

# Journal of Humanistic Mathematics

---

Volume 13 | Issue 2

July 2023

---

## Responsible Data Science for Genocide Prevention

Victor Piercey  
*Ferris State University*

Follow this and additional works at: <https://scholarship.claremont.edu/jhm>



Part of the [Data Science Commons](#), [Mathematics Commons](#), [Other International and Area Studies Commons](#), and the [Social Justice Commons](#)

---

### Recommended Citation

Victor Piercey, "Responsible Data Science for Genocide Prevention," *Journal of Humanistic Mathematics*, Volume 13 Issue 2 (July 2023), pages 64-85. DOI: 10.5642/jhummath.DNFZ3077. Available at: <https://scholarship.claremont.edu/jhm/vol13/iss2/6>

©2023 by the authors. This work is licensed under a Creative Commons License.

JHM is an open access bi-annual journal sponsored by the Claremont Center for the Mathematical Sciences and published by the Claremont Colleges Library | ISSN 2159-8118 | <http://scholarship.claremont.edu/jhm/>

The editorial staff of JHM works hard to make sure the scholarship disseminated in JHM is accurate and upholds professional ethical guidelines. However the views and opinions expressed in each published manuscript belong exclusively to the individual contributor(s). The publisher and the editors do not endorse or accept responsibility for them. See <https://scholarship.claremont.edu/jhm/policies.html> for more information.

---

## Responsible Data Science for Genocide Prevention

### Cover Page Footnote

The author would like to thank Carrie Diaz Eaton for all of the conversations that helped shape this and future work, Sadie Witkowski for providing communication workshops where some of the approaches to this paper took shape, the editors and referees for their thoughtful comments, and Catherine Buell and Rochelle Tractenberg for their collaboration on ethical mathematical practice.

# Responsible Data Science for Genocide Prevention

Victor Piercey

*Ferris State University, Big Rapids, Michigan, USA*

VictorPiercey@ferris.edu

---

## Abstract

The term “genocide” emerged out of an effort to describe mass atrocities committed in the first half of the 20<sup>th</sup> century. Despite a convention of the United Nations outlawing genocide as a matter of international law, the problem persists. Some organizations (including the United Nations) are developing indicator frameworks and “early-warning” systems that leverage data science to produce risk assessments of countries where conflict is present. These tools raise questions about responsible data use, specifically regarding the data sources and social biases built into algorithms through their training data. This essay seeks to engage mathematicians in discussing these concerns.

**Keywords:** data science, ethics, genocide, social justice, international affairs, political conflict.

---

## 1. The *Holodomor*

In the early 1930s, the Soviet Union embarked on the collectivization of agriculture, forcing individual farmers onto state collective farms. Those who resisted, and those simply accused of being a “rich farmer” (called a *kulak*), were either executed or sent to the Gulag. Resistance to collective farming and the state siphoning of grain “surpluses” (which often included the grain needed for subsistence) led to famine and general economic disaster in the Soviet agricultural sector in 1932-33 [18, 24].

The famine was particularly acute in Ukraine, where NKVD officers (the forerunner of the KGB) set up roadblocks to prevent Ukrainians from leaving Ukraine while

organized teams of policemen and [Communist] party activists, motivated by hunger, fear and a decade of hateful and conspiratorial rhetoric, entered peasant households and took everything edible: potatoes, beets, squash, beans, peas, anything in the oven and anything in the cupboard, farm animals and pets. [2, page *xxvi*].

By 1933, the famine in Ukraine had become so severe that cannibalism became widespread, including parents eating their own children [2, pages 256-261]. Overall estimates of the death toll range from 13% to 25% of the Ukrainian population [22, 28]. Ukrainians call this the *Holodomor*, which is Ukrainian for “death by hunger.”

Having studied survivor testimonies and archival evidence, Applebaum [2] concluded

[n]either crop failure nor bad weather caused the famine in Ukraine. Although the chaos of collectivization helped create the conditions that led to famine, the high numbers of deaths in Ukraine between 1932 and 1934, and especially the spike in the spring of 1933, were not caused directly by collectivization either. Starvation was the result, rather, of the forcible removal of food from people’s homes; the roadblocks that prevented peasants from seeking work or food; the harsh rules of the blacklists imposed on farms and villages; the restrictions on barter and trade; and the vicious propaganda campaign designed to persuade Ukrainians to watch, unmoved, as their neighbors died of hunger. [2, page 347].

Applebaum’s findings are consistent with those of a Congressional Committee [33], whose findings laid the blame for the famine at the policies and actions of the Soviet government. Applebaum is among those who consider the *Holodomor* to be a genocide. The most notable scholar to refer to the famine in Ukraine as genocide is Lemkin [20], who established the term itself.

What if the *Holodomor* occurred in an environment with today’s technology, today’s data science, and today’s international legal infrastructure? Suppose

that cell phones and computing technology existed and were as widespread in the 1930s as in 2023, that data science techniques and algorithms along with computing languages were present, and that the United Nations and all relevant treaties existed. The question is – would this have made a difference?

