[W]hat lies beneath: Using Latent Networks to Improve Spatial Predictions of Political Violence

Abstract

Spatial interdependencies commonly drive the spread of violence in civil conflict. To address such interdependence, scholars often use spatial lags to model the diffusion of violence, but this requires an explicit operationalization of the connectivity matrices that represent the spread of conflict. Unfortunately, in many cases, there are multiple competing processes that facilitate the spread of violence making it difficult to identify the true data-generating process. We show how a network driven methodology can allow us to account for the spread of violence, even in the cases where we cannot directly measure the factors that drive diffusion. To do so, we estimate a latent connectivity matrix that captures a variety of possible diffusion patterns. We use this procedure to study intrastate conflict in eight conflict-prone countries and show how our framework enables substantially better predictive performance than canonical spatial lag measures. We also investigate the circumstances under which canonical spatial lags suffice, and those under which a latent network approach is beneficial.

Word Count: 8245

Introduction

Patterns of violence tend to exhibit strong degrees of spatial correlation, clustering, and diffusion (Dorff, Minhas and Ward, 2017). Violence can diffuse spatially across numerous pathways: intuitively, it might propagate through geographic proximity wherein violence in one province, neighborhood, or country is liable to spread to a neighboring area (Buhaug and Gleditsch, 2008). Similarly, socio-economic conditions across regions, such as horizontal inequalities and shared ethnicity might also drive trends behind violent events (Cederman, Weidmann and Gleditsch, 2011). Scholars have shown that climate factors such as precipitation and temperature have a demonstrated relationship to the risk of violence (O'Loughlin et al., 2012) and geographic barriers linking vital borders such as mountain ranges, or infrastructure like roads, can influence the spread of conflict (Braithwaite, 2010). Modeling these diffusion pathways has become increasingly common through a spatial lag framework ($y_t \sim \rho W y_{t-1}$) in which scholars specify a set of weighting matrices, W, to account for suspected drivers of conflict diffusion. Yet, given the diverse plausible drivers of civil conflict diffusion highlighted above, the task of determining and operationalizing an appropriate set of weighting matrices is difficult for scholars to accurately identify *ex ante*.

The question of how to identify diffusion processes is especially relevant due to the increased interest in prediction and use of machine learning techniques in the study of violence. It is easy to understand why prediction has taken on particular significance for conflict scholars (Guo, Gleditsch and Wilson, 2018). For researchers and policy practitioners alike, preventing loss of human life means successfully identifying locations where political violence is likely to be high in the near future. This goal of prediction, predicated on the prevention of violence, has led researchers to turn towards machine learning techniques that can typically handle more complex data generating processes and out-perform traditional approaches in forecast accuracy.

Of course, machine-learning techniques have important limitations. First, such models are only as good as the features the user is able to construct for the model's input. Dealing with this problem in the context of conflict prediction is especially relevant as we need to think carefully about how to account for the variegated pathways through which conflict may spread across a country. Second, high quality, fine-grained data is often not available at the sub-national unit of analysis in a time series context. If our goal is to predict violence within a relatively small time-frame at any sub-national level, these data problems limit the benefit of adopting these approaches.

To overcome these obstacles, we propose an unsupervised network approach which estimates a weighting matrix with probabilistic measurements of diffusion pathways within a country. This approach enables us to take advantage of previous methodological advances in the spatial modeling literature and supplement them with recent advances in network science to generate a set of features that embed sub-national geographic units onto a latent conflict space. Sub-national units more proximate in the latent space are more likely to spread conflict to one another. We apply this technique to eight countries that have experienced particularly high counts of intrastate conflict during the 21st century according to the Uppsala Conflict Data Project (UCDP). Each of the countries ranks among the most violent civil conflicts in Africa during the 21st century. For each of these eight countries, we show that simply integrating our network based feature into machine learning pipelines for conflict prediction enables us to better predict the occurrence of conflicts in an out-of-sample context.

Spatial Dimensions of Conflict

For decades, scholars have recognized the importance of spatial factors in the study of interstate violence (Richardson, 1960; Bremer, 1992), with an early focus on shared borders and a later focus on both distance and the regional context of regime type (Ward and Gleditsch, 2002). While Ward and Kirby (1987) find significant spatial autocorrelation in conflict, they argue that treating borders as a cause of conflict misses the strategic and contingent nature of borders. Gartzke (2006) shows that the effect of contiguity on violence depends on the relative development of states, where distance has less of a pacifying effect on richer states.

In more recent years, scholars have investigated not just the spatial diffusion of conflict across states, but also how conflict spreads within a country. Weidmann and Ward (2010) find strong spatial correlation in patterns of violence in Bosnia during the 1990s, and Townsley, Johnson and Ratcliffe (2008) find similar spatial and temporal correlations in the use of improved explosive devices in Iraq. O'Loughlin and Witmer (2011) show how conflicts spread from hotspots in the caucuses and highlight how conflicts are driven, not just by distance, but by shared religious affiliations between groups. Schutte and Weidmann (2011) distinguish between conflict flows due to relocation (violence that is no longer in location 1 but has moved into location 2) and escalation (when violence appears in location 1 and then, in the next time period, also in location 2) and find that escalation is more common than relocation in several civil conflicts.

Literature on the spatial diffusion of intrastate conflict has offered multiple mechanisms that might explain the spread of violence across sub-state regions (Gleditsch, 2007). Salehyan and Gleditsch (2007) find that refugees are an important vector in the spread of conflict due to shared networks and shifts in population composition in the host state. Similarly, a number of studies have shown that countries that share ethnic kin are more likely to be linked in the spread of conflict (Buhaug and Gleditsch, 2008; Cederman, Gleditsch and Girardin, 2009; Wucherpfennig et al., 2011; Cederman, Hug and Krebs, 2010; Metternich, Minhas and Ward, 2017).

Focusing on governments' role in enabling or constraining violence, Kathman (2010), Beardsley (2011), and Danneman and Ritter (2014) find that the spread of civil conflict is conditional on states' policy choices—intervention, peacekeeping, and repression respectively. Research on the micro-foundations of conflict and peace suggest that both violent and non-violent actions are often driven by localized conditions, such as historical legacies of repression (Osorio, Schubiger and Weintraub, 2018), competition between anti-government groups (Metternich et al., 2013) and the distribution of civilian collective action efforts (Dorff and Braithwaite, 2018).

Undoubtedly, previous scholarship has consistently found that spatially distributed conditions are substantively important to our understanding of conflict processes, but the particular ways in which these spatial factors manifest can change substantially from conflict to conflict. (Franzese and Hays, 2008) provide a useful typology of the ways that spatial effects can manifest. Spatially-driven patterns can arise when an exogeneous common-shock has similar effects in a number of areas, or they can occur because the processes we study are actually interdependent. For example, Franzese and Hays (2008) point to five different mechanisms that might drive interdependence among actors: coercion, competition, strategic learning, emulation, and migration. Of course, multiple factors may be at play in a given conflict. Thus, the state of the literature suggests that there are many barriers to identifying appropriate measures for the spatial dimensions of conflict.

