

JAKUB TUSZYŃSKI

Effectiveness of selected AI models in predicting victims of terrorist attacks¹

Abstract

Terrorism continues to be a problem for many countries around the world. The article compares the effectiveness of selected machine learning algorithms in predicting the victims of terrorist attacks in order to answer the question of whether they can serve as one of the anti-terrorist tools. An exploratory data analysis was carried out, and selected trends and characteristics of terrorist attacks were discussed. Some measures for evaluating the classification algorithms used in the study are presented, and potential directions for further research are indicated.

Keywords:

AI,
machine learning,
terrorism,
victims,
classification

¹ The article is based on a master's thesis entitled *Effectiveness of selected AI models in predicting victims of terrorist attacks*, defended at the Faculty of Journalism, Information and Bibliology, University of Warsaw. The author used excerpts from chapters 3 and 6. The thesis was awarded in the 12th edition of the competition of the Head of the Internal Security Agency for the best doctoral, master's or bachelor's thesis concerning state security in the context of intelligence, terrorist, economic threats.

Investigation of terrorist attacks

Due to definitional disputes, for the purposes of this article, a terrorist attack is assumed to be (...) *an intentional act or threat of violence by a non-state actor*². As part of this article, a study was conducted to compare the effectiveness of different artificial intelligence algorithms in predicting the victims of terrorist attacks. The study used the Global Terrorism Database (hereafter: GTD) maintained by researchers from the START consortium³, which contains information on terrorist attacks.

The GTD has adopted three criteria, at least two of which must be met for an event to be considered a terrorist attack. These are:

- the act of violence was intended to achieve a political, economic, religious or social objective;
- the act of violence contained evidence of intent to coerce, intimidate or otherwise convey a message to a wider audience other than the immediate victims;
- the act of violence fell outside the scope of international humanitarian law⁴.

Events where the amount of information was insufficient to clearly determine whether the event was a terrorist attack or not and are filterable by the user were also flagged.

Assumptions

In order to eliminate, as accurately as possible, cases in which an event was misclassified as a terrorist attack, those observations that did not meet all three criteria described above and those about which the authors of the database had doubts were excluded.

Victims of an attack are considered to be any non-terrorist person injured or killed as a result of the incident. Several machine learning models were built and compared using appropriate metrics.

² *Data Collection Methodology*, Global Terrorism Database, <http://www.start-dev.umd.edu/gtd/using-gtd/> [accessed: 21 V 2022].

³ *History of the GTD*, Global Terrorism Database, <https://start.umd.edu/gtd/about/History.aspx> [accessed: 11 V 2022].

⁴ *Data Collection Methodology...*

Exploratory data analysis

During the exploratory data analysis, the focus was on understanding the dataset under study. First, a structural analysis was carried out.

Figure 1 distinguishes the following features of the collection:

- more than 200 000 lines of recorded attacks;
- 135 columns containing characteristics describing the event in question;
- The *dtypes* describe the data types of the individual columns. These are categorical data, having a finite number of categories - 9, columns containing floating point numbers - 53, integers occurring in 24 columns and 49 columns containing data that can be both strings and numbers. The data type *object* is assigned when no other data type can be explicitly assigned;
- the collection takes up approximately 200 megabytes of memory.

```
Int64Index: 201183 entries, 0 to 201182
Columns: 135 entries, eventid to related
dtypes: category(9), float64(53), int64(24), object(49)
memory usage: 197.1+ MB
```

Figure 1. Basic statistics describing the dataset.

Source: own elaboration.

Next, data that did not meet all three criteria for a terrorist attack and those about which the authors of the database had doubts were filtered out. Columns in which 50 per cent or more of the rows were empty were also removed. The purpose of reducing the dataset was to speed up the operations performed on it. In addition, most of the machine learning algorithms used in the study require the dataset to have no empty values. Complementing them with the mean, median or most frequent value with such a large number of empty observations would result in an unauthorised generalisation from a small amount of data. This operation reduced the dataset to 154 260 rows and 60 columns, with a file size of just under 70 megabytes.

In the next step, the basic values of the individual columns were checked. The x-axis shows the column names and the y-axis shows

the calculated statistics: number of observations, mean, standard deviation, minimum value, first quartile, median, third quartile and maximum value.

In Figure 2, it can be seen that the *latitude* and *longitude* variables contain missing values, which were appropriately reprocessed by removing the rows with missing values in these variables, so they were used during modelling. It is also worth noting the minimum value of the *vicinity* variable, which is -9 (this is visible in the last row of the *min* column). The authors of the database mark cases of missing data in this way. This is described in the so-called *Codebook*⁵.

	count	mean	std	min	25%	50%	75%	max
eventid	154260.0	2.805463e+11	1.294940e+09	1.970000e+11	1.995042e+11	2.012022e+11	2.015091e+11	2.019123e+11
year	154260.0	2.005397e+03	1.294940e+01	1.970000e+03	1.995000e+03	2.012000e+03	2.015000e+03	2.019000e+03
imonth	154260.0	6.443965e+00	3.392222e+00	0.000000e+00	4.000000e+00	6.000000e+00	9.000000e+00	1.200000e+01
iday	154260.0	1.563523e+01	8.803117e+00	0.000000e+00	8.000000e+00	1.500000e+01	2.300000e+01	3.100000e+01
extended	154260.0	5.527032e-02	2.285079e-01	0.000000e+00	0.000000e+00	0.000000e+00	0.000000e+00	1.000000e+00
country	154260.0	1.297654e+02	1.116024e+02	4.000000e+00	7.800000e+01	9.700000e+01	1.400000e+02	1.004000e+03
region	154260.0	7.321360e+00	2.438782e+00	1.000000e+00	6.000000e+00	8.000000e+00	1.000000e+01	1.200000e+01
latitude	151329.0	2.369732e+01	1.798608e+01	-5.315461e+01	1.184079e+01	3.153024e+01	3.451689e+01	7.463355e+01
longitude	151329.0	3.227138e+01	5.485520e+01	-1.578583e+02	9.735686e+00	4.414823e+01	6.914701e+01	1.793667e+02
specificity	154259.0	1.447591e+00	9.567426e-01	1.000000e+00	1.000000e+00	1.000000e+00	1.000000e+00	5.000000e+00
vicinity	154260.0	6.333463e-02	2.803160e-01	-9.000000e+00	0.000000e+00	0.000000e+00	0.000000e+00	1.000000e+00

Figure 2. Extract from statistics on numerical variables.

