A Machine Learning Approach to Modeling Human Migration

Caleb Robinson

School of Computational Science and Engineering
Georgia Institute of Technology
Atlanta, Georgia
dcrobins@gatech.edu

Bistra Dilkina

Department of Computer Science
University of Southern California
Los Angeles, California
dilkina@usc.edu

ABSTRACT

Human migration is a type of human mobility, where a trip involves a person moving with the intention of changing their home location. Predicting human migration as accurately as possible is important in city planning applications, international trade, spread of infectious diseases, conservation planning, and public policy development. Traditional human mobility models, such as gravity models or the more recent radiation model, predict human mobility flows based on population and distance features only. These models have been validated on commuting flows, a different type of human mobility, and are mainly used in modeling scenarios where large amounts of prior ground truth mobility data are not available. One downside of these models is that they have a fixed form and are therefore not able to capture more complicated migration dynamics. We propose machine learning models that are able to incorporate any number of exogenous features, to predict origin/destination human migration flows. Our machine learning models outperform traditional human mobility models on a variety of evaluation metrics, both in the task of predicting migrations between US counties as well as international migrations. In general, predictive machine learning models of human migration will provide a flexible base with which to model human migration under different what-if conditions, such as potential sea level rise or population growth scenarios.

CCS CONCEPTS

• **Social and professional topics** → **Sustainability**; • **Computing methodologies** → *Modeling methodologies*; Supervised learning by regression;

KEYWORDS

migration modeling, machine learning, computational sustainability

ACM Reference Format:

Caleb Robinson and Bistra Dilkina. 2018. A Machine Learning Approach to Modeling Human Migration. In *COMPASS '18: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, June 20–22, 2018, Menlo Park and San Jose, CA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3209811.3209868>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

COMPASS '18, June 20–22, 2018, Menlo Park and San Jose, CA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5816-3/18/06...\$15.00

<https://doi.org/10.1145/3209811.3209868>

1 INTRODUCTION

Models of human mobility in their different forms are important for many reasons. Models of human commuting can help reduce traffic congestion and pollution, and can be used to drive land use policy and development choices [5]. Models of human migration are equally important to policy makers as they can give broader estimates of how the population of an area will change in upcoming years, how labor markets might be affected [7], how infectious diseases spread [1, 27], and how international trade will change [8]. Much recent research focuses on modeling human commuting flows [16, 20]; however little has focused on explicitly modeling human migration.

Human mobility has been traditionally modeled with the so-called gravity model, which posits that the probability of a trip between two locations decays directly as a function of the distance between them. This model was introduced in its modern form in 1946 [31] and has been used in many applications since [3, 11, 14, 18, 21, 25]. More recently, the radiation model [26] has been shown to capture long range trips better than gravity based models, and is described as ‘a universal model for mobility and migration patterns’. The radiation model posits that the probability of a trip will decay indirectly with distance and directly with the amount of intervening opportunities, a notion first proposed by Stouffer [28]. The radiation model has been extended several times since being proposed in 2012 [20, 24, 30]. In general, gravity models have been shown to be more capable of reproducing commuting flows, i.e. human mobility at small spatial scales [16], while radiation models have been shown to be better at reproducing migration flows, i.e. human mobility at larger spatial scales [26]. Additionally, human migrations have been estimated by fitting generalized linear models derived from the gravity model [4, 6]. Both gravity and radiation models are analytical models with crafted functional forms and limited input data requirements. These models are focused on explaining human migration, rely on linear relationships between independent variables, and use hand crafted features for each zone. These approaches, while useful for explaining human migration, trade predictive power for interpretability. Data sources such as the World Bank and US Census provide many zone-based features that can be algorithmically combined in a non-linear manner by tree or neural network based models to best predict human migration.

Our key contributions are as follows:

- (1) We develop the first general machine learning formulation of the human migration prediction problem.
- (2) We develop a pipeline for training machine learning models to tackle this problem that includes procedures to deal with

dataset imbalance, hyperparameter tuning, and performance evaluation.

- (3) We develop a custom loss function for training artificial neural networks that is more suitable for the migration prediction task.
- (4) We compare the performance of machine learning models to traditional models of human migration on two datasets, and show that the machine learning models outperform the traditional models in all cases.