Conflict Prediction

While research identifying the drivers of diffusion has led to important advances in the field, the focus on prediction is a welcomed addition to the literature. Increasingly, scholars have begun to pay attention to prediction both as a way of evaluating the effects of important variables and as an end in and of itself. Both Ward, Greenhill and Bakke (2010) and Schrodt (2014) have criticized the emphasis on statistical significance in the field and suggest that scholars instead aim to generate models that can predict *out-of-sample*. As a number of scholars have pointed out, many extant empirical models in conflict are quite lacking when it comes to predicting conflict out-of-sample (Ward, Siverson and Cao, 2006; Beck, King and Zeng, 2000; Gleditsch and Ward, 2010). Recently, however, more conflict researchers have developed models explicitly focused on the prediction of conflict related processes, such as humanitarian crises (Gurr and Harff, 1996), civil war onset (Hegre and Sambanis, 2006; Ward et al., 2013), civil war duration (Bennett and Stam, 2009), incidences of armed conflict over a 40 year period (Hegre et al., 2013), mediation in civil conflict Clayton and Gleditsch (2014), and dissident groups' turns towards violence (Shellman, Levey and Young, 2013), among others. Some scholars have even attempted to use predictive accuracy as a substitute for "p-values" by showing which factors improve on the ability to predict out-of-sample, such as (Brandt, Colaresi and Freeman, 2008)'s assessment of the effect of public opinion on conflict in Israel. This research demonstrates the value of prediction for predicting and–hopefully–ameliorating civil conflict.

The rise in scholarship on the benefits of prediction has led to a growing interest in the ability of machine learning techniques to aid in the prediction of violent processes. Notably, Colaresi and Mahmood (2017) argue that machine learning could help to develop better predictors for the study of violent conflict. In early attempts, results were mixed: in an initial effort to predict interstate conflict using a neural network, Schrodt (1990) was unable to out-predict more conventional linear models. The prediction of interstate conflict using a similar technique by Beck, King and Zeng (2000) was the subject of controversy over whether it could really predict accurately out-of-sample (de Marchi, Gelpi and Grynaviski, 2004). However, in the decades since, advances in computing power and algorithmic sophistication have increased our ability to use machine learning techniques to predict conflict. For example, Hill and Jones (2014) have used random forests to distinguish which factors best predict government repression and Jones and Lupu (2018) implement a random-forest like algorithm to explore the relationship between regime type and violence. Muchlinski et al. (2016) employ a random forest model to predict civil war onsets, and Perry (2013) also utilize a random forest (along with a Naive Bayes classifier) to predict the incidence of battles at the district level. One of the biggest issues with these techniques, however, is that they are quite data intensive, and require large numbers of features to yield their promised predictive accuracy. This becomes difficult when the goal is to predict phenomena at increasingly fine resolutions— both temporally and spatially.

Latent Diffusion

Clearly, conflict processes are exceedingly complex. To accomplish accurate prediction, we turn to a latent variable approach. Latent variable frameworks generate variables that are not directly observable, but can be constructed from observable information. This allows for a compression, or reduction, of complex data making it easier to understand multidimensional data generating processes. We suggest that this is a particularly useful strategy for conflict prediction, where conditions that predict violence are often shifting and difficult to identify over time. Further, a singular, dominant 'cause' or predictor of diffusion might not actually exist, producing even greater barriers for prediction. To return to Franzese and Hays (2008)'s typology, a civil conflict might experience common shocks (due to economic factors, for example, which effect some but not all regions of a country) or there could be *coercion* by international actors to channel or suppress violence in certain regions. In some cases we have even seen *learning*, where different non-state actors share tactical and strategic innovations, and civil conflict often leads to *migration* where refugees from one region repeat certain conflict behaviors in a new region. All of these phenomenon will not necessarily map onto a proximity-based weight matrix and some of them will cut in different directions. A latent approach allows us to capture the latent connections between provinces in a given country, based on past diffusion patterns, and leverage this measurement for prediction.

These latent connections might be driven by an an assortment of important processes even at the sub-state level, such as communal relations between warring parties, shared ethnic ties, or state capacity and reach. For example, in the case of Nigeria, changes in the location of violent events in the northeast, where Boko Haram engages in terroristic behavior against civilians, is arguably driven by different factors than conflict in the Niger delta where militia groups vie to capture oil production centers. In South Sudan, though oil revenues and resource management has played a key role in sustaining conflict, the behavior and presence of international actors is also significant to the conflict's development (Johnson, 2014).

Notably, it is beyond the goals and scope of this article to explicitly characterize our latent variable for each conflict to offer a new, named concept of diffusion that specifies certain processes over others. But focusing too much on perfectly specifying and measuring one process, or a collection of processes, over others as key determinants of diffusion, will likely hinder our ability to forecast violence.

Data and Features

Case Selection

In order to study patterns of spatial diffusion and violence we will rely on the UCDP's Geolocated Event Data (GED) developed by Sundberg and Melander (2013). This dataset has 142,902 records of violent events in 118 countries between 1989 and 2017 located across space and time. To begin, we rank all countries in the data by order of conflict intensity, by which we define as countries with the highest number of violent events per year of conflict, we then examine the longest continuous period of violence.[1] We first examine the top 20 most violent countries, but choose to drop one case, and thus added the 21st most violent country.[2] UCDP GED data is used to create our measure of conflict where $y_{il,t} = 1$ indicates that a conflict occurred in country i's region l at time t ($y_{il,t} = 0$ if no conflict occurred).

We obtain data on administrative boundaries in these twenty countries from the United Nations Organization for Coordination of Humanitarian Affairs, which has made shape files for these countries available on the Humanitarian Data Exchange. In each of these cases we attempt to choose the administrative divisions that would provide between ten and fifty units. We do this to avoid having too few units for our network analysis and to prevent using an overly sparse matrix. In the end, this means we utilize the level two divisions for South Africa, Sierra Leone, Rwanda, and Senegal, and the level one divisions for the other 16 countries. In Table 1 we display descriptive information for each case. *Years* is recorded as the number of years with a GED event in the country; *Conflict Years* are the periods of time where each year had at least one conflictual event; *# of Events* is the total number of GED events for that country; and *Events per Conflict Year* is the number of events divided by the number of years with at least one event. This collection of cases displays variability across these measures, ranging from countries that have an average of 140 GED events per year to those that only have around 6. In the remaining sections, we explain the details of our approach. To do so, we select South Sudan as an illustrative case in order to clearly demonstrate each step of our analysis. Of course, all of these steps are applied to all 20 cases in our study.

[1]The most consecutive years where there was at least one violent event per year.

[2]We dropped Uganda because the country is in the process of persistent adminstrative boundary changes, making the process of predicting violence in a given district fraught.

Country	Conflict Years	# of Events	Events per Conflict Year	Admin
Algeria	**1990 - 2016**	3786	140.22	1
Somalia	**1989 - 2016**	3685	131.61	1
Libya	2011 - 2012,**2014 - 2016**	621	103.50	1
South Africa	**1989 - 2000**,2004 - 2004,2016 - 2016	2232	79.71	2
Nigeria	1990 - 1994,**1996 - 2016**	1981	73.37	1
Sudan	**1989 - 2016**	1801	64.32	1
Sierra Leone	**1991 - 2001**	644	58.55	2
Angola	**1989 - 2005**,2007 - 2011,2013 - 2016	1586	56.64	1
Ethiopia	**1989 - 2016**	1568	56.00	1
South Sudan	**2011 - 2016**	262	43.67	1
DR Congo (Zaire)	**1993 - 2004**,2006 - 2016	1008	42.00	1
Uganda	1989 - 1992,1994 - 2009,2016 - 2016	1110	39.64	Not Used
Burundi	1990 - 1992,**1994 - 2008**,2012 - 2012,2014 - 2016	882	32.67	1
Central African Republic	2001 - 2003,2006 - 2007,**2009 - 2016**	313	19.56	1
Kenya	1989 - 1989,**1991 - 2016**	537	19.18	1
Liberia	**1989 - 1996**,2000 - 2003	180	12.00	1
Mali	1990 - 1991,**1994 - 1994**,1997 - 1997,1999 - 1999,2004 - 2005,2007 - 2016	293	10.85	1
Rwanda	**1990 - 1994**,1996 - 1998,2001 - 2001,2004 - 2004,2012 - 2013	243	10.12	2
Senegal	1989 - 1990,1992 - 1993,**1995 - 2006**,2008 - 2013	176	7.04	2
Mozambique	**1989 - 1992**,2004 - 2005,2012 - 2014,2016 - 2016	189	6.75	1
Chad	1989 - 1995,**1997 - 2010**,2015 - 2016	185	6.61	1

Table 1: Country Case Information, years used are in **bold**.

