The Secret Lives of Names? Name Embeddings from Social Media

Junting Ye
Stony Brook University
Stony Brook, NY
juyye@cs.stonybrook.edu

Steven Skiena
Stony Brook University
Stony Brook, NY
skiena@cs.stonybrook.edu

ABSTRACT

Your name tells a lot about you: your gender, ethnicity and so on. It has been shown that name embeddings are more effective in representing names than traditional substring features. However, our previous name embedding model is trained on private email data and are not publicly accessible. In this paper, we explore learning name embeddings from public Twitter data. We argue that Twitter embeddings have two key advantages: (i) they can and will be publicly released to support research community. (ii) even with a smaller training corpus, Twitter embeddings achieve similar performances on multiple tasks comparing to email embeddings.

As a test case to show the power of name embeddings, we investigate the modeling of lifespans. We find it interesting that adding name embeddings can further improve the performances of models using demographic features, which are traditionally used for lifespan modeling. Through residual analysis, we observe that fine-grained groups (potentially reflecting socioeconomic status) are the latent contributing factors encoded in name embeddings. These were previously hidden to demographic models, and may help to enhance the predictive power of a wide class of research studies.

ACM Reference Format:

Junting Ye and Steven Skiena. 2019. The Secret Lives of Names? Name Embeddings from Social Media. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3292500.3330709>

1 INTRODUCTION

Your name tells a lot about you. It commonly reveals your gender (male or female) and ethnicity (White, Black, Hispanic, or Asian/Pacific Islander). It can reveal your religion and your country of family origin. It can even inform on your marital status (is it hyphenated?), age (e.g. the generational differences between Fannie and Caitlin), or socioeconomic class (consider Archibald vs. Jethro).

Name embeddings are distributed representations which encode the cultural context of name parts (i.e. given name and surname) in 100-dimension vectors learned through an unsupervised technique. It has been shown that name embeddings are more effective

Male	1th NN	2nd NN	3rd NN	4th NN
Andy	Pete	Stuart	Craig	Will
Dario	Giovanni	Luigi	Francesco	Claudio
Hilton	Jefferson	Maryellen	Jayne	Brock
Lamar	Ty	Reggie	Jada	Myles
Mohammad	Abdul	Ahmad	Hassan	Ahmed
Rocco	Francesca	Carlo	Giovanni	Luigi
Female	1th NN	2nd NN	3rd NN	4th NN
Adrienne	Aimee	Brittany	April	Kristen
Aisha	Maryam	Fatima	Ayesha	Fatimah
Brianna	Brooke	Kayla	Kaylee	Megan
Chan	Ka	Cherry	Yun	Sha
Cheyenne	Hannah	Kayla	Madison	Kelsey
Gabriella	Isabella	Dario	Cecilia	Paola

Table 1: Four nearest neighbors of representative names in Twitter embedding space, showing how they preserve gender and ethnicity associations. Notes: Asian (Chinese, Korean, Japanese, Vietnamese), British, European (Spanish, Italian), Middle Eastern (Arabic, Hebrew), North American (African-American, Native American, Contemporary).

representations than substrings on various tasks [18, 37]. Table 1 presents a representative set of name parts, each with their four nearest neighbors in name embedding space. It is clear that they preserve associations of gender and ethnicity. Unfortunately, previous embeddings were trained on private email data and are not publicly accessible to research community.

In this paper, we propose to learn name embeddings from public Twitter data. Our motivation is that name embeddings perform well because of homophily, i.e. the tendency for people to associate with those similar to themselves. These associations are reflected by communication patterns, which explains why large-scale email networks proved so effective at elucidating them. We argue that homophily in communication is universal, and also exists social media [4]. Two major properties make Twitter embeddings a better alternative: (i) Twitter name embeddings can and will be released to support research community. (ii) Twitter embeddings achieve similar performances on gender, ethnicity and nationality identification as Email embeddings, even though the training corpus for Email is two times larger than that for Twitter. We observe that Twitter embeddings have better performances on gender prediction, while Email embeddings achieve higher scores on ethnic predictions.

A second focus of our work is to demonstrate the predictive power of name embeddings on lifespan modeling, where gender, ethnicity and nationality are all contributing features. Average lifespan is one of the most critical measurements associated with quality

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](https://permissions.acm.org).

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330709>

of life across different demographic groups. Mortality prediction for individuals from available features is the foundation of life insurance industry. Here we demonstrate how an individual's most readily available features (names and corresponding embeddings) can be used to improve the accuracy over comparable demographic models. It is an amazing testament to the power of homophily that contemporary communication patterns can account for mortality in people born over a century ago.

We summarize our primary contributions in this paper as following:

- *Twitter name embeddings.* We explore and evaluate nine versions of Twitter name embeddings (see Table 2). We get interesting observations via performance comparisons: (i) *Mention* embeddings outperform *Email* embeddings and other Twitter embeddings on gender recognition, indicating stronger gender homophily in Twitter mentions. (ii) *Followers* embeddings work better than *Followee*, because ordinary users' followers tend to be family members and/or close friends, while there are more celebrities among followees. The performance is improved after removing celebrities' names from followee lists. (iii) *Aggregated** embeddings perform the best among nine Twitter versions. They have similar vocabulary size and achieve comparable performances on gender, ethnicity and nationality classification as *Email* embeddings. Twitter name embeddings are shared for research community (www.name-prism.com).
- *Demonstrating the power of name embeddings to improve lifespan modeling.* To demonstrate the power of name embeddings, we train a series of models to predict lifespan as a function of five traditional demographic variables (birth year, state, gender, ethnicity and nationality) and name embedding features. We construct 32 (i.e. 2^5) different sets of linear regression models containing specific subsets of demographic variables, with and without Twitter/email name embeddings. Incorporating name embeddings in all cases improves the underlying models significantly (p values smaller than 0.01).
- *Uncovering latent factors encoded in name embeddings.* Implicit feature models, like name embeddings, do not come with natural explanations of exactly what effective properties they are encoding. However, we can gain insight by identifying the names which contribute most strongly to the final model. By conducting residual analysis, we get the most favorable/unfavorable names that increase/decrease lifespan most from the latent factors in name embeddings. We observe that fine-grained groups are the latent contributing factors encoded in name embeddings. For example, our results show that a class-based life-expectancy bias against diminutive names (e.g. Wm, Dan and Guy) as compared to their formal forms (William, Daniel and Guido). In addition, 17 out of 20 most favorable last names have Jewish origins, which agrees with existing observation that Jewish have long average lifespan [3].

