A Tutorial on Stance Detection

ACM WSDM 2022 Tutorial

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

[Part 2]

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

Part 2 Coverage: Outstanding Issues, Applications & Conclusions

1. Outstanding Issues
    1. Stance Detection in Data streams
    2. Context-sensitive stance detection
    3. Cross-lingual and multilingual stance detection
    4. Stance detection on non-textual data and robots

2. Application Areas
    1. Opinion surveys/polling
    2. Information retrieval
    3. Stance summarization
    4. Rumour classification
    5. Fake news detection

3. Concluding Remarks

Outstanding Issues in Stance Detection

1. **Stance Detection in Data streams**
2. **Context-sensitive stance detection**
3. **Cross-lingual and multilingual stance detection**
4. **Stance detection on non-textual data and robots**

Static Stance Detection

- Static model/hypothesis creation
- Training, Verification, and Testing
- K-folding: Cross validation
- Static hypothesis during an iteration
- Evaluation: Accuracy and other measures
- If used with temporal data: Do not use the future data to train for the past.

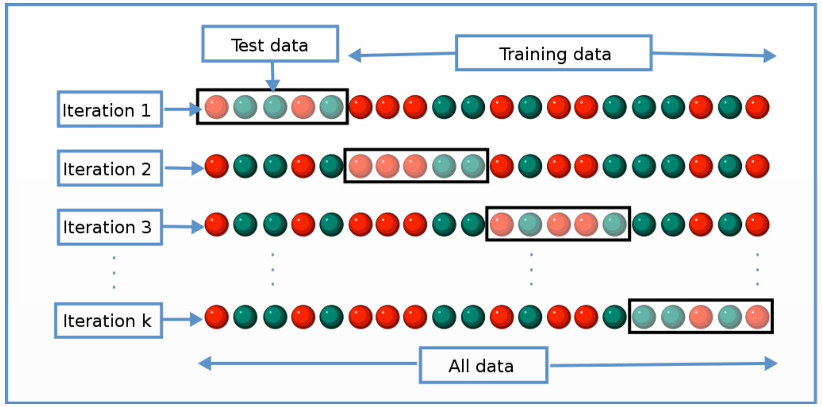

https://en.wikipedia.org/wiki/Cross-validation_(statistics)

Stance Detection in Data Streams

- Longer portion of Part 2
  - Data Stream Properties
  - Static Stance Detection
  - Data stream stance detection
  - Prequential evaluation
  - Concept Drift
  - Online stance detection scenarios

Data Streams

- Tweets
- User Comments/Engagements
- News articles
- Stock tickers
- Sensor data
- Intelligence reports
- Web clicks
- …

Data Stream Properties

- A sequence of non-stopping temporal data
- 3 Vs:
  - Velocity
  - Volume
  - Variety
- Limitless data
- Limited memory
- Limited processing time
- Data distribution may change over time

Stance Detection in a Data Stream

- No separate train and test sets
- Use interleaved-test-then-train method: use each instance first to **test** the model, and **then** to **train** the model
- Dynamic change of the classification model (variations possible)
- Stance detection in tweets stream

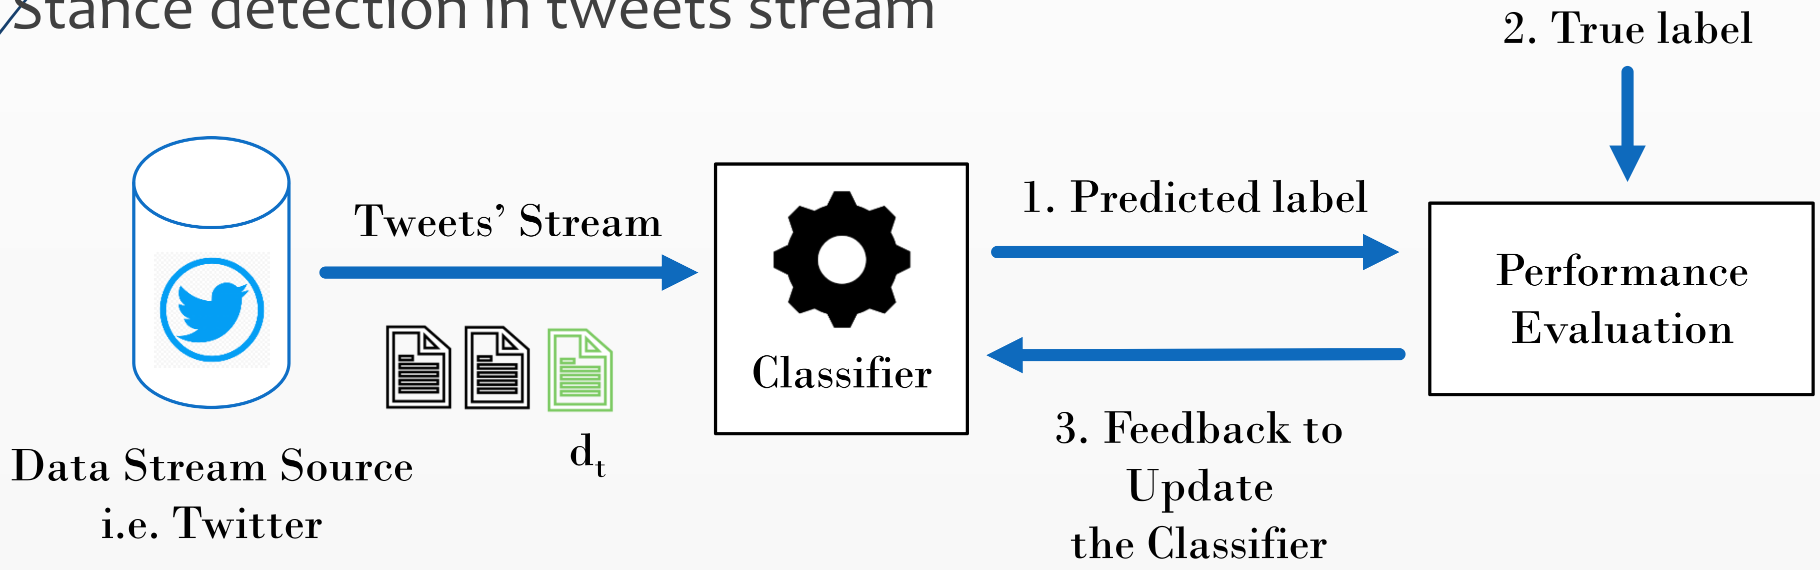

Prequential Evaluation vs. Overall Evaluation

- It assumes that the correct label is available right after testing.
- Interleaved-test-then-train evaluation.
- While calculating prequential accuracy, each data instance is used for two purpose: First testing then for training.
- Evaluation window size.
- Overall accuracy at time t, is the accumulated prequential accuracy calculated for all the stream data until time t.