Maybe, and maybe not. But what would this have looked like?

Perhaps émigré groups would have escaped and notified authorities in the United Nations. Social media posts might have alerted authorities that something was happening in Ukraine, tracked by identifying relevant hashtags. Algorithms might be used to estimate the probability that what was occurring was a genocide and based on such algorithms the international community may have intervened.

Where would the data for these algorithms come from? They could be scraped from social media posts. They could also come from hacking Soviet government cell phones and computers, and possibly individual cell phones. NKVD officers could have been geolocated from hacked cell phone data. The Soviet government would object, arguing that they were targeted or profiled by the “imperialist” countries because they were a communist country. There would have been a debate in the international community about the role of human rights and state sovereignty, and whether émigré groups from Ukraine could be trusted.

This scenario paints a picture of how data science could be used to prevent genocide, but also raises concerns. Concerns include respecting privacy and data autonomy, the influence of social biases on the international stage, and the role of local communities and national sovereignty in the world order.

These concerns can be framed in terms of ethical practices. This essay raises these concerns as potential applications of ethical guidelines to the use of data science in genocide prevention. While of interest to professional data scientists, I hope this essay will be read by a general mathematical audience as an example of how ethical concerns should shape their actions if they move into data science work motivated by social justice.

## 2. Genocide in the Post-War World

### 2.1. *The United Nations and the Genocide Convention*

The term “genocide” was coined by Lemkin [19] in 1944 to describe what would come to be known as the Holocaust. In 1948, the United Nations General Assembly adopted the United Nations Convention on the Prevention and Punishment of the Crime of Genocide [35] (the “UN Genocide Convention”). The UN Genocide Convention went into force in 1951, and since its adoption, has been ratified by 152 countries.

Examples of genocides that have occurred since the UN Genocide Convention went into effect include mass killings in Cambodia between 1975 and 1979, El Salvador in 1981, and Rwanda in 1994. Historical events that predate the convention and are often considered genocide include the treatment of Native Americans by the United States in the 19<sup>th</sup> century, the treatment of Armenians in Turkey in 1916, and the *Holodomor* described above. Recent events that may be considered genocide include the treatment of the Rohingya in Myanmar (since 2016) and the Uighurs by China (since 2018).

The UN Genocide Convention defines “genocide” as:

any of the following acts committed with *intent to destroy, in whole or in part*, a national, ethnical, racial, or religious group, as such:

- (i) Killing members of the group;
- (ii) Causing serious bodily or mental harm to members of the group;
- (iii) Deliberately inflicting on the group conditions of life calculated to bring about its physical destruction in whole or in part;
- (iv) Imposing measures intended to prevent births within the group;
- (v) Forcibly transferring children of the group to another group.

[35, Article 2, emphasis added]

The term “intent” is critical. In criminal law, there not only has to be a criminal act, but a criminal intent to commit the crime.

The convention has rarely been enforced and has not yet functioned as a prevention tool. The only case in which a violation of the convention has been found was against Serbia and Montenegro at the International Court of Justice in 2007 [8]. There are currently two cases before the International Court of Justice based on the convention, one against Myanmar on behalf of the Rohingya, and another against Russia on behalf of Ukraine.

### *2.2. Framework for Atrocity Crimes*

As part of an effort to combat genocide, and in response to ethnic cleansing in the former Yugoslavia and the genocide against the Tutsis in Rwanda, the United Nations established the Office on Genocide Prevention and the Responsibility to Protect. This office developed a framework for analysis of atrocity crimes [34] (the “UN Framework”). Atrocity crimes is a collective term for genocide, war crimes, and crimes against humanity. The framework lists indicators intended to help identify atrocity crimes. Regarding genocide, the framework uses the term “protected groups” to refer to national, ethnical, racial, or religious groups (as specified in Article 2 of the UN Genocide Convention).

The UN Framework is divided into fourteen “risk factors” whose presence is determined by a set of indicators specific to each. The first eight risk factors are common to all three atrocity crimes. This is followed by two for genocide, two for crimes against humanity, and two for war crimes.

The risk factors are listed in Table 1. Some sample indicators, which we will return to below, are in Table 2.

### *2.3. A Sample Data Science Model*

By and large, the UN indicators (and those in other models) are binary. For these indicators to be useful, it would be helpful to have a way to process their values.

How would a data science model work and what role would these indicators play in prevention? A model could function like a car’s “check-engine” light.