It is important to note that name embeddings encode homophily as features without explicit labels of gender, ethnicity and nationality. Models which discriminate on such criteria are a growing social concern [25]. Name embeddings have the potential to help

identify biases, as name embedding-based classifiers [37] are already widely used by over 100 social scientists and economists to study discrimination and homophily [14, 34, 35]. For example, Gornall and Strebulaev find that Asian entrepreneurs received a 6% higher rate of interested replies than White, after sending 80,000 pitch emails introducing promising but fictitious start-ups to 28,000 venture capitalists [17]. AlShebli et. al. study the effect of diversity on scientific impact, as reflected in citations. They find that ethnic diversity has the strongest correlation with scientific impact [5]. Therefore, we believe a public and sharable name embeddings will help to enhance the predictive power of a wide class of research studies.

2 RELATED WORK

2.1 Names and Mortality

There have been several previous studies of the impact of names on lifespans. Compared to our work, these have generally been performed on smaller datasets (hundreds or perhaps thousands of individuals), versus the 85 million names in our study. Further, they have generally studied surface features of names as opposed to the latent properties exposed by our name embeddings. In particular, Abel and Kruger [2] observed that several categories of people whose first name began with 'D' appeared to die earlier than those with other names. This effect did not show up in a larger-scale study [31], and an independent study by Pinzur and Smith [26] concludes that first name and life expectancy are not related.

Among athletes, Abel and Kruger [1] observe that having nicknames increases longevity. Shin and Cho [30] report that self-reported stress declines after people legally change their names, demonstrating that there can be genuine physiological effects associated with undesired names. Pena's analysis of SSDI data suggests that people with more frequent names have shorter average and median lifespans.

Nelson and Simmons [24] identify several surprising impacts of names, including that students whose names begin with C or D achieve lower GPAs and attend lower-ranked law schools than do students whose names begin with A or B. Jones, et al. [21] find that people disproportionately marry others whose first or last name resembles their own.

2.2 Gender, Nationality and Ethnicity Detection

Nationality and ethnicity are important demographic categorizations of people, standing in as proxies to represent a range of cultural and historical experiences. Names are important markers of cultural diversity, and have often served as the basis of automatic nationality classification for biomedical and sociological research. In the medical literature, nationality from names has been used as a proxy to reflect genetic differences [9, 12] and public health disparity [10, 28] between groups. Nationality identification is also important in ads targeting, and academic studies of political campaigns and social media [7, 13]. Name analysis is often the only practical way to gather ethnicity/nationality annotations because of privacy or legal concerns.

Name ethnicity classifiers often make use of characteristic substrings in names as features [6, 13, 33]. Ambekar et al. combine

decision tree and Hidden Markov Model to conduct hierarchical classification on a taxonomy with 13 leaf classes [6]. Treeratpituk et al. utilize both alphabet and phonetics sequences in names to improve performance [33]. Chang et al. use Bayesian methods to infer ethnicity of Facebook users with US census data and study the interactions between ethnic groups [13]. The linguistics features from users' tweets also reveal their ethnicities [27]. Other relevant efforts are binary ethnicity classifiers on names, e.g. Hispanic vs. Non-Hispanic [11], Chinese vs. Non-Chinese [15], South Asian vs. Non-South Asian [19].

The ethnicity/nationality classifier, *NamePrism*, consists of a 39-class name nationality classifier and a 6-class ethnicity classifier. It uses Naive Bayes model for training and testing on 74 million names labeled with country of residence. Extensive experiments [37] demonstrate that it achieves a better classification performance (F1 score) on names drawn from Wikipedia (0.651) and Email/Twitter (0.795) than competing classifiers HMM [6]. We adopt *NamePrism* to the ethnicity/nationality classification for our experiments.

Gender classifiers typically classify names according to statistics on the ratio of males to females observed in the U.S. Census. More specifically, we use data from the 1990 U.S. Census data to label popular first names by gender. We use these names' labels to approximate that of less common names. In particular, for a given first name f , we find its k nearest neighbors in name embedding space and use a majority vote to decide the gender of f .

3 NAME EMBEDDINGS

Distributed representations are feature encodings where objects are represented by points in an abstract d -dimensional space, such that similar objects are represented by points close in space. Such representations are a fundamental aspect of Deep Learning [16], a recent approach to machine learning which has proven to lead to improved results on many computer vision and natural language processing tasks. Word embeddings are a particularly important type of distributed representation, where each word is denoted by a single point, so that words which play similar roles tend to be represented by nearby points [23].

Inspired by word embeddings, Ye et al. develop name embeddings as a form of distributed representation to capture the semantic meaning of first-name and last-name parts [37]. These new representations were trained on the email contact lists of 57 million people. The use of contact lists is motivated by the principle of homophily: that people generally communicate with people similar to themselves [22]. In other words, people disproportionately associate with others of the same gender, ethnicity, nationality, and class. More formally, name embedding algorithm tries to maximize following objective:

$$\log \sigma(v_n^T v_{n'}) + \sum_{i=1}^k \mathbb{E}_{n_i \sim D} \log \sigma(-v_n^T v_{n_i}) \quad (1)$$

where v_n is the embedding of name part n , and $v_{n'}$ is the embedding of a nearby name part n' that co-occur with n in the same contact list. n_i is a random sample from name part distribution D . $\sigma(x)$ is sigmoid function, i.e. $\sigma(x) = 1/(1 + e^{-x})$.

