



OPEN

Learning future terrorist targets through temporal meta-graphs

Gian Maria Campedelli^{1✉}, Mihovil Bartulovic² & Kathleen M. Carley²

In the last 20 years, terrorism has led to hundreds of thousands of deaths and massive economic, political, and humanitarian crises in several regions of the world. Using real-world data on attacks occurred in Afghanistan and Iraq from 2001 to 2018, we propose the use of temporal meta-graphs and deep learning to forecast future terrorist targets. Focusing on three event dimensions, i.e., employed weapons, deployed tactics and chosen targets, meta-graphs map the connections among temporally close attacks, capturing their operational similarities and dependencies. From these temporal meta-graphs, we derive 2-day-based time series that measure the centrality of each feature within each dimension over time. Formulating the problem in the context of the strategic behavior of terrorist actors, these multivariate temporal sequences are then utilized to learn what target types are at the highest risk of being chosen. The paper makes two contributions. First, it demonstrates that engineering the feature space via temporal meta-graphs produces richer knowledge than shallow time-series that only rely on frequency of feature occurrences. Second, the performed experiments reveal that bi-directional LSTM networks achieve superior forecasting performance compared to other algorithms, calling for future research aiming at fully discovering the potential of artificial intelligence to counter terrorist violence.

After peaking in 2014, terrorism activity worldwide has been on the decline in the last five years, as a consequence of several major defeats suffered by the Islamic State and Boko Haram, two of the world's most prominent jihadist organizations. Nonetheless, terrorism remains a persistent threat to the population of many areas of the world. Despite the decline in attacks and fatalities, in 2018 alone terrorist attacks worldwide led to 15,952 deaths¹.

The multifaceted nature of terrorism, characterized by a myriad of ideologies, motives, actors, and objectives, poses a challenge to governments, institutions, and policy-makers around the world. Terrorism, in fact, undermines states' stability, peace, and cooperation between countries, in addition to economic development and basic human rights. Given its salience and relevance, the United Nations includes the prevention of terrorism (along with violence and crime) as a target of the sixteenth Sustainable Development Goal, which specifically frames the promotion of peaceful and inclusive societies.

Scholars have called for the development of a dedicated scientific field focusing on the computational study of conflicts, civil wars, and terrorism^{2,3}. However, to date, attempts to exploit artificial intelligence for such purposes have been few and scattered. Whilst terrorism remains characterized by high levels of uncertainty and unpredictability⁴, trans-disciplinary research can help in providing data-driven solutions aimed at countering this phenomenon, exploiting the promising juncture of richer data, powerful computational models, and solid theories of terrorist behavior.

In light of this, the present study aims at bridging artificial intelligence and terrorism research by proposing a new computational framework based on meta-graphs, time-series, and forecasting algorithms commonly used in the fields of machine and deep learning. Retrieving event data from the Global Terrorism Database, we focus on all the attacks that occurred in Afghanistan and Iraq from 2001 to 2018 and construct 2-day-based meta-graphs representing the operational connections emerging from three event dimensions: utilized weapons, deployed tactics, and chosen targets. Once meta-graphs are created, we derive time-series mapping the centrality of each feature in each dimension. The generated time series are then utilized to learn the existing recurring patterns between operational features to forecast the next most likely central - and therefore popular - targets.

A baseline approach assuming no changes in terrorist dynamics over time and five deep learning models (i.e., Feed-Forward Neural Networks, Long Short-Term Memory Networks, Convolutional Neural Networks, Bidirectional Long Short-Term Memory Networks, and Convolutional Long Short-Term Memory Networks) are assessed in terms of forecasting performance. The outcomes are compared in by means of Mean Squared Error and two metrics that we introduce for this case study: Element-wise and Set-wise Accuracy. Furthermore, our graph-based feature engineering framework is compared against models that exploit shallow time-series simply

¹Department of Sociology and Social Research, University of Trento, 38122 Trento, Italy. ²School of Computer Science, Carnegie Mellon University, Pittsburgh 15123, USA. ✉email: gianmaria.campedelli@unitn.it

reporting the aggregate count of each tactic, weapon, and target in each 2 day-based time unit. The comparison aims at demonstrating that incorporating operational inter-dependencies through network metrics provides more information than merely considering event characteristics as independent from one another.

The statistical results signal that time-series gathered from temporal meta-graphs are better suited than shallow time-series for forecasting the next most central targets. Furthermore, Bidirectional Long Short-Term Memory networks achieve higher results compared to other modeling alternatives in both datasets. Forecasting outcomes are promising and stimulate future research designed to exploit the strength of computational sciences and artificial intelligence to study terrorist events and behaviors. Our work and presented outcomes pave the path for further collaboration among different disciplines to combine the practical necessity to forecast and predict as well as the need to theoretically and etiologically understand how terrorist groups act to strategically maximize their payoffs.

Background

Related work. The study of terrorist targets holds a prominent role in the literature: beyond its theoretical relevance, shining a light on individuals or entities at high risk of being hit can indeed help in designing prevention policies and allocating resources to protect such targets^{5,6}.

Studies investigating the characteristics and dynamics behind terrorist target selection have mainly employed traditional time-series methods, relying on yearly- or monthly-based observations^{7–10}, mostly framing research in the spirit of inference, rather than forecasting or prediction. These works highlighted the high-level patterns occurring globally, signaling how, over the decades, terrorist actors have substantially changed their operational and strategic behaviors. Nonetheless, these analytical approaches have limited ability to provide actionable knowledge for practically solving the counter-terrorism problem of resource allocation and attack prevention, given their meso or macro temporal focus.

More recently, an increasing use of computational approaches favored by higher availability of data fostered the diffusion of works that focused on temporal micro scales. Scholars have applied computational models investigating attack sequences, analyzing the spatio-temporal concentration of terrorist events and testing novel algorithmic solutions aimed at predicting future activity, with a particular emphasis on hotspots or violent eruptions. Among the tested algorithmic solutions are the use of point process modeling^{11–13}, network-based approaches^{14,15}, Hidden Markov models¹⁶, near-repeat analysis¹⁷, and early-warning statistical solutions using partial attack sequences¹⁸. In this computationally-intensive strand of research, the attention on terrorist targets, however, has been overlooked.

Overall, most literature has modeled terrorist attacks treating all events without discriminating them by their substantial features. Yet, this simplification greatly underestimates the multi-layered complexity of terrorist dynamics. Besides being patterned in their temporal characterization, terrorist attacks may follow patterns also in their essential operational nature^{19,20}. Ignoring this information and assuming all attacks are uni-dimensional fail to consider the hidden connections between temporally close events and the recurring operational similarities of distinct campaigns or strategies.

Machine and deep learning algorithms can help to overcome the limitations of the extant research in terms of paucity of attention to the fine-grained temporal analysis of terrorist targets, answering the call for scientific initiatives that should develop stronger connections between methodology and theory, rather than merely privileging one of the two²¹. The power, flexibility, ability to detect and handle non-linearity of these algorithmic architectures represent promising advantages that the field of terrorism research should explore. In the last years, few studies have attempted to exploit the strengths of artificial intelligence in this domain. Among these, Liu et al.²² presented a novel recurrent model with spatial and temporal components, and used data on terrorist attacks as one of two distinct experiments to evaluate the method's performance. However, the authors do not address how data have been processed before the proper modeling part, nor sufficiently clarify the implications of their forecasts, largely affecting the theoretical value of the experiment for terrorism research purposes. Ding et al.²³, instead, used data on terrorist attacks from 1970 to 2015 to forecast event locations in 2016, comparing the ability of three different machine learning classifiers in solving the task. While the authors interestingly combine several data sources trying to connect the methodological aspect with theory, the yearly scale of their predictions limits the usefulness of the results from a practical point of view, in line with issues already described in the literature employing more traditional statistical approaches.

More recently, Jain et al.²⁴ presented the results of a study aimed at highlighting the promises of Convolutional Neural Networks in predicting long-term terrorism activity. However, the authors do not employ existing real-world data, but instead evaluate their approach using artificially generated data. Furthermore, the models only employ univariate signals, excluding potential correlated signals that may impact forecasts.

In light of the paucity of works in this domain, this work on the one hand proposes a computational framework aimed at bridging the two disciplines of terrorism research and artificial intelligence for forecasting most likely targets and fostering the use of machine and deep learning for social good, using real-world data at a fine-grained temporal resolution. On the other hand, we seek to contribute to the theoretical study of terrorism by framing the problem in the context of the strategic theories of terrorist behavior.

