

SAGE: A Hybrid Geopolitical Event Forecasting System

Fred Morstatter¹, Aram Galstyan¹, Gleb Satyukov¹, Daniel Benjamin¹, Andres Abeliuk¹, Mehrnoosh Mirtaheri¹, KSM Tozammel Hossain¹, Pedro Szekely¹, Emilio Ferrara¹, Akira Matsui¹, Mark Steyvers², Stephen Bennet², David Budescu³, Mark Himmelstein³, Michael Ward⁴, Andreas Beger⁴, Michele Catasta⁵, Rok Sosic⁵, Jure Leskovec⁵, Pavel Atanasov⁶, Regina Joseph⁶, Rajiv Sethi⁷ and Ali Abbas⁸

¹USC Information Sciences Institute

²University of California, Irvine

³Fordham University

⁴Predictive Heuristics, Inc.

⁵Stanford University

⁶Pytho, LLC.

⁷Columbia University

⁸University of Southern California

{fredmors, galstyan, gleb, benjamin, aabeliuk, mehrnoosh, tozammel, pszekely, ferrarae, akira}@isi.edu, {mark.steyvers, stbennet}@uci.edu, {budescu, mhimmelstein}@fordham.edu, {michael.don.ward, adbeger}@gmail.com, {pirroh, rok, jure}@cs.stanford.edu, pdatanasov@gmail.com, rj@sibylink.com, rs328@columbia.edu, aliabbas@price.usc.edu

Abstract

Forecasting of geopolitical events is a notoriously difficult task, with experts failing to significantly outperform a random baseline across many types of forecasting events. One successful way to increase the performance of forecasting tasks is to turn to crowdsourcing: leveraging many forecasts from non-expert users. Simultaneously, advances in machine learning have led to models that can produce reasonable, although not perfect, forecasts for many tasks. Recent efforts have shown that forecasts can be further improved by “hybridizing” human forecasters: pairing them with the machine models in an effort to combine the unique advantages of both. In this demonstration, we present Synergistic Anticipation of Geopolitical Events (SAGE), a platform for human/computer interaction that facilitates human reasoning with machine models.

1 Introduction

Forecasting is extremely difficult, even for experts [Camerer and Johnson, 1997]. When experts’ accuracy is tracked over time, it has been shown to have predictive accuracy comparable to a random guess [Tetlock, 2017]. One approach to improve forecasts is to crowdsource forecasts. The benefit of crowdsourcing is that it allows for a large number of forecasts that can be then aggregated into a single probability estimate. This “wisdom of the crowds” effect has been a successful approach in diverse domains [Mellers *et al.*, 2014].

A more ambitious approach is to create hybridized forecasting systems. A hybridized forecasting platform contains not only human input, but also that from machine models. These machine models can produce forecasts about the outcome of an event, or be used as tools to help the human comb through large amounts of data to find evidence of historical events. Generally, humans are better at reasoning about qualitative data, while machines can process and aggregate quantitative data at a scale beyond human capacity. The hypothesis is that combining the strengths of these two forecasting sources will result in a system that is able to produce more accurate forecasts than any one input.

In this demo we present Synergistic Anticipation of Geopolitical Events (SAGE). SAGE is a hybridized forecasting platform for geopolitical events. In this demo we will focus on two components: 1) the human/computer interaction components of SAGE and the functionality through which humans interact with them in generating forecasts; and 2) the machine models we have developed, and the different ways in which we develop models for questions that have a scarcity or complete lack of data.

2 Geopolitical Event Forecasting

In SAGE, all forecasting problems are posed as multiple choice questions, consisting of the question text and a set of possible answer options. The answer options are the possible outcomes, separated into mutually-exclusive choices. Forecasts are provided by assigning probabilities to each answer option. Forecasts are scored using the Brier score [Brier, 1950], which rewards forecasts assigning the highest probability to the correct outcome.



Figure 1: One hybridization component in SAGE. A time series model carries the value of a commodity price forward.