In a nutshell, the objective aims to maximize the similarities of nearby name part pairs (the first term) and minimize random pair similarities (the second term, i.e. negative sampling). Therefore, the locality properties of name embeddings reflect underlying similarities between name parts, e.g. gender, ethnicity and nationality.

However, *Email* embeddings and *NamePrism* are corporate property and not shareable [37]. In this section, we discuss how to learn powerful name embeddings from public Twitter data. In particular, we focus on comparing embeddings trained on different user associations from Twitter. We appreciate generous assistance from *NamePrism* team in preparing the experiments.

3.1 Learning Name Embeddings from Twitter

We explore the potential of learning name embeddings from Twitter, one of the most popular social media in the world. Its API enables us to access public Tweets and users profiles. In this paper, we are interested in two types of data: (i) Tweets containing user associations, including the ones with user mentions and retweets. (ii) follower and followee lists of ordinary users (numbers of followers/followees range between 50 and 500). We assume that follower/followee lists of these users tend to encode more homophily signals than those from celebrities or inactive users.

3.1.1 Nine Training Corpora. Nine different Twitter training corpora are prepared to compare strength of different embeddings. Their definitions are as follows. All Names are extracted from Twitter profiles using user IDs. We expect names in the pairs/lists are statistically similar because of homophily. For the convenience of description, let $U_r(u_0) = \{u_1, u_2, \dots, u_n\}$ be the follower list of user u_0 , $U_e(u_0) = \{u_1, u_2, \dots, u_m\}$ be u_0 's followee list.

- **Retweet:** Twitter user pairs are extracted from retweets, i.e. (u_0, u_1) . u_0 posts the retweet and u_1 is the original Tweet author.
- **Mention:** List of users extracted from Tweets with user mentions, i.e. $(u_0, u_1, \dots, u_n)$. u_0 is the user posting the Tweet. u_1 to u_n are the users mentioned in the Tweet.
- **Follower:** List of users who follow user u_0 (i.e. $U_r(u_0)$).
- **Followee:** List of users whom user u_0 follows (i.e. $U_e(u_0)$).
- **Followee*:** We removed celebrities with more than 10,000 followers from followee lists. We assume less homophily between celebrities and fans.
- **Friend:** Users whom u_0 follows and also who follow u_0 (i.e. $U_f(u_0) = U_r(u_0) \cap U_e(u_0)$).
- **NonFriend:** Users who are either followers or followees of u_0 but not both (i.e. $U_f(u_0) = U_r(u_0) \Delta U_e(u_0)$).
- **Aggregated:** Aggregation of *Retweet*, *Mention*, *Follower* and *Followee*. *Friend* and *NonFriend* are excluded due to redundancy.
- **Aggregated*:** An aggregation of *Retweet*, *Mention*, *Follower* and *Followee**.

3.1.2 Data Cleaning. Raw data from Twitter can be noisy. Following rules are used to clean data from Twitter API: the first two rules filter out low quality user associations, and the last one normalizes name strings from user profiles.

Figure 1: Follower count distributions of *seed users*’ followers and followees. We characterize Twitter users with r , the ratio of follower over followee count. **Celebrity: $r > 10$. **Ordinary**: $r \leq 10$. More celebrities among followees. Homophily between fans and celebrities is not as strong as that between families and friends. So *Followee** removes names of celebrities to strengthen homophily among followee lists.**

- **Tweets:** Twitter API provides a small sample of real-time public Tweets¹. On average, we collect about 3.5M Tweets everyday. ~17% (0.6M) are retweets. ~54% (1.9M) contains at least one user mention. Remaining Tweets are filtered out.
- **Users:** In order to get lists of followers and followees, we choose a random set of Twitter users meeting following standards as *seed users*: (i) number of followers in range [50, 500]. (ii) number of followees in range [50, 500]. (iii) daily average posts less than 10. The motivation is to select Twitter users with enough social links but not celebrities nor social bots [20].
- **Names:** Twitter user names can be very noisy, e.g. random strings, misspelled words, emoji and notations. Therefore, we remove special symbols, punctuation and notations in various languages from names. We also filter out names without separators because it is not certain whether they are first or last names. Uncommon names with less than 5 occurrences are also removed.

3.1.3 Followers vs. Followees. After aggregating followers and followees separately, we find these two user groups are fundamentally different. As shown in Figure 1, we use a simple but effective way to characterize user, measuring the ratio of follower over followee (referred as r). User are assigned label *celebrity* if r is greater than 10, otherwise *ordinary*.

As shown in the right, almost half of the followees are celebrities. These celebrities tend to have more than 10,000 followers. We argue that the reason is Twitter allows one-way relation instead of reciprocal relations for *Facebook*. Therefore, ordinary users can follow celebrities they like, as well as their friends and family. As a consequence, these users have more celebrities among their followees and more friends/family among followers. We argue that homophily among friends/family is stronger than that among celebrities and their fans. We will show in Section *Experiments* that performances are improved after removing the celebrities.

¹https://developer.twitter.com/en/docs/tweets/sample-realtime/overview/GET_status_sample

Embedding	Vocab. Size	Corpus Size
<i>Retweet</i>	0.67M	53.61M
<i>Mention</i>	1.19M	174.30M
<i>Follower</i>	1.39M	140.69M
<i>Followee</i>	1.21M	204.40M
<i>Followee*</i>	1.19M	94.20M
<i>Friends</i>	0.77M	60.89M
<i>NonFriends</i>	1.30M	223.25M
<i>Aggregated</i>	3.01M	573.00M
<i>Aggregated*</i>	2.99M	508.13M
<i>Email</i>	4.10M	1140.00M

Table 2: Nine training corpora for Twitter name embeddings. *Email* is baseline corpus. Corpus size of *Followee is much smaller than *Followee*, while vocabulary size does not change much. *Aggregated** has similar vocabulary and corpus sizes as *Email*.**

3.1.4 Hyper Parameters. One of our goals is to compare the performance of Twitter embeddings with email embeddings learned in [37]. Therefore, the same experimental settings are used: skip-gram model with negative sample. Each name part consists of 100 dimensions. The size of moving window for context is 5 and 10 examples for negative sampling. We learn the embeddings for 20 epochs. Strings with less than five occurrences in corpus are ignored.

