reported on this list: <https://github.com/mrblasco/genderNamesITA>

social media population with rule-based classifiers, or even with a simple dictionary, as we did.

Lastly, we identified each profile’s approximate locations by matching the profile location information against a list of Italian regions and municipalities. All three procedures for annotating reference labels were designed to achieve high accuracy at the expense of potentially annotating fewer users. Nonetheless, we are left with close to 20K users for which we have reference labels for both gender and age. In the final dataset, we focus on users with both gender and age information to have a consistent sample for the prediction across models.

2.2.1. Data quality

To ensure the accuracy of the coding, three members of the research team manually verified the reference labels for gender and age in a random subset of 1,200 users for gender and 350 users for age. For gender, the three annotators marked between 99.33% and 99.42% of the labels as correct, with an average inter-annotator raw agreement of 99.94%. For age, the annotators marked between 93.55% and 97.10% of the labels as correct, with an average inter-annotator raw agreement of 94.41%. Two out of three annotators extended the verification of age labels to all the 1,200 users that were verified for gender and then discarded from the dataset the users with incorrect labels (either gender or age) according to at least one human annotator, keeping the remaining as a test set for the following experiments on gender and age classification. Hence, our test set guarantees 100% quality (according to two annotators) labels for each user.

To make the data even more useful for follow-up research, we also evaluate the representativeness of our data relative to the population of Italian Twitter users. Overall, the results, which we discuss and visualize in greater detail in Appendix A, suggest that our data closely mimic the characteristics of the overall Italian Twitter population.

	Train	Valid	Test
# Users	17262	1919	1119
% w/ Pic	0.80	0.78	1.00
% w/ Bio	0.47	0.45	0.46
% w/ Location	43.0	41.1	47.2
% Female	0.37	0.37	0.36
Average Age	42.3	42.7	42.3
# Tweets	24.5M	2.8M	1.6M
% RT	23.8	24.6	19.5
# Tweets/User	338	355	274

Table 1: DADIT overview

3. Gender and Age Classification

In this section, we analyze whether 1) tweets constitute a valuable source of information for the classification of the gender and age of users; 2) text-only classifiers can compete with standard multimodal models like M3; 3) all the modalities available in our dataset (text and pictures) can be leveraged in a simple way to improve performance.

3.1. Experimental setting

We use the constructed dataset to train and compare a number of gender and age classifiers. For our analysis, we divide our database of 20K Twitter users into a training (85.0%), validation (9.5%), and test set (5.5%) set. Table 1 provides an overview of each split’s users, tweets, and demographic characteristics. Note that we split the data to ensure that all users in the test set have profile images. Otherwise, all samples are identical based on demographic characteristics.

For the classification task, we are interested in cases where gender and age can not be inferred with the rule-based methods that we used for the automatic annotation of DADIT, because of the lack of the necessary clues. We simulate this situation by avoiding to provide models with users’ full names as a feature, and we mask age and birthyear information from bios and tweets. We apply a light preprocessing to tweets: removing mentions, URLs, and any tweet shorter than three words. For each user, we concatenate their bio and most recent tweets (until the maximum input length is reached) to form one contiguous text to be input as features in the text classifiers. We compare the performance of a dummy classifier that predicts the majority class with four models.

M3 As a strong baseline, we use the approach of Wang et al. (2019), which is one of the standard models for inferring the age and gender characteristics of Twitter users (Bianchi et al., 2022). M3 can use users’ usernames, biographies, and profile images as features, but not tweets. In our analysis, we avoid providing usernames as we experiment in a setting where this information is unavailable (otherwise, gender could be inferred with high confidence as we did in Section 2.2). M3 classifies ages into four classes: (0-19), [19-30), [30-40), and [40-100). We adopt the same age classes for all the other models to make them comparable. We compare the performance of M3 with and without the profile image. Omitting the profile image tells us about M3’s performance when users do not provide a profile image or the image offers no identifiable characteristics.

CV As an alternative baseline classifier, we use another popular, open-source model based on computer vision (CV) (Levi and Hassner, 2015).

This model predicts gender and age solely based on users’ profile pictures.

XLM We leverage our dataset to train a classifier based on twitter-XLM-roberta-base (Barbieri et al., 2021)³, a multilingual language model finetuned on twitter data. We train a classification head for gender and age, separately, and we also finetune the weights of the backbone model.

