FORECASTING FATALITIES FROM STATE BASED CONFLICTS USING MARKOV MODELS

A PREPRINT

David Randahl

Department of Peace and Conflict Research, Uppsala University

david.randahl@pcr.uu.se

Johan Vegelius

Department of Medical Science, Uppsala University

June 27, 2024

ABSTRACT

In this contribution to the VIEWS 2023 prediction challenge, we propose using a set of different Markov-type latent state models to make prediction of fatalities from state-based conflicts on the country-month level. Partly building on the Markov modeling strategy from the VIEWS 2020 prediction contest, we propose three types of Markov-style models. First, we use an observed Markov model (OMM) which utilizes domain knowledge about conflict states to define observed states through which countries can move over time. The OMM is flexible as it does not require any parametric assumptions, and can be viewed as a set of classification and regression problems. Second, we propose a hidden (pseudo-) Markov model (HPMM) which utilizes unknowable, latent or hidden, states which the countries can move through over time. The HPMM is not strictly Markovian as we relax the assumption that the transition matrices are conditional on discrete states and instead model transitions conditional on weighted states from the posterior state probabilities. Finally, we propose aa Gaussian process continuous Markov model (GPCMM) which utilizes a continuous observed Markov 'state' through which countries move over time.

Keywords Markov models • Forecasting • Gaussian process • Fatalities

Federal Foreign Office

PREVIEW
Prediction
Visualisation
Early Warning

![VIEWS_logo_new.png]

1 Introduction

This contribution to the VIEWS 2023/24 prediction challenge[1] proposes using a set of different Markov-style models to model and make forecasts for the number of fatalities from state-based conflicts for all countries of the world. The approach builds on our contribution from the previous VIEWS competition (Randahl and Vegelius 2022) but extends this to making distributional forecasts, and the type of Markov-style models used. We continue to believe that the real-world dynamics of armed conflict closely resemble a Markov model with a large number of states. The fundamental idea of a Markov process is that for each point in time, t, an individual. i, has a specific state, s, which may generate some observed outcome, y. The individual i may move through different states, for instance 'peace', 'escalation', or 'war', over time. Crucially in a Markov process, however, is that the likelihood of the future state(s) is conditional only on the current state and not the previous states of the Markov chain (Karlin 2014). ^[For a more comprehensive description of Markov models, see (Karlin 2014)

In order to make the distributional predictions of the prediction target, we generalize away from the traditional Markov model (MM), focused on transition between observed Markov states, and the hidden Markov model (HMM), focused on transitions between unobserved, hidden, Markov states as well as the observed *emissions* from these states, to three Markov-style prediction models, the observed Markov model (OMM), the hidden pseudo Markov model (HPMM), and the Gaussian process continuous Markov model (GPCMM). We call these 'Markov-style' models since they are not setup in the same way as the traditional MM and HMM models, and are not necessarily strictly Markovian. The three Markov-style prediction models are similar to traditional Markov models in that they assume that each unit (country) at each point in time (month) has an underlying *latent state* and that this latent state affects the data generating process of the outcome we are predicting. In the OMM, we define observable latent states and condition transitions between states as well as the prediction of the outcome on the actual observed state of the process. In the HPMM we use three pre-defined parametric distributions as the Markov states and their observed posterior probabilities as a mixture state for each unit at each point in time. Transitions are not modeled as discrete transition between underlying hidden states but rather as changes in the posterior state probabilities between the two time points. In the GPCMM the states are observed but continuous and transitions and predictions are modeled as a joint gaussian process over the states and features.

In this submission for the prediction competition, all of our models use a combination of thematic principal components and transforms of the lagged dependent variable as predictor features. More specifically we use one principal component from the water-related variables in the input data, two from the vdem data, two on the history of violence, two on indicators relating to socioeconomic development, one on military expenditure, one on population, and two on neighborhood characteristics, as well as the log1p sum of the total number of fatalities from state-based violence in the last two years, and a decaying log1p sum of the total number of fatalities from state-based violence since the start of the time-series with a 12 month half-time.