3.2 Experiments

3.2.1 Dataset. Two raw datasets have been collected from Twitter for experiments: (i) 286 million Tweets are collected from real-time stream sample from Jan. 15 to Mar. 21, 2018. (ii) a collection of 922,140 *seed users*’ full lists of followers and followees. *Seed users* are collected from real-time Tweet stream. 89 million unique user profiles are gathered to extract names of the followers and followees. As shown in Table 2, we prepared nine training corpora from this data. *Email* is the dataset used in [37] and it is collected from 57 million email users.

3.2.2 Performance Comparison. Name embeddings prove extremely useful for various tasks because they encode cultural signals of name parts implicitly in the distributed representations. Among the many latent signals, gender, ethnicity and nationalities are major ones that can easily be evaluated. We use ground truth labels from U.S. Census Bureau to measure whether same-gender and same-ethnicity names sit together in embedding space. 74 million name labels from [37] are used to compare classification performances on a 39-leaf nationality taxonomy. 80% of labels are used for training while 20% for testing.

Census 1990 contains ground truth labels of 1,219 male and 4,275 female first names. *Census 2000* provides ethnic distribution of 151,671 last names. We use the names that exist in vocabularies of all embeddings for fair comparison, resulting in 878 male and 3,479 female names for gender evaluation. 58,407 White, 2,519 Black, 4,521 API (Asian and Pacific Islander) and 5,346 Hispanic names are collected for ethnicity in the same manner.

Figure 2 compares performances on gender. *Mention* embeddings consistently outperform other embeddings by significant

Figure 2: Ratio of same-gender names among top k nearest neighbors ($k \in [1, 10, 50, 100]$) in name embedding spaces. *Mention* performs the best (avg. on female: 0.94, male: 0.74), reflecting stronger gender homophily in Twitter mentions. *Aggregated* outperforms *Email* on average (female: 0.94 vs. 0.91, male: 0.67 vs. 0.59). *Random* performances are proportional to the ratio of labeled female name count over male.

margins on both females and males. This suggests that gender bias in Twitter “mention” is much stronger than that in “retweet” and follower/followee relations. In other words, Twitter users are more likely to “@” others of same gender, who probably share similar interests or opinions. *Aggregated** has similar vocabulary size as *Email* and achieves better performances than *Email* for both genders. *Followee** gets a slightly smaller ratio than *followee* on female (avg. 0.88 vs. 0.91) but significantly better on male (avg. 0.65 vs. 0.50). *Random* embeddings mean that each name part is assigned a random name embedding such that names of each gender uniformly distributed in embedding space. Given a male name, for example, its nearest neighbors have almost the same distribution as the overall gender distribution of the label set. Therefore, we expect performances of male names to be lower than female, because there are far less male name labels (29% vs. 71%). We also use similar random embeddings for ethnicity evaluation.

Table 3 shows the ratios of same-ethnicity last names among their nearest neighbors. It is interesting to see *Mention* gets higher scores than *Retweet*, indicating more ethnic homophily in mentions. One possible explanation is users are more likely to mention or raise attention from their friends while retweeting or quoting from the famous ones. As we have shown in Figure 1, there are more celebrities among followees. So *Followee* has lower same-ethnicity ratios than *Followers*. After removing celebrity names, *Followee** outperforms similarly as *Follower*. The superior performances of *Followee** over *Followee* on both gender and ethnicities validate less homophily among celebrity-fan pairs and the effectiveness of removing celebrities. Therefore, *Aggregate** performs best among all Twitter embeddings, after combining training examples from *Followee** instead of *Followee*. *Email* gets highest ratio among all. *Black* names are harder to classify because they only take up 3.5% of all labels.

Embedding	White	Black	API	Hisp.	Avg.
<i>Random</i>	0.82	0.04	0.06	0.08	0.25
<i>Retweet</i>	0.92	0.20	0.57	0.64	0.58
<i>Mention</i>	0.93	0.22	0.61	0.71	0.62
<i>Follower</i>	0.94	0.31	0.77	0.86	0.72
<i>Followee</i>	0.92	0.27	0.72	0.81	0.68
<i>Followee*</i>	0.94	0.31	0.77	0.84	0.72
<i>Friends</i>	0.93	0.28	0.74	0.81	0.69
<i>NonFriends</i>	0.92	0.26	0.71	0.82	0.68
<i>Aggregated</i>	0.93	0.32	0.76	0.83	0.71
<i>Aggregated*</i>	0.94	0.33	0.79	0.86	0.73
<i>Email</i>	0.96	0.47	0.83	0.87	0.78

Table 3: Ratios of same-ethnicity names among nearest neighbors. *Aggregated achieves highest ratios among all Twitter embeddings and gets comparable performance comparing to *Email*. *Follower* outperforms *Followee*, while *Followee** has the same average ratio as *Follower*, which validates that removing celebrities is effective. (API: Asian and Pacific Islander)**

To make a fair comparison on nationality performance, we adopt the same classification method, experiment settings and label data as in [37]. Table 4 shows that *Aggregated** (*Tw.*) has similar performance as *Email* (*Em.*). For some classes, like Thailand, Baltic and Pakistan, *Aggregated** outperforms *Email* embeddings. *Email* performs slightly better than *Twitter* w.r.t. weighted average F1 score on 39 leaf classes. We also noticed that the performances are highly dependent on the size of data. For less developed places like Cambodia and countries in central Asian and Maghreb, very limited user associations and labels are collected. Therefore, their F1 scores are much below average performance.

