Project Jetson: Predicting Migration Patterns During the Somali Conflict

Alina Holmstrom

Introduction

In 2011, UNHCR refugee camps in Ethiopia were scrambling. Neighboring Somalia was experiencing a severe drought, which, in combination with lasting conflict, left Somalia's people impacted by the worst famine the country had seen in 25 years (UNHCR, 2019). Untenable living conditions drove hundreds of thousands of Somalis to seek refuge in neighboring countries, desperate to save themselves and their families.

Unprepared for the influx of people, UNHCR refugee camps were left struggling to meet the needs of existing camp residents as well as the nearly 2,000 daily arrivals. Humanitarian workers, perpetually stretched thin during regular times, were rushing to process new arrivals, provide food to starving children and parents, and ensure everyone had a place to sleep. The Dollo Ado camp team recalls spending their entire yearly budget for nutritional screening and protection support within a few short months (UNHCR, 2019). Eventually the crisis abated, but in 2017 yet another drought seized the country. Haunted by their experience in 2011, teams on the ground requested UNHCR's assistance in developing a method to predict migration flows so refugee camp workers could ensure adequate staff and resources to meet the needs of every individual that came through their doors.

Understanding the urgency of the situation, UNHCR began developing Project Jetson, a machine learning program that utilizes supervised machine learning to predict Somali refugee and IDP migration patterns up to one month in advance.

Background

Since the Cold War, countries, international organizations, and researchers have been attempting to predict migration patterns. Increased duration of refugee status and protracted conflicts led to states' desires to predict migration patterns, some with the intent to ensure a proper welcome for refugees and others with the intent of keeping refugees out. Early prediction methods took the form of surveys, conflict analysis, and modelling using quantitative data (Schmeidl, 2004). These methods were often labeled more generally as early warning systems, or the close monitoring of certain information to produce short-term predictions of impending mass migration (OECD, 2019). Old-school early warning systems factored in the complexity of migration and identified several drivers of migration.

These drivers were separated into 3 main categories: root causes, proximate causes, and intervening factors (Schmeidl, 2004; Clark, 1989). Root causes include historical, cultural, or other long-term impacts, such as border disputes or inter-religious conflicts, that may impact migration (Schmeidl, 2004; Clark, 1989). Proximate causes include current social, political, economic, or other conditions that combine with root causes to influence migration (Schmeidl, 2004; Clark, 1989). Intervening factors include any other events, not included in the first two groups, that may influence migration, such as climate, transportation availability, or destination country perceptions (Schmeidl, 2004; Clark, 1989).

These factors were utilized to create models for predicting migration patterns. Common models included threshold models, used to identify a threshold of human rights violations that would push people to migrate, and sequential models, which attempted to prove conflicts followed a sequence of events leading to migration (Schmeidl, 2004). A third type of model, pattern-recognition, became popular when computers started being developed with the capacity and complexity needed to use supervised machine learning to learn all possible combination of factors that have led and could lead to migration (Schmeidl, 2004). While somewhat effective, these models were rudimentary and lacked the detail and precision necessary to accurately predict future migration (Schmeidl, 2004).

Today, there are 3 main types of prediction methods: early warning systems, foresight, and forecasting. Early warning systems monitor qualitative and quantitative data in real time to create short-term predictions of migration patterns in crisis contexts (Future Migration Trends, 2020; OECD, 2019). Foresight refers to using qualitative scenario methods to predict future migration trends (Future Migration Trends, 2020; OECD, 2019). Forecasting refers to predicting medium or long-term migration trends using traditional quantitative modeling methods (Future Migration Trends, 2020; OECD, 2019). Given the differences between methods, predictions vary greatly. Even within the same prediction method, predictions often change based on differences in factors considered, prediction time period, assumptions, differing weights given to factors, and overall methodology (Future Migration Trends, 2020).

Given the differences in context, there is no standard prediction formula or method that can be applied to every migration-producing event. In fact, many experts argue that in order to effectively predict migration trends, models should focus on specific migration situations (Future Migration Trends, 2020; OECD, 2019; Suleimenova, 2017; Schmeidl, 2004). Project Jetson falls under the forecasting method.