2 The observed Markov model

The observed Markov model (OMM) we propose for this prediction competition is setup similar to the OMM model proposed in our contribution to the VIEWS 2020 prediction competition (Randahl and Vegelius 2022). We use the same four states as in our previous contribution, namely, 'war' for countries which observe more than 0 fatalities in the current month and previous month, 'escalation' for countries which observe more than 0 fatalities in the current month but no fatalities in the previous month, 'peace' for countries which observe 0 fatalities in the current month an previous month, and 'deescalation' for countries which observe 0 fatalities in the current month but more than 0 fatalities in the previous month. As these states are observed, transitions between the states can be modeled using any type of classifier which produce probabilities that can be used to create transition matrices conditional on the states. In this submission we use probability forests (Athey, Tibshirani, and Wager 2019; Tibshirani et al. 2023) to model these transitions, with the full set of principal components and transforms of the lagged prediction target as input features.

Given that the states in the OMM are observed, the prediction of the target itself reduces to a set of regression problems conditional on the observed state. Similar to the transitions, this means that any regressor can be used to model these. As the goal of this prediction competition is to produce a distribution, or random sample, of values for each individual in the test sets we use quantile forests (Athey, Tibshirani, and Wager 2019; Tibshirani et al. 2023) in order to approximate a distribution of plausible values. We only estimate these quantiles for the 'war' and 'escalation' states, as the number of fatalities is known to be zero for the 'peace' and 'deescalation' states.

[1] This paper documents a contribution to the VIEWS Prediction Challenge 2023/2024. Financial support for the Prediction Challenge was provided by the German Ministry for Foreign Affairs. For more information on the Prediction Challenge please see Hegre et. al (forthcoming) and https://viewsforecasting.org/research/prediction-challenge-2023

To produce distributional predictions of the prediction target for each country-month 15 months into the future from the end of the training data we simulate the Markov chain for this time period for each individual country based on the last available predictor features. Conditional on the simulated Markov state, we then make a random draw of the predicted number of fatalities. For the 'peace' and 'deescalation' states, this always corresponds to zero fatalities. For the 'war' and 'escalation' states, we make a random uniform draw of the quantile for that individual country month and then use the quantile forest to predict the specific quantile. We repeat this process 1,000 times to produce 1,000 random draws for each country-month.

3 The hidden pseudo Markov model

The hidden pseudo Markov model (HPMM) we propose for this prediction competition is similar to standard hidden Markov models (HMM) (see for instance Karlin 2014; Jackson 2011) and the HMM we proposed for the previous prediction competition (Randahl and Vegelius 2022). However, while a standard HMM could work for this prediction problem, these HMMs often suffer from problems of identification and specification (see for instance Rabiner 1989; Mattila 2018). Initial tests using standard HMMs failed to converge unless the number of predictor features was restrictively low and the number of states were set to 2. In this submission we propose a more flexible version of the HMM which we call the hidden pseudo Markov model in which we relax the assumption that the future state(s) of the process are conditional on the discrete current state of the process. Instead, we condition the future states of the process on the posterior probability of the different states in the current time point. This reduces the computational complexity of the HMM algorithm as there is no longer a need to compute the Markov chain path through the forward-backward algorithm in the training state. One way of viewing this algorithm is as a mixture between the OMM and HMM as it uses the posterior state probabilities as observed non-discrete states. Another upside of this approach is that this again also reduces the problem of transition into a regression problem [2] which increases the range of potential models which can be used to model transitions. We note that the use of the HPMM should be considered to be very experimental.

