RESEARCH

Open Access

Forecasting the locational dynamics of transnational terrorism: a network analytic approach

Bruce A Desmarais^{1*} and Skyler J Cranmer²

Abstract

Efforts to combat and prevent transnational terrorism rely, to a great extent, on the effective allocation of security resources. Critical to the success of this allocation process is the identification of the likely geopolitical sources and targets of terrorism. We construct the network of transnational terrorist attacks, in which source (sender) and target (receiver) countries share a directed edge, and we evaluate a network analytic approach to forecasting the geopolitical sources and targets of terrorism. We integrate a deterministic, similarity-based, link prediction framework into a probabilistic modeling approach in order to develop an edge-forecasting method. Using a database of over 12,000 transnational terrorist attacks occurring between 1968 and 2002, we show that probabilistic link prediction is not only capable of accurate forecasting during a terrorist campaign, but is a promising approach to forecasting the onset of terrorist hostilities between a source and a target.

Introduction

The accurate forecasting of transnational terrorism is among the most pressing problems of contemporary security policy. Counter-terrorism efforts may be greatly aided if resources, ranging from analytic focus to prevention capabilities, can effectively target the source of terrorist threats before terrorists can launch attacks. Traditional efforts to statistically forecast terrorism are complicated by the fact that most forecasting models perform best when there is a long series of data to study, thus rendering them impotent for identifying new sources of threat and forecasting attacks from sources that have not attacked before. For example, a protracted terrorist campaign may provide enough data for accurate time-series forecasting, but, from a policy perspective, may not be useful because the target government already understands that it is under attack from that source. The major challenge in forecasting transnational terrorism, as we see it, is identifying sources of threat (countries who's nationals conduct terrorist attacks against a given target state)

before any attacks from that source have been observed and quantifying the extent of the threat.

We approach the problem of forecasting transnational terrorism from a network analytic perspective with the supposition that the structure of the transnational terrorist network may be the best predictor of its own evolution. We construct the global transnational terrorism network, in which a directed edge exists from the state that produces a transnational terrorist to the state attacked by that terrorist, using the ITERATE dataset [1]. The data cover more than 12,000 transnational terrorist attacks between 1968 and 2002. We create an edge forecasting framework for the transnational terrorist network by integrating a deterministic similarity-based edge prediction method developed by Liben-Nowell and Kleinberg [2] with a model-based probabilistic approach developed by Hanneke, Fu, and Xing [3]. The result is a likelihood-based forecasting model capable of quantifying the transnational terrorist threat one state poses to another, even before any attacks have occurred between the two states. Specifically, we predict edges in the transnational terror network by substituting the network structure embedded in the recent patterns of transnational terrorism into the model that best predicted the network up to time $t - 1$. Thus, the predictive models for t are not based on the data

*Correspondence: desmarais@polsci.umass.edu

¹ Department of Political Science, University of Massachusetts at Amherst, Amherst, Massachusetts 01003, USA

Full list of author information is available at the end of the article

from t . The result is a forecasting model that predicts edges that occur in the next time period (year) with probabilities orders of magnitude greater than those that do not occur. As such, our model provides forecasts of the locational dynamics of terrorist threats both before and after the first attacks have been realized. Our model lifts a major limitation of traditional approaches to forecasting transnational terrorist events and can provide early warnings of emerging threats.

Background

Though transnational terrorism has been a highly visible policy problem for the United States since the terrorist attacks of September 11, 2001, transnational terrorism is an older phenomena that, in many ways, came into its heyday with skyjacking, hostage taking, and bombing campaigns in the 1960's and 1970's. Quantitative research on transnational terrorist violence^a has focused largely on two related themes: the predictors of terrorist violence and trend analysis of the violence itself. The first thread of literature, developed largely in the field of political science, has linked higher probabilities of transnational terrorist violence to target states that have democratic governments [4], politically left governments [5], more veto-players in their governments [6], are perceived to be more likely to grant concessions to terrorists [7], and have further economic reach [8]. While this literature does much to shed light on the factors that may make a state more likely to suffer transnational terrorist attacks, it does little to provide forecasts of terrorist violence with any degree of precision. A second thread of literature has focused on trends in transnational terrorist violence itself. This literature, based at the intersection of political science and economics, has established that terrorist attacks exhibit cycling behavior [9], the number of terrorist attacks is decreasing but their lethality is increasing [10], substantial increases in levels of violence are usually unsustainable for the terrorist group [11], terrorist attacks are persistent following shocks in states suffering from low levels of terrorism (not for states suffering high levels of terrorism) [12], and terrorists tend to substitute targets when one type of target is hardened [13].

The literature that might help forecast terrorist events, however, has a major limitation: existent findings are either overly broad (i.e. democracies are more likely to be attacked by terrorists) or rely on voluminous attack data before trends can be identified and forecasts made. The lack of specificity from most statistical models of terrorist attacks is problematic from a policy perspective because, while it may warn a government that it is at elevated risk of terrorist attacks based on its attributes, it does little to inform the government about the source or timing of the threat. The trend analyses, while useful for predicting the ebb and flow of terrorist campaigns, are limited

by not being able to make accurate predictions about the onset of hostilities. Furthermore, the literature on why discontented groups resort to terrorism [14,15] is dedicated more to *explaining* the onset of terrorist campaigns than *predicting* such onset.