2 TRADITIONAL MIGRATION MODELS

In human mobility modeling, we are usually given n zones with the goal of predicting the number of people, T_{ij} , that move from every zone i to every other zone j . Traditional mobility models, such as the radiation model [26] and the gravity model [16], break the problem of estimating T_{ij} into two pieces: estimating the total number of people G_i that leave zone i (also referred to as a production function), and estimating the probability P_{ij} of a move occurring from i to j . The predicted number of migrants from i to j would be $\hat{T}_{ij} = G_i P_{ij}$. As a convention, the probabilities are normalized such that, given an origin i , the probabilities of traveling to all other destinations sum to 1, i.e. $\sum_{j=1}^n P_{ij} = 1$. If prior information about the number of incoming and outgoing travelers per zone is known, then a constrained framework, such as the one described in Lenormand et al [16], can be used. If this information is not known, which is the case in predicting migrations, then a *production function* must be used to estimate the number of outgoing travelers per zone instead. A simple *production function* for a dataset can be found by expressing the number of outgoing migrants of a zone as a constant fraction of the population of that zone (for counties in the US, this percentage is ≈ 0.03 [26]). Given prior timestep's data on the population m_i and the corresponding number of outgoing migrants O_i in the current timestep for every zone i , the *production function* is expressed as $G_i = M(m_i) = \alpha m_i$, where α is the slope of the line of best fit through the pairs (m_i, O_i) .

The traditional models of human mobility that we include in this study are: the radiation model [26], extended radiation model [30], and gravity models with both power and exponential distance decay functions [16]. The only information used by these models is: m_i , population of a zone i for both the origin and destination zones; d_{ij} , the distance between two zones i and j ; and s_{ij} , a metric of intervening opportunities measured as the total population of all intervening zones between i and j , defined as all zones whose centroid falls in the circle centered at i with radius d_{ij} (not including zones i or j). See Table 1 for a description of each model.

All of these models have two parameters that must be tuned using historical data: the production function parameter, α , that determines what fraction of the population of a zone will migrate away in a given year, and a model parameter, β . From a socio-economic point of view, traditional models have the advantage of being interpretable, however we will show that this interpretability comes at a cost of predictive accuracy, as machine learning models can use similar historic data to achieve better results.

3 EVALUATION METHODS

To evaluate how well alternative models perform, we use four metrics that compare how well a predicted migration matrix, $\hat{T}$, recreates the ground truth values, T . The first two of these (*CPC*, *CPC_d*) have been used in previous literature to evaluate human mobility models [16, 17], and the other two are standard regression metrics:

Common Part of Commuters (CPC) This metric directly compares numbers of travelers between the predicted and ground truth matrices. It will be 0 when the two matrices have no entries in common, and 1 when they are identical. We note that this metric, as used in previous studies of commuting flows, is identical to the Bray-Curtis similarity score used to compare abundance data in ecological studies [9, 15].

$$CPC(T, \hat{T}) = \frac{2 \sum_{i,j=1}^n \min(T_{ij}, \hat{T}_{ij})}{\sum_{i,j=1}^n T_{ij} + \sum_{i,j=1}^n \hat{T}_{ij}} \quad (1)$$

Common Part of Commuters Distance Variant (CPC_d) This metric measures how well a predicted migration matrix recreates trips at the same distances as the ground truth data. In this definition, N is a histogram where a bin N_k contains the number of migrants that travel between $2k - 2$ and $2k$ kilometers. It will be 0 when the two matrices do not have any migrations at the same distance, and 1 when all fall within the same distances.

$$CPC_d(T, \hat{T}) = \frac{2 \sum_{k=1}^{\infty} \min(N_k, \hat{N}_k)}{\sum_{k=1}^{\infty} N_k + \sum_{k=1}^{\infty} \hat{N}_k} \quad (2)$$

Root mean squared error (RMSE) This is a standard regression measure that will “punish” larger errors more than small errors. This score ranges from 0 in a perfect match, to arbitrarily large values as the predictions become worse.

$$RMSE(T, \hat{T}) = \sqrt{\frac{1}{n} \sum_{i,j=1}^n (T_{ij} - \hat{T}_{ij})^2} \quad (3)$$

Coefficient of determination (r^2) This score measures the goodness of fit between a set of predictions and the ground truth values. This score ranges from 1, in a perfect fit, to arbitrarily negative values as a fit becomes worse, and is 0 when the predictions are equivalent to the expectation of the ground truth values.

$$r^2(T, \hat{T}) = 1 - \frac{\sum_{i,j=1}^n (T_{ij} - \hat{T}_{ij})^2}{\sum_{i,j=1}^n (T_{ij} - \bar{T})^2} \quad (4)$$

In addition to the previous four metrics, we compare the ground truth number of incoming migrants and the predicted number of incoming migrants per zone using mean absolute error (MAE) and r^2 . The predicted number of incoming migrants for a zone, i , is calculated as $\hat{v}_i = \sum_{j=1}^n \hat{T}_{ji}$. We argue that it is important to explicitly measure how well each model performs at estimating the number of incoming migrants, because the number of incoming migrants to a location will be the most important measure for policy makers in that area. Incoming migrant predictions can inform population growth estimates and hence infrastructure planning and job analysis.