V. M. A. Souza, D. F. Silva, J. Gama, G. E. A. P. A. Batista, "Data stream classification guided by clustering on nonstationary environments and extreme verification latency," *SDM*, pp. 873-881, 2015

Prequential Evaluation: Example – Window Size: 4

0: Wrong prediction , 1: Correct prediction

Prediction	1,0,0,1	1,0,1,1	1,1,1,1	0,0,1,1	0,1,0,0

Prequential acc	50%	75%	100%	50%	25%
Overall acc	50%	62.50%	75%	68.75%	60%

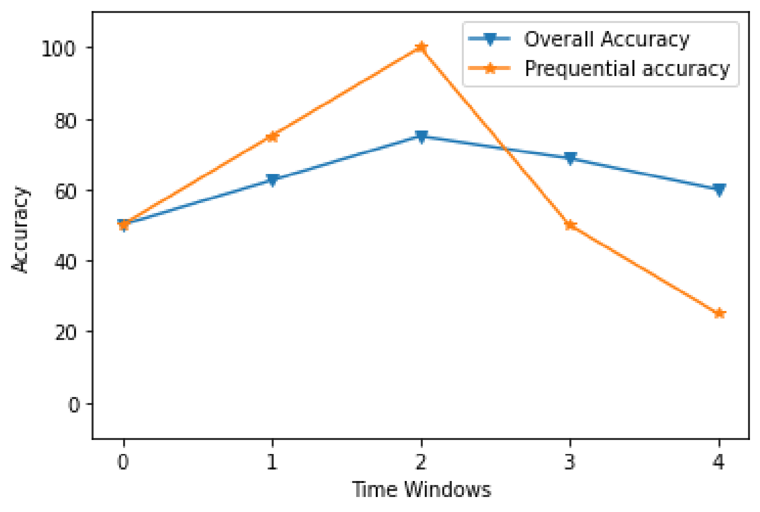

Analysis Tools: MOA and Scikit-multiflow

- MOA (Massive Online Analysis) is one of the popular frameworks for data stream analysis in JAVA.
- Scikit-multiflow is a framework which is the equivalent of MOA in Python
- Both contain some of the commonly known machine learning algorithms e.g. Hoeffding Tree, Naïve Bayes,…
- They can simulate the stream environment by generating temporal data with desired characteristics.
- https://moa.cms.waikato.ac.nz/
- https://scikit-multiflow.github.io/

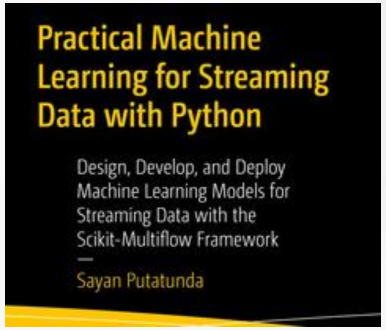

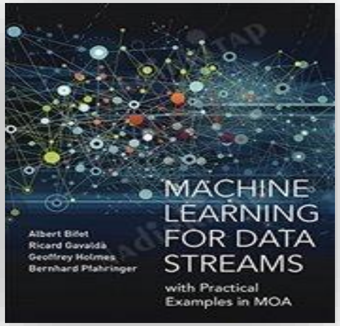

Dynamic Online Processing in Data Streams

1. **Concept Drift**: changes in the data stream.
2. Be able to react to concept drift: Learn the variations in data
3. Types of concept drift (what happens?): Real, Virtual
4. Types of concept drift (how it happens?): Abrupt, Gradual, …
5. Concept drift detection and handling

Bifet, Albert, et al. MOA: Massive online analysis. *Journal of Machine Learning Research* 11. 1601-1604 (2010).

Concept Drift: What Happens?

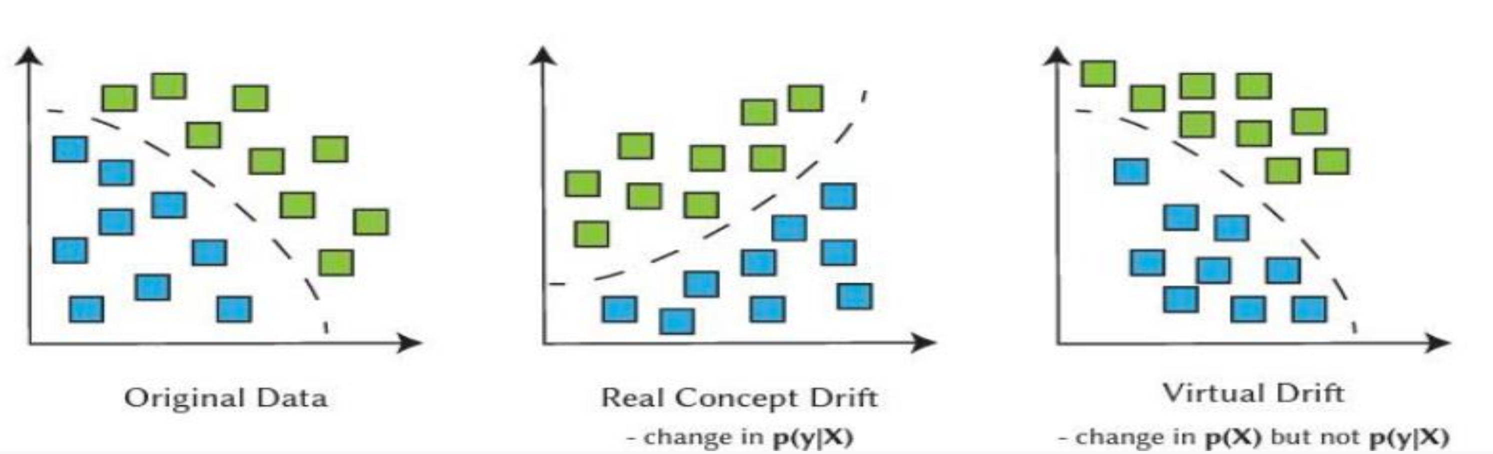

Decision boundary changes

Input characteristics changes

J. Gama, I. Žliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia. 2014. A survey ˙ on concept drift adaptation. ACM Computing Surveys (CSUR) 46, 4 (2014), 44.