Flan-T5 Flan-T5 is a multilingual, instruction-finetuned language model achieving state-of-the-art performance in several zero-shot tasks (Chung et al., 2022). We test its largest checkpoint (flan-t5-xxl)⁴ in a zero-shot classification setting, as in Plaza-del Arco et al. (2023). We use it to classify the gender and age of users based on simple prompts reported in Appendix B.

GPT3.5 Finally, we test the gpt-3.5-turbo through OpenAI’s API, in a few-shot setting. GPT3.5 is one of the largest LMs available to date, the successor of the popular GPT3 (Brown et al., 2020). We use it to classify the gender and age of users based on simple prompts reported in Appendix Appendix B. In our preliminary experiments, GPT performed consistently better when provided with labeled examples in the prompt, consistently with Lupo et al. (2023). Therefore, we test it in a few-shot setting with two examples for each age class, and five for each gender.

3.2. Results

The results of the models’ evaluation on our Italian test set are in Tables 2. Three main findings stand out. First, we document that including information from tweets consistently and significantly improves model performance. For example, the F1-scores for gender prediction of both Flan-T5 and GPT3.5 increase by over 10 points. We observe even larger performance gains for the prediction of age.

Second, we find that the finetuned XLM stands out as the top performer in both tasks, particularly for age prediction, which is most crucial for users’ demographic classification, as discussed earlier. Even with profile images, M3 achieves an F1-score of 41.4 relative to an F1-score of 60.2 of the finetuned model. The F1-score of the finetuned model is also nearly twice as high as the score achieved by other competitors.

Third, we show that additional performance gain for the prediction of gender can be achieved by a simple ensemble learner that combines the predictions from M3 and XLM, weighting their output probabilities with their F1-scores. Surprisingly, however, ensemble results do not improve over XLM on age prediction.

³<https://huggingface.co/cardiffnlp/twitter-xml-roberta-base>

⁴<https://huggingface.co/google/flan-t5-xxl>

Model	Features	Italian				German			
		Gender		Age		Gender		Age	
		F1	Acc	F1	Acc	F1	Acc	F1	Acc
Dummy	majority class	39.2	64.4	15.9	46.8	44.1	78.9	17.0	51.3
M3	bio	50.9	67.7	16.6	16.4	-	-	-	-
M3	bio+pic	79.3	83.1	41.4	53.7	56.2	80.2	25.9	31.7
CV	pic	70.7	76.0	27.7	35.2	64.7	79.4	28.2	35.6
XLM	bio	81.6	72.5	33.3	52.2	-	-	-	-
XLM	bio+twit	86.8	88.1	60.2	66.4	<u>74.3</u>	<u>82.0</u>	<u>45.8</u>	<u>58.0</u>
Flan-T5	bio	<u>62.1</u>	<u>70.5</u>	<u>16.3</u>	<u>16.2</u>	-	-	-	-
Flan-T5	bio+twit	75.4	76.8	34.0	58.0	80.8	87.4	29.0	35.1
GPT3.5	bio	64.9	72.8	21.1	20.3	-	-	-	-
GPT3.5	bio+twit	77.4	78.0	31.4	32.0	48.2	72.2	31.4	41.2
XLM+M3	bio+twit+pic-	90.8	91.9	60.0	67.2	<u>74.8</u>	83.6	48.1	58.8

Table 2: Comparison of gender and age classifiers on Italian and German Twitter users. Highest performance per column in bold. Results outperforming our strong baseline (M3 bio+pic) with statistical significance ($p < 0.01$) are underlined.

Lastly, we compare the best Italian classifiers on a German test set, collected and annotated following the same procedure described in Section 2.2.⁵ Note that, after fine-tuning XLM on the Italian training data, we do not further finetune it on any German training data, since we lack them. Instead, we leverage XLM’s multilingual abilities to solve the German classification tasks without further fine-tuning. The findings from the German test set validate our results in Italian. Interestingly, we find that even though XLM was only finetuned on Italian Twitter data, it is still among the best-performing models for gender prediction in German, and achieves the highest F1-score in the prediction of age. This finding highlights again the importance of annotated datasets like DADIT.