Model	Equation
Radiation	$\hat{T}_{ij} = M(m_i) \frac{m_i m_j}{(m_i + s_{ij})(m_i + m_j + s_{ij})}$
Extended Radiation	$\hat{T}_{ij} = M(m_i) \frac{[(m_i + m_j + s_{ij})^\beta - (m_i + s_{ij})^\beta](m_i^\beta + 1)}{[(m_i + s_{ij})^\beta + 1][(m_i + m_j + s_{ij})^\beta + 1]}$
Gravity with Power Law	$\hat{T}_{ij} = M(m_i) \frac{m_i m_j d_{ij}^{-\beta}}{\sum_{k=1}^n m_k d_{ik}^{-\beta}}$
Gravity with Exponential Law	$\hat{T}_{ij} = M(m_i) \frac{m_i m_j e^{-\beta d_{ij}}}{\sum_{k=1}^n m_k e^{-\beta d_{ik}}}$

Table 1: Traditional migration models

4 LEARNING MIGRATION MODELS

Formally, the problem of modeling human migration is as follows: given n zones, d_1 features describing each zone, $\mathbf{F} \in \mathbb{R}^{n \times d_1}$, and d_2 joint features describing features between a pair of zones, $\mathbf{J} \in \mathbb{R}^{n \times n \times d_2}$, at some timestep t , the objective is to predict the origin/destination migration matrix $\hat{\mathbf{T}} \in \mathbb{N}^{n \times n}$ at the next timestep, $t + 1$, where an entry $\hat{T}_{ij}$ represents the estimated number of migrants relocating from zone i to zone j . Our goal is to estimate a function $f(F_i, F_j, J_{ij}) = \hat{T}_{ij}$, which takes the features of zone i and j , as well as the joint features between them, as input, and directly outputs the estimated number of migrants that travel from i to j . This approach is different from how the traditional migration models work as it does not require a *production function*, but instead directly predicts $\hat{T}_{ij}$. This formulation contains the simplifying assumption that migrant flows are static in time, meaning that they can be entirely determined by the features from the previous timestep. In reality migrant flows will be dependent on temporal features, such as long term developmental trends, however many places will not have enough data to take advantage of these patterns. With this formulation, our models can be applied more broadly to predict future migration patterns in locations that have only collected a single year of data.

Hyperparameter optimization

To fit f for a given dataset, we will train two machine learning models, “extreme” gradient boosting regression (XGBoost model) [2], and an artificial neural network model (ANN model) [13]. Each of these models contains several hyperparameters that must be tuned to obtain good performance on a given learning task. Our first model, the XGBoost model, is a standard machine learning model based on gradient boosting trees [19] that often performs very well on many regression and classification tasks. One benefit of this model is that it gives a ranking of the relative feature importances [10]. The parameters of the XGBoost model that we consider for hyperparameter tuning are the maximum tree depth, number of estimators, and learning rate. Our ANN model is composed of densely connected layers with ReLU activation layers¹. We tune the following ANN parameters: loss function, number of layers, layer width, number of training epochs, and training mini-batch size.

¹Our model is implemented in Python with the Keras library: <https://keras.io/>

Dealing with zero-inflated data

We observe that migration data is heavily zero-inflated, where in any given year, most pairs of zones do not have any migrants traveling between them, i.e. $T_{ij} = 0$ for most (i, j) pairs. Considering migrations between US counties [23], less than 1% of the possible pairs of counties have migrations between them. This imbalance will cause problems for machine learning models. To address this problem, when creating a training dataset we undersample “negative” samples between pairs of zones for which there are no observed migrations. This is a naïve technique that will necessarily throw out available information [12]. To offset this, we introduce a hyperparameter k that determines the number of “negative” examples of migrations to train with. If there are n_t pairs of zones where there are observed migrations, “positive examples”, we include all n_t , and an additional $n_t k$ randomly chosen zone pairs where there are no observed migrants. This hyperparameter is included in both the XGBoost and ANN model parameter searches. We give further details on the hyperparameter tuning process in Section 5.

Custom ANN loss function

Previously we mentioned that our ANN model will consider different loss functions as a hyperparameter. Common loss functions for regression tasks include “mean squared error” (MSE), “mean absolute error” (MAE), and “mean absolute percentage error” (MAPE). Our preliminary experimental results show that MAE and MAPE loss functions perform poorly, partially due to their inability to punish large errors and deal with many zeros respectively. To contrast with the aforementioned zero-inflation problem, we observe that the distribution of migrant counts has a long tail, whereby few pairs of zones consistently experience large volumes of migrations. Considering these observations, and because the CPC metric is one of the key metrics of interest (described in detail in Section 3), we derive a new loss function based on CPC to train the ANN model with. This loss function is given as:

$$L(y, \hat{y}) = 1 - \frac{2 \sum_{i=1}^n \min(y_i, \hat{y}_i)}{\sum_{i=1}^n y_i + \sum_{i=1}^n \hat{y}_i} \quad (5)$$

Where y_i is a migration flow entry from $\mathbf{T}$. The gradient update for this loss function is:

$$\frac{\partial L(y, \hat{y})}{\partial \hat{y}_j} = \frac{2 \sum_{i=1}^n \min(y_i, \hat{y}_i)}{(\sum_{i=1}^n y_i + \sum_{i=1}^n \hat{y}_i)^2} - \begin{cases} 2 & \hat{y}_j < y_j \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Intuitively, this loss will ‘reward’ predictions where the number of migrants matches the ground truth, while also enforcing per link absolute error to be minimized. This loss function is not exactly equivalent to the *CPC*, as during the ANN training it will only consider a single mini-batch worth of samples at a time (in our case $|y| = |\hat{y}| = |\text{mini batch}|$, where as the *CPC* metric is a function of the entire migration matrix). Empirically this custom loss function results in better validation performance and faster training times than MSE loss.