4 LIFESPAN MODELING

The strength of name embeddings lies in the implicit signals encoded in distributed representations. These signals come from concurrences of names, or more accurately, social interactions between individuals (e.g. Tweets). These signals are useful for many downstream tasks. In this section, we demonstrate the power of name embeddings in modeling lifespan, where gender, ethnicity and nationality all are contributing factors.

4.1 Social Security Death Index Dataset

The Social Security Death Index (SSDI) is maintained and distributed by the Social Security Administration to prevent identity fraud associated with using identifiers of deceased individuals. The SSDI has also been employed in hundreds of academic research associated with medical and demographic analysis, such as [8, 32]. The research applicability of the SSDI compared to other resources has been studied in [29, 36].

Each record in the SSDI consists of an individual’s full name, their date of birth and death, and their social security number (SSN). The dataset we studied contains 85,822,194 death records. Our analysis

Nationality	Name#	Em.	Tw.	Nationality	Name#	Em.	Tw.	Nationality	Name#	Em.	Tw.
CelticEnglish*	3505K	0.73	0.70	Muslim	1475K	0.74	0.73	Jewish*	11K	0.40	0.37
SouthAsian*	2623K	0.89	0.88	African	606K	0.59	0.56	EastAsian	6157K	0.92	0.91
Hispanic	6892K	0.91	0.89	Greek*	259K	0.89	0.87	Nordic	195K	0.73	0.70
Europe	5371K	0.84	0.81								
EastEurope*	65K	0.49	0.49	Nubian*	577K	0.65	0.62	Maghreb*	47K	0.15	0.14
SouthKorea*	68K	0.86	0.83	Malay	2596K	0.86	0.84	Chinese*	2901K	0.93	0.92
Portuguese*	2683K	0.89	0.87	Turkic	78K	0.68	0.66	Pakistanis	179K	0.51	0.50
Philippines*	1137K	0.72	0.69	Persian*	423K	0.66	0.64	Spanish*	3072K	0.85	0.83
Scandinavian	165K	0.70	0.67	Finland*	30K	0.74	0.72	German*	1278K	0.74	0.70
WestAfrican*	315K	0.56	0.54	Baltics*	12K	0.41	0.42	Japan*	65K	0.84	0.78
SouthAfrican*	66K	0.37	0.36	Russian*	121K	0.72	0.72	Arabia*	172K	0.51	0.51
EastAfrican*	225K	0.57	0.53	French*	2674K	0.83	0.80	Indochina	528K	0.90	0.87
SouthSlavs*	68K	0.57	0.54	Italian	1153K	0.75	0.72				
Cambodia*	1K	0.16	0.05	Turkey*	75K	0.69	0.68	Sweden*	74K	0.61	0.58
Bangladesh*	78K	0.58	0.56	Vietnam*	502K	0.91	0.89	Thailand*	18K	0.59	0.67
Malaysia*	242K	0.48	0.45	Pakistan*	101K	0.45	0.50	Denmark*	49K	0.66	0.63
CentralAsian*	3K	0.20	0.16	Italy*	825K	0.71	0.68	Romania*	329K	0.66	0.64
Indonesia*	2354K	0.87	0.84	Norway*	42K	0.62	0.59	Myanmar*	7K	0.61	0.58
Weighted Avg.	—	0.81	0.79								

Table 4: Nationality classification performances (f1 scores) of *Email (Em.)* and *Twitter (aggregated*, Tw.)* embeddings on a 39-leaf nationality taxonomy. The taxonomy has three levels, which are separated with bolder lines. ‘*’ marks leaf nationalities. *Weighted Avg.* is count-weighted average F1 score of leaf nationalities. Twitter embeddings achieve comparable performances on nationality classification.

B	S	G	E	N	NoEbd	ShEbd	EmEbd	TwEbd
					13.418	13.423	12.781	12.747
√					8.052	8.049	7.792	7.800
	√				13.373	13.369	12.765	12.742
		√			13.150	13.135	12.768	12.745
			√		13.314	13.309	12.761	12.721
				√	13.271	13.268	12.765	12.734
√	√	√	√	√	7.775	7.777	7.739	7.744

Table 5: Average Prediction Error (in years) of seven sets of models using different features. The demographic features are: birth year (B), state (S), gender (G), ethnicity (E), nationality (N). Extra features include: no embedding (NoEbd), shuffled embedding (ShEbd), Email embedding (EmEbd), Twitter embedding (TwEbd). Each number (or prediction error) is the average of 20 runs. Birth year is the most important feature due to survivorship bias. Using name embeddings improves performance significantly.

was performed on the master file of November 30, 2011², using a random sample of 2,991,927 records for experiments.

Figure 3 (Top) presents the number of SSDI records by birth year, further broken down by gender. The peak of the distribution was born between 1910 and 1930. Before this peak, women outnumber men in the database, a consequence of more of them surviving to be issued social security cards. Men and women have represented with equal frequency since approximately 1945. Figure 3 (Bottom) presents the average lifespan of SSDI records by birth year, further

broken down by gender. Survivorship biases account for this strange distribution. The earliest records have an average lifespan above 90, reflecting that they had to live long enough to receive identification numbers (i.e. survivorship bias). The average lifespan has decreased almost linearly since 1940, and equally for women as for men. We anticipate that these totals will increase and diverge with time, as the distribution moves beyond the prematurely deceased.

4.2 Demographic Features

We extracted/inferred following discrete demographic features from each SSDI record, of the type which are traditional for lifespan models. The classifiers for gender, ethnicity and nationality predictions are introduced in Section *Related Work*.