Source: own elaboration.

Figure 3 shows the statistics for the text variables: number of observations, number of unique values, most frequent value and its frequency.

In the dataset, some of the variables appear in both numeric and text form. This is, for example, the variable *region_txt* (Figure 3) and the variable *region* (Figure 2). This was taken into account before modelling because of the possible correlation between the same variables and the unnecessary complexity of the dataset, which translates into a slower model training process.

⁵ *Codebook: Inclusion Criteria and Variables*, Global Terrorism Database, August 2018, <https://www.start.umd.edu/gtd/downloads/Codebook.pdf> [accessed: 30 V 2022].

	count	unique	top	freq
country_txt	154260	202		Iraq 23407
region_txt	154260	12		Middle East & North Africa 43858
provstate	154260	2380		Baghdad 7563
city	153874	34443		Unknown 7594
summary	112205	109189	09/00/2016: Sometime between September ...	100
attacktype1_txt	154260	9		Bombing/Explosion 79879
targettype1_txt	154260	22		Private Citizens & Property 43145
targetsubtype1_txt	144441	112		Unnamed Civilian/Unspecified 11599
corp1	123010	32224		Unknown 18458
target1	153807	73022		Civilians 7489
natlty1_txt	152648	209		Iraq 23077

Figure 3. Extract from statistics on categorical variables.

Source: own elaboration.

There is also an apparent problem with the high number of unique values of some variables, such as *city*, which may have affected the performance of the models built. This problem was resolved during data processing by removing such variables from the dataset.

In the next step, empty values were removed from the variables affecting the number of injured and killed. The problem of giving the total number of victims killed and injured by the incident with the number of terrorists killed and injured was then solved. The number of terrorists was subtracted from the total number of killed and injured to obtain the number of non-terrorist casualties.

After this operation, new variables were created: *ncasualites*, which is the sum of killed and injured, and *cas_class*, where zero was used to denote cases where there were no casualties and one to denote those events where there were casualties. The preprocessed data were then saved to a new file.

Trends in the number of terrorist attacks and the number of victims

Visualisations of some of the variables were produced, allowing for a more in-depth analysis. The charts show the number of terrorist attacks carried out between 1970 and 2019 and their victims.

From the second half of the 1970s to the early 1990s, there was an increase in the number of attacks (Chart 1) as well as their victims (Chart 2). The marked increase in the number of victims in 2001 is due to the 11 September attack on the World Trade Center (WTC). Another noticeable upward trend in both occurred in 2005 and continued until

2014-2015. With, as already mentioned, the increase around 2012 is partly due to a change in data collection methodology⁶, however, this increase started even before 2005. Since 2015, a decreasing trend in the number of both attacks and victims can be seen. As of 4 June 2022⁷, researchers from the START consortium have not published data from 2020-2021, so the impact of the COVID-19 pandemic on the dynamics of the incidence of terrorist attacks is unknown.



Chart 1. Number of terrorist attacks 1970-2019.

Source: own elaboration.



Chart 2. Number of victims of terrorist attacks 1970-2019.

Source: own elaboration.

⁶ Identifying terrorist incidents to the GTD prior to 2012 required the use of approximately 300 unique news sources. After the 2012 update, more than 1,500 unique news sources were used. These sources included international news agencies and English translations of local newspapers published in various foreign languages.

⁷ Update - as of 15 July 2023, data from the first half of 2021 are available.

The number of terrorist attacks and their victims was then visualised by region (Charts 3 and 4).




Chart 3. Number of terrorist attacks 1970-2019 by region.

Source: own elaboration.

In Western Europe, most attacks took place in the second half of the 1970s, which did not translate into an increase in the number of victims. A similar situation can be observed in South America between 1980 and 1995, when there was a dynamic increase in the number of attacks (also not translated into casualties). In the 1990s, there was a decline in the number of attacks in all regions. Nevertheless, an increase in casualties was observed in Sub-Saharan Africa, the Middle East and North Africa, and South Asia and East Asia. The sharp increase in casualties in 2001 in North America is not a data entry error, as that was when the attack on the WTC took place.

The Middle East and North Africa region, South Asia and Sub-Saharan Africa also saw an increase in the number of attacks and casualties in the early 2000s. The growth curve in the Middle East stands out in particular, being steeper than the curves of the two previously mentioned regions. This may be related to the US intervention in Afghanistan and Iraq. In the Middle East, the number of casualties in 2000 was around 900, and in 2005 it was almost 9,000. Since 2015, there has been an apparent downward trend in the number of both terrorist attacks and victims (with the exception of South Asia, where upward trend in the number of casualties is observed).




Chart 4. Number of victims of terrorist attacks 1970-2019 by region.

Source: own elaboration.

It is worth noting the Western European area during this period, as despite the migration crisis of 2015, there has been no marked increase in terrorist attacks there (Chart 5).




Chart 5. Number of terrorist attacks in Western Europe 1970-2019.


Source: own elaboration.

