

RESEARCH

Open Access



Identification of subgroups of terror attacks with shared characteristics for the purpose of preventing mass-casualty attacks: a data-mining approach

Gonen Singer* and Maya Golan

Abstract

Security and intelligence agencies around the world invest considerable resources in preventing terrorist attacks, as these may cause strategic damage, national demoralization, infringement of sovereignty, and government instability. Recently, data-mining techniques have evolved to allow identification of patterns and associations in criminal data that were not apparent using traditional analysis. The aim of this paper is to illustrate how to use interpretable classification algorithms to identify subgroups (“patterns”) of terrorist incidents that share common characteristics and that result in mass fatalities. This approach can produce insights far beyond those of conventional macro-level studies that use hypothesis-testing and regression models. In addition to this methodological contribution, from a practical perspective, exploring the characteristics identified in the “patterns” can lead to prevention strategies, such as alteration of the physical or systemic environment. This is in line with situational crime prevention (SCP) theory. We apply our methodology to the Global Terrorism Database (GTD). We present three examples in which terror attacks that are described by a particular pattern (set of characteristics) resulted in a high probability of mass casualties, while attacks that differ in just one of these characteristics (i.e., month of attack, geographical area targeted, or type of attack) resulted in far fewer casualties. We propose exploration of the differentiating characteristic as a means of reducing the probability of mass-fatality terrorist incidents.

Keywords: Global Terrorism Database, Mass-casualty terror attack, Interpretable classification models

Introduction

In an effort to better understand the logic of terrorists’ actions and to develop more efficient prevention strategies, criminologists suggest different typologies of terror attack based on various dimensions. A widely explored variable is the number of victims. Coordinated bombings, suicide attacks, explosions in confined spaces, and attacks during rush hour are some of the components of terrorist attacks that are often considered as “new”, aimed at maximizing the number of injuries and fatalities

(Hoffman 2006; Simon and Benjamin 2000). The sarin gas attack in Tokyo’s subway system in 1995, the coordinated bombings on commuter systems in Madrid in 2004 and Mumbai in 2006, the suicide attacks on the London subway system in 2005, and the bomb attack on a train between Moscow and St. Petersburg in 2009 are all examples of this new terror type (Regens et al. 2015).

Large-scale terror attacks often have devastating consequences. The attacks of 9/11 had significant repercussions on the global economy, with the passenger aviation industry absorbing much of the shock (Drakos 2004). Mass terror also damages a country’s economy, as it reduces foreign direct investment and lowers the confidence of domestic investors (Shahbaz et al. 2013). Terrorism and

*Correspondence: singerg@afeka.ac.il
Department of Industrial Engineering & Management, Tel-Aviv Afeka
Academic College of Engineering, 69988 Tel-Aviv, Israel



threats to national security are also documented to have impacts on tourism (Araña and León 2008). Psychological effects, such as public fear and stress, are a further consequence. Additionally, terrorism can affect political tolerance (Peffley et al. 2014) and can cause infringement of sovereignty, potentially leading to military confrontations. These devastating effects motivate governments to develop tools to assess and prevent large-scale terrorist attacks (e.g., LaFree and Bersani 2014).

In an effort to understand mass-causality terror acts, scholars attempt to identify their unique characteristics. A prevalent argument is that religiously-guided acts produce more victims. Indeed, according to the Terrorism Knowledge Base (Piazza 2009), there are an average of 38.1 casualties per attack in religious acts, compared to 9 in leftist or nationalist acts. Mierau (2015) supports this argument by claiming that ideological factors are the main determinant of lethality. Asal and Rethemeyer (2008) find that groups who combine ethnonationalist with religious ideologies are the most lethal, followed by purely religious groups. Several explanations have been offered for the tendency of religious groups to engage in mass-casualty attacks. These include the dehumanization of the victims in the eyes of religious terrorists (Berman and Laitin 2006); the spiritual, rather than practical, goals of such terrorists (Enders and Sandler 2000); the large scope of the target (i.e., society rather than individuals); and the conceptualization of violence as a desired goal rather than as a tool to achieve something else (Hoffman and McCormick 2004).

Piazza (2009) argues that the traditional assumption that Islamic terrorists mainly engage in mass-causality attacks is not valid and provides an alternative argument: differences in the number of casualties are generated by differences in group organizational features and goal structures. Similarly, Heger et al. (2012) explore organizational features and demonstrate that functional differentiation, clear command and control structures, and accountability are associated with increased numbers of victims. Asal et al. (2015) show that technical expertise within a terrorist organization minimizes civilian casualties while increasing the ability to kill high-value targets. Nemeth (2014) uses the Global Terrorism Database (GTD) to explore the effect of competition among terror groups. He shows that nationalist and religious groups respond to competition with more terrorism, while the converse is the case for left-wing organizations. Another research direction focuses on the mechanical aspects of mass-causality attacks. For example, Arnold et al. (2004) investigate the outcome of an attack according to bombing type, while Parachini (2001) compares outcomes for attacks involving conventional and unconventional weapons. Nilsson (2018) demonstrates that the lethality of suicide bombings is

greatest when there are many hard targets. Aylwin et al. (2006) approach the problem from the perspective of the response services and identify the system processes that reduce critical mortality after an urban attack.

Most studies that attempt to uncover patterns of terror attacks use clustering analysis, spatial and temporal statistics, and geographical information systems (GIS). White et al. (2013) apply a temporal approach to the GTD, and publish a self-exciting model of the risk and volatility of terror events, as well as resilience to them. Morris and Slocum (2012) also use the GTD and apply latent class growth analysis and general mixture modeling to identify country-level patterns of terrorism. Other studies that examine or model geographical and/or temporal patterns include: LaFree et al. (2012), Behlendorf et al. (2012), Berrebi and Lakdawalla (2007), Webb and Cutter (2009), Siebeneck et al. (2009), Brown et al. (2004), Liu et al. (2001), Inyaem et al. (2010), Reed et al. (2013), and Mohler et al. (2015).