- **Birth year:** Birth years are represented by 130 binary features. Each one corresponds to a birth year.
- **States:** We infer states using first three digits of SSN. In total, 59 possible binary state/territory features are extracted.
- **Gender:** Gender is inferred with a classifier based on U.S. census data.
- **Ethnicity:** We use *NamePrism* to predict ethnicity based on names.
- **Nationality:** *NamePrism* is also used to predict nationality based on names.

4.3 Linear Regression Models

To evaluate the power of name embeddings for predicting lifespan, we build 32 (i.e. 2⁵) sets of models using linear regression. Each set is trained on a particular subset of the 5 demographic features described above. The four models of each set are distinguished by

²Dataset available: <http://ssdmf.info/>

Figure 3: Distributions of SSDI records. *Top*: the number of records sorted by birth year. Most were born between 1910 and 1930. *Bottom*: the average lifespan by birth year. Survivorship bias causes unusually long lifespan in the beginning, while prematurely deceased ones make decreasing lifespans at the end of the curves.

whether they use no embedding features (*NoEbd*), Twitter name embeddings (*TwEbd*), Email name embeddings (*EmEbd*), or a randomly shuffled permutation of Twitter embeddings to add dimensionality without additional information (*ShEbd*), as a control.

Let X notate the feature vectors, and Y be the ground-truth lifespans. y_i denotes the lifespan of i_{th} record and $\hat{y}_i$ is the predicted lifespan using i_{th} feature vector x_i . Then $|y_i - \hat{y}_i|$ is the error made by prediction (in years). We seek the coefficients w to optimize following loss function:

$$\min_w \sum_i (y_i - wx_i)^T (y_i - wx_i) + \lambda ||w||_2^2 \quad (2)$$

Here λ is the constant governing the strength of the regularization term, to guard against overfitting. We observed that the performances are not sensitive to λ and it is empirically assigned 0.003 for all regression models.

4.4 Performance Analysis

We use 90% records as training data and use the rest as testing data. Table 5 presents the average test error of 20 runs after random divisions of training and testing data. Due to survivorship bias, the most powerful single feature is the birth year, which yields an absolute error of 8.052 years. The strength of birth year feature separates models into two groups, with/without birth year (Figure 4).

Figure 4: Visualization of prediction errors (in years) of 32 sets of models (two zoom-in inset figures). Using name embeddings (blue and green dots) improve performances significantly (all p values smaller 0.01 under Welch’s t-test). Red dots are on the line, $y = x$, reflecting training without embedding has similar performances as training with shuffled embeddings. Therefore, the gains of name embeddings come from the latent signals captured by name embeddings, instead of the increase of feature dimensions.

To understand the effect of embeddings on lifespan models, we visualize results of the 32 model sets as points in Figure 4. The performance is strikingly linear, and adding name embeddings improves the models of each of the 32 possible variable settings. We also conduct four significance tests using *Welch’s t-test* with following null hypotheses: (i) the means of *NoEbd* and *ShEbd* are the same; (ii) the means of *TwEbd* and *EmEbd* are the same (iii) the means of *NoEbd* and *EmEbd* are the same; (iv) the means of *NoEbd* and *TwEbd* are the same. The p values of first two tests under all variable settings are larger than 0.01, and p-values of the last two tests are smaller than 0.01.

The significance tests results show that (i) name embeddings further improve the performances of demographic lifespan models; (ii) the improvements come from latent signals encoded in name embeddings instead of the increase of feature dimensions. (iii) *Twitter* embeddings have similar performances with *Email* embeddings on lifespan modeling.

4.5 Latent Factors

Name embeddings are powerful at capturing latent properties of class and cultural group dynamics, but the nature of these properties remains hidden within unlabeled dimensions. This makes it difficult to determine exactly what properties they are keying on for a particular model. To provide some insight into how names affect lifespans, we identified the most favorable/unfavorable first and last names through residual analysis.

4.5.1 Residual Analysis. Name embeddings encodes various demographic signals, including gender, ethnicity, nationalities and other latent signals. In our best lifespan model, we combine the explicit demographic features with name embeddings. Therefore,

Short	Count	Gain	Formal	Count	Gain
Gust	5429	-1.965	Gustav	8639	-1.353
Wm	6684	-1.623	William	773086	-0.581
Gus	10439	-1.597	Angus	1965	-0.091
Hans	10599	-1.322	Johannes	1153	-1.053
Alex	24520	-1.297	Alexander	34265	-1.210
Dan	12368	-1.296	Daniel	55567	-0.559
Guy	28664	-1.204	Guido	1694	-0.417
Effie	33844	-1.195	Euphemia	753	0.073
Average		-1.437	Average		-0.649

Table 6: 8 out of 20 most unfavorable first names are in diminutive forms. In contrast, their corresponding formal names have larger gains (in years). A systematic study on 155 diminutive/formal name pairs proves bias favoring formal names, suggesting two groups representing different socioeconomic classes.

there are redundant signals in the input features. In order to identify latent signals encoded in name embeddings, we conduct residual analysis with following steps: (i) train a linear regression model (referred as “demographic model”) using demographic features and ground-truth lifespan; (ii) train linear regression model (referred as “residual model”) using name embeddings as features and residuals (i.e. prediction errors) as target values.

More formally, the demographic model tries to minimize loss function $L_d(\mathbf{w}_d | \mathbf{X}_d, Y_{ls})$, where Y_{ls} is the ground truth lifespans and $\mathbf{X}_d$ is demographic feature vectors. Let $\hat{Y}_{ls}$ be the prediction lifespan made by demographic model, i.e. $\hat{Y}_{ls} = \mathbf{w}_d \cdot \mathbf{X}_d$ (the intercept term is ignored for brevity), then the residual model minimizes $L_r(\mathbf{w}_r | \mathbf{X}_e, Y_{ls} - \hat{Y}_{ls})$. $\mathbf{X}_e$ is name embedding features. $L_d(\cdot)$ and $L_r(\cdot)$ are in the same form as Equation 2. Finally, we use $\hat{Y}_r = \mathbf{w}_r \cdot \mathbf{X}_e$ to compute gains of name parts. If a name part gets positive gain (i.e. *favorable name*), it means individuals with this name tend to live longer. In the opposite, *unfavorable names* get negative gains.