To date, the literature is moot in terms of statistical models designed not only to forecast attacks during terrorist campaigns, but to forecast emerging terrorist threats, *and* do so with a high degree of locational specificity: predicting not only when but from which country transnational terrorist attacks will emanate. Such a tool would provide insight from an academic perspective and prove useful from a policy perspective. It is with this aim that we develop a hybrid methodology from both deterministic and stochastic edge prediction techniques to predict the timing and locational dynamics of transnational terrorism.

Data

The data for our study are drawn from the "International Terrorism: Attributes of Terrorist Events" (ITERATE) dataset [1]. These data are well suited to our aims as they cover all transnational terrorist attacks over a 34 year period (1968–2002). The operational definition of transnational terrorism used for data collection is "the use, or threat of use, of anxiety-inducing, extra-normal violence for political purposes by any individual or group, whether acting for or in opposition to established governmental authority, when such action is intended to influence the attitudes and behavior of a target group wider than the immediate victims and when, through the nationality or foreign ties of its perpetrators, its location, the nature of its institutional or human victims, or the mechanics of its resolution, its ramifications transcend national boundaries." ([1], p. 2) The ITERATE data are one of the most comprehensive and commonly used data sets on transnational terrorism (for example, see [4,5,8,10–12,16]).

The ITERATE data, among other variables, codes the known nationalities of the terrorists who participate in a given transnational attack and the location of the attack. From these variables, we create the network of transnational terrorist attacks by year. We code a directed edge as existing from country i to country j if a national of country i participates in a terrorist attack on location j ; we call these *terrorist edges*. Note that many attacks, such as the September 11, 2001 attacks on the United States, are perpetrated by non-state actors (e.g., al-Qaeda). In such instances, we code a directed edge from every known nationality of the attackers involved in the attack implementation to the target vertex. One may object that transnational terrorists sometimes emigrate from their homelands, developing doctrine, training, and establishing support networks in states that are neither their home

nor their target. However, when we consider the process of radicalization, this is less concerning. When a terrorist emigrates to develop their plans and capabilities, the process of radicalization proceeded the emigration and, thus, will typically take place in their home country. It is the case that some of the better-known terrorist masterminds spend substantial portions of their lives abroad, hopping from place to place, but this is more typical of senior group leadership and less typical of those who actually execute the attacks. As such, many of the globe-hoppers would not be included in the dataset anyway (if they did not participate in an identifiable way in the attack). Second, attempting to consider where a terrorist has spent their “important years” would introduce substantial coding problems. For example, many of the 9/11 hijackers received critical training (i.e. flight training) in the United States, but drawing a looping edge from the U.S. to the U.S. makes little substantive sense. Furthermore, data on where terrorists have spent their time prior to their attacks is not generally available and is prone to error and uncertainty; terrorists typically try to conceal their training camps and hideouts. As it stands, no information is available in our dataset beyond the known nationalities of attackers, and so we use that as a criteria for constructing the network.

To construct the network, we must also define the universe of possible terrorist edges in a given year, which requires a definition of membership in the international system. We rely principally on the definition of the state system provided by the Correlates of War Project, which keeps a database of all countries that belong to the international community in a given year.^b We augment membership in the international community by adding a few non-state actors that have been either the source or the target of transnational terrorism (e.g. Palestine and the United Nations). Though such entities are not technically countries, including them in the network is preferable to excluding them because contested territories are important producers of transnational terrorism and not including international governmental organizations would misclassify the target of the attack (i.e. where it occurred rather than who was targeted).

All together, we produce a network with a median of 175 vertices over the time span that we study. A feature of these data that is somewhat atypical in network analysis is that self-ties are present. It is common for terrorists from state i to commit an attack in state i . While we include these “loops” in our analysis, it is important to point out that all of these loops are affiliated in some way with a *transnational* terrorist incident. The ITER-ATE data do not code domestic terrorist attacks and, as such, loops represent acts in which a native of the target country collaborates with foreign terrorists to launch an attack.

Two empirical features of this network are particularly relevant to our forecasting approach: edge innovation and transitivity. By edge “innovation,” we refer to terrorist edges that did not exist in the previous time period or periods. Substantively, edge innovation means that a terrorist from country i attacked a target in country j , where there had not been an attack from i to j for some time. The number of new and recurrent edges in the network over time is depicted in Figure 1. The shading indicates the degree to which the current edges also occurred in the past. What we see here is that, over the period 1980 – 2002, edges in the current network are just as likely to be innovations in the network (i.e., those that have not formed in the previous 10 years), as they are to be recurrent edges from the previous year. The relatively low degree of recurrence in the network implies that forecasting based on the dynamics of edge stability/memory will fail to capture a substantial portion of the year-to-year activity in the network. We can also see the importance of forecasting models designed to predict edge innovations; traditional time-series approaches to forecasting these innovations would prove fruitless.

The second feature we focus on, transitivity, is a measure of how important the network proximity of two vertices is to the likelihood of terrorist edge forming between those vertices (i.e., if the network is highly transitive, then configurations in which $a \rightarrow b \cap a \rightarrow c \Rightarrow b \rightarrow c$ will be common). The foundational work on edge prediction in networks uses measures of network proximity to forecast future ties [2]. If transitivity is a persistent feature in a network, then proximity-based forecasts should be relatively successful. However, if, in the terrorism network, the neighbor of a neighbor is not a neighbor, then another approach to link forecasting would be necessary.

