Research Article

Optimal Proxy Selection for Socioeconomic Status Inference on Twitter

Jacob Levy Abitbol ¹, Eric Fleury,² and Márton Karsai ¹

¹Univ Lyon, Inria, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, F-69007 Lyon, France

²Inria, F-75012 Paris, France

Correspondence should be addressed to Márton Karsai; marton.karsai@ens-lyon.fr

Received 25 January 2019; Accepted 15 April 2019; Published 19 May 2019

Guest Editor: Xin Huang

Copyright © 2019 Jacob Levy Abitbol et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Individual socioeconomic status inference from online traces is a remarkably difficult task. While current methods commonly train predictive models on incomplete data by appending socioeconomic information of residential areas or professional occupation profiles, little attention has been paid to how well this information serves as a proxy for the individual demographic trait of interest when fed to a learning model. Here we address this question by proposing three different data collection and combination methods to first estimate and, in turn, infer the socioeconomic status of French Twitter users from their online semantics. We assess the validity of each proxy measure by analyzing the performance of our prediction pipeline when trained on these datasets. Despite having to rely on different user sets, we find that training our model on professional occupation provides better predictive performance than open census data or remote sensed expert annotation of habitual environments. Furthermore, we release the tools we developed in the hope it will provide a generalizable framework to estimate socioeconomic status of large numbers of Twitter users as well as contribute to the scientific discussion on social stratification and inequalities.

1. Introduction

Over the last decade the emergence of online social services changed the way we diffuse or acquire information and how we interact with each other. Every day billions of individuals use such services, while their penetration in our everyday lives seems ever-growing. In turn, online activities generate a massive volume of publicly available data, which are open to analysis, and fuel new data-driven developments in industry and research. These advances have led to the paradigm shift in the design of marketing strategies [1], emergence of new services, and open the door for data-driven reasoning on social phenomena relying on society-large observations [2]. The digital footprints left across these multiple media platforms provide us with a unique source to study patterns in human behavior down to the individual level, e.g., to understand how the linguistic phenotype of a given user is related to social attributes such as socioeconomic status (SES).

The quantification and inference of SES of individuals is a long-standing question in the social sciences. It is a

rather difficult problem as it may depend on a combination of individual characteristics and environmental variables [3]. Some of these features can be easier assessed like income, gender, or age whereas others, relying to some degree on self-definition and sometimes entangled with privacy issues, are harder to assign like ethnicity, occupation, education level, or home location. Furthermore, individual SES correlates with other individual or network attributes, as users tend to build social links with others of similar SES, a phenomenon known as status homophily [4], arguably driving the observed stratification of society [5]. At the same time, shared social environment, similar education level, and social influence have been shown to jointly lead socioeconomic groups to exhibit stereotypical behavioral patterns, such as shared political opinion [6] or similar linguistic patterns [7]. Although these features are entangled and causal relation between them is far from understood, they appear as correlations in the data.

Datasets recording multiple characteristics of human behavior are more and more available due to recent developments in data collection technologies and increasingly popular online platforms and personal digital devices. The

automatic tracking of online activities (commonly associated with profile data and meta-information), the precise recording of interaction dynamics and mobility patterns collected through mobile personal devices, together with the detailed and expert annotated census data all provide new grounds for the inference of individual features or behavioral patterns [2]. The exploitation of these data sources has already been proven to be fruitful as cutting edge recommendation systems, advanced methods for health record analysis, or successful prediction tools for social behavior heavily rely on them [8]. Nevertheless, despite the available data, some inference tasks, like individual SES prediction, remain an open challenge.

The precise inference of SES would contribute to overcome multiple scientific challenges and could potentially have multiple commercial applications [9]. Further, robust SES inference would provide unique opportunities to gain deeper insights into socioeconomic inequalities [10], social stratification [5], and into the mechanisms driving network evolution, such as status homophily or social segregation.