5 EXPERIMENTS

Datasets

We perform experiments comparing the performance of traditional models to our machine learning models on two datasets, the *USA Migration* dataset and the *Global Migration* dataset.

The *USA Migration* dataset consists of yearly intra-county migrations in the USA between 3106 counties from the IRS Tax-Stats data [23] for the 11 years in the range from 2004 to 2014. We supplement the migration data with the following 7 per-county features (taken from the US Census estimates and calculated from the Census TIGER line maps of US county boundaries): population, land area, population density, median household income, county water area, is a coastal county, and number of neighboring counties. In addition to these 7 per-county features, we add the following between-county features: distance, intervening population, intervening land area, intervening number of counties, intervening population density, and intervening median income. The intervening features are calculated based on the idea of “intervening opportunities” presented in the radiation model. For any given county-level variable, x , e.g. population, the intervening amount of that variable between counties i and j is defined as $s_{i,j}^x$, the sum of all x in the intervening counties that fall within the circle centered at county i with a radius to county j (excluding x_i and x_j).

The *Global Migration* dataset consists of decadal inter-country migration data between 207 countries from the World Bank Global Bilateral Migration Database [22]. The *Global Migration* dataset contains 5 timesteps, one every 10 years from 1960 to 2000. In the *Global Migration* dataset we add the following 5 per-country features (taken from World Bank World Development Indicators data [29]): population, population density, population growth, agricultural emissions, and land area. Additionally, we include 3 between-country features: distance, intervening population, and intervening land area.

For each year of data in the *USA Migration* and *Global Migration* datasets we create an ‘observation’ for each pair of zones, an origin zone and destination zone (counties in the *USA Migration* dataset and countries in the *Global Migration* dataset). Each observation consists of the per-zone features for both the origin zone and destination zone (population of origin, population of destination, etc.) and the between-zone features of the origin and destination. This corresponds exactly to the $F_{i\cdot}, F_{\cdot j}, J_{ij}$ of the function f (described in Section 4) that we want to learn. An observation is associated with the target number of migrants, T_{ij} , traveling from the origin to the destination. Formally, for a given year, t , number of zones, n , number of per-zone features, d_1 , and number of between-zone features, d_2 , we create a matrix of observations $\mathbf{X}_t \in \mathbb{R}^{n^2 \times (2d_1 + d_2)}$ and

vector of targets $Y_t \in \mathbb{R}^{n^2}$, capturing the migration flows observed in year $t + 1$.

Experimental Setup

To select the hyperparameters of the models that we described in Sections 2 and 4, we consider triplets of “years” of data as training, validation, and testing sets. Specifically, for three years of data $\{(X_{t-2}, Y_{t-2}), (X_{t-1}, Y_{t-1}), (X_t, Y_t)\}$, we call (X_{t-2}, Y_{t-2}) the training set, (X_{t-1}, Y_{t-1}) the validation set, and (X_t, Y_t) the test set. We tune the hyperparameters of the models using a randomized grid search with 50 sampled hyperparameters using the training and validation sets. We select the best set of hyperparameters according to the *CPC* score, then use those parameters to train a model on the validation set and record its performance on the test set. We repeat this process for each $(t - 2, t - 1, t)$ triplet of years in each dataset. Our final results are reported as averages over the test set results.

A hyperparameter present in both XGBoost and ANN models is the downsampling rate, k . As a preprocessing step for a given year of training data (X_t, Y_t) , we include all observations, X_i where $Y_i > 0$ (let this number of samples be m), and choose $k * m$ random samples with replacement from the remaining observation (where $Y_i = 0$). This sampling process only takes place when training a model. When testing a model on the validation or test sets, the full datasets are always used. For experiments with the *USA Migration* dataset we consider values of k in a uniform distribution of integers from 5 to 100, while for experiments with the *Global Migration* dataset we consider the uniform distribution of integers from 1 to 5, because the average percentage of non-zeros is $< 1\%$ and 20% respectively.