4.5.2 Diminutive vs Formal First Names. Among the records with birth year between 1880 to 1910 (less influenced by survivorship bias, see Figure 3), 8 out of 20 most unfavorable first names occurring more than 5000 times are in diminutive form (see Table 6). It is interesting that we find the gains of responding long-version names are significantly larger. We suspect that the distinction captured here is one of socioeconomic class because formal names might be generally expected to appear in official documents more often.

More systematically, we test on 155 pairs of diminutive and formal English names from Wikipedia³. It turned out that 114 pairs (74%) agrees with this observation when using email embeddings, namely names in formal forms get larger gains compared to diminutives. Similarly, 60% pairs favor formal names using Twitter embeddings. Under the null hypothesis that there is no bias toward short forms or formal ones, email and Twitter embeddings both get p values smaller than 0.01. The null hypothesis is rejected.

4.5.3 Fine-grained Subgroups: Table 7 shows that, among birth year between 1880 to 1910, 17 out of the 20 most favorable last

Jewish			Scandinavian		
Name	Count	Gain	Name	Count	Gain
Katz	7143	1.12	Svensson	127	1.23
Bernstein	5122	1.11	Olsson	647	1.16
Shapiro	6156	1.06	Johansson	494	1.14
Solomon	6232	1.03	Persson	507	1.13
Goldman	5502	1.03	Karlsson	129	1.13
Levy	8739	1.01	Nilsson	682	1.01
Feldman	5031	1.01	Larsson	207	0.90
Friedman	8865	0.97	Karlsen	145	0.79
Rosenberg	6448	0.95	Kristiansen	175	0.57
Goldstein	8837	0.95	Andersen	6196	0.43
Cohen	22273	0.92	Christensen	10877	0.43
Stern	5537	0.89	Rasmussen	5813	0.33
Greenberg	6020	0.88	Pedersen	3782	0.33
Goldberg	9306	0.86	Larsen	9488	0.31
Levine	8518	0.86	Hansen	23578	0.29
Rosen	5363	0.85	Nielsen	7192	0.28
Kessler	5240	0.65	Olsen	11080	0.23

Table 7: 17 out of 20 most favorable last name with more than 5000 occurrences have Ashkenazi Jewish origin. Most popular Scandinavian last names get positive gains (in years). Both populations have longer lifespans than the average in U.S. Name embeddings are able to capture such fine-grained distinctions between groups.

names are Ashkenazi Jewish. This phenomenon is interesting because Jewish people have long lived in many countries so no single ethnicity or nationality feature could capture this group well. However in name embedding space, similar names have similar representations, because of communication homophily. The observation that these Jewish people had longer life expectancy agrees with the observation made by Institute for Jewish Policy Research[3]. It is also interesting to see that popular Scandinavian last names⁴ all get positive gains.

5 CONCLUSION

Name embeddings prove more effective feature representations of names than traditional substrings. However, existing *Email* name embeddings are not publicly accessible. In this paper, we present a new way to learn name embeddings from Twitter. Extensive experiment results show the power of Twitter embeddings on gender, ethnicity, nationality. We release Twitter name embeddings to support research communities (www.name-prism.com).

We also demonstrate that name embeddings can improve the accuracy of lifespan models. Extrapolating from these results, we believe they can be used to strengthen predictive models for related tasks in the social, economic, and medical sciences. This is particularly true in large-scale but data-poor studies, where the name must serve as a proxy for reported gender, ethnicity, or nationality. The exact nature of the hidden factors implicitly encoded within our name embeddings that provide this predictive power is an exciting open question for further research. We presume that this

³<https://en.wikipedia.org/wiki/Hypocorism>

⁴https://en.wikipedia.org/wiki/Scandinavian_family_name_etymology

includes subtle class-based distinctions (e.g. socioeconomic status and fine-grained groups) which are hidden by the coarse categorical variables traditionally observed and recorded.

ACKNOWLEDGMENTS

The authors thank the reviewers for their useful comments. This work was partially supported by NSF grant IIS-1546113. Any conclusions expressed in this material are of the authors' and do not necessarily reflect the views, either expressed or implied, of the funding party.

