



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

X-stance: A Multilingual Multi-Target Dataset for Stance Detection

Citation for published version:

Vamvas, J & Sennrich, R 2020, X-stance: A Multilingual Multi-Target Dataset for Stance Detection. in *Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS)*, 9, CEUR Workshop Proceedings, vol. 2624, CEUR-WS.org, 5th SwissText & 16th KONVENS Joint Conference 2020, Virtual conference, Switzerland, 23/06/20. <<http://ceur-ws.org/Vol-2624/>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



X-stance: A Multilingual Multi-Target Dataset for Stance Detection

Jannis Vamvas¹ Rico Sennrich^{1,2}

¹Department of Computational Linguistics, University of Zurich

²School of Informatics, University of Edinburgh

{vamvas, sennrich}@cl.uzh.ch

Abstract

We extract a large-scale stance detection dataset from comments written by candidates of elections in Switzerland. The dataset consists of German, French and Italian text, allowing for a cross-lingual evaluation of stance detection. It contains 67 000 comments on more than 150 political issues (*targets*). Unlike stance detection models that have specific target issues, we use the dataset to train a single model on all the issues. To make learning across targets possible, we prepend to each instance a natural question that represents the target (e.g. “Do you support X?”). Baseline results from multilingual BERT show that zero-shot cross-lingual and cross-target transfer of stance detection is moderately successful with this approach.

1 Introduction

In recent years many datasets have been created for the task of automated stance detection, advancing natural language understanding systems for political science, opinion research and other application areas. Typically, such benchmarks (Mohammad et al., 2016a) are composed of short pieces of text commenting on politicians or public issues and are manually annotated with their stance towards a *target entity* (e.g. Climate Change, or Trump). However, they are limited in scope on multiple levels (Küçük and Can, 2020).

First of all, it is questionable how well current stance detection methods perform in a cross-lingual setting, as the multilingual datasets avail-

able today are relatively small, and specific to a single target (Taulé et al., 2017, 2018). Furthermore, specific models tend to be developed for each single target or pair of targets (Sobhani et al., 2017). Concerns have been raised that cross-target performance is often considerably lower than fully supervised performance (Küçük and Can, 2020).

In this paper we propose a much larger dataset that combines multilinguality and a multitude of topics and targets. X-stance comprises more than 150 questions about Swiss politics and more than 67k answers given by candidates running for political office in Switzerland. Questions are available in four languages: English, Swiss Standard German, French, and Italian. The language of a comment depends on the candidate’s region of origin.

We have extracted the data from the voting advice application *Smartvote*. Candidates respond to questions mainly in categorical form (yes / rather yes / rather no / no). They can also submit a free-text comment to justify or explain their categorical answer. An example is given in Figure 1.

We transform the dataset into a stance detection task by interpreting the question as a natural-language representation of the target, and the commentary as the input to be classified.

The dataset is split into a multilingual training set and into several test sets to evaluate zero-shot cross-lingual and cross-target transfer. To provide a baseline, we fine-tune a multilingual BERT model (Devlin et al., 2019) on X-stance. We show that the baseline accuracy is comparable to previous stance detection benchmarks while leaving ample room for improvement. In addition, the model can generalize to a degree both cross-lingually and in a cross-target setting.

We have made the dataset and the code for reproducing the baseline models publicly available.¹

Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0)

