Towards Anticipatory Analytics: Forecasting Instability Across Countries from Dynamic Knowledge Graphs

Suraj Maharjan, Prasha Shrestha
University of Houston, Houston, TX 77004
smaharjan2,pshrestha3@uh.edu

Katherine Porterfield, Dustin Arendt and Svitlana Volkova
Pacific Northwest National Laboratory, Richland, WA 99354
firstname.lastname@pnnl.gov

1 Introduction

Protests, civil unrest, and instability around the world not only cause physical damages but also human casualties. It would be possible to mitigate these losses if there were a way to forecast such events happening world-wide. With the plethora of information coming from a variety of news sources, a system can be designed to alert about such events in advance. The ability to forecast protests, fights, and assaults can help governments to make plans to safeguard citizens from potential casualties, and monitor essential supplies.

2 Problem Statement

Prior work on protest prediction have used event data collected from Global Database of Events Language and Tone (GDELT),[1] Recorded Future,[2] and Integrated Crisis Early Warning System (ICEWS).[3] We propose to take advantage of the existing knowledge graph databases e.g., GDELT to develop deep learning models with reasoning capabilities that learn from dynamic knowledge graph representations of events, entities, and sentiments to forecast instability and future event distributions across countries. Our models are trained on sequences of daily knowledge graphs that encode the relationships across entities (nodes), event types (edges), and times. Most of the earlier work in this area focuses on applying deep learning to static graphs. Unlike the existing approaches, we propose to apply deep learning models to dynamic graphs of events and actors that vary with time to forecast the future.

[1] https://www.gdeltproject.org
[2] https://www.recordedfuture.com
[3] https://www.lockheedmartin.com/us/products/W-ICEWS.html

3 Dynamic Knowledge Graph Dataset

Our experiments rely on global events extracted from a public dataset GDELT 1.0. GDELT monitors news and social media, and extracts actors, events, locations, and other information, which is summarized and stored into an event database. We use events collected between 01/01/00 to 12/31/13 as our training data and events between 01/01/14 to 01/21/17 as our test data. We only consider a subset of fields from GDELT: SQLDATE, Actor1Name, Actor1CountryCode, Actor2Code, Actor2Name, EventRootCode, GoldsteinScale, NumMentions, AvgTone, and ActionGeo_CountryCode. The EventRootCode classifies the event that took place in a country defined by ActionGeo_CountryCode into one of the 20 possible categories. We limit our experiments to 15 countries: *Afghanistan, Australia, China, Egypt, India, Iran, Iraq, Israel, Nigeria, Pakistan, Russia, Syria, Turkey, United Kingdom, and United States* and 7 most frequent event types: *protest, threaten, assault, fight, disapprove, demand, and reject*. In addition, we take advantage of the GoldsteinScale that quantifies the theoretical potential impact of the event on countries' stability, and other fields e.g., AvgTone and NumMentions that measure the average tone of all reports containing one or more mentions of the event.

4 Approach

We experiment with the state-of-the-art machine learning (ML) and deep learning (DL) models that either take into account (a) the sequence of historical instability scores or (b) entire knowledge graph representations over time. For the task of forecasting the change in the Goldstein (aka instability score), we first find the average instability score per day, and then take the difference between instability scores for two adjacent days.

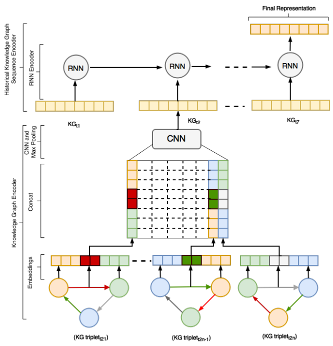

Figure 1: CNN-RNN model used to forecast instability score from dynamic knowledge graph sequences.

Our baseline models take the sequence of daily Goldstein scores of previous n days as input to ML algorithms and Recurrent Neural Networks (RNNs). Since these models take advantage of the historical sequences of instability scores, the results represent an upper bound for this task.

Our next models rely exclusively on dynamic knowledge graph representations. First, we take daily distributions of actors, events, and average tones from previous n days and feed them into the RNN model. Second, we build a hierarchical Convolution Neural Network (CNN-RNN) model as shown in Figure 1 that relies on daily knowledge graphs of previous n days. From daily knowledge graphs we learn a summary vector per day, and then pass the sequences of these summary vectors through an RNN to get the final vector. Our CNN-RNN model correlates with a hierarchical document classification model (Yang et al., 2016; Chen et al., 2016) that first learns sentence representations from sequences of words, and then learns a final document representation from the sequences of sentence representations.

5 Results

We report lower root mean square error (RMSE) obtained using a vanilla RNN model as compared to other ML algorithms: ElasticNet, Linear Regression, and AdaBoost. The RNN model yields the lowest RMSE for 10 out of 15 countries. Pearson correlation coefficient also showed the same trend as we obtained the highest Pearson correlation with vanilla RNN. Our CNN-RNN model did not perform as well as the vanilla RNN model. We suspect that this happened due to sparsity in the knowledge graph input. Therefore, we plan to take advantage of recently emerged Graph Convolutional Networks (GCNs) that address this problem (Kipf and Welling, 2017).

Moreover, we observe the lowest RMSE scores for the countries with lots of data e.g., US, Israel, and Syria, regardless of the model used. Likewise, countries with less training data such as Afghanistan and Iran yield higher RMSE scores. Although these methods only use the historical Goldstein scores and it seems unlikely for them to be affected by the number of events in the dataset, the Goldstein scores themselves are directly computed from the events and the scores are probably more reliable, the more data there is.

6 Conclusions and Future Work

This work presents preliminary results on developing deep learning models to forecast the future from dynamic knowledge graph representations. As expected, we report the best performance obtained using a vanilla RNN model trained on sequences of historical instability scores. To boost model performance trained exclusively on sparse dynamic knowledge graph representations we plan to rely on GCNs instead of CNNs. GCNs can directly take into account knowledge graphs as adjacency matrices, which is rather difficult with the traditional CNNs. In the future, we plan to test our models on a variety of forecasting tasks including predicting future events, sentiments, and event stock market prices. We also plan to test generalizability of our deep learning models trained on dynamic graph representations across real-world datasets including Twitter, YouTube, and GitHub.

References

Huimin Chen, Maosong Sun, Cunchao Tu, Yankai Lin, and Zhiyuan Liu. 2016. Neural sentiment classification with user and product attention. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1650–1659, Austin, Texas. Association for Computational Linguistics.

Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.

Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.