The HPMM still requires, just as the HMM, the researcher to specify a number of parametric, hidden, states for the process. Inspired by the extreme value and zero inflated regression model (Randahl and Vegelius forthcoming) we for this submission use a HPMM with three hidden states, one zero state, one negative binomial state, and one extreme-value, pareto, state. The parameters for the states and transitions are found using the EM-algorithm (Dempster, Laird, and Rubin 1977). To produce distributional predictions of the prediction target for each country-month 15 months into the future from the end of the training data we simulate the Markov chain for this time period for each individual country based on the last available predictor features. In each time-step we draw simulated fatalities and then update the posterior probabilities for the latent states which are then used to make predictions of both fatalities and latent states in the next time step. We run this procedure 1,000 times to produce 1,000 HPMM chains for each country in each time period.

4 The Gaussian process continuous Markov model

The Gaussian process continuous Markov model (GPCMM) we propose for this prediction competition is similar to OMM in that the state is observed. The state in GPCMM is defined by the value of a predefined 'state variable'. As in OMM and HPMM the state defines the distribution of the number of fatalities, e.g., in terms of the effect of covariates and the variance conditional on the covariates. In opposition to OMM and HPMM the state variable can be continuous providing a potentially infinite space of states. Observations of similar values of the state variable are assumed to be associated with similar distributions of the number of fatalities as a function of covariates, whereas observations with large distances in the state space can have very different distributions. The parameters determining the distribution of the number of fatalities are assumed to follow a Gaussian process (Schulz, Speekenbrink, and Krause 2018; Bishop 2006) with a kernel defined by the distance in the state space. In analog with the OMM, the GPCMM reduces to a continuous set of regression problems. One for each value of the state variable.

The state transition is defined by another Gaussian process where similar states are associated with similar transition probabilities among states. Hyper parameters in the kernel are estimated by maximizing the model evidence by the Laplace approximation.

The distributional predictions are produced by propagating the state variable 15 steps into the future from the last data point in the training. At each time point, a prediction of the number of fatalities is drawn from its distribution conditional on the drawn state at that time point. This procedure is repeated 1,000 times to produce 1,000 random draws for each country-month.

[2] As the prediction target here are probabilities rather than classes it is no longer a classification problem

The GPCMM is still under development and in this submission we are only using the 'state variable', defined as a decay function of the number of fatalities observed in the country, as well as its lag as prediction features when simulating the GPCMM chains.

References

Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. "Generalized Random Forests."

Bishop, Christopher M. 2006. *Pattern Recognition and Machine Learning*. springer.

Dempster, Arthur P, Nan M Laird, and Donald B Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society: Series B (Methodological)* 39 (1): 1–22.

Hegre, Håvard, and et.al. forthcoming. "The 2023/24 VIEWS Prediction Competition" XXX (forthcoming).

Jackson, Christopher. 2011. "Multi-State Models for Panel Data: The Msm Package for r." *Journal of Statistical Software* 38: 1–28.

Karlin, Samuel. 2014. *A First Course in Stochastic Processes*. Academic press.

Mattila, Robert. 2018. "Hidden Markov Models: Identification, Control and Inverse Filtering." PhD thesis, KTH Royal Institute of Technology.

Rabiner, Lawrence R. 1989. "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition." *Proceedings of the IEEE* 77 (2): 257–86.

Randahl, David, and Johan Vegelius. forthcoming. "Inference with Extremes: Accounting for Extreme Values in Count Regression Models." *International Studies Quarterly* x (x): x.

———. 2022. "Predicting Escalating and de-Escalating Violence in Africa Using Markov Models." *International Interactions* 48 (4): 597–613.

Schulz, Eric, Maarten Speekenbrink, and Andreas Krause. 2018. "A Tutorial on Gaussian Process Regression: Modelling, Exploring, and Exploiting Functions." *Journal of Mathematical Psychology* 85: 1–16.

Tibshirani, Julie, Susan Athey, Rina Friedberg, Vitor Hadad, David Hirshberg, Luke Miner, Erik Sverdrup, Stefan Wager, Marvin Wright, and Maintainer Julie Tibshirani. 2023. "Package 'Grf'."