Features

Using this data, we generate a set of standard features to account for temporal and geographic spatial dependence. In particular, we include variables measuring the occurrence of battles in surrounding provinces in the previous period, following work by Neumayer and Plümper (2010), Hays, Kachi and Franzese (2010), Ward and Gleditsch (2002) among others. This is done using both a binary weight matrix, where province i and j have a value of 1 if they are directly contiguous and 0 otherwise, and weights using the distance between each province's centroid. We also create one to five year lags for battles in the province itself, and use a cubic spline of these lagged values to capture more complex and non-monotonic patterns of temporal dependency (Carter and Signorino, 2010). Finally, we include indicator variables for each province to capture the differential likelihood of violence in different regions of a country. We will compare these features to our network approach which we explain in the following section. These features are shown below in Table 2.

One issue with adopting this approach is that it leads to the generation of a large number of features that are likely correlated with one another. Correlations between the features would not be surprising as we are generating multiple lagged versions of

Features	Dependence	Baseline	Networks
Spatial Lag using Contiguity Matrix	Spatial	✓	✓
Spatial Lag using Centroid Matrix	Spatial	✓	✓
1-5 year Lag of DV	Temporal	✓	✓
Cubic splines of DV	Temporal	✓	✓
State Level Indicator	Heterogeneity	✓	✓
Spatial Lag using Latent Diffusion Matrix	Unobserved Spatial		✓
Total Number of Features		9	10

Table 2: Feature construction of available dependence measures in both baseline and network models

the same variable. To deal with this, we conduct a principal component analysis (PCA) over our set of spatial and temporal features. PCA reduces the dimensionality of a dataset by finding a set of linearly uncorrelated principal components while still retaining key patterns in the original dataset. The new, uncorrelated variables (the principal components) are what we pass as features to the classifiers in all of the analyses presented in the following sections.

Network Based Feature Construction

The set of spatial variables we include in the model all operate under the principle that there is an explicit geographic connection which explains how conflict diffuses from one province to another. However, there might be factors beyond just geography that explain why conflict diffuses within a country. Measuring these kind of latent connections is possible through a network based approach. The first step in doing this is to construct a set of socio-matrices from of our events of interest. Specifically, the measure we construct should represent the ways in which conflict may diffuse across provinces. To do this, we employ a simple decision rule:

- given province i has a value of 1 at time t & province j has a value of 0 at time t
- if at time $t+1$ j also experiences an event, then we code the i, j^{th} value in our diffusion matrix as 1

- alternatively, if at time $t+1$, j does not experience an event, then we code the i, j^{th} value of our diffusion matrix for that value as 0

We generate socio-matrices using this decision rule for every time period and country. [3]. Next we generate our low-dimensional representation of connectivity between provinces using the multiplicative effects portion of the Additive & Multiplicative Effects (AME) model (Hoff, 2005; Minhas, Hoff and Ward, 2019). Generally, the AME model can be used to represent network dependencies through a set of random effects. We use the multiplicative effects portion of this model to develop a latent factor space that measures how likely an event is to spread from one province to another. The model is specified as follows:

$$\begin{aligned} y_{ij} &= f(\theta_{ij}), \text{ where} \\ \theta_{ij} &= \mathbf{u}_i^\top \mathbf{D} \mathbf{v}_j \end{aligned} \tag{1}$$

The multiplicative term here is: $\mathbf{u}_i^\top \mathbf{D} \mathbf{v}_j = \sum_{k \in K} d_k u_{ik} v_{jk}$. K denotes the dimensions of the latent space. The construction of the LFM here is actually quite similar to the recommender systems that companies like Amazon and Netflix have used to model customer behavior (Resnick and Varian, 1997; Bennett and Lanning, 2007). This model posits a latent vector of characteristics $\mathbf{u_i}$ and $\mathbf{v_j}$ for each sender i and receiver j. The similarity or dissimilarity of these vectors will then influence the likelihood of activity, and provide a representation of third-order interdependencies. The LFM parameters are estimated by a variant of the singular value decomposition (SVD) of the observed

[3]aggregated to a monthly level due to sparsity when attempting to use a weekly level

network.[4]

These latent factors are calculated to account for homophily and stochastic equivalence in relational data. These factors map well onto the phenomenon that we are attempting to measure through our latent diffusion matrix. Homophily captures the idea that certain unobserved factors make two provinces more or less similar, and that we are more likely to see the diffusion of violence across similar provinces. We can again turn to the example of regions with similar populations of an ethnic minority, in which violence in either region is likely to spread to the other through ethnic mobilization. The essential message here is that while we may not know exactly which characteristics across provinces influence processes of diffusion, the AME model, by using the SVD, is very good at inferring latent similarity by using previously observed instances of diffusion.

Stochastic equivalence captures the idea that two provinces might play similar roles in the diffusion of violence within a state. In particular, this means that violence is equally likely to diffuse from either province to a common third province: for example violence is likely to diffuse from either state i or state j to states k but not to state l. We could see an example of this in cases of separatist violence against the government. Violence is likely to begin in potential breakaway provinces on the periphery and move towards the center. So we might observe that different peripheral states with large separatist groups are stochastically equivalent, as violence is unlikely to diffuse between these states, but is very likely to move from these states to the country's capital. On the other hand, in cases of violence between rival groups in the core of a country, we should see the central provinces of that country exhibite stochastic equivalence, as

[4]Unlike in traditional SVD, in the latent factor model, the singular values are not restricted to be positive, this allows us to account for both positive and negative homophily.

violence is likely to diffuse between them first, and later spread to the periphery.

To compute the SVD we factorize our observed network into the product of three matrices: $\mathbf{U}, \mathbf{D},$ and $\mathbf{V}$. This provides us with a low-dimensional representation of our original network.[5] Values in $\mathbf{U}$ provide a representation of how stochastically equivalent certain provinces are as signal markers for where conflict will occur next. $\hat{\mathbf{u}}_i \approx \hat{\mathbf{u}}_j$ would indicate that provinces i and j experience conflict at the same time as a specific set of other provinces. $\mathbf{V}$ provide a similar representation but from the perspective of how similar actors are as receivers. The values in $\mathbf{D}$, a diagonal matrix, represent levels of homophily in the network.

These third order interdependencies can capture the different patterns and pathways in which violence spreads within a given country. It is important to note here that while we will see high variance between states based on differential patterns of homophily and stochastic equivalence, our technique allows us to account for these interdependencies in our latent diffusion matrix without needing to specify them beforehand or determine the appropriate covariates.

This model is estimated through a Bayesian probit framework. The algorithm proceeds as follows until convergence:

- For each $k \in K$:
  - Sample $\mathbf{U}_{[,k]} \mid \mathbf{U}_{[,-k]}, \mathbf{V}$ (Normal)
  - Sample $\mathbf{V}_{[,k]} \mid \mathbf{U}, \mathbf{V}_{[,-k]}$ (Normal)
  - Sample $\mathbf{D}_{[k,k]} \mid \mathbf{U}, \mathbf{V}$ (Normal)[6]

Once we estimate these models for every country and time period, we calculate the

[5] The dimensions of $\mathbf{U}$ and $\mathbf{V}$ are $n \times K$ and $\mathbf{D}$ is a $K \times K$ diagonal matrix.

