A Tutorial on Stance Detection

ACM WSDM 2022 Tutorial

https://dkucuk.github.io/stancedetection/

[Part 1]

DİLEK KÜÇÜK
Energy Institute
TÜBİTAK MRC, Ankara-Turkey
dilek.kucuk@tubitak.gov.tr

FAZLI CAN
Computer Engineering Department
Bilkent University, Ankara-Turkey
canf@cs.bilkent.edu.tr

Outline

- Introduction
- Core Concepts and Related Problems
- Stance Detection Competitions (Shared Tasks)
- Historical and Contemporary Approaches
- Common Stance Detection Datasets

Introduction

- **Stance detection** is also known as:
  - stance classification
  - stance identification
  - stance prediction
  - stance analysis
  - debate-side classification
  - debate stance classification

- Stance detection is at the crossroads of:
  - Natural language processing
  - Social media analysis
  - Information retrieval

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.

Introduction

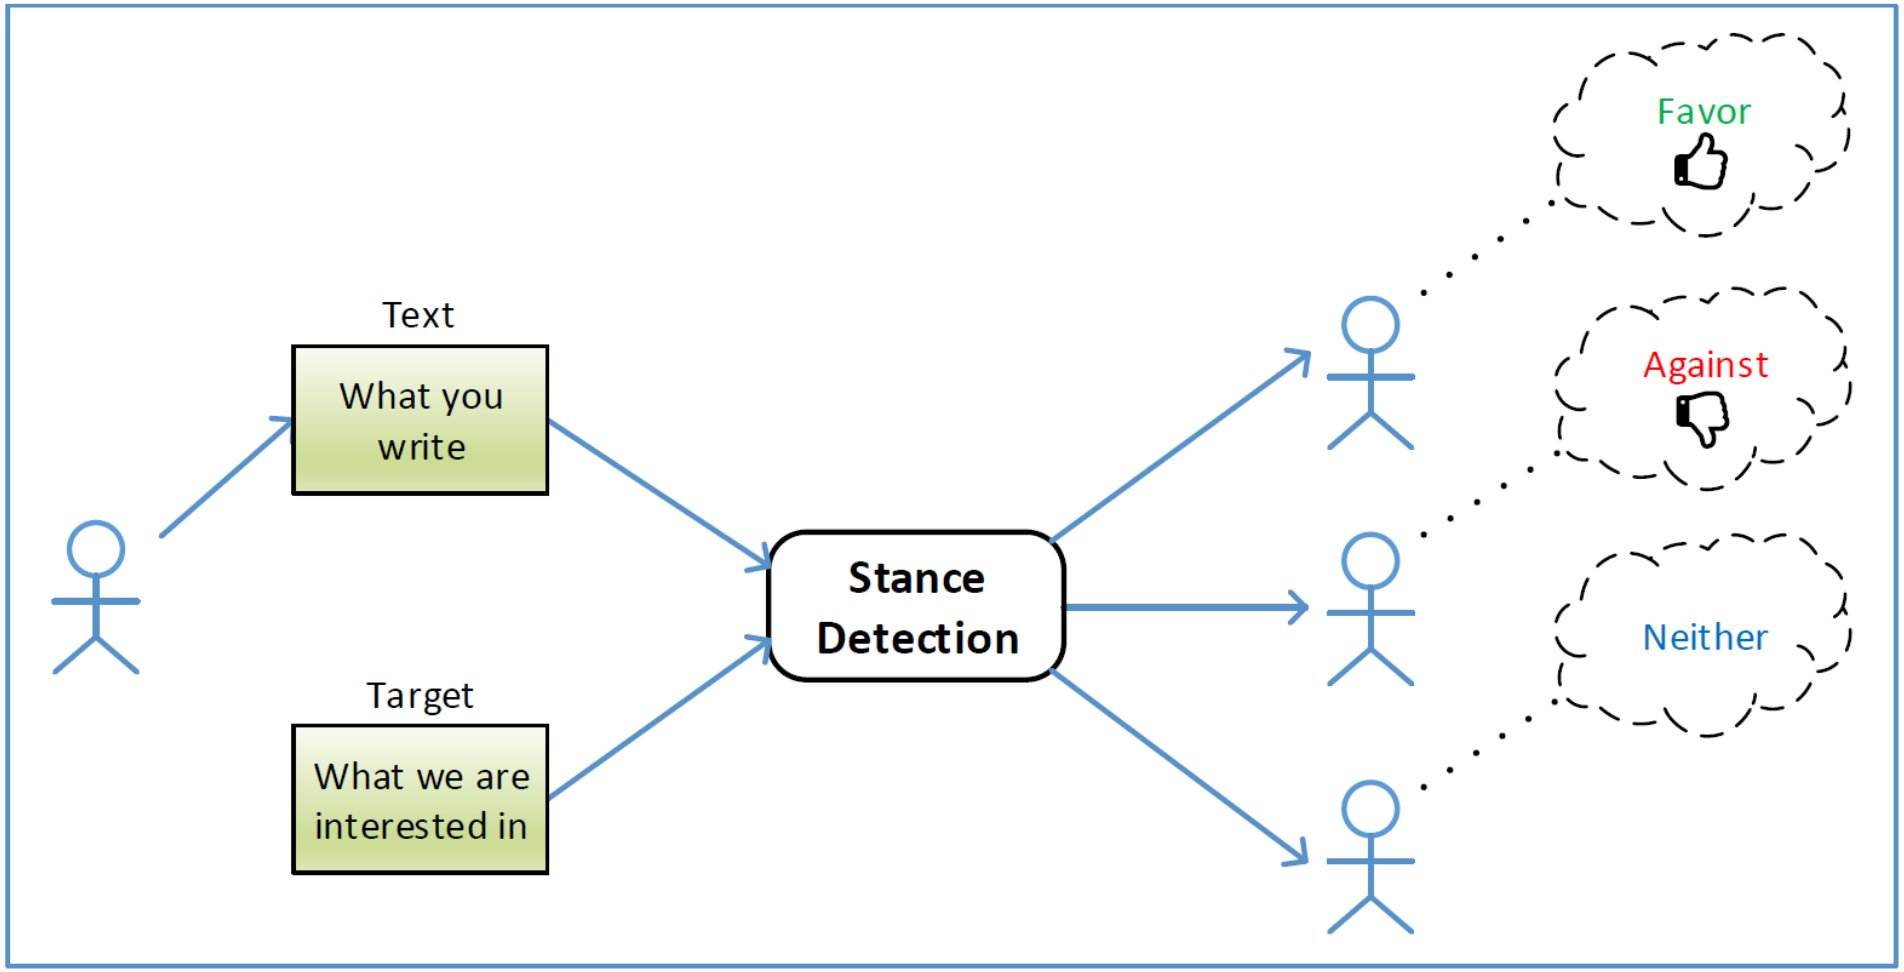

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.

Introduction

- Other **common class labels** used in stance detection research:
  - Favor, Against, None
  - Favor, Against, Neutral
  - For, Against
  - Support, Oppose
  - Pro, Con
  - Pro, Anti

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.

Core Concepts

Definition 1.1 (Stance Detection). For an input in the form of a piece of text and a target pair, stance detection is a classification problem where the stance of the author of the text is sought in the form of a category label from this set: {*Favor*, *Against*, *Neither*}. Occasionally, the category label of *Neutral* is also added to the set of stance categories [Mohammad et al. 2016b] and the target may or may not be explicitly mentioned in the text [Augenstein et al. 2016a; Mohammad et al. 2016b].

Definition 1.2 (Multi-target Stance Detection). For an input in the form of a piece of text and a set of related targets, multi-target stance detection is a classification problem where the stance of the text author is sought as a category label from this set: {*Favor*, *Against*, *Neither*} for each target and each stance classification (for each target) might have an effect on the classifications for the remaining targets [Sobhani 2017].

Definition 1.3 (Cross-target Stance Detection). Cross-target stance detection is a classification problem where the stance of the text author is sought for a specific target as a category label from this set: {*Favor*, *Against*, *Neither*}, in a settings where stance annotations are available for (though related but) different targets, i.e., there is not enough stance-annotated training data for the target under consideration [Augenstein et al. 2016a; Xu et al. 2018].