3.3. Scaling up

In this section we explore ways to push even further the performance of our age and gender classification approach, with a focus on the Italian dataset. Two strategies are evaluated: employing a larger pre-trained encoder for the classifier and incorporating additional user-level features. We transition from `twitter-XLM-roberta-base` to the larger `twitter-XLM-roberta-large-2022` (XLM-large - Barbieri et al. (2021)⁶). Table 3 outlines the observed changes in classification performance, revealing a modest enhancement in age prediction and a more notable improvement in gender classification.

We next incorporate extra user-level features as input to our classifiers. In the age classifier, we added users’ names, their Twitter usernames, account creation years, and counts of tweets, follow-

Model	Target	Features	F1	Acc
XLM	Age	bio+twit	60.2	66.4
XLM-large	Age	bio+twit	61.1	67.3
XLM-large	Age	all	63.5	69.9
XLM	Gender	bio+twit	86.8	88.1
XLM-large	Gender	bio+twit	93.0	90.9
XLM-large	Gender	all-name	91.8	89.0

Table 3: Age and gender prediction with XLM classifiers on the Italian dataset.

ings, and followers. For gender classification, we used all these features but omitted users’ names for the reasons explained above.⁷ The inclusion of these features markedly improved the accuracy of age classification. However, the gender classification model showed a minor decline in performance.

3.4. Quality-Coverage Tradeoff

For settings requiring higher classification quality, we can comply at the expense of some coverage of classified users by restricting the classification to those users that the model classifies with higher confidence. Typically, observations are classified based on the highest probability class output by the model. Therefore, we can restrict the classification to the subset of users where the model’s output probability for the most likely class exceeds a decision threshold. In Figure 1, we show that the decision threshold (minimum probability for the predicted class) is positively correlated to the F1 score but negatively correlated to the share of users that can be predicted, for both gender and age classification. For instance, in gender classification,

⁵More details on the German test set in Appendix A.

⁶<https://huggingface.co/cardiffnlp/twitter-xlm-roberta-large-2022>

⁷The gender classifier trained with users’ names achieves nearly perfect accuracy, as expected




Figure 1: F1 score and prediction coverage tradeoff.

setting a threshold of 0.8 improves the F1 score by approximately 3 points against the standard 0.5 threshold.⁸ However, the 0.8 threshold also reduces the classification coverage of users by over 10%. In other words, more than 10% of the users can not be classified at the given confidence level. In age classification, the same threshold boosts the F1 score by around 5 percentage points, yet it halves the coverage of classified users.

4. Related Work

Accurately predicting the demographic characteristics of social media users has become central for computational social science. The wealth of information on online behavior has been used to infer, among others, demographic characteristics, income, political ideology, religion, and health outcomes (Al Zamal et al., 2012; Culotta, 2014; Barberá, 2015; Preoȃiuc-Pietro et al., 2017; Preoȃiuc-Pietro and Ungar, 2018; Matz et al., 2019; Jaidka et al., 2020; Fujiwara et al., 2023; Chaturvedi and Chaturvedi, 2023). Especially the demographic characteristics of users (e.g., gender, age, race, ethnicity) are used to characterize heterogeneity in the behavior of users groups (e.g., women and men) (e.g., Holmberg and Hellsten, 2015; Usher et al., 2018; Bathina et al., 2021; Zhang et al., 2021; Burel et al., 2021; Allen et al., 2022). Further, it has been shown that demographic characteristics can improve classification performance on other attributes (Volkova et al., 2013; Hovy, 2015; Lynn et al., 2017; Li et al., 2018; Hovy and Yang, 2021). As a result, models of user classification have become a staple of modern social media analysis.

Existing approaches to classify social media user data encompass various techniques that leverage the information available in online interactions.

⁸We used this decision threshold, which is equivalent to using none, in all our previous experiments on gender.

Overall, three main approaches are used for the classification of users. First and most frequently, information is extracted from the user profile information (e.g., screen name, bio) (Wang et al., 2019; Lockhart et al., 2023). Secondly, the profile information can be enriched using information from the profile image (Alowibdi et al., 2013; Wang et al., 2019). The profile image is particularly informative for classifying users' gender, age, and ethnicity. A third approach is analyzing the network connections on the social media platform. These approaches have been particularly successful in classifying the age and political ideology of users (Al Zamal et al., 2012; Barberá, 2015).