Figure 1 Number of country-country edges in the transnational terrorism network by year. Shading indicates the number of edges that are recurrent from previous years and how many are new edges, or innovations in the network.

We measure the degree of transitivity in the transnational terrorism network over time using a conditional uniform graph (CUG) test. The CUG test for transitivity [17,18] evaluates the observed degree of transitivity (the proportion of potential transitive triads that are actually closed) against a null distribution computed on a random sample of networks without transitivity. The null distribution of networks places a uniform probability on the possible networks, which have the same number of mutual and asymmetric dyads as the observed networks (i.e., the null distribution of networks is a uniform distribution conditional upon having the same dyad census as the observed network). The results of the CUG tests are depicted in Figure 2. In 21 of the 23 years depicted, the observed value of transitivity is larger than any of the 1,000 simulated null values. This indicates that there is a substantial degree of transitivity in the terrorism network. The exceptions are 1999 and 2001. It is not clear why the network would not be noticeably transitive in this period. The lower overall density of the network may make it more difficult to identify this feature.

These two results, (1) that many edges represent innovations with respect to the recent past and (2) the network exhibits significant transitivity, indicate that proximity-based link prediction should be a fruitful exercise on the transnational terrorism network. The persistence of innovation means that any information that can be leveraged about the indirect ties in the network will be quite valuable. The transitivity indicates that indirect ties will have predictive power with respect to edge formation.

Why is transnational terrorism transitive?

It is not immediately intuitive why the network of transnational terrorist attacks would exhibit the high degree of transitivity we observe with the CUG test. Conflictual international networks, such as international war [19], exhibit substantial intransitivity; whereas transitivity is commonly a feature of cooperative international networks (e.g. military alliances) [20,21].

When approaching this question theoretically, it is useful to distinguish between two common types of terrorist groups by motivation: ethno-nationalist groups and religious/ideological groups. Ethno-nationalist groups have specific goals of creating a state for their ethnic group or liberating their nation from occupation (real or perceived) and, as such, their attacks tend to be confined to only one target. Examples of such groups include the Provisional Irish Republican Army (PIRA) and the Kurdistan Workers' Party (PKK). For ethno-nationalist groups, terrorist violence is tied directly to the attainment of their nationalistic goals. Conversely, religious/ideological groups are organized around a religious or political/ideological ideal and tend to lack a clear policy objective.^c However, such groups tend to have broader, multinational, support bases

Figure 2 Conditional uniform graph tests for transitivity in the transnational terrorism network. For each year, the histogram of the null distribution of transitivity is depicted in gray, and the observed value of transitivity is located at the red line. The null distribution is constructed by measuring the transitivity of 1,000 networks simulated from a uniform graph distribution conditional on the dyad census.

and to attack targets in a number of countries. They tend to target countries that exhibit religious/ideological cultures that are counter to the groups' ideals. Groups such as al-Qaeda and the Red Army Faction fall into this category.

We offer two theoretical explanations for the transitivity of the transnational terrorist network, the first of which applies more to religious/ideological groups than to ethno-nationalist groups. These explanations are not mutually exclusive and both may be occurring simultaneously. First, consider a religious/ideological terrorist group X headquartered in state A . Suppose group X 's ideology makes states B and C targets for attacks or campaigns. Assuming the group has a support base in state A , it is likely to recruit and train heavily there. During the course of its campaigns against B and C , group X continues to recruit from A , but also recruits from B and C . This pattern is both feasible and reasonably common for religious/ideological groups who may be headquartered in one country, but whose support base spans several. After some time, group X will have operational units in each of the three countries. Early (and continued) recruiting in A for attacks in B and C result in an out-2-star from A (i.e., two edges that form the start of a transitive triad). However, when recruits from B participate in attacks on C (or recruits from C participating in attacks on B), a transitive triad is formed. This process, through which ideological/religious groups garner multinational support bases, is illustrated in the left column of Figure 3.

A second situation that can give rise to transitivity in the network, for either ethno-nationalist or religious/ideological groups, is an environment with multiple terrorist groups. For example, extreme right and extreme left groups may be active in the same region and their overlapping attacks may form transitive patterns. This two-group process is illustrated in the right column of Figure 3. Furthermore, a multi-group region need not be populated with groups of opposing ideals in order to form transitive triads. Bapat [22] provides an explanation for why infighting between groups with a common opponent is common, especially as peace processes progress. When a target state makes a move towards peace with its attackers, that move will usually involve active negotiations with one or several groups. In order to get their desired concessions, the negotiating group must moderate its behavior. This, however, is often not welcome by other groups or by radicals within the negotiating group. If the negotiating group is successful, it will, at best, seize power away from rival groups or, at worst, seize power away from rivals while negotiating a settlement that is unacceptable to a broad swath of rivals and radicals; either way, rival groups have an incentive to resist the change and radicals within the negotiating group have an incentive to form splinter groups. Resisting groups have been known to launch

spoiler campaigns against the target state (attacks designed to provoke the target state and make peace less likely) or launch attacks directly on the negotiating group. At the same time, the negotiating group needs a monopoly on terrorist violence in order to be a reliable negotiating partner for the government, so it has an incentive to crack down on dissidents and rivals.