Technical Aspects of Project Jetson

Project Jetson uses supervised machine learning and forecasting methods to predict future migration trends. Supervised machine learning refers to training a computer to recognize data types using labeled data (Turner, 2018). More specifically, Jetson uses a type of supervised machine learning called regression. In a regression, the computer uses labeled data fed to it to make predictions about the future (Turner, 2018). Examples of regression supervised machine learning used in daily life include algorithms predicting the price of a house based on its features, or weather apps predicting the weather based on past weather trends. Jetson uses past data such as climate patterns and rates of violence to make predictions about future migration trends.

Jetson was created by a UNHCR team of data scientists, AI engineers, and regional experts, including input from Somali refugees themselves (UNHCR, 2022). Using an earlier UNHCR initiative, European Winter Cell, as a blueprint for Jetson's open data sharing and prediction models, the team was able to create the basis of Jetson's programming (UNHCR, 2019). With their existing knowledge of the region and accompanying migration factors, the team created a list of independent and dependent variables which it used to "teach" the program. Ten independent variables were selected, including rates of violent conflict, rain and water levels, historic migration patterns, and market prices (UNHCR, 2022; IOM, 2022). The dependent variable was forced displacement both within and outside of Somalia (UNHCR, 2022). One of the most striking qualities of Project Jetson is the way in which it collects this data. As mentioned before, accurate, timely, and comprehensive data is crucial in making accurate predictions. To get this data, UNHCR leveraged its partnerships with other organizations and UN agencies to employ the use of open data (UNHCR, 2022). Open data includes data shared on partners' websites, joint research efforts between UN agencies, or data that was shared through a data-sharing agreement (UNHCR, 2022). Both individual-level and regional-level data and trends are used to create regional-level predictions of migration. In compliance with UNHCR's data protection policies, the data collected from partners is anonymized and aggregated per month and per region to ensure data and privacy protection (UNHCR, 2022; UNHCR, 2019). To ensure relevance of the data, data is updated monthly (UNHCR, 2022).

Results of Jetson's algorithm are published on UNHCR's Project Jetson website in the form of an interactive map. The map shows Somalia divided into regions and color-coded by the predicted level of displacement. A picture of Jetson's most recent map can be found in Exhibit A. The same site shows Jetson's predictions from April 2018 through April 2019 in the form of an interactive chart, which allows you to select a region and view Jetson's predictions and how they compare to the actual number of people who migrated (UNHCR, 2022; UNHCR, 2019). A picture of the interactive chart can be found in Exhibit B.

Successes and Challenges

UNHCR defines success as Project Jetson's ability to "predict the number of arrivals (IDPs and refugees); improve research for protective purposes; [and] develop a methodology for forced migration predictions" (UNHCR, 2022). According to UNHCR's website, Jetson's predictions were accurate for eleven of Somalia's eighteen regions as of 2018. This enabled UNHCR to preemptively increase or reappropriate resources to different camps depending on the predictions. The ability to appropriate resources in advance likely saved UNHCR tens of thousands, if not hundreds of thousands, of dollars, and enabled the organization to provide food and shelter for thousands of individuals that otherwise would have been left with limited or even no access to these resources. By contrast, when Jetson did not make accurate predictions, UNHCR likely lost funds in misappropriation. It could be argued, though, that UNHCR would likely have lost funds without Jetson, so they could only gain from the implementation of this technology.

By UNHCR's metrics, Jetson is marked a success. However, there are three main challenges to Project Jetson's approach: reliability, generalizability, and privacy concerns. While Jetson's predictions may have been accurate for the majority of the country's regions, lack of consistency has led to Jetson's predictions having not been published since April of an unspecified year. As seen in Exhibit A, Jetson's map comes with a yellow warning text at the bottom alerting users that the map information is outdated. The Project Jetson team admitted to challenges in acquiring the data needed to make accurate predictions, as data is often limited, delayed, or completely missing from active conflict zones or areas with no humanitarian access (UNHCR, 2022). When timely and comprehensive data is crucial for making accurate predictions, how well does a tool like Jetson fair in practice over the long term? Potentially not well. Jetson's lack of consistent updates points to a larger performance failure of the program. Programs like Jetson might work well over the short term, but in protracted conflicts with long-term migration effects, as seen in Somalia, a different approach may be necessary.