5. Conclusion

Social media platforms have provided a goldmine of data for machine learning and NLP models to analyze user behaviour and interactions. We introduce DADIT: a rigorously curated, demographically annotated data set for the social sciences. It contains over 20K representative Italian Twitter users, with over 30M tweets, useful for the study of political attitudes and beliefs. As social media analysis usually relies on the demographic characteristics of the users for stratification, we compare the performance of different methods for the automated prediction of users' gender and age in contexts where this information can not be easily retrieved from users' profiles. We show the latent potential of tweet content to enrich user classification approaches for all tweeters, regardless of profile information. Second, we demonstrate that finetuned models based on user bios and tweets systematically outperform multimodal state-of-the-art methods based on bios and profile images. Lastly, we suggest a lightweight solution to take advantage of the multimodality of our data.

Acknowledgements

The authors are grateful for the financial support of the Cariplo Foundation as part of the MENTALISM Project (MEasuriNg, Tracking, and AnaLyzing Inequality using Social Media). We are grateful to Derya Durmush for her excellent research assistance. We also thank Debora Nozza for her insightful comments. We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support. Dirk Hovy is a member of the Bocconi Institute for Data Science and Analytics (BIDSA). Carlo Schwarz is a member of the PERICLES and IGIER research centres at Bocconi as well as a Research Affiliate of the CEPR.

Ethical Considerations

We anonymized Twitter handles as part of the preprocessing to ensure GDPR compliance with at least pseudo-anonymity (motivated intruders are still likely to be able to recover identities with some effort, given that the data is in the public domain). In the reproduction data, tweet IDs are only shared in the dehydrated form in compliance with Twitter's Developer Policies: "Individuals redistributing tweet IDs and/or User IDs on behalf of an academic institution for the sole purpose of non-commercial research are permitted to redistribute an unlimited number of tweet IDs and/or User IDs." <https://developer.twitter.com/en/developer-terms/more-on-restricted-use-cases>. Hydrated and non-anonymized data will only be made available for research purposes upon request. We are aware that the non-anonymized data could be used to target individuals. As all data are already in the public domain, the additional risk is minimal, and we believe that the value of the data for the scientific community outweighs the risk.

Limitations

There are two minor limitations to the external validity of our findings. First, while we confirm the representativeness of our dataset relative to the overall Italian Twitter population, our sample is likely not a random set of Italian Twitter users. Further, as we label gender based on names, our gender coding is binary and might not fully represent the gender expression of users. We nonetheless believe that the DADIT data hold great promise for the study of attitudes and beliefs of the Italian population, as Twitter data are never representative of the overall population in any case.

A second limitation regards the multimodal potential of DADIT: we trained text classifiers on it,

but only ran pre-trained vision classifiers in inference mode. Finetuning vision models on our data could bring further classification improvements, although likely only to a limited extent, as they have already been trained on larger datasets. Perhaps more interestingly, future works could achieve performance improvements with a stronger ensemble method, i.e., training a classification head on top of the image and text representations output by the vision and text models.

Data and Code availability

We publicly release DADIT and our best gender and age classifiers through Hugging Face, at this address: <https://huggingface.co/lorelupo>. The software to train and run the classifiers, or to replicate our results is available at https://github.com/lorelupo/twitter_user_classification.

6. Bibliographical References

- Faiyaz Al Zamal, Wendy Liu, and Derek Ruths. 2012. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 6.1, pages 387–390.
- Righi Alessandra, Mauro M Gentile, and Domenico M Bianco. 2019. Who tweets in italian? demographic characteristics of twitter users. In *New Statistical Developments in Data Science: SIS 2017, Florence, Italy, June 28-30*, pages 329–344. Springer.
- Jennifer Allen, Cameron Martel, and David G Rand. 2022. Birds of a feather don't fact-check each other: Partisanship and the evaluation of news in twitter's birdwatch crowdsourced fact-checking program. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Jalal S Alowibdi, Ugo A Buy, and Philip Yu. 2013. Language independent gender classification on twitter. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*, pages 739–743.
- Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political analysis*, 23(1):76–91.
- Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2021. Xlm-t: Multilingual language models in twitter for sentiment analysis and beyond. *arXiv preprint arXiv:2104.12250*.