<b>Common Risk Factors</b>
Risk Factor 1: Situations of armed conflict or other forms of instability
Risk Factor 2: Record of serious violations of international human rights and humanitarian law.
Risk Factor 3: Weakness of State structures
Risk Factor 4: Motives or incentives
Risk Factor 5: Capacity to commit atrocity crimes
Risk Factor 6: Absence of mitigating factors
Risk Factor 7: Enabling circumstances or preparatory action
Risk Factor 8: Triggering factors
<b>Genocide Risk Factors</b>
Risk Factor 9: Intergroup tensions or patterns of discrimination against protected groups
Risk Factor 10: Signs of intent to destroy, in whole or in part, a protected group
<b>Crimes Against Humanity Risk Factors</b>
Risk Factor 11: Signs of a widespread or systematic attack against any civilian population
Risk Factor 12: Signs of a plan or policy to attack any civilian population
<b>War Crimes Risk Factors</b>
Risk Factor 13: Serious threats to those protected under international humanitarian law
Risk Factor 14: Serious threats to humanitarian or peacekeeping operations

*Table 1: Risk factors in the United Nations Framework for Atrocity Crimes*

With sufficient conditions based on the indicators, a “check engine” light might come on as a warning. Those of us with experience owning cars know that we do not automatically do something to the car when the light comes on — we bring it to a professional to be examined. Someone reads the code, a mechanic might look at the engine and other appropriate parts of the car, recommend a course of action, and then it is up to the car owner to decide. At no point is human judgment abandoned. So it should be with an early-warning data science model for genocide.

For a first pass, logistic regression seems suited to the task of processing genocide indicators and flashing a warning light when appropriate. Logistic regression assumes a Bernoulli response variable and is used to estimate the probability of a “success” in a Bernoulli trial. In our case, the response



<b>Risk Factor 1: Situations of armed conflict or other forms of instability</b>
Indicator 1.7: Economic instability caused by scarcity of resources or disputes over their use or exploitation.
Indicator 1.9: Economic instability caused by acute poverty, mass unemployment, or deep horizontal inequalities.
Indicator 1.11: Social instability caused by exclusion or tensions based on identity issues, their perception, or extremist forms.
<b>Risk Factor 2: Record of serious violations of international human rights and humanitarian law.</b>
Indicator 2.2: Past acts of genocide, crimes against humanity, war crimes, or their incitement.
Indicator 2.8: Widespread mistrust in State institutions or among different groups as a result of impunity.
<b>Risk Factor 4: Motives or incentives</b>
Indicator 4.1: Political motives, particularly those aimed at the attainment or consolidation of power.
Indicator 4.7: Ideologies based on the supremacy of a certain identity or on extremist versions of identity.
Indicator 4.8: Politicization of past grievances, tensions, or impunity.
<b>Risk Factor 9: Intergroup tensions or patterns of discrimination against protected groups</b>
Indicator 9.2: Denial of the existence of protected groups or of recognition of elements of their identity.
<b>Risk Factor 10: Signs of intent to destroy, in whole or in part, a protected group</b>
Indicator 10.7: Expressions of public euphoria at having control over a protected group and its existence.

*Table 2: Sample indicators for risk factors 1, 2, 4, 9, and 10.*

variable would be the presence of a genocide (1 if so, 0 if not), and the regression would estimate the probability, conditioned on the values of the indicators, that an event is a genocide.

One drawback of logistic regression is that with rare events, such as genocide, they tend to carry a bias toward false negatives. That is, they are less likely to predict a genocide when one is in progress or likely to occur. There are modifications to logistic regression to account for this bias, such as “rare-events logistic regression” [15] and the use of “elastic-net regularization” with logistic regression [37]. These modified logistic regression techniques are

common in the literature on predicting and analyzing political violence and conflict. For example, Collier and Hoeffler [11] used rare-events logistic regression to predict the onset of civil wars. An example directly related to genocide prevention is “The Early Warning Project” [31], a system that uses logistic regression with elastic-net regularization. The model uses thirty features, several of which align with the indicators of the UN Framework.

Another issue with logistic regression is the underlying assumption of linearity. Muchlinski and colleagues [25] recommend a technique called random forests — using randomly selected subsets of the training data to produce a series of decision trees, then aggregating the results from each tree. Decision trees and random forests are non-parametric, meaning they do not require a functional form. The authors found that in addition to avoiding an assumption of a functional form, random forests performed better (in terms of reducing false negatives) than various forms of modified logistic regression. Muchlinski and colleagues tested random forest techniques in the prediction of civil war onsets, and their work has been applied to predicting other types of civil strife (see, e.g., [12, 29, 26]).

For a more thorough peek into what can operate behind a “check-engine” light for political conflict, see [16]. The authors describe a pipeline intended to be used for any type of political conflict, including civil strife and genocide. They tested their pipeline using the “targeted mass killing” dataset [7] for genocide and politicicide. The forecasting algorithm that they found applicable for this purpose was to combine the results of random forests, logistic regression, and a more general regression model in which the prediction variable is Bernoulli but the functional form can be any smooth function.

In the following, we will use the basic logistic regression algorithm as a stand-in for any appropriate technique. While it is important to use the strongest forecasting method possible, the discussion below does not depend on the choice of methodology. Logistic regression is a sufficiently simple “check-engine” light technology for our purposes.

As with a “check-engine” light, it is critical that the use of such a model is advisory only. We are not interested in having a model that acts like a credit score with automated decisions. Rather, the model should inform decisions about resource allocation and where to investigate, monitor, or look closer.

Put differently, we do not want the process to be *data-driven* — where the data dictates or “drives” the decisions — but rather *data-informed*.