In this work, we take a horizontal approach to this problem and explore various ways to infer the SES of a large sample of social media users. We propose different data collection and combination strategies using open, crawlable, or expert annotated socioeconomic data for the prediction task. Specifically, we use an extensive Twitter dataset of 1.3M users located in France, all associated with their tweets and profile information, 32,053 of them having inferred home locations. Individual SES is estimated by relying on three separate datasets, namely, socioeconomic census data, crawled profession information, and expert annotated Google Street View images of users' home locations. Each of these datasets is then used as ground-truth to infer the SES of Twitter users from profile and semantic features similar to [11]. We aim to explore and assess how the SES of social media users can be obtained and how much the inference problem depends on annotation and the user's individual and linguistic attributes. In addition, to demonstrate the power of our location inference method, we group users into nine distinct socioeconomic classes to identify correlations between the predictability of mobility [12] of geolocated users and their socioeconomic status. We observe that as a user's SES increases, so does his/her radius of gyration which in turn lowers the upper bound of predictability of his/her whereabouts.

We provide in Section 2 an overview of the related literature to contextualize the novelty of our work. In Section 3 we provide a detailed description of the data collection and combination methods including analysis of Twitter, census, mobility, occupation, and home location data. In Section 4 we introduce the features extracted to solve the SES inference problem, with results summarized in Section 5. Finally, in Sections 6 and 7 we conclude our paper with a brief discussion of the limitations and perspectives of our methods. Note that this paper is partially based on results published in a proceeding paper [13], while the source code of the reported methods has recently been released [14]. In terms of novelty, we outline the following contributions. (i) We introduced a new analytical framework to understand

the relationship between mobility and socioeconomic status as well as its effect on our inference task; (ii) we provided a detailed analysis of the performance of our residential location filtering procedure; (iii) we delved into the study of our predictive performance, studying the set of features that were most determinant in inferring socioeconomic status; and finally (iv) we released to the scientific community all the pipelines used in this study to ease the collection and study of similar datasets as well as to motivate further study in this area.

2. Related Works

There is a growing effort in the field to combine online behavioral data with census records, and expert annotated information to infer social attributes of users of online services. The predicted attributes range from easily assessable individual characteristics such as age [15], or occupation [11, 16–18], to more complex psychological and sociological traits like political affiliation [19], personality [20], or SES [11, 21].

Predictive features proposed to infer the desired attributes are also numerous. In case of Twitter, user information can be publicly queried within the limits of the public API [22]. User characteristics collected in this way, such as profile features, tweeting behavior, social network, and linguistic content, have been used for prediction, while other inference methods relying on external data sources such as website traffic data [23] or census data [24, 25] have also proven effective. Nonetheless, only recent works involve user semantics in a broader context related to social networks, spatiotemporal information, and personal attributes [11, 17, 18, 26].

In this framework, aggregated studies of user spatial data have been particularly useful in fueling the analysis of human mobility patterns. Early work on mobile communication datasets [12, 27] showed that individuals tend to return to a few highly frequented locations leading to a high predictability in human travelling patterns. Analogous behavior was later exposed in Twitter [28], enabling the use of this social media platform as a proxy for tracking and predicting human movement.

Akin to this strand of research, user features collected from online platforms have also been used in the inference of individual demographic traits. The tradition of relating SES of individuals to their language dates back to the early stages of sociolinguistics where it was first shown that social status reflected through a person's occupation is a determinant factor in the way language is used [29]. This line of research was recently revisited by Lamos et al. to study the SES inference problem on Twitter. In a series of works [11, 17, 18, 26], the authors applied Gaussian Processes to predict user income, occupation, and socioeconomic class based on demographic, psycholinguistic features, and a standardized job classification taxonomy, which mapped Twitter users to their professional occupations. The high predictive performance has proven this concept with $r = 0.633$ for income prediction, and a precision of 55% for 9-ways SOC classification, and 82% for binary SES classification. Nevertheless, the models developed by the authors are learned by relying on datasets, which were manually labeled through an annotation

process crowdsourced through Amazon Mechanical Turk at a high monetary cost. Although the labeled data has been released and provides the base for new extensions [15], it has two potential shortfalls that need to be acknowledged. First, the method requires access to a detailed job taxonomy, in this case specific to England, which hinders potential extensions of this line of work to other languages and countries. Furthermore, the language to income pipeline seems to show some dependency on the sample of users that actively chose to disclose their profession in their Twitter profile. Features obtained on this set might not be easily recovered from a wider sample of Twitter users. This limits the generalization of these results without assuming a costly acquisition of a new dataset.

3. Data Collection and Combination

Our first motivation in this study was to overcome earlier limitations by exploring alternative data collection and combination methods. We study here three ways to estimate the SES of Twitter users by using (a) open census data, (b) crawled and manually annotated data on professional skills and occupation, and (c) expert annotated data on home location Street View images. We provide here a collection of procedures that enable interested researchers to introduce predictive performance and scalability considerations when interested in developing language to SES inference pipelines. In the following we present in detail all of our data collection and combination methods.