For the XGBoost models, we sample the following parameters: maximum tree depth from $\mathcal{U}\{2, 7\}^2$, number of estimators from $\mathcal{U}\{25, 275\}$, and learning rate uniformly in the range from 0 to 0.5. For the ANN models we sample the following parameters: network loss function uniformly from $\{\text{‘CPC Loss’}, \text{‘MSE’}\}$, number of layers uniformly from $\mathcal{U}\{1, 5\}$, layer width from $\mathcal{U}\{16, 128\}$, number of training epochs from $\mathcal{U}\{10, 50\}$, and training mini-batch size uniformly from $\{2^9, 2^{10}, 2^{11}, 2^{12}, 2^{13}, 2^{14}\}$.

We calibrate the parameters of the traditional models in a similar manner. Every traditional model, except for the radiation model, has two parameters, α and β , that must be calibrated to give useful results (where the radiation model only uses α). For each pair of years of data, $\{(X_{t-1}, Y_{t-1}), (X_t, Y_t)\}$, that we refer to as training and testing sets respectively, we find the value of α that gives the best production function on the training set. Similarly, we find the value of β that maximizes the *CPC* score of each traditional models on the training set. We then use these α and β values to run each model on the testing set, and report the results as averages over all test set results.

To directly compare how the ML models and traditional models perform under the same conditions, we perform experiments where the ML models are used with the same production functions as the traditional models. This imposes an artificial constraint on the ML models, as these models are able to directly estimate the number of migrations between two zones, without supplemental information on the number of outgoing migrants from each zone.

² $\mathcal{U}\{a, b\}$ is the discrete uniform distribution.

USA Migrations	Metrics on full matrix				Metrics on incoming migrants vector (Average Incoming Migrants = 3,196)	
Production Function	CPC	CPC _d	RMSE	r ²	MAE	r ²
Gravity Model Exponential Decay	0.53 +/- 0.01	0.66 +/- 0.02	87.4 +/- 9.0	-1.48 +/- 0.28	1,216 +/- 128	0.67 +/- 0.03
Gravity Model Power Law Decay	0.56 +/- 0.01	0.78 +/- 0.02	75.7 +/- 8.0	-0.86 +/- 0.26	1,129 +/- 129	0.72 +/- 0.04
Radiation Model	0.53 +/- 0.01	0.76 +/- 0.02	47.6 +/- 5.0	0.26 +/- 0.09	1,346 +/- 148	0.80 +/- 0.02
Extended Radiation Model	0.58 +/- 0.01	0.83 +/- 0.01	35.6 +/- 3.0	0.59 +/- 0.03	1,123 +/- 117	0.86 +/- 0.02
XGBoost model - traditional features	0.51 +/- 0.08	0.74 +/- 0.07	28.6 +/- 5.4	0.72 +/- 0.10	1,151 +/- 249	0.86 +/- 0.04
ANN model - traditional features	0.63 +/- 0.01	0.86 +/- 0.02	35.1 +/- 3.2	0.60 +/- 0.03	911 +/- 107	0.91 +/- 0.01
XGBoost model - extended features	0.58 +/- 0.03	0.78 +/- 0.02	24.2 +/- 1.4	0.81 +/- 0.02	968 +/- 56	0.89 +/- 0.02
ANN model - extended features	0.68 +/- 0.01	0.89 +/- 0.02	29.8 +/- 2.7	0.71 +/- 0.02	935 +/- 98	0.91 +/- 0.02
No Production Function	CPC	CPC _d	RMSE	r ²	MAE	r ²
XGBoost model - traditional features	0.54 +/- 0.11	0.99 +/- 0.02	18.5 +/- 6.1	0.88 +/- 0.08	3,091 +/- 1,740	0.41 +/- 0.85
ANN model - traditional features	0.63 +/- 0.02	0.88 +/- 0.06	35.3 +/- 3.5	0.60 +/- 0.04	1,188 +/- 259	0.84 +/- 0.16
XGBoost model - extended features	0.62 +/- 0.04	0.99 +/- 0.02	13.0 +/- 1.5	0.94 +/- 0.02	2,060 +/- 622	0.76 +/- 0.28
ANN model - extended features	0.69 +/- 0.01	0.93 +/- 0.05	28.0 +/- 3.6	0.75 +/- 0.03	909 +/- 48	0.92 +/- 0.04

Table 2: USA Migration results. Comparison of the ANN and XGBoost models with and without a production function to traditional migration models. The values shown in the table are the average and standard deviations of the models' test performance on 2006 through 2014 data. Bold values indicate the best values per column.