[6] Subsequent to estimation, $\mathbf{D}$ matrix is absorbed into the calculation for $\mathbf{V}$ as we iterate through K.

likelihood that conflict diffuses from one province to another based on this model. We then proceed to use the predicted probabilities from our model as a new weighting matrix, through which we generate a new set of spatial lags that we include in our set of features for predicting conflict across our sample of twenty countries. Figure 1 summarizes our framework for developing weighting matrices to account for diffusion patterns. For each country, we include all of the features discussed in table 2, but we also include this network derived measure of spatial dependence in the Principle Component Analysis.

Figure 1: Summary of process to create network based features for South Sudan.

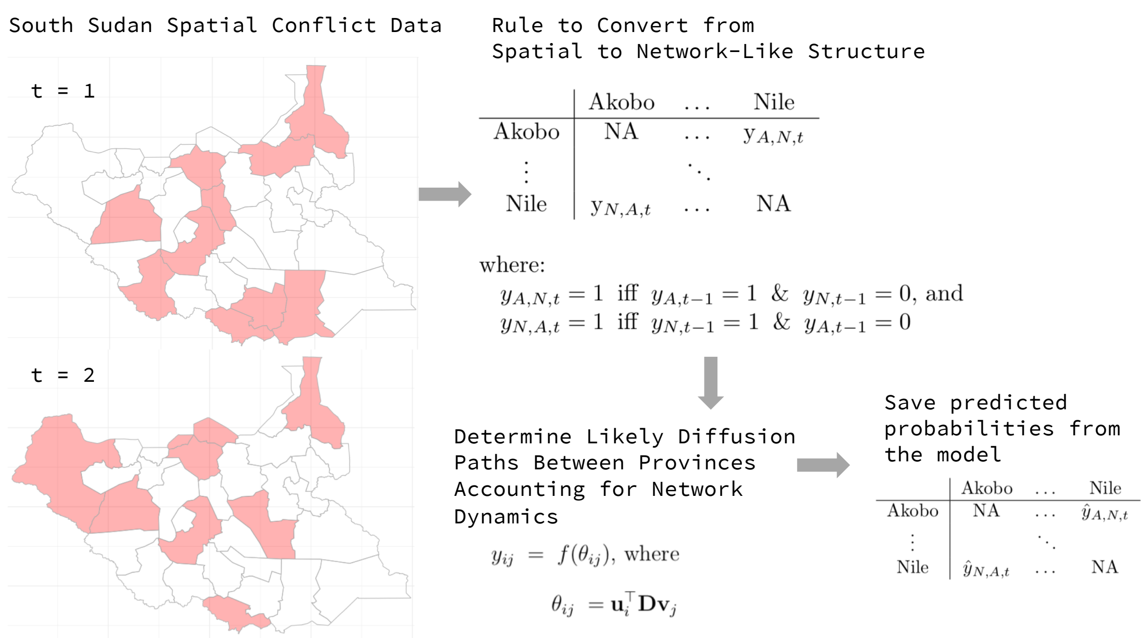

Simple Machine Learning Pipeline for Modeling Conflict

To provide a framework to test whether or not our network based features improve our ability to predict the occurrence of conflict, we construct a straightforward machine learning pipeline. To begin, we utilize a wide array of models common to the machine learning literature that use the features described above as inputs. Each model will have two versions, one that uses a PCA including the first nine features detailed in table 2 and one using use those nine features *as well as* our network measure of spatial dependence. These set of models are listed below:

- Logistic regression via penalized maximum likelihood (Friedman, Hastie and Tibshirani, 2010)
- Support vector machines (Meyer and Wien, 2001; Chang and Lin, 2011)
- Random Forest (Liaw and Wiener, 2002)
- Regularized boosted regression models (Chen and Guestrin, 2016)

These models are run separately on each of the countries in our sample. After having run these models for each country we then work towards leveraging the predictions from each model using Bayesian Model Averaging (BMA) to calibrate a forecast ensemble. Our logic for doing so is straightforward and flows from a long recognized suggestion in the prediction literature that multiple models often provide a better description of the data generating process (Hoeting et al., 1999). Relying on any one model comes with the downside of having to rely on that model's particular distribution or assumptions. BMA is an extension of Bayesian inference to the problem of model selection (Raftery et al., 2005). With BMA, the likelihood of a conflict assigned to a given subnational month unit by the overall model is determined by taking the probability of an event as predicted by a single model, multiplying by the probability that the model is true model given a sample of the data, and summing these values across all models in the hypothesis space. Say that we have data, D, which is comprised of a set of features, X, and an outcome vector, y. The model space is approximated by a set of learners, L, with l representing an individual hypothesis in that space. Equation 2 then describes how the probability of conflict is determined for a given sub-national month unit:

$$p(y_i|x_i, D, L) = \sum_{l \in L} p(y_i|x_i, l)p(l|D) \tag{2}$$

Utilizing Bayes' theorem, the posterior probability that l is the true model ($p(l|D)$) can be estimated by Equation 3, where $p(l)$ represents the prior probability of l and the product of $p(d_i|l)$ determines the likelihood.[7]

$$p(l|D) \propto p(l) \prod_{i=1}^{n} p(d_i|l) \tag{3}$$

We follow the typical approach in the BMA literature of assuming a uniform class noise model to determine the likelihood in Equation 3. The uniform class noise model assumes that the possible value of each observation is corrupted with probability ϵ, thus $p(d_i|l)$ is $1 - \epsilon$ if the learner, l, correctly predicts the value of y_i and ϵ otherwise. This enables us to reformulate Equation 3 as Equation 4, where s is the number of correct predictions determined by l and ϵ can be approximated by the average error rate of the model.

$$p(l|D) \propto p(l)(1-\epsilon)^s(\epsilon)^{n-s} \tag{4}$$

Accordingly, this approach enables us to generate a prediction from a set of learners, such that the final prediction from the BMA is a weighted, linear combination of

[7]The prior probability of $p(D)$ is the same for each model and can thus be ignored.

each learner's probability distribution. We summarize this approach in Figure 2.

Figure 2: Summary of how we combine classifiers using BMA to generate predictions.

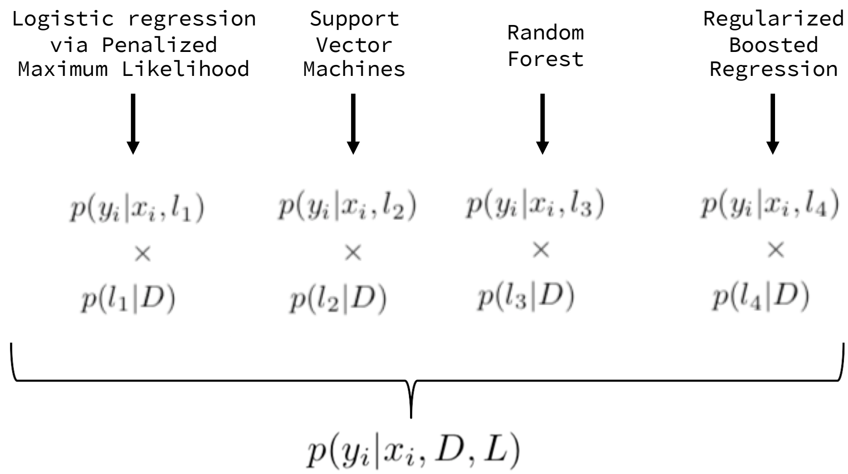

For each country we end up having two sets of predictions from BMAs, one that includes all ten features, and one that excludes the latent network measure of spatial dependence.