¹<http://doi.org/10.5281/zenodo.3831317>

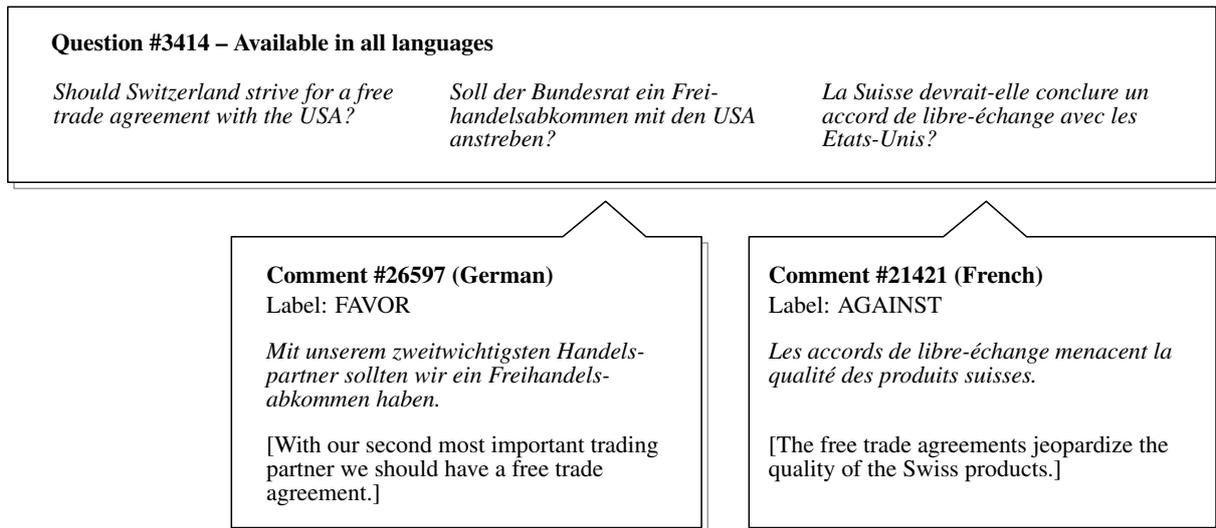


Figure 1: Example of a question and two answers in the X-stance dataset. The answers were submitted by electoral candidates on a voting advice website. The author of the German comment was in favor of the issue; the author of the French comment against. Both authors use comments to explain their respective stance.

2 Related Work

Multilingual Stance Detection In the context of the IberEval shared tasks, two related multilingual datasets have been created (Taulé et al., 2017, 2018). Both are a collection of annotated Spanish and Catalan tweets. Crucially, the tweets in both languages focus on the same issue (Catalan independence); given this fact they are the first truly multilingual stance detection datasets known to us.

With regard to the languages covered by X-stance, only monolingual datasets seem to be available. For French, a collection of tweets on French presidential candidates has been annotated with stance (Lai et al., 2020). Similarly, two datasets of Italian tweets on the occasion of the 2016 constitutional referendum have been created (Lai et al., 2018, 2020). With regard to German, a corpus of 270 sentences has been annotated with fine-grained stance and attitude information (Clematide et al., 2012). Furthermore, fine-grained stance detection has been qualitatively studied on a large corpus of Facebook posts (Klenner et al., 2017).

Multi-Target Stance Detection The SemEval-2016 task on detecting stance in tweets (Mohammad et al., 2016b) offers data concerning multiple targets (Atheism, Climate Change, Feminism, Hillary Clinton, and Abortion). In the supervised subtask A, participants tended to develop a *target-specific* model for each of those targets. In sub-

task B cross-target transfer to the target “Donald Trump” was tested, for which no annotated training data were provided. While this required the development of more universal models, their performance was generally much lower.

Sobhani et al. (2017) introduced a *multi-target* stance dataset which provides two targets per instance. For example, a model designed in this framework is supposed to simultaneously classify a tweet with regard to Clinton and with regard to Trump. While in theory the framework allows for more than two targets, it is still restricted to a finite and clearly defined set of targets. It focuses on modeling the dependencies of multiple targets within the same text sample, while our approach focuses on learning stance detection from many samples with many different targets.

Representation Learning for Stance Detection

In a target-specific setting, Ghosh et al. (2019) perform a systematic evaluation of stance detection approaches. They also evaluate BERT (Devlin et al., 2019) and find that it consistently outperforms previous approaches.

However, they only experiment with a single-segment encoding of the input, preventing cross-target transfer of the model. Augenstein et al. (2016) propose a conditional encoding approach to encode both the target and the tweet as sequences. They use a bidirectional LSTM to condition the encoding of the tweets on the encoding of the target, and then apply a nonlinear projection on

Topic	Questions	Answers
Digitisation	2	1168
Economy	23	6899
Education	16	7639
Finances	15	3980
Foreign Policy	16	4393
Immigration	19	6270
Infrastructure & Environment	31	9590
Security	20	5193
Society	17	6275
Welfare	15	8508
Total (training topics)	174	59 915
Healthcare	11	4711
Political System	9	2645
Total (held-out topics)	20	7356

Table 1: Number of questions and answers per topic.

the conditionally encoded tweet. This allows them to train a model that can generalize to previously unseen targets.

3 The x-stance Dataset

3.1 Task Definition

The input provided by x-stance is two-fold: (A) a natural language question concerning a political issue; (B) a natural language commentary on a specific stance towards the question.

The label to be predicted is either ‘favor’ or ‘against’. This corresponds to a standard established by [Mohammad et al. \(2016a\)](#). However, x-stance differs from that dataset in that it lacks a ‘neither’ class; all comments refer to either a ‘favor’ or an ‘against’ position. The task posed by x-stance is thus a binary classification task.