Types of Concept Drift: How it Happens?

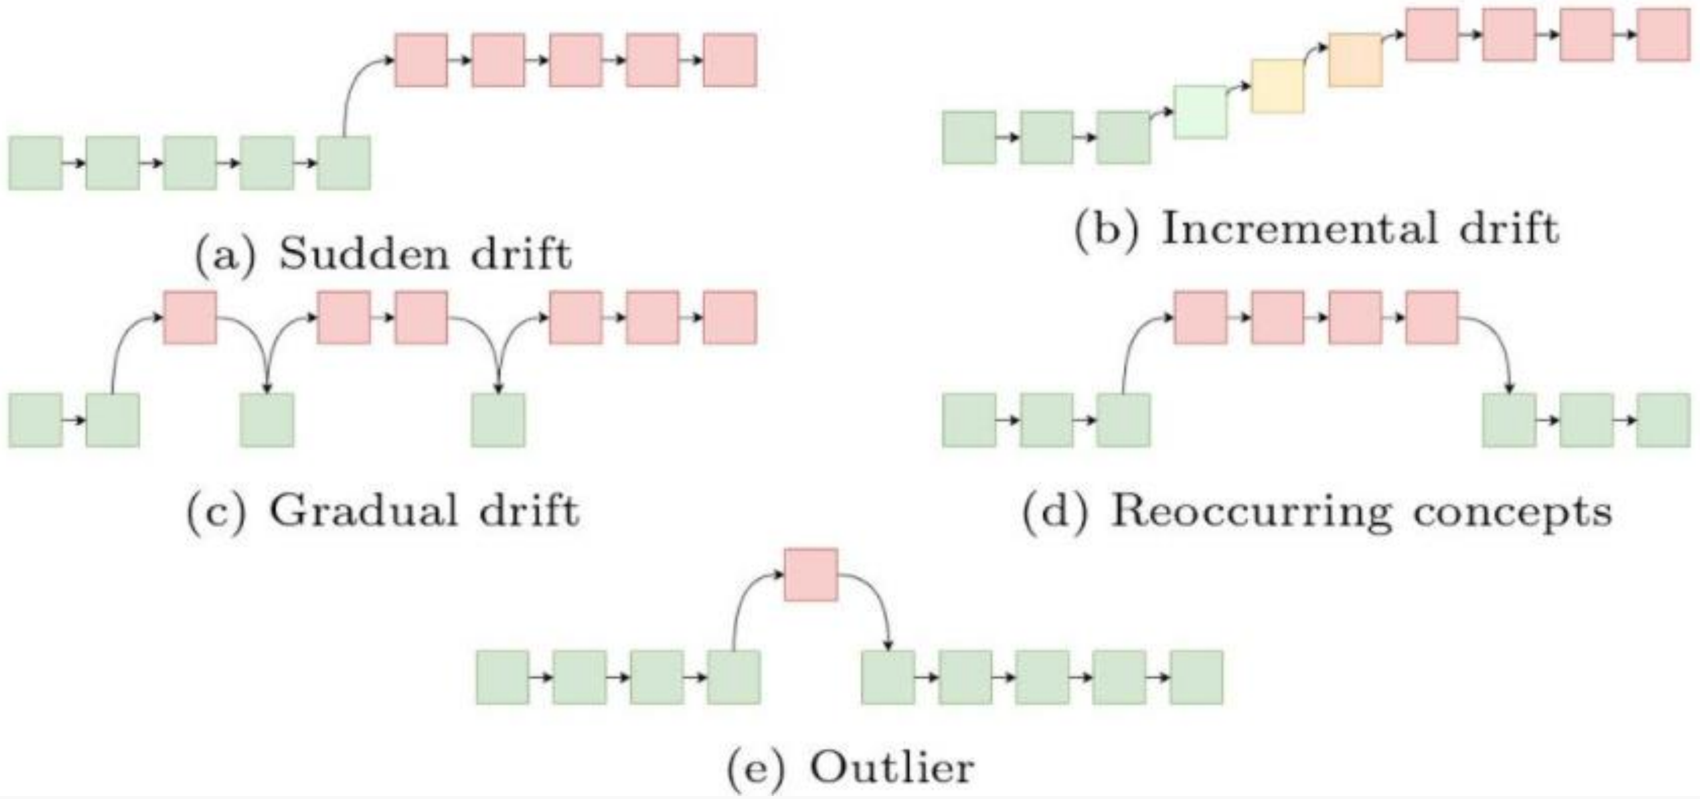

What to do for Concept Drift?

- Concept drift detection:
  - Prediction performance decrease
  - Using an Algorithm: supervised vs. unsupervised
- Concept drift handling
  - Active: Detect and retrain
  - Passive: Let the system handle

Lu, A. Liu, F. Dong, F. Gu, J. Gama and G. Zhang. Learning under Concept Drift: A Review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346-2363 (2019).

Concept Drift Detection

- Concept drift detection algorithms find the drift point in a data stream.
- They usually detect the change based on the changes in the distribution of data over time.
- Detectors send an alarm in case of any changes and the system adapts to the changes based on a concept drift adaptation algorithm.