Theoretical framework. Different theories have been proposed to describe and explain terrorist actions. These can be mainly divided into three perspectives: (1) psychological, (2) organizational and (3) strategic. Psychological theories of terrorism aim at explaining the individual causes leading to join terrorist actions and are mostly concerned with considerations covering motivations, individual drivers, and personal traits. Organizational theories, in turn, focus on the internal structure and the formal symbolism of each group as a way to read their behavior. Finally, strategic theories—which are derived from rationalist philosophy—address terrorist

Country/ dataset	Original N of attacks	Filtered N of attacks	Daily average	Daily St. Dev.	Daily Min.	Daily Max.	N of targets	N of weapons	N of tactics
Afghanistan	14,371	12,106	1.841	2.722	0	68	18	5	9
Iraq	25,886	22,764	3.462	4.670	0	107	20	5	8

Table 1. Descriptive statistics of the Afghanistan and Iraq datasets reporting the total number of attacks that occurred between 2001 and 2018 and the total number of targets, weapons and tactics represented in the considered time frame.

groups' decision-making and originate in the study of conflicts. Within this latter field, Schelling²⁵ posited that the parties engaging in a conflict are adaptive strategic agents: they hence try to find the most suitable ways to win, ruling out the opponent, as in a game or a contest. This straightforward consideration has been widely adopted by terrorism researchers who have formalized terrorism as an instrumental type of activity carried out to achieve a given set of long and short-run objectives²⁶. As noted by McCormick²⁷, terrorist groups are organizations that aim at maximizing their expected political returns or minimizing the expected costs related to a set of objectives. Notably, besides this adaptive and adversarial characteristic, the strategic frame assumes that terrorist groups act with a collective rationality^{28,29}: a terrorist group can be thought of as a unique actor, existing a unitary entity *per se*, in spite of its distinct internal components. Although this assumption simplifies reality, as terrorist groups can be structured in very different ways and these organizational features may impact decision-making processes, when considering historical events and their multidimensional characteristics, the assumption of collective rationality originated from Schelling in his studies on conflict adaptivity holds and actually helps in interpreting the life-cycle and behavior of a group.

Many constraints severely limit the strategic decision-making of a group (i.e., limited manpower³⁰), and such constraints have an impact on the type of attacks (as the ultimate and visible step of a decision-making process) that a terrorist group will plot. The strategic theoretical approach helps in unfolding some of the visible dynamics that data can reveal, including behavioral variations in combinations of tactics, weapons, and targets²⁷.

Recently, empirical research has corroborated the strategical perspective, showcasing for instance that terrorist violence follows specific patterns. To exemplify, terrorism is often characterized by self-excitability and self-propagation^{12,13,31}. The occurrence of an attack increases the probability of subsequent attacks in the same area within a limited time window, similar to what happens with earthquakes and their aftershocks, as a way to rationally maximize the inflicted damages of the attack waves. Nonetheless, empirical evaluation of whether the non-random nature of attacks can be extended also to events' operational characteristics is lacking.

The intuition behind this work builds on these theoretical propositions and seek to further scrutinize their ability to shed light on terrorism: we hypothesize that, given the complex adaptive and strategic decision-making processes in terrorist violence, we can exploit the temporal multi-dimensionality of the hidden operational connections among temporally close events to learn what type of targets will be at highest risk of being hit in the immediate future.

Data

The analyses in this work rely on data drawn from the Global Terrorism Database (GTD), maintained by the START research center at the University of Maryland³². The GTD is the world's most comprehensive and detailed open-access dataset on terrorist events and START releases an updated version of the dataset every year. The dataset includes now data on more than 200,000 real-world events. To be included in the dataset, an event has to meet specific criteria³³. These criteria are divided into two different levels.

There are primarily three first-level criteria that all have to be verified. These are related to (1) the intentionality and the violence (or immediate threat of violence) of the incident and (2) the sub-national nature of terrorist actors. There are also three second-level criteria, but the condition is that at least two of them are respected. Second level criteria relate to (1) the specific political, economic, religious, or social goal of each act, (2) the evidence of an intention to coerce, intimidate or convey messages to larger audiences than the immediate victims, (3) the context of action which has to be outside of legitimate warfare activities. Finally, although an event respects these two levels and is included in the dataset, an additional filtering mechanism (variable *doubter*) is introduced to control for conflicting information or acts that may not be of exclusive terrorist nature. Each event is associated with dozens of variables, mapping geographic and temporal information, event characteristics, consequences in terms of fatalities and economic damages, and attack perpetrators.

Besides considering the information on the country where an attack has occurred and the day in which the attack was plotted, this work specifically considers three core dimensions describing each event, namely (1) tactics, (2) weapons, and (3) targets. Tactics, weapons and targets that have been rarely chosen or employed in terrorist attacks that occurred in Afghanistan and Iraq have been excluded by the analysis. For a detailed explanation of the rationale of this filtering, see the Supplementary Material. Descriptive statistics on terrorist attacks in both countries are reported in Table 1.

Tactics. In the GTD, each attack can be characterized by up to three different tactics. Specifically, a single attack may be plotted using a mix of different tactics, and this information generally pinpoints a certain amount of logistical complexity. In the period under consideration, attacks in Afghanistan and Iraq have deployed using "Bombing/Explosion", "Hijacking", "Armed Assault", "Facility/Infrastructure Attack", "Assassination", "Hostage Taking (Kidnapping)", "Hostage Taking (Barricade Incident)", and "Unknown".

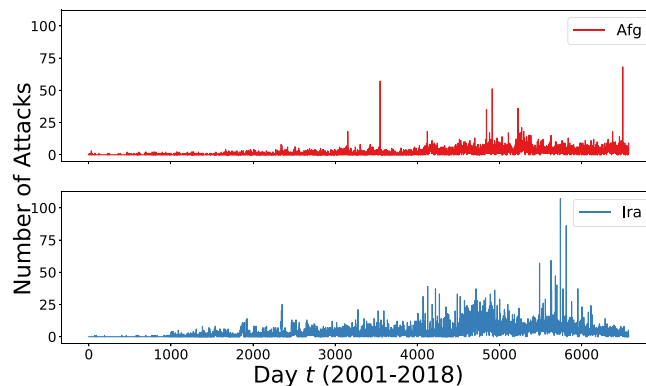


Figure 1. Time-series of terrorist attacks at the day level in Afghanistan (top) and Iraq (bottom). The figure has been created using Matplotlib version 3.1.3. Url: <https://matplotlib.org/3.1.3/contents.html>.

Weapons. For every event, the GTD records up to four different weapons. The higher the number of weapons utilized in a single attack, the higher the probability that the actor possesses a high amount of resources. In the Iraq and Afghanistan datasets, the represented weapon types are “Firearms”, “Incendiary”, “Explosives”, “Melee”, and “Unknown”.

Targets. Finally, each event can be associated with up to three different target categories. In the two datasets, the represented target types are the following: “Private Citizens and Property”, “Government (Diplomatic)”, “Business”, “Police”, “Government (General)”, “NGO”, “Journalists and Media”, “Violent Political Party”, “Religious Figures/Institutions”, “Transportation”, “Unknown”, “Terrorists/Non-State Militia”, “Utilities”, “Military”, “Telecommunication”, “Educational Institution”, “Tourists”, “Other”, “Food or Water Supply”, “Airports & Aircraft”.

The analytical experiments are performed using all the data regarding terrorist attacks that occurred in Afghanistan and Iraq from January 1st 2001 to December 31st 2018. The time series reporting the daily number of attacks in the countries under consideration are visualized in Fig. 1.

Temporal meta-graphs and graph-derived time-series

The main technical contribution of this work regards the representation of terrorist events through temporal meta graphs. The literature has shown that terrorist attacks do not occur at random. The associated core intuition is that, besides temporal clustering, there exist operational recurring patterns that can be learned to infer future terrorist actions. To capture the interconnections between events and their characteristics, framing the problem as a traditional time-series one is not sufficient.

In light of this, we introduce a new framework that exploits the advantages of graph-derived time series. First, per each time unit weighted graphs representing the meta-connections existing within the three data dimensions under consideration (tactics, weapons, and targets) are generated. Once this step is completed, we calculate, for each dimension and for each time unit, the normalized degree centrality of all the features. Normalized degree centrality maps the popularity of a certain weapon, tactic, or target in a given 2-day temporal window, by encapsulating it in a 1-dimensional space of complex information that emerged from a clustered series of attacks. By employing graph-derived time series we couple two layers of interdependence among events: the temporal and the operational one. Centrality not only portrays a certain target popularity: it may also aid in understanding the topological structure of a given set of attacks from the operational point of view, facilitating wide and distributed public safety strategies.