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey.** ACM Computing Surveys, 53, 1.

Core Concepts

Definition 1.4 (Rumour Stance Classification). For an input in the form of a piece of text and a rumour pair, rumour stance classification is a problem where the position of the text author towards the veracity of the rumour is sought for, in the form of a category label from this set: {*Supporting*, *Denying*, *Querying*, *Commenting*}. As the set of possible category labels, a subset of this set such as {*Supporting*, *Denying*} is occasionally employed [Zubiaga et al. 2018].

Definition 1.5 (Fake News Stance Detection). For an input in the form of news headline and a news body pair (where the headline and body parts may belong to different news articles), this is a classification problem where the stance of the body towards the claim of the headline is sought for, in the form of a category label from this set: {*Agrees*, *Disagrees*, *Discusses* (the same topic), *Unrelated*}. This problem is defined in order to facilitate the task of fake news detection [FNC 2017].

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey.** ACM Computing Surveys, 53, 1.

Core Concepts

Table 1. Sample Tweets from SemEval 2016 Stance Dataset [Mohammad et al. 2016b].

Tweet	Stance Target	Stance	Sentiment
RT @TheCLF: Thanks to everyone in Maine who contacted their legislators in support of #energyefficiency funding! #MEpoli #SemST	Climate Change is a Real Concern	Favor	Positive
We live in a sad world when wanting equality makes you a troll... #SemST	Feminist Movement	Favor	Negative
I don't believe in the hereafter. I believe in the here and now. #SemST	Atheism	Favor	Neither
@violencehurts @WomenCanSee The unborn also have rights #defendthe8th #SemST	Legalization of Abortion	Against	Positive
I'm conservative but I must admit I'd rather see @SenSanders as president than Mrs. Clinton. #stillvotingGOP #politics #SemST	Hillary Clinton	Against	Negative
I have my work and my faith... If that's boring to some people, I can't tell you how much I don't care. ~Madonna Ciccone #SemST	Atheism	Against	Neither
@BadgerGeno @kreichert27 @jackbahlman Too busy protesting :) #LoveForAll #BackdoorBadgers #SemST	Hillary Clinton	Neither	Positive
@ShowTruth You're truly unwelcome here. Please leave. #ygk #SemST	Legalization of Abortion	Neither	Negative
@Maisie_Williams everyone feels that way at times. Not just women #SemST	Atheism	Neither	Neither

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey.** ACM Computing Surveys, 53, 1.