Unsupervised Concept Drift Detection: Old vs. New

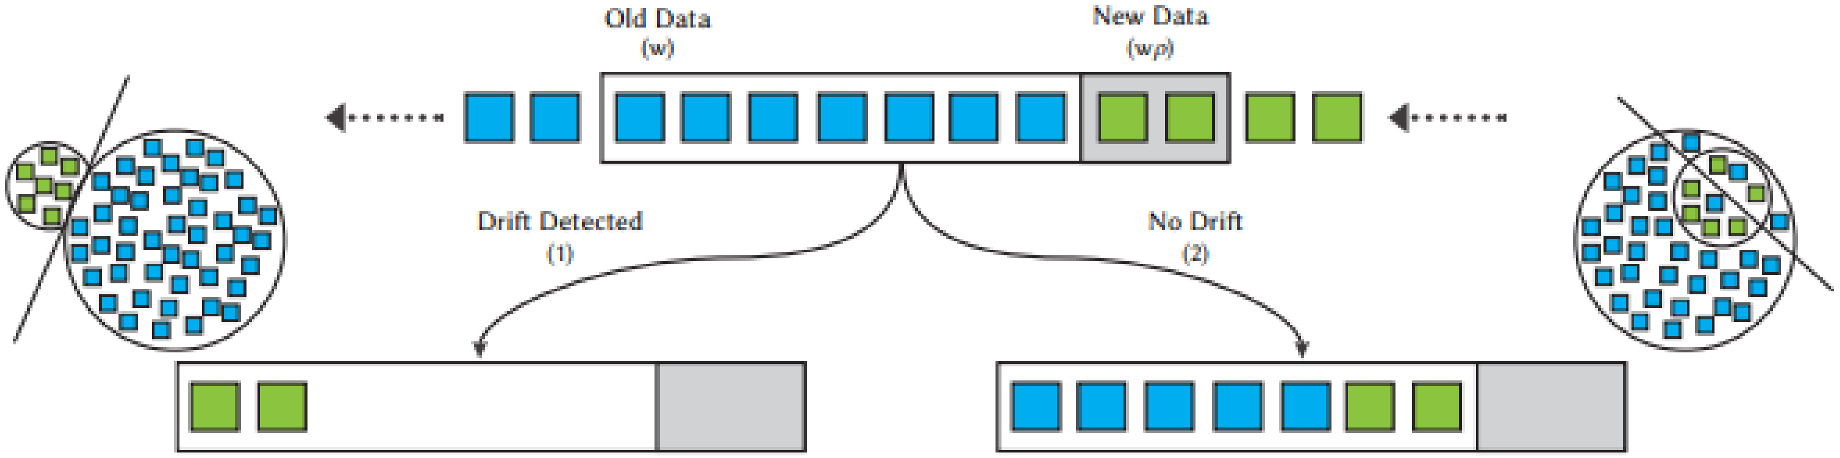

Figure 1: Drift detection workflow: (1): Drift detected. The old and the new data are separable. Samples from the old portion are discarded and partially filled with the samples from the new. (2): No drift. These sets are nested. The oldest $w\rho$ samples are removed and the window is shifted to left where the samples from the new fill the space that becomes empty.

Ö Gözüaçık, Baç üyükçakır, H. Bonab, F. Can, F. "Unsupervised concept drift detection with a discriminative classifier." (Short paper.) *The 28th ACM International Conference on Information and Knowledge Management*, 2019.

Unsupervised Concept Drift Detection: From Us or Not

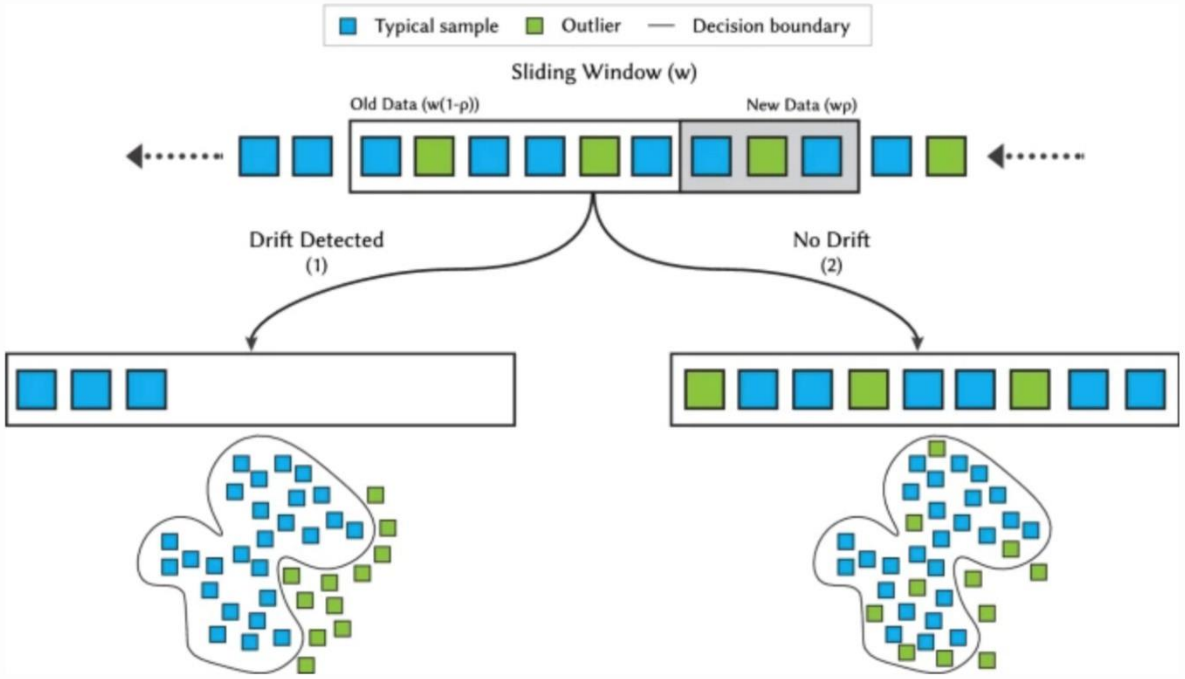

Drift detection workflow: (1): Drift detected. The percentage of outliers exceed the threshold (ρ). There is a change in the distribution of the data. Samples from the old portion are discarded and are partially filled with samples from the new data window. (2): No drift. There is no change in the data distribution. The oldest sample is removed and the window is shifted to the left, filling the empty space

Ö. Gözüaçık and F. Can, "Concept learning using one-class classifiers for implicit drift detection in evolving data streams." *Artificial Intelligence Review* 54 (5), 3725-3747.

Concept Drift Handling

1. Update the model over time in certain time order. For example, update the model every week and learn the new data.
2. Incrementally learn the new data over time.
3. Assign more weights to incoming data right after the drift.
4. Restart learning from the drift point.
5. Use ensemble of weak classifiers to adjust to the changes.

Using Ensemble Methods

- Ensemble learning is a paradigm where multiple machine learning algorithm results are combined to get more accurate results.
- Each component can use a different model, different types of features (text, speech, video,…)
- Some Ensemble methods:
  1. Bagging: sampling with replacement and averaging
  2. Boosting: combine weak learners (better than random) to obtain a strong learner
  3. Stacking: a learner that learns classifiers: GOOWE
- Using these methods and voting systems like majority voting, the system can adapt to new changes.
- GOOWE: An example ensemble method.