3.1. Twitter Corpus. Our central dataset was collected from Twitter, an online news and social networking service. Through Twitter, users can post and interact by “tweeting” messages with restricted length. Tweets may come with several types of metadata including information about the author’s profile and the detected language as well as where and when the tweet was posted. Specifically, we recorded 90,369,215 tweets written in French, posted by 1.3 Million users in the time zones GMT and GMT+1 over one year (between August 2014 and July 2015) (the collection of this dataset was approved by the Ethic Committee of the hosting academic institute of the authors.). These tweets were obtained via the Twitter Powertrack API provided by Datasift with an access rate of 15%. Using this dataset we built various other corpora.

3.1.1. Geolocated Users. To find users with a representative home location we followed the method published in [30, 31]. As a bottom line, we concentrated on 127,614 users who posted at least five geolocated tweets with valid GPS coordinates, with at least three of them within a valid census cell (for definition see later), and over a longer period than seven days. Applying these filters we obtained 1,000,064 locations from geolocated tweets. By focusing on the geolocated users, we kept those with limited mobility, i.e., with median distance between locations not greater than 30 km, with tweets posted at places and times, which did not require travel faster than 130 km/h (maximum speed allowed within France), and with no more than three tweets within a two seconds window. We

further filtered out tweets with coordinates corresponding to locations referring to places (such as “Paris” or “France”). Thus, we removed locations that did not exactly correspond to GPS-tagged tweets and also users, which were most likely bots.

Home location was estimated by the most frequent location for a user among all coordinates she visited. This way we obtained 32,053 users, each associated with a unique home location. Finally, we collected the latest 3,200 tweets from the timeline of all of geolocated users using the Twitter public API [22]. Note that by applying these consecutive filters we obtained a more representative population as the Gini index, indicating overall socioeconomic inequalities, was 37.3% before filtering becoming 36.4% due to the filtering methods, which is closer to the value reported by the World Bank (33.7%) [32].

To verify our results, we computed the average weekday and weekend distance from each recorded location of a user to her inferred home location defined either as her most frequent location overall or among locations posted outside of work-hours from 9AM to 6PM (see Figures 1(a) and 1(b)). This circadian pattern displays great similarity to earlier results [31] with two maxima, roughly corresponding to times at the workplace, and a local minimum at 1PM due to people having lunch at home for locations posted in weekdays. Moreover, when comparing the weekday and weekend patterns of behavior, we notice that the average distance to the inferred home location seemed to be smaller when considering only geolocations posted during the weekend, most likely due to the absence of the home-to-work commuting during the weekend. We found that this circadian pattern was more consistent with earlier results [31] when we considered all geolocated tweets (“All” in Figure 1(a)) rather than only tweets including “home-related” expressions (“Night” in Figure 1(a)). To further verify the inferred home locations, for a subset of 29,389 users we looked for regular expressions in their tweets that were indicative of being at home [31], such as “chez moi”, “bruit”, “dormir” or “nuit”. In Figure 1(c) we show the temporal distribution of the rate of the word “dormir” at the inferred home locations. This distribution appears with a peak around 10PM, which is very different from the overall distribution of geolocated tweets throughout the day considering any location (see Figure 1(b)).

3.1.2. Linguistic Data. To obtain meaningful linguistic data we preprocessed the incoming tweet streams in several ways. As our central question here deals with language semantics of individuals, retweets do not bring any additional information to our study, thus we removed them by default. We also removed any expressions considered to be semantically meaningless like URLs, emoticons, mentions of other users (denoted by the @ symbol), and hashtags (denoted by the # symbol) to simplify later postprocessing. In addition, as a last step of textual preprocessing, we downcased and stripped the punctuation from the text of every tweet.

3.2. Census Data. Our first method to associate SES to geolocated users builds on an open census income dataset at intraurban level for France [33]. Obtained from 2010

FIGURE 1: Average distance from home of active users per hour of the day depending on whether the geolocated tweet was posted during a weekday (a) or the weekend (b) and the most frequent location was chosen among all (blue) or only the night ones (red). (c) Hourly rate of all geolocated tweets and (d) geolocated tweets mentioning ‘dormir’ averaged over all weekdays.