We start our data processing procedure by introducing a dataset $\mathcal{D}_{A \times z}$ that contains $|A|$ terrorist attacks and $|z|$ variables associated with each attack, exactly corresponding to the original format of the GTD. At this point, we filter out separately all the attacks that occurred in Afghanistan and Iraq in the time frame under consideration and we obtain two separate datasets: $\mathcal{D}_{t \times z}^{\text{AFG}}$ and $\mathcal{D}_{t \times z}^{\text{IRA}}$. The two new datasets are composed of $|t|$ observations (in this particular case one observation represents 1 day) and $|z|$ features. Here, the value corresponding to each feature is simply the number of times that feature was present in attacks plotted in that time unit, i.e. a single day. By doing so, we transitioned from an event-based dataset $\mathcal{D}_{A \times z}$ to a time-based one. At this point, $\mathcal{D}_{t \times z}^{\text{AFG}}$ and $\mathcal{D}_{t \times z}^{\text{IRA}}$ can be subset into $m \leq |z|$ theoretical dimensions. As anticipated, the dimensions here used are $m = 3$: the set of tactics features (X), the set of weapon features (W), and the set of target features (Y). Using the general C superscript indicating the country of reference, the subsetting leads to $D^C = \{\mathcal{D}_X, \mathcal{D}_W, \mathcal{D}_Y\}$.

Further, for each $\{\mathcal{D}_X, \mathcal{D}_W, \mathcal{D}_Y\}$ we create U temporal slices, such that $u = 2t$. In other words, for each dimension, we collapse the data describing the attacks in time units made of 2 days each by summing the count of each feature in the same 2 days. The reason behind the creation of 2-day-based time units is two-fold. On the one hand, relying on single day-based time series raises the risk of having overly sparse series, with very small graphs that would carry little to no relational information. On the other hand, in the real world resource allocation problems require time to be addressed, and having a forecasting system that operates day by day would produce knowledge that would be hard to transform into concrete and meaningful decisions. Thus, for this application,

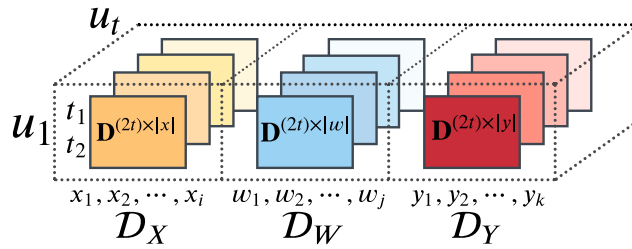


Figure 2. Graphic visualization sample of the 4d tensor resulting from the data processing. The figure has been created using Lucidchart. Url: <https://www.lucidchart.com>.

a 2-day architecture represents a good compromise. It reduces the sparsity guaranteeing that each time input is sufficiently rich in information and it provides predictions for the next 2 days, such that policymakers or intelligence decision-makers would strive less in changing resource allocation strategies too often. The temporal slicing leads to a 4-dimensional tensor D^C , as depicted in Fig. 2.

The tensor is composed of three tensors of rank 3, each representing one dimension i.e., tactics, weapons and targets. Each dimension is in turn composed of U matrices. For instance, for the weapons dimension W , each matrix is composed of $2t$ rows and $|W|$ columns, mapping all the weapons that have at least one occurrence over the entire history. To further exemplify, $D_W[1]$ is mapping the first temporal slice in the weapon dimension may be represented by:

$$t_1 \begin{bmatrix} \text{Explosives} & \text{Firearms} & \dots & \text{Incendiary} \\ 5 & 1 & \dots & 0 \\ t_2 \begin{bmatrix} 1 & 0 & \dots & 2 \end{bmatrix} \end{bmatrix}$$

At this point, to obtain the centrality of each feature in each matrix in the 4D-tensor, we first compute:

$$G_I[u] = D_I[u]^T D_I[u] \tag{1}$$

where $I \in \{X, W, Y\}$. By multiplying the transpose $D_I[u]^T$ by $D_I[u]$, we obtain a $|I| \times |I|$ square matrix whose entries represent the number of times every pair of features has been connected in the time unit under consideration. This matrix is interpreted as a meta-graph in which the connections between entities are not directly physical or tangible. Instead, the meta-graph represents a flexible abstract conceptualization aimed at linking together entities, such as particular tactics, that have been employed together in a specific time frame, within a limited set of attacks and that can be part of a logistically complex terrorist campaign.

The final step of this procedure is the computation of $\psi_{i, \text{Norm}}[u]$ which is the normalized centrality of each feature $i \in I$ for each 2-day temporal unit u . Given that $G_I[u]$ can be interpreted as a weighted square graph, then the weighted degree centrality of the feature i in $G_I[u]$ is computed as:

$$\psi_i[u] = \sum_{\substack{i \\ j \neq i}}^{|I|} G_{i,j}[u] \tag{2}$$

where $G_{i,j}[u]$ denotes the (i, j) entry of $G_I[u]$. Consequently, the normalized value is obtained through:

$$\psi_{i, \text{Norm}}[u] = \frac{\psi_i[u]}{\max_{i \in I} \psi_i[u]} \tag{3}$$

Normalizing the degree centrality allows to relatively compare the importance of each feature across time units that may present high variation in terrorist activity.

Finally, this leads to the creation of multivariate time-series in the form:

$$\Psi[U] = \{\psi_{i, \text{Norm}}[u]\}_{u=0}^U, i = 1, 2, \dots, F = |I| \tag{4}$$

where F is equal to the total number of features across all dimensions $|X| + |W| + |Y|$. For instance, in Afghanistan during the 2001–2018 time period attacks involved 5 weapons types, 9 tactics, and 18 targets making $F = 32$. Each $\psi_{i, \text{Norm}}[u]$ maps the relative importance of each feature in its respective dimension in a specific u (a sample visualization of this process is reported in Fig. 3). Instead of only using the simple counts of occurrences of each feature, the computed centrality value embeds how prevalent was a specific feature compared to the others, taking into account the meta-connections resulting from the complex logistic operations put in place by the terrorist actors active in Afghanistan and Iraq.

Reshaping the data structure from a sequence of graphs to a sequence of continuous values in the range $[0, 1]$ simplifies the problem while preserving the relevant relational information emerging from each meta-graph. Figure 4 depicts the Pearson's correlation among all features in all dimensions, while the temporal evolution of centrality values for features in the targets dimension \mathcal{D}_Y is visualized in Fig. 5. For both countries, plots

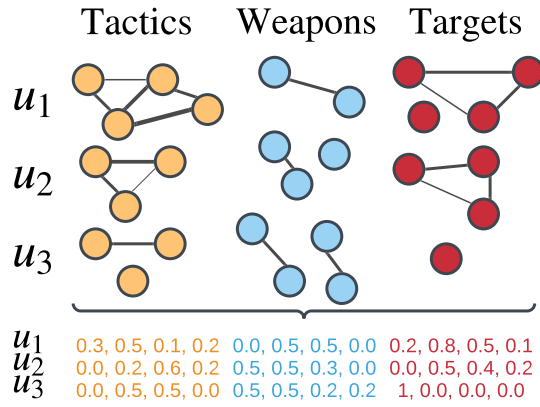


Figure 3. Sample visual depiction of the transformation of temporal meta-graphs in each dimension across u time-units in multivariate time-series capturing the normalized degree centrality of each feature in each dimension across the same u time units. The figure has been created using Lucidchart. <https://www.lucidchart.com>.