Several recent studies have used supervised machine-learning algorithms to detect patterns of crimes. Caines et al. (2018) classify posts in online hacking-related forums according to their intent. Kuang et al. (2017) use topic modeling to investigate the relationships between the full-array formal crime type classifications used by police and narrative texts associated with crime events. Fernando et al. (2018) use a random-forest algorithm to detect hate-speech messages in cyberspace based solely on metadata. The aim of this study is to demonstrate the utility of applying data-mining algorithms to identify subgroups of terrorist incidents that share common characteristics and that result in mass fatalities. These patterns produce insights that cannot necessarily be derived using conventional methods, such as hypothesis-testing or regression models, since the latter focus on identifying features that are correlated with the number of casualties across the whole sample. In contrast, our approach allows the identification of features that are correlated with lethality in particular *subgroups* of attack. Furthermore, contrary to previous studies (including those that use data mining), we do not explicitly assume dependencies on time and space. Rather, we use interpretable classification models to distinguish between sets of features (not necessarily related to time or space) that characterize low-versus multi-casualty terror attacks. These patterns may increase the likelihood of preventing terror attacks by extracting clear rules and actions for security forces to follow. Our approach is in line with situational crime prevention (SCP) theory (Clarke 1980, 1997, 2010), which states that one should reduce opportunities for crime, instead of attempting to understand or change the underlying dispositions and motives of criminals.

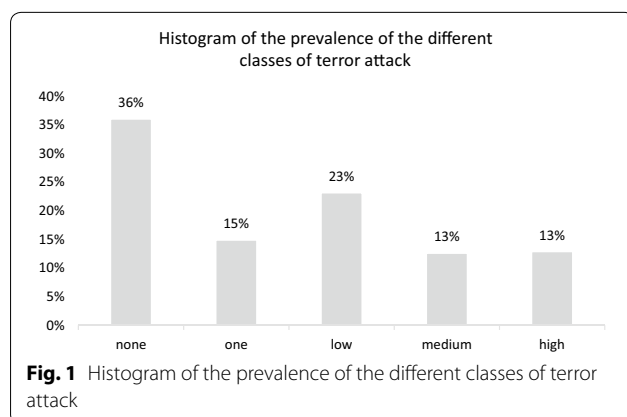
Data and methodology

Data

This study used interpretable classification models to identify patterns of terrorist attacks, according to known characteristics derived from historical data. A “pattern” refers to a set of values for particular features in a database. For this purpose, we used the Global Terrorism Database (GTD) (LaFree and Dugan 2007; LaFree 2010), which is an open-source database on terrorist attacks around the world from 1970 to 2016. The GTD is provided by the National Consortium for the Study of Terrorism and Response to Terrorism (START) at the University of Maryland. It contains data on more than 170,000 domestic and international terrorist incidents, including dozens of features on location, tactics, perpetrators, targets, and outcomes of the events. The full set of features, separated into categories, is presented in [Appendix A](#) (for further details, see the GTD 2016).

We limit our analysis to the most recent terror attacks: 45,000 cases between 2014 and 2016. Thus, we assume that the characteristics of terror attacks have changed over the years and that attacks before 2014 are less relevant to our analysis. The aim was to differentiate between low- and mass-casualty attacks, which is a supervised problem, i.e., build up a logic that assign each training example of terror attack to a different level of casualty attacks and then eventually uses that logic on unlabeled data. The identification of a mass-casualty event was achieved by creating a classifier field, which was based on four features: “nkill”—total number of confirmed fatalities from the incident, “nkillter”—number of perpetrator fatalities, “nwound”—number of non-fatal injuries, “nwoundter”—number of perpetrators injured. The classifier field was calculated as follows:

$$class = nkill - nkillter + nwound - nwoundter.$$

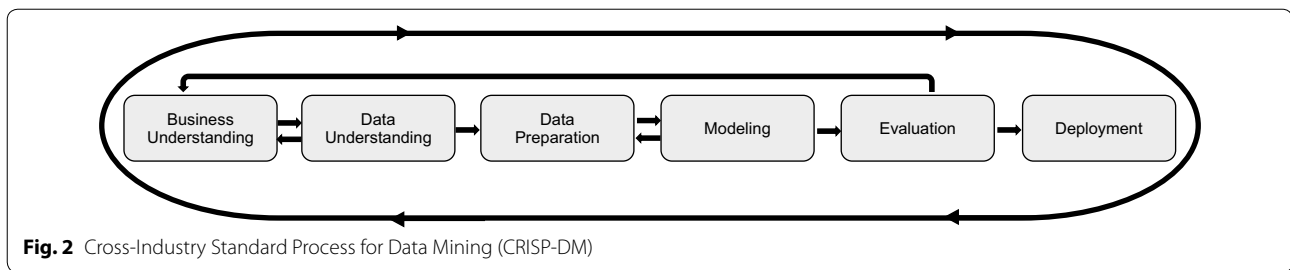


To simplify the analysis, we discretized the class field into five categories: “none”=0, “one”=1, “low”=2–5, “medium”=6–10, “high”=11 and above (see Fig. 1). The rationale for this discretization method was twofold: (1) to distinguish between incidents without casualties and other events; (2) to produce a histogram of the prevalence of the different classes such that the number of terror attacks in each category (except for “none”) was similar.

Methodology

The project was conducted in accordance with the Cross-Industry Standard Process for Data Mining (CRISP-DM), proposed by Chapman et al. (2000). The six stages of CRISP-DM are presented in Fig. 2. The first stage, business understanding, involves understanding the project objectives and requirements from a business perspective. The main business objective was to identify mass-casualty attacks based on their known characteristics. In this stage, we familiarized ourselves with the data source; in particular, we examined the different definitions of a terrorist attack used in the database. We also examined definitions of a mass-casualty attack that were employed in previous studies (Arnold et al. 2004; Parachini 2001; Aylwin et al. 2006). As a result of all these exploratory processes and taking into consideration the goal of the current research, we devised the discretization method described above to define the target field. In the second stage—data understanding—we delved further into the database, learning its features and identifying problems related to data quality.

In the data preparation stage, we applied some pre-processing actions, such as data cleaning, data reduction and data transformation. In data cleaning, tuples with irrelevant or exceptional data were dealt with as follows: either the entire tuple was deleted, or the unusable values were replaced with the most probable value. For example, tuples that did not state the number of casualties were considered irrelevant and thus deleted, while tuples with “exceptional” values for latitude and longitude (i.e., inconsistent with the country) were modified by replacing these values with average values for other incidents at the same place. In the data reduction step, in cases where two features were strongly inter-correlated (e.g., region code and region name), one of the features was removed. An example of data transformation was as follows: in the case where a categorical feature had a large number of categories (e.g., the feature “city”, which included 130 different cities, many of which appeared with low frequency), the low-frequency categories were grouped into a single category called “other”. This type of transformation was performed to prevent overfitting of the models.