Method

Our method of forecasting is based on a probabilistic modeling framework, and incorporates proximity measures from a deterministic edge prediction method. For the probabilistic base of our method, we use the temporal exponential random graph modeling (TERGM) framework developed by Hanneke, Fu, and Xing [3]. To define the specific TERGMs, we incorporate vertex-similarity measures drawn from Liben-Nowell and Kleinberg [2], who originally used those measures in a deterministic manner. We first review the TERGM approach to modeling networks, then we describe the vertex-similarity measures and how we incorporate them into the TERGM modeling framework.

The exponential random graph model (ERGM) [23,24] is defined by a flexible discrete, exponential family joint distribution for a network N . The likelihood function is

$$\mathcal{P}(N, \theta) = \frac{\exp\{\theta' \Gamma(N)\}}{\sum_{\text{all } N^* \in \mathcal{N}} \exp\{\theta' \Gamma(N^*)\}}, \quad (1)$$

where $\theta \in \mathbb{R}^p$ is the parameter vector, $\Gamma: \mathcal{N} \rightarrow \mathbb{R}^p$ is a vector of statistics that are computed on the network, and the summation in the denominator is a normalizing constant computed on the support of N (denoted $\mathcal{N}$). This model is flexible in that Γ can be specified to capture virtually any form of interdependence among the edges in the networks, as well as dependence of the edges on exogenous features.

Hanneke, Fu, and Xing [3] extended the ERGM to model a network observed at numerous discrete time points, which extends the first-order dependence of the original model formulation [25]. This is accomplished mathematically by including past realizations of N in the Γ that defines the ERG distribution for the current network. The network is observed in T discrete time periods. Let N^t be the observed network at time t . Temporal interdependence of the networks can be built into the model by conditioning N^t on K previous realizations of the network.

The probability of observing N^t in the temporally interdependent ERGM (TERGM) of order K is written as

$$\mathcal{P}(N^t | K, \theta) = \frac{\exp\{\theta' \Gamma(N^t, N^{t-1}, \dots, N^{t-K})\}}{C(\theta, N^{t-K}, \dots, N^{t-1})}. \quad (2)$$

Recall that the denominator in Equation 1 is a normalizing constant (i.e., a partition function). We update the

Figure 3 Illustration of single- and two-group processes that result in transitivity in the transnational terrorism network. The left column illustrates a single-group process that results in transitive closure. The right column illustrates a two-group process that also results in transitive closure.

notation in 2 to reflect this property. The likelihood function is defined as $\prod_{t=K+1}^T \mathcal{P}(N^t|K, \theta)$.^d In our implementation, a new θ is estimated for each t (i.e., θ^t). We allow parameters to vary by period due to temporal heterogeneity in the network structure, which is indicated by the considerable decrease in the density of the network over time and is visible in Figure 1.

In what can safely be called the foundational work on edge prediction in complex networks, Liben-Nowell and Kleinberg [2] capitalize the insight that vertices close or similar to each other in a network are likely to link in the future. They describe multiple measures of “proximity” in a network. Each measure of proximity results in a score $\delta(i, j)$ for each dyad of vertices ij in the network. These scores are then computed on a training network defined over an interval of the past. Dyads are then ranked with

respect to these scores. For prediction, the potential new edges are constrained to be those meeting a certain degree threshold in the training network. The subset of dyads defined on this network are then ranked with respect to their proximity scores and dyads that have high δ are predicted to be edges in the next time interval. Each proximity measure is evaluated with respect to its predictive performance.

We build upon this work by integrating proximity measures into the Γ of the TERGM. In our approach, the individual proximity measures are combined into a single model using the estimated weights (θ). A single best performing proximity measure need not be selected, a feature which provides considerable flexibility in the proximity terms to include. Also, the TERGM estimates permit us to forecast the probability of edges in the future. Each

proximity measure δ is integrated into the TERGM by adding

$$\Gamma(N^t, N^{t_0, t-1}) = \sum_{ij} N_{ij}^t \delta(i, j)^{t_0, t-1}.$$

If states i or j were not in the international system during the period taken as the training network interval $[t_0, t-1]$, then we set $\delta(i, j)^{t_0, t-1} = 0$. Considering Figure 1, there appears to be significant dependence of the current network on the network from $t-1$ and that in $[t-5, t-2]$. We therefore perform our analyses setting t_0 at $t-1$ and $t-5$.

We consider 7 measures of the proximity of vertices. Figure 4 provides illustration of the sub-graph configurations that constitute 4 of the proximity measures.