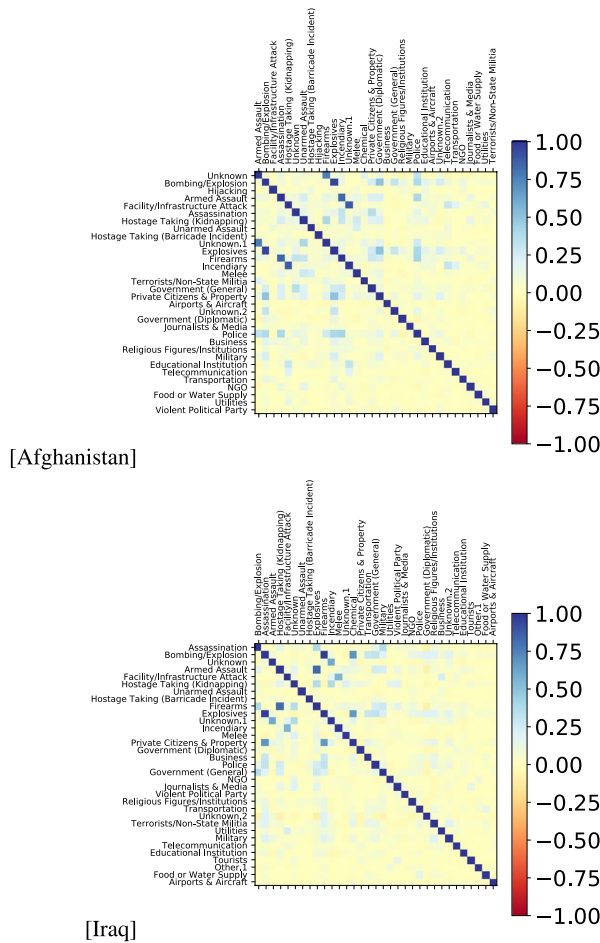
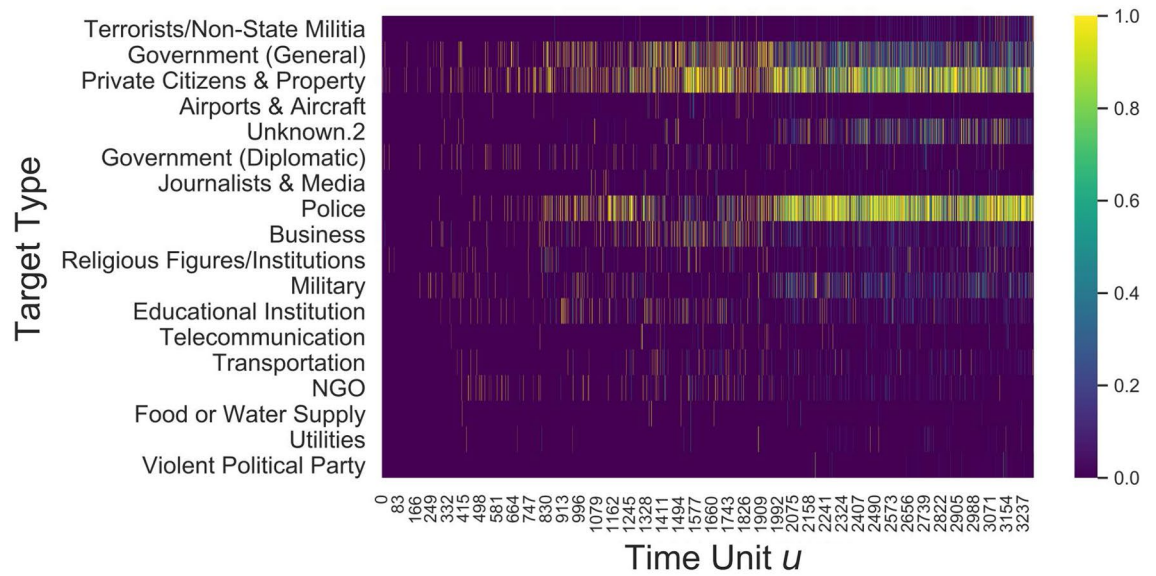
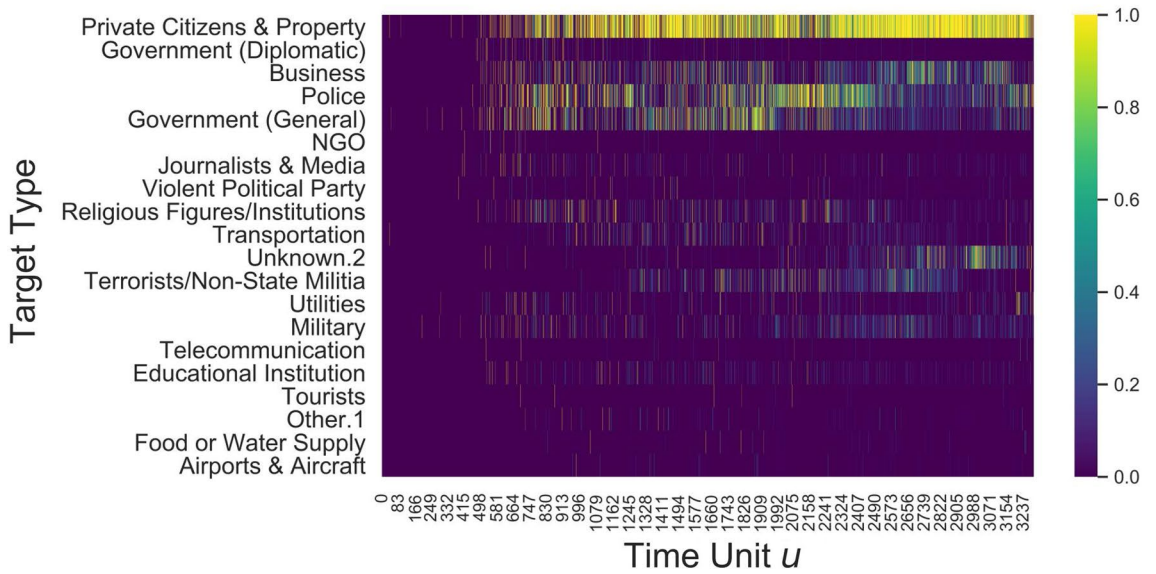


Figure 4. Pearson's correlation of centrality values among all features $F = |X| + |W| + |Y|$ for the entire time-span U . The figures have been created using Matplotlib version 3.1.3. <https://matplotlib.org/3.1.3/contents.html>.



(a) Afghanistan



(b) Iraq

Figure 5. Temporal evolution of centrality in the target dimensions (\mathcal{D}_Y) for the Afghanistan and Iraq cases. The figures have been created using Matplotlib version 3.1.3. <https://matplotlib.org/3.1.3/contents.html>.

describing the characteristics and distributions of tactics, weapons and targets are available in the Supplementary Materials (Figures S1–S4 for Afghanistan, and Figures S5–S8 for Iraq).

Methods

Algorithmic setup and models. We test the performance of six different types of models on our multivariate time-series forecasting task. Besides the computational contribution of the work in proposing the use of meta-graphs to model terrorist activity overtime, we are also interested in addressing relevant theoretical questions from a terrorism research standpoint. For this reason, the models are trained and fit using different input widths, i.e., a different number of time units. Algorithms performing better with shorter input widths would suggest that terrorist actors rapidly change their strategies and operations, indicating very low memory and thus making it inefficient to rely on large amounts of data. Contrarily, longer input widths would denote that algorithms have to be trained by taking into account relevant portions of recent events’ history given a relative degree of stability in attack operations.

In addition, to demonstrate the utility of our meta-graph learning framework, we train the same identical models (both in terms of architectures and input widths) on as many models using a shallow framework that

learns target centralities from a feature space engineered simply using aggregate counts of every single weapon and tactics in every time unit. We will then compare the overall results of the two approaches to empirically assess which one better captures the inherent dynamics occurring across terrorist events.

All the models have been trained using TensorFlow 2.5³⁴ and have been initialized with the same random seed. For both datasets, the 70% of the time units ($u = 2,300$; ~ 12.6 years) are used to train the model, the 20% ($u = 658$, ~ 3.6 years) for validation and the remaining 10% ($u = 329$, ~ 1.8 years) for testing. Additional modeling details besides the ones presented below are available in the Supplementary Materials in Section S2.

Baseline. The first tested model is a simple baseline scenario in which the forecasted centrality values $\hat{\psi}_{i,\text{Norm}}[u + 1]$ in the next time unit will be equal to the current ground-truth centrality values $\psi_{i,\text{Norm}}[u]$. This trivial case assumes that terrorist dynamics do not change over time, pointing in the direction of low tendency of terrorist actors to innovate their operational behaviors.

Feedforward neural network. The second model tested was a feed-forward dense neural network (FNN) that has two fully connected hidden layers. FNN is the most simple and straightforward class of neural networks and FNNs are not explicitly designed for applications that have time-series or sequence data³⁵. In fact, FNNs are not able to capture the temporal interdependencies existing across inputs and thus treat each input sequence as independent of one another. Their inclusion in our experiments is motivated by the interest in investigating whether our multivariate time series actually possess temporal autocorrelations or not through a comparison with other architectures that are designed to handle ordered inputs.

Long short-term memory network. Recurrent Neural Networks (RNN) are one of the algorithmic standards in research problems involving time series data or, more in general, sequences. In the present work, we employ an RNN with Long Short-Term Memory (LSTM) units³⁶. LSTM networks have been proposed to solve the well-known problem of vanishing (or exploding) gradients in traditional RNN. They include memory cells that can maintain information in memory for longer periods, thus allowing the algorithm to more efficiently learn long-term dependencies in the data. We employ LSTM networks with dropout to avoid over-fitting³⁷, using a two hidden layer structure.