### 3. Responsibilities and Ethical Issues

When building a model, it is important to think about unintended consequences. Even with the best of intentions, the implementation of a model could still cause harm. In addition, the model could be used for purposes that go beyond the original intention.

For the purposes of this essay, I raise two concerns: (1) the potential for violations of privacy and data autonomy, and (2) the risk of perpetuating social biases through algorithms. Thinking about the “check engine” light analogy, one might be concerned that the data feeding into the system may monitor your driving, and that the light may react differently in different car models but under similar conditions.

The concerns we raise align with results from a non-scientific study of the beliefs of the mathematics community about the ethical responsibilities of mathematicians. While the Association for Computing Machinery (ACM) and the American Statistical Association (ASA) each have their own codes of professional responsibility (see [3], and [5] respectively), mathematicians generally have very limited professional responsibility standards (see, e.g., [4, 23]). Since they are finding themselves interested in and doing data science work, it is important that mathematicians recognize or adopt more robust professional responsibility guidelines than they currently have in place.

To that end, Buell, Piercey, and Tractenberg [6] conducted a survey in which they modified items from the ACM and ASA codes and asked participants to determine whether the modified item constituted an ethical obligation of a mathematician. Table 3 contains a sampling of the items which garnered over 85% endorsement and are relevant to the concerns below. I refer to the items in Table 3, along with others with endorsement rates above 50%, as the “proto-guidelines.” In the discussion below, these items will be referred to by their item number.

Addressing ethical matters will ultimately require human judgment, high standards for action, and the careful balancing of harms and benefits.

Item Number	Item: “The ethical mathematics practitioner ...”	Percent “Yes”
7	“is candid about any known or suspected ... biases when working with data, methods, or models.”	97.9%
11	“conforms to ... any restrictions on [the data’s] use established by the data provider (to the extent legally required)”	97.6%
23	“... promotes the equal dignity and fair treatment of all people ...”	93.5%
32	“respects others and promotes justice, and inclusiveness in all work. Fosters fair participation of all people. Avoids and mitigates bias and prejudice. Does nothing to limit fair access.”	85.3%
33	“Minimizes the possibility of indirectly or unintentionally harming others ...”	85.3%

Table 3: Sample proto-guidelines with endorsement at or above 85%

Below I offer suggestions relating to the processes and algorithms that can be used to reduce harm. As one might expect, social scientists in conflict forecasting take their ethical responsibilities seriously. For example, the authors of [16] note that the pipeline they introduce will maximize true positives, which they think of as an ethical duty. However, expecting algorithms alone to address ethics problems is, as Maurice Chiodo said at a 2018 conference on ethics in mathematics, “like trying to put a fire out by pouring fire on it” [10]. In other words, there is a limit to what we can do with algorithms no matter how good our intentions are.

### 3.1. Privacy and Data Autonomy

Our first concern is for the privacy and ownership of the data collected for a genocide early-warning algorithm. Some data are clearly observable — such as anything published in a newspaper. Other data may not be and may be private.

This concern is most acute when considering data used to document intent. As noted above, the definition of genocide in the UN Convention requires an “intent to destroy, in whole or in part,” some protected group. UN indicators

such as 4.1 and 10.7, which hinge on intent and attitudes, may not be easy to observe since they require evidence of private thoughts and intentions.

A possible starting point might be surveying victims and survivors. Gatewood and Price [13], in describing how social network analysis can be used to analyze community resources for women in conflict zones, suggested collecting data for analysis through surveys. In conditions of ongoing conflict or genocide, it may be difficult to obtain survey responses. Depending on who is surveyed, responses may not be able to be taken at face-value. Methodologies to handle surveying populations that are difficult to reach and handling incomplete data include respondent-driven sampling, in which initial members surveyed hand “tickets” to others in the population to connect them with researchers [14], and multiple system estimation, in which researchers triangulate multiple populations and data sources with attention to observations that are common to more than one source [21]. These techniques are more useful in estimating the impact of a conflict such as genocide after the event concludes, rather than identifying ongoing genocide (see, e.g., [30]).

A potential data source for preventative action may be social media posts. The use of hashtags may help in identifying relevant data, although the language and the hashtags may be modified over time to avoid detection. Indeed, this occurs with child sexual exploitation, where deep learning algorithms are used to update search algorithms [17]. But beyond the practicalities of tracking social media posts, who owns those data? Who *should* own those data? Are they public? Is there an expectation of privacy? From a legal perspective, this might depend on the terms of service. But from an ethical perspective, the question is more complicated (see, e.g., [1]).

Documenting or discovering intent may require information from situations in which there is an unambiguous expectation of privacy, such as meetings or private conversations. These data might be collected, for example, through means such as phone hacking, satellite surveillance, and the use of intelligence assets. While it is not my purpose here to evaluate ethically the use of intelligence tradecraft in data and evidence gathering, those of us designing or using algorithms need to be mindful of the source of the data and whether or not the circumstances warranted such action.