Related Problems

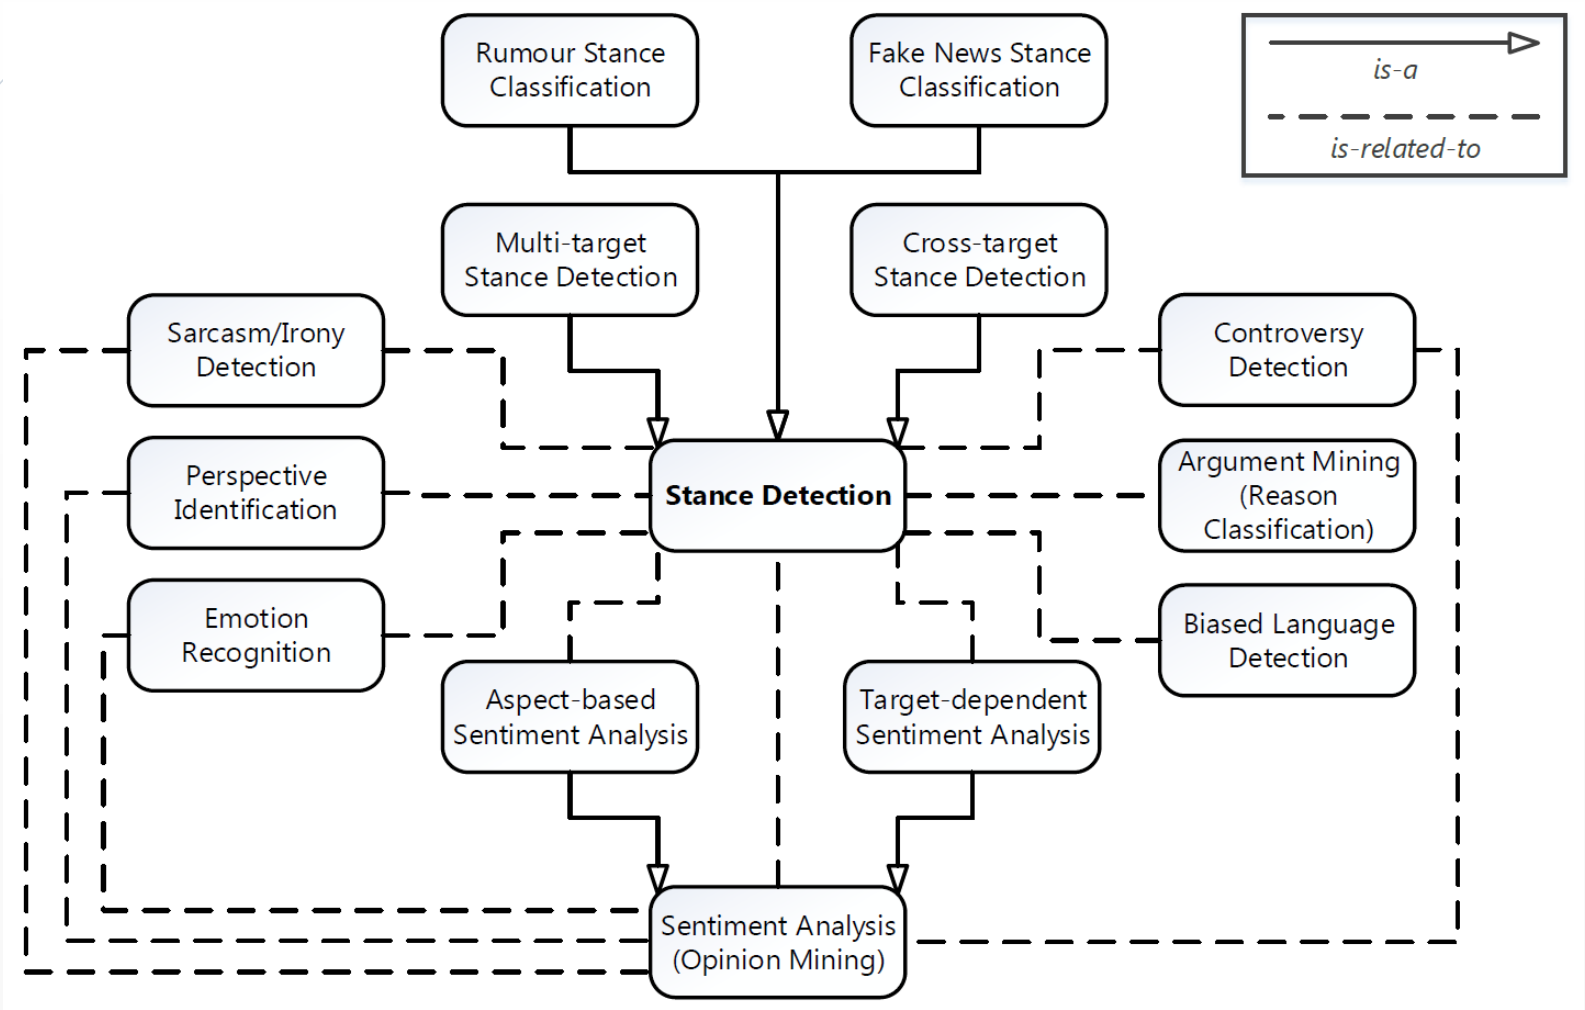

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.

