

Improving Stance Detection by Leveraging Measurement Knowledge from Social Sciences: A Case Study of Dutch Political Tweets and Traditional Gender Role Division

Qixiang Fang^[0000-0003-2689-6653], Anastasia Giachanou^[0000-0002-7601-8667],
and Ayoub Bagheri^[0000-0001-6366-2173]

Department of Methodology and Statistics, Utrecht University, Utrecht, Netherlands
{q.fang, a.giachanou, a.bagheri}@uu.nl

Abstract. Stance detection concerns automatically determining the viewpoint (i.e., in favour of, against, or neutral) of a text’s author towards a target. Stance detection has been applied to many research topics, among which the detection of stances behind political tweets is an important one. In this paper, we apply stance detection to a dataset of tweets from official party accounts in the Netherlands between 2017 and 2021, with a focus on stances towards traditional gender role division, a dividing issue between (some) Dutch political parties. To implement and improve stance detection of traditional gender role division, we propose to leverage an established survey instrument from social sciences, which has been validated for the purpose of measuring attitudes towards traditional gender role division. Based on our experiments, we show that using such a validated survey instrument helps to improve stance detection performance.

Keywords: Zero-shot learning · Bertje · Reliability.

1 Introduction

Stance detection is an important natural language processing (NLP) task. Stance detection aims to automatically determine the viewpoint of a text’s author towards a target [16]. This viewpoint, also known as stance, can be *in favour of*, *against* or *neutral*. Typical targets include various social and political topics (e.g., abortion, feminism), political figures (e.g., Trump, Clinton) and events (e.g., referendums) [11,12,1]. Stance detection has been extensively applied to various research areas, including political discourse analysis [8,12]. Political tweets, in particular, serve as a rich source of information for understanding political stances and attitudes, making stance detection particularly relevant in this domain.

Detecting stances behind political tweets is a multifaceted task that involves analysing the linguistic nuances and contextual factors embedded within tweets

to determine the viewpoint expressed towards a specific issue or target. Traditionally, this process has been approached using supervised learning methods, where a classifier is trained on a dataset of labelled tweets, where each tweet is associated with a predefined stance category [13,17,12]. Unsupervised learning methods for stance detection also exist, which do not rely on labelled data, but instead, seek to identify latent patterns and structures within a dataset of unlabelled tweets [7,4,20]. They typically employ techniques such as topic modelling, sentiment analysis, and distributional semantics to extract latent features from the tweets. These features are then analysed to identify clusters or groupings of tweets that share similar linguistic characteristics and semantic associations. These clusters are then assigned stance labels based on their perceived sentiment or topic orientation.

Recently, stance detection is increasingly conceptualised as textual entailment recognition (TER) [1,16]. TER, a more generic natural language understanding task, concerns determining whether a hypothesis follows from a premise, where three outcomes are possible: “entailment”, “contradiction” and “neutral” [3]. By reformulating stance detection as judging whether a text (i.e., a premise) entails a stance towards a target (i.e., a hypothesis), a stance detection problem becomes a TER one. This reformulation offers two potential benefits. First, TER-based models are not target-specific and may be applied to unseen targets. Second, TER datasets can be used for training stance detection models. Currently, labelled stance detection datasets are limited to a few pre-defined targets and languages [1]. In contrast, TER datasets are not target-specific and are available in more languages. Therefore, we can use TER datasets (in addition to stance detection ones) for model training when considering stance detection as TER. Research has shown that training on additional TER datasets improves downstream stance detection performance [18].

In this paper, we investigate the application of stance detection as TER to a dataset of tweets from political parties and politicians in the Netherlands between 2017 and 2021, focusing specifically on stances towards traditional gender role division. This issue holds significant political and social implications, often dividing various Dutch political parties. To further improve stance detection performance and capture nuanced stances, we propose utilising a validated survey instrument from social sciences specifically designed to measure attitudes towards traditional gender role division. Our experiments demonstrate that incorporating this instrument indeed leads to improved stance detection performance. To the best of our knowledge, our study is the first to incorporate survey instruments for stance detection.