Convolutional neural network. Convolutional neural networks (CNNs) are extensively used in computer vision problems, operating on images and videos for prediction and classification purposes³⁸. To handle such tasks, CNNs mostly work on 2-dimensional data such as images and videos. However, CNNs can also handle unidimensional signals such as time series³⁹. A 1D-CNN incorporates a convolutional hidden layer that applies and slides a filter over the sequence, which can be generally seen as a non-linear transformation of the input.

Bidirectional LSTM. Bidirectional LSTM (Bi-LSTM) are an extension of traditional LSTM models⁴⁰. They train two distinct LSTM layers: the first layer uses the sequence in the traditional forward order, while the other layer is trained on the sequence passed backward. This mechanism allows the network to preserve both information on the past and information on the future, thus fully exploiting the temporal dynamics of the time series under consideration.

CNN-LSTM. To complement the strengths of FNNs, CNNs, and LSTMs, Sainath et al.⁴¹ have proposed an architecture combining the three to solve speech recognition tasks. Since then, the approach has demonstrated many promises also in applications involving time-series data. We thus test the performance of a CNN-LSTM inspired by the work of Sainath et al.: the network has a first 1d convolutional layer, followed by a max pooling layer. The information is passed through a dense layer and an LSTM one, respectively. Finally, the network involves an additional dense layer that processes the output.

Performance evaluation. To evaluate the performance of the compared algorithms we first use a standard metric, i.e., mean squared error (MSE), computed as:

$$\text{MSE} = \frac{1}{U} \sum_{i=1}^u \left(\psi_{i,\text{Norm}}[u] - \hat{\psi}_{i,\text{Norm}}[u] \right)^2 \quad (5)$$

It is worth specifying that the evaluation is done on the test data and that in Eqs. (5)–(9), U denotes the length of the test data and u denotes one specific time slice of the test data. We keep this general notation to avoid confusion. We relied on MSE because we aimed at penalizing bigger errors. Furthermore, we then propose two alternative metrics. While centrality values are capturing correlational patterns across different attack dimensions, they are not directly interpretable in real-world terms. For this reason, we shift to a set perspective and we test the models on their ability to learn the most central targets, i.e. the most popular, frequent, and connected ones, in two different ways.

Element-wise accuracy. Element-wise accuracy Φ (EWA) is the simplest metric between the two. Given the sequence of test data, the sets $S[u]$ and $\hat{S}[u]$ represent the actual set of two most central targets and the predicted set of two most central targets at u . We define the element-wise accuracy $\phi[u]$ as:

$$\phi[u] := \begin{cases} 1 & \text{if } \hat{S}[u] \cap S[u] \neq \emptyset \\ 0 & \text{if } \hat{S}[u] \cap S[u] = \emptyset \end{cases} \quad (6)$$

Equation (6) means that if the sets have at least one element in common, then ϕ_{u_i} is equal to 1, while if the two sets are disjoint the value will be equal to zero. For the entire history of considered time units U , then, the overall EW accuracy Φ_U is computed as:

$$\Phi_U = \frac{1}{U} \sum_{i=1}^U \phi[u] \quad (7)$$

with Φ_U being the ratio between the sum of single unit binary accuracies $\phi[u]$ and the total number of time units U .

Set-wise accuracy. Set-wise accuracy Γ (SWA) is more challenging and further tests the ability of the deep learning models to identify and predict the correct set of targets $S[u]$. The cardinality of $S[u]$ is bounded in the range $0 < |S[u]| \leq 2$. Thus, for a given time unit u , single $\gamma[u]$ is defined as:

$$\gamma[u] := \frac{|\hat{S}[u] \cap S[u]|}{|S[u]|} \quad (8)$$

In our particular case, $\gamma[u]$ is equal to 1 if the two sets are perfectly identical (as in any set, it is worth noting, the order does not matter), is 0 when the two sets are disjoint and 0.5 when there is an intersection between S and \hat{S} . Finally, the overall metric Γ for the sequence U is given by:

$$\Gamma_U = \frac{1}{U} \sum_{u=1}^U \gamma[u] \quad (9)$$

Γ_U is then simply computed as the average value of all $\gamma[u]$ over the entire sequence of time units U . The metric aims at providing more comprehensive information to researchers and intelligence analysts potentially interested in designing logistically wider prevention strategies.

In the analyzed time-series, there is a considerable amount of cases in which time units did not record any attack in both datasets. In Afghanistan, no-attack units account for a 20.54% of the total time units u , while for Iraq this percentage is higher (21.2%). No-attack units are most prevalent in the first years of the considered time span, but we needed however to consider this eventuality in the performance evaluation of the proposed models.

Correctly forecasting no-attack units is crucial. In a counter-terrorism scenario, avoiding forecasting likely targets (and therefore attacks) in time units that do not experience any terrorist action reduces the costs of policy and intelligence strategy deployment. A wrong prediction in a no-attack unit is defined as a case in which: $S[u] = \emptyset \wedge \hat{S}[u] \neq \emptyset$, meaning that the set of predicted targets at risk of being hit $\hat{S}[u]$ has at least one target y , while the actual set of hit targets $S[u]$ is empty because no attacks were recorded.

Our forecasting framework involves a complex forecasting task. It is complex for two main interconnected reasons. First, we train our models to predict a relatively large set of time series characterized by sparsity and non-stationarity. Second, the number of time units is limited. Deep learning for time series prediction or classification has been deployed in many domains, including high-frequency finance applications^{42,43} or weather and climate prediction⁴⁴⁻⁴⁷. However, unlike these forecasting settings, open-access data on terrorist events cannot be measured using micro-scales such as minutes, seconds or hours, because data are collected at the daily level, hence leading to a significantly limited number of 2-day based time units u .

Therefore, the models have to quickly learn complex temporal and operational inter-dependencies without relying on massive amounts of data. Given these premises, it is highly unlikely that the forecasts can exactly predict centrality values $\psi_{i, \text{Norm}}[u]$ that are equal to 0 for each one of the features in the X , W and Y dimensions. Forecasts will contain noise, leading to very small centrality estimates for several features. When those very small $\psi_{i, \text{Norm}}[u]$ are produced as estimates, the element-wise and set-wise accuracy metrics will produce erroneous results. To accommodate the necessity to correctly forecast no-attack days, we have defined a rule that overcomes the likely presence of noise-driven values. Given \ddot{u} , that is a time unit u with zero attacks, we have thus set a threshold $\xi = 0.1$, such that:

$$\phi[\ddot{u}], \gamma[\ddot{u}] := \begin{cases} 1 & \hat{\psi}_{i, \text{Norm}} < \xi, \forall i \in I \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Equation (10) means that the element-wise accuracy $\phi[\ddot{u}]$ and the set-wise accuracy $\gamma[\ddot{u}]$ of the time unit \ddot{u} are both equal to 1 if all the predicted centrality values $\hat{\psi}_{i, \text{Norm}}[\ddot{u}]$ are below the threshold $\xi = 0.1$.

Conversely, if $\exists i$ such that $\hat{\psi}_{i, \text{Norm}}[\ddot{u}] > \xi$, then the evaluation of both the metrics is equal to 0. Given the distributional ranges of all the values (see Supplementary Materials), the ξ threshold set at 0.1 resulted in the right compromise between an excessive penalization of the algorithm performance and a shallow assessment of the actual forecasting capabilities of the proposed framework. Figure 6 shows the temporal trend of counts of non-zero centrality values in the target dimension for Afghanistan and Iraq.

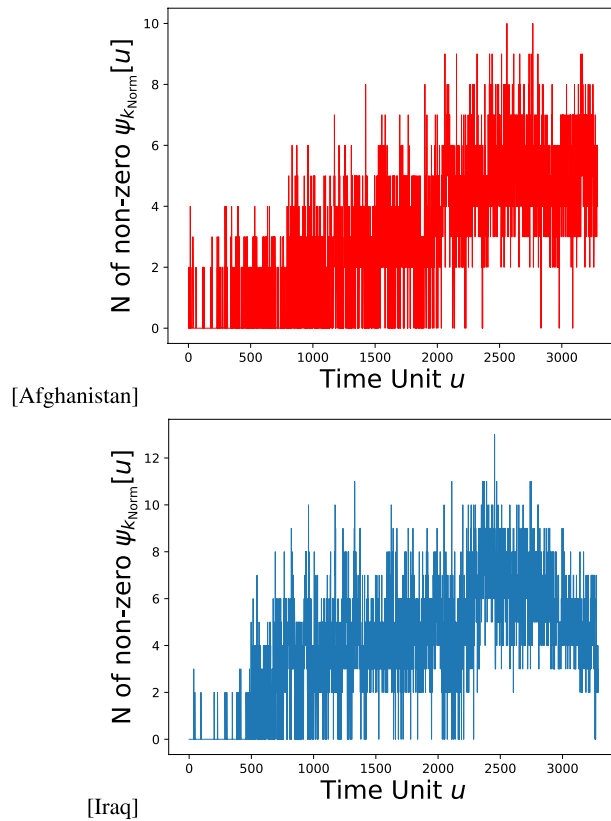


Figure 6. Temporal Trend of the count of non-zero centrality values $\psi_{k_{\text{Norm}}}[u]$ in the target dimension for Afghanistan and Iraq. The figures have been created using Matplotlib version 3.1.3. <https://matplotlib.org/3.1.3/contents.html>.