As an evaluation metric we report the macro-average of the F1-score for ‘favor’ and the F1-score for ‘against’, similar to [Mohammad et al. \(2016b\)](#). We use this metric mainly to strengthen comparability with the previous benchmarks.

3.2 Data Collection

Provenance We downloaded the questions and answers via the *Smartvote* API². The downloaded data cover 175 communal, cantonal and national elections between 2011 and 2020.

All candidates in an election who participate in *Smartvote* are asked the same set of questions, but

²<https://smartvote.ch>

depending on the locale they see translated versions of the questions. They can answer each question with either ‘yes’, ‘rather yes’, ‘rather no’, or ‘no’. They can supplement each answer with a comment of at most 500 characters.

The questions asked on *Smartvote* have been edited by a team of political scientists. They are intended to cover a broad range of political issues relevant at the time of the election. A detailed documentation of the design of *Smartvote* and the editing process of the questions is provided by [Thurman and Gasser \(2009\)](#).

Preprocessing We merged the two labels on each pole into a single label: ‘yes’ and ‘rather yes’ were combined into ‘favor’; ‘rather no’, or ‘no’ into ‘against’. This improves the consistency of the data and the comparability to previous stance detection datasets. We did not further preprocess the text of the comments.

Language Identification As the API does not provide the language of comments, we employed a language identifier to automatically annotate this information. We used the *langdetect* library ([Shuyo, 2010](#)). For each responder we classified all the comments jointly, assuming that responders did not switch code during the answering of the questionnaire.

We applied the identifier in a two-step approach. In the first run we allowed the identifier to output all 55 languages that it supports out of the box, plus Romansh, the fourth official language in Switzerland³. We found that no Romansh comments were detected and that all unexpected outputs were misclassifications of German, French or Italian comments. We further concluded that little or no Swiss German comments are in the dataset; otherwise, some of them would have manifested themselves via misclassifications (e.g. as Dutch).

In the second run, drawing from these conclusions, we restricted the identifier’s set of choices to English, French, German and Italian.

Filtering We pre-filtered the questions and answers to improve the quality of the dataset. To keep the domain of the data surveyable, we set a focus on national-level questions. Therefore, all

³Namely the Rumantsch Grischun variety; the language profile was created using resources from the Zurich Parallel Corpus Collection ([Graën et al., 2019](#)) and the *Quotidiana* corpus (<https://github.com/ProSvizraRumantscha/corpora>).

	Intra-target (New answers to known questions)		Cross-question (New questions within known topics)		Cross-topic	
DE	Train:	33 850	Test:	3143	Test:	5269
	Test:	2871				
	Valid:	3479				
FR	Train:	11 790	Test:	1170	Test:	1914
	Test:	1055				
	Valid:	1284				
IT	Test:	1173	Test:	(110)	Test:	(173)

Table 2: Number of answer instances in the training, validation and test sets. The *upper left corner* represents a multilingually supervised task, where training, validation and test data are from exactly the same domain. The *top-to-bottom axis* gives rise to a cross-lingual transfer task, where a model trained on German and French is evaluated on Italian answers to the same questions. The *left-to-right axis* represents a continuous shift of domain: In the *middle column*, the model is tested on previously unseen questions that belong to the same topics as seen during training. In the *right column* the model encounters unseen answers to unseen questions within an unseen topic. The two test sets in parentheses are too small for a significant evaluation.

questions and corresponding answers pertaining to national elections were included.

In the context of communal and cantonal elections, candidates have answered both local questions and a subset of the national questions. Of those elections, we only considered answers to the questions that also had been asked in a national election. They were only used to augment the training set while the validation and test sets were restricted to answers from national elections.

We discarded the fewer than 20 comments classified as English. Furthermore, we discarded instances that met any of the following conditions:

- Question is not a closed question or does not address a clearly defined political issue.
- No comment was submitted by the candidate or the comment is shorter than 50 characters.
- Comment starts with “but” or a similar indicator that the comment is not self-contained.
- Comment contains a URL.

In total, a fifth of the comments were filtered out.

Topics The questions have been organized by the *Smartvote* editors into categories (such as “Economy”). We further consolidated the pre-defined categories into 12 broad topics (Table 1).

Compliance The dataset is shared under a CC BY-NC 4.0 license. Copyright remains with www.smartvote.ch.

Given the sensitive nature of the data, we increase the anonymity of the data by hashing the respondents’ IDs. No personal attributes of the respondents are included in the dataset. We provide a data statement (Bender and Friedman, 2018) in Appendix B.