French tax returns, it was released in December 2016 by the National Institute of Statistics and Economic Studies (INSEE) of France. This dataset collects detailed socioeconomic information of individuals at the census block level (called IRIS), which are defined as territorial cells with varying size but corresponding to blocks of around 2,000 inhabitants, as shown in Figure 2 for greater Paris. For each cell, the data records the deciles of the income distribution of inhabitants. Note that the IRIS data does not provide full coverage of the French territory, as some cells were not reported to avoid identification of individuals (in accordance with current privacy laws), or to avoid territorial cells of excessive area. Nevertheless, this limitation did not hinder our results significantly as we only considered users who posted at least three times from valid IRIS cells, as explained in Section 3.1.1.

To associate a single income value to each user, we identified the cell of their estimated home locations and assigned them with the median of the corresponding income distribution. Thus we obtained an average socioeconomic indicator for each user, which was distributed heterogeneously in

accordance with Pareto’s law [34]. This is demonstrated in Figure 6(a), where the $C(f)$ cumulative income distributions as the function of population fraction f appears as a Lorentz-curve with area under the diagonal proportional to socioeconomic inequalities. As an example, Figure 2 depicts the spatial distribution of 2,000 users with inferred home locations in IRIS cells located in central Paris and colored as the median income.

3.3. Mobility Analysis. In order to further assess the validity of the assignment of location to socioeconomic status, we studied the mobility traces generated by the set of geolocated users. Specifically, mirroring previous work [12, 28], we focused on the predictability of individual trajectories by analyzing the visitation patterns of the top $n = 10$ locations. To do so, given a user having visited at least $k = 5$ different census blocks, we computed $\pi_k(i)$ the fraction of time a user spends in the top i most visited location as

$$\pi_k(i) = \frac{N_i}{\sum_{j=1}^n N_j}, \quad i = 1, \dots, n \quad (1)$$

FIGURE 2: IRIS area cells in central Paris colored according to the median income of inhabitants, with inferred home locations of 2,000 Twitter users.

with N_i the number of times the user appeared in the i -th location. Note that in the above definition, $\forall i \leq k$, $N_i \in \mathbb{N}^*$. This metric was shown by Song et al. [12] to be an upper bound to an individual's predictability. At the same time, taking geolocated users' median income inferred from census cells we associated them into one of nine socioeconomic classes (1-poorest, 9-richest) following the social stratum model introduced in [5]. This procedure sorts users by their income, takes the CDF of income (shown in Figure 6(a)), and divides users into groups such that each group represents the same total sum of income. This partitioning, due to the Lorentz-curve shape of the CDF, provides socioeconomic classes with decreasing size with increasing income, as expected from theory and other real observations [5, 35].

As previously pointed out [12, 28] we noticed that, for all socioeconomic classes, the first and second most visited locations concentrate between 60 and 74% of all geolocations, suggesting the high likelihood of an individual tweeting preferentially from his/her home or office [28]. Furthermore, by aggregating users per socioeconomic class, an interesting trend was exposed: the higher the socioeconomic class in question, the lower the (upper bound of) predictability of the considered users was for any of the top n locations. Indeed, when studying how $\pi_5(1)$ is related to a user's socioeconomic status, we see that the higher one's socioeconomic class is, the less frequently he is seen at the most visited location (see Figure 3(a)). The robustness of this trend was further assessed by repeating our analysis for different values of k , always recovering this decreasing trend (see Figure 3(b)). The

underlying explanation to this pattern may actually lie in the correlation between the socioeconomic status of people and the diversity of locations they visit. For instance, greater diversity may lead to a lower predictability of movement, which in turn might cause the observed trends. To control this, we computed the average radius of gyration r_g , describing the typical range of a user's trajectory per socioeconomic class, defined as

$$r_g = \sqrt{\frac{1}{L} \sum_{i=1}^L (\vec{r}_i - \vec{r}_{cm})^2}. \quad (2)$$

Here $\vec{r}_i$ represents the position at time i , $\vec{r}_{cm} = (1/L) \sum_{i=1}^L \vec{r}_i$ is the center of mass of the trajectory, and L is the total number of recorded points for the user's location. As we see in Figure 3(c), r_g seems to increase on average as higher socioeconomic status is considered. Hence, high SES users tend to have more diverse mobility patterns than low SES ones, which in turn leads to lower predictability of their whereabouts. These results may relate to previous work [5, 36], which explain this trend by means of the positive payoff between commuting farther for better jobs, while keeping better housing conditions. As a consequence, the inferred home location for high SES users might be less precise due to their more dispersed mobility patterns and the lower predictability of their whereabouts. This is one reason why we define our SES inference later as a two-way inference problem, dividing users into a "rich" and a "poor" class.