2 Background

2.1 Stance Detection as Textual Entailment Recognition

TER, as mentioned in section 1, concerns judging whether a hypothesis (h) can be inferred from a premise (p). If yes, we conclude that p entails h ($p \Rightarrow h$).

Typically, three outcomes are possible: “entailment”, “contradiction” and “neutral”. Note that the relation between p and h is directional, where $p \Rightarrow h$ does not imply $h \Rightarrow p$.

The connection between the two tasks, TER and stance detection, becomes clear when we reformulate stance detection as judging whether a text (t) entails a stance to a target (i.e., $t \Rightarrow stt$). For instance, say the text t is “I appreciate more women in government positions”, the target is “traditional gender role division”, and the stance is “in favour of”. Then, we can formulate the stance to the target stt as “I am in favour of traditional gender role division”. As such, the task becomes judging whether “I appreciate more women in government positions” *entails, contradicts or is neutral* towards “I am in favour of traditional gender role division”. The three possible TER outcomes “entailment”, “contradiction” and “neutral” correspond to the three stance detection outcomes: “in favour of”, “against” and “neutral”.

One benefit of considering stance detection as TER is that a model can be easily trained on data involving multiple targets and can incorporate information about the targets during training. This may allow the model to predict stances towards unseen targets, also known as zero-shot learning [26]. Another benefit is that TER datasets can be used for model training. For languages where existing stance detection datasets are available, additional TER datasets may help to further improve model performance. When stance detection datasets do not exist for a language but TER datasets do, we can still train a model based on TER datasets to predict stances. It is, however, important to note that current TER datasets are quite different from stance detection datasets [2,25]. The former mainly consists of descriptions of concrete scenes (e.g., “two dogs are fighting”) which assume little world knowledge (e.g., social and political topics), while the latter are typically user-generated texts (e.g., social media posts) and concern diverse targets that can be abstract and ambiguous. This data discrepancy may limit the benefit of modelling stance detection as TER.

2.2 Stance Detection as Attitudinal Measurements in Social Sciences

In social sciences, stance detection is treated as a measurement problem and typically relies on the use of survey instruments. For instance, to measure stances towards traditional gender role division, one can ask a participant to indicate on a Likert scale from 1 to 5 how much they agree with the following statement: “*Overall, family life suffers the consequences if the mother has a full-time job.*” [19]. The more they agree with this statement, the more they are in favour of the target (i.e., traditional gender role division).

Survey instruments need to be both “*valid*” and “*reliable*”. “Valid” means that they measure what they are supposed to, while “reliable” means that the measurements do not suffer from large random variations [21]. One common way to establish the validity of survey instruments is to ensure that they cover all the aspects of a given target. This is also known as “content validity” [21]. For instance, when measuring stances to traditional gender role division, it is

important to cover relevant subdomains of this target like in career, childcare, housework and education [23]. Furthermore, higher reliability can be achieved by averaging the responses to multiple survey instruments concerning the same target.

Established survey instruments exist for many stance detection targets (e.g., feminism [6], abortion [9], climate change [5] and politicians [15]). They can potentially benefit the modelling of stance detection as TER. For instance, when using a model trained on TER datasets to detect stances towards traditional gender role division, the hypothesis can be as simple as “I am in favour of traditional gender role division”. This hypothesis, however, does not explicitly cover all the relevant subdomains of the target (i.e., in career, childcare, housework and education). Consequently, the model may fail to identify texts that entail these subdomain stances. We may mitigate this issue by using established survey items (that have demonstrated content validity) to construct multiple hypotheses. Averaging the model’s predictions across these multiple hypotheses can also lead to more reliable, less random results. Therefore, constructing multiple hypotheses based on established survey instruments for a given target may improve prediction validity and reliability.

3 Experiments

We investigate the approach of modelling stance detection as TER by applying it to a dataset of tweets from Dutch political parties and politicians in the Netherlands, with the goal to measure stances towards traditional gender role division.