1. **Flow.** This measure generalizes preferential attachment [2] to the directed case. Preferential attachment is implemented as a measure of proximity by Liben-Nowell [2] as $\delta(i, j) = k_i k_j$, where k is the degree (i.e., the number of edges) of a given vertex. To take advantage of the direction of the ties, we conceive of a process whereby an attack from i to j is likely if i sends many attacks and/or j receives many attacks. In other words, we posit that frequent attackers are close to regular targets in the network. The measure of flow is $\delta(i, j) = k_i^o k_j^i$, where k^o and k^i are the out and in-degrees respectively.
2. **CTarget.** The number of common targets shared by two countries: $\delta(i, j) = \sum_h N_{ih} N_{jh}$.
3. **CAttacker.** The number of common attackers shared by two countries: $\delta(i, j) = \sum_h N_{hi} N_{hj}$.
4. **JacSim.** The Jaccard similarity between two countries is $\delta(i, j) = [\text{CTarget} + \text{CAttacker}] / [k_i + k_j]$, which normalizes the measure of common neighbors by the total number of neighbors of the vertices in the dyad.
5. **AASim.** The Adamic/Adar similarity adjusts the measure of common neighbors for the rarity of the neighbors to which the two countries tie. This measure is defined as $\delta(i, j) = \sum_h [\ln(k_h)]^{-1} (N_{ih} N_{jh} + N_{hi} N_{ji})$.
6. **SameCom.** Common community membership. We partition the countries into communities using the random walk modularity optimization algorithm "Walktrap" [26] and create an indicator, $\delta(i, j) = \mathbf{1}(c_i == c_j)$, of whether i and j are members of the same community.
7. **Distance.** Lastly, we include the minimum path length between i and j . We set $\delta(i, j)$ equal to the number of countries in the network plus one if there is no path from i to j .

In each model we include a count of the number of edges in the network to model the network's density. In addition to the proximity measures, and following Hanneke, Fu, and Xing [3], we include a memory term (**PrevAttack**) to capture persistence in the ties between the training network and the current network. The memory term at time t is specified as $\sum_{ij} N_{ij}^t N_{ij}^{t_0, t-1} + (1 - N_{ij}^t)(1 - N_{ij}^{t_0, t-1})$. We try each statistic computed on the networks over the interval $[t-1, t-1]$ and $[t-5, t-1]$. In the interest of comparing the performance of models on edges that did not occur in the past 10 years, we start our analysis at 1980 and go through the end of the dataset in 2002. Each model contains the edges term and some subset, including the empty set, of the memory and proximity terms, at both the one and five year intervals. The memory term and each of the proximity terms is (a) included computed on the one year training interval, (b) included computed on the five year training interval, and (c) excluded from the model. This leads to a total of $3^8 = 6,561$ models estimated at each t .

The forecast model for t is selected as the best performing model up to $t-1$. Performance is judged based on the predictive log score (i.e., the forecast log-likelihood) [27]. By using the predictive log score, in expectation, we use the model with the minimum Kullback-Leibler divergence from the actual model that generated the data [27]. We use θ^{t-1} to perform the forecast of N^t , which was estimated to fit N^{t-1} based on $N^{t_0-1, t-2}$. Thus, it is a true forecast in that the TERGM used to predict N^t has only been trained on the series of networks up to $t-1$.

The forecasting algorithm we employ is summarized as:

1. Estimate each of the 6,561 forecasting models for each time point from 1980 up to the previous year. Denote the structural measures in model M as Γ_M . Let θ_M^t be the parameters estimated on N^t using Γ_M .
2. Select as the forecasting model (M^*) for time t to be that model maximizing

$$\sum_{i=1981}^{t-1} \ln [\mathcal{P}(N^i | \theta_{M^*}^{i-1}, \Gamma_M)].$$

3. Forecast the next network from the distribution

$$\frac{\exp\{[\theta_{M^*}^{t-1}]' \Gamma_{M^*}(N, N^{t_0, t-1})\}}{C(\theta_{M^*}^{t-1})}.$$

4. Draw many forecast networks from the distribution in item 3 and compute the mean edge value in order to estimate the probability of any particular edge.

We apply this algorithm to all 23 years of the transnational terrorist network under consideration.

Figure 4 Sub-graph configurations underlying predictive measures. (a) Flow, (b) Geodesic Distance, (c) Common Attacker, (d) Common Target.

Results

There are a number of questions to be asked of our results. First, what is the overall performance of our approach to forecasting transnational terrorism? Second, does our approach offer leverage in predicting innovations in transnational terrorist relationships? Third, can

we see any patterns in the proximity and memory features that predict transnational terrorism? We address these questions in turn.

We begin by considering the overall performance of our forecasting method. The overall predictive performance is evaluated using the area under the receiver operating

characteristic curve (AUC) [28]. The receiver operating characteristic (ROC) curve gives the relationship between the false positive and true positive rates in predicting the value of a dichotomous outcome (e.g., the presence or absence of an edge in a network). A perfect classifier has an AUC of 1 and the closer to 1 an AUC is, the better the model is predicting. In Figure 5, we contrast the AUCs for the one-year-ahead forecasts of the best-predicting specification and two specifications that include only memory terms of 1 and 5 years respectively. Comparing to the “just memory” models allows us to identify the contribution of adding the indirect network proximity terms to the forecasting model. The proximity model with the highest log score up to $t - 1$ performs much better than the memory models. The AUC is approximately 0.95 on average, which compares to 0.83 and 0.72 for the five and one year memory models respectively. The implication is that future edges in the transnational terrorism network can be forecast based on network proximity in the recent past.

