

Demoting Outdated ‘Truth’ With Machine Learning

Originally published at <https://www.unite.ai/demoting-outdated-truth-with-machine-learning/>



Sometimes the truth has an expiry date. When a time-limited claim (such as ‘masks are obligatory on public transport’) emerges in search engine rankings, its apparent ‘authoritative’ solution can outstay its welcome even by many years, outranking later and more accurate content on the same topic.

This is a by-product of search engine algorithms’ determination to identify and promote ‘long-term’ definitive solutions, and of their proclivity to prioritize well-linked content that maintains traffic over time – and of an increasingly circumspect attitude to newer content in the emerging age of [fake news](#).

Alternately, devaluing valuable web content simply because the timestamp associated with it has passed an arbitrary ‘validity window’ risks that a generation of genuinely useful content will be automatically demoted in favor of subsequent material that may be of a lower standard.

Towards redressing this syndrome, a new [paper](#) from researchers in Italy, Belgium and Denmark has used a variety of machine learning techniques to develop a methodology for *time-aware evidence ranking*.

Beyond Outdated Answers

The paper is authored by researchers from the European Commission at the Joint Research Centre (JRC) in Ispra, the Katholieke Universiteit at Leuven, and the University of Copenhagen.

The work considers four temporal ranking methods applied over three fact-checking methodologies, each with a different approach to evidence ranking, and offers a novel methodology for ranking that uses evidence timestamps as a ‘gold standard’. The study shows that time-aware evidence ranking improves the perspicacity of results, and also improves the authority and veracity predictions of time-sensitive facts and claims.

The research is offered as a possible adjunct to later or existing systems, and is designed to aid in research, and as a possible additional factor for inclusion in developing new and evolved search engine algorithms.

The work models the temporal dynamics of evidence for content-based fact-checking, and outperforms the ‘semantic similarity’ approaches adopted by typical search engine ranking algorithms. The model trained by the researchers uses an optimized [learning-to-rank](#) function which can be easily superimposed into an existing fact-checking architecture. The researchers contend that the system is a novel contribution to automated fact-checking.


Amending Multiple Fact-Checking Architectures

The researchers imposed their time-constrained factoring onto three existing fact-checking architectures. The first of these is the Bidirectional Long Short Term Memory (BiLSTM) model proposed in the [MultiFC dataset](#) released in 2019.

The second is a modification to the first, with a unidirectional Recurrent Neural Network (RNN) replacing the LSTM component.

The third model the researchers used is a [DistilBERT transformer](#) from the Hugging Faces library, a distilled version of Google’s [BERT](#) NLP model.

Across all three architectures, the researchers applied a ListMLE loss, from research led by Microsoft, which has contributed consistently to novel fact-checking research over the last two decades.



The two primary fact-checking models to which the research team has added a temporal component as a filter for authority and subsequent ranking values. Source: <https://arxiv.org/pdf/2009.06402.pdf>

Timestamp values were extracted from the training metadata, and included as ranking factors in each model.

Testing

Experimental evaluation for the system involved the use of the MultiFC dataset, since it is currently the only high-volume open source dataset available for this particular research interest. MultiFC contains 34,924 real-world claims obtained from 26 different fact-checking domains, including Snopes and the Washington Post.

Prediction of each claim's veracity are augmented by ten evidence snippets provided by the Google Search API, and predictions obtained via a confluence of elements, including speaker, tags and categories.

Very often the relevant time-stamp is not necessarily the one that's contained in the metadata; an article may refer to events from previous times, and in this case the researchers' systems had to take care to extract and convert that data directly from the text. Without this process, a 're-run' of outdated news will tend to give it a new gloss, particularly in the case of high authority sites, propagating the outmoded data.

The dates were extracted with a Python routine, and the official metadata dates tested for consistency of formatting (since, for instance, US and UK date-stamp formatting is different). When manually verified, zero errors were found in the time-stamp metadata.

Results

Against a manual check of the automated results, the researchers found that time-aware evidence ranking improved notably on relevance assumptions that were based on pure semantic similarity or SERPs rankings. They also ascertain that their method improves veracity predictions for time-sensitive claims (i.e. circumstances where a news situation may be changing rapidly, and where it's essential that up-to-date information be prioritized without merely brute-forcing the prioritization of the most recent results on a topic).

The researchers note that this approach will be of high value in improving ranking models for volatile topics such as politics and entertainment, where information changes rapidly, and high-ranking developments require a framework for automatic demotion from the top spots in ranking that they may have achieved on release.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More