3.1 Datasets

We use Twitter (now X)’s API to collect a dataset of tweets from the official Twitter accounts of 13 major Dutch political parties (see Table 1) and the accounts of 10 politicians per party. We collect up to 3,200 most recent tweets per account (API limit) and in total, 311,362 tweets. We lower-case the tweets, remove urls and special characters, leave out tweets shorter than five words, and retain only the ones posted between 2017 and 2021. The resulting dataset contains 247,902 tweets. To reduce the number of tweets unrelated to the target of interest, we further apply a simple filtering strategy using keywords related to the target,¹ leading to a much smaller final dataset of 2,601 tweets.

Previous research suggests that a party’s stance towards an issue aligns with the average opinion of its voters [10]. This inspires us to use the 2017-2021 data from the LISS panel² to provide ground truth labels on the party level. The LISS panel is a representative sample of Dutch households that participate in monthly surveys about various topics like health, work and political views.

¹ The Dutch equivalent of “woman”, “man”, “mother”, “father”, “boy” and “girl”: “vrouw”, “man”, “moeder”, “vader”, “jongen”, “meisje”.

² <https://www.lissdata.nl/about-panel>

The LISS data contains responses to 11 established survey items about stances towards traditional gender role division (see Section B), as well as the party for which the participant last voted. We group the participants by the party they voted for and average their responses according to each participant’s respective sampling weight across the 11 survey items for a given year and across all years. These scores are used to rank the parties in terms of their year(s)-specific stances towards traditional gender role division (see Table 1). For tweet-level ground truth labels, we rely on manual annotation (see Section 3.2 for more information, and Appendix C for annotated examples).

Table 1. Parties Ranked by LISS-based Scores on Stances towards Traditional Gender Role Division. The higher the scores, the more in favour of traditional gender role division. Note that these scores are averaged across the years 2017-2021.

Party	Score
SGP	2.86
DENK	2.56
PVV	2.47
PLUS50	2.38
CU	2.35
CDA	2.28
FVD	2.26
SP	2.16
VVD	2.14
PVDD	2.03
PVDA	1.98
D66	1.96
GROENLINKS	1.83

Furthermore, we use the SICK-NL dataset [24], which is the Dutch translation of the original English SICK dataset [14]. It contains 9,840 pairs of sentences that describe concrete scenes (e.g., “Two dogs are fighting”), including 5,595 neutral pairs, 1,424 contradiction pairs and 2,821 entailment pairs. To the best of our knowledge, SICK-NL is the only Dutch TER dataset. Also, no labelled stance detection dataset exists for the Dutch language.

3.2 Experiment 1: Zero-shot Stance Detection with BERTje

Due to the lack of labeled Dutch stance detection datasets, we choose to fine-tune a TER classifier on the available SICK-NL dataset, and subsequently apply the trained classifier to perform zero-shot prediction of stances towards traditional gender role division. Specifically, we fine-tune BERTje, a pretrained Dutch BERT model [22], on the SICK-NL dataset, using 9,345 pairs for training and the rest for validation. We use the same hyperparameters as [24]. See Appendix A for more computational details. Next, we apply the fine-tuned BERTje model to our

dataset of Dutch political tweets, where the goal is to predict whether a tweet entails a hypothesis (i.e., a stance towards traditional gender role division). We use two types of hypotheses. The first is a simple statement (in Dutch): “I am in favour of traditional gender role division”. The second consists of the 11 survey items about traditional gender role division from the LISS panel (also in Dutch). See Section B for more information about the construction of hypotheses.

We evaluate the stance detection performance at both the individual tweet level and the party level. Due to the large number of tweets in the dataset and lack of resources for manual annotation of all the tweets, we focus on evaluating only the top k tweets with the highest predicted entailment scores, where $k = 10, 50, 100$. The more out of the k tweets are, indeed, true entailments, the more evidence that the model has encoded information useful for (zero-shot) stance prediction.

For the tweet-level evaluation, we obtain the labels for the k tweets by asking two Dutch-speaking PhD researchers to manually annotate the tweets. Where inconsistent annotations occur, the annotators discuss them and agree on one correct label. Then, we calculate the percentages of correctly classified entailments among the k tweets (i.e., top- k precision) across two conditions (i.e., whether a simple hypothesis or survey items are used). We contrast these results with a random baseline, which is based on a random sample of 100 tweets from all the tweets in the final dataset. The same two PhD researchers from before manually annotate these 100 tweets. The percentage of entailments is used as the baseline precision score.