Table 1 and Chart 6 provide information on the ratio of casualties to attacks, hereafter referred to as the casualty ratio, which is a more meaningful indicator than directly comparing data.

Table 1. Comparison of terrorist attack casualty rates by region.

Region	Casualty rate
East Asia	19.6:1
North America	10.7:1
Middle East and North Africa	7.0:1
Sub-Saharan Africa	6.9:1
Central Asia	6.2:1
South Asia	5.0:1
Eastern Europe	4.0:1
South-East Asia	3.0:1
Central America and the Caribbean	2.7:1
South America	2.3:1
Western Europe	1.8:1
Australasia and Oceania	1.7:1

Source: own elaboration.

**Chart 6.** Number of terrorist attacks and number of victims by region.

Source: own elaboration.

Attacks in the Middle East and North Africa and South Asia - despite the highest absolute numbers - are not the most deadly. The highest casualty

rate occurs in East Asia and almost doubles that of North America, which is in second place. Sub-Saharan Africa is in third place. On average, attacks in Australasia and Oceania and Western Europe have the lowest casualty rate. It is clear from these statistics that, when it comes to terrorism, the divide between the rich north and the poor south does not translate directly. Some of the poorer regions, such as South America, South-East Asia or Central America and the Caribbean, do not have a significantly higher casualty rate than the richer regions. This may indicate non-economic factors that influence the effectiveness of terrorists.

After analysing the data, the ten countries with the highest number of terrorist attacks carried out between 1970 and 2019 (Chart 7) and the 10 countries with the highest number of victims of these attacks (Chart 8) were identified in turn.

These countries (by region⁸) are:

- Middle East and North Africa: Iraq, Yemen, Syria;
- South Asia: Afghanistan, Pakistan, India, Sri Lanka;
- Southeast Asia: Philippines, Thailand;
- South America: Colombia, Peru;
- North America: United States;
- Sub-Saharan Africa: Nigeria.




Chart 7. Number of terrorist attacks in countries ranking in the top 10 in terms of number of such attacks.

Source: own elaboration.

⁸ Countries are assigned to their regions according to the *Codebook*.




Chart 8. Number of victims of terrorist attacks in countries ranking in the top 10 in terms of such victims.

Source: own elaboration.

Due to the relatively low number of attacks and casualties compared to other countries, no East Asian country is on the list (despite a high casualty rate). Iraq is the most affected, with almost twice as many terrorist attacks as Afghanistan, which translates into more than twice as many victims. The other countries included already have more similar values in terms of both the number of terrorist attacks and the victims of these incidents.

Table 2 indicates the casualty rate for the countries with the highest number of casualties. The highest ratio are 12.8 for Syria, 12.0 for the United States and 11.2 for Sri Lanka. On average, the least number of victims of the attack was in Colombia and the Philippines.

Table 2. Comparison of the casualty rate of terrorist attacks for the 10 countries with the highest number of such casualties.

Country	Casualty rate
Syria	12.8:1
United States	12.0:1
Sri Lanka	11.2:1
Iraq	8.2:1
Nigeria	6.6:1
Afghanistan	5.6:1
Pakistan	4.7:1
India	3.8:1
Colombia	2.8:1
Philippines	2.8:1

Source: own elaboration.

It can therefore be concluded that geographical factors have a bearing on whether or not a terrorist attack will result in casualties. They were therefore taken into account during the modelling and their impact on the prediction result has been analysed.

Activity of selected terrorist groups

An analysis of the activity of the groups responsible for terrorist attacks was then carried out. Chart 9 shows the most active terrorist groups between 1970-2019. The top ten places are occupied by groups linked to Islamic terrorism: Taliban, Al-Shabaab, Boko Haram, Islamic State (ISIS) also known as Islamic State of Iraq and the Levant (ISIL) and far-left terrorism: Basque Fatherland and Freedom (ETA), the New People's Army (NPA), the Revolutionary Armed Forces of Colombia (FARC), the Maoists, Shining Path (Spanish: Sendero Luminoso (SL), Communist Party of India-Maoist (CPI-Maoist). This may indicate some correlation between a group's motivation to carry out an attack and its high level of activity; however, in the GTD, individual groups are not assigned characteristics. This may be a clue for the START consortium researchers to update the GTD in this regard.




Chart 9. Number of terrorist attacks carried out between 1970 and 2019 by the 10 most active terrorist groups.

Source: own elaboration.

The oldest group of all those mentioned above is ETA, which, during the period under study, was most active between 1970 and 1980 and has been increasingly inactive since the early 1990s. The last recorded attack in the GTD carried out by this group was in 2011. This is most likely related to the end of the organisation's armed activities in October of the same year⁹. ETA finally self-dissolved in 2018¹⁰.

Also active to a fairly high degree was the Shining Path, whose activity was mainly between 1980 and 1990. Thereafter, there was a marked decline in the number of attacks carried out by this group. This was probably due to the arrest of its successive leaders in the 1990s and in 2012¹¹. One faction of the Shining Path, the Militarised Communist Party of Peru, which in 2018 splintered from the SL, remained active¹².

Another group whose increased activity occurred in the 1980s is the New People's Army operating in the Philippines. It was formed in 1969 as the armed arm of the Communist Party of the Philippines¹³. The decline in the group's activity after 1990 may be due to the arrest of key figures in the organisation and internal purges¹⁴. Various ceasefires and negotiations from 1986 and 2010, among others, were broken by the NPA¹⁵. Another increase in the group's activity has been recorded since 2013. The NPA, unlike the SL and ETA, remains active.