Global Migrations	Metrics on full matrix				Metrics on incoming migrants vector (Average Incoming Migrants = 674,858)	
Production Function	CPC	CPC _d	RMSE	r ²	MAE	r ²
Gravity Model Exponential Decay	0.16 +/- 0.00	0.16 +/- 0.00	62,218 +/- 5,341	0.02 +/- 0.03	651,194 +/- 80,220	0.00 +/- 0.02
Gravity Model Power Law Decay	0.16 +/- 0.00	0.15 +/- 0.00	61,523 +/- 5,278	0.05 +/- 0.00	628,678 +/- 79,474	0.03 +/- 0.03
Radiation Model	0.16 +/- 0.00	0.14 +/- 0.00	62,173 +/- 5,277	0.02 +/- 0.00	614,483 +/- 79,378	0.04 +/- 0.02
Extended Radiation Model	0.16 +/- 0.00	0.14 +/- 0.00	62,108 +/- 5,299	0.03 +/- 0.00	618,576 +/- 76,150	0.03 +/- 0.02
XGBoost model - traditional features	0.18 +/- 0.01	0.14 +/- 0.01	58,377 +/- 5,141	0.14 +/- 0.01	597,478 +/- 79,178	0.10 +/- 0.01
ANN model - traditional features	0.19 +/- 0.01	0.15 +/- 0.01	60,272 +/- 4,610	0.08 +/- 0.02	589,789 +/- 79,542	0.26 +/- 0.03
XGBoost model - extended features	0.21 +/- 0.01	0.16 +/- 0.01	57,909 +/- 5,409	0.16 +/- 0.02	573,090 +/- 56,987	0.15 +/- 0.02
ANN model - extended features	0.22 +/- 0.02	0.17 +/- 0.01	58,887 +/- 4,477	0.12 +/- 0.02	563,259 +/- 74,127	0.24 +/- 0.06
No Production Function	CPC	CPC _d	RMSE	r ²	MAE	r ²
XGBoost model - traditional features	0.33 +/- 0.02	0.59 +/- 0.03	52,729 +/- 5,455	0.26 +/- 0.26	938,905 +/- 172,834	0.18 +/- 0.28
ANN model - traditional features	0.33 +/- 0.01	0.37 +/- 0.04	56,005 +/- 882	0.20 +/- 0.11	537,351 +/- 44,034	0.53 +/- 0.16
XGBoost model - extended features	0.43 +/- 0.03	0.64 +/- 0.02	47,329 +/- 5,073	0.42 +/- 0.13	577,473 +/- 77,315	0.48 +/- 0.34
ANN model - extended features	0.40 +/- 0.02	0.43 +/- 0.02	50,921 +/- 3,556	0.33 +/- 0.13	459,841 +/- 55,479	0.52 +/- 0.30

Table 3: Global Migration results. Comparison of the ANN and XGBoost models with and without a production function to traditional migration models. The values shown in the table are the average and standard deviations of the models' test performance on 2006 through 2014 data. Bold values indicate the best values per column.

To apply a production function, M , to the predictions made by a ML model, $\hat{T}$, we create a new set of predictions, $\hat{T}'$, where an entry $\hat{T}'_{ij} = M(m_i) \frac{\hat{T}_{ij}}{\sum_{k=1}^n \hat{T}_{ik}}$.

Results

Tables 2 and 3 show the average results over all years of data of the ML models and traditional models in the *USA Migration* and *Global Migration* datasets respectively. From these tables we observe that the best traditional model for the *USA Migration* dataset is the Extended Radiation model, beating the other traditional models

in all metrics. None of the traditional models are able to capture the migration dynamics in the *Global Migration* dataset; they all have an r^2 score near 0, meaning that a model which predicts the average number of migrants for every link would perform just as well. The ML models perform much better. In the case where the ML models are constrained to the same conditions as the traditional models, using traditional features and a production function, the ANN model beats all of the traditional models in 5 out of the 6 measures in the *USA Migration* datasets, and the XGBoost model beats all the traditional models in 4 out of the 6 metrics. Similarly, the ML models considerably outperform the traditional models in the

USA Features	Importance	Global Features	Importance
Intervening number of counties	25.3% +/- 2.4%	Population growth of origin	19.5% +/- 16.0%
Population of origin (trad)	15.7% +/- 1.7%	Intervening population (trad)	12.3% +/- 3.7%
Population of destination (trad)	14.2% +/- 0.9%	Agricultural emissions of destination	10.6% +/- 5.8%
Intervening population (trad)	6.1% +/- 1.2%	Intervening land area	8.7% +/- 5.6%
Is destination coastal	4.3% +/- 4.6%	Population growth of destination	7.9% +/- 6.2%
Distance between counties (trad)	3.7% +/- 0.9%	Population of destination (trad)	6.9% +/- 0.8%
Intervening area	3.6% +/- 0.9%	Distance between counties (trad)	6.6% +/- 1.9%
Area of destination	3.5% +/- 0.5%	Population of origin (trad)	6.1% +/- 1.6%
Number of neighbors destination	3% +/- 1.6%	Population density of destination	5.7% +/- 4.5%
Water area of origin	3% +/- 1.3%	Land area of origin	5.2% +/- 3.1%

Table 4: Top 10 most important (extended) features in both the *USA Migration* and *Global Migration* datasets. The values in the table show the average and standard deviations of the information gain feature importances from an XGBoost model trained on the extended feature set for each timestep of data.