For the party-level evaluation, we compute the average entailment probabilities for each party and year combination. Then, we calculate the Spearman’s correlation (ρ) between the average entailment probabilities and the LISS-based ranks of parties (across year and party). Scores close to 1 indicate better agreement.

Table 2. Top- k precision scores for entailments, as well as top- k Spearman’s ρ between the estimated entailment probabilities and the LISS-based scores averaged across party and year. Top- k refers to the tweets with the top 10/50/100 predicted entailment probabilities. The random baseline is based on a random sample of 100 tweets. The two non-baseline conditions concern whether a simple hypothesis (i.e., Without Survey) or survey items (i.e., With Survey) are used.

	Random Baseline	Without Survey	With Survey
Precision ₁₀	0.11	0.10	0.70
Precision ₅₀	0.11	0.20	0.52
Precision ₁₀₀	0.11	0.22	0.42
Spearman’s ρ_{10}	-	-0.04	-0.07
Spearman’s ρ_{50}	-	0.06	-0.07
Spearman’s ρ_{100}	-	0.15	-0.08
Spearman’s ρ_{all}	-	0.13	-0.09

Results Table 2 summarises the results on top- k precision and Spearman’s ρ . For the rows pertaining to top- k precision, the scores are the percentages of tweets correctly classified as entailments. We can make two observations. First, on the tweet level, the best precision scores consistently occur when survey items are used, indicating that the use of survey items to construct multiple hypotheses can be beneficial for stance detection. Second, however, the Spearman’s ρ scores paint a different picture: all the scores are very far from 1, indicating poor performance on the party level. These findings suggest that our Bertje model encodes useful information for zero-shot prediction of an unseen stance target, but only when survey items are used.

3.3 Experiment 2: Zero-shot Stance Detection with GPT-3.5 Turbo

To further validate our findings from Experiment 1, we conduct a second experiment where we use prompts with GPT-3.5 Turbo for zero-shot stance detection of traditional gender role division. We use a balanced dataset of 200 political tweets that were labelled during the annotation procedure of Experiment 1. We use the OpenAI API with GPT-3.5 Turbo, with the following prompt template, where {premise} and {hypothesis} refer to two separate string variables:

““““

Given the following premise and hypothesis (both in Dutch), determine whether the premise entails, contradicts or is neutral about the hypothesis.

Premise: {premise}

Hypothesis: {hypothesis}

Note that the premise is an official tweet by a political party in the Netherlands, while the hypothesis is related to a positive stance towards traditional gender role division. If the relationship between the premise and the hypothesis is entailment, this means that the political party behind the tweet/premise likely agrees with the hypothesis.

Give your answer response as one of the three outcomes: “entails”, “contradicts” or “neutral”.

””””

We use an English template because in this case, GPT-3.5 Turbo gives responses that are consistent with the requirement of the prompt (i.e., returning only “entails”, “contradicts” or “neutral”). In contrast, using a Dutch template would throw off the model, leading it to return sentences (e.g., “The premise entails the hypothesis”; “There is no relationship between the premise and the hypothesis”) and sometimes, responses in other languages like English. To improve reproducibility, we set the temperature parameter to be 0 and the random seed to 1.

Results Using only a simple hypothesis, we obtain an accuracy score of 0.58; using all the survey items, however, we obtain a higher accuracy score of 0.725. This further suggests evidence for the efficacy of using survey items for stance detection.

4 Conclusion & Discussion

In this paper, we apply stance detection to a dataset of tweets from official party accounts in the Netherlands between 2017 and 2021, with a focus on detecting stances towards traditional gender role division. We also investigate whether using established survey instruments can be helpful for stance detection.

In Experiment 1, we are faced with the challenge of lacking labelled stance detection datasets in Dutch. We thus opt to reformulate our stance detection task as TER, where we train a BERTje model on a Dutch TER dataset and subsequently apply it to zero-shot stance detection. We find using survey items consistently outperforms using a simple hypothesis for tweet-level stance detection. However, our model fails to achieve good stance detection on the party level. In Experiment 2, we further test the benefit of using survey instruments by repeating the tweet-level analysis from Experiment 1 using prompts with OpenAI’s GPT-3.5 Turbo API. We again find using survey questions beneficial for stance detection of traditional gender role division.