Regardless of the lack of consistency, what can other organizations or states learn from Project Jetson? How generalizable is this tool? In other words, how much of Jetson can be replicated in different contexts? In line with expert opinions on migration prediction techniques, Jetson was designed specifically for the Somali context. The variables selected, the data used to generate predictions, and the data used to create the tool are all specific to Somalia and the Somali conflict. While the weights in the algorithm could be adjusted for other contexts, the data used to train the model and the expert input into what factors are included would need to be redone for every different context. The time-intensive tasks of collecting data, identifying variables, assigning weights to those variables, testing and retesting the program (though the testing process may be replicable), and creating a data sharing network to collect future data from which predictions can be made would constitute a significant drain on resources for successful replication. These are all intensive tasks requiring time and human resources that are often limited or unavailable in times of crises, which is when prediction technology like Jetson is needed the most. Is this a useful endeavor? Is it worth putting this amount of time into a program just to have to start over again in a different context? In the case of protracted conflict situations, one would hope that this process would only need to be done once with the goal of continued prediction accuracy for the duration of the conflict. Given Jetson's limitations in its long-term prediction generation capabilities, the amount of time and data needed to create and sustain such a tool may not be worth the predictions it generates.

A more effective method may be to apply this predictive model outside of conflict contexts. As mentioned before, predictive migration models have been successfully deployed by states to predict migration inflows for decades. Jetson's model could be applied to the U.S. context, for example, to predict the migration flows from South America to ensure adequate resource allocation to border facilities. It is worth noting, however, that use of Jetson technology in this way could spark additional concerns surrounding the potential use of this technology to allocate resources to block individuals from entering rather than to accommodate them.

UNHCR paints Project Jetson as a success not only for its predictive abilities, but also as evidence of what open data sharing within the humanitarian sphere can produce (UNHCR, 2022). Jetson's early predictions serve as evidence that open data sharing between organizations results in a more robust pool of data that can mitigate some of the issues surrounding lack of data access in conflict zones. If UNHCR did not have data for a certain indicator in a region, a partner agency may have that data, and through the open sharing network established by Jetson UNHCR was able to access this data with the timeliness necessary to be able to use it. This would not be possible in a closed data environment. In addition to these positives, open data sharing comes with its own challenges in the form of privacy concerns. While open data sharing often makes use of publicly available data, some of the data is sourced through private bilateral networks or through legal data sharing agreements. This brings up data privacy concerns surrounding confidentiality and consent of data sharing. Did the refugees whose data is so critical to Jetson's operations agree to their data being circulated between agencies? What about the potential for data breaches, given the open nature of the data sharing network? UNHCR takes measures to protect data privacy by anonymizing and aggregating data received, but the original data is still stored on their network. It is likely that data privacy concerns are unavoidable when dealing with the open data sharing networks that are necessary for the success of programs like Jetson, but do the benefits outweigh the risks?

Implications

Given Project Jetson's successes and challenges, it is safe to say that the AI migration prediction technology code for humanitarian organizations has yet to be cracked. The ever-changing nature of conflicts leads to unpredictability that mitigates humanitarian organizations' ability to accurately and consistently predict migration trends as compared to states, who have been able to do so with general success due to the more stable nature of their long-term migration trends. Can humanitarian organizations overcome these challenges and create migration prediction models that are sufficiently replicable, reliable, and secure? Do the benefits of doing so outweigh the risks? If not, other actors such as states or NGOs may consider using Project Jetson as a model to predict migration flows outside of conflict contexts, where the availability of data and time constraints may pose less of an issue.