Related Problems

- **Sentiment analysis** is usually considered as polarity detection: a classification output as **Positive**, **Negative**, or **Neutral** is expected.
  - **Aspect-oriented (aspect-based, or aspect-level) sentiment analysis:** sentiment polarities towards a target entity and different aspects of this entity are considered (Pontiki et al., 2015)
  - **Target-based (target-dependent) sentiment analysis**: sentiment polarity towards the target is explored within the text, given a text and a target pair (Jiang et al., 2011)

Related Problems

- Differences between **stance detection** and **target-based sentiment analysis:**
  - Stance target may not be explicitly given in the input text
  - Stance target may not be the target of the sentiment in the text
  - Stance target may be an event while the target is usually an entity or an aspect in (aspect-based or target-based) sentiment analysis.

Related Problems

- **Perspective identification** is the automatic determination of the point-of-view of the author of a piece of text, from the textual content
  - **Democrats** or **Republicans** in the context of US elections
- **Sarcasm/irony detection** is a classification problem where the existence of sarcasm/irony in a given text is explored
  - **Sarcasm** is defined as the verbal form of an **irony**

Related Problems

- In **controversy detection**, a (relevant) controversy score is generally calculated and associated with each unit of content

- **Biased language detection** explores the existence of an inclination/tendency towards a particular perspective within text (Recasens et al., 2013)
  - In terms of stance, biased language detection can be defined as the detection of textual content **which includes a particular non-neutral stance**

Related Problems

- **Argument (argumentation) mining** is the extraction of possible argument structure in a given text (Lippi and Torroni, 2016)
  1. Detection of the argumentative sentences
  2. Extraction of argument components (claims and evidences/premises)
  3. Forming the final argument graph by connecting the extracted components
- **Emotion recognition** determines emotions in a given text
  - Joy, sadness, fear, etc.

Stance Detection Competitions (Shared Tasks)

1. SemEval-2016 Task 6: Detecting Stance in (English) Tweets
2. Shared Task of Stance Detection in Chinese Microblogs at NLPCC-ICCPOL-2016
3. Shared Task of Stance Detection in Spanish and Catalan Tweets at IberEval-2017
4. SardiStance: Stance Detection Task in Italian Tweets at EVALITA-2020
5. VaxxStance: Going Beyond Text in Cross-lingual Stance Detection at IberLeF-2021

Stance Detection Competitions (Shared Tasks)

1. SemEval-2016 Shared Task (on Twitter)

- By Mohammad et al. (2016) on English tweets (4,870 tweets)
- **Targets:** Atheism, Climate change is a real concern, Feminist movement, Hillary Clinton, Legalization of abortion, Donald Trump
- Two subtasks:
  - A. supervised stance detection
  - B. weakly supervised stance detection
- An **RNN-based system** outperforms other participants in **subtask-A** (**F-score: 67.82%**)
- **SVM-based approach** (baseline system) by the organizers attains an **F-score of 68.98% for subtask A**
- A **CNN-based system** outperforms others in **subtask-B** (**F-score: 56.28%**)

Stance Detection Competitions (Shared Tasks)

2. NLPCC-ICCPOL-2016 shared task (on Weibo)

- By Xu et al. (2016) on Chinese microblog posts (4,000 annotated, 2400 unannotated)
- **Targets:** iPhone SE, Set off firecrackers in the Spring Festival, Russia's anti terrorist operations in Syria, Two child policy, Prohibition of motorcycles and restrictions on electric vehicles in Shenzhen, Genetically modified food, Nuclear test in DPRK
- Supervised and weakly supervised stance detection subtasks.
  - The best performing system for **subtask-A** is based on **SVM** and **random forest** (**maximum F-score is 71.06%**)
  - The best system for **subtask-B** achieves an **average F-score of 46.87%.**

Stance Detection Competitions (Shared Tasks)

3. IberEval-2017 shared task (on Twitter)

- By Taulé et al. (2017) on Spanish and Catalan tweets (5,400 tweets in Spanish, 5,400 tweets in Catalan)
- **Target:** Independence of Catalonia
- Best performing system
  - for **Spanish** is based on **SVM**
  - for **Catalan** is based on **logistic regression**

Stance Detection Competitions (Shared Tasks)

4. SardiStance-2020 shared task (on Twitter)

- By Cignarella et al. (2020) on Italian tweets (a total of 3,242 tweets)
- **Target:** Sardines Movement
- 2 subtasks: (A) textual stance detection, (B) contextual stance detection
- Best performing systems for tasks utilize pre-trained transformer-based (deep learning) models
  - Versions of **BERT*** (Devlin et al., 2018) trained on Italian tweets

*** Bidirectional Encoder Representations from Transformers**

Stance Detection Competitions (Shared Tasks)

5. VaxxStance shared task at IberLeF-2021

- By Agerri et al. (2021) on Basque and Spanish tweets (1,384 tweets in Basque and 2,697 tweets in Spanish) including 4 tracks:
  - 1. **Close Track**: (A) textual stance detection, (B) contextual stance detection
    - Only the provided data for each language is allowed
  - 2. **Open Track**: Any can of data including additional tweets can be used
  - 3. **Zero-shot Track**: Tweets of the target language cannot be used for training
- **Target**: Vaccines
- Best performing system utilizes **logistic regression**
- Contextual features (retweets, friends etc.) improve performance.