FIGURE 3: (a) The fraction of time a user spends in his/her top 10 most visited locations per socioeconomic class. (b) $\pi_k(1)$, fraction of time spent at the most visited location per socioeconomic class for users with at least k different locations. (c) Average radius of gyration per socioeconomic class (error bars are computed as $\sigma(r_g)/\sqrt{N}$, with N number of samples).

3.4. Occupation Data. Earlier studies [11, 17, 18] demonstrated that annotated occupation information can be effectively used to derive precise income for individuals and to infer therefore their SES. However, these methods required a somewhat selective set of Twitter users as well as an expensive annotation process by hiring premium annotators, e.g., from Amazon Mechanical Turk. Our goal here was to obtain the occupations for a general set of Twitter users without the involvement of annotators, but by collecting data from parallel online services.

As a second method to estimate SES, we took a sample of Twitter users who mentioned their LinkedIn [37] profile URL in their tweets or Twitter profile. Using these pointers we collected professional profile descriptions from LinkedIn by relying on an automatic crawler mainly used in Search Engine Optimization (SEO) tasks [38]. We obtained 4,140 Twitter/LinkedIn users all associated with their job title, professional skills, and profile description. Apart from the advantage of working with structured data, professional information extracted from LinkedIn is significantly more reliable than Twitter's due to the high degree of social scrutiny to which each profile is exposed [39].

To associate income to Twitter users with LinkedIn profiles, we matched them with a given salary based on their reported profession and an occupational-salary classification table provided by INSEE [40]. Due to the ambiguous naming of jobs and to acknowledge permanent/nonpermanent,

senior/junior contract types we followed three strategies for the matching. In 40% of the cases we directly associated the reported job titles to regular expressions of an occupation. In 50% of the cases we used string sequencing methods [41] to associate reported and official names of occupations with at least 90% match. For the remaining 10% of users we directly inspected profiles. The distribution of estimated salaries reflects the expected income heterogeneities as shown in Figure 6(a). Users were eventually assigned to one of two SES classes based on whether their salary was higher or lower than the average value of the income distribution. Also note that LinkedIn users may not be representative of the whole population. We discuss this and other types of potential biases in Section 6.

3.5. Expert Annotated Home Location Data. Finally, motivated by recent remote sensing techniques, we sought to estimate SES via the analysis of the urban environment around the inferred home locations. Similar methodology has been lately reported by the remote sensing community [42] to predict sociodemographic features of a given neighborhood by analyzing Google Street View images to detect different car models, or to predict poverty rates across urban areas in Africa from satellite imagery [43]. Driven by this line of work, we estimated the SES of geolocated Twitter users as follows.

FIGURE 4: Confusion matrix in the test set obtained with a ResNet50 trained on the UCMerced for land use inference.

3.5.1. Preselection of Home Locations. Using geolocated users identified in Section 3.1.1, we further filtered them to obtain a smaller set of users with more precise inferred home locations. We screened all of their geotagged tweets and looked for regular expressions determining whether or not a tweet was sent from home [31]. As explained in Section 3.1.1, we exploited that “home-suspected” expressions appeared with a particular temporal distribution (see Figure 1(c)) since these expressions were used during the night when users are at home. This selection yielded 28,397 users mentioning “home-suspected” expressions regularly at their inferred home locations.

3.5.2. Identification of Urban/Residential Areas. In order to filter out inferred home locations not in urban/residential areas, we downloaded via Google Maps Static API [44] a satellite view in a 100m radius around each coordinate (for a sample see Figure 5(a)). To discriminate between residential and nonresidential areas, we built on land use classifier [45] using aerial imagery from the UC Merced dataset [46]. This dataset contains 2100 256×256 1m/px aerial RGB images over 21 classes of different land uses (for a pair of sample images see Figure 5(b)). To classify land use, a CaffeNet architecture was trained, which reached accuracy over 95%. Here, we instantiated a ResNet50 network using keras [47] pretrained on ImageNet [48], where all layers except the last five were

frozen. The network was then trained with 10-fold cross-validation achieving a 93% accuracy after the first 100 epochs (cf. Figure 4). We used this model to classify images of the estimated home location satellite views (cf. Figure 5(a)) and kept those which were identified as residential areas (see Figure 5(b), showing the activation of the two first hidden layers of the trained model). This way, 5,396 inferred home locations were discarded.