Our study has limitations. For instance, we focus on a single stance detection target and only a subset of the political tweet dataset. Including more stance detection targets and more datasets is necessary to further support the generalisability of our findings. Nevertheless, we hope that our study will inspire stance detection researchers to explore the use of survey instruments for stance detection research further. Furthermore, our second experiment, which utilised GPT-3.5, is now somewhat outdated given the recent emergence of more capable large language models like GPT-4 and GPT-4o. We also do not incorporate with other prompting engineering techniques such as chain-of-thought and few-shot prompting, which have the potential to improve the model’s task performance.

For future research, we suggest, among others, converting existing stance detection datasets into TER-compatible formats for model training, developing more general TER datasets (beyond concrete scenes), using multilingual models (where datasets in different languages may be combined), and comparing different strategies for constructing TER hypotheses from survey items.

Acknowledgments. We thank Dong Nguyen for her constructive and helpful feedback. This work was partially supported by the Dutch Research Council (NWO): grant number: VI.Vidi.195.152 to D. L. Oberski.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aldayel, A., Magdy, W.: Stance detection on social media: State of the art and trends. *Information Processing & Management* **58**(4), 102597 (2021). <https://doi.org/10.1016/J.IPM.2021.102597>

2. Bowman, S.R., Angeli, G., Potts, C., Manning, C.D.: A large annotated corpus for learning natural language inference. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. pp. 632–642. Association for Computational Linguistics, Lisbon, Portugal (2015). <https://doi.org/10.18653/v1/D15-1075>
3. Dagan, I., Glickman, O.: Probabilistic textual entailment: Generic applied modeling of language variability. *Learning Methods for Text Understanding and Mining* **2004**(26-29), 2–5 (2004)
4. Darwish, K., Stefanov, P., Aupetit, M., Nakov, P.: Unsupervised user stance detection on Twitter. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 14, pp. 141–152 (2020), <https://ojs.aaai.org/index.php/ICWSM/article/download/7286/7140/10516>
5. Dijkstra, E., Goedhart, M.: Development and validation of the ACSI: Measuring students’ science attitudes, pro-environmental behaviour, climate change attitudes and knowledge. *Environmental Education Research* **18**(6), 733–749 (2012). <https://doi.org/10.1080/13504622.2012.662213>
6. Frieze, I.H., McHugh, M.C.: Measuring feminism and gender role attitudes. *Psychology of Women Quarterly* **22**(3), 349–352 (1998). <https://doi.org/10.1111/j.1471-6402.1998.tb00159.x>
7. Gambini, M., Senette, C., Fagni, T., Tesconi, M.: From tweets to stance: An unsupervised framework for user stance detection on Twitter. In: International Conference on Discovery Science. pp. 96–110. Springer (2023). https://doi.org/10.1007/978-3-031-45275-8_7
8. Gunhal, P., Bashyam, A., Zhang, K., Koster, A., Huang, J., Haresh, N., Singh, R., Lutz, M.: Stance detection of political tweets with transformer architectures. In: 2022 13th International Conference on Information and Communication Technology Convergence (ICTC). pp. 658–663. IEEE (2022). <https://doi.org/10.1109/ICTC55196.2022.9952951>
9. Hendriks, J., Fyfe, S., Styles, I., Skinner, S.R., Merriman, G.: Scale construction utilising the rasch unidimensional measurement model: A measurement of adolescent attitudes towards abortion. *The Australasian Medical Journal* **5**(5), 251 (2012). <https://doi.org/10.4066/AMJ.2012.952>
10. Ibenskas, R., Polk, J.: Congruence and party responsiveness in western europe in the 21st century. *West European Politics* **45**(2), 201–222 (2022). <https://doi.org/10.1080/01402382.2020.1859756>
11. Lahoti, P., Garimella, K., Gionis, A.: Joint non-negative matrix factorisation for learning ideological leaning on Twitter. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. pp. 351–359 (2018). <https://doi.org/10.1145/3159652.3159669>
12. Lai, M., Cignarella, A.T., Farías, D.I.H., Bosco, C., Patti, V., Rosso, P.: Multilingual stance detection in social media political debates. *Computer Speech & Language* **63**, 101075 (2020). <https://doi.org/10.1016/j.csl.2020.101075>
13. Li, Y., Sosea, T., Sawant, A., Nair, A.J., Inkpen, D., Caragea, C.: P-stance: A large dataset for stance detection in political domain. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021. pp. 2355–2365 (2021). <https://doi.org/10.18653/v1/2021.findings-acl.208>
14. Marelli, M., Menini, S., Baroni, M., Bentivogli, L., Bernardi, R., Zamparelli, R.: A SICK cure for the evaluation of compositional distributional semantic models. In: Proceedings of the Ninth International Conference on Language Resources