In terms of the proto-guidelines above in Table 3, the item most directly applicable is item 11, which requires conformity with data-providers expressing restrictions on the use of their data. This item is limited to the extent “legally required.” On the other hand, items such as 23 and 32, which describe a broader concern for dignity and respect for people, could extend to respecting privacy beyond the limits imposed by law.

One way to reduce the potential for privacy violations is through identifying ethically observable correlates with intent, where by “ethically-observable” we mean observable without violating privacy or data autonomy. There are several genocide datasets available for training. The most recent such dataset is called the “Targeted Mass Killing” dataset (the “TMK dataset”), introduced in [7]. The TMK dataset contains 207 unique events. Among the features, there is a binary feature for the presence of evidence of stated intent, a binary feature for the presence of evidence of organizational intent (logistical operations or troop movements, for example), and the number of deaths. Based on these three features, a severity measure called a “TMK ordinal” is assigned to each event. This is a value between 1 and 8. An event with an ordinal at or above 4 (which requires at least 1000 annual deaths and evidence of either stated or organizational intent) is deemed a genocide (or a politicide, as politically identified groups are not protected under the UN Genocide Convention).

The TMK dataset can be used to explore what ethically-observable features correlate with the presence of genocide. One way to explore these data is through clustering. Using binary features, the metric induced by the  $\ell_1$  norm tells us how many features are different for two observations, and as such is a sensible way to measure how far apart the observations are. Using this metric with the events in the TMK database yields the cluster diagram in Figure 1. Here, each endpoint is an event, and we can see about five “clusters” of events — events that are close to one another. I am currently analyzing these clusters to determine how to characterize them.

Correlates can also be found by running logistic regression (possibly modified as described above), where the outcome is whether or not the TMK ordinal is 4 or larger (i.e., whether an event constitutes a genocide according to the standard in [7]). Then one can use backward selection — a process by




Figure 1: A cluster diagram using hierarchical clustering with the  $\ell_1$  norm on the TMK dataset. Diagram generated by the author in R.

which all features are initially used in the regression, and then the regression is rerun repeatedly, after removing the least significant feature, one at a time. A model can be selected by identifying the lowest of a measure called “Akaike Information Criterion” (or AIC). To illustrate what this looks like using only basic logistic regression, the R output for the selected model is displayed in Figure 2. In this case, the correlates are the nature of the actor perpetrating the action (government actor or non-government-actor that is connected to the government), whether there was foreign military intervention, and whether there was a triggering event.

Methodologies such as random forests can also be used to identify correlates, whether on their own or in combination with other algorithms. Note that [25] and [16] point out that their methodologies work not only for prediction and forecasting but also to find correlates.

This does not solve the problem completely. How was intent determined in [7]? For example, were logistic movements detected by satellite surveillance?

```

Call:
glm(formula = is.genocide ~ ., family = binomial(link = "logit"),
    data = tmk_sup_log_3)

Deviance Residuals:
    Min       1Q   Median       3Q      Max
-1.7774 -1.1363 -0.7114  1.2190  1.7309

Coefficients:
                Estimate Std. Error z value Pr(>|z|)
(Intercept)      -0.3235    0.3169  -1.021  0.30734
is.non.gov.actor.connected.to.gov  0.8193    0.4206   1.948  0.05142 .
is.government.actor      1.1475    0.3589   3.198  0.00139 **
mil.int                0.8530    0.3671   2.323  0.02016 *
trig                  -0.9214    0.3417  -2.696  0.00701 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

    Null deviance: 286.92  on 206  degrees of freedom
Residual deviance: 269.16  on 202  degrees of freedom
AIC: 279.16

Number of Fisher Scoring iterations: 4

```

Figure 2: The R output from running logistic regression on the TMK dataset using the best model chosen by backward selection. The resulting regression equation can be used to estimate that a probability that a TMK event is a genocide (or politicide) given the presence of any, all, or none of the indicators in the model. Generated by the author in R.

If so, are we comfortable using those data to identify correlates? Is it clear that the features identified in the logistic model are ethically observable?

Despite these questions about past data collection practices, this process, if successful and within acceptable parameters for inference, could reduce future harm by reducing the need for private data. This depends on whether the benefits of using the model outweigh the risk of using data collected by harmful means, a judgment that cannot be built into an algorithm.

### 3.2. Social Biases

Another concern is whether a genocide prediction algorithm encodes social biases about countries and regions. For example, of the countries in the TMK database, 81% are part of the global South. As another example, countries on a current Ongoing Mass Killing Watch List [31] are displayed in Figure 3.





Figure 3: World map (Mercator) showing countries in the Ongoing Mass Killing Watch List. Map generated by author using [mapchart.net](https://mapchart.net).

Figure 3 shows that every country on the list is in Africa or Asia. This may reflect biases or training data that incorporate social biases about countries in these regions – that they are “not capable” of self-government.

One might think broadening the scope beyond individual countries could help. Verdeja [36, page 19] suggests that genocide prevention may actually be more effective if focused on regions rather than individual countries. Examining this map, we note that outside of a few isolated countries, the regions indicated are Central Africa, East Africa, and the Indian Subcontinent, each of which may be subject to the social biases we are concerned with.