Results

The results reported in Table 2 overall showcase two relevant sets of results. First, the models clearly illustrate that structuring the problem of forecasting the next most central terrorist targets through a meta-graph learning framework outperforms an alternative scenario in which targets are forecasted simply using the shallow count of feature occurrences. This suggests that engineering the feature space by exploiting inter-dependencies among events in a network fashion allows capturing more information regarding the patterned nature of terrorist activity. In both the Afghanistan and Iraq cases, Γ and Φ are always higher when using graph-derived time-series. One exception is given by the MSE, which appeared to be lower in the shallow-learning case for the Iraq dataset and *ex-aequo* in the Afghanistan one. While a shallow learning approach works reasonably well in terms of regression, the meta-graph scenario works consistently better when forecasted targets need to be correctly ranked in terms of their ground-truth centrality. The comparison between meta-graph learning and shallow learning performances is provided in Fig. 7.

Second, the outcomes indicate that Bi-LSTM models always reach higher results in terms of Φ and Γ , regardless of the considered dataset and feature engineering approach. This convergence demonstrates that fitting a model that can rely on both backward and forward inputs enables the neural network to gather much richer contextual information, leading to superior forecasting performance.

Third, the outcomes also reveal that a baseline model with no learning component assuming no changes in terrorist strategies is highly inaccurate and, furthermore, the performance of Feedforward Dense Networks underscores that treating centrality sequences without taking into consideration the underlying temporal connections across different time units leads to sub-optimal predictions.

We were also interested in assessing the model results not only across algorithmic architectures but also across different input widths, contributing both to the computational and the theoretical study of terrorism. Highly similar results emerged, with a notable exception. A Bi-LSTM model trained using 5 time units u as input widths have reached the highest performance in both datasets for what concerns Φ and in the Iraq case also in terms of Γ . However, in the Afghanistan dataset, the same model architecture using 30 time-units u as input width obtained the highest Γ . The need for six times more data points for obtaining optimal forecasts denotes the higher stability of terrorism in Afghanistan, and a more microcycle-like structure of events in Iraq⁴⁸. This difference (although minimal) between the Afghanistan and Iraq cases might be explained by a higher heterogeneity of strategies and, ultimately, actors involved and opens new lines of inquiry to understand what factors are driving these dynamics (e.g., higher creativity, more resources, internal fighting).

Model	Input Width	META-GRAPH LEARNING						SHALLOW LEARNING					
		Afghanistan			Iraq			Afghanistan			Iraq		
		Γ	Φ	MSE	Γ	Φ	MSE	Γ	Φ	MSE	Γ	Φ	MSE
Baseline	1	0.1570	0.3140	0.7733	0.0701	0.1371	0.2371	0.1570	0.3140	0.7733	0.0701	0.1371	0.2371
FNN	1	0.6478	0.9146	0.0399	0.5594	0.9176	0.0275	0.6036	0.8658	0.0448	0.5183	0.8841	0.0527
LSTM	1	0.6753	0.9237	0.0395	0.5716	0.9237	0.0272	0.6265	0.8963	0.0410	0.5320	0.9024	0.0281
CNN	1	0.6204	0.8780	0.0414	0.5564	0.9176	0.0266	0.5807	0.8658	0.0485	0.5182	0.8963	0.0787
Bi-LSTM	1	0.6692	0.9238	0.0372	0.5442	0.9112	0.0261	0.6265	0.8871	0.0406	0.4969	0.8963	0.0268
CNN-LSTM	1	0.6494	0.9116	0.0393	0.5457	0.9146	0.0269	0.6189	0.8871	0.0382	0.5259	0.8993	0.0269
FNN	5	0.6204	0.8920	0.0415	0.5663	0.9228	0.0261	0.5432	0.8333	0.0569	0.5216	0.9012	0.0486
LSTM	5	0.6651	0.9136	0.0392	0.5771	0.9290	0.0254	0.5633	0.8395	0.0407	0.5370	0.9135	0.0248
CNN	5	0.6404	0.8889	0.0395	0.5524	0.9105	0.0265	0.5462	0.8364	0.0591	0.5000	0.8858	0.0541
Bi-LSTM	5	0.6836	0.9352	0.0366	0.5787	0.9290	0.0246	0.6404	0.9104	0.0370	0.5401	0.9104	0.0238
CNN-LSTM	5	0.6327	0.8920	0.0403	0.5370	0.8920	0.8920	0.5802	0.8488	0.0389	0.5201	0.8981	0.0249
FNN	15	0.6210	0.8885	0.0442	0.5478	0.8949	0.0286	0.6369	0.9044	0.0450	0.5207	0.8949	0.0567
LSTM	15	0.6768	0.9236	0.0376	0.5732	0.9267	0.0263	0.6257	0.8949	0.0370	0.5286	0.9012	0.0255
CNN	15	0.6322	0.8949	0.0441	0.5478	0.8980	0.0317	0.5859	0.8630	0.0503	0.5254	0.8949	0.0531
Bi-LSTM	15	0.6863	0.9331	0.0370	0.5573	0.9172	0.0240	0.6496	0.9171	0.0366	0.5000	0.8949	0.0248
CNN-LSTM	15	0.6433	0.8949	0.0440	0.5398	0.8917	0.0246	0.6114	0.8757	0.0384	0.4952	0.8789	0.0265
FNN	30	0.5335	0.8361	0.0520	0.5669	0.9163	0.0311	0.5785	0.8127	0.0485	0.5301	0.8996	0.0429
LSTM	30	0.6756	0.9298	0.0377	0.5735	0.9231	0.0263	0.6054	0.8795	0.0431	0.5301	0.9063	0.0250
CNN	30	0.6187	0.8829	0.0478	0.5619	0.9130	0.0331	0.5619	0.8361	0.0592	0.5317	0.8963	0.0451
Bi-LSTM	30	0.6890	0.9333	0.0369	0.5769	0.9230	0.0268	0.6572	0.9130	0.0376	0.5017	0.8862	0.0252
CNN-LSTM	30	0.5852	0.8729	0.0482	0.5585	0.9164	0.0250	0.5953	0.8662	0.0452	0.5183	0.8996	0.0301

Table 2. Performance in terms of set-wise accuracy (Γ), event-wise accuracy (Φ) and mean squared error (MSE) across chosen algorithms with varying input widths. For Γ and Φ , higher values mean better performance. The contrary holds for MSE. Values in bold indicate the best performance obtained in each country overall (i.e., taking into consideration also Shallow datasets).

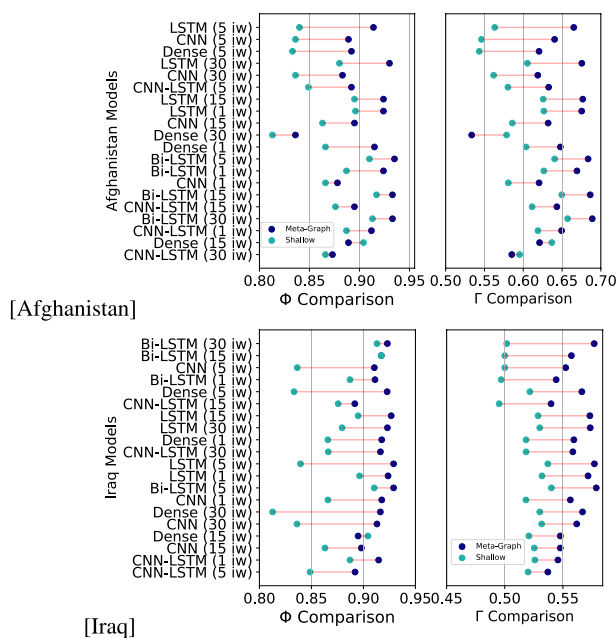


Figure 7. Comparison between Meta-Graph and Shallow Learning for Afghanistan and Iraq. Models are ordered based on descending difference in Γ (Set-wise Accuracy) performance. “IW” indicates “input width”. The figures have been created using Matplotlib version 3.1.3. <https://matplotlib.org/3.1.3/contents.html>.