Above, we make the claim that network proximity based forecasting will allow us to leverage the considerable transitivity in transnational terrorism and forecast edge innovations: to predict a terrorist edge from one state to another where no such edge existed within a given window of time. Here, we evaluate the performance of our method on edges that did not occur in the recent past. A straightforward way to evaluate the predictive performance of our forecasting method is to consider the difference between the probability of edge formation assigned to those dyads that do form edges versus the probability of edge formation among those that do not form edges. Figure 6 shows forecasted probabilities assigned to edges that experience attacks and those that do not experience

Figure 5 One-year-ahead forecast AUC for various models. To compute the forecast AUC, we use the parameters estimated by modeling N^{t-1} based on $N^{t_0-1,t-2}$ to forecast N^t based on $N^{t_0,t-1}$. The “Best” model is the one that has had the highest average predictive log score up to t . “Memory 1” only has a memory effect from the previous year and an edges term. The “Memory 5” model only has a memory effect from the previous 5 years and an edges effect.

Figure 6 Forecasted probabilities for edges that experience attacks and those that do not experience attacks. Forecasted probabilities for edges that experience attacks and those that do not experience attacks. $\bar{\pi}$ is the mean forecast probability in the respective year. The probability of an attack from one location to another in year t is computed with the TERGM estimates from the model that, in terms of predictive log score, most accurately predicted N^{t-1} from $N^{t_0-1,t-2}$. In the interest of clarity, the y-axis is depicted on the log scale. Box plots depict the distribution of $\bar{\pi}$ among those directed dyads that experienced an attack in t .

attacks (potential edges that are not realized). We consider innovations of edges that have not occurred in the previous 1, 5 and 10 years, respectively. The results are similar for each of the three innovation intervals. The model assigns consistently low probabilities (in the range of 0.002) to potential edges that result in non-attacks.

More importantly, the edges that do form are, on average, forecast to do so with one to two *orders of magnitude* higher probability than are non-edges. The strong performance of our forecasting model in predicting edge innovations is a major contribution to the policy problem of identifying and addressing threats from the myriad of potential sources of transnational terrorism.

Lastly, we consider the contributions of the individual proximity and memory measures to the fit of the model. Figure 7 displays the ratio of the mean one-year-ahead forecast areas under the ROC curve with and without the given measure. A value greater than unity indicates that the average forecast AUC is higher when the respective term is included in the model. The plots show results from one (red) and five (blue) year memory models and all results are shown over time. Those statistics that consistently produce ratios greater than one can be said to make consistent contributions to the predictive performance of the model. Consistently highly performing measures include **PrevAttack**, **Flow**, **CAttacker** and **AASim**. The **SameCom** and **Distance** measures contribute substantial predictive performance in some years, but their effects are more volatile. For all of the measures that consistently add to the predictive performance of the model, the measure is computed on the five year interval. The superior performance of the measures computed with five year memories reinforces the result from Figure 5 that the transnational terrorism network exhibits long memory.

Case test: The Saudi link to the U.S. in 2001

The terror attacks of September 11, 2001 are the most spectacular and terrible the world has yet seen. They were also unexpected by policy and intelligence analysts. Among the surprises of the 9/11 attacks was the fact that 15 of the 19 hijackers were Saudi; indeed, it was the first time a Saudi citizen had committed any sort of attack on U.S. soil in almost three decades. As a test case, we examine what information our forecasting model provides about that 9/11 link from Saudi Arabia to the U.S.

To begin the case analysis, consider the ranked list, reported in Table 1, of the ten most highly predicted

Table 1 Top ten predicted sources of terrorism against targets inside the U.S., 2001

Rank	Country	P(Attack)
1	Algeria	0.126
2	Pakistan	0.055
3	Iraq	0.044
4	Jordan	0.037
5	Cuba	0.037
6	Canada	0.029
7	Romania	0.024
8	Saudi Arabia	0.012
9	Egypt	0.011
10	Iran	0.011

The rank-ordered sources of transnational terrorist threat most highly predicted by our forecasting model for the year 2001.

sources of attacks on the U.S. in 2001. Saudi Arabia is the 8th highest predicted source of attack and the list generally suggests a high risk posed by Middle Eastern countries; with Algeria, Pakistan, Iraq, Jordan, Egypt, and Iran also making the top 10. Canada may seem like an odd country to make the top 10 threat list, but attacks on the US from Canada have occurred, the most recent attack by a Canadian citizen preceding 9/11 occurred in 1999. This list, we believe, would be a useful guide for intelligence analysts and law enforcement officials attempting to efficiently divide their counter-terrorism resources between a larger number of possible threats.

The comparatively high probability assigned to attacks by Saudi citizens bodes well for our model, but how does the prediction for 2001 compare to other years and which statistics in our model produce the prediction? Figure 8 shows the percentile rank of Saudi Arabian citizens among those from all other countries in the world in a given year, with respect to the predicted probability of an attack on the U.S. We can see from the Figure that, at the time of the 9/11 attacks, the probability of attack on U.S. soil by Saudi citizens was the highest it had ever been (and the second

Figure 7 Relative forecasting performance of models with and without respective effects. Depicted is the ratio of the average one-year-ahead forecast areas under the ROC curves with and without the particular term. Terms computed on the one (five) previous year(s) are in red (blue).