3.5.3. Home Location Data with Expert Annotated SES. Next we aimed to estimate SES from architectural/urban features associated with the home locations. Given that our goal here was to lean on socioeconomic labels that were not estimated by the census, we forfeit the use of deep learning models to infer SES by relying upon human annotation. Thus, for each home location we collected two additional satellite views at different resolutions as well as six Street View images, each with a horizontal view of approximately 90°. We randomly selected a sample of 1,000 locations and involved architects to assign a SES score (from 1 to 9) to a sample set of selected locations based on the satellite and Street View around it (both samples had 333 overlapping locations). For validation, we took users from each annotated SES class and computed the distribution of their incomes inferred from the IRIS census data (see Section 3.2). Violin plots in Figure 5(d) show that in expert annotated data, as expected,

FIGURE 5: Top: ResNet50 Output: (a) original satellite view; (b) first two hidden layers activation; (c) final top-3 most frequent predicted area types; (d) architect SES score agreement with census median income for the sampled home locations. It is shown as violin plots of income distributions for users annotated in different classes (shown on x-axis and by color).

the inferred income values were positively correlated with the annotated SES classes. Labels were then categorized into two socioeconomic classes for comparison purposes. All in all, both annotators assigned the same label to the overlapping locations in 81.7% of samples.

To solve the SES inference problem we used the described three datasets (for a summary see Table 1). We defined the inference task as a two-way classification problem by dividing the user set of each dataset into two groups. For the census and occupation datasets the lower and higher SES classes were separated by the average income computed from the whole distribution, while in the case of the expert annotated data we assigned people from the lowest five SES labels to the lower SES class in the two-way task. The relative fractions of people assigned to the two classes are depicted in Figure 6(b) for each dataset and summarized in Table 1.

4. Feature Selection

Using the user profile information and tweets collected from every account’s timeline, we built a feature set for each user,

similar to Lampos et al. [11]. We categorized features into two sets, one containing shallow features directly observable from the data, while the other was obtained via a pipeline of data processing methods to capture semantic user features.

4.1. User Level Features. The user level features are based on the general user information or aggregated statistics about the tweets [17]. We therefore include general ordinal values such as the number and rate of retweets, mentions, and coarse-grained information about the social network of users (number of friends, followers, and ratio of friends to followers). Finally we vectorized each user’s profile description and tweets and selected the top 450 and 560 1-grams and 2-grams, respectively, observed through their accounts (where the rank of a given 1-gram was estimated via *tf-idf* [49]).

4.2. Linguistic Features. To represent textual information, in addition to word count data, we used topic models to encode coarse-grained information on the content of the tweets of

FIGURE 6: Cumulative distributions of income as a function of sorted fraction f of individuals. Dashed line corresponds to the perfectly balanced distribution. Distributions appear similar in spite of dealing with heterogeneous samples.

TABLE 1: Number of users and estimated fractions of low and high SES in each dataset.

	Census	Occupation	Expert
Size	32,053	4,140	1,000
Low SES	0.54	0.46	0.58
High SES	0.46	0.54	0.42

a user, similar to [11]. This enabled us to easily interpret the relation between semantic and socioeconomic features. Specifically, we started by training a *word2vec* model [50] on the whole set of tweets (obtained in the 2014-2015 time-frame) by using the skip-gram model and negative sampling with parameters similar to [15, 17]. To scale up the analysis, the number of dimensions for the embedding was kept at 50. This embedded words in the initial dataset in a $\mathbb{R}^{50}$ vector space.