Other Competitions Related to Stance Detection

Fake News Challenge-2017 [http://www.fakenewschallenge.org/]

- The first stage in this competition is **fake news stance detection**
  - Given a headline and a body text
  - An output classification label from this set was expected from the participants: {**agrees**, **disagrees**, **discusses**, **unrelated**}
  - **49,972 annotated and 25,413 unannotated headline-body pairs** exist in the dataset.
    - https://github.com/FakeNewsChallenge/fnc-1

Other Competitions Related to Stance Detection

RumourEval-2019 shared task: determining rumour veracity and support for rumours

- Subtask A of RumourEval-2019 is on rumour stance detection
  - Given a rumour (as **target**) in a source tweet, a number of tweets are classified
  - The class labels are: **support**, **deny**, **query**, and **comment**.
- The competition was previously conducted in 2017 as RumourEval-2017.

Gorrell, G., Kochkina, E., Liakata, M., Aker, A., Zubiaga, A., Bontcheva, K., & Derczynski, L. (2019). **SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours.** In Proceedings of SemEval-2019.

A Historical Perspective

Earlier Work on Stance Detection [2006 - 2015]

- **Earlier work are carried out on**
  - Congressional-floor debates
  - Company internal discussions
  - Online social, political, and ideological debates (in public forums)
  - Online debates about products
  - Spontaneous speech (a single study by Levow et al. (2014))
  - Student essays
  - Tweets (few studies)

- **Approaches in earlier work**
  - Few rule-based methods
  - Supervised learning methods (SVM, decision tree, random forest, HMM, CRF, ILP, …)

Approaches to Stance Detection

Table 5. Temporal Distribution of Published Papers on Stance Detection

Publication Year	Number of Papers
2006 – 2010	5
2011 – 2014	8
2015 – 2016	38
2017 – 2019	78

* We have come across **more than 40 papers** related to stance detection **published in 2020 and 2021**