- Krishna C Bathina, Marijn Ten Thij, Lorenzo Lorenzo-Luaces, Lauren A Rutter, and Johan Bollen. 2021. Individuals with depression express more distorted thinking on social media. *Nature human behaviour*, 5(4):458–466.
- Federico Bianchi, Vincenzo Cutrona, and Dirk Hovy. 2022. Twitter-demographer: A flow-based tool to enrich twitter data. *arXiv preprint arXiv:2201.10986*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Grégoire Burel, Tracie Farrell, and Harith Alani. 2021. Demographics and topics impact on the co-spread of covid-19 misinformation and fact-checks on twitter. *Information Processing & Management*, 58(6):102732.
- Rochana Chaturvedi and Sugat Chaturvedi. 2023. It’s all in the name: A character-based approach to infer religion. *Political Analysis*, pages 1–16.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Aron Culotta. 2014. Estimating county health statistics with twitter. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1335–1344.
- Thomas Fujiwara, Karsten Müller, and Carlo Schwarz. 2023. The effect of social media on elections: Evidence from the united states. *Forthcoming Journal of the European Economic Association*.
- Kim Holmberg and lina Hellsten. 2015. Gender differences in the climate change communication on twitter. *Internet research*, 25(5):811–828.
- Dirk Hovy. 2015. Demographic factors improve classification performance. 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 752–762.
- Dirk Hovy and Diyi Yang. 2021. The importance of modeling social factors of language: Theory and practice. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 588–602.
- Kokil Jaidka, Salvatore Giorgi, H Andrew Schwartz, Margaret L Kern, Lyle H Ungar, and Johannes C Eichstaedt. 2020. Estimating geographic subjective well-being from twitter: A comparison of dictionary and data-driven language methods. *Proceedings of the National Academy of Sciences*, 117(19):10165–10171.
- Gil Levi and Tal Hassner. 2015. Age and gender classification using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 34–42.
- Yitong Li, Timothy Baldwin, and Trevor Cohn. 2018. Towards robust and privacy-preserving text representations. *arXiv preprint arXiv:1805.06093*.
- Jeffrey W Lockhart, Molly M King, and Christin Munsch. 2023. Name-based demographic inference and the unequal distribution of misrecognition. *Nature Human Behaviour*, pages 1–12.
- Lorenzo Lupo, Oscar Magnusson, Dirk Hovy, Elin Naurin, and Lena Wängnerud. 2023. [How to use large language models for text coding: The case of fatherhood roles in public policy documents](#).
- Veronica Lynn, Youngseo Son, Vivek Kulkarni, Niranjan Balasubramanian, and H Andrew Schwartz. 2017. Human centered nlp with user-factor adaptation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1146–1155.
- Sandra C Matz, Jochen I Menges, David J Stillwell, and H Andrew Schwartz. 2019. Predicting individual-level income from facebook profiles. *PloS one*, 14(3):e0214369.
- Pew. 2022. Jobs, hobbies top the list of things u.s. adults put in their twitter profiles; references to politics relatively rare, by regina widjaya. <https://www.pewresearch.org/short-reads/2022/05/05/jobs-hobbies-top-the-list-of-things-u-s-adults-put-in-their-twitter-profiles-references-to-politics-relatively-rare/>.
- Flor Miriam Plaza-del Arco, Debora Nozza, and Dirk Hovy. 2023. Leveraging label variation in large language models for zero-shot text classification. *arXiv preprint arXiv:2307.12973*.
- Daniel Preoțiu-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.
- Daniel Preoțiu-Pietro and Lyle Ungar. 2018. User-level race and ethnicity predictors from twitter

- text. In *Proceedings of the 27th international conference on computational linguistics*, pages 1534–1545.
- Statista. 2023. Social media: Twitter users in Italy. <https://www.statista.com/study/73547/social-media-twitter-users-in-italy/>.
- Zachary C Steinert-Threlkeld. 2018. *Twitter as data*. Cambridge University Press.
- Nikki Usher, Jesse Holcomb, and Justin Littman. 2018. Twitter makes it worse: Political journalists, gendered echo chambers, and the amplification of gender bias. *The international journal of press/politics*, 23(3):324–344.
- Svitlana Volkova, Theresa Wilson, and David Yarowsky. 2013. Exploring demographic language variations to improve multilingual sentiment analysis in social media. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1815–1827.
- Zijian Wang, Scott Hale, David Ifeoluwa Adelani, Przemyslaw Grabowicz, Timo Hartman, Fabian Flöck, and David Jurgens. 2019. Demographic inference and representative population estimates from multilingual social media data. In *The world wide web conference*, pages 2056–2067.
- Quanzeng You, Sumit Bhatia, Tong Sun, and Jiebo Luo. 2014. [The eyes of the beholder: Gender prediction using images posted in online social networks](#). In *2014 IEEE International Conference on Data Mining Workshop*, pages 1026–1030.
- Yipeng Zhang, Hanjia Lyu, Yubao Liu, Xiyang Zhang, Yu Wang, and Jiebo Luo. 2021. Monitoring depression trends on twitter during the covid-19 pandemic: observational study. *JMIR infodemiology*, 1(1):e26769.