GOOWE: Geometrically Optimum…

- Notations and Assumptions:
  - **p class labels** as C = (C1, C2… Cp)
    multi-class problem (multi-labels is also possible)
  - **m classifier** systems as CS = (CS1, CS2… CSm)
  - For each of these Instances (I_i), every classifier system CS_j returns a set of scores as $S_{ij} = (S_{ij}^1, S_{ij}^2 \ldots S_{ij}^p)$

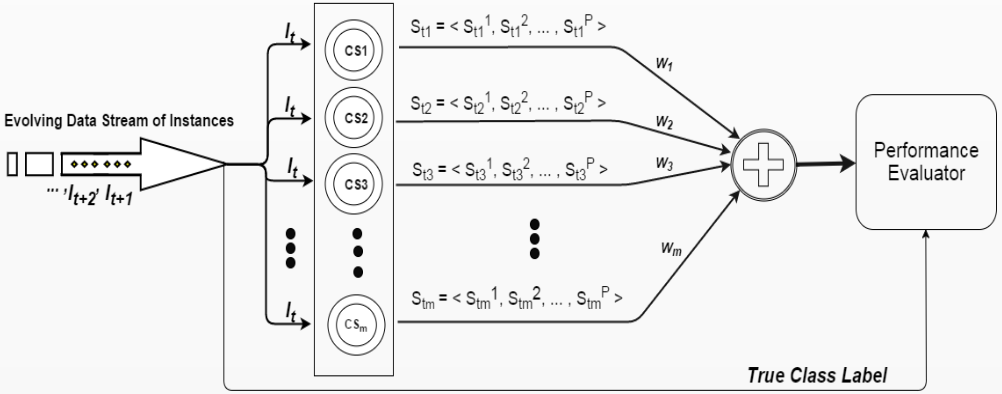

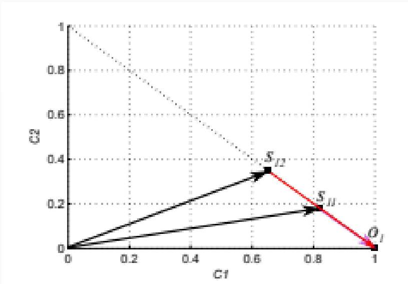

H. Bonab and F. Can, "GOOWE: Geometrically optimum and online-weighted ensemble classifier for evolving data streams", *ACM Transactions on Knowledge Discovery from Data*, 12(2), 1-25:33., 2018.

Law of Diminishing Returns (m = p): Is It Counter Intuitive?

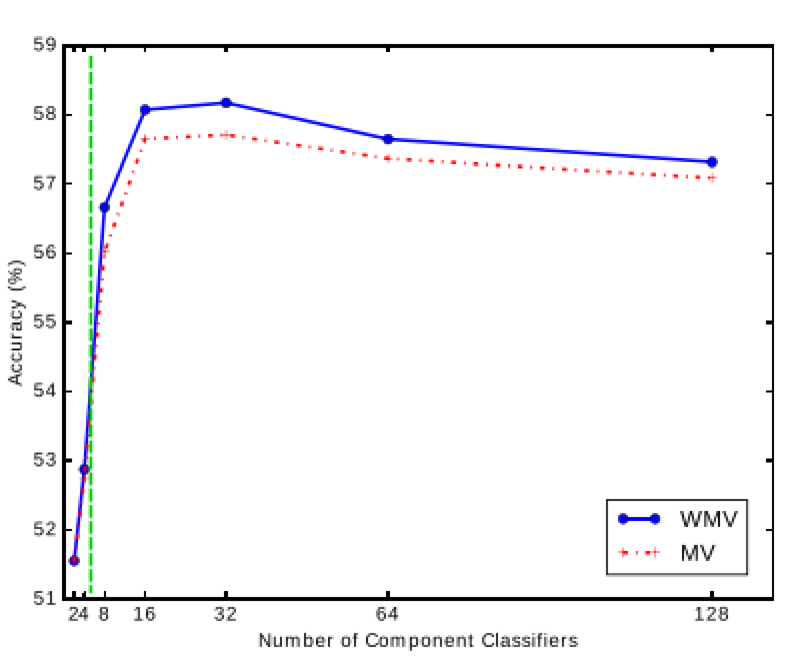

The highest effectiveness is observed much closer to the theoretically ideal m=p green line rather than the maximum number of components.

H. Bonab and F. Can, 2018.

Online Stance Detection Scenario

- No true labels are available: Extreme verification latency
- A set of initial data items with labels are available: Initial supervision
- An incremental concept drift environment: slow change in the input data items
- Use clustering and classification together

incremental change

Vinícius M. A. de Souza, Diego Furtado Silva, João Gama, Gustavo E. A. P. A. Batista:
Data Stream Classification Guided by Clustering on Nonstationary Environments and Extreme Verification Latency. SDM 2015: 873-881

Online Stance Detection Scenario (cont.)

1. Assign labels to a pool of data items by human annotation (c: number of labels: approve, disapprove, neutral)
2. Build an initial classification model using the initial labelled data
3. Divide the initial labelled data into c clusters (approve, disapprove, neutral)
4. Receive unlabelled new data items and insert them into a new pool
5. Detect the stance label of pooled data items with the available classification model
6. Use a clustering algorithm and obtain clusters for the pool
7. Map newly formed clusters to previously labelled clusters: By this way assign labels to new clusters (their members)
8. Replace the classification model using the labelled data items of the new clusters
9. Go to step 4

Vinícius M. A. de Souza, Diego Furtado Silva, João Gama, Gustavo E. A. P. A. Batista:
Data Stream Classification Guided by Clustering on Nonstationary Environments and Extreme Verification Latency. SDM 2015: 873-881

Use of Concept Drift Detection in This Scenario

If no concept drift is detected no need to renew the model.

Concept drift detection runs concurrently.

Advantage of Concept Drift Detection: Time efficiency: No delay in the results.