REFERENCES

- [1] Ernest L Abel and Michael L Kruger. 2006. Nicknames increase longevity. *OMEGA-Journal of Death and Dying* 53, 3 (2006), 243–248.
- [2] Ernest L Abel and Michael L Kruger. 2009. Athletes, doctors, and lawyers with first names beginning with "D" die sooner. *Death studies* 34, 1 (2009), 71–81.
- [3] Sarah Abramson, David Graham, and Jonathan Boyd. 2011. Key trends in the British Jewish community: A review of data on poverty, the elderly and children. (2011).
- [4] Faiyaz Al Zamil, Wendy Liu, and Derek Ruths. 2012. Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors. *ICWSM* 270 (2012), 2012.
- [5] Bedoor K AlShebli, Talal Rahwan, and Wei Lee Woon. 2018. The preeminence of ethnic diversity in scientific collaboration. *Nature communications* 9, 1 (2018), 5163.
- [6] Anurag Ambekar, Charles Ward, Jahangir Mohammed, Swapna Male, and Steven Skiena. 2009. Name-ethnicity classification from open sources. In *SIGKDD*. ACM, 49–58.
- [7] Osei Appiah. 2001. Ethnic identification on adolescents' evaluations of advertisements. *Journal of Advertising Research* 41, 5 (2001), 7–22.
- [8] Eric Backlund, Paul D Sorlie, and Norman J Johnson. 1996. The shape of the relationship between income and mortality in the United States: evidence from the National Longitudinal Mortality Study. *Annals of Epidemiology* 6, 1 (1996), 12–20.
- [9] Yambazi Banda, Mark N Kvale, Thomas J Hoffmann, Stephanie E Hesselson, Dilrini Ranatunga, Hua Tang, Chiara Sabatti, Lisa A Croen, Brad P Dispensa, Mary Henderson, et al. 2015. Characterizing race/ethnicity and genetic ancestry for 100,000 subjects in the Genetic Epidemiology Research on Adult Health and Aging (GERA) cohort. *Genetics* 200, 4 (2015), 1285–1295.
- [10] Donald A Barr. 2014. *Health disparities in the United States: Social class, race, ethnicity, and health*. JHU Press.
- [11] Robert W Buechley. 1976. Generally useful ethnic search system: GUESS. In *Annual Meeting of the American Names Society*.
- [12] Esteban González Burchard, Elad Ziv, Eliseo J Pérez-Stable, and Dean Sheppard. 2003. The importance of race and ethnic background in biomedical research and clinical practice. *The New England journal of medicine* 348, 12 (2003), 1170.
- [13] Jonathan Chang, Itamar Rosenn, Lars Backstrom, and Cameron Marlow. 2010. ePluribus: Ethnicity on Social Networks. In *ICWSM*, Vol. 10. 18–25.
- [14] Alexander Cochardt, Stephan Heller, and Vitaly Orlov. 2018. In Military We Trust: The Effect of Managers' Military Background on Mutual Fund Flows. *Available at SSRN 3303755* (2018).
- [15] Andrew J Coldman, Terry Braun, and Richard P Gallagher. 1988. The classification of ethnic status using name information. *Journal of epidemiology and community health* 42, 4 (1988), 390–395.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- [17] Will Gornall and Ilya A Strebulaev. 2018. Gender, Race, and Entrepreneurship: A Randomized Field Experiment on Venture Capitalists and Angels. *Available at SSRN 3301982* (2018).
- [18] Shuchu Han, Yifan Hu, Steven Skiena, Baris Coskun, Meizhu Liu, Hong Qin, and Jaime Perez. 2017. Generating Look-alike Names For Security Challenges. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*. ACM, 57–67.
- [19] Seeromanie Harding, Howard Dews, and Stephen Ludi Simpson. 1999. The potential to identify South Asians using a computerised algorithm to classify names. *Population Trends London* (1999), 46–49.
- [20] Xia Hu, Jiliang Tang, Yanchao Zhang, and Huan Liu. 2013. Social Spammer Detection in Microblogging. In *IJCAI*, Vol. 13. 2633–2639.
- [21] John T Jones, Brett W Pelham, Mauricio Carvallo, and Matthew C Mirenberg. 2004. How do I love thee? Let me count the Js: implicit egotism and interpersonal attraction. *Journal of personality and social psychology* 87, 5 (2004), 665.
- [22] Jure Leskovec and Eric Horvitz. 2008. Planetary-scale views on a large instant-messaging network. In *WWW*. ACM, 915–924.
- [23] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [24] Leif D Nelson and Joseph P Simmons. 2007. Moniker maladies when names sabotage success. *Psychological Science* 18, 12 (2007), 1106–1112.
- [25] Catherine O'Neill. 2016. Weapons of Math Destruction. *How Big Data Increases Inequality and Threatens Democracy* (2016).
- [26] Laura Pinzur and Gary Smith. 2009. First names and longevity. *Perceptual and motor skills* 108, 1 (2009), 149–160.
- [27] Daniel Preotiuc-Pietro and Lyle H. Ungar. 2018. User-Level Race and Ethnicity Predictors from Twitter Text. In *Proceedings of the 27th International Conference on Computational Linguistics*, Emily M. Bender, Leon Derczynski, and Pierre Isabelle (Eds.). Association for Computational Linguistics, 1534–1545.
- [28] James Quesada, Laurie Kain Hart, and Philippe Bourgois. 2011. Structural vulnerability and health: Latino migrant laborers in the United States. *Medical Anthropology* 30, 4 (2011), 339–362.
- [29] Janet W Rich-Edwards, Karen A Corsano, and Meir J Stampfer. 1994. Test of the national death index and equifax nationwide death search. *American Journal of Epidemiology* 140, 11 (1994), 1016–1019.
- [30] Sang-Chun Shin and Sung-Je Cho. 2014. The Impact of Names upon the Stress and Self-esteem Before and After Renaming. *Journal of the Korea Academia-Industrial cooperation Society* 15, 5 (2014), 2662–2670.
- [31] Gary Smith. 2012. Do People Whose Names Begin with âĀĬJDâĀĬ Really Die Young? *Death studies* 36, 2 (2012), 182–189.
- [32] Ian M Thompson Jr, Phyllis J Goodman, Catherine M Tangen, Howard L Parnes, Lori M Minasian, Paul A Godley, M Scott Lucia, and Leslie G Ford. 2013. Long-term survival of participants in the prostate cancer prevention trial. *New England Journal of Medicine* 369, 7 (2013), 603–610.
- [33] Pucktada Treeratpituk and C Lee Giles. 2012. Name-ethnicity classification and ethnicity-sensitive name matching. In *AAAI*.
- [34] Michal Vaanunu and Chen Avin. 2018. Homophily and Nationality Assortativity Among the Most Cited Researchers' Social Network. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 584–586.
- [35] Buvaneshwaran Venugopal. 2017. Homophily, Information Asymmetry and Performance in the Angels Market. (2017).
- [36] Brent C Williams, Lucy B Demitrack, and Brant E Fries. 1992. The accuracy of the National Death Index when personal identifiers other than Social Security number are used. *American Journal of Public Health* 82, 8 (1992), 1145–1147.
- [37] Junting Ye, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven Skiena. 2017. Nationality Classification Using Name Embeddings. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1897–1906.