Following data preparation, the dataset comprised 45,200 terror attacks and 26 features.

The aims of the modeling stage were to (a) classify multi-casualty terror attacks and (b) identify significant patterns leading to multi-casualty terror attacks. To these ends, we chose to use interpretable classification models, as opposed to black-box models, which have poor interpretability (e.g., neural networks and support vector machine; see Letham et al. 2015; Wang and Rudin 2015; Ribeiro et al. 2016). Interpretable models, such as decision trees and Bayesian models (Kotsiantis 2007; Singer and Golan 2019), may be used to find the most influential parameters that differentiate between mass- and low-casualty terror attacks. Furthermore, interpretable models result in a set of clear rules that can be acted upon without the need for technical or machine-learning knowledge. In the present study, these actions might be able to prevent large-scale terror attacks. Four interpretable classification models from the Weka data-mining software were selected for comparison: C4.5, Bayesian Network, PART and Naïve Bayes. We selected these algorithms for the following reasons. Bayes networks and decision trees are the most common machine-learning interpretable models for classifications problems. From decision trees, we chose the C4.5 decision tree (Singer and Golan 2019), in preference to other algorithms (such as ID3), since our dataset contains a mixture of different types of variable, in the sense that some variables span a wide range of values (e.g., city) while others are simple binary indicators (e.g., suicide attack). Secondly, we included the PART algorithm (Frank and Witten 1999), which uses the C4.5 decision tree and thus includes all its advantages. However, instead of identifying a set of rules in the first phase and then refining it by discarding rules in an optimization phase (as in the C4.5 algorithm), PART identifies one rule at a time by repeatedly generating partial decision trees, without the need for an optimization phase. PART is thus faster and more efficient, and in some cases can achieve better results. Another class of interpretable model that can be used to describe the relationship among features is the Bayesian Network (BN) model (Ben-Gal 2007). Several studies claim that BNs are competitive with

state-of-the-art classifiers such as the C4.5 decision tree (Friedman and Goldszmidt 1996; Janssens et al. 2006). Finally, Naïve Bayes is a specific type of Bayesian network classifier with strong assumptions of independence among features.

In the evaluation stage, the interpretable classification models were constructed on a training dataset and then evaluated on a testing set (using a fivefold cross-validation paradigm), after which the best-performing model was chosen (see “[Comparison of classification algorithms](#)” section). The cross-validation method is an extension of the known training–testing method. Both of these methods test the model’s ability to predict new data that were not used in constructing it. However, in order to reduce variability in the performance measures, bias in selecting the training and testing data, and unreliable estimates of future performance (Moore 2001), the cross-validation method performs multiple rounds using different partitions such that each instance in the database appears in both the training and the testing procedures. Thus, the data were divided into five subsets, where four-fifths were used for training the classification models and the remaining fifth was used as the test set. For a given interpretable model, the classification performance measures were calculated five times, each time using a different sub-sample as the test sample. This technique is commonly used with classification solutions (Maimon and Rokach 2005). The next section presents the comparison between the classification algorithms, followed by an illustration of how the best classifier may be used to derive practical insights.

Results and discussion

Comparison of classification algorithms

Firstly, we evaluated the performance of the classifiers in identifying the class category “high”, which was termed a “positive” event. This category is particularly important since the aim of this research was to identify patterns that characterize mass-casualty terror attacks. Our overall performance measure for the “high” class category was AUC, the area under the receiver operating characteristic

(ROC) curve. This curve was created by plotting the true positive rate (TPR), i.e., the proportion of positive events that were identified as such by the classifier, against the false positive rate (FPR), i.e., the proportion of negative events that were wrongly categorized as positive. Each point on the ROC curve represents the performance of the model at a different discrimination threshold of the classification algorithm (Green and Swets 1966).

In addition to the performance measure just described, which pertained to “positive” events only, we calculated the overall AUC and Cohen’s Kappa Statistic (KE) for all five classes. The AUC of the classifier model was calculated as a weighted average measure of all 5 AUC values, each one calculated from an ROC curve for one of the 5 categories, according to the proportion of all attacks that belong to this category among all attacks. KE is a metric that compares an observed accuracy (in this case, the discrimination ability of the classifier for each of the categories) with an expected accuracy (random chance). For example, an observed accuracy of 90% is much less impressive when the expected accuracy is 85% than when it is 40%. Table 1 presents results for the aforementioned measures for the different interpretable classification

Table 1 Performance measures for different interpretable classification algorithms

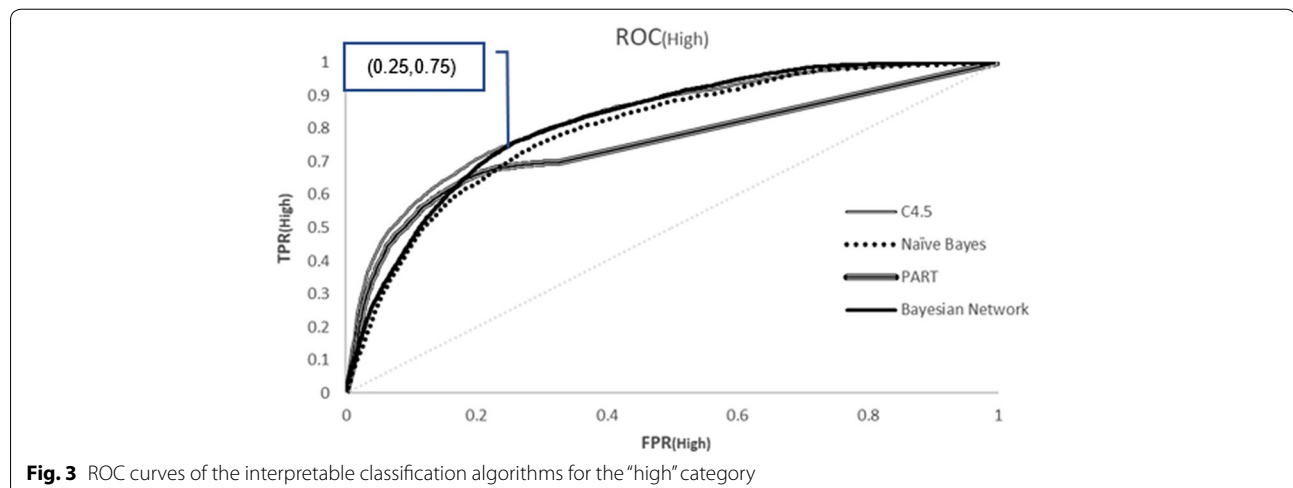
Interpretable classification algorithm	Performance measure for “positive” events	Overall performance measures	
	AUC_{High}	AUC	KE
C4.5	0.832	0.841	0.468
Naïve Bayes	0.797	0.781	0.343
Bayesian Network	0.818	0.800	0.375
PART	0.752	0.765	0.435