Appendix

Exhibit A: Project Jetson Interactive Map of Predictions

Source: UNHCR Project Jetson website: https://jetson.unhcr.org/tech.html

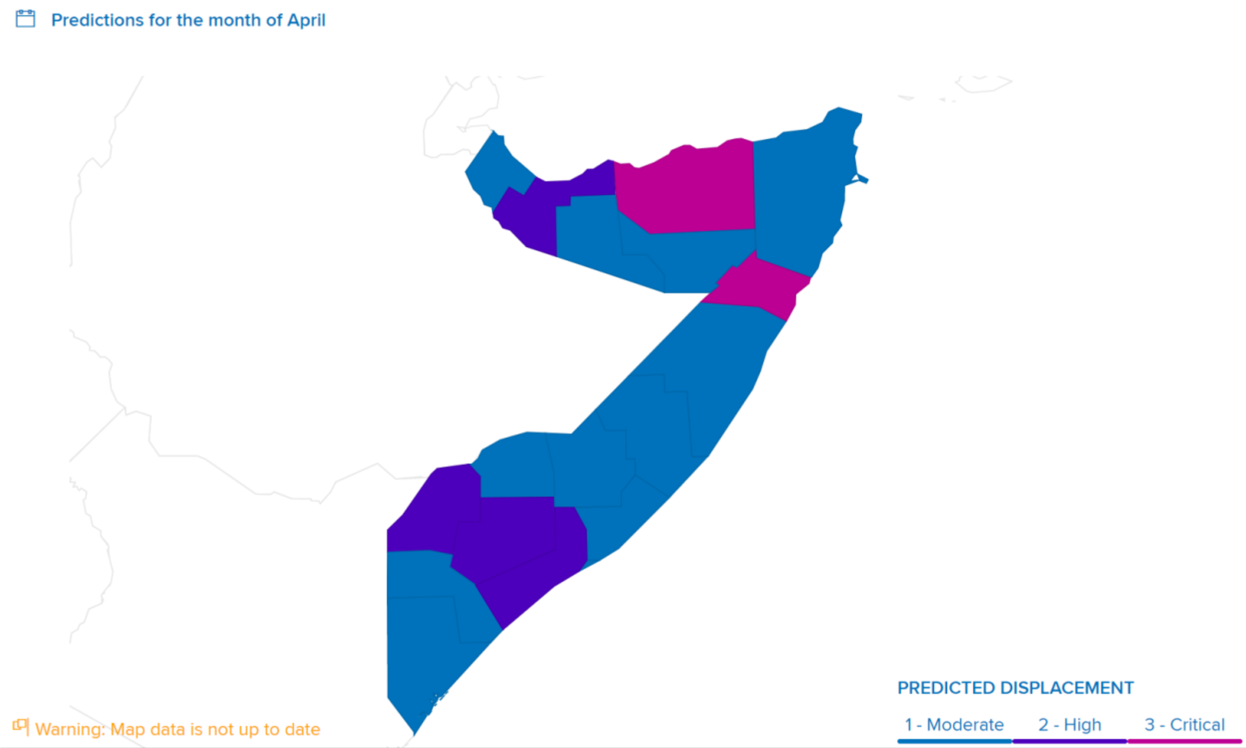

Exhibit B: Project Jetson Interactive Chart Predictions

UNHCR Project Jetson website: https://unhcrinnovation.shinyapps.io/Somalia/

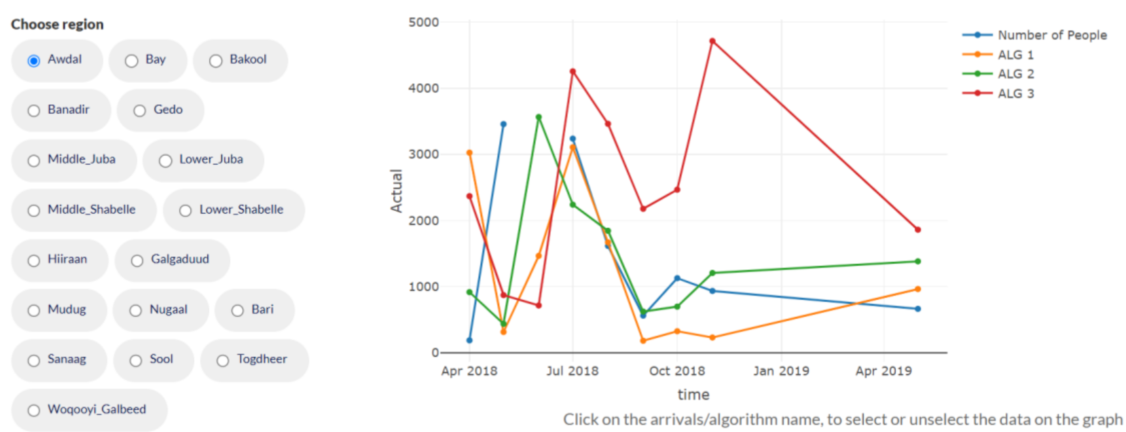

Exhibit C: Project Jetson Testing Results

Project Jetson: Predicting Forced Displacement in Somalia

Rebeca Moreno Jiménez[1], Sofia Kyriazi[1], and Katherine Hoffmann Pham[2,3]

[1] UNHCR Innovation [2] NYU Stern School of Business [3] United Nations Global Pulse

Motivation

Predicting arrivals of refugees and Internally Displaced Persons (IDPs) is of critical interest in humanitarian emergencies. Field operations teams are interested in anticipating these arrivals in order to **provide adequate assistance** in the form of food, shelter, and protection-related services.

