

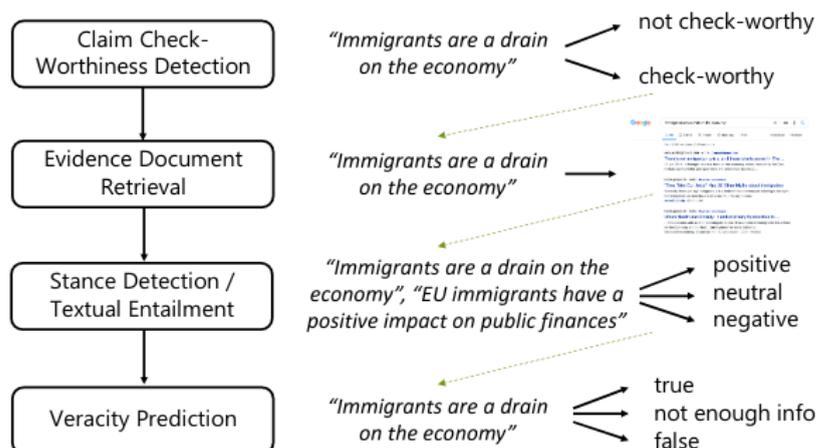
Fake News and Fact-checking: Inference, Multilingual and Few-shot Approaches

Proposers: Rodrigo Agerri

Contact: rodrigo.agerri@ehu.eus

Description

Determining the veracity of a given document or story, namely, whether it is fake or legitimate, is a very complex task, even for expert fact-checkers. Thus, it is common to break down the fake news detection task in different stages, as shown by the following figure (Augenstein 2021):



This illustrates a typical content-based fact checking pipeline, starting with the detection of check-worthy claims, and ending with the verification of a claim's veracity.

We are particularly interested in the stance detection and veracity prediction (claim verification) tasks. The first task consists of establishing what other news sources are saying about the given document or story (whether they agree, disagree, etc. with the news story), namely, determining their stance with respect to that document or news story. This decision is supported by two main ideas:

1. A Stance Detection system should allow a human fact checker to enter a document (headline, message, claim, etc.) and retrieve the top documents from other news sources that agree, disagree or discuss the given document.
2. Based on the previous step, it would be possible to build a "truth-labelling" system based on the weighted credibility of the various news organizations from which the stance has been retrieved.

Thus, given a headline and a follow-up news document, the task consists of determining whether the follow-up document agrees, disagrees, discusses, or comments something unrelated with respect to the given news headline. As an example of the task based on this definition, consider the following news snippets:

- **Headline:** “Robert Plant Ripped up \$800M Led Zeppelin Reunion Contract.”
- **Follow-up 1 Agrees:** “. . . Led Zeppelin’s Robert Plant turned down £500 MILLION to reform supergroup . . .”
- **Follow-up 2 Disagrees:** “. . . No, Robert Plant did not rip up an \$800 million deal to get Led Zeppelin back together . . .”
- **Follow-up 3 Discusses:** “. . . Robert Plant reportedly tore up an \$800 million Led Zeppelin reunion deal . . .”
- **Follow-up 4 Unrelated:** “. . . Richard Branson’s Virgin Galactic is set to launch SpaceShipTwo today . . .”

The second task, veracity prediction or claim verification, aims to establish whether a given claim is supported/unsupported by a given set of contextual evidence, or whether there is not enough to perform the inference (Thorne et al. 2018).

In this project we propose to study Natural Language Inference approaches to detect the stance of a given text with respect to the original trigger headline. Furthermore, we would also like to extend this task to other languages of interest, for which data is very scarce. Taking these two components, we would like to mostly experiment with few-shot learning approaches (Schick and Schütze 2021).

Objectives

The candidate may choose between the following objectives:

1. (Semi-) Automatic development of a dataset of fake news detection in languages other than English.
2. Experiment with deep learning approaches for fake news, claim verification and fact-checking, including Natural Language Inference and few-shot learning based on prompting.
3. Experiment on crosslingual approaches to claim verification or stance detection.

Tasks and Plan

- Month 1: Start of the project, defining the objectives and tasks.
- Month 2: Start experiments. Optionally, it is recommended for the candidates to attend the "Seminar on language technologies. Deep Learning (LAP 18). https://ixa.si.ehu.es/master/programa_html
- Months 3-5: Experiments and final development.
- Final month: Writing up.

References

Isabelle Augenstein (2021). [Towards Explainable Fact-checking](#). In arxiv.

Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk and Veselin Stoyanov [XNLI: Cross-lingual Sentence Understanding through Inference \(corpus page\)](#). In EMNLP, 2018.

Adina Williams, Nikita Nangia, and Samuel R. Bowman. [A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference \(corpus page\)](#) In NAACL, 2018

Hanselowski et al. (2018) [A Retrospective Analysis of the Fake News Challenge Stance Detection Task](#). In COLING.

Schick, Timo and Hinrich Schütze. "[True Few-Shot Learning with Prompts -- A Real-World Perspective](#)." (2021). In arxiv.

Thorne, James, Andreas Vlachos, Christos Christodoulopoulos and Arpit Mittal. "[FEVER: a Large-scale Dataset for Fact Extraction and VERification](#)." In NAACL (2018).