models. The numbers in bold signify the maximum value for a given performance measure.

From Table 1, it is clear that C4.5 outperforms the other classification algorithms for all performance measures. Figure 3 presents the ROC curves for the “high” category. It can be seen that below a certain false positive rate (0.25), the C4.5 algorithm achieves the best results, while for higher FPR values, the performances of C4.5 and Bayesian Network are similar. Given these results, we chose the C4.5 algorithm as the classifier to predict multi-casualty terror attacks.

Decision trees can be overfitting (Singer et al. 2019), mainly when: (i) the number of levels of the features is too high, (ii) there is large variance in the number of levels among different features, or (iii) there is large variation between the training and the testing dataset. We took the following precautions to reduce the likelihood of overfitted: (i) In the data preparation stage, we grouped low-frequency categories together, resulting in nominal features with fewer than tens of values each. (ii) In the modeling stage, we used the C4.5 decision tree, which overcomes the limitation of some decision trees that are overly sensitive to features with large numbers of values as well as large variance in the number of values. (iii) In the evaluation stage, we checked that the performance of the training dataset was close to the performance of the testing dataset.

The final stage of the data-mining process, known as deployment, involves organizing the knowledge gained and presenting it such that others can use it. After choosing the classifier with the best performance (C4.5), we used this model to derive practical insights that can help criminologists explore strategies for preventing multi-causality terror attacks. This approach is described in the following subsection.



Practical insights

A decision tree is a map consisting of “nodes” and “branches” that represents the possible outcomes of a series of related choices. Starting from the root node (containing all tuples in the dataset), the tree is built by branching the dataset into possible outcomes, with the interior nodes containing subsets of tuples. The branching follows a set of splitting rules that are based on classification features and their values. A leaf is a terminal node that specifies the value distribution of the target variable; thus, the expected value of the target is the value with the highest probability. The paths from root to leaf represent patterns. In this study, we derive patterns (sets of features and values) that define subgroups of terror attacks resulting in a given casualty level. The C4.5 model yielded 5,000 such patterns.

We present three patterns that differentiate between terror attacks with and without a “high” number of casualties. For each pattern, we were able to recommend an investigation process, the outcome of which should contribute to the prevention of mass-casualty terror

attacks. Table 2 summarizes these examples. The second column (“Lift”) shows the increase in the probability of a “high” casualty-level (given the specific pattern) relative to the inherent probability of a “high” event in the data (13%; see Fig. 1). Thus, the lift is a measure of the performance of the model.

Example 1: Terror attacks in Mosul of type “kidnapping”

The first pattern concerns attacks in the city of Mosul, Iraq that did not involve suicide, were directed against military or police, and included hostage-taking. These yielded a high number of casualties (86% of events classed as “high”). The prevalence of a “high” casualty-level in other types of attack (e.g., bombing, assassination, or armed assault) was only 4%. These results may reflect the common strategy of Daesh to kidnap soldiers and police officers and to later kill them. Accordingly, a possible direction for investigation that may reduce the number of casualties would be to learn the kidnapping tactics used by Daesh relative to other attack types, and propose actions to prevent or respond to hostage-taking, so as to avoid multiple casualties.

Table 2 Examples of multi-casualty terror attack patterns yielded by the C4.5 classifier

Pattern of multi-casualty terror attack	Lift $\left(\frac{\% \text{ in the pattern}}{\% \text{ in the data}}\right)$	% of multi-casualty events given a different value of a specific feature in the pattern	Recommended investigation that could contribute to the prevention of terror attacks
Suicide = “No” Target = Military or Police City = “Mosul” Attack Type = Hostage-taking (kidnapping) ↓ “High” number of casualties = 6 events out of 7 (86%)	$\left(\frac{86}{13}\right) = 6.6$	Suicide = “No” Target = Military or Police City = “Mosul” Attack Type = Bombing/Assassination/ Armed Assault ↓ “High” number of casualties = 6 events out of 146 (4%)	Learn the methodology of the act of hostage-taking against police and military in Mosul, and how it differs from other types of attack Propose actions to prevent or respond to hostage-taking, so as to avoid multiple casualties
Suicide = “No” Target = “Checkpoint” Attack Type = Armed Assault Nationality of Target = “Afghanistan” Province = “Helmand” Month of event < 5 ↓ “High” number of casualties = 22 events out of 26 (85%)	$\left(\frac{85}{13}\right) = 6.5$	Suicide = “No” Target = “Checkpoint” Attack Type = Armed Assault Nationality of Target = “Afghanistan” Province = “Helmand” Month of event = 5 ↓ “High” number of casualties = 3 events out of 17 (18%)	Try to engage in negotiations and resolve things nonviolently during May and/or use this as an opportunity to evacuate people so as to avoid multiple-causality terror attacks during the rest of the year.
Suicide = “No” Target = “village” Attack Type = Armed Assault Country = “Nigeria” Perpetrator Group Name = Fulani extremist Weapon type = Firearms Latitude > 10 Longitude > 7.4 ↓ “High” number of casualties = 5 events out of 5 (100%)	$\left(\frac{100}{13}\right) = 7.7$	Suicide = “No” Target = “village” Attack Type = Armed Assault Country = “Nigeria” Perpetrator Group Name = Fulani extremist Weapon type = Firearms Latitude > 10 and Longitude ≤ 7.4 or Latitude < 10 ↓ “High” number of casualties = 59 events out of 303 (19%)	Investigate the defense and response abilities of villages in Nigeria located at Kaduna (latitude > 10 and longitude > 7.4), relative to those of villages located in other areas Apply the approaches to tackling violent conflicts used at other locations to Kaduna—approaches that are rooted in the prevention of multiple-causality terror attacks