Eventually we extracted conversation topics by running a spectral clustering algorithm on the word-to-word similarity matrix $M \in \mathbb{R}^{V \times V}$ with V vocabulary size and elements defined as the $M_{ij} = \langle u_i, u_j \rangle / \|u_i\| \|u_j\|$ cosine similarity between word vectors. Here $u_i \in \mathbb{R}^{50}$ is a vector of a word $i \in V$ in the embedding, $\langle \cdot \rangle$ is the dot product of vectors, and $\| \cdot \|$ is the L^2 norm of a vector. This definition allows for negative entries in the matrix to cluster, which were set to null in our case. This is consistent with the goal of the clustering procedure as negative similarities should not encode dissimilarity between pairs of words but orthogonality between the embeddings. This procedure was run for 50, 100, and 200 clusters and allowed the homogeneous distribution of words among clusters (hard clustering). The best results were obtained with 100 topics in the topic model. Finally, we manually labeled topics based on the words assigned to them and computed the topic-to-topic correlation matrix shown in

Figure 7. There, after block diagonalization, we found clearly correlated groups of topics which could be associated with larger topical areas such as communication, advertisement, or soccer.

As a result we could compute a representative topic distribution for each user, defined as a vector of normalized usage frequency of words from each topic. Also note that the topic distribution for a given user was automatically obtained as it depends only on the set of tweets and the learned topic clusters without further parametrization.

To demonstrate how discriminative the identified topics were in terms of the SES of users we associated with each user the 9th decile value of the income distribution corresponding to the census block of their home location and computed for each labeled topic the average income of users depending on whether or not they mentioned the given topic. Results in Figure 8 demonstrate that topics related to politics, technology, or culture are more discussed by people with higher income, while other topics associated with slang, insults, or informal abbreviations are more used by people of lower income. These observable differences between the average income of people, who use (or not) words from discriminative topics, demonstrate well the potential of word topic clustering used as features for the inference of SES. All in all, each user in our dataset was assigned with a 1117 feature vector encoding the lexical and semantic profile she displayed

FIGURE 7: Clustered topic-to-topic correlation matrix: topics are generated via the spectral clustering of the word2vec word cosimilarity matrix. Row labels are the name of topics while column labels are their categories. Blue cells (resp. red) assign negative (resp. positive) Pearson's correlation coefficients.

on Twitter. We did not apply any further feature selection as the distribution of importance of features appeared rather smooth (not shown here). It did not provide evident ways to identify a clear set of particularly determinant features, but rather indicated that the combination of them was important.

5. Results

In order to assess the degree to which linguistic features can be used for discriminating users by their socioeconomic class, we trained with these feature sets different learning algorithms. Namely, we used the XGBoost algorithm [51], an implementation of the gradient-boosted decision trees for this task. Training a decision tree learning algorithm involves the generation of a series of rules, split points or nodes

ordered in a tree-like structure enabling the prediction of a target output value based on the values of the input features. More specifically, XGBoost, as an ensemble technique, is trained by sequentially adding a high number of individually weak but complementary classifiers to produce a robust estimator: each new model is built to be maximally correlated with the negative gradient of the loss function associated with the model assembly [52]. To evaluate the performance of this method we benchmarked it against more standard ensemble learning algorithms such as AdaBoost, Logistic Regression, SVM, and Random Forest.

For each socioeconomic dataset, we trained our models by using 75% of the available data for training and the remaining 25% for testing. During the training phase, the training data undergoes a k -fold inner cross-validation, with $k = 5$, where all splits are computed in a stratified manner

FIGURE 8: Average income for users who tweeted about a given topic (blue) vs. those who did not (red). Label of the considered topic is on the left.

TABLE 2: Classification performance (5-CV): AUC scores (mean $\pm$ STD) of five different classifiers on each dataset.

	Census	Occupation	Expert
AdaBoost	0.549 $\pm$ 0.009	0.628 $\pm$ 0.022	0.575 $\pm$ 0.013
Logistic Reg.	0.658 $\pm$ 0.013	0.778 $\pm$ 0.058	0.571 $\pm$ 0.033
SVM	0.657 $\pm$ 0.016	0.788 $\pm$ 0.041	0.600 $\pm$ 0.012
Random Forest	0.677 $\pm$ 0.011	0.783 $\pm$ 0.017	0.593 $\pm$ 0.049
XGBoost	0.700 $\pm$ 0.011	0.798 $\pm$ 0.015	0.605 $\pm$ 0.029

to get the same ratio of lower to higher SES users. The four first blocks were used for inner training and the remainder for inner testing. This was repeated ten times for each model so that in the end each model's performance on the validation set was averaged over 50 samples. For each model, the parameters were fine-tuned by training 500 different models over the aforementioned splits. The selected one was that which gave the best performance on average, which was then applied to the held-out test set. This is then repeated through a 5-fold outer cross-validation.