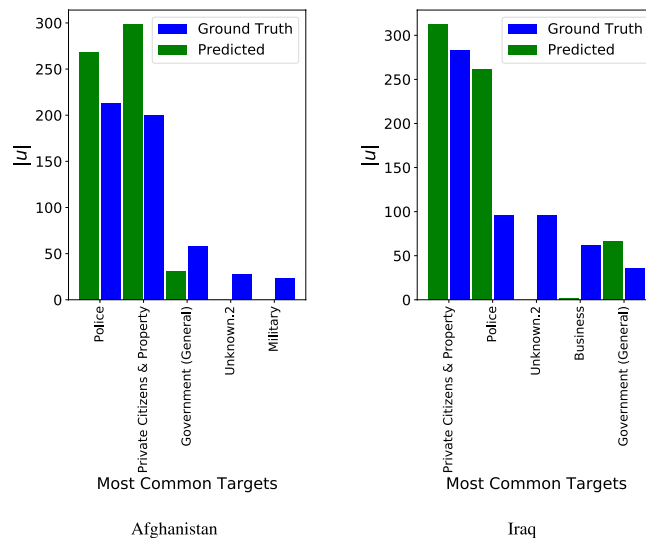


Figure 8. Bar chart comparing the number of time units $|u|$ a feature was empirically ranked among the two highest centrality values in the test set and number of times the same feature was forecasted among the two highest centrality values. The graph reports the result for the model with the highest Γ for Afghanistan and Iraq, respectively (Afghanistan: Bi-LSTM with 30 as input width; Iraq: Bi-LSTM with 5 as input width) and shows the five most common targets in each dataset. The figures have been created using Matplotlib version 3.1.3. <https://matplotlib.org/3.1.3/contents.html>.

To further inspect the quality and characteristics of the forecasting models, Fig. 8 conveys target-level information for both datasets. Specifically, the bar charts report a comparison between the empirical number of time units $|u|$ in which a certain target was among the highest two centrality values and the corresponding predicted number of time units in which the same target was forecasted to be in the top-two set, limiting the comparison to the outcomes of the model achieving the best Γ for each country. For Afghanistan and Iraq, Figures S9 and S10 in the Supplementary Materials also show the correlation between each vector $\Psi_{norm}[u]$ and $\hat{\Psi}_{norm}[u]$, representing respectively the empirical centrality values of each target feature at each time-stamp u and the corresponding predicted centrality values.

One common feature appears for both datasets: the Bi-LSTM models fail when dealing with time unit vectors having Unknown.2 as highly relevant target type (“Unknown.2” is used to distinguish unknown targets from unknown tactics, “Unknown”, and unknown weapons, “Unknown.1”, in the dataset). This is probably due to the peculiar mix of patterns associated with this specific target feature. Additionally, in Afghanistan the model poorly performs in predicting a central role for Military targets and in Iraq the Bi-LSTM model reaches unsatisfactory results when Business appears to be empirically relevant.

It is worth noting how the Afghanistan distribution of the top five empirical targets appears slightly more unbalanced compared to the Iraq one, possibly influencing the overall model results commented in Table 2, which are generally higher in Afghanistan. In both case studies, the forecasting models overestimate the prevalence of the most important empirical targets, calling for future efforts to engineer models that are more capable of capturing the nuances hidden under the different inter-connected dimensions of terrorist actions. Improving the model ability to handle these nuances will lead to a higher propensity of disentangling anomalous combinations of tactics, weapons and targets that may be hard to learn for the algorithms in their present form (possibly due to the low amount of data used compared to standard deep learning applications and the underlying non-stationarity of certain signals used in the datasets). This, in turn, would lead to forecasting distributions increasingly resembling empirical ones, which is the ultimate aim as results could be meaningfully used and deployed also in reference to rare events, intended as the outcomes of rare operational combinations.

In spite of these limits, the satisfactory performance recorded for the two best models (Afghanistan: $\Gamma = 0.6890$; Iraq: $\Gamma = 0.5787$) along with the visualized distribution of the most central targets corroborates the position of Clarke and Newman⁹ who repeatedly claimed the importance of protecting a restricted number of possible terrorist targets as a meaningful way to counter and prevent terrorist violence. While, in principle, terrorists have an almost infinite number of possible targets to choose from, their choice is limited by a number of constraints and motivations (both material and immaterial): this translates into the shrinking of the options’ spectrum, and in the recurring consistency of a very limited number of target types. In line with Clarke and Newman’s argument, the presented models underscore how suboptimal models can still provide effective intelligence knowledge given the patterned nature of terrorist actors, a mixed effect of the strategic character of their actions and the bounded portfolio of resources and options at their disposal.

Discussion and conclusions

Artificial intelligence and computational approaches are increasingly gaining momentum in the study of societal problems, including crime and terrorism. To contribute to this developing area of research, this paper proposed a novel computational framework designed to investigate terrorism dynamics and forecast future terrorist targets. Relying on real-world data gathered from the GTD, we presented a method for representing the complexity and interdependence of terrorist attacks that relied on the extraction of temporal meta-graphs from thousands of events that occurred from 2001 to 2018. We further tested the ability of six different modeling architectures to correctly predict the targets to be at the highest risk of being chosen in the next 2 days. The work presented the outcomes of multiple experiments performed focusing on terrorist activity in Afghanistan and Iraq. We first compared our approach using temporal meta-graphs against a shallow learning scenario in which forecasts are obtained using a feature space that only considers the count of occurrences across operational features of terrorist events. The comparison demonstrated that using temporal meta-graphs to construct time-series leads to superior forecasting performance, suggesting that embedding event inter-dependencies into temporal sequences offers richer context and information to capture terrorist complexity.

Additionally, results hint that Bidirectional LSTM (Bi-LSTM) models outperform the other architectures in both datasets, although differences arise across the two in terms of forecasting performance and optimal input width.

This work addressed an unexplored research problem in the growing literature that applies artificial intelligence for social impact, highlighting the potential of machine and deep learning solutions in forecasting terrorist strategies. These promising results call for future research endeavors that should address some of the limitations of the current work, including the generalizability of the results in other geographical contexts, the discrimination between actors operating in the same country, and the use of alternative contextual information such as data on military campaigns or civil conflicts. In terms of limitations, we specifically highlight that our approach does not take into account geo-spatial information. This decision is justified by the assumption that events occurring in the same country are part of a high-level strategic decision-making process, and that this high-level process is decoded through unified event dynamics. While this assumption is more easily justifiable in countries where one or very few terrorist actors are active or where, in spite of large pools of active actors most events are associated with one or few of them (e.g., Afghanistan), we recognize that it becomes more problematic for countries experiencing activity from a higher number of groups or organizations, as it is the case for Iraq. Patterns that exist at a national level may hide existing dynamics at the regional or provincial level. Future work should thus consider a localized geographical dimension to provide forecasts that have a deeper practical value, as they would produce heterogeneous probabilistic risk scales for different territories, helping intelligence agencies in setting up more tailored counter-terrorism strategies.

Data availability

The data and code used to conduct the analyses here presented are made available at the following GitHub repository: <https://github.com/gcampede/terrorism-metagraphs>.