Example 2: Terror attacks at Helmand checkpoints in the month of May

The second pattern refers to armed assaults at Helmand checkpoints in Afghanistan, which are characterized by fewer casualties in May than in other months. Since it is unlikely that this is due to an irregularity in the activity of security forces in this month, a cultural explanation may shed light on this finding. Indeed, the religious festival of “Ramadan”, in which Muslims fast and are required to refrain from smoking, drinking, engaging in sexual relations and fighting (except for self-defense), takes place during the month of May. Studies show that crime rates decrease during Ramadan (e.g., Tavakoli 2012). Thus it could be worthwhile investigating the option of trying to negotiate or resolve things nonviolently during May, or alternatively, to use this as a time to evacuate people, for the purpose of avoiding multiple-causality terror attacks during the rest of the year.

Example 3: Terror attacks by the Fulani group against Nigerian villages in specific locations

The third pattern concerns the Fulani militant group, a terrorist group operating in Nigeria and parts of the Central African Republic. Tensions between the Fulani, the majority of whom are Muslim, and Nigerian farmers, the majority of whom are Christian, are largely driven by economic causes. Our data show that, between the years of 2014 and 2016, each of the 5 armed assaults on villages in a specific area of Kaduna (latitude > 10 and longitude > 7.4) resulted in a “high” number of casualties, while attacks with the same characteristics in other areas in Nigeria only had a 19% probability of a “high” casualty-level. In Plateau, for example, only 5 out of 29 attacks resulted in a “high” number of casualties, despite the fact that in most of these attacks, more than 20 houses were damaged. Investigation into differences in the security environment between Kaduna and other places in Nigeria leads to interesting insights. Firstly, the security environment in Kaduna is complex and challenging due to the presence of armed groups, high crime rates, and the risk of kidnapping. Secondly, in other places in Nigeria, the approach to tackling violent conflicts has been anchored in preventive mechanisms. For example, in Plateau, a Council of Elders has been established to mediate in latent conflicts between or among ethnic or religious groups, and work is being carried out to foster peaceful co-existence between religious and civil society.

This pattern illustrates the advantages of machine-learning algorithms in identifying insights related to features that have not previously been considered in the literature. Latitude and longitude are environmental variables that would not generally be considered as a component of the attack, but rather an indication of a

problematic place. Indeed, 1950 incidents in the database take place in Nigeria, most of them characterized by the attack type “Armed assault” (50%) and the weapon type “Firearms” (85%). In most of these attacks, farmers who were not armed were attacked by the Fulani group equipped with firearms. Thus, the interesting insight lies not in these well-studied features (e.g., weapon type), but in the finding that attacks in a specific area resulted in a “high” casualty-level.

Given that there is a strong correlation between the feature ‘country’ and the longitude and latitude, it would not be unreasonable to suggest manually dropping the pair longitude and latitude from the model (as sometimes happens in machine-learning projects; see Duda et al. 2012). By dropping these variables, the finding would be as follows: Terror attacks by the Fulani group in Nigeria of type “Armed assault” against villages using firearms result in a “high” number of casualties in 26% of cases, while the prevalence of a “high” casualty-level when using other weapons, for an attack with the same characteristics, is only 14%. In our estimation, this insight would be far less interesting.

Discussion

The aim of this study was to demonstrate how data-mining algorithms can provide insights that are not necessarily apparent using other methodologies, by identifying subgroups of terror attacks that share common characteristics. The results revealed that it is possible to discover explanatory variables that are influential for a particular subgroup. In contrast, when correlation or regression analysis is applied to a large sample, it is only possible to uncover a small number of features that are strongly correlated with the number of casualties. Thus, data mining appears to be a suitable approach with which to carry out research from an SCP perspective, since the analysis may yield situational results that cannot be explained by other criminology theories.

The C4.5 algorithm achieved the best performance. Three patterns were presented that lent themselves to practical insights and suggested future investigation directions. In these examples, terror attacks with a particular set of characteristics resulted in a high probability (between 83% and 100%) of high-level casualties. Meanwhile, terror attacks that differed in just one characteristic (i.e., the month of attack, the specific location, or the type of attack), resulted in much lower probabilities of high-level casualties (< 20%). It is interesting to mention that when regression analysis is applied to our dataset, the features suicide, target type and attack type are found to have the highest correlation with the number of casualties, as already shown in the literature (Arnold et al. 2004; Hoffman 2006; Simon and Benjamin 2000;

Hoffman and McCormick 2004). It can be seen that these same features were also common in the patterns used to derive practical insights from our interpretable classification algorithm (see Table 2), which is unsurprising, as the latter is aligned with conventional statistical analysis. However, the correlation values for some of the differentiating features (i.e., month, longitude and latitude) are substantially lower, confirming that they are not influential across the full sample. Furthermore, these variables have hardly been mentioned in the literature.

The main contribution of this study is that it demonstrates the ability to use analytical tools to investigate mass-casualty terrorist attacks. We argue that classification algorithms are highly suitable for this purpose, since they are not influenced by hidden human assumptions, past reasoning or common knowledge—all of which may be misleading with regard to predicting the next mass-casualty terror attack. The outcomes of these tools are patterns that describe reality without the biases of human judgment. The current methodology can be adapted to analyze criminal data in the context of a specific theory, rather than adopting a situational approach. The classifier we used was the identification of an attack as a mass-casualty event. One could also utilize different classifiers that represent other theoretical assumptions. For example, Clarke and Newman's (2006) four pillars of opportunity (i.e., target, weapon, tools or training, and facilitating condition) could be used to create four classifiers, each representing one pillar. The patterns that would emerge from such an analysis would be aligned with this specific theory.