Approaches to Stance Detection

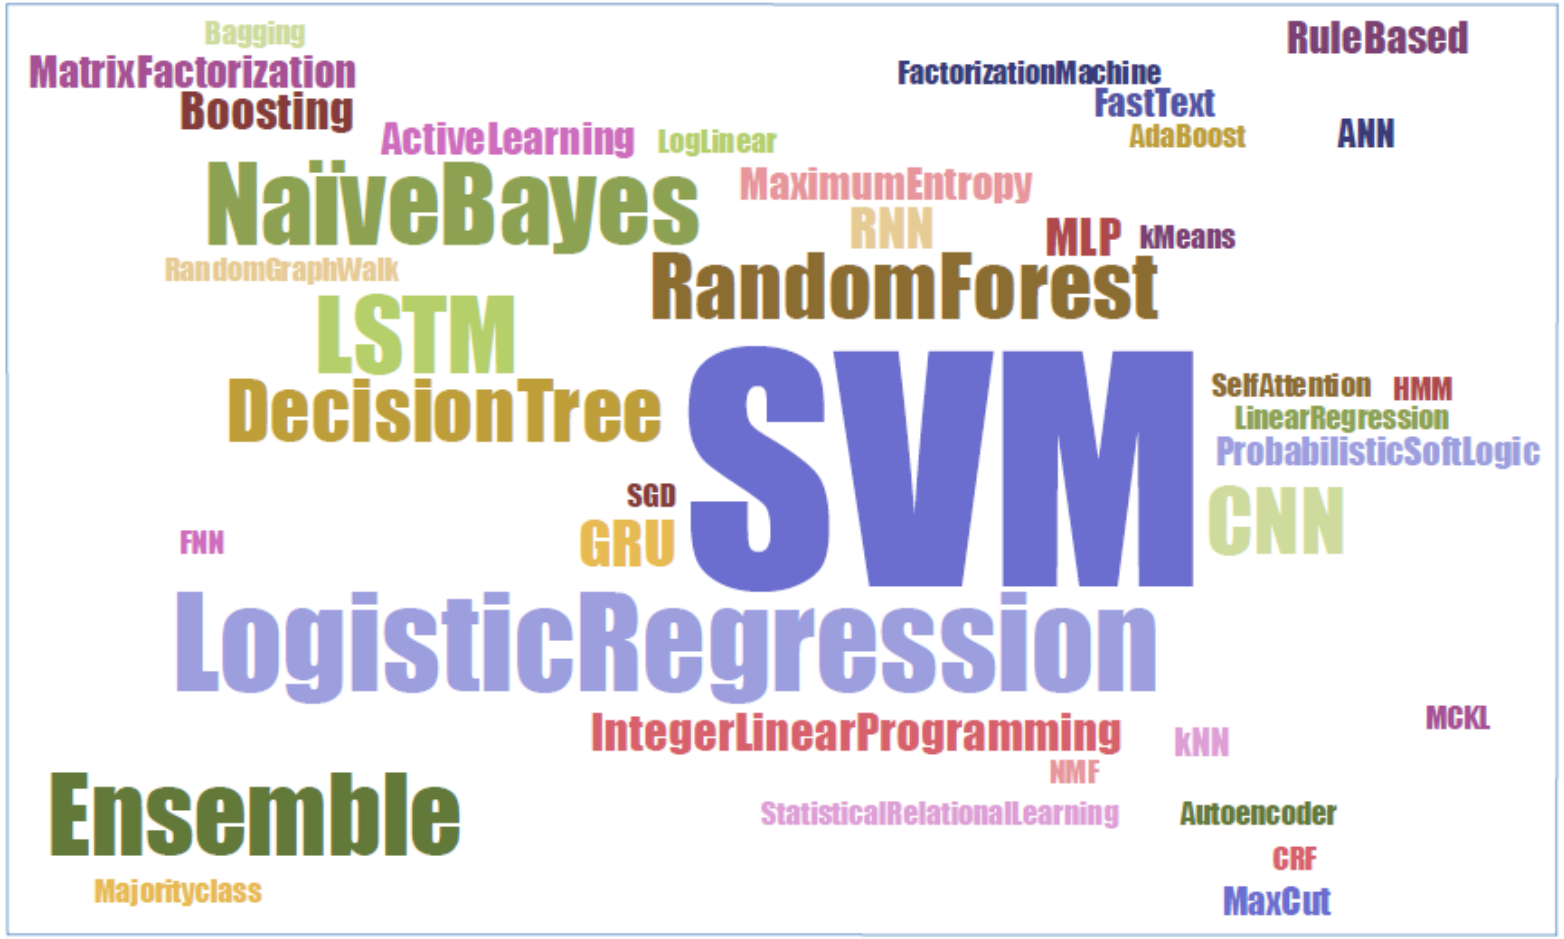

Fig. 4. A word cloud of the algorithms used for stance detection problem in the published papers included in this survey paper.

Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.

Contemporary Approaches to Stance Detection

- **Feature-based machine learning approaches**
  - SVM, Logistic Regression, Naïve Bayes, Decision Trees, ANN, ILP, kNN, …
- **Deep learning approaches*****
  - LSTM (RNN), Bi-LSTM, RNN, GRU (RNN), CNN
  - Pre-trained deep learning based language models such as BERT
- **Ensemble learning approaches**
  - Random Forest, Majority Voting, Proprietary ensemble learners, Boosting, Bagging, …

Contemporary Approaches to Stance Detection

- **Common Features Used by Learning Systems**
  - **Lexical features** (bag-of-words, word and character n-grams, skip-grams, hashtags, stance-indicative words, theme and context words, …)
  - **Features based on interactions among posts and users** (retweets, replies, agreement/disagreement links, quotes, …), and temporal information regarding the posts,
    - Also referred to as **contextual features**
  - **Features based on sentiment, subjectivity, arguing/argumentation lexicons, emotion indicator words…**

Contemporary Approaches to Stance Detection

Common Features Used by Learning Systems (cont'd)

- **Word vector representations** such as **word2vec** [Mikolov et al., 2013] and **GloVe** [Pennington et al., 2014] vectors
- **Topic modeling related features** such as those based on Latent Dirichlet Allocation (LDA), Latent Semantic Analysis (LSA), and TF-IDF vectors of lexical features
- **Features based on POS tags, named entities, dependency relations, syntactic rules, and coreference resolution**

Contemporary Approaches to Stance Detection

- The topics in stance detection datasets used by related research include **controversial topics, political/ideological debates, and elections/referendums**, among others.

- If the stance detection problem is stated in the form of three-way classification from the set {Favor, Against, Neither}, then a **two-phase pipelined approach** is commonly employed:

  1. classifying the stance of the tweet as Relevant (Favor or Against) and Irrelevant (Neither) towards the target,
  2. further classifying the Relevant tweets as having the Favor or Against stance towards the target.

- If there are more than one target in the stance detection settings, it is a common and effective practice to **train a separate classifier for each stance target**.

Contemporary Approaches to Stance Detection

- **Cross-target (or weakly-supervised) stance detection is a harder problem** to solve
  - the corresponding detection performance is often lower than supervised stance detection.
- In addition to the actual textual content under consideration (post, tweet, etc.), **it is beneficial to represent user information and conversational interactions** (**contextual features**) when performing stance detection.
- Traditional feature-based ML approaches, deep learning methods, and ensemble classifiers are all employed for stance detection in different studies.
  - It can be concluded that **all of these types of learning systems are reported to attain favorable and competitive performance** for the stance detection problem

Common Stance Detection Datasets

Authors	Domain	Annotation Classes	Target(s)	Size
[Mohammad et al. 2016a]	Tweets (English)	Favor, Against, Neither	Atheism, Climate change is a real concern, Feminist movement, Hillary Clinton, Legalization of abortion, Donald Trump	4,870 tweets
[Mohammad et al. 2017]	Tweets (English)	Favor, Against, Neither for stance; Positive, Negative, and Neither for sentiment	Atheism, Climate change is a real concern, Feminist movement, Hillary Clinton, Legalization of abortion, Donald Trump	4,870 tweets
[Xu et al. 2016b]	Microblogs (Chinese)	Favor, Against, None	iPhone SE, Set off firecrackers in the Spring Festival, Russia's anti terrorist operations in Syria, Two child policy, Prohibition of motorcycles and restrictions on electric vehicles in Shenzhen, Genetically modified food, Nuclear test in DPRK	4,000 annotated and 2,400 unannotated tweets
[Taulé et al. 2017]	Tweets (Catalan & Spanish)	Favor, Against, None	Independence of Catalonia	5,400 tweets in Spanish and 5,400 tweets in Catalan