highest achieved during the timeframe of our dataset). Interestingly, we also see that, based on the connectivity of the network, the Saudi threat begins mounting seriously in the late 1980's; during this time, attackers from Saudi Arabia strike in Israel. The drop noticed in 1992 occurs because, for that year, Saudi Arabia has no outgoing edges during the five year training period and the U.S. and Saudi Arabia have no common attackers. We see the probability of an attack jump back up again, elevate, and stay high following 1992. During this time, Iran becomes a common attacker for the U.S. and Saudi Arabia and a surge of international attacks by Saudi citizens begins: Saudis participate in attacks in Egypt, Jordan, the Philippines, Kenya and Tanzania, Pakistan, Uruguay, Albania, and Cambodia. Note, while not edges to the U.S. territorially, and therefore not recorded as edges to the U.S. in our networks, the 1998 Kenyan and Tanzanian attacks were orchestrated by al-Qaida and targeted U.S. embassies in those countries. This is also the time during which Osama bin Laden turned his resources against the U.S.; a response to his outrage over U.S. military forces being allowed in Saudi Arabia during the Gulf War.

We can see in Table 2, that flow, common attacker, and distance are producing the prediction for 2001, since these statistics comprised the model that best predicted the network in 2000. Flow and distance are both intuitive. Flow captures the fact that simultaneously Saudi Arabia is a (increasingly) large sender of attacks and the U.S. is (increasingly) a large receiver of attacks. Distance simply indicates that the geodesic distance between the U.S. and Saudi Arabia is large. The common attacker tie, as recorded in the ITERATE data, is a bit less intuitive, because the states are tied by Canada. This is not as odd as it might initially seem and actually reinforces our theoretical query into the transitivity of the network:

Table 2 Predictive model for 2001

δ	t_0	θ	δ %tile*
PrevAttack	1996	1.64	0
Flow	1996	0.027	99.99
CAttacker	1996	0.24	98.46
AASim	2000	0.5	0
SameComm	2000	0.441	0
Distance	1996	4.07	98.89

* δ %tile is the Percentile rank of $\delta(SA, US)$. The model used to predict the network in 2001. This is the model that best predicted the network in 2000 based on the θ estimated on the 1999 network. The θ in this table were estimated on the 2000 network. The percentile rank is based on a comparison with all of the other predictive scores (i.e., δ) of the other directed pairs of countries for which 2001 predictive scores are computed. The t_0 denote the beginning of the interval on which the predicting network is defined for that statistic, with all predictive network intervals ending in 2000.

Islamic extremists, and bin Laden in particular, were taking advantage of geographic proximity and soft borders by using Canada as a base for attacks on the U.S. [29]. In the meantime, a Canadian citizen was recruited by bin Laden, and participated in an attack on westerners (British and Irish) in Saudi Arabia [30]. This activity in Canada also explains the fact that Canada appears on our 2001 top-threat list for the U.S.

Conclusion

Our study of the transnational terrorist network makes at least three major contributions. First, we contribute to the literature on edge forecasting in complex networks. We integrate deterministic proximity-based forecasting into the probabilistic TERGM modeling of the evolution of a network based on its own topography. This approach can be applied to the edge prediction problem in many areas. This general method may prove useful for predicting the occurrence of edges in any variety of other networks such as international conflict (war), scientific collaboration and friendship.

We contributed to the state of knowledge about transnational terrorism by identifying that (1) there is an even mix of memory and innovation in the transnational terrorism network and (2) the network exhibits substantial transitivity. The transitivity we observe is likely due to economic, political, linguistic, and religious clustering that are subsumed under country labels. In other words, the way in which we define the vertices in the networks acts as a catch-all for the features likely to drive terrorism. Future research should address what specific features of countries predict terrorist link formation.

Lastly, but perhaps most importantly, we have advanced terrorism forecasting models in two critical directions. Our approach provides the necessary source-target specificity required to be useful for protecting the target *and*

addressing the source. Second, by leveraging the information on indirect ties in the network, our method is able to predict the occurrence of new terrorist edges. The ability to predict edge innovation is critical from a policy perspective because early warning is essential for the allocation of security resources that can, potentially, save lives.

Endnotes

^aIt is worth noting that there is a deep literature on psychological traits that make an individual more likely to become a terrorist and on the psychological group dynamics that affect a group's cohesion, but we do not consider this literature extensively here because its focus relates to individual behaviors rather than the amount of violence states suffer at the hands of terrorists. See [14,31] for reviews of this literature.

^bThese data are freely available at <http://www.correlatesofwar.org/>.

^cThey tend to at least lack realistic policy objectives, though they may have stated objectives such as the imposition of global Islamic law or communism.

^dNote that the computation of the likelihood requires the omission of the first K networks. Hanneke, Fu, and Xing [3] point out that one could also specify a separate probability model for these networks, but it is unclear what benefit would come of such an exercise, since this convenience model would, by construction, be misspecified.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

BAD developed the models in this article and drafted the manuscript. SJ developed the models and drafted the manuscript. Both authors read and approved the final manuscript.

Acknowledgements

This work was supported, in part, by a research grant from the College of Social and Behavioral Sciences, University of Massachusetts Amherst and by a grant from the University of North Carolina at Chapel Hill's University Research Council.

Author details

¹Department of Political Science, University of Massachusetts at Amherst, Amherst, Massachusetts 01003, USA. ²Department of Political Science, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, USA.