There are limitations to our method. First, the classifier establishes machine-oriented rather than human-oriented rules. Thus, it can only be used as a support tool and not as an automatic one, suggesting future research directions that involve building human-oriented models. Second, challenges remain with regard to increasing the classification accuracy and converting the insights from the models into actionable recommendations. Enrichment of the database using additional information (e.g., religious holidays) may help to meet these challenges. Finally, the model needs to be updated dynamically, as terrorist modes of operation are constantly changing.

Further research could evaluate the performance of black-box models, such as neural networks and support vector machines, instead of interpretable classification models. If black-box models are shown to achieve superior performance, they would be better able to predict whether a terror attack will result in a high number of casualties, but they would not uncover the most influential parameters, nor would they provide descriptive information about such attacks. An additional research direction could be to use the proposed methodology

to answer different questions. These could include, for example, the identification of successful versus unsuccessful terror attacks according to indications other than the number of casualties, and identification of the conditions under which a certain weapon or strategy is employed by different terrorist groups.

Conclusions

This study aimed to demonstrate how data-mining algorithms can be used to identify subgroups of terrorist incidents that share common characteristics and that result in mass fatalities. The C4.5 algorithm was found to achieve the best performance when applied to the Global Terrorism Database. We presented three examples in which the output of the algorithm was used to obtain practical insights. In each example, there was one characteristic (i.e., month of attack, geographical area, or type of attack) that differentiated between attacks with and without a high casualty-level. From a theoretical perspective, we provide insights into the complex multi-dimensional nature of mass-causality terror acts. From a methodological perspective, we demonstrate the value of using a data-mining approach alongside classical statistical methods. Finally, from a practical perspective, we provide a tool for deriving complex, unexpected insights that can be used to develop recommendations for reducing the probability of a mass-casualty terrorist attack.

Abbreviations

GTD: Global Terrorism Database; GIS: geographical information systems; START : National Consortium for the Study of Terrorism and Response to Terrorism; CRISP-DM: cross-industry standard process for data mining; AUC: area under curve; ROC: receiver operating characteristic curve. TPR: true positive rate; FPR: false positive rate; KE: Cohen's kappa statistic.

Acknowledgements

Not applicable.

Authors' contributions

GS and MG conceived of the presented idea, developed the theory and performed the computations discussed the results and wrote paper. Both authors read and approved the final manuscript.

Authors' information

Gonen Singer, Ph. D., is a Senior Lecturer of Industrial Engineering and Management at AFEKA-Tel-Aviv Academic College of Engineering. He joined the Faculty at 2008, shortly after its establishment. He was appointed Head of the Department in the years 2009–2015. Dr. Singer has extensive expertise in machine learning techniques and stochastic optimal control and their application to real-world problems in different areas, such as retail, manufacturing and Criminology

Maya Golan, Ph.D., is a Senior Lecturer of Industrial Engineering and Management at AFEKA-Tel-Aviv Academic College of Engineering. She was appointed Head of the Department in the years 2016–2018. She received her Ph.D. in behavioral sciences and management from the Technion—Israel Institute of Technology in 2009. Current research interests include employee well-being, human-machine interfaces and human behavior via analytical techniques.

Funding

Not applicable.

Availability of data and materials

The Global Terrorism Database (GTD) dataset supporting the conclusions of this article is available in <https://www.kaggle.com/START-UMD/gtd>.

Competing interests

Not applicable.

Appendix A: Detailed explanation of the dataset variables by category

Category	Features	Explanation
GTD ID AND DATE	Date	Date at which the incident occurred
	Extended incident	Indication of whether the duration extended beyond 24 h
Incident information	Incident summary	A brief narrative summary of the incident
	Inclusion criteria	Which of the inclusion criteria are met, such as political, economic etc.
	Doubt terrorism proper?	Indication whether the incident is an act of terrorism
	Alternative designation	Only if there is doubt as to whether the incident is an act of terrorism, this feature identifies the most likely categorization of the incident other than terrorism
Incident location	Part of multiple incident	Indication if the incident is part of several other attacks
	Country	The country or location where the incident occurred
	Region	The region in which the incident occurred
	Province/Administrative Region/State	The name of the 1st order subnational administrative region in which the event occurs
	City	The name of the city, village, or town in which the incident occurred
	Vicinity	Indication if the incident occurred in the immediate vicinity of the city or in the city itself
	Latitude	The latitude of the city in which the event occurred
	Longitude	The longitude of the city in which the event occurred
Attack information	Geocoding specificity	The geospatial resolution of the latitude and longitude fields
	Attack type (1–3)	The general method of attack (kidnapping, bombing, assassination, etc.). Up to three attack types can be recorded for each incident
	Successful attack	Indication of the success of a terrorist strike (defined according to the tangible effects of the attack)
	Suicide attack	Indication if there is evidence that the perpetrator did not intend to escape from the attack alive

Category	Features	Explanation
Weapon information	Weapon type (1–4)	The general type of weapon used in the incident. Up to four weapon types are recorded for each incident
	Weapon sub-type (1–4)	A more specific value for most of the weapon types identified. For example, weapon type “chemical” can be sub-type poison or explosive. Up to four weapon sub-types are recorded for each incident
Target/Victim information	Target/Victim type (1–3)	The general type of target/victim (business, government, police etc.). Up to three target/victim types are recorded for each incident
	Target/Victim sub-type (1–3)	The more specific target category, providing the next level of designation for each target type (for example, target type “business” can be sub-type bank/hotel/farm etc.). Up to three target/victim sub-types are recorded for each incident
Perpetrator information	Nationality of target/victim (1–3)	The nationality of the target that was attacked. Up to three nationalities of target/victim are recorded for each incident
	Perpetrator group name (1–3)	The name of the group that carried out the attack. Up to three perpetrator groups are recorded for each incident
Perpetrator information	Perpetrator sub-group name (1–3)	This field contains any additional qualifiers or details about the name of the group that carried out the attack. Up to three perpetrator sub-groups are recorded for each incident
	Perpetrator group suspected/unconfirmed (1–3)	Indication whether or not the information reported by sources about the Perpetrator group name(s) is based on speculation or dubious claims of responsibility. Up to three perpetrator group suspected/unconfirmed indications are recorded for each incident
	Unaffiliated individual(s)	Indication whether or not the attack was carried out by an individual or several individuals not known to be affiliated with a group or organization
Perpetrator information	Number of perpetrators	The total number of terrorists participating in the incident
	Number of perpetrators captured	The number of perpetrators taken into custody
	Claim of responsibility?	Indication whether a group or person(s) claimed responsibility for the attack
	Mode for claim of responsibility	One of 10 modes used by claimants to claim responsibility; might be useful to verify authenticity and track trends in behavior (letter, call, E-mail etc.)
	Competing claims of responsibility?	Indication whether more than one group claimed separate responsibility for the attack