3.3 Data Split

We held out the topics “Healthcare” and “Political System” from the training data and created a separate *cross-topic* test set that contains the questions and answers related to those topics.

Furthermore, in order to test *cross-question* generalization performance within previously seen topics, we manually selected 16 held-out questions that are distributed over the remaining 10 topics. We selected the held-out questions manually because we wanted to make sure that they are truly unseen and that no paraphrases of the questions are found in the training set.

We designated Italian as a test-only language, since relatively few comments have been written in Italian. From the remaining German and French data we randomly selected a percentage of respondents as validation or as test respondents.

As a result we received one training set, one validation set and four test sets. The sizes of the sets are listed in Table 2. We did not consider test sets that are cross-lingual and cross-target at the same time, as they would have been too small to yield significant results.

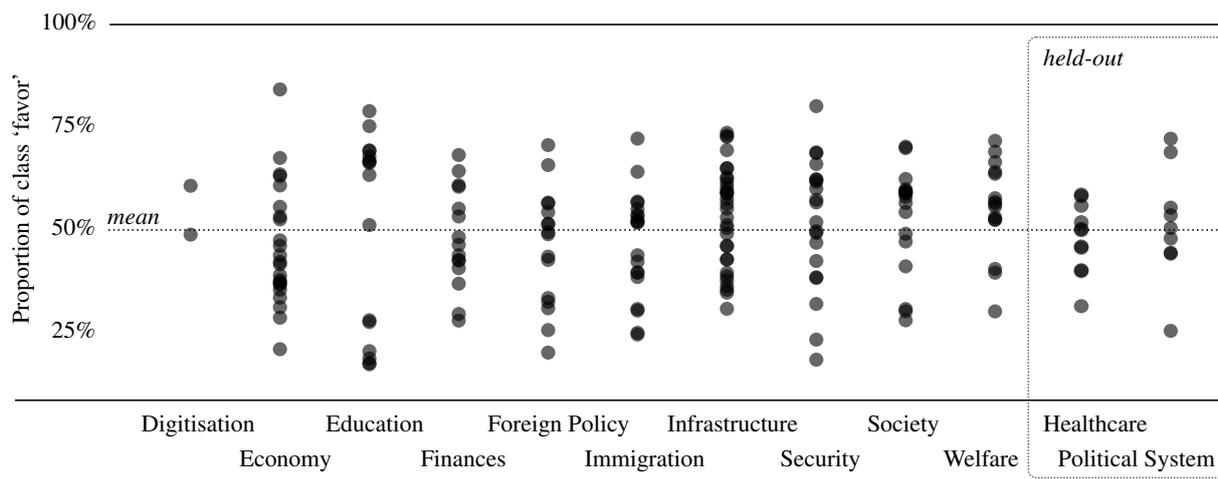


Figure 2: Proportion of ‘favor’ labels per question, grouped by topic. While the proportion of favorable answers varies from question to question, it is balanced overall.

3.4 Analysis

Some observations regarding the composition of X-stance can be made.

Class Distribution Figure 2 visualizes the proportion of ‘favor’ and ‘against’ stances for each target in the dataset. The ratio differs between questions but is relatively equally distributed across the topics. In particular, the questions in the held-out topics (with a ‘favor’ ratio of 49.4%) have a similar class distribution as the questions in other topics (with a ‘favor’ ratio of 50.0%).

Linguistic Properties Not every question is unique; some questions are paraphrases describing the same political issue. For example, in the 2015 election, the candidates were asked: “*Should the consumption of cannabis as well as its possession for personal use be legalised?*” Four years later they were asked: “*Should cannabis use be legalized?*” However, we do not see any need to consolidate those duplicates because they contribute to the diversity of the training data.

We further observe that while some questions in the dataset are quite short, some questions are rather convoluted. For example, a typical long question reads:

Some 1% of direct payments to Swiss agriculture currently go to organic farming operations. Should this proportion be increased at the expense of standard farming operations as part of Switzerland’s 2014-2017 agricultural policy?

Such longer questions might be more challenging to process semantically.

Languages The x-stance dataset has more German samples than French samples. The language ratio of about 3:1 is consistent across all training and test sets. Given the two languages it is possible to either train two monolingual models or to train a single model in a multi-source setup (McDonald et al., 2011). We choose a multi-source baseline because M-BERT is known to benefit from multilingual training data both in a supervised and in a cross-lingual scenario (Kondratyuk and Straka, 2019).

4 Baseline Experiments

We evaluate four baselines to obtain an impression of the difficulty of the task.