In sum, our full pipeline to generate predictions for the two approaches is shown below:

- Baseline: 9 Features[8] $\rightarrow$ PCA[9] $\rightarrow$ 4 ML Models [10] $\rightarrow$ BMA[11] $\rightarrow$ Prediction
- Baseline +Network : 10 Features[12] $\rightarrow$ PCA $\rightarrow$4 ML Models $\rightarrow$ BMA $\rightarrow$ Prediction

Assessment Strategy

Using our machine learning pipeline, we conduct two out-of-sample exercises. The first involves a 30-fold cross-validation exercise in which we set observations of our

[8] See table 2
[9] See pages 9-10
[10] See page 15
[11] See figure 2
[12] Including the network features described on pages 10-14.

dependent variables for each country to missing, and then see whether we can reconstruct those missing observations.[13] The second out-of-sample exercise is a forecasting one in which we set all the observations from the last month of conflict for each country to missing, and attempt to predict the missing values.

Cross-Validation Results

To evaluate the predictive ability of our model we utilize Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves. ROC curves depicting our model's performance for South Sudan are shown on the left-most plot of Figure 3 and PR curves on the right. To highlight the improvement our framework provides, we compare a version of our machine learning pipeline that includes network features (in blue) against a "base" version of the same pipeline excluding network features (but that includes traditional spatial measures, this model is shown in red).[14] For both ROC and PR curves we also provide area under the curve statistics in the bottom right. Last, we include separation plots to provide a visual summary of how well the two models fair in predicting conflict (Greenhill, Ward and Sacks, 2011).

[13]The Bayesian approach that we use to generate our network based diffusion features can accommodate missing data.

[14]The tuning parameters for each of the classifiers were kept static across these two versions.

Figure 3: Assessments of out-of-sample performance for the prediction of conflict occurrence in South Sudanese states using ROC curves, separation plots, and PR curves for 30-fold cross-validation.

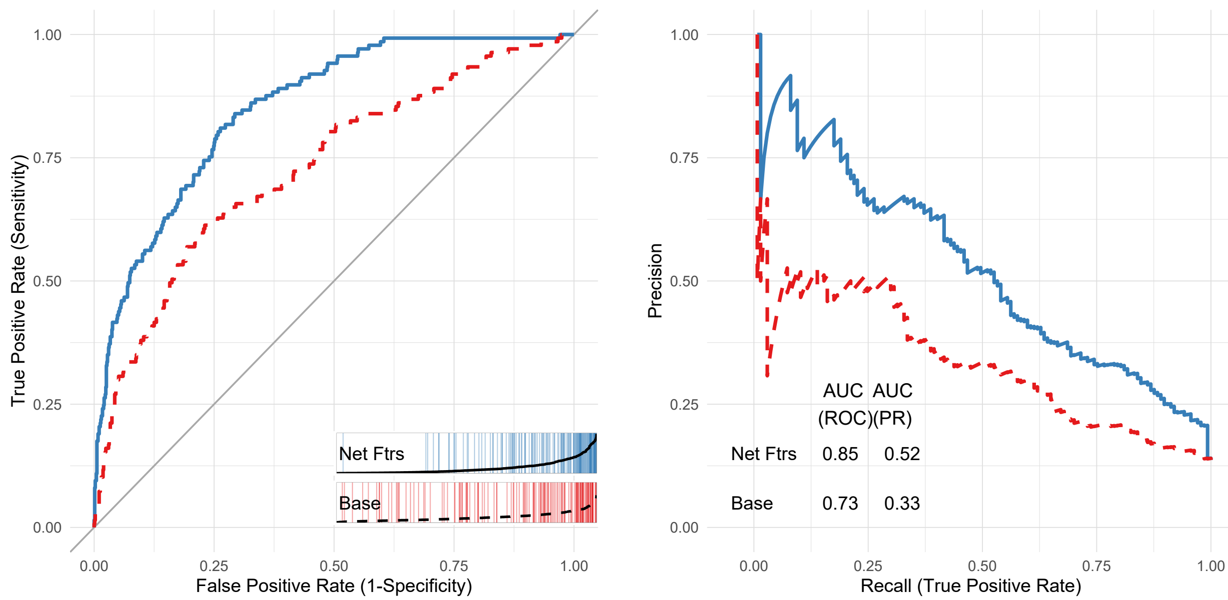

Across each of these tests it is clear that in the case of South Sudan utilizing our network based features as inputs notably improves the predictive ability of the pipeline. For the remaining seven countries, we summarize the performance results in Figure 4. For each country we achieve substantively better predictive ability when we include our network based features. Most importantly, the difference in performance is even more stark when we examine the area under the PR curves, which indicate that incorporating our network based features enables us to better capture actual instances of conflict.

Temporal Forecasting

Last, we examine the ability of our pipeline to generate forecasts of future events. To do this, we compare the two versions of our pipeline in their ability to predict conflictual events during the final month of our data. The results are summarized in Figure 5. In general, we find similar patterns as to what was shown from the cross-validation analysis. Particularly, we see that incorporating network based features al-

Figure 4: Out-of-sample AUC statistics from a 30-fold cross-validation exercise across the countries in our sample. Cases are sorted by the precision-recall (PR) performance of the Base + Network Features model, from low (left) to high (right).

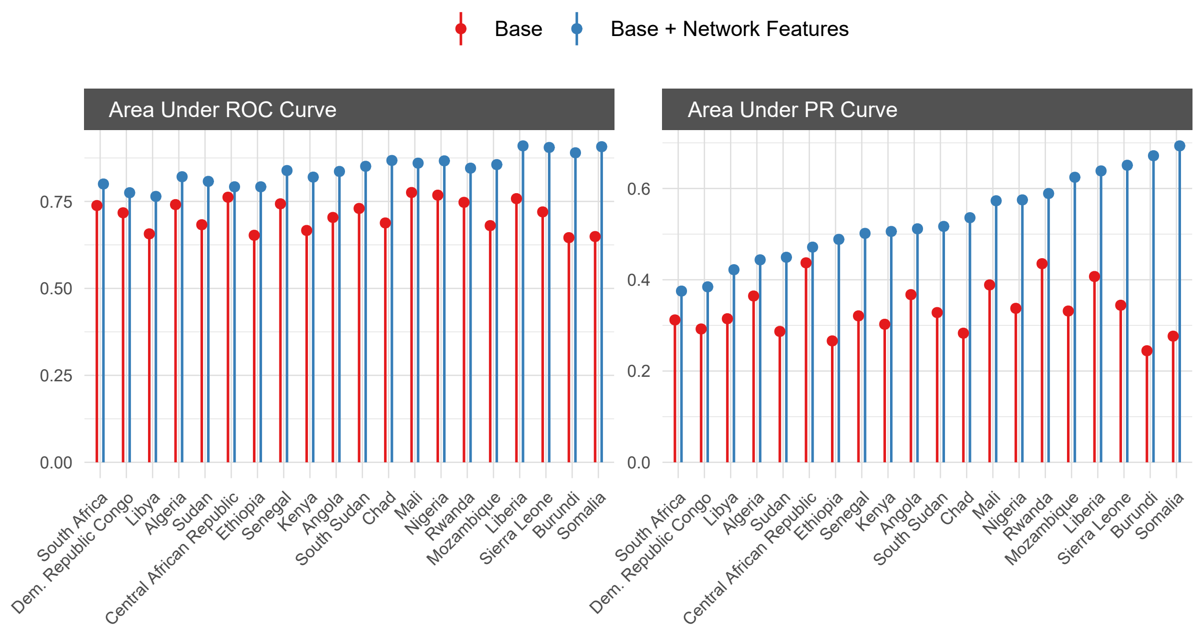

ways leads to an improvement in performance, and that improvement is most notable when we examine our ability to predict actual instances of conflict.