Common Stance Detection Datasets

Authors	Domain	Annotation Classes	Target(s)	Size
[Sobhani et al. 2017]	Tweets (English)	Favor, Against, Neither	{Clinton-Sanders}, {Clinton-Trump}, {Cruz-Trump}	4,455 tweets
[Küçük 2017b]	Tweets (Turkish)	Favor, Against	Galatasaray, Fenerbahçe	700 tweets
[Küçük and Can 2018]	Tweets (Turkish)	Favor, Against	Galatasaray, Fenerbahçe	1,065 tweets
[Murakami and Raymond 2010]	Online debates (Japanese)	Support, Oppose	Selected five ideas	481 comments about five ideas
[Darwish et al. 2017]	Tweets (Arabic)	Favor (Positive), Against (Negative)	Transfer of two islands from Egypt to Saudi Arabia	33,024 tweets
[Hercig et al. 2017]	News comments (Czech)	In Favor, Against, Neither	Miloš Zeman, Smoking ban in restaurants	5,423 news comments

Recent Datasets

- **X-Stance**: dataset of 67,000 comments in German, French, and Italian
  - Annotated for stance detection
  - Targets: 150 political issues
  - Stance classes:
    - favor, against

Vamvas, J., & Sennrich, R. (2020). **X-stance: A multilingual multi-target dataset for stance detection**. *arXiv preprint arXiv:2003.08385*.

Recent Datasets

- **E-FRA and R-ITA:** distinct tweet datasets in French and Italian
  - Annotated for stance detection (1,116 French tweets and 833 Italian tweets)
  - **Targets**: election candidates (Macron & Le Pen) for E-FRA and constitutional reform for R-ITA.
  - Stance classes:
    - favor, against, none

Lai, M., Cignarella, A. T., Farías, D. I. H., Bosco, C., Patti, V., & Rosso, P. (2020). **Multilingual stance detection in social media political debates**. *Computer Speech & Language*, *63*, 101075

Recent Datasets

- **WT-WT**: financial dataset of 51,284 tweets in English
  - Annotated for rumour stance detection
  - **Targets**: Five mergers and acquisition operations
  - Rumour stance classes:
    - support, refute, comment, unrelated

Conforti, C., Berndt, J., Pilehvar, M. T., Giannitsarou, C., Toxvaerd, F., & Collier, N. (2020). **Will-They-Won't-They: A Very Large Dataset for Stance Detection on Twitter**. In Proceedings of ACL.

Recent Datasets

- **P-Stance**: A dataset of 21,574 tweets in English from the political domain
  - Annotated for stance detection
  - **Targets**: Donald Trump, Joe Biden, Bernie Sanders
  - Stance classes:
    - Favor, Against
    - None class removed due to issues about annotator agreement.

Li, Y., Sosea, T., Sawant, A., Nair, A. J., Inkpen, D., & Caragea, C. (2021). **P-Stance: A Large Dataset for Stance Detection in Political Domain**. In *Findings of the ACL: ACL-IJCNLP 2021* .

Resources: Recent Datasets

- **Stance-annotated datasets exist for several languages:**
  - Arabic
  - Basque
  - Catalan
  - Chinese
  - Czech
  - English
  - English-Hindi
  - French
  - German
  - Italian
  - Japanese
  - Russian
  - Spanish
  - Turkish

Data Augmentation

- **Data augmentation** is used to cope with training data scarcity in ML research by extending the dataset with modified versions of existing data.

- Even simple methods like synonym replacement and random swap are known to help improve performance.

  Wei, J., & Zou, K. (2019). **EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks.** In *Proceedings of the EMNLP-IJCNLP*.

- **Data augmentation techniques** can also be used to generate data **to improve stance detection performance**

  Li, Y., & Caragea, C. (2021). **Target-Aware Data Augmentation for Stance Detection**. In Proceedings of NAACL-HLT.

Evaluation Metrics

$$F = \frac{F_{Favor} + F_{Against}}{2} \qquad F_{Favor} = \frac{2 * P_{Favor} * R_{Favor}}{P_{Favor} + R_{Favor}} \qquad F_{Against} = \frac{2 * P_{Against} * R_{Against}}{P_{Against} + R_{Against}}$$

$$P_{Favor} = \frac{Correct_{Favor}}{Correct_{Favor} + Spurious_{Favor}} \qquad P_{Against} = \frac{Correct_{Against}}{Correct_{Against} + Spurious_{Against}}$$

$$R_{Favor} = \frac{Correct_{Favor}}{Correct_{Favor} + Missing_{Favor}} \qquad R_{Against} = \frac{Correct_{Against}}{Correct_{Against} + Missing_{Against}}$$

$$F = \frac{F_{Favor} + F_{Against} + F_{Neither}}{3}$$

$$Accuracy = \frac{Correct\ classifications}{All\ classifications}$$

Resources: Software and Tools

- Few papers present visualization systems/tools for stance detection.
- Many papers use the following machine learning tools, libraries in their stance detection experiments:
  - **Weka**
  - **Scikit-learn package**
  - **Keras**
  - **Theano**
  - **Gensim**
  - **SVM^{light}**
  - **FastText**
  - **Brainy**

- Pre-trained deep learning based models like **BERT** (Devlin et al., 2018) are also commonly used for stance detection.