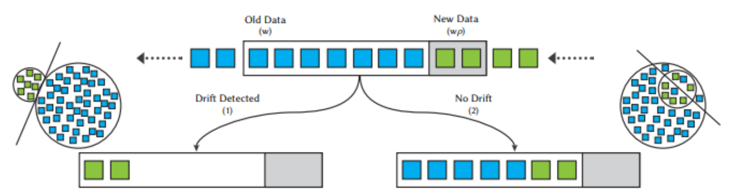

Figure 1: Drift detection workflow: (1): Drift detected. The old and the new data are separable. Samples from the old portion are discarded and partially filled with the samples from the new. (2): No drift. These sets are nested. The oldest $w\rho$ samples are removed and the window is shifted to left where the samples from the new fill the space that becomes empty.

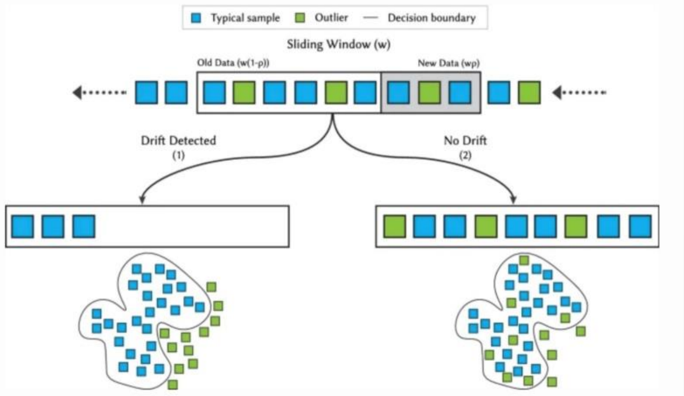

Drift detection workflow: (1): Drift detected. The percentage of outliers exceed the threshold (ρ). There is a change in the distribution of the data. Samples from the old portion are discarded and are partially filled with samples from the new data window. (2): No drift. There is no change in the data distribution. The oldest sample is removed and the window is shifted to the left, filling the empty space

Vinícius M. A. de Souza, Diego Furtado Silva, João Gama, Gustavo E. A. P. A. Batista:
Data Stream Classification Guided by Clustering on Nonstationary Environments and Extreme Verification Latency. SDM 2015: 873-881

How to Incorporate Other Types of Concept Drifts

A small portion of labelled data is available spread in the data stream (1%, 5%, ... 100%).

Availability of labelled data at different time instances makes system adaptable to different concept drifts.

Clusters are constructed with maintenance.

Label propagation is used to label clusters.

Classification using clustering results.

An ensemble of k-NN classifiers are used.

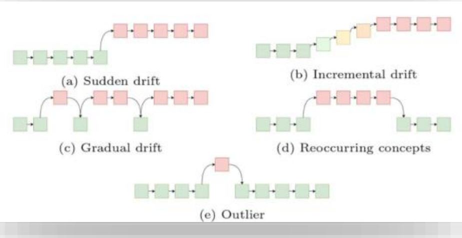

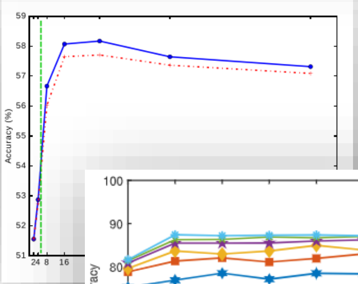

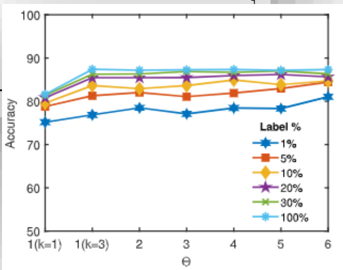

S. U. Din, J. Shao, J. Kumar, W. Ali, J. Liu, J., and Y. Ye, "Online reliable semi-supervised learning on evolving data", *Information Sciences*, 525 (2020) pp. 153-171, 2020.

Evaluation of Online Stance Detection Applications

- Prequential Evaluation
- User study: User satisfaction study
- SDR: Scalability-Dynamism-Robustness
- User satisfaction is the final decision level.

Katarzyna et al. How to design the fair experimental classifier evaluation. Applied Soft Computing Journal, 104(2021),

Classifier Performance Evaluation: SDR

Scalability

- Data streams with different size & characteristics

Dynamism

- Different types of concept drifts

Robustness

- Missing labels, incorrect labels
- Concept evolution: Brandnew labels
- Missing/Incorrect features

Outstanding Issues on Stance Detection

1. **Stance Detection in Data streams**
2. **Context-sensitive stance detection**
3. **Cross-lingual and multilingual stance detection**
4. **Stance detection on non-textual data and robots**

Context-sensitive Stance Detection

- How we see the world: Determined by our context.
- Context:
  - Google definition: “the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood.”
  - Oxford Learner’s Dictionaries: “the situation in which something happens and that helps you to understand it.”
- Principle of locality:
  - Physics: “an object is directly influenced only by its immediate surroundings.”
  - Computer Science: Temporal locality, Spatial locality.
- How we make decisions
  - How we proceed is affected by our locality/context.

Shu, K. et al. Beyond News Contents: The Role of Social Context for Fake News Detection. WSDM 2019: 312-320

Cross-lingual and Multilingual Stance Detection

- Experiments need datasets.
- Datasets: Annotation ➔ difficult!
- How about borrowing from languages with datasets.
  - Annotated dataset in a given language (e.g., English) can be automatically translated into the target languages .
  - Train a sentiment analysis model using recurrent neural networks with reviews in English. We then translate reviews in other languages and reuse this model to evaluate the sentiments (Can et al, 2018).

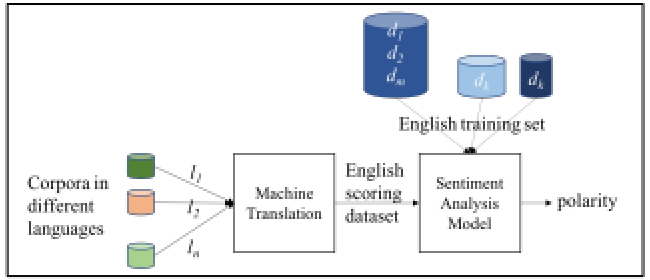

Figure 1: Multilingual sentiment analysis approach.

Can, E., F., Ezen-Can, A., and Can, F. "Multilingual sentiment analysis: An RNN-based framework for limited data," arXiv preprint arXiv:1806.04511 (2018).

Stance Detection on Non-textual Data and Robots

- Sentiment Analysis using facial, verbal, vocal input (Shrivastava et al, 2018) ➔ Stance detection!