Category	Features	Explanation
Casualties and consequence	Total number of fatalities	Number of total confirmed fatalities for the incident
	Number of perpetrator fatalities	Limited to only perpetrator fatalities
	Total number of injured	Number of confirmed non-fatal injuries to both perpetrators and victims
	Number of perpetrators injured	Limited to only perpetrator fatalities
	Property damage?	Indication if there is evidence of property damage from the incident
	Extent of property damage	If "Property damage?" is "Yes", then one of three categories describes the extent of the property damage
	Value of property damage	If "Property damage?" is "Yes", then the exact U.S. dollar amount (at the time of the incident) of total damages is listed
	Hostages or kidnapping victims	This field records whether or not the victims were taken hostage
	Total number of hostages/kidnapping victims	Total number of hostages or kidnapping victims
	Hours/Days of kidnapping/hostage incident	Duration of the incident in case of kidnapping
	Ransom demanded	Indication if the incident involved a demand of monetary ransom
	Total ransom amount demanded	If a ransom was demanded, then the amount (in U.S. dollars) is listed in this field
	Total ransom amount paid	If a ransom amount was paid, then the amount (in U.S. dollars) is listed in this field
	Kidnapping/Hostage outcome	This field captures the eventual fate of hostages and kidnap victims
	Number released/escaped/rescued	Number of hostages who survived the incident

Received: 18 July 2019 Accepted: 4 December 2019

Published online: 18 December 2019

References

- Araña, J. E., & León, C. J. (2008). The impact of terrorism on tourism demand. *Annals of Tourism Research*, 35, 299–315.
- Arnold, J. L., Halpern, P., & Smithline, H. (2004). Mass casualty terrorist bombings: A comparison of outcomes by bombing type. *Annals of Emergency Medicine*, 43, 263–273.
- Asal, V., Gill, P., Rethemeyer, R. K., & Horgan, J. (2015). Killing range: Explaining lethality variance within a terrorist organization. *Journal of Conflict Resolution*, 59(3), 401–427.
- Asal, V., & Rethemeyer, R. K. (2008). The nature of the beast: Organizational structures and the lethality of terrorist attacks. *The Journal of Politics*, 70(2), 437–449.
- Aylwin, C., Koing, T. C., Brennan, N. W., Shirley, P. J., Davies, G., Walsh, M., et al. (2006). Reduction in critical mortality in urban mass casualty incidents: analysis of triage, surge, and resource use after the London Bombings on July 7, 2005. *Lancet*, 368, 2219–2225.
- Behlendorf, B., LaFree, G., & Legault, R. (2012). Predicting microcycles of terrorist violence: Evidence from FMLN and ETA. *Journal of Quantitative Criminology*, 28, 49–75.
- Ben-Gal, I. (2007). Bayesian networks. In F. Ruggeri, F. Faltn, & R. Kenett (Eds.), *Encyclopedia of statistics in quality and reliability*. New York: Wiley.
- Berman, E., & Laitin, D. (2006). Rational martyrs versus hard targets: Evidence on the tactical use of suicide attacks. In E. Meyerson (Ed.), *Suicide bombing from an interdisciplinary perspective*. Princeton, NJ: Princeton University Press.
- Berberi, C., & Lakdawalla, D. (2007). How does terrorism risk vary across space and time? An analysis based on the Israeli experience. *Defence and Peace Economics*, 18, 113–131.
- Brown, D., Dalton, J., & Hoyle, H. (2004). Spatial forecast methods for terrorist events in urban environments. *Lecture Notes in Computer Science*, 3073, 426–435.
- Caines, A., Pastrana, S., Hutchings, A., & Buttery, P. J. (2018). Automatically identifying the function and intent of posts in underground forums. *Crime Science*, 7(1), 19.
- Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., et al. (2000). *CRISP-DM: Step-by-step Data Mining Guide*. USA: SPSS Inc.
- Clarke, R. V. (1980). Situational crime prevention: Theory and practice. *British Journal of Criminology*, 20, 136–147.
- Clarke, R. V. (1997). *Situational crime prevention: Successful case studies*. Guildersland, New York: Harrow and Heston.
- Clarke, R. V. (2010). Situational crime prevention: Theoretical background and current practice. In M. D. Krohn, A. J. Lizotte, & G. P. Hall (Eds.), *Handbook on crime and deviance*. New York: Springer.
- Clarke, R. V., & Newman, G. (2006). *Outsmarting the terrorists*. Westport: Praeger Security International.
- Drakos, K. (2004). Terrorism-induced structural shifts in financial risk: Airline stocks in the aftermath of the September 11th terror attacks. *European Journal of Political Economy*, 20, 435–446.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2012). *Pattern classification*. New York: Wiley.
- Enders, W., & Sandler, T. (2000). Is transnational terrorism becoming more threatening?: A time-series investigation. *Journal of Conflict Resolution*, 44(3), 307–332.
- Fernando, M. L., Asier, M., & Miriam, E. (2018). Hate is in the air! But where? Introducing an algorithm to detect hate speech in digital microenvironments. *Crime Science*, 7(1), 15.
- Frank, E., & Witten, I. H. (1999). Generating accurate rule sets without global optimization. In *International Conference on Machine Learning* (pp. 144–151).
- Friedman, N., & Goldszmidt, M. (1996). Building classifiers using Bayesian networks. In *Proceedings of the National Conference on Artificial Intelligence* (pp. 1277–1284).
- Global Terrorism Database. (2016). *Codebook: Inclusion criteria and variables* (pp. 1–62).
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Heger, L., Jung, D., & Wong, W. H. (2012). Organizing for resistance: How group structure impacts the character of violence. *Terrorism and Political Violence*, 24(5), 743–768.
- Hoffman, B. (2006). *Inside terrorism*. New York: Columbia University Press.
- Hoffman, B., & McCormick, G. (2004). Terrorism, signaling and suicide attack. *Studies in Conflict and Terrorism*, 27, 243–281.
- Inyaem, U., Haruechaiyasak, C., Meesad, P., & Tran, D. (2010). Terrorism event classification using fuzzy inference systems. *International Journal of Computer Science and Information Security*, 7(3), 247–256.
- Janssens, D., Wets, G., Brijis, T., Vanhoof, K., Arentze, T., & Timmermans, H. (2006). Integrating Bayesian networks and decision trees in a sequential rule-based transportation model. *European Journal of Operational Research*, 175(1), 16–34.