4.1 Majority Class Baselines

The first pair of baselines uses the most frequent class in the training set for prediction. Specifically, the *global* majority class baseline predicts the most frequent class across all training targets while the *target-wise* majority class baseline predicts the class that is most frequent for a given target question. The latter can only be applied to the intra-target test sets.

4.2 Bag-of-Words Baseline

As a second baseline, we train a *fastText* bag-of-words linear classifier (Joulin et al., 2017). For each comment, we select the translation of the question that matches its language, and concatenate it to the comment. We tokenize the text using the Europarl preprocessing tools (Koehn, 2005).

The ‘against’ class was slightly upsampled in

the training data so that the classes are balanced when summing over all questions and topics.

We use the standard settings provided by the fastText library.⁴ Optimal hyperparameters from the following range were determined based on the validation accuracy:

- Learning rate: 0.1, 0.2, 1
- Number of epochs: 5, 50

The word vectors were set to a size of 300. We do not initialize them with pre-trained multilingual embeddings since preliminary experiments did not show a beneficial effect.

4.3 Multilingual BERT Baseline

As our main baseline model we fine-tune multilingual BERT (M-BERT) on the task (Devlin et al., 2019) which has been pre-trained jointly in 104 languages⁵ and has established itself as a state of the art for various multilingual tasks (Wu and Dredze, 2019; Pires et al., 2019). Within the field of stance detection, BERT can outperform both feature-based and other neural approaches in a monolingual English setting (Ghosh et al., 2019).

Architecture In the context of BERT we interpret the X-stance task as *sequence pair classification* inspired by natural language inference tasks (Bowman et al., 2015). We follow the procedure outlined by Devlin et al. (2019) for such tasks. We designate the question as segment A and the comment as segment B. The two segments are separated with the special token [SEP], and the special token [CLS] is prepended to the sequence. The final hidden state corresponding to [CLS] is then classified by a linear layer.

We fine-tune the full model with a cross-entropy loss, using the AllenNLP library (Gardner et al., 2018) as a basis for our implementation.

Training As above, we balanced out the number of classes in the training set. We use a batch size of 16 and a maximum sequence length of 512 subwords, and performed a grid search over the following hyperparameters based on the validation accuracy:

- Learning rate: 5e-5, 3e-5, 2e-5
- Number of epochs: 3, 4

⁴<https://github.com/facebookresearch/fastText>

⁵<https://github.com/google-research/bert/blob/master/multilingual.md>

	DE	FR	IT
Majority class (global)	33.1	34.8	34.4
Majority class (target-wise)	60.8	65.1	59.3
fastText	69.9	71.2	53.7
M-BERT	76.8	76.6	70.2

Table 3: Baseline scores in the *cross-lingual* setting. No Italian samples were seen during training, making this a case of zero-shot cross-lingual transfer. The scores are reported as the macro-average of the F1-scores for ‘favor’ and for ‘against’.

The grid search was repeated independently for every variant that we test in the following subsections. Furthermore, the standard recommendations for fine-tuning BERT were used: Adam with $\beta_1 = 0.9$ and $\beta_2 = 0.999$; an L2 weight decay of 0.01; a learning rate warmup over the first 10% of the steps; and a linear decay of the learning rate. A dropout probability of 0.1 was set on all layers.

Results Table 3 shows the results for the cross-lingual setting. M-BERT performs consistently better than the previous baselines. Even the zero-shot performance in Italian, while significantly lower than the supervised scores, is much better than the target-wise majority class baseline.

Results for the cross-target setting are given in Table 4. Similar to the cross-lingual setting, model performance drops in the cross-target setting, but M-BERT remains the strongest baseline and easily surpasses the majority class baselines. Furthermore, the cross-question score of M-BERT is slightly lower than the cross-topic score.

4.4 How Important is Consistent Language?

The default setup preserves *horizontal language consistency* in that the language of the questions always corresponds to the language of the comments. For example, the Italian test instances are combined with the Italian version of the questions, even though during training the model has only ever seen the German and French version of them.

An alternative concept is *vertical language consistency*, whereby the questions are consistently presented in one language, regardless of the comment. To test whether horizontal or vertical consistency is more helpful, we train and evaluate M-BERT on a dataset variant where all questions are in their English version. We chose English as a lingua franca because it had the largest share of data during the pre-training of M-BERT.