- and Evaluation (LREC'14). pp. 216–223. European Language Resources Association (ELRA), Reykjavik, Iceland (2014), http://www.lrec-conf.org/proceedings/lrec2014/pdf/363_Paper.pdf
15. Martin, E.A., Traugott, M.W., Kennedy, C.: A review and proposal for a new measure of poll accuracy. *Public Opinion Quarterly* **69**(3), 342–369 (2005). <https://doi.org/10.1093/poq/nfi044>
 16. Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: SemEval-2016 Task 6: detecting stance in tweets. In: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. pp. 31–41. Association for Computational Linguistics (2016). <https://doi.org/10.18653/v1/S16-1003>
 17. Samih, Y., Darwish, K.: A few topical tweets are enough for effective user stance detection. In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. pp. 2637–2646 (2021). <https://doi.org/10.18653/v1/2021.eacl-main.227>
 18. Schiller, B., Daxenberger, J., Gurevych, I.: Stance detection benchmark: How robust is your stance detection? *KI-Künstliche Intelligenz* **35**, 329–341 (2021). <https://doi.org/10.1007/s13218-021-00714-w>
 19. Stam, K., Verbakel, E., de Graaf, P.M.: Do values matter? The impact of work ethic and traditional gender role values on female labour market supply. *Social Indicators Research* **116**, 593–610 (2014). <https://doi.org/10.1007/s11205-013-0287-x>
 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 (2020). <https://doi.org/10.18653/v1/2020.acl-main.50>
 21. Trochim, W.M.K., Donnelly, J.P., Arora, K.: *Research Methods: The Essential Knowledge Base*. Cengage Learning, Boston (2015)
 22. de Vries, W., van Cranenburgh, A., Bisazza, A., Caselli, T., Noord, G.v., Nissim, M.: BERTje: A Dutch BERT model. *arXiv* (2019), <http://arxiv.org/abs/1912.09582>
 23. Walter, J.G.: The adequacy of measures of gender roles attitudes: A review of current measures in omnibus surveys. *Quality & Quantity* **52**, 829 – 848 (2018). <https://doi.org/10.1007/s11135-017-0491-x>
 24. Wijnholds, G., Moortgat, M.: SICK-NL: A dataset for dutch natural language inference. In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Online (2021). <https://doi.org/10.18653/v1/2021.eacl-main.126>
 25. Williams, A., Nangia, N., Bowman, S.: A broad-coverage challenge corpus for sentence understanding through inference. In: Walker, M., Ji, H., Stent, A. (eds.) *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 1112–1122. Association for Computational Linguistics, New Orleans, Louisiana (2018). <https://doi.org/10.18653/v1/N18-1101>
 26. Yin, W., Hay, J., Roth, D.: Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. pp. 3914–3923. Association for Computational Linguistics, Hong Kong, China (2019). <https://doi.org/10.18653/v1/D19-1404>

A Additional Reproducibility Information

Code and Data Availability All the code and (anonymised) datasets will be made public on GitHub upon paper publication.

Computing Infrastructure All analyses were done on one of the researchers’ personal laptop, which runs on an 11th Gen Intel(R) Core(TM) i7-11800H processor with 32 GB of RAM.

Python 3.9, PyTorch 1.10.2, Huggingface Transformers 4.16.2 and CUDA 11.3 were used for finetuning BERTje and the downstream stance detection. R 4.1.2 and Tidiverse 1.3.1 were used for data cleaning and wrangling.