Explaining Differences in Performance

We observe significant variation in the benefit of using this network technique to account for spatial interdependencies across the different cases. While the inclusion of network features improves predictive performance in all countries, in some cases (notably Ethiopia, Burundi, and Somalia) the improvement is particularly stark, while in other cases (South Africa and the Central African Republic being noteworthy examples), the difference is slight. To try to explain this divergence between countries, we look at the relationship between our latent spatial lags and more traditional geographically based features. Figure 6 shows the proportion of variation in our latent diffusion matrix

Figure 5: Forecasting occurrences of **battles** by state. Cases are sorted by the precision-recall (PR) performance of the Base + Network Features model, from low (left) to high (right).

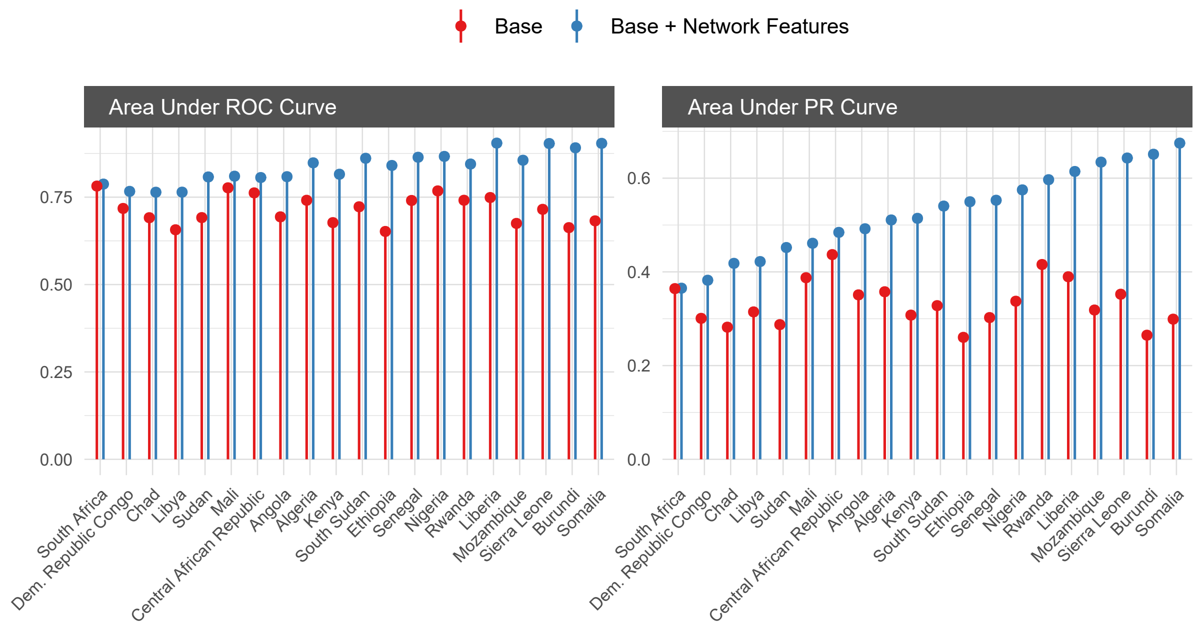

that is explained by contiguity and centroid distance in each country.

Evidence shown in Figure 6 suggests that geography does not primarily explain variation in the latent diffusion matrix, in fact only one of our 20 countries has geography explaining more than half of the variance in our latent diffusion matrix. This is consistent with findings from the interstate conflict literature on diffusion. More interestingly, the three countries where the correlation is strongest are South Africa, the Democratic Republic of the Congo, and the Central African Republic, notable for being three countries where our technique yielded the least improvement in forecasting accuracy. The countries with the smallest correspondence between spatial weights matrices are the three countries where the network model had some of the most notable overperformance. If a country has relatively small patterns of spatial interdependence, or if these patterns are relatively close to the geographic patterns we included as spatial weights in the models, then we should expect to see convergence in predictive performance between the models with and without these network features. On the other hand, when there are strong patterns of spatial diffusion that do not map well onto adjacency matrixes, models that only include classic weight matrices will fall behind.

Figure 6: Visualization shows the proportion of variation explained in the latent diffusion matrices by contiguity and centroid distance for each country.