- Joint stance detection from these different modalities
  - Presidential debate: Audience is give buttons for approve/disapprove. A facial sentiment/stance analysis?
  - Using such information by emotional robots: Involves several ethical concerns.
  - Combining stance obtained from different modalities using an ensemble approach: Lie detection.

V. Shrivastava, V. Richhariya, and V. Richhariya, "Puzzling Out Emotions: A Deep-Learning Approach to Multimodal Sentiment Analysis," in 2018 International Conference on Advanced Computation and Telecommunication (ICACAT), pp. 1-6, 2018.

Application Areas

1. **Opinion surveys/polling**
2. **Information retrieval**
3. **Stance summarization**
4. **Rumour classification**
5. **Fake news detection**

Opinion Surveys/Polling

- Topics include
  - Political/ideological/social debates
  - Product reviews
  - Elections/referendums
- Traditional
- Using social media, e.g. using Twitter (Sen et al., 2020)

  Emphasize the increasing importance of digital traces in polling
  - Different datasets different targets
  - Reliability
  - Direct vs. Indirect stance
- By means of automatic stance detection, whether a community is in favor of or against a topic of interest can be estimated, replacing (or complementing) the traditional practices of performing surveys/polls.

You can never be sure!

Sen, I., Flöck, F., Wagner, C. On the reliability of detecting approval of political actors in tweets. Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing, 2020, 1413-1426.

Information Retrieval

1. Stance detection for search result personalization:
   - Determine user stance on different issues and provide results that would match their preferences.
   - Result: Eco room, Filter Bubble.
   - It is possible to change user opinions.
2. Answering multi-perspective queries:
   - “Is treatment x effective for disease y?” Answering such queries requires stances (support or oppose). (Sen et al., 2018).

Eli Pariser. 2011. The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think. The Penguin Press, New York.
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. 273–281.

Stance Summarization

- Observations:
  - Social media data items: continuous, endless, fast!
  - Hard to follow: Summarization of stances is needed (textual & visual)
- Studies:
  - Jang, Allan (2018): For tweets, defined as a ranking task, and use representative tweets as stance summary.
  - Wei et al. (2021): For movie critics and debates, first selects summary worthy documents, second stage uses MMR (maximal marginal relevance). Aims to find salient and non-redundant opinions.
- Promising research area:
  - Google Scholar search: “Stance Summarization” *returns 21 results* (February 1, 2022).

Myungha Jang and James Allan. 2018. Explaining controversy on social media via stance summarization. In Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval.

Penghui Wei, Jiaho Zhao, Wenji Mao, A Graph-to-Sequence Learning Framework for Summarizing Opinionated texts. IEEE/ACM Transactions on Speech and Language Processing Vol 29, 2021.

Rumour Classification

Rumour is defined as a piece of information that has not yet been verified.

They are misinformation with no intention of deceive. They can be spreaded for entertainment.

- The work by Zubiaga et al. (2018a) has four basic components in rumour processing
  - rumour identification,
  - tracking,
  - classification of rumour stance {Supporting, Denying, Querying, Commenting}
  - Veracity {True, False, Unverified}

Rumour stance detection and fake news stance detection can employ shared datasets.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018a. Detection and resolution of rumours in social media: A survey. ACM Computing Surveys (CSUR) 51, 2 (2018), 32.

Fake News Detection

They can be detected from their (Zhou, Zafarani, 2020)

- False knowledge they carry
- Writing style
- Propagation pattern
- Credibility of its resources

Immediate detection can be important (Ksieniewicz et al., 2020)

- Fake news detection in data streams
- Uses different numbers of features
- Ensembles

Zhou, X, Zafarani, R. A survey of fake news: Fundamental theories, detection methods, and opportunities, ACM Computing Surveys, 53(5): Article 109. September 2020.

Ksieniewicz, P. et al., Fake news detection from data streamsIJCNN, 2020.

Fake News Detection

Growth of fake news? ➔ Effects on democracy and public trust.

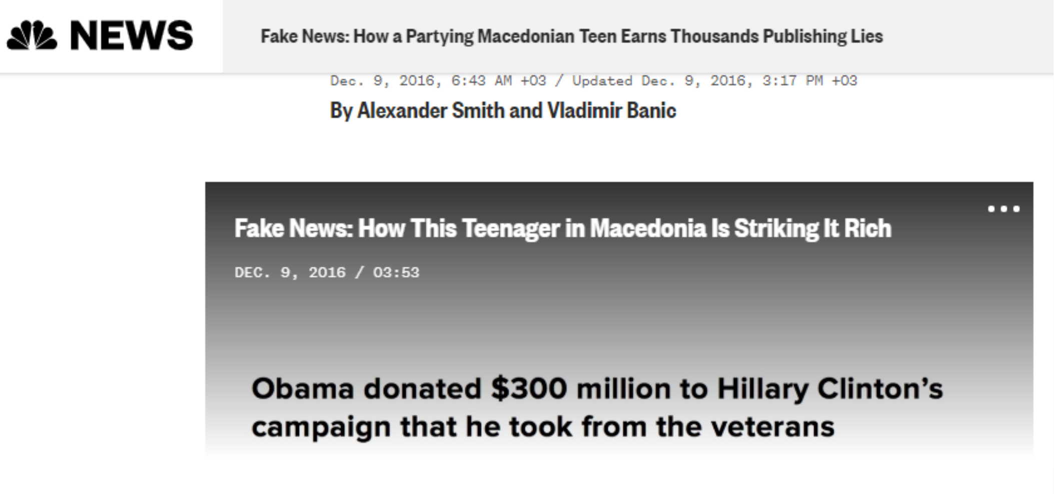

A view of Veles GEORGI LICOVSKI / EPA

Orson Welles, War of the Words, 1938

NBC News: https://www.nbcnews.com/news/world/fake-news-how-partying-macedonian-teen-earns-thousands-publishing-lies-n692451

Concluding Remarks

- Progress So Far
- Stance detection and social responsibility
- Resilient Systems
- Future Work

Progress So Far

- Several approaches
- Annotation Guidelines
- Datasets
- Evaluation Metrics
- Software and Tools
- Several Application Areas
- Several Future Research Possibilities

Stance Detection and Social Responsibility

- Computing technologies and ethics
  - “With great power comes great responsibility. Technology is now one of the most powerful forces shaping society, and we are responsible for it.” Moshe Vardi
  - Thinkers of the book *Architects of Intelligence*
  - What is waiting for us?
- Unintended) consequences of our research…
  - Following people determining their stance
  - Pushing items in the social media for changing their stance
- What we **should not** create?”