Global Migration, outperforming them in all metrics. Considering the extended feature set results, the ML models perform even better. The ANN and XGBoost models without a production function outperform the same models with a production function in 5 out of the 6 metrics. The XGBoost model outperforms the ANN model in 3 out of the 4 metrics that evaluate the models' per link predictions, however the ANN model performs better on the two metrics that evaluate the aggregate incoming migrant prediction performance.

These results suggest that more features than those which are used by the traditional models, are needed to accurately predict human migrations. The ability of ML models to incorporate any number of additional features is one of the key motivations for using them to obtain more accurate results. Considering this, it will be insightful to understand, which of the features are most informative to the ML models. Since feature importance analysis for ANNs is quite challenging, we report in Table 4 the top 10 most important features for both datasets (based on information gain) in the XGBoost model trained on the *extended feature set*, averaged over all years of data. In both datasets, the intervening population feature is in the top 4 important features which validates the intuition that intervening opportunities are important in predicting migrations. The most important feature in the *USA Migration* dataset is the number of intervening counties between two locations, a simpler form of the intervening population idea. In the *Global Migration* dataset, the population growth of the origin is the most important feature on average, with a large standard deviation. In some years this feature is very important, however in other years it is less so. Intuitively, population growth will be correlated with the amount of incoming migration. During relatively stable years, with small population growth, other features will be more predictive of migration.

In Figure 1 we show the difference between the actual and predicted numbers of incoming migrants per county for the two best traditional models, and all of the ML models without production functions. From these maps we can see that between ML models, those trained with the extended feature set perform better than those trained with only the traditional features. Specifically, without the extended features, the ML models underpredict the number of migrations to the western portion of the United States. When

the extended features are taken into account, the models are able to correct for this spatial bias. Holding with the experimental results, we can see that the ANN model with extended features best captures the incoming migrant distributions per county. The ANN model is able to more accurately match the number of migrants that travel to rural areas (e.g. to the midwestern US), compared to the traditional models that consistently over estimate the numbers of migrants to rural areas. In general, these maps agree with our empirical results, that the ANN model (with the lowest average incoming migrants *MAE*) is able to best predict migrations.

6 CONCLUSION

With the increasing availability of high resolution socio-economic data in countries that record human migrant flows, it is possible to use machine learning models of human migration rather than traditional gravity or radiation models. Machine learning models offer greater levels of modeling flexibility, as they can combine many input features in non-linear ways that can not be captured by static equations. Furthermore, machine learning models can be easily customized to the problem or country at hand.

We develop two machine learning based models for the task of predicting human migration flows, for both between counties in the US and between countries across the world. We compare these models to traditional human migration models using two sets of features and show that our models outperform the traditional models in most of the evaluation metrics.

We would like to extend this work to better explain human migration through a more complete analysis of features included in the model, and incorporating different models. While the XGBoost model can provide a ranking of feature importances, this does not fully explain the dynamics that drive human migration. Additionally we would like to study how these migration models could be specialized to predict migrations under extreme weather events. Hurricanes and other natural disasters can displace large populations, and determining where these populations will resettle is the first step towards providing a unique planning tool for policy makers. To achieve these goals, higher resolution migration data, on both spatial and temporal scales, will need to be obtained. Extreme weather events, by definition, are short lived and their effects

Figure 1: USA Migrations modeling error. These maps show the difference between the ground truth number of incoming migrants and predicted number of incoming migrants per county for 6 models in 2014. Blue corresponds to overestimation by the model, red to underestimation by the model, and white if the model accurately predicts the correct number of incoming migrants. Top row shows the results for the Extended Radiation model and Gravity model with power law distance decay. Middle row shows the results for ANN models trained with the extended and common feature sets. Bottom row shows the results for XGBoost models trained with the extended and common feature sets.

are averaged out in coarser resolution datasets such as the ones used in this study. Finally, we would like to study the connections between human migration and other processes primarily driven by aggregate human behaviors such as inter- and intra-national trade.

ACKNOWLEDGMENTS

This work is supported by the NSF grant CCF-1522054 (COMPUNET: Expanding Horizons of Computational Sustainability).