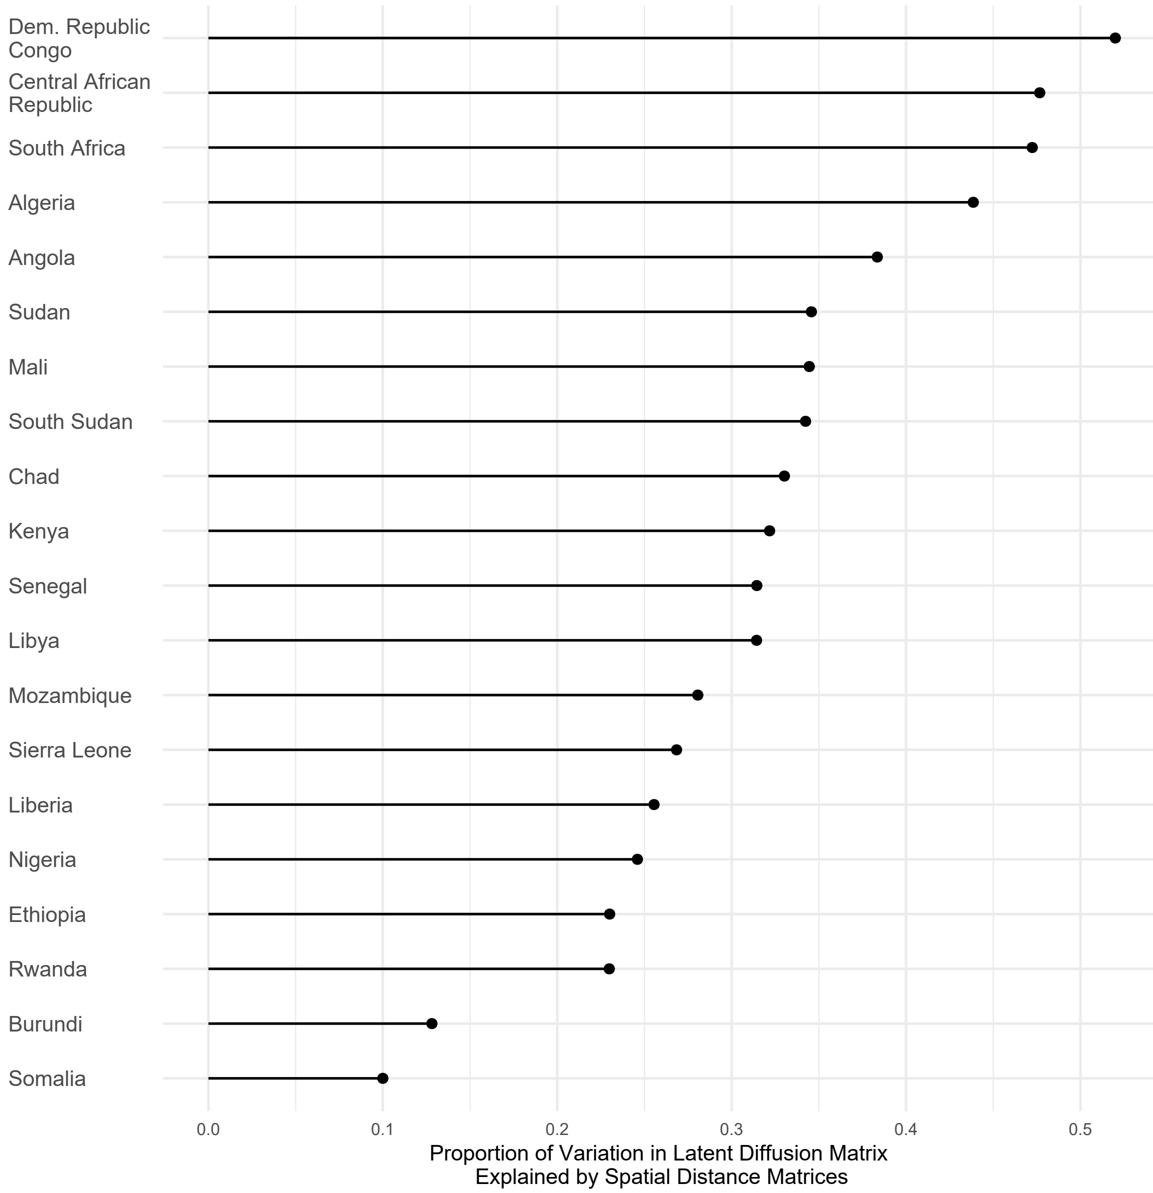

Discussion

As has been established in the literature on civil conflict, violence diffuses spatially, but the way it diffuses is conditional on many factors—factors that are difficult to observe in a granular way. To help account for these varying spatial factors, we turn to a network approach that allows us to infer a latent diffusion matrix based on the history of conflict in a given country. When we apply this technique to civil conflict in twenty African countries, we consistently find that we are able to better predict violence out-of-sample. Importantly, we are able to both more accurately predict where there will be an absence of conflict, and when and where battles will occur. This improvement holds when we look at countries with high levels of violence, and those where violence is only moderate. In addition, by investigating the correlation between the geographic and latent spatial weights matrices, our study offers new insights into the types of cases that will benefit the most from this new approach.

Our study highlights the potential of our approach to substantially enhance the prediction of conflict at fine-grained levels of analysis. Though it makes an important contribution, a critical question remains: what conditions make conflict more likely to spillover from one region to another? In the future, we hope to examine not only the regional level covariates that make a particular sub-national unit more or less likely to spread conflict but also the role played by shared attributes of those units. Doing so will enable us to both test existing theories of conflict contagion and also provide insight into how to limit the spread of violence in the future. Most importantly, we can combine these factors with measures of third-party interventions in civil conflict – peacekeepers, election-monitors, and military personnel – to determine the extent to which policy interventions can contain and roll back civil conflict. This will generate an important contribution not just for scholarship, but to broader audiences including policy practitioners.

While this approach has been quite fruitful in predicting conflict, we believe it can be applied more broadly to different phenomenon that diffuse spatially. The utility of this approach depends on a few key requirements. First, there must be a phenomenon of interest that displays some spatial interdependencies – if these are not present, the latent network measure will simply be adding noise to an estimator. Second, there must be multiple variegated pathways in which spatial interdependence can effect the phenomenon of interest, since if there were only one pathway, it would be preferable to simply measure that pathway directly. Finally, there should be some reason that you cannot simply use data on each of the different potential pathways to measure diffusion directly. If these requirements hold, we believe this method can provide significant improvements on using a traditional geographic lag, in our ability to predict not just the spread of conflict, but of many different international phenomena. Notably, our approach might be particularly useful for studying the diffusion of international trade agreements and the cross-national spread of civil war. Both political phenomena express known patterns of interdependence with multiple pathways but often elude straightforward measurement strategies. As our study demonstrates, it is in these areas of political science, where measurement is difficult, interdependence is known, and competing mechanisms might be at play, where a latent network approach will prove most beneficial.

References

Beardsley, Kyle. 2011. "Peacekeeping and the Contagion of Armed Conflict." *The Journal of Politics* 73(04):1051–1064.

Beck, Neal, Gary King and Langche Zeng. 2000. "The Problem with Quantitative Studies of Conflict: A Conjecture." *American Political Science Review* 94(1):21–36.

Bennett, D. Scott and Allan C. Stam. 2009. "Revisiting Predictions of War Duration." *Conflict Management and Peace Science* 26(3):256–267.

Bennett, James and Stan Lanning. 2007. The Netflix Prize. In *Proceedings of KDD cup and workshop*.

Braithwaite, Alex. 2010. "MIDLOC: Introducing the Militarized Interstate Dispute Location dataset." *Journal of Peace Research* 47(1):91–98.

Brandt, Patrick T., Michael Colaresi and John R. Freeman. 2008. "The Dynamics of Reciprocity, Accountability, and Credibility." *Journal of Conflict Resolution* 52(3):343–374.

Bremer, Stuart A. 1992. "Dangerous dyads conditions affecting the likelihood of interstate war, 1816-1965." *Journal of Conflict Resolution* 36(2):309–341.

Buhaug, Halvard and Kristian Skrede Gleditsch. 2008. "Contagion or Confusion? Why Conflicts Cluster in Space." *International Studies Quarterly* 52(2):215–233.

Carter, David B. and Curtis S. Signorino. 2010. "Back to the Future: Modeling Time Dependence in Binary Data." *Political Analysis* 18(3):271–292.

Cederman, Lars-Erik, Kristian Skrede Gleditsch and Luc Girardin. 2009. "Ethnic Triads: Assessing the Influence of Kin Groups on Civil Wars." *World Politics* 61(3):403–437.

Cederman, Lars-Erik, Nils B Weidmann and Kristian Skrede Gleditsch. 2011. "Horizontal inequalities and ethnonationalist civil war: A global comparison." *American Political Science Review* 105(03):478–495.

Cederman, Lars-Erik, Simon Hug and Lutz F. Krebs. 2010. "Democratization and civil war: Empirical Evidence." *Journal of Peace Research* 47(4):377–394.

Chang, Chih-Chung and Chih-Jen Lin. 2011. "LIBSVM: A library for support vector machines." *ACM transactions on intelligent systems and technology (TIST)* 2(3):27.

Chen, Tianqi and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM pp. 785–794.

Clayton, Govinda and Kristian Skrede Gleditsch. 2014. "Will we see helping hands? Predicting civil war mediation and likely success." *Conflict Management and Peace Science* 31(3):265–284.

Colaresi, Michael and Zuhaib Mahmood. 2017. "Do the robot: Lessons from machine learning to improve conflict forecasting." *Journal of Peace Research* 54(2):193–214.

Danneman, Nathan and Emily Hencken Ritter. 2014. "Contagious rebellion and preemptive repression." *Journal of Conflict Resolution* 58(2):254–279.

de Marchi, Scott, Christopher Gelpi and Jeffrey D. Grynaviski. 2004. "Untangling Neural Nets." *American Political Science Review* 98(02):371–378.

Dorff, Cassy and Jessica M. Braithwaite. 2018. "Fear of Nonviolent Organizing in Mexico's Criminal Conflict." *Journal of Global Security Studies* 3(3):271–284.

Dorff, Cassy L., Shahryar Minhas and Michael D. Ward. 2017. Latent Networks and Spatial Networks in Politics. In *Oxford Handbook of Political Networks*, ed. Jennifer Nicoll

Victor, Mark Lubell and Alexander H. Montgomery. Oxford, England: Oxford University Press chapter 9.

Franzese, Robert J. and Jude C. Hays. 2008. "Interdependence in Comparative Politics: Substance, Theory, Empirics, Substance." *Comparative Political Studies* 41(4/5):742–780.

Friedman, Jerome, Trevor Hastie and Rob Tibshirani. 2010. "Regularization paths for generalized linear models via coordinate descent." *Journal of statistical software* 33(1):1.

Gartzke, Erik. 2006. "The (Other) Liberal Peace: Capitalism, Common Interests, Development, and Democracy." Columbia University, unpublished PDF via author.

Gleditsch, Kristian Skrede. 2007. "Transnational dimensions of civil war." *Journal of Peace Research* 44(3):293–309.