Vardi M. Y. (2018) Computer Professionals for Social Responsibility Commun ACM 61(1): 9.
Ford, Martin, *Architects of Intelligence: The Truth about AI from the People Building it*, 2018, UK: Packt Publishing. (Interviews with 23 scientists.)

Abstraction & Resilience

- Formalization
  - Solving a problem using a formal approach
  - Transferring a solution to another similar field
- Systems must be resilient
  - Adaptable to disruptive environments
  - Robust to manipulation: Secure

Vardi MY (2020) Efficiency vs. resilience: what COVID-19 teaches computing. Commun ACM 63(5):9.
Aho, A., Ullman J. Abstractions, their algorithms, and their compilers, Commun ACM 65(2): 76-91.

Future Work

- Common test collections: Larger and Different Languages
  - Replicable research
  - Objective comparison of results
- Synthetic data stream generation
  - Difficult: Natural language
  - Difficult: explicit vs. implicit stance, difficult for humans too

Future Work

- Libraries/Frameworks
  - Provides baselines
  - Open to additions in terms of algorithms and datasets
- Development of Online Systems
  - Online Stance Detection Evaluation
  - Definition & Evaluation Metrics
  - How generalizable are the test results: theory to practice
  - Lab Evaluation vs. Practical Evaluation

Future Work

- From Individualized Implementations to More General Ones
  - Context-sensitive
  - Cross-lingual and multilingual
  - Non-textual data and robots

Ford, Martin, Architects of Intelligence: The Truth about AI from the People Building it, 2018, UK: Packt Publishing.

References

1. Aho, A., Ullman J. Abstractions, their algorithms, and their compilers, Commun ACM 65(2): 76-91.
2. A. Bifet, G. Holmes,R. Kirkby, and B. Pfahringer, "Massive online analysis," Journal of Machine Learning *Research* 11, May (2010), pp. 1601-1604.
3. B. Bharathi, J. Bhuvana, and N. N. A Balaji, "SSNCSE-NLP @ EVALITA2020: Textual and Contextual Stance Detection from Tweets Using Machine Learning Approach," *CEUR-WS*, Vol. 2765, Paper 151, 2020.
4. H. Bonab, and F. Can, "Less is more: A comprehensive framework for the number of components of ensemble classifiers," *IEEE Transactions on Neural Networks and Learning Systems*, 30(9), pp. 2735-2745, 2019.
5. H. Bonab, and F. Can, "GOOWE: Geometrically optimum and online-weighted ensemble classifier for evolving data streams," *ACM Transactions on Knowledge Discovery from Data*, 12(2), 25:1-25:33, 2018.
6. E. F. Can, A. Ezen-Can, A., and F. Can, "Multilingual sentiment analysis: An RNN-based framework for limited data," *arXiv preprint arXiv:1806.04511*, 2018.
7. S. U. Din, J. Shao, J. Kumar, W. Ali, J. Liu, J., and Y. Ye, "Online reliable semi-supervised learning on evolving data," *Information Sciences*, 525 (2020) pp. 153-171, 2020.
8. M. Ford, *Architects of Intelligence: The Truth about AI from the People Building it*, UK: Packt Publishing, 2018.
9. J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, 46(4), 44:1–44:37, 2014.
10. Ö Gözüaçık, Baç üyükçakır, H. Bonab, and F. Can, F. "Unsupervised concept drift detection with a discriminative classifier." (Short paper.) *The 28th ACM International Conference on Information and Knowledge Management*, 2019.mber
11. Ö. Gözüaçık and F. Can, "Concept learning using one-class classifiers for implicit drift detection in evolving data streams." *Artificial Intelligence Review* 54 (5), 3725-3747.
12. Ahmed Hamdi et al. Multilingual Dataset for Named Entity Recognition, EntityLinking and Stance Detection in Historical Newspapers, ACM SIGIR 2021 Proceedings

References

9. Katarzyna et al. “How to design the fair experimental classifier evaluation,” *Applied Soft Computing Journal*, 104, 2021.
10. Ksieniewicz, P. et al., “Fake news detection from data streams,” IJCNN, 2020.
11. J. Lu, A. Liu, F. Dong, J. Gu, J. Gama, and G. Zhang, 2019, Learning under Concept Drift: A Review, *IEEE Transactions on Knowledge and Data Engineering*, 31(12), pp. 2346-2363, 2019.
12. MOA: https://moa.cms.waikato.ac.nz/
13. E. Pariser, *The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think*. The Penguin Press, New York, 2012.
14. Sakala Venkata Krishna Rohit and Navjyoti Singh. 2018. Analysis of speeches in Indian parliamentary debates. arXiv preprint arXiv:1808.06834 (2018).
15. Scikit multiflow: https://scikit-multiflow.github.io/
16. A. Sen, M. Sinha, S. S. Mannarswamy, and S Roy, “Stance classification of multi-perspective consumer health information,” *Proceedings of the ACM India Joint International Conference on Data Science and Management*, 2018.
17. I. Sen, F. Flöck, and C. Wagner, “On the reliability of detecting approval of political actors in tweets,” Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing, pp. 1413-1426, 2020.
18. V. Shrivastava, V. Richhariya, and V. Richhariya, "Puzzling Out Emotions: A Deep-Learning Approach to Multimodal Sentiment Analysis," *Proceedings of International Conference on Advanced Computation and Telecommunication (ICACAT)*, pp. 1-6, 2018.
19. Shu, K. et al. Beyond News Contents: The Role of Social Context for Fake News Detection. WSDM 2019: 312-320.
20. V. M. A. Souza, D. F. Silva, J. Gama, G. E. A. P. A. Batista, “Data stream classification guided by clustering on nonstationary environments and extreme verification latency,” *SDM*, pp. 873-881, 2015
21. G. I. Webb, R. Hyde, H. Cao, H.-L. Nguyen, and F. Petitjean,”Characterizing concept drift,” *CoRR abs/1511.03816, 2015.*
22. X. Zho and R. Zafarani, “A survey of fake news: Fundamental theories, detection methods, and opportunities,” *ACM Computing Surveys*, 53(5): Article 109. September 2020.

Thank You