Runtime The runtime for finetuning the BERTje model is about four GPU hours. The runtime for the downstream stance detection is about three GPU hours. The runtime for the remaining analyses is negligible.

Number of Parameters The total number of parameters in BERTje is 109,139,715.

Validation Performance The best, selected finetuned BERTje model has a loss of 0.41114.

Hyperparameter Search For the BERTje model, we use the same hyperparameters and search strategy as [24], including:

- num_train_epochs=20
- per_device_train_batch_size=16
- per_device_eval_batch_size=64
- warmup_steps=250
- weight_decay=0.01

Among the 20 epochs, we select the one that has the lowest validation loss.

Cost of OpenAI’s GPT-3.5 Turbo API 1.13 US dollars.

B The Survey Instruments and Hypotheses

The following survey items from LISS are used to construct the hypotheses in TER. English translation are in *italics*:

1. Een werkende moeder kan niet zo’n warme en hechte relatie met haar kinderen hebben als een moeder die niet werkt. *A working mother’s relationship with her children cannot be just as close and warm as that of a non-working mother*
2. Een kind dat nog niet naar school gaat zal er waarschijnlijk onder lijden als zijn of haar moeder werkt. *A child that is not yet attending school is likely to suffer the consequences if his or her mother has a job.*
3. Al met al lijdt het gezinsleven er onder als de vrouw een volledige baan heeft. *Overall, family life suffers the consequences if the mother has a full-time job.*

4. Alleen de man moet bijdragen aan het gezinsinkomen. *Only men should contribute to the family income.*

5. De man moet het geld verdienen, de vrouw moet voor het huishouden en het gezin zorgen. *The father should earn money, while the mother takes care of the household and the family.*

6. Mannen zouden niet een groter deel van het huishoudelijk werk moeten doen dan nu het geval is. *Fathers ought not to do more in terms of household work than they do at present.*

7. Mannen zouden niet meer moeten doen aan de verzorging van de kinderen dan nu het geval is. *Fathers ought not to do more in terms of childcare than they do at present.*

8. Een vrouw is geschikter om kleine kinderen op te voeden dan een man. *A woman is more suited to rearing young children than a man.*

9. Voor een meisje is het eigenlijk toch niet zo belangrijk als voor een jongen om een goede schoolopleiding te krijgen. *It is actually less important for a girl than for a boy to get a good education.*

10. Jongens kun je in het algemeen vrijer opvoeden dan meisjes. *Generally speaking, boys can be reared more liberally than girls.*

11. Het is onnatuurlijk als vrouwen in een bedrijf leiding uitoefenen over mannen. *It is unnatural for women in firms to have control over men.*

Where survey items are not used, the following statement is adopted:

Ik ben voorstander van de traditionele rolverdeling tussen mannen en vrouwen. *I am in favour of traditional gender role division.*

C Examples of Tweets of Different Stances towards Traditional Gender Role Division

See below for examples of Tweets implying different stances towards traditional gender role division. English translation in *italics*.

Against traditional gender role division: “Ik zal me altijd ten volle inzetten voor de rechtvaardige strijd voor het vervolmaken van de gelijkwaardige positie van de vrouw in onze samenleving. Dan reken ik ook op de steun van alle mannen.” *I will always be fully committed to the just struggle to perfect the equal position of women in our society. Then I also count on the support of all men.*

In favour of traditional gender role division: “Zien we nog het voorrecht en het bijzondere van het moederschap? Ik zou bijna geen werk van een man kunnen bedenken wat zo hoog reikt als het moederschap.” *Do we still see the privilege and specialness of motherhood? I could hardly think of any work of a man that reaches as high as motherhood.*

Neutral towards or unrelated to traditional gender role division: “Opnieuw het misverstand dat legalisering van abortus tot meer abortus leidt. Het tegendeel is het geval. Legalisering leidt niet tot meer abortus, maar wel tot meer veilige abortus en minder sterfte onder vrouwen.” *Again the misunderstanding that legalisation of abortion leads to more abortion. The opposite is the case.*

Legalisation does not lead to more abortion, but to more safe abortion and less death among women.