Objectives

We identify three types of information needs in humanitarian emergencies: (1) description, (2) prediction, and (3) simulation. We focus on **predicting** arrivals by region 1 and 3 months in advance.

Context

- **70.8m** people displaced worldwide (58% IDP)
- **2.6m** IDPs in Somalia alone (21% of population)

Source: UNHCR 2019, OCHA 2019

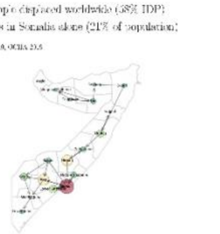

Figure 1: Inflows averaging at least 100 people per month

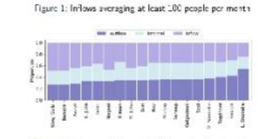

Figure 2: Inflows, outflows and within-region displacement

Research Questions

- **How best to forecast** displacement across 18 regions of Somalia using historical data?
- **What level of performance** is achievable?
- **How does performance vary** w/ forecast horizon, # features, # historical observations?

Data

We aggregate the data at the level of region and month. We observe 107 months' worth of data on 18 regions, covering the period from January 2011 through November 2019. We collect a diverse set of political, economic, environmental, health, and geographic features.

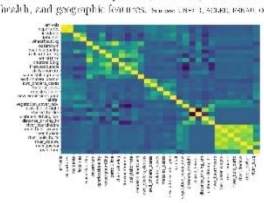

Figure 3: Correlations between variables in the dataset

Practical Challenges

- Integrating/updating data from **various sources**
- Target variable has missing values (57% in Sanaag)
- Average displacement varies by region (~600-13,000)
- Variability is **growing over time**
- Recording of arrivals can be **unreliable**

Models

Data through December 2018 is used for training, and data from 2019 forms the test set. We use 10-fold time series cross-validation for tuning the regression models. We test:

- **Naïve models:** lags, historical means, exponentially weighted means, expanding means
- **Classic regression models:** lasso/ridge regressions, decision trees, random forests, XGBoost, AdaBoost
- **Neural network models:** perceptrons, LSTM

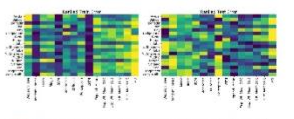

Figure 4: Ranked error by model on the train and test set (1 month)

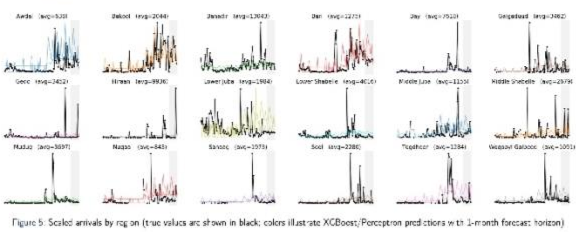

Figure 5: Scaled arrivals by region (true values are shown in black; colors illustrate XGBoost/Perceptron predictions with 1-month forecast horizon)

Work in Progress

- Further **tune** models to improve performance
- Explore **sub-regional** and **gravity models**
- Add additional models (e.g. econometric)
- Reframe as **classification** problem (red flags)

Preliminary Results

The best 1 month RMSE achieved was 18,021, relative to a mean of 3,555 arrivals, a standard deviation of 11,278, and a maximum of 235,069. Currently, the models struggle with overfitting and naive benchmarks are competitive. Interestingly, we don't see performance declining much with the forecast horizon.