References

- Agerri, R., Centeno, R., Espinosa, M., de Landa, J. F., & Rodrigo, A. (2021). **VaxxStance@IberLEF 2021: Overview of the Task on Going Beyond Text in Cross-Lingual Stance Detection** Procesamiento del Lenguaje Natural, 67, 173-181.
- Augenstein, I., Rocktäschel, T., Vlachos, A., & Bontcheva, K. (2016). **Stance detection with bidirectional conditional encoding**. *arXiv preprint arXiv:1606.05464*.
- Cignarella, A. T., Lai, M., Bosco, C., Patti, V., & Paolo, R. (2020**). SardiStance@EVALITA2020: Overview of the Task on Stance Detection in Italian Tweets**. In *Proceedings of* EVALITA-2020.
- Conforti, C., Berndt, J., Pilehvar, M. T., Giannitsarou, C., Toxvaerd, F., & Collier, N. (2020). **Will-They-Won't-They: A Very Large Dataset for Stance Detection on Twitter**. In Proceedings of ACL.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). **BERT: Pre-training of deep bidirectional transformers for language understanding**. *arXiv preprint arXiv:1810.04805*.
- Gorrell, G., Kochkina, E., Liakata, M., Aker, A., Zubiaga, A., Bontcheva, K., & Derczynski, L. (2019). SemEval-2019 task 7: **RumourEval, determining rumour veracity and support for rumours.** In *Proceedings of SemEval-2019*.

References

- Jiang, L., Yu, M., Zhou, M., Liu, X., & Zhao, T. (2011). **Target-dependent twitter sentiment classification**. In *Proceedings of the ACL: HLT*.
- Küçük, D. and Can, F. (2020). **Stance Detection: A Survey**. ACM Computing Surveys, 53, 1.
- Lai, M., Cignarella, A. T., Farías, D. I. H., Bosco, C., Patti, V., & Rosso, P. (2020). **Multilingual stance detection in social media political debates**. *Computer Speech & Language*, *63*, 101075.
- Levow, G. A., Freeman, V., Hrynkevich, A., Ostendorf, M., Wright, R., Chan, J., Luan, Y. & Tran, T. (2014). **Recognition of stance strength and polarity in spontaneous speech**. In *IEEE Spoken Language Technology Workshop (SLT)*.
- Li, Y., & Caragea, C. (2021). **Target-Aware Data Augmentation for Stance Detection**. In Proceedings of NAACL-HLT.
- Li, Y., Sosea, T., Sawant, A., Nair, A. J., Inkpen, D., & Caragea, C. (2021). **P-Stance: A Large Dataset for Stance Detection in Political Domain**. In *Findings of the ACL: ACL-IJCNLP 2021* .
- Lippi, M., & Torroni, P. (2016). **Argumentation mining: State of the art and emerging trends**. ACM Transactions on Internet Technology (TOIT), 16(2), 1-25.

References

- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). **Efficient estimation of word representations in vector space**. arXiv preprint arXiv:1301.3781.
- Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., & Cherry, C. (2016). **Semeval-2016 task 6: Detecting stance in tweets**. In *Proceedings of SemEval-2016*.
- Mohammad, S. M., Sobhani, P., & Kiritchenko, S. (2017). **Stance and sentiment in tweets**. ACM Transactions on Internet Technology (TOIT), 17(3), 26.
- Pennington, J., Socher, R., & Manning, C. (2014). **GloVe: Global vectors for word representation**. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP).
- Pontiki, M., Galanis, D., Papageorgiou, H., Manandhar, S., & Androutsopoulos, I. (2015). **Semeval-2015 task 12: Aspect based sentiment analysis**. In Proceedings of SemEval-2015.
- Recasens, M., Danescu-Niculescu-Mizil, C., & Jurafsky, D. (2013). **Linguistic models for analyzing and detecting biased language**. In Proceedings of ACL.
- Sobhani, P., Inkpen, D., & Zhu, X. (2017). **A dataset for multi-target stance detection**. In *Proceedings of EACL.*

References

- Taulé, M., Martí, M. A., Rangel, F. M., Rosso, P., Bosco, C., & Patti, V. (2017). **Overview of the task on stance and gender detection in tweets on Catalan independence at IberEval 2017.** In *Proceedings of IberEval-2017.*
- Vamvas, J., & Sennrich, R. (2020). **X-stance: A multilingual multi-target dataset for stance detection**. arXiv preprint arXiv:2003.08385.
- Wei, J., & Zou, K. (2019). **EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks**. In *Proceedings of the EMNLP-IJCNLP*.
- Xu, C., Paris, C., Nepal, S., & Sparks, R. (2018). **Cross-target stance classification with self-attention networks**. *arXiv preprint arXiv:1805.06593.*
- Xu, R., Zhou, Y., Wu, D., Gui, L., Du, J., & Xue, Y. (2016). **Overview of NLPCC shared task 4: Stance detection in Chinese microblogs**. In *Natural Language Understanding and Intelligent Applications* (pp. 907-916).

A Tutorial on Stance Detection

[End of Part 1]