	Intra-target			Cross-question			Cross-topic		
	DE	FR	Mean	DE	FR	Mean	DE	FR	Mean
Majority class (global)	33.1	34.8	33.9	36.4	37.9	37.1	32.1	33.8	32.9
Majority class (target-wise)	60.8	65.1	62.9	-	-	-	-	-	-
fastText	69.9	71.2	70.5	62.0	65.6	63.7	63.1	65.5	64.3
M-BERT	76.8	76.6	76.6	68.5	68.4	68.4	68.9	70.9	69.9

Table 4: Baseline scores in the *cross-target* setting. For each test set we separately report a German and a French score, as well as their harmonic mean.

Results are shown in Table 5. While the effect is negligible in most settings, cross-lingual performance increases when all questions are in English.

4.5 How Important are the Segments?

In order to rule out that only the questions or only the comments are necessary to optimally solve the task, we conduct some additional experiments:

- Only use a single segment containing the comment, removing the questions from the training and test data (*missing questions*).
- Only use the question and remove the comment (*missing comments*).

In both cases the performance decreases across all evaluation settings (Table 5). The loss in performance is much higher when comments are missing, indicating that the comments contain the most important information about stance. As can be expected, the score achieved without comments is only slightly different from the target-wise majority class baseline.

But there is also a loss in performance when the questions are missing, which underlines the importance of pairing both pieces of text. The effect of missing questions is especially strong in the supervised and cross-lingual settings. To illustrate this, we provide in Table A8 some examples of comments that occur with multiple different targets in the training set. Those examples can explain why the target can be essential for disambiguating a stance detection problem. On the other hand, the effect of omitting the questions is less pronounced in the cross-target settings.

The above single-segment experiments tell us that both the comment and the question provide crucial information. But it is possible that the M-BERT model, even though trained on both segments, mainly looks at a single segment at test

time. To rule this out, we probe the model with randomized data at test time:

- Test the model on versions of the test sets where the comments remain in place but the questions are shuffled randomly (*random questions*). We make sure that the random questions come from the same test set and language as the original questions.
- Keep the questions in place and randomize the comments (*random comments*). Again we shuffle the comments only within test set boundaries.

The results in Table 5 show that the performance of the model decreases in both cases, confirming that it learns to take into account both segments.

4.6 How Important are Spelled-Out Targets?

Finally we test whether the target really needs to be represented by natural language (e.g. “Do you support X?”). An alternative is to represent the target with a trainable embedding instead.

In order to fit target embeddings smoothly into our architecture, we represent each target type with a different reserved symbol from the M-BERT vocabulary. Segment A is then set to this symbol instead of a natural language question.

The results for this experiment are listed in the bottom row of Table 5. An M-BERT model that learns target embeddings instead of encoding a question performs clearly worse in the supervised and cross-lingual settings. From this we conclude that spelled-out natural language questions provide important linguistic detail that can help in stance detection.

5 Discussion

Our experiments show that M-BERT achieves a reasonable accuracy on X-stance, outperforming majority class baselines and a fastText classifier.

	Supervised	Cross-Lingual	Cross-Question	Cross-Topic
M-BERT	76.6	70.2	68.4	69.9
— with English questions	76.1	71.7	68.5	69.4
— with missing questions	73.2	67.1	67.8	69.3
— with missing comments	64.2	60.5	51.1	48.6
— with random questions	56.0	52.5	47.7	48.5
— with random comments	50.7	50.7	48.2	48.7
— with target embeddings	70.1	66.0	68.4	69.0

Table 5: Results for additional experiments. The cross-lingual score is the F1-score on the Italian test set. For the supervised, cross-question and cross-topic settings we report the harmonic mean of the German and French scores.

Dataset	Evaluation	Score
SemEval-2016	Ghosh et al. (2019)	75.1
MPCHI	Ghosh et al. (2019)	75.6
X-stance	this paper	76.6

Table 6: Performance of BERT-like models on different supervised stance detection benchmarks.

To put the supervised score into context we list scores that variants of BERT have achieved on other stance detection datasets in Table 6. It seems that the supervised part of X-stance has a similar difficulty as the SemEval-2016 (Mohammad et al., 2016a) or MPCHI (Sen et al., 2018) datasets on which BERT has previously been evaluated.

On the other hand, in the cross-lingual and cross-target settings, the mean score drops by 6–8 percentage points compared to the supervised setting; while zero-shot transfer is possible to a degree, it can still be improved.

The additional experiments (Table 5) validate the results and show that the sequence-pair classification approach to stance detection is justified.

It is interesting to see what errors the M-BERT model makes. Table A7 presents instances where it predicts the wrong label with a high confidence. These examples indicate that many comments express their stance only on a very implicit level, and thus hint at a potential weakness of the dataset. Because on the voting advice platform the label is explicitly shown to readers in addition to the comments, the comments do not need to express the stance explicitly.