- Kotsiantis, S. (2007). Supervised learning: A review of classification techniques. *Informatica*, 31, 249–268.
- Kuang, D., Brantingham, P. J., & Bertozzi, A. L. (2017). Crime topic modeling. *Crime Science*, 6(1), 12.
- LaFree, G. (2010). The Global Terrorism Database: Accomplishments and challenges. *Perspectives on Terrorism*, 4, 24–46.
- LaFree, G., & Bersani, B. E. (2014). County-level correlates of terrorist attacks in the United States. *Criminology & Public Policy*, 13, 455–481.
- LaFree, G., & Dugan, L. (2007). Introducing the global terrorism database. *Terrorism and Political Violence*, 19, 181–204.
- LaFree, G., Dugan, L., Xie, M., & Singh, P. (2012). Spatial and temporal patterns of terrorist attacks by ETA 1970 to 2007. *Journal of Quantitative Criminology*, 28, 7–29.
- Letham, B., Rudin, C., McCormick, T. H., & Madigan, D. (2015). Interpretable classifiers using rules and Bayesian analysis: Building a better stroke prediction model. *Annals of Applied Statistics*, 9, 1350–1371.
- Liu, Y. Y., Yang, M., Ramsay, M., et al. (2001). A comparison of logistic regression, classification and regression tree, and neural networks models in predicting violent re-offending. *Journal of Quantitative Criminology*, 27(4), 547–573.
- Maimon, O., & Rokach, L. (2005). *The data mining and knowledge discovery handbook*. Heidelberg: Springer.
- Mierau, J. O. (2015). The activity and lethality of militant groups: Ideology, capacity, and environment. *Dynamics of Asymmetric Conflict*, 8(1), 23–37.
- Mohler, G., Short, M. B., Malinowski, S., Johnson, M., Tita, G. E., Bertozzi, A. L., et al. (2015). Randomized controlled field trials of predictive policing. *Journal of the American Statistical Association*, 110(512), 1399–1411.
- Moore, A. W. (2001). *Cross-validation for detecting and preventing overfitting*. Carnegie: School of Computer Science, Carnegie Mellon University.
- Morris, N. A., & Slocum, L. A. (2012). Estimating country-level terrorism trends using group based trajectory analyses: Latent class growth analysis and general mixture modeling. *Journal of Quantitative Criminology*, 28, 103–139.
- Nemeth, S. (2014). The effect of competition on terrorist group operations. *Journal of Conflict Resolution*, 58(2), 336–362.
- Nilsson, M. (2018). Hard and soft targets: the lethality of suicide terrorism. *Journal of International Relations and Development*, 21(1), 101–117.
- Parachini, J. V. (2001). Comparing motives and outcomes in mass casualty terrorism involving conventional and unconventional weapons. *Studies in Conflict and Terrorism*, 24(5), 389–406.
- Peffley, M., Hutchison, M. L., & Shamir, M. (2014). The impact of persistent terrorism on political tolerance: Israel, 1980 to 2011. *American Political Science Review*, 109, 817–832.
- Piazza, J. A. (2009). Is Islamist terrorism more dangerous?: An empirical study of group ideology organization, and goal structure. *Terrorism and Political Violence*, 21(1), 62–88.
- Reed, G. S., Colley, W. N., & Aviles, S. M. (2013). Analyzing behavior signatures for terrorist attack forecasting. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 10, 203–213.
- Regens, J. L., Schultheiss, A., & Mould, N. (2015). Regional variation in causes of injuries among terrorism victims for mass casualty events. *Frontiers in Public Health*, 3(198), 1–6.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Model-agnostic interpretability of machine learning. In *ICML Workshop on Human Interpretability in Machine Learning (WHI)*, New York, NY.
- Shahbaz, M., Shabbir, M. S., Malik, M. N., & Wolters, M. E. (2013). An analysis of a causal relationship between economic growth and terrorism in Pakistan. *Economic Modeling*, 35, 21–29.
- Siebeneck, L. K., Medina, R. M., Yamada, I., & Hepner, G. F. (2009). Spatial and temporal analyses of terrorist incidents in Iraq, 2004–2006. *Studies in Conflict & Terrorism*, 32, 591–610.
- Simon, S., & Benjamin, D. (2000). America and the New Terrorism. *Survival*, 42, 59–75.
- Singer, G., & Golan, M. (2019). Applying data mining algorithms to encourage mental health disclosure in the workplace. *International Journal of Business Information Systems*. <https://doi.org/10.1504/IJBIS.2020.10019486>.
- Singer, G., Golan, M., Rabin, N., & Kleper, D. (2019). Evaluation of the effect of learning disabilities and accommodations on the prediction of the stability of academic behavior of undergraduate engineering students using decision trees. *European Journal of Engineering Education*. <https://doi.org/10.1080/03043797.2019.1677560>.
- Tavakoli, N. (2012). Effect of spirituality on decreasing crimes and social damages: A case study on Ramadan. *International Research Journal of Applied and Basic Sciences*, 3(3), 518–524.
- Wang, F., & Rudin, C. (2015). Falling rule lists. *JMLR Workshop and Conference Proceedings, San Diego CA*, 38, 1013–1022.
- Webb, J. J., & Cutter, S. L. (2009). The geography of US terrorist incidents, 1970–2004. *Terrorism and Political Violence*, 21, 428–449.
- White, G., Porter, M. D., & Mazerolle, L. (2013). Terrorism risk, resilience and volatility: A comparison of terrorism patterns in three Southeast Asian countries. *Journal of Quantitative Criminology*, 29, 295–320.

Publisher's Note

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