2.1 Examples of Geopolitical Forecasting Questions in SAGE

Below are actual questions that have been posed to forecasters in SAGE. They are organized by quantitative (those having a numeric outcome from a base rate dataset), and qualitative (those where a base rate is unavailable).

Quantitative: What will be the daily closing price of Japan’s Nikkei 225 index on 11 January 2019?

Answer Options: More than 21,750; Between 20,550 and 21,750, inclusive; Less than 20,550.

Qualitative: Will Pakistan execute or be targeted in an acknowledged national military attack before 1 February 2019?

Answer Options: Yes; No.

3 Hybridization in SAGE

SAGE performs hybridization in two distinct ways. The first is through the use of time series forecasting. The second is a data visualization approach that maps the question to ICEWS [Ward *et al.*, 2012] event data.

3.1 Time Series Forecasting

In quantitative questions, such as stock market values and commodity prices, we have time series data available from various data sources such as Wikipedia and Trading Economics.¹ This time series data can be used for two purposes: to give the human forecaster a sense of a base rate (how often the historical value has fallen within each answer option), and to give the machine model historical data to create a time series forecast. An example of the time series visualization, complete with machine model prediction, can be seen in Figure 1. In this example, the forecaster can view the historical data, and the machine prediction. The dropdowns at the top allow the forecaster to tune several parameters, including selecting between time series models,² and different periods of training data.

¹<https://tradingeconomics.com>

²For example, ARIMA, Random Walk, and RNN-based models.

3.2 CAMEO Base Rate Visualization

For qualitative questions, such as leadership disruptions and elections, historical data is rarely available to generate a time series forecast. In these cases, SAGE obtains a base rate by mapping these questions to their CAMEO representations. CAMEO [Gerner *et al.*, 2002] is a structured representation of geopolitical events. By mapping the question to this language, SAGE generates a base rate from historical occurrences of similar CAMEO events in ICEWS data. We visualize this information using a heatmap showing the frequency of events over time. In the heatmap, each square represents the frequency of each related CAMEO code (vertical axis) at a given time point (horizontal axis). This informs the human forecaster of the frequency of such events.

4 Demonstration

SAGE is publicly available online at <https://sage-platform.isi.edu/>. Through public recruitment, SAGE has been able to recruit over 500 human forecasters who have predicted over 450 questions pertaining to geopolitics, sports, medicine and climate, among others.

Hybrid forecasting is a new task in human/computer interaction that requires humans to interact with machine models to produce more accurate visions of the future. In the course of this demo, we will provide the participants an overview of hybrid forecasting. We will then provide them with an introduction to SAGE, showing them how they can use hybridization components to refine their forecasts about real geopolitical events.

Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2017-17071900005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [Brier, 1950] Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- [Camerer and Johnson, 1997] Colin F Camerer and Eric J Johnson. The process-performance paradox in expert judgment: How can experts know so much and predict so badly. *Research on judgment and decision making: Currents, connections, and controversies*, 342, 1997.
- [Gerner *et al.*, 2002] Deborah J Gerner, Philip A Schrodt, Omur Yilmaz, and Rajaa Abu-Jabr. The creation of cameo (conflict and mediation event observations): An event data framework for a post cold war world. In *annual meeting of the American Political Science Association*, volume 29, 2002.
- [Mellers *et al.*, 2014] Barbara Mellers, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gurcay, Katrina Fincher, Sydney E. Scott, Don Moore, Pavel Atanasov, Samuel A. Swift, Terry Murray, Eric Stone, and Philip E. Tetlock. Psychological strategies for winning a geopolitical forecasting tournament. *Psychological Science*, 25(5):1106–1115, 2014. PMID: 24659192.
- [Tetlock, 2017] Philip E Tetlock. *Expert political judgment: How good is it? How can we know?* Princeton University Press, 2017.
- [Ward *et al.*, 2012] Michael D Ward, Nils W Metternich, Christopher Carrington, Cassy Dorff, Max Gallop, Florian M Hollenbach, Anna Schultz, and Simon Weschle. Geographical models of crises: Evidence from icews. *Advances in Design for Cross-Cultural Activities*, 429, 2012.