Manual annotation could eliminate very implicit samples in a future version of the dataset. However, the sheer size and breadth of the dataset could not realistically be achieved with manual an-

notation, and, in our view, largely compensates for the implicitness of the texts.

6 Conclusion

We have presented a new dataset for political stance detection called X-stance. The dataset extends over a broad range of topics and issues regarding national Swiss politics. This diversity of topics opens up an opportunity to further study multi-target learning. Moreover, being partly Swiss Standard German, partly French and Italian, the dataset promotes a multilingual approach to stance detection.

By compiling formal commentary by politicians on political questions, we add a new text genre to the field of stance detection. We also propose a question–answer format that allows us to condition stance detection models on a target naturally.

Our baseline results with multilingual BERT show that the model has some capability to perform zero-shot transfer to unseen languages and to unseen targets (both within a topic and to unseen topics). However, there is some gap in performance that future work could address. We expect that the X-stance dataset could furthermore be a valuable resource for fields such as argument mining, argument search or topic classification.

Acknowledgments

This work was funded by the Swiss National Science Foundation (project MUTAMUR; no. 176727). We would like to thank Isabelle Augenstein, Anne Göhring and the anonymous reviewers for helpful feedback.

References

- Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. [Stance detection with bidirectional conditional encoding](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 876–885, Austin, Texas. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Simon Clematide, Stefan Gindl, Manfred Klenner, Stefanos Petrakis, Robert Remus, Josef Ruppenhofer, Ulli Waltinger, and Michael Wiegand. 2012. [MLSA — a multi-layered reference corpus for German sentiment analysis](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3551–3556, Istanbul, Turkey. European Language Resources Association (ELRA).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Taffjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Shalmoli Ghosh, Prajwal Singhania, Siddharth Singh, Koustav Rudra, and Saptarshi Ghosh. 2019. [Stance detection in web and social media: a comparative study](#). In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 75–87. Springer.
- Johannes Graën, Tannon Kew, Anastassia Shaitarova, and Martin Volk. 2019. [Modelling large parallel corpora: The zurich parallel corpus collection](#). In *Proceedings of the 7th Workshop on Challenges in the Management of Large Corpora (CMLC)*, pages 1–8. Leibniz-Institut für Deutsche Sprache.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Manfred Klenner, Don Tuggener, and Simon Clematide. 2017. [Stance detection in Facebook posts of a German right-wing party](#). In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pages 31–40, Valencia, Spain. Association for Computational Linguistics.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). *Machine Translation Summit, 2005*, pages 79–86.
- Dan Kondratyuk and Milan Straka. 2019. [75 languages, 1 model: Parsing universal dependencies universally](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Dilek Küçük and Fazli Can. 2020. [Stance detection: A survey](#). *ACM Comput. Surv.*, 53(1).
- Mirko Lai, Alessandra Teresa Cignarella, Delia Irazú Hernández Farías, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. [Multilingual stance detection in social media political debates](#). *Computer Speech & Language*, page 101075.
- Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. [Stance evolution and twitter interactions in an italian political debate](#). In *International Conference on Applications of Natural Language to Information Systems*, pages 15–27. Springer.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. [Multi-source transfer of delexicalized dependency parsers](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 62–72, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016a. [A dataset for detecting stance in tweets](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3945–3952, Portorož, Slovenia. European Language Resources Association (ELRA).
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016b. [SemEval-2016 task 6: Detecting stance in tweets](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.

- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Anirban Sen, Manjira Sinha, Sandya Mannarswamy, and Shourya Roy. 2018. [Stance classification of multi-perspective consumer health information.](#) In *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, pages 273–281.
- Nakatani Shuyo. 2010. [Language detection library for java.](#)
- Parinaz Sobhani, Diana Inkpen, and Xiaodan Zhu. 2017. [A dataset for multi-target stance detection.](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 551–557, Valencia, Spain. Association for Computational Linguistics.
- Mariona Taulé, M Antònia Martí, Francisco Rangel, Paolo Rosso, Cristina Bosco, and Viviana Patti. 2017. [Overview of the task on stance and gender detection in tweets on catalan independence at ibereval 2017.](#) In *2nd Workshop on Evaluation of Human Language Technologies for Iberian Languages, IberEval 2017*, volume 1881, pages 157–177.
- Mariona Taulé, Francisco Rangel, M Antònia Martí, and Paolo Rosso. 2018. [Overview of the task on multimodal stance detection in tweets on catalan #1oct referendum.](#) In *3rd Workshop on Evaluation of Human Language Technologies for Iberian Languages, IberEval 2018*, volume 2150, pages 149–166.
- James Thurman and Urs Gasser. 2009. [Three case studies from switzerland: Smartvote.](#) *Berkman Center Research Publications*.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.