A. Dataset Details

Italian Data The data collection of the Italian Twitter data started from a list of all Italian senators and deputies in the XVIII Legislature (2018-2022). We manually searched for Twitter accounts for each of the politicians. It is worth noting that a few politicians do not have a Twitter account, while less than a handful have more than one account associated with their name. In total, we obtained Twitter accounts for 802 Italian politicians. Table 4 provides an overview of the number of accounts that were collected for each of the Italian parties. In total, these accounts are followed by 3.2 million unique users, for which we then collected user-profiles and tweets.

Party	# Accounts	# Followers
Movimento 5 Stelle	191	2,548,834
Lega	172	2,413,691
Partito Democratico	131	5,209,329
Forza Italia	97	1,441,865
Misto	74	3,013,421
Fratelli d'Italia	55	1,526,940
Italia Viva	39	4,728,160
Coraggio Italia	16	29,787
Liberi e Uguali	8	1,051,836
Indipendente	7	479,298
C.A.L.	6	32,622
Per le Autonomie	3	166,589
Azione	1	349,375
Europa Verde	1	9,969
Südtiroler Volkspartei	1	3,033

Table 4: Summary of Italian source accounts.

German Data The German dataset results from a similar data collection as the Italian data. The authors collected the Twitter followers of eight German parties (CDU, CSU, SPD, FDP, Greens, The Left, AfD, NPD). Table 5 lists the number of followers for each of the collected accounts. We then collected the user profiles and tweets of the followers. In total, we obtained around 965,000 user profiles. Afterward, we label users' gender based on the "full name" and age based on the user bios. This leaves us with 388 user profiles for which we have both gender and age coding. For this subsample, we then collected user images.

A.1. Representativeness

We evaluate the representativeness of the Italian data relative to the demographic characteristics of Italian Twitter users. Unfortunately, there is no census of Italian Twitter users from which we could obtain the gender and age of the universe of Italian Twitter users. Instead, we rely on information

Party	# Followers
CDU	297,417
CSU	196,119
SPD	370,956
FDP	350,234
Greens	488,959
The Left	276,713
AfD	147,942
NPD	5,184

Table 5: Summary of German source accounts.

obtained from research by [Alessandra et al. \(2019\)](#) and a report from [Statista \(2023\)](#). The former provides information on the gender and location of Twitter users, while the latter contains information on the age of the Twitter user base. As the report by [Statista \(2023\)](#) splits the age of users by generation, we convert the age of the users to the same generations. Note that these generations differ from the age bins we use in the main paper.

Overall, the demographic characteristics of users in our data align well with existing work (see Figure 2). 63% of our users are male relative to 67% in the work by [\(Alessandra et al., 2019\)](#). Also, the age distribution of our users aligns closely with the reported age distributions of generations in [\(Statista, 2023\)](#). The main difference is that we appear to observe more baby boomers among Twitter users in our data. This actually makes the age distribution of our Twitter users more representative of the age distribution of the Italian population overall.

B. Experimental Details

Model	# Tokens
M3	na
XLM	512
Flan-T5	512
GPT3.5	2048

Table 6: Models' max input lengths used in the experiments.

M3 We employed the PyTorch implementation of the M3 model to infer the age and gender of Twitter users [Wang et al. \(2019\)](#). M3 uses names, usernames, biographies, and profile images, but it does not utilize tweet content. For our analysis, we specifically excluded usernames from the input, replacing both names and usernames with empty strings, as detailed in our discussion. An essential requirement for M3 is the language of the user's biography; for this, we utilized the *langdetect* package to identify the bio's language within M3's




Figure 2: Representativeness of Twitter users in DADIT.

supported range. If the detected language was unsupported by M3 (rear cases), Italian was chosen as the default language. The model outputs probabilities for different gender and age groups. Our approach was to consider the group with the highest probability as the definitive prediction made by the M3 model.