The potential for bias can be seen in the indicators and features in a given model. For example, in the UN Framework, indicators 1.7 and 1.9 reference economic instability related to national and individual poverty, reflecting a potential bias that poorer countries may be less suitable for self-government. UN Framework indicator 2.2 references past atrocity crimes, making it more likely that countries with problematic pasts will be monitored regardless of whether such actions are likely to continue. While monitoring may be warranted in the short-term, at some point it should cease. For example, the international community no longer monitors Germany for potential atrocity crimes.

Among the proto-guidelines discussed above in Table 3, item 7 specifically indicates an obligation to be candid about biases. This means, at a minimum, disclosing biases. However, item 32 asserts an obligation to avoid and

mitigate bias and prejudice. Although item 32 garnered less support than item 7 (85.3% versus 97.9%), it is still worth considering.

So, what can be done to mitigate bias? We could omit indicators that encode social biases. But is that enough? And what if those indicators are actually important and meaningful?

Another approach might be to introduce a “penalty term” that reduces the predicted probability that an event is a genocide if the event is occurring in a country that may be subjected to social bias.

To see what this could look like in the simplified example with basic logistic regression, the regression equation starts as:

$$\ln \left( \frac{\mu_i}{1 - \mu_i} \right) = \vec{x}_i^T \vec{\beta}$$

In this equation,  $\mu_i$  is the estimated probability that event  $i$  is a genocide,  $\vec{x}_i$  is the vector of the values of the features for event  $i$ , and  $\vec{\beta}$  is the vector of coefficients estimated from training data. Note that the first coordinate of  $\vec{x}_i$  will be 1, making the first coordinate of  $\vec{\beta}$  the intercept.

The crudest way to build a penalty into the model starts with making a list of countries that could be subjected to social bias and creating a variable  $y_i$  that is 1 if the country where the event is occurring is on the list and 0 otherwise. Then we let  $\lambda$  be a positive parameter reflecting the severity of the penalty selected. The new equation is:

$$\ln \left( \frac{\mu_i}{1 - \mu_i} \right) = \vec{x}_i^T \vec{\beta} - \lambda y_i$$

The value of  $\lambda$  can be chosen by performing a “tuning” process such as “ $k$ -fold cross validation” on training data. Solving for  $\mu_i$  shows that so long as  $\lambda > 0$ , if  $y_i = 1$  (i.e., the country is in the social bias list) then the predicted probability will be lower than it would be without the penalty.

This raises several questions. Does the imposition of a penalty effectively defeat the purpose of the process, making it nearly impossible to detect potential genocide? Are there structural reasons for correlations between the presence of potential genocide and the global South (see, e.g., [27, pages *xxix-xxviii*])? Can these be attended to without invoking a deficit perspective?

I do not suggest that I have answers to these questions, but rather that mathematicians working in this space consider them in their work.

#### 4. Where We Go From Here

All this is not to say that we should not use data science techniques to help inform genocide prevention. We should. This endeavor can help save many lives. Rather, this work requires care to do so responsibly. This means we mitigate harms in the data collection process by respecting privacy and data autonomy and that we are mindful of the potential to perpetuate social biases. There are likely a multitude of other ethical considerations, and we should all participate in those conversations.

It is also important to remain attentive to potential issues when working with partner communities. Working with partners from the community is an important part of social justice scholarship [9]. When it comes to genocide prevention and the application of data science models, this should be part of a collaboration with those with the lived experience relevant to understand and interpret what the data say and mean.

Such community partners may be émigré and refugee organizations. This could involve political opponents of the current regime. As such, it requires some caution. Those subjected to repression now could subject others to conflict and further repression in the future. For example, in Rwanda, the local response to the 1994 genocide ignited over a decade of conflict in Central Africa (see, e.g., [27]). This is not limited to Africa or the global South. The French Revolution led to the Reign of Terror, and the Russian Revolution replaced Tsarism with Bolshevism. While working with local communities is important and arguably necessary, we do not want to perpetuate a cycle of violence.

Finally, we need to be mindful that international action could mean violating a nation's sovereignty. While the details and alignments in the world order have shifted over time, the fundamental basis is that of national sovereignty, first established in the Treaty of Westphalia ending the Thirty Years War in 1648. This principle is implicit in the fact that the main sources for international law are treaties. In particular, the UN Genocide Convention is a treaty, and relies upon signatory states to enact their own legislation for enforcement [35, Article 6].

As mathematicians and data scientists, it is critical that we approach this work with humility. Genocide prevention is a difficult problem. In the end, this requires human decision and human action. Decisions can be difficult. Indeed, as Totten [32] mentions, even if there were a perfect and comprehensive early warning system, prevention would require overcoming the challenge of finding the political will to act. Any data science work in this space needs to be supportive and used to inform the process, not to drive it. In short, we should be mindful of a line in the movie *Thirteen Days* about the Cuban Missile Crisis, in which President Kennedy says, “[t]here’s something immoral about abandoning your own judgment.”