A Examples

Question	Comment	Gold Label	Prob.
<i>Befürworten Sie eine vollständige Liberalisierung der Geschäftsöffnungszeiten?</i> [Are you in favour of a complete liberalisation of business hours for shops?]	<i>Ausser Sonntag. Dies sollte ein Ruhetag bleiben können.</i> [Except Sunday. That should remain a day of rest.]	FAVOR	0.001
<i>Soll die Schweiz innerhalb der nächsten vier Jahre EU-Beitrittsverhandlungen aufnehmen?</i> [Should Switzerland embark on negotiations in the next four years to join the EU?]	<i>In den nächsten vier Jahren ist dies wohl unrealistisch.</i> [For the next four years this is probably unrealistic.]	FAVOR	0.005
<i>Befürworten Sie einen Ausbau des Landschaftsschutzes?</i> [Are you in favour of extending landscape protection?]	<i>Wenn es darum geht erneuerbare Energien zu fördern, ist sogar eine Lockerung angebracht.</i> [When it comes to promoting renewable energy, even a relaxation is appropriate.]	AGAINST	0.006
<i>La Suisse devrait-elle engager des négociations pour un accord de libre échange avec les Etats-Unis?</i> [Should Switzerland start negotiations with the USA on a free trade agreement?]	<i>Il faut cependant en parallèle veiller à ce que la Suisse ne soit pas mise de côté par les Etats-Unis!</i> [At the same time it must be ensured that Switzerland is not sidelined by the United States!]	AGAINST	0.010

Table A7: Some classification errors where the predicted probability of the correct label is especially low. The examples have been taken from the validation set.

Comment ...	is favorable towards target ...	but against target ...
<i>Ich will offene Grenzen für Waren und selbstverantwortliche mündige Bürger. Der Staat hat kein Recht, uns einzuschränken.</i> [I want open borders for goods and responsible citizens. The state has no right to restrict us.]	<i>Soll die Schweiz mit den USA Verhandlungen über ein Freihandelsabkommen aufnehmen?</i> [Should Switzerland start negotiations with the USA on a free trade agreement?]	<i>Soll die Schweiz das Schengen-Abkommen mit der EU kündigen und wieder verstärkte Personenkontrollen direkt an der Grenze einführen?</i> [Should Switzerland terminate the Schengen Agreement with the EU and reintroduce increased identity checks directly on the border?]
<i>Hier gilt der Grundsatz der Eigenverantwortung und Selbstbestimmung des Unternehmens!</i> [The principle of personal responsibility and corporate self-regulation applies here!]	<i>Sind Sie für eine vollständige Liberalisierung der Ladenöffnungszeiten?</i> [Are you in favour of the complete liberalization of shop opening times?]	<i>Würden Sie die Einführung einer Frauenquote in Verwaltungsräten börsenkotierter Unternehmen befürworten?</i> [Would you support the introduction of a woman's quota for the Boards of Directors of listed companies?]

Table A8: Two comments that imply a positive stance towards one target issue but a negative stance towards another target issue. Such cases can be found in the dataset because respondents have copy-pasted some comments. These examples have been extracted from the training set.

B Data Statement

Curation rationale In order to study the automatic detection of stances on political issues, questions and candidate responses on the voting advice application `smartvote.ch` were downloaded. Mainly data pertaining to national-level issues were included to reduce variability.

Language variety The training set consists of questions and answers in Swiss Standard German and Swiss French (74.1% de-CH; 25.9% fr-CH). The test sets also contain questions and answers in Swiss Italian (67.1% de-CH; 24.7% fr-CH; 8.2% it-CH). The questions have also been translated into English.

Speaker demographic (answers)

- Candidates for communal, cantonal or national elections in Switzerland who have filled out an online questionnaire.
- Age: 18 or older – mixed.
- Gender: Unknown – mixed.
- Race/ethnicity: Unknown – mixed.
- Native language: Unknown – mixed.
- Socioeconomic status: Unknown – mixed.
- Different speakers represented: 7581.
- Presence of disordered speech: Unknown.

Speech situation

- The questions were edited and translated by political scientists for a public voting advice website.
- The answers were written between 2011 and 2020 by the users of the website.

Text characteristics Questions, answers, arguments and comments regarding political issues.