Received: 18 December 2020; Accepted: 1 April 2021

Published online: 20 April 2021

References

- Institute for Economics and Peace. *Global Terrorism Index 2019: Measuring the Impact of Terrorism* (Tech. Rep, Sidney, 2019).
- Guo, W., Gleditsch, K. & Wilson, A. Retool AI to forecast and limit wars. *Nature* **562**, 331–333. <https://doi.org/10.1038/d41586-018-07026-4> (2018).
- McKendrick, K. *Artificial Intelligence Prediction and Counterterrorism* (Tech. Rep, Chatham House, 2019).
- Schiermeier, Q. Attempts to predict terrorist attacks hit limits. *Nat. News* **517**, 419. <https://doi.org/10.1038/517419a> (2015).
- Clarke, R. V. G. & Newman, G. R. *Outsmarting the Terrorists* (Greenwood Publishing Group, 2006).
- Bier, V., Oliveros, S. & Samuelson, L. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *J. Public Econ. Theory* **9**, 563–587. <https://doi.org/10.1111/j.1467-9779.2007.00320.x> (2007).
- Enders, W. & Sandler, T. Is transnational terrorism becoming more threatening?: A time-series investigation. *J. Confl. Resolut.* **44**, 307–332. <https://doi.org/10.1177/0022002700044003002> (2000).
- Enders, W. & Sandler, T. Patterns of transnational terrorism, 1970–1999: Alternative time-series estimates. *Int. Stud. Q.* **46**, 145–165. <https://doi.org/10.1111/1468-2478.00227> (2002).
- Asal, V. H. *et al.* The softest of targets: A study on terrorist target selection. *J. Appl. Secur. Res.* **4**, 258–278. <https://doi.org/10.1080/19361610902929990> (2009).
- Santifort, C., Sandler, T. & Brandt, P. T. Terrorist attack and target diversity: Change-points and their drivers. *J. Peace Res.* **50**, 75–90. <https://doi.org/10.1177/0022343312445651> (2013).
- Zammit-Mangion, A., Dewar, M., Kadiramanathan, V. & Sanguinetti, G. Point process modelling of the Afghan War Diary. *Proc. Natl. Acad. Sci.* **109**, 12414–12419. <https://doi.org/10.1073/pnas.1203177109> (2012).
- Tench, S., Fry, H. & Gill, P. Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army. *Eur. J. Appl. Math.* **27**, 377–402 (2016).
- Clark, N. J. & Dixon, P. M. Modeling and estimation for self-exciting spatio-temporal models of terrorist activity. *Ann. Appl. Stat.* **12**, 633–653. <https://doi.org/10.1214/17-AOAS1112> (2018).
- Desmarais, B. A. & Cranmer, S. J. Forecasting the locational dynamics of transnational terrorism: A network analytic approach. *Secur. Inform.* <https://doi.org/10.1186/2190-8532-2-8> (2013).
- Campedelli, G. M., Cruickshank, I. & Carley, K. M. Complex networks for terrorist target prediction. In *Social, Cultural, and Behavioral Modeling* Vol. 10899 (eds Thomson, R. *et al.*) 348–353 (Springer, 2018). https://doi.org/10.1007/978-3-319-93372-6_38.
- Petroff, V. B., Bond, J. H., Bond, D. H. & Bond, D. H. Using hidden Markov models to predict terror before it hits (again). In *Handbook of Computational Approaches to Counterterrorism* (ed. Subrahmanian, V.) 163–180 (Springer, 2013). https://doi.org/10.1007/978-1-4614-5311-6_8.
- Chuang, Y.-L., Ben-Asher, N. & D'Orsogna, M. R. Local alliances and rivalries shape near-repeat terror activity of al-Qaeda, ISIS, and insurgents. *Proc. Natl. Acad. Sci.* **116**, 20898–20903. <https://doi.org/10.1073/pnas.1904418116> (2019).

18. Yang, Y., Pah, A. R. & Uzzi, B. Quantifying the future lethality of terror organizations. *Proc. Natl. Acad. Sci.* **116**, 21463–21468. <https://doi.org/10.1073/pnas.1901975116> (2019).
19. Campedelli, G. M., Cruickshank, I. M. & Carley, K. A complex networks approach to find latent clusters of terrorist groups. *Appl. Netw. Sci.* **4**, 59. <https://doi.org/10.1007/s41109-019-0184-6> (2019).
20. Polo, S. M. The quality of terrorist violence: Explaining the logic of terrorist target choice. *J. Peace Res.* **57**, 235–250. <https://doi.org/10.1177/0022343319829799> (2020).
21. Bakker, E. Forecasting Terrorism: The need for a more systematic approach. *J. Strat. Secur.* **5**, 69–84 (2012).
22. Liu, Q., Wu, S., Wang, L. & Tan, T. Predicting the next location: A recurrent model with spatial and temporal contexts. In: *Thirtieth AAAI Conference on Artificial Intelligence* (2016).
23. Ding, F., Ge, Q., Jiang, D., Fu, J. & Hao, M. Understanding the dynamics of terrorism events with multiple-discipline datasets and machine learning approach. *PLoS One* **12**, e0179057. <https://doi.org/10.1371/journal.pone.0179057> (2017).
24. Jain, A. K., Grumber, C., Gelhausen, P., Häring, I. & Stolz, A. A toy model study for long-term terror event time series prediction with CNN. *Eur. J. Secur. Res.* <https://doi.org/10.1007/s41125-019-00061-w> (2019).
25. Schelling, T. C. *The Strategy of Conflict* (Harvard University Press, 1980).
26. Corsi, J. R. Terrorism as a desperate game: Fear, bargaining, and communication in the terrorist event. *J. Confl. Resolut.* **25**, 47–85. <https://doi.org/10.1177/002200278102500103> (1981).
27. McCormick, G. H. Terrorist decision making. *Annu. Rev. Polit. Sci.* **6**, 473–507. <https://doi.org/10.1146/annurev.polisci.6.121901.085601> (2003).
28. Crenshaw, M. Theories of terrorism: Instrumental and organizational approaches. *J. Strat. Stud.* **10**, 13–31. <https://doi.org/10.1080/01402398708437313> (1987).
29. Sandler, T. & Lapan, H. E. The calculus of dissent: An analysis of terrorists' choice of targets. *Synthesis* **76**, 245–261. <https://doi.org/10.1007/BF00869591> (1988).
30. Drake, C. J. M. *Terrorists' Target Selection* (St. Martin's Press, 1998).
31. Porter, M. D. & White, G. Self-exciting hurdle models for terrorist activity. *Ann. Appl. Stat.* **6**, 106–124. <https://doi.org/10.1214/11-AOAS513> (2012).
32. LaFree, G. & Dugan, L. Introducing the global terrorism database. *Terror. Polit. Violence* **19**, 181–204. <https://doi.org/10.1080/09546550701246817> (2007).
33. START. GTD Codebook: Inclusion Criteria and Variables. Tech. Rep., University of Maryland (2017).
34. Abadi, M. *et al.* TensorFlow: Large-scale machine learning on heterogeneous systems (2015). Software available from <https://www.tensorflow.org>.
35. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016). <http://www.deeplearningbook.org>.
36. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735> (1997).
37. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
38. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444. <https://doi.org/10.1038/nature14539> (2015).
39. Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L. & Muller, P.-A. Deep learning for time series classification: A review. [arXiv: 1809.04356](https://arxiv.org/abs/1809.04356) [cs, stat]. <https://doi.org/10.1007/s10618-019-00619-1> (2019).
40. Schuster, M. & Paliwal, K. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **45**, 2673–2681. <https://doi.org/10.1109/78.650093> (1997).
41. Sainath, T. N., Vinyals, O., Senior, A. & Sak, H. Convolutional, Long Short-Term Memory, fully connected Deep Neural Networks. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4580–4584. <https://doi.org/10.1109/ICASSP.2015.7178838> (2015). ISSN: 2379-190X.
42. Bao, W., Yue, J. & Rao, Y. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLoS One* **12**, e0180944. <https://doi.org/10.1371/journal.pone.0180944> (2017).
43. Sezer, O. B., Gudelek, M. U. & Ozbayoglu, A. M. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Appl. Soft Comput.* **90**, 106181. <https://doi.org/10.1016/j.asoc.2020.106181> (2020).
44. Shi, X. *et al.* Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems* Vol. 28 (eds Cortes, C. *et al.*) 802–810 (Curran Associates Inc, New York, 2015).
45. Qing, X. & Niu, Y. Hourly day-ahead solar irradiance prediction using weather forecasts by LSTM. *Energy* **148**, 461–468. <https://doi.org/10.1016/j.energy.2018.01.177> (2018).
46. Alemany, S., Beltran, J., Perez, A. & Ganzfried, S. Predicting Hurricane trajectories using a recurrent neural network. *Proc. AAAI Conf. Artif. Intell.* **33**, 468–475. <https://doi.org/10.1609/aaai.v33i01.3301468> (2019).
47. Wang, D., Liu, B., Tan, P.-N. & Luo, L. OMuLeT: Online multi-lead time location prediction for hurricane trajectory forecasting. *Proc. AAAI Conf. Artif. Intell.* **34**, 963–970. <https://doi.org/10.1609/aaai.v34i01.5444> (2020).
48. Behlendorf, B., LaFree, G. & Legault, R. Microcycles of violence: Evidence from terrorist attacks by ETA and the FMLN. *J. Quant. Criminol.* **28**, 49–75. <https://doi.org/10.1007/s10940-011-9153-7> (2012).

Author contributions

G.M.C., M.B. and K.M.C. designed research and experiments; G.M.C. and M.B. performed research and experiments; G.M.C. and M.B. analyzed the data and results; G.M.C. and M.B. wrote the paper. All authors read, reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-021-87709-7>.

Correspondence and requests for materials should be addressed to G.M.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021