## 5. Why Me?

I would like to conclude by sharing why I am interested in this work.

One reason for my interest is that I was once a lawyer and am now a mathematician. As a lawyer, I studied international law and learned to look beyond formalisms to underlying scenarios. While my dissertation for my mathematics PhD was in algebraic geometry, my interest since has moved toward actuarial and data science, mostly motivated by my teaching. In addition, I have always been a student of history. Putting these interests together feels only natural.

But there are personal reasons for my interest as well. I am a Jewish-American, an inheritor of the losses that occurred in the Holocaust. It is difficult to watch antisemitism and other targeted hatred grow more common (or maybe just more visible) in my home country of the United States and feel powerless.

Beyond my heritage, I am also the last survivor of my immediate family. My father died of cancer in 1999 as did my mother in 2003, while my only brother Eric died in 2000 in an alcohol-related incident. At the time, Eric was in medical school, and I am convinced would have done great for the world had he lived. For me, if I can do work that is helpful, and especially if I can help prevent even one unnecessary death, then I can live up to the burden left for me by the loss of my family.

## References

- [1] Ali M. Al-Khouri, “Data ownership: Who owns ‘my data’?”, *International Journal of Management and Information Technology*, Volume **2** (2012), pages 1–8.
- [2] Anne Applebaum, *Red Famine: Stalin’s War on Ukraine*, Doubleday, New York NY, 2017.
- [3] Association for Computing Machinery (ACM), *Code of Ethics* (2018), available at <https://www.acm.org/about-acm/code-of-ethics>, last accessed on July 10, 2023.
- [4] American Mathematical Society (AMS), *Ethical Guidelines of the American Mathematical Society*, (2005, updated 2019 and 2022), available at <https://www.ams.org/about-us/governance/policy-statements/sec-ethics>, last accessed on July 10, 2023.
- [5] American Statistical Association (ASA), *ASA Ethical Guidelines for Statistical Practice*, (2022), available at <https://www.amstat.org/ASA/Your-Career/Ethical-Guidelines-for-Statistical-Practice.aspx>, last accessed on July 10, 2023.
- [6] Catherine Buell, Victor I. Piercey, and Rochelle E. Tractenberg, “Leveraging guidelines for ethical practice of statistics and computing to develop standards for ethical mathematical practice: A white paper” (2022), available at <https://arxiv.org/abs/2209.09311>, last accessed on July 10, 2023.
- [7] Charles Butcher, Benjamin E. Goldsmith, Sascha Nanlohy, Arcot Sowmya, and David Muchlinksi, “Introducing the targeted mass killing data set for the study and forecasting of mass atrocities,” *Journal of Conflict Resolution*, Volume **64** (2020), pages 1524–1547.
- [8] International Court of Justice, *Bosnia and Herzegovina v. Serbia and Montenegro*, International Court of Justice Reports of Judgements, Advisory Opinions, and Orders, Number 91, (2007), available at <https://www.icj-cij.org/case/91/judgments>, last accessed on July 10, 2023.
- [9] Alice K. Johnson Butterfield and Tracy M. Soska, “University-community partnerships: An introduction,” *Journal of Community Practice*, Volume **12** (2004), pages 1–11.

- [10] Maurice Chiodo, Opening Address, *Ethics in Mathematics Conference 1*, Cambridge, UK, April 2018. Video available at [https://www.youtube.com/watch?v=3e-E\\_SYK8Dg&t=7s](https://www.youtube.com/watch?v=3e-E_SYK8Dg&t=7s), last accessed on July 10, 2023.
- [11] Paul Collier and Anke Hoeffler, “Greed and grievance in civil war,” *Oxford Economic Papers*, Volume **54** Number 4 (2004), pages 563–595.
- [12] Mahdi Fawaz and Erwan LeQuellec, “Indirect rivalries and civil wars: Empirical evidence,” *Defence and Peace Economics* (2022), pages 1–28. doi:[10.1080/10242694.2022.2129350](https://doi.org/10.1080/10242694.2022.2129350)
- [13] James R. Gatewood and Candice R. Price, “Utilizing social network analysis to study communities of women in conflict zones,” *Journal of Humanistic Mathematics*, Volume **7** Issue 1 (2017), pages 3–21. doi:[10.5642/jhummath.201701.03](https://doi.org/10.5642/jhummath.201701.03)
- [14] Krista J. Gile, Isabelle S. Beaudry, Mark S. Hancock, and Miles Q. Ott, “Methods for Inference from Respondent Driven Sampling Data,” *Annual Review of Statistics and Its Application*, Volume **5** (2018), pages 65–93.
- [15] Gary King and Langche Zheng, “Logistic regression in rate events data,” *Political Analysis*, Volume **9** Number 2 (2001), pages 137–163.
- [16] Alfred Krzywicki, David Muchlinski, Benjamin E. Goldsmith, and Arcot Sowmya, “From academia to policy makers: A methodology for real-time forecasting of infrequent events,” *Journal of Computational Science*, Volume **5** (2022), pages 1489–1510.
- [17] Hee-Eun Lee, Tatiana Ermakova, Vasilis Ververis, and Benjamin Fabian, “Detecting child sexual abuse material: A comprehensive survey,” *Forensic Science International: Digital Investigation*, Volume **34** (2020). doi:[10.1016/j.fsidi.2020.301022](https://doi.org/10.1016/j.fsidi.2020.301022)
- [18] Moshe Lewin, “The disappearance of planning in the plan,” *Slavic Review*, Volume **32** (1973), pages 271–287.
- [19] Raphael Lemkin, *Axis Rule in Occupied Europe: Laws of Occupation, Analysis of Government, Proposals for Redress*, Columbia University Press, New York NY, 1944.