REFERENCES

- [1] Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J Ramasco, and Alessandro Vespignani. 2009. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the National Academy of Sciences* 106, 51 (2009), 21484–21489.
- [2] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 785–794.
- [3] Gordon L Clark and Kenneth P Ballard. 1980. Modeling out-migration from depressed regions: the significance of origin and destination characteristics. *Environment and Planning A* 12, 7 (1980), 799–812.
- [4] Joel E Cohen, Marta Roig, Daniel C Reuman, and Cai GoGwilt. 2008. International migration beyond gravity: A statistical model for use in population projections. *Proceedings of the National Academy of Sciences* 105, 40 (2008), 15269–15274.
- [5] Andrea De Montis, Alessandro Chessa, Michele Campagna, Simone Caschili, and Giancarlo Deplano. 2010. Modeling commuting systems through a complex network analysis: A study of the Italian islands of Sardinia and Sicily. *Journal of Transport and Land Use* 2, 3 (2010).
- [6] Adam Dennett and Alan Wilson. 2013. A multilevel spatial interaction modelling framework for estimating interregional migration in Europe. *Environment and Planning A* 45, 6 (2013), 1491–1507.
- [7] Taryn Dinkelman and Martine Mariotti. 2016. The Long-Run Effects of Labor Migration on Human Capital Formation in Communities of Origin. *American Economic Journal: Applied Economics* 8, 4 (2016), 1–35.
- [8] Giorgio Fagiolo and Marina Mastrorillo. 2013. International migration network: Topology and modeling. *Physical Review E* 88, 1 (2013), 012812.
- [9] Daniel P Faith, Peter R Minchin, and Lee Belbin. 1987. Compositional dissimilarity as a robust measure of ecological distance. *Vegetatio* 69, 1-3 (1987), 57–68.
- [10] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [11] Michael J Greenwood. 1985. Human migration: Theory, models, and empirical studies*. *Journal of regional Science* 25, 4 (1985), 521–544.
- [12] Xinjian Guo, Yilong Yin, Cailing Dong, Gongping Yang, and Guangtong Zhou. 2008. On the class imbalance problem. In *Natural Computation, 2008. ICNC'08. Fourth International Conference on*, Vol. 4. IEEE, 192–201.
- [13] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
- [14] Everett S Lee. 1966. A theory of migration. *Demography* 3, 1 (1966), 47–57.
- [15] Pierre Legendre and Loïc J J Legendre. 2012. *Numerical ecology*. Vol. 24. Elsevier.
- [16] Maxime Lenormand, Alexis Bassolas, and José J Ramasco. 2016. Systematic comparison of trip distribution laws and models. *Journal of Transport Geography* 51 (2016), 158–169.
- [17] Maxime Lenormand, Sylvie Huet, Floriana Gargiulo, and Guillaume Deffuant. 2012. A universal model of commuting networks. *PLoS one* 7, 10 (2012), e45985.
- [18] Emmanuel Letouzé, Mark Purser, Francisco Rodríguez, Matthew Cummins, and others. 2009. Revisiting the migration-development nexus: a gravity model approach. *Human Development Research Paper* 44 (2009).
- [19] Llew Mason, Jonathan Baxter, Peter L Bartlett, and Marcus R Frean. 1999. Boosting Algorithms as Gradient Descent. In *NIPS*. 512–518.
- [20] A Paolo Masucci, Joan Serras, Anders Johansson, and Michael Batty. 2013. Gravity versus radiation models: On the importance of scale and heterogeneity in commuting flows. *Physical Review E* 88, 2 (2013), 022812.
- [21] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. 2012. A tale of many cities: universal patterns in human urban mobility. *PLoS one* 7, 5 (2012), e37027.
- [22] Çağlar Özden, Christopher R Parsons, Maurice Schiff, and Terrie L Walmsley. 2011. Where on earth is everybody? The evolution of global bilateral migration 1960–2000. *The World Bank Economic Review* 25, 1 (2011), 12–56.
- [23] Kevin Pierce. 2015. *SOI migration data. A new approach: Methodological improvements for SOI's United States population migration data, calendar years 2011–2012*. Technical Report. Statistics of Income, Internal Revenue Service.
- [24] Yihui Ren, Mária Ercsey-Ravasz, Pu Wang, Marta C González, and Zoltán Toroczkai. 2014. Predicting commuter flows in spatial networks using a radiation model based on temporal ranges. *Nature communications* 5 (2014).
- [25] Morton Schneider. 1959. Gravity models and trip distribution theory. *Papers in Regional Science* 5, 1 (1959), 51–56.
- [26] Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. 2012. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96–100.
- [27] Alessandro Sorichetta, Tom J Bird, Nick W Ruktanonchai, Elisabeth zu Erbach-Schoenberg, Carla Pezzulo, Natalia Tejedor, Ian C Waddock, Jason D Sadler, Andres J Garcia, Luigi Sedda, and others. 2016. Mapping internal connectivity through human migration in malaria endemic countries. *Scientific Data* 3 (2016).
- [28] Samuel A Stouffer. 1940. Intervening opportunities: a theory relating mobility and distance. *American sociological review* 5, 6 (1940), 845–867.
- [29] World Bank. 2017. <http://databank.worldbank.org/data/reports.aspx>. (2017).
- [30] Yingxiang Yang, Carlos Herrera, Nathan Eagle, and Marta C González. 2014. Limits of predictability in commuting flows in the absence of data for calibration. *Scientific reports* 4 (2014).
- [31] George Kingsley Zipf. 1946. The P 1 P 2/D hypothesis: on the intercity movement of persons. *American sociological review* 11, 6 (1946), 677–686.