Gleditsch, Kristian Skrede and Michael D. Ward. 2010. "Contentious Issues and Forecasting Interstate Disputes." Presented to the 2010 Annual Meeting of the International Studies Association.

Greenhill, Brian, Michael D. Ward and Audrey Sacks. 2011. "The Separation Plot: A New Visual Method for Evaluating the Fit of Binary Data." *American Journal of Political Science* 55(4):991–1002.

Guo, Weisi, Kristian Gleditsch and Alan Wilson. 2018. "Retool AI to forecast and limit wars.".

Gurr, Ted Robert and Barbara Harff. 1996. "Early warning of communal conflict and humanitarian crisis." *Tokyo: United Nations University Press, Monograph Series on Governance and Conflict Resolution* .

Hays, Jude C., Aya Kachi and Robert J. Franzese. 2010. "A Spatial Model Incorporating Dynamic, Endogenous Network Interdependence: A Political Science Application." *Statistical Methodology* 7(3):406–428.

Hegre, Håvard, Joakim Karlsen, Håvard Mokleiv Nygård, Håvard Strand and Henrik Urdal. 2013. "Predicting Armed Conflict, 2011–2050." *International Studies Quarterly* 57(2):250–270.

Hegre, Havard and Nicholas Sambanis. 2006. "Sensitivity Analysis of Empirical Results on Civil War Onset.pdf." *Journal of Conflict Resolution* 50(4):508–535.

Hill, Daniel W. and Zachary M. Jones. 2014. "An Empirical Evaluation of Explanations for State Repression." *American Political Science Review* 108(03):661–687.

Hoeting, Jennifer A., David Madigan, Adrian E. Raftery and Christopher T. Volinsky. 1999. "Bayesian Model Averaging: A Tutorial." *Statistical Science* 14(4):382–417.

Hoff, Peter D. 2005. "Bilinear Mixed-Effects Models for Dyadic Data." *Journal of the American Statistical Association* 100(4690):286–295.

Johnson, Douglas H. 2014. "Briefing: the crisis in South Sudan." *African Affairs* 113(451):300–309.

Jones, Zachary M and Yonatan Lupu. 2018. "Is There More Violence in the Middle?" *American Journal of Political Science* 62(3):652–667.

Kathman, Jacob D. 2010. "Civil war contagion and neighboring interventions." *International Studies Quarterly* 54(4):989–1012.

Liaw, Andy and Matthew Wiener. 2002. "Classification and Regression by randomForest." *R news* 2(3):18–22.

Metternich, Nils, Shahryar Minhas and Michael Ward. 2017. "Firewall? Or Wall on Fire? A Unified Framework of Conflict Contagion and the Role of Ethnic Exclusion." *Journal of Conflict Resolution* 61(6):1151–1173.

Metternich, Nils W., Cassy Dorff, Max Gallop, Simon Weschle and Michael D. Ward. 2013. "Antigovernment Networks in Civil Conflicts: How Network Structures Affect Conflictual Behavior." *American Journal of Political Science* 57(4):892–911.

Meyer, David and FH Technikum Wien. 2001. "Support vector machines." *R News* 1(3):23–26.

Minhas, Shahryar, Peter D. Hoff and Michael D. Ward. 2019. "Inferential Approaches for Network Analysis: AMEN for Latent Factor Models." *Political Analysis* 27(2):208–222.

Muchlinski, David, David Siroky, Jingrui He and Matthew Kocher. 2016. "Comparing Random Forest with Logistic Regression for Predicting Class-Imbalanced Civil War Onset Data." *Political Analysis* 24(1):87–103.

Neumayer, Eric and Thomas Plümper. 2010. "Spatial effects in dyadic data." *International Organization* 64(01):145–166.

O'Loughlin, John, Frank D. W. Witmer, Andrew M. Linke, Arlene Laing, Andrew Getterman and Jimy Dudhia. 2012. "Climate Variability and Conflict Risk in East-Africa, 1990-2009." *Proceedings of the National Academy of Sciences* 109(45):18344–18349.

O'Loughlin, John and Frank DW Witmer. 2011. "The localized geographies of violence in the North Caucasus of Russia, 1999–2007." *Annals of the Association of American Geographers* 101(1):178–201.

Osorio, Javier, Livia Schubiger and Michael Weintraub. 2018. "Disappearing dissent?

Repression and state consolidation in Mexico." *Journal of Peace Research* 55(2):252–266.

Perry, Chris. 2013. "Machine learning and conflict prediction: a use case." SSRN Working Paper.

Raftery, Adrian E., Tillman Gneiting, Fadoua Balabdaoui and M. Polakowski. 2005. "Using Bayesian Model Averaging to Calibrate Forecast Ensembles." *Monthly Weather Review* 133:1155–1174.

Resnick, Paul and Hal R. Varian. 1997. "Recommender Systems." *Communications of the ACM* 40(3):56–58.

Richardson, Lewis F. 1960. *Arms and Insecurity*. Chicago and Pittsburgh, PA: Quadrangle/Boxwood.

Salehyan, Idean and Kristian Skrede Gleditsch. 2007. "Refugees and the Spread of Civil War." *International Organization* 60(2):335–366.

Schrodt, Philip A. 1990. "A Methodological Critique of a Test of the Effects of the Maharishi Technology of the Unified Field." *Journal of Conflict Resolution* 34(4):745–755.

Schrodt, Philip A. 2014. "Seven deadly sins of contemporary quantitative political analysis." *Journal of Peace Research* 51(2):287–300.

Schutte, Sebastian and Nils B. Weidmann. 2011. "Diffusion patterns of violence in civil wars." *Political Geography* 30(3):143–152.

Shellman, Stephen M., Brian P. Levey and Joseph K. Young. 2013. "Shifting Sands: Explaining and Predicting Phase Shifts by Dissident Organizations." *Journal of Peace Research* 50(3):319–336.

Sundberg, Ralph and Erik Melander. 2013. "Introducing the UCDP Georeferenced Event Dataset." *Journal of Peace Research* 50(4):523–532.

Townsley, Michael, Shane D Johnson and Jerry H Ratcliffe. 2008. "Space time dynamics of insurgent activity in Iraq." *Security Journal* 21(3):139–146.

Ward, Michael D. and Andrew Kirby. 1987. "Re-examining Spatial Models of International Conflicts." *Annals of American Association of Geographers* 31(1):86–105.

Ward, Michael D., Brian D. Greenhill and Kristin M. Bakke. 2010. "The Perils of Policy by p-value: Predicting Civil Conflicts." *Journal of Peace Research* 47(4):363–375.

Ward, Michael D. and Kristian Skrede Gleditsch. 2002. "Location, Location, Location: An MCMC Approach to Modeling the Spatial Context of War and Peace." *Political Analysis* 10(3):244–260.

Ward, Michael D., Nils W. Metternich, Cassy L. Dorff, Max Gallop, Florian M. Hollenbach, Anna Schultz and Simon Weschle. 2013. "Learning from the Past and Stepping into the Future: Toward a New Generation of Conflict Prediction." *International Studies Review* 16(4):473–644.

Ward, Michael D., Randoph M. Siverson and Xun Cao. 2006. "Everybody Out of the Pool: Reconstructing the Democratic Peace." Working Paper 55, Center for Statistics and Social Science, University of Washington.

Weidmann, N.B. and M.D. Ward. 2010. "Predicting conflict in space and time." *Journal of Conflict Resolution* 54(6):883–901.

Wucherpfennig, Julian, Nils B. Weidmann, Luc Girardin, Lars-Erik Cederman and Andreas Wimmer. 2011. "Politically Relevant Ethnic Groups Across Space and Time: Introducing the GeoEPR Dataset." *Conflict Management and Peace Science* 28(5):423–437.