Computer vision model We apply a popular open-source computer vision classifier (Levi and Hassner, 2015) for inferring age and gender from users' profile pictures. In a first step, the classifier identifies faces in the images and then classifies each face in terms of gender and age. The original application of the model uses the following age categories: [0, 2], [4, 6], [8, 12], [15, 20], [25, 32], [38, 43], [48, 53], [60, 100]. We convert these to the classes used by M3 by aggregating predicted probabilities for appropriate classes. If the profile image shows no face, we assign equal probability to all classes and then use the majority class ("male" for gender, "above 40" for age) as default. Finally, if there are multiple faces on the image, we assign

the average predicted age and gender to the user.

XMLM For our classification task, we utilized the twitter-XLM-roberta-base model (Barbieri et al., 2021), a multilingual language model specifically fine-tuned on Twitter data. We trained two separate classification models for gender and age with the same fine-tuning settings for both models. Initially, we allocated 10% of our training data as a validation set. The fine-tuning process employed the Adam optimizer, with an initial learning rate set at 2e-5 and a linear learning rate scheduler. The initial two epochs focused on warming up the gradients by training only the classification head. Following this, we unfroze the entire model for further training. We employed an early stopping mechanism, using F1 score as the criterion for improvement. Training was halted if there was no improvement in the F1 score on the validation set after three epochs. The model iteration that achieved the best F1 score on the validation set was retained as the final model.

Flan-T5
Classify the gender (male or female) of the author of these texts: Bio: [input bio] Tweets: [input tweets] Gender:

Classify the age of the author of the following texts across 4 age groups: "0-19", "20-29", "30-39", "40-100". Texts: Bio: [input bio] Tweets: [input tweets] Age group:
GPT3.5
Classify the gender of the author of these texts (only answer with the label - "male" or "female" - even if uncertain!): [few shot examples] Bio: [input bio] Tweets: [input tweets] Gender:

Guess the age of the author of the following texts across 4 age groups: "0-19", "20-29", "30-39", "40-100". Only answer with an age group, even if uncertain. [few shot examples] Bio: [input bio] Tweets: [input tweets] Age group:

Table 7: Language models’ prompts for gender and age classification. Flan-T5 was prompted in a zero-shot setting, GPT3.5 in a few-shot settings (8 exemplars).

# tokens	# shots	F1	Acc
512	0	75.2	76.1
1028	0	76.4	77.2
2056	0	76.9	77.7
2056	8	77.4	78.0

Table 8: GPT3.5 performance on Italian gender prediction with varying context lengths (# tokens) and exemplars (# shots). The longer the context, the more tweets from the user are included.

Flan-T5 We experiment in a zero-shot setting with Flan-T5 because, in preliminary tests, it did not benefit from including exemplars in the prompt. When adding tweets as a feature, we do so until we saturate the maximum input length of 512 tokens. The prompts used are shown in Table 7. Even if the input text is in Italian, prompts are in English as the model has been trained mostly with English instructions. Indeed, Italian prompts performed less well in our preliminary experiments for both Flan-T5 and GPT3.5. The generated output is parsed to find patterns corresponding to one of the labels. In case multiple labels are generated, we only consider the first one, while we assign the majority class label as the default in case nothing valid is generated.

GPT3.5 In this case, we experiment in a few-shot setting, which proved to be significantly superior to zero-shot classification. When adding tweets as a feature, we do so until we saturate the maxi-

imum input length of 512 tokens. The prompts used are shown in Table 7. The model benefits from longer context and few-shots, as shown in Table 8. Although gpt-3.5-turbo has a maximum context length of 4096, we decided to run prompts with 2048 tokens maximum as a compromise between performance, monetary cost, and time. In fact, when adding tweets up to 4096 tokens, we needed to decrease the frequency of API requests so as not to exceed the maximum number of tokens allowed per minute. This would result in a significant slowdown of the classification process for a non-significant increase in performance. As in the case of Flan-T5, the generated output is parsed to find patterns corresponding to one of the labels. In case multiple labels are generated, we only consider the first one, while we assign the majority class label as the default in case nothing valid is generated.

When adding tweets as a feature, we do so until we saturate the maximum input length of 512 tokens. The generated output is parsed to find patterns corresponding to one of the labels. In case multiple labels are generated, we only consider the first one, while we assign the majority class label as the default in case nothing valid is generated.