In terms of prediction score, we followed a standard procedure in the literature [53] and evaluated the learned models by considering the area under the receiver operating characteristic curve (AUC). This metric can be thought as the probability that a classifier ranks a randomly chosen positive instance higher than a randomly chosen negative one [52]. This procedure was applied to each of our datasets. The obtained results are shown in Figure 9 and in Table 3.

As a result, we first observed that XGBoost consistently provided top prediction scores when compared to AdaBoost and Random Forest (all performance scores are summarized in Table 2). We hence used it for our predictions in the remainder of this study. We found that the LinkedIn data was the best, with $AUC = 0.80$, to train a model to predict SES of people based on their semantic features. It provided a 10% increase in performance as compared to the census based inference with $AUC = 0.70$, and 19% relative to expert annotated data with $AUC = 0.61$. Thus we can conclude that there seems to be a trade-off between

FIGURE 9: ROC curves for 2-way SES prediction using tuned XGBoost in each of the 3 SES datasets. AUC values are reported in the legend. The dashed line corresponds to the line of no discrimination. Solid lines assign average values over all folds while shaded regions represent standard deviation.

scalability and prediction quality, as while the occupation dataset provided the best results, it seems unlikely to be subject to any upscaling due to the high cost of obtaining a

TABLE 3: Detailed average performance (5-CV) on test data for the binary SES inference problem for each of the 3 datasets.

Dataset	SES Class	Performance on test set		
		Precision	Recall	F1-score
Census	Low	0.652	0.596	0.624
	High	0.628	0.682	0.652
LinkedIn	Low	0.700	0.733	0.717
	High	0.735	0.702	0.720
Architect	Low	0.622	0.598	0.607
	High	0.550	0.573	0.556

FIGURE 10: Top (red) and bottom (blue) five topics ranked in terms of their predictive performance in the XGBoost model trained on LinkedIn. Error bars indicate s.d. across the 10 cross-validation samples.

clean dataset. Relying on location to estimate SES seems to be more likely to benefit from such an approach, though at the cost of an increased number of mislabeled users in the dataset. Moreover, the annotator’s estimation of SES using Street View at each home location seems to be hindered by the large variability of urban features. Note that even though interagreement is 76%, the Cohen’s kappa score for annotator interagreement is low at 0.169. Furthermore, we remark that the expert annotated pipeline was also subject to noise affecting the home location estimations, which potentially contributed to the lowest predictive performance.

We also report the top and bottom five topics ranked by their importance when the XGBoost model was trained on the best performing proxy, i.e., on occupation (see Figure 10). Perhaps unsurprisingly, topics related to professional occupations are the ones recognized by the model as most important. Nevertheless, syntax remains an important feature too. Furthermore, topics associated with particular communities (German/Turkish) or general interest (Soccer) seem to be less useful in terms of SES discrimination. This could in turn be explained by the sparsity of individuals using them or inversely, by the breadth of users discussing them.

Finally, it should also be noted that following recent work by Aletras and Chamberlain in [26], we tested our model by extending the feature set with the *node2vec* embedding of users computed from the mutual mention graph of Twitter.

Nevertheless, in our setting, it did not significantly increase the overall predictive performance of the inference pipeline. We hence did not include it in the feature set for the sake of simplicity.

6. Limitations

In this work we combined multiple datasets collected from various sources. Each of them came with some bias due to the data collection and posttreatment methods or the incomplete set of users. These biases may limit the success of our inference; thus, their identification is important for the interpretation and future developments of our framework.

(i) *Location Data.* Although we designed very strict conditions for the precise inference of home locations of geolocated users, this process may have some uncertainty due to outlier behavior. Further bias may be induced by the relatively long time passed between the posting of the location data and of the tweets collection of users.

(ii) *Census Data.* As we already mentioned the census data does not cover the entire French territory as it reports only cells with close to 2,000 inhabitants. This may introduce biases in two ways: by limiting the number of people in our sample living in rural areas and by associating income

with large variation to each cell. While the former limit had marginal effects on our predictions, as Twitter users mostly live in urban areas, we addressed the latter effect by associating the median income to users located in a given cell.

(iii) *Occupation Data.* LinkedIn as a professional online social network is predominantly used by people from IT, business, management, marketing, or other expert areas, typically associated with