The FARC was founded in 1964. The original aim of the organisation was to overthrow the government of Colombia¹⁶. Between 2008 and 2015 there was an increase in its terrorist activity. Although the government has attempted to negotiate with the FARC since 2012, including ceasefires,

⁹ *Basque group Eta says armed campaign is over*, BBC News, 20 X 2011, <https://www.bbc.com/news/world-europe-15393014> [accessed: 6 VI 2022].

¹⁰ I. Binnie, *Basque separatist group ETA says it has "completely dissolved"*, Reuters, 2 V 2018, <https://www.reuters.com/article/us-spain-eta-idUSKBN1I31TP> [accessed: 6 VI 2022].

¹¹ S. Saffón, *Peru in Familiar Stalemate With Shining Path Rebels*, InSight Crime, 4 IX 2020, <https://insightcrime.org/news/brief/peru-stalemate-shining-path/> [accessed: 8 VI 2022].

¹² Ibid.

¹³ *Communist Part of the Philippines – New People's Army*, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/communist-party-philippines-new-peoples-army> [accessed: 11 VI 2022].

¹⁴ Ibid.

¹⁵ Ibid.

¹⁶ *Revolutionary Armed Forces of Colombia (FARC)*, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/revolutionary-armed-forces-colombia-farc> [accessed: 8 VI 2022].

the group has regularly broken them¹⁷. In 2016, an agreement was reached between the government and the FARC, which transformed itself into a political party and ceased its armed activities¹⁸. According to GTD data, not a single attack has been carried out by this organisation since 2016.

The Communist Party of India (Maoist), on the other hand, is a group outlawed by the Indian government¹⁹. The increase in CPI-Maoist activity in 2009 and its subsequent decline may be linked to “Green Hunt” counter-terrorism operation targeting the organisation²⁰. Despite the efforts made, the CPI-Maoist could not be dismantled. The Maoists, in turn, is the collective name for terrorist far-left groups not part of the CPI-Maoist²¹.

The Taliban emerged as an organisation in 1994 and ruled Afghanistan from 1996 to 2001²². A surge in their activity can be seen after 2001, when they were defeated militarily, but the organisation was not dismantled. The Taliban returned to power in 2021.

Al-Shabab is a Somali-origin organisation operating in East Africa and Yemen. Activity peaked in 2014. However, there was a decline a year later, which can be linked to the killing of one of the organisation’s leaders by the US in 2014, a major loss for the group²³. Despite efforts to combat the organisation by the Somali government and the US, it remains active.

Boko Haram was established in Nigeria in 2002²⁴ and began its activities as a terrorist organisation in 2009, when a shootout took place between

¹⁷ Ibid.

¹⁸ Ibid.

¹⁹ *Left Wing Extremism Division*, Ministry of Home Affairs, https://web.archive.org/web/20220707070953/https://www.mha.gov.in/division_of_mha/left-wing-extremism-division [accessed: 8 VI 2022].

²⁰ A. Sethi, *Green Hunt: the anatomy of an operation*, *The Hindu*, 6 II 2010, <https://www.thehindu.com/opinion/op-ed/Green-Hunt-the-anatomy-of-an-operation/article16812797.ece>. [accessed: 8 VI 2022].

²¹ *Deaths in Maoist attacks down by 21%: Shah at CMs’ meeting*, *The Times of India*, 27 IX 2021, <https://timesofindia.indiatimes.com/india/deaths-in-naxal-attacks-down-by-21-shah-at-cms-meeting/articleshow/86543018.cms> [accessed: 8 VI 2022].

²² *The Afghan Taliban*, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/afghan-taliban> [accessed: 8 VI 2022].

²³ *Pentagon confirms death of Somalia terror leader*, *The Washington Times*, 5 IX 2014, <https://www.washingtontimes.com/news/2014/sep/5/pentagon-confirms-death-of-somalia-terror-leader/> [accessed: 10 VI 2022].

²⁴ H. Matfess, *Boko Haram: History and Context*, in: *Oxford Research Encyclopedia of African History*, Oxford University Press 2017, p. 1.

its members and the police²⁵. It is difficult to predict trends in the group's activity, as it is influenced by the actions of the authorities in the form of, for example, the use of armed forces, but also by Boko Haram's resilience to these actions²⁶.

The last group discussed is the Islamic State. Originally founded in 1999 under the name Jama'at al-Tawhid wal-Jihad, the organisation was renamed Al-Qaeda in Iraq after the US intervention in Iraq²⁷. The group was strengthened after the American withdrawal from Iraq. Taking advantage of this fact and the outbreak of civil war in Syria, it began successively taking over areas of both countries in 2013 and 2014²⁸. During this period, the organisation changed its name first to Islamic State in Iraq and Syria, then in June 2014 it announced the creation of a caliphate and finally adopted the name Islamic State²⁹. The dynamics of change in the group's activity are clearly linked to its progress in taking over the territories of the above-mentioned countries. It reached its peak in 2014 and was highly active until 2017, at times surpassing the Taliban. The apparent decline in activity in 2018 and 2019 is most likely linked to the group's loss of most territories. In 2017, ISIS was pushed out of the urban centres it controlled, and two years later the organisation lost control of its last territories in Baghuz province in Syria³⁰. With the loss of territories, and thus the means to conduct operations, ISIS activity has gradually decreased. There is a clear downward trend in this respect, but despite the efforts of Kurdish, Iraqi and Syrian forces or the involvement of the US army, the group remains active.

It is noteworthy that all of the most active groups, despite the various actions and measures taken, continue to be active, excluding those that have self-dissolved as a result of the negotiations undertaken with these groups by their respective governments.

A casualty rate was calculated for the most active terrorist groups. It is shown in Table 3.

²⁵ Ibid, p. 7.

²⁶ Ibid, p. 15.

²⁷ *The Islamic State*, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/islamic-state> [accessed: 11 VI 2022].

²⁸ Ibid.

²⁹ Ibid.

³⁰ Ibid.

Table 3. Casualty rates for the most active terrorist groups.

Terrorist group	Casualty rate
Islamic State	11.0:1
Boko Haram	10.0:1
Al-Shabab	5.8:1
Taliban	5.8:1
Revolutionary Armed Forces of Colombia	3.6:1
Shining Path	2.7:1
Communist Party of India (Maoist)	1.9:1
New People's Army	1.8:1
Basque Country and Freedom	1.6:1
Maoists	1.4:1

Source: own elaboration.

It can be seen that the casualty rate for Islamist groups is significantly higher than for far-left groups. This could be another argument for adding a new feature to the dataset - indicating the religious or political affiliation of the group in question and exploring this further. This could have a positive impact on the improvement of machine learning models for the prediction of victims of terrorist attacks.

Most common targets of terrorist attacks

The characteristics of terrorist attacks including their targets and types were then visualised. Chart 10 shows the most common targets of terrorist attacks carried out between 1970 and 2019, which were, in order: civilians and property, police, government institutions (general), business, public transport, energy infrastructure, military, religious leaders or religious institutions, schools, government institutions (diplomats).




Chart 10. Most frequent targets of terrorist attacks by number of attacks carried out against them.

Source: own elaboration.

A high number of attacks targeting civilians and property may reflect a desire to intimidate public opinion in a country and thus influence the actions of its government. Conversely, attacks on the police, the military or government institutions, may reflect the frequent political motivation of terrorist groups that regard state authorities as an enemy to be fought. Attacks on the private sector in the broadest sense, carried out for example by the FARC, may indicate the extreme left-wing motivations of such groups, whose aim in the longer term may be to bring about the abolition of private property.

Types of terrorist attacks

It also examined which type of terrorist attack was carried out most frequently (Chart 11). A regional breakdown was made, which allowed the specifics of attacks carried out in different parts of the world to become more apparent.




Chart 11. Percentage of different types of terrorist attacks by region.

Source: own elaboration.

Attacks using explosives accounted for the largest proportion in all regions surveyed, excluding Australasia and Oceania. In addition, armed robbery and homicide and, in some regions, attacks on infrastructure were common. This can provide an indication to state governments as to what type of attack state institutions should be prepared for. They can then review whether the services competent to deal with the terrorist threat have developed adequate procedures for the type of event, and whether the health services will be efficient enough to effectively care for the injured (this could reduce the number of fatalities). Such an evaluation of the current capacities of state institutions would indicate certain gaps in resilience to these types of events.

Data processing

In the next step, data processing was carried out. First, the final variables used for modelling were selected. These are presented in Table 4.

Table 4. Description of variables used in the survey.

Variable name	Variable description
<i>Extended</i>	determines whether the duration of the incident exceeded 24 hours
<i>Country_txt</i>	identifies the country in which the event occurred
<i>Region</i>	identifies the region in which the event occurred

<i>Latitude</i>	determines the latitude of the place where the event occurred
<i>Longitude</i>	determines the longitude of the place where the event occurred
<i>Specificity</i>	defines the geospatial resolution of the latitude and longitude fields. The most detailed resolution available across the dataset is the centre of the city, village or town where the attack occurred. Higher resolution coordinates, although possible, are not systematically included in the database
<i>Vicinity</i>	determines whether the event occurred in the immediate vicinity of the city concerned
<i>Multiple</i>	determines whether a terrorist attack is linked to other attacks
<i>Success</i>	the success of a terrorist attack is defined by its tangible effects. Success is not assessed in terms of the perpetrators' broader goals. For example, a bomb that exploded in a building would be considered a success even if it did not succeed in destroying the building or provoking government retaliation
<i>Suicide</i>	determines whether an attack was a suicide bombing
<i>Attack type1</i>	identifies the type of terrorist attack
<i>Targ type1</i>	identifies the type of terrorist target
<i>Targ subtype1</i>	defines the objective category in more detail
<i>Natly1</i>	is the nationality of the target attacked, not necessarily the same as the country in which the incident occurred, although in most cases this is the case. In the case of aircraft hijacking, the nationality of the aircraft is recorded, not the nationality of the passengers
<i>Gname</i>	includes the name of the group that carried out the attack
<i>Guncertain1</i>	determines whether information provided by sources about the group responsible for the attack is based on speculation or questionable claims of responsibility
<i>Individual</i>	determines whether the attack was carried out by a person or several persons not known to be associated with a terrorist group or organisation

<i>Nperps</i>	determines the total number of terrorists involved in the incident
<i>Claimed</i>	is used to indicate whether a group or individual(s) has admitted to an attack
<i>Weaptype1</i>	identifies the type of weapon used
<i>Weapsubtype1</i>	defines the category of weapons in more detail
<i>Property</i>	determine whether property has been damaged as a result of the incident
<i>Ishostkid</i>	determines whether the victims were taken hostage or abducted during the incident
<i>Int_log</i>	indicates whether, in order to carry out the attack, the group of perpetrators crossed the border
<i>Int_misc</i>	indicates whether the group of perpetrators attacked a target of a different nationality
<i>Int_any</i>	determines whether all conditions for variables with the prefix <i>int</i> are met
<i>Cas_class</i>	determines whether an incident has caused casualties

Source: own elaboration based on: Codebook: Inclusion Criteria and Variables, Global Terrorism Database, August 2018, <http://www.start-dev.umd.edu/gtd/downloads/Codebook.pdf> [accessed: 30 V 2022].

This reduced the number of columns from 60 to 28. Constraint the number of features will speed up the training process for the machine learning models. Due to the fact that some of the columns are in numeric or text form such as *region* and *region_txt*, it was decided to select only one column in such a case, in order to simplify the dataset. Columns containing metadata such as the unique event identifier or the original data source were also removed. Variables specifying the conditions for the inclusion of an event in the GTD were not taken into account because, as already mentioned, those attacks that did not meet all the conditions and those for which there was doubt were filtered out. Thus, these variables do not carry meaningful information for the machine learning models to influence prediction, and would only prolong the training process. The last group of features removed are variables that specify the number of victims or injured in order to avoid data leakage, which would have made the results of the study unreliable.

The number of terrorist attacks that ended in casualties was then examined. About 59 per cent of cases resulted in at least one non-terrorist being killed or injured. This means that there is an imbalance of classes in the predicted variable, which can negatively affect the prediction result. This problem was addressed when building the machine learning models by setting the *class weight* parameter to a *balanced* value. This allows the models to pay more attention to the lesser class, which helps to balance the impact of each class on the model, increasing the overall prediction performance.

A further split was made into a training set (80 per cent of the data) and a test set (the remaining 20 per cent). Stratified sampling was used to ensure that the classes of the predicted variable were similarly distributed throughout the dataset.

In the next step, a pipeline was created performing the final transformations on the dataset.

The first step is to transform the *country_txt* and *gname* variables from textual to numeric form, as the machine learning models built can only work on data in this form. Due to the large number of unique values in both columns, it was decided to code them based on the number of occurrences. This method - as opposed to the 1-of-n (one-hot encoding) method, which results in as many columns as there are unique values - does not have the side-effect of increasing the dimensionality of the data.

The blank values were then replaced with -9. This is how the GTD codes values for which the researchers did not have sufficient information to explicitly assign a specific value to an event characteristic³¹.

In the final step, the data was normalised due to the fact that some of the models built, such as logistic regression and support vector machines, are sensitive to extremely different value scales. This method solves this problem and can translate into better performance of these models and speed up the learning process.

It is worth pointing out that the calculation of the number of occurrences of a given value should only take place on the training set. On the test set, on the other hand, transformations are made on the basis of calculations from the training set. This is an important point, because a different procedure leads to data leakage and thus affects the test results. For this reason, it was decided to use the pipeline available in the Scikit-

³¹ *History of the GTD...*

learn library, which makes it easy to control the transformation steps and thus reduce the risk of error (Figure 4).




Figure 4. Pipeline performing transformations on a data set.

Source: own elaboration.

At the end of the processing stage, the training and test sets were saved to separate files in order to save the transformations carried out and to be able to proceed directly to modelling afterwards.

The process of training machine learning models

All models were trained on a desktop computer with parameters: 16 GB RAM, AMD Ryzen 5 3600 processor. Each model was trained in an analogous way: hyperparameters were queried using the Optuna framework, and metrics were written using the MLFlow library. To avoid unnecessary repetition, the training process will only be shown using the decision tree as an example (Figure 5).

```

def objective(trial):
    params = {
        "max_depth": trial.suggest_int("max_depth", 15, 50),
        "min_samples_leaf": trial.suggest_int("min_samples_leaf", 1, 40),
        "class_weight": trial.suggest_categorical("class_weight", ["balanced"]),
        "criterion": trial.suggest_categorical("criterion", ["gini", "entropy"])
    }

    model = DecisionTreeClassifier(**params)

    scoring = ["accuracy", "precision", "recall", "f1"]

    preds = cross_validate(model, X_train, y_train, cv=5, n_jobs=-1, scoring=scoring)

    accuracy = np.mean(preds["test_accuracy"])
    precision = np.mean(preds["test_precision"])
    recall = np.mean(preds["test_recall"])
    f1 = np.mean(preds["test_f1"])

    return accuracy, precision, recall, f1

```

Figure 5. Code fragment responsible for searching for hyperparameters.

Source: own elaboration.

The first line defines a function whose name and accepted arguments follow the convention adopted in Optun. In lines 2-7, the space of hyperparameters was defined, which in a later stage was searched to find the best possible combination of them. In the case of lines 2-3, the interval in which the values of these hyperparameters are to be searched was marked and in this case it will be an integer. Particularly noteworthy is line 5, which indicates that the class weight should be balanced. This is one of the methods to solve the previously mentioned problem of unbalanced classes. This was followed by the initialisation of the model on line 9, and then the metrics used to evaluate the models were specified, namely accuracy, precision, recall and F1. These will be discussed in detail further on in the article. In line 13, the model is trained on the training set using cross validation, and then the scores for the individual variables are calculated, which are finally returned by this function.

In the next step, a so-called *study* was created, in which its name and the direction of optimisation of the metrics are specified. As the metrics used relate to the classification problem, they have been maximised. This is followed by a search of the hyperparameters (Figure 6), which are stored using MLFlow (Figure 7).

```

1 study = optuna.create_study(study_name="decision_tree",
2                             directions=["maximize", "maximize", "maximize", "maximize"])
3 study.optimize(objective, n_trials=100, callbacks=[mlflow_callback])

```

Figure 6. Initialisation of hyperparameter searches by Optuna.

Source: own elaboration.

Metrics <		Parameters <							
<input type="checkbox"/>	accuracy	f1	precision	recall	criterion	max_depth	max_features	min_samples_leaf	min_samples_split
<input type="checkbox"/>	0.849	0.868	0.892	0.845	entropy	21	-	38	-
<input type="checkbox"/>	0.848	0.868	0.887	0.849	entropy	37	-	15	-
<input type="checkbox"/>	0.848	0.868	0.89	0.846	entropy	33	None	20	82
<input type="checkbox"/>	0.848	0.867	0.888	0.848	entropy	25	-	17	-
<input type="checkbox"/>	0.848	0.867	0.89	0.846	entropy	18	None	34	31
<input type="checkbox"/>	0.848	0.867	0.89	0.846	entropy	18	None	34	31

Figure 7. Excerpt from the MLflow dashboard with stored metrics and decision tree parameters.

Source: own elaboration.

For each model, several hundred iterations were performed to search for the optimal hyperparameters. All models with the optimal hyperparameters for them were then trained on the entire training set and validated on the test set. Those models that achieved the highest values of the F1 metric were selected for training. The parameters of these models and their results will be presented later. The individual quality metrics of the classification models used during the study are described earlier.

Selected quality measures of classification models

Due to the multitude of different quality measures used in the evaluation of classification models, only those measures used during the study are described.

The accuracy of a classifier is a measure of how many cases are classified correctly³². It can be represented by the following formula:

³² S. Raschka et al., *Machine Learning with PyTorch and Scikit-Learn: Develop machine learning and deep learning models with Python*, Birmingham 2022, p. 13.

$$\text{accuracy} = \frac{\text{TP} + \text{TN}}{\text{FP} + \text{FN} + \text{TP} + \text{TN}}$$

where:

- TP (true positive), where the model has correctly classified an event as causing casualties;
- TN (true negative) cases, where the model has correctly classified an event as not causing casualties;
- FP (false positive) cases, where the model has incorrectly classified an event as causing casualties;
- FN (false negative) cases, where the model has incorrectly classified an event as not causing casualties³³.

Two further measures, precision and recall, are directly related. Precision promotes a situation in which the classifier is confident in its decisions and makes as few false-positive errors as possible, but at the cost of this is an increase in false-negative predictions³⁴. The opposite is true with regard to recall, since the situation is promoted in which the classifier makes as few false-negative errors as possible, but at the expense of increasing false-positive cases³⁵. Thus, if we optimise the classifier so that it minimises the chances of incorrectly classifying an event as a non-victory, it will have high recall. The formula of recall is as follows:

$$\text{recall} = \frac{\text{TP}}{\text{FN} + \text{TP}}$$

However, this will come at the expense of precision. The formula of precision is as follows:

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

³³ Ibid, p. 195.

³⁴ Ibid, p. 196.

³⁵ Ibid.

In order to balance precision and sensitivity, the F1 measure is used, which is the harmonic mean of precision and recall³⁶. This means that in order to achieve a high F1 measure, the classifier must have high scores in both precision and sensitivity, as the harmonic mean attaches more importance to low values³⁷.

$$F1 = 2 \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

Test results for individual models

In Tables 5–8 the hyperparameters for each machine learning model are indicated. Table 9 presents the results of their validation on the test set. It should be noted that only those hyperparameter values that were previously searched for are shown. If a hyperparameter is not in the table, it means that it takes a default value according to the documentation of the relevant library. Since the voting classifier and the stacked classifier are composed of other built models, their hyperparameters are identical to those shown in tables.

Table 5. Hyperparameter values for logistic regression.

Hyperparameter name	Value
<i>C</i>	9.58649376280703
<i>class_weight</i>	balanced
<i>max_iter</i>	500

Source: own elaboration.

Table 6. Hyperparameter values for a linear support vector machine.

Hyperparameter name	Value
<i>C</i>	0.0036775852394361204

³⁶ A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, Sebastopol 2019, p. 140.

³⁷ Ibid.

<i>class_weight</i>	balanced
<i>dual</i>	false
<i>penalty</i>	L1

Source: own elaboration.

Table 7. Hyperparameter values for the decision tree.

Hyperparameter name	Value
<i>criterion</i>	entropy
<i>class_weight</i>	balanced
<i>max_depth</i>	21
<i>min_samples_leaf</i>	38

Source: own elaboration.

Table 8. Hyperparameter values for the random forest.

Hyperparameter name	Value
<i>criterion</i>	entropy
<i>class_weight</i>	balanced
<i>max_depth</i>	42
<i>min_samples_split</i>	6
<i>n_estimators</i>	551
<i>max_features</i>	sqrt

Source: own elaboration.

Table 9. Hyperparameter values for the XGBoost model.

Hyperparameter name	Value
<i>colsample_bylevel</i>	0.6783261477402747
<i>colsample_bytree</i>	0.23127225599162296
<i>gamma</i>	0.4906870500968865
<i>learning_rate</i>	0.0675784773135259
<i>max_delta_step</i>	5
<i>max_depth</i>	22
<i>min_child_weight</i>	1
<i>n_estimators</i>	1475

<i>reg_alpha</i>	0.12263684424466229
<i>reg_lambda</i>	1.9559489540115411
<i>scale_pos_weight</i>	0.9676022078596858
<i>subsample</i>	0.976706655475712

Source: own elaboration.

Table 10 shows the results of each model along with their training times. The best results from each category are shown in bold.

Table 10. Results of individual machine learning models.

Model name	Training duration	Accuracy	Precision	Sensitivity	F1
Logistic regression	2.8 s	0.765	0.823	0.762	0.792
Support vector machines	10.4 s	0.765	0.824	0.763	0.792
Decision tree	1.0 s	0.845	0.894	0.835	0.864
Random forest	1.3 min	0.876	0.891	0.898	0.895
XGBoost	1.4 min	0.879	0.890	0.905	0.898
Voting classifier	3.6 min	0.870	0.888	0.891	0.889
Stacked classifier	7.8 min	0.879	0.904	0.889	0.896

Source: own elaboration.

The best model in terms of quality measure values is XGBoost. Evidently inferior results were achieved by linear models, namely logistic regression and support vector machines. It is therefore probably better to focus on tree models in further research. It is worth noting that the difference in performance between XGBoost and the stacked classifier is small, however, XGBoost's training time is more than five times shorter. Such results suggest that further research should focus on tree-based models based on gradient enhancement.

No model has managed to exceed the 90 per cent performance threshold on the F1 metric, despite many iterations when looking for hyperparameters. It is likely that an expansion of the dataset with new variables, such as the political/religious affiliation of the terrorist group in question, the incidence of ethnic/political/religious tensions in the country, the distance of the incident site from the nearest hospital, would be needed to achieve results of around 95 per cent.

It was also decided to calculate the feature importance for the model based on the XGBoost library, in order to determine which features were most relevant to the model when making the prediction. This is presented in Chart 12, in which the y-axis shows the features described earlier and the x-axis shows the value of the feature for the prediction. The most important features were targets and weapon subtype. When evaluating contingency plans for a terrorist attack, these two factors should be considered particularly carefully. A set of geographical factors, such as region or longitude and latitude, for example, indicate that terrorism is a problem for some areas of the world and has different characteristics.




Chart 12. Key features for the XGBoost model.

Source: own elaboration.

Conclusions

A study of terrorist attacks occurring between 1970 and 2019 was conducted, including an in-depth analysis of trends in the activity of selected terrorist groups. The collaborative models built performed significantly better than linear models.

The results indicate that future focus should be on models based on collaborative learning, as linear models, such as logistic regression, performed noticeably worse. In particular, tree-based models based on gradient reinforcement should be looked at, as XGBoost achieved better results with shorter training times compared to the voting classifier and the stacked classifier. Further considerations could also investigate the effectiveness of models similar to XGBoost, such as CatBoost, LightGBM, or see how deep neural networks would handle such a task.

When the START consortium researchers publish the full data for 2020-2021, it would be worthwhile to analyse whether terrorist activity has changed, and to characterise their activities during the COVID-19 pandemic. Also, supplementing the GTD with new variables, such as, for example, information on ethnic or religious tensions, the economic condition of a country, or the religious or political affiliation of a given terrorist group, could positively influence the classification results, including helping to surpass the 90 per cent barrier.

This article shows that it is possible to use machine learning effectively in the field of security. This can assist the relevant authorities in developing crisis management plans in the event of a terrorist incident. The state can also take educational measures in the form of information campaigns or teaching in schools how to behave during a terrorist attack depending on the type of attack. As shown through the significance values of the variables, it was the target of the attack that was the most important factor for the XGBoost model, and as indicated earlier, civilians are most often attacked. With educational measures it would probably be possible to reduce the number of victims of terrorist attacks to some extent. The second most important factor turned out to be the subtype of weapons, so that after an in-depth analysis by the relevant services of the means used by particular terrorist groups in a given area, new regulations could be developed to make it more difficult for terrorists to obtain weapons.

Bibliography

Géron A., *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, Sebastopol 2019.

Matfess H., *Boko Haram: History and Context*, in: *Oxford Research Encyclopedia of African History*, Oxford University Press 2017.

Raschka S. et al., *Machine Learning with PyTorch and Scikit-Learn: Develop machine learning and deep learning models with Python*, Birmingham 2022.

Internet sources

Basque group Eta says armed campaign is over, BBC News, 20 X 2011, <https://www.bbc.com/news/world-europe-15393014> [accessed: 6 VI 2022].

Binnie I., *Basque separatist group ETA says it has “completely dissolved”*, Reuters, 2 V 2018, <https://www.reuters.com/article/us-spain-eta-idUSKBN1I31TP> [accessed: 6 VI 2022].

Codebook: Inclusion Criteria and Variables, Global Terrorism Database, August 2018, <https://www.start.umd.edu/gtd/downloads/Codebook.pdf/> [accessed: 30 V 2022].

Communist Part of the Philippines – New People’s Army, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/communist-party-philippines-new-peoples-army> [accessed: 11 VI 2022].

Data Collection Methodology, Global Terrorism Database, <http://www.start-dev.umd.edu/gtd/using-gtd/> [accessed: 21 V 2022].

Deaths in Maoist attacks down by 21%: Shah at CMs’ meeting, The Times of India, 27 IX 2021, <https://timesofindia.indiatimes.com/india/deaths-in-naxal-attacks-down-by-21-shah-at-cms-meeting/articleshow/86543018.cms> [accessed: 8 VI 2022].

History of the GTD, Global Terrorism Database, <https://start.umd.edu/gtd/about/History.aspx> [accessed: 11 V 2022].

Left Wing Extremism Division, Ministry of Home Affairs, https://web.archive.org/web/20220707070953/https://www.mha.gov.in/division_of_mha/left-wing-extremism-division [accessed: 8 VI 2022].

Pentagon confirms death of Somalia terror leader, The Washington Times, 5 IX 2014, <https://www.washingtontimes.com/news/2014/sep/5/pentagon-confirms-death-of-somalia-terror-leader/> [accessed: 10 VI 2022].

Revolutionary Armed Forces of Colombia (FARC), Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/revolutionary-armed-forces-colombia-farc> [accessed: 8 VI 2022].

Saffón S., *Peru in Familiar Stalemate With Shining Path Rebels*, InSight Crime, 4 IX 2020, <https://insightcrime.org/news/brief/peru-stalemate-shining-path/> [accessed: 8 VI 2022].

Sethi A., *Green Hunt: the anatomy of an operation*, The Hindu, 6 II 2010, <https://www.thehindu.com/opinion/op-ed/Green-Hunt-the-anatomy-of-an-operation/article16812797.ece> [accessed: 8 VI 2022].

The Afghan Taliban, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/afghan-taliban> [accessed: 8 VI 2022].

The Islamic State, Stanford University, <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/islamic-state> [accessed: 11 VI 2022].

Jakub Tuszyński

University of Warsaw graduate with a master's degree in Big Data Management.