Received: 15 March 2012 Accepted: 30 January 2013

Published: 15 March 2013

References

1. EF Mickolus, T Sandler, JM Murdock, PA Flemming, International Terrorism: Attributes of Terrorist Events (ITERATE), 1968-2007. Vinyard Software. (Dunn Loring, 2008)
2. D Liben-Nowell, J Kleinberg, in *Proceedings of the Twelfth International Conference on Information and Knowledge Management*. The link prediction problem for social networks, (2003), pp. 556-559
3. S Hanneke, W Fu, EP Xing, Discrete Temporal Models of Social Networks. *Electron. J. Stat.* **4**, 585-605 (2010)

4. Q Li, Does democracy promote or reduce transnational terrorist incidents? *J. Confl. Resolution.* **49**(2), 278-297 (2005)
5. MT Koch, SJ Cranmer, Testing the "Dick Cheney" hypothesis: do governments of the left attract more terrorism than governments of the right? *Confl. Manag. Peace Sci.* **24**(4), 311-326 (2007)
6. JK Young, L Dugan, Veto players and terror. *J. Peace Res.* **48**, 19-33 (2011)
7. AH Kydd, BF Walter, The strategies of terrorism. *Int. Secur.* **31**, 49-80 (2006)
8. Q Li, D Schaub, Economic globalization and transnational terrorism. A pooled time-series analysis. *J. Confl. Resolution.* **48**(2), 230-258 (2004)
9. W Enders, GF Parise, T Sandler, A time-series analysis of transnational terrorism: trends and cycles. *Defense Peace Econ.* **3**(4), 305-320 (1992)
10. W Enders, T Sandler, Is transnational terrorism becoming more threatening? A time-series investigation. *J. Confl. Resolution.* **44**(3), 307-332 (2000)
11. W Enders, T Sandler, Patterns of transnational terrorism, 1970-1999: alternative time-Series estimates. *Int. Stud. Q.* **46**(2), 145-165 (2002)
12. W Enders, T Sandler, Transnational terrorism 1968-2000: thresholds, persistence, and forecasts. *South Econ. J.* **71**(3), 467-482 (2005)
13. PT Brandt, T Sandler, What do transnational terrorists target? Has it changed? Are we safer? *J. Confl. Resolution.* **54**(2), 214-236 (2010)
14. M Crenshaw, *Explaining Terrorism: Causes, Processes, and Consequences* (Routledge, London, 2010)
15. M Crenshaw, The causes of terrorism. *Comp. Polit.* **13**(4), 379-399 (1981)
16. AB Krueger, Malečková, Education, poverty and terrorism: Is there a causal connection? *J. Econ. Perspect.* **17**(4), 119-144 (2003)
17. BS Anderson, CT Butts, KM Carley, The interaction of size and density with graph-level indices. *Soc. Netw.* **21**(3), 239-267 (1999)
18. CT Butts, Social networks: a methodological introduction. *Asian J. Soc. Psychol.* **11**, 13-41 (2008)
19. SJ Cranmer, BA Desmarais, Inferential network analysis with exponential random graph models. *Pol. Anal.* **19**(1), 66-86 (2011)
20. SJ Cranmer, BA Desmarais, EJ Menninga, Complex dependencies in the alliance network. *Confl. Manag. Peace Sci.* **29**(3), 279-313 (2012)
21. SJ Cranmer, BA Desmarais, JH Kirkland, Toward a network theory of alliance formation. *Int. Interact.* **38**(3), 295-324 (2012)
22. NA Bapat, State bargaining with transnational terrorist groups. *Int. Stud. Q.* **50**(2), 215-232 (2006)
23. S Wasserman, P Pattison, Logit models for social networks: I. an introduction to Markov graphs and p. *Psychometrika.* **61**(3), 401-425 (1996)
24. J Park, M Newman, Statistical mechanics of networks. *Phys. Rev. E.* **70**(6), 66117-66130 (2004)
25. G Robins, P Pattison, Random graph models for temporal processes in social networks. *J. Math. Sociol.* **25**, 5-41 (2001)
26. P Pons, M Latapy, in *Computer and Information Sciences - ISCIS 2005, Volume 3733 of Lecture Notes in Computer Science*. Computing communities in large networks using random walks, (2005), pp. 284-293
27. J Geweke, G Amisano, Optimal prediction pools. *J. Econometrics.* **164**(1), 130-141 (2011)
28. ZC Qin, in *Proceedings of 2005 International Conference on Machine Learning and Cybernetics, Volume 5*. ROC analysis for predictions made by probabilistic classifiers, (2005), pp. 3119-3124
29. BBC News, Bin Laden 'using Canada as base' (2000). <http://news.bbc.co.uk/2/hi/americas/793178.stm>
30. L Gordon, Briton Admits Blasts. (The Guardian 4 February, 2001). <http://www.guardian.co.uk/world/2001/feb/05/saudiarabia>
31. M Crenshaw, in *Terrorism: Roots, Impact, Responses*, ed. by L Howard. How Terrorists think: psychological contributions to understanding terrorism, (1992), pp. 71-80

doi:10.1186/2190-8532-2-8

Cite this article as: Desmarais and Cranmer: Forecasting the locational dynamics of transnational terrorism: a network analytic approach. *Security Informatics* 2013 **2**:8.