Model	Test - 1mo	Train - 1mo	Test - 3mo	Train - 3mo
Perceptron	**18,021**	11,423	18,229	11,508
Expand. Mean	16,083	12,336	18,319	12,871
Ridge	16,271	11,065	18,376	11,366
Lasso	16,328	11,043	18,391	11,182
Exp. Wt. Mean	16,341	11,776	19,015	12,920
Random Forest	16,388	9,624	18,851	9,087
XGBoost	16,454	9,479	**18,027**	11,700
SVM	16,460	11,043	18,264	11,171
Decision Tree	16,544	11,432	19,321	9,662
LSTM	16,539	**4,905**		
Hist. Mean	16,654	12,167	19,269	12,929
AdaBoost	16,571	8,265	18,292	**9,147**
Lag	25,349	12,033	26,349	12,383

Table 1: Root Mean Square Error (RMSE) by model

Implementation

We have designed a dashboard for exploring the data and models, and plan to hand the project over to the regional team for long term development. Open questions include:

1. What level of performance should be required before release? How to quantify/communicate **uncertainty**?
2. How to empower users to explore and interrogate? Can we use **explainability** tools to aid understanding?
3. When should a **warning** be triggered? What predicted % increase in arrivals is critical?

As we learn more about the problem, we are developing a broader literature review on forecasting displacement. We are also supporting discussions of how predictive analytics can be applied in other critical contexts (e.g. the Sahel).

References

"A Goat Story." Web log. *Medium* (blog). UNHCR Innovation Service, May 8, 2019. https://medium.com/unhcr-innovation-service/a-goat-story-3ed6bdd2b237.

Fournier-Tombs, Eleonore. "Towards a United Nations Internal Regulation for Artificial Intelligence." *Big Data & Society* 8, no. 2 (2021): 205395172110394. https://doi.org/10.1177/20539517211039493.

"Future Migration Trends." Migration data portal. IOM, October 20, 2020. https://www.migrationdataportal.org/themes/future-migration-trends#recent-trends.

Lin, Larry, Kathleen M. Carley, and Shih-Fen Cheng. "An Agent-Based Approach to Human Migration Movement." *2016 Winter Simulation Conference (WSC)*, 2016. https://doi.org/10.1109/wsc.2016.7822380.

OECD. "Anticipating, Monitoring and Reacting to Inflows of Refugees and Other Vulnerable Migrants." *Ready to Help?*, 2019, 83–87. https://doi.org/10.1787/9789264311312-8-en.

Raymer, James. "The Estimation of International Migration Flows: A General Technique Focused on the Origin-Destination Association Structure." *Environment and Planning A: Economy and Space* 39, no. 4 (2007): 985–95. https://doi.org/10.1068/a38264.

Schmeidl, Susanne. "The Early Warning of Forced Migration: State or Human Security" Chapter. In *Refugees and Forced Displacement: International Security, Human Vulnerability, and the State*, 130–55. New Delhi: United Nations University, 2004.

Suleimenova, Diana, David Bell, and Derek Groen. "A Generalized Simulation Development Approach for Predicting Refugee Destinations." *Scientific Reports* 7, no. 1 (2017). https://doi.org/10.1038/s41598-017-13828-9.

UNHCR. "Home · UNHCR/Jetson Wiki." GitHub. UNHCR, 2019. https://github.com/unhcr/Jetson/wiki.

UNHCR Innovation Service. "Is It Possible to Predict Forced Displacement?" Web log. *Medium* (blog). UNHCR Innovation Service, May 13, 2019. https://medium.com/unhcr-innovation-service/is-it-possible-to-predict-forced-displacement-58960afe0ba1.

UNHCR. "Project Jetson." Project Jetson, 2022. https://jetson.unhcr.org/.

UNHCR. "Project Jetson." UNHCR Innovation. UNHCR, October 27, 2021. https://www.unhcr.org/innovation/project-jetson/.

"United Nations High Commissioner for Refugees (UNHCR)." AI for Good. UNHCR, December 14, 2021. https://aiforgood.itu.int/about-ai-for-good/un-ai-actions/unhcr/.

"Using Artificial Intelligence to Model Displacement in Somalia • UN Global Pulse." UN Global Pulse. Pulse Lab New York, April 23, 2020. https://www.unglobalpulse.org/project/using-artificial-intelligence-to-model-displacement-in-somalia/.

IOM. "'Project Jetson': An Experiment for Predicting Movements of Displaced People in Somalia Using Machine Learning." Migration data portal, 2022. IOM. Accessed September 26, 2022. https://www.migrationdataportal.org/data-innovation/Project-Jetson-experiment-predicting-movements-displaced-Somalia-machine-learning.

Turner, L. (2018, June 27). *Machine learning: A Primer*. Medium. Retrieved October 9, 2022, from https://medium.com/@lizziedotdev/lets-talk-about-machine-learning-ddca914e9dd1.