- [20] Raphael Lemkin, “Soviet Genocide in the Ukraine,” unpublished talk, 1953. Text available in Raphael Lemkin Papers, The New York Public Library, Manuscripts and Archives Division, Astor, Lenox and Tilden Foundations, Raphael Lemkin ZL-273. Reel 3. Also available online at <https://willzuzak.ca/tp/holodomor2013/oliver20171004Lemkin.pdf>, last accessed on July 10, 2023.
- [21] Kristian Lum, Megan Emily Price, and David Banks, “Applications of multiple-systems estimation in human rights research,” *The American Statistician*, Volume **67** (2013), pages 191–200.
- [22] Cheryl A. Madden, “The Holodomor: 1932–1933,” *Canadian-American Slavic Studies*, Volume **37** (2003), pages 13–26.
- [23] Mathematical Association of America (MAA). *Welcoming Environment, Code of Ethics, and Whistleblower Policy* (2017), available at <https://www.maa.org/about-maa/policies-and-procedures/welcoming-environment-code-of-ethics-and-whistleblower-policy>, last accessed on July 10, 2023.
- [24] James R. Millar, “Mass collectivization and the contribution of Soviet agriculture to the first five-year plan: A review article,” *Slavic Review*, Volume **33** (1974), pages 750–766.
- [25] David Muchlinski, David Siroky, Jingrui He, and Matthew Kocher, “Comparing random forest with logistic regression for predicting class-imbalanced civil war onset data,” Volume **24** (2016), Number 1, pages 87–103.
- [26] Jonathan Pinckney and Babak RezaeeDaryakeneri, “When the levee breaks: A forecasting model of violent and nonviolent dissent,” *International Interactions*, Volume **49**, Number 5 (2022), pages 997–1026.
- [27] Gérard Prunier, *Africa’s World War: Congo, the Rwandan Genocide, and the Making of a Continental Catastrophe*, Oxford University Press, Oxford UK, 2009.
- [28] Omelian Rudnytskyi, Nataliia Levchuk, Oleh Wolowyna, Pavlo Shevchuk, and Alla Kovbasiuk, “Demography of a man-made human catastrophe: The case of massive famine in Ukraine, 1932–33,” *Canadian Studies in Population*, Volume **42** (2015), pages 53–80.

- [29] Ali Hamza Syed, *Disentangling the Relationship Between Remittances and Peace*, M.A. Thesis, University of Chicago, 2022.  
doi:[10.6082/uchicago.4340](https://doi.org/10.6082/uchicago.4340)
- [30] Romesh Silva, Jasmine Marwaha, and Jeff Klingner, *Violent Deaths and Enforced Disappearances During the Counterinsurgency in Punjab, India*, Human Rights Data Analysis Group, 2009. Report available at <https://ensaaf.org/publications/reports/descriptiveanalysis/>, last accessed on July 10, 2023.
- [31] The United States Holocaust Memorial Museum and Dartmouth College, *Early Warning Project: Ongoing Mass Killing*, available at <https://earlywarningproject.ushmm.org/ongoing-mass-killing>, last accessed on July 10, 2023.
- [32] Samuel Totten, *The Prevention and Intervention of Genocide: An Annotated Bibliography*, Routledge Publishing, Abingdon UK, 2006.
- [33] US Congress and Commission on the Ukraine Famine, *Investigation of the Ukrainian Famine*, 1988.
- [34] United Nations, *Framework of Analysis for Atrocity Crimes - A tool For Prrevention*, 2014, available at <https://www.refworld.org/docid/548afd5f4.html>, last accessed on July 10, 2023.
- [35] UN General Assembly, *Convention on the Prevention and Punishment of the Crime of Genocide, 9 December 1948*, United Nations, Treaty Series, Volume **78** (1948),, page 277, available at <https://www.refworld.org/docid/3ae6b3ac0.html>, last accessed on July 10, 2023.
- [36] Ernesto Verdeja, “Predicting genocide and mass atrocities,” *Genocide Studies and Prevention: An International Journal*, Volume **9** (2016), pages 13–32.
- [37] Hui Zou and Trevor Hastie, “Regularization and variable selection via the elastic net,” *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, Volume **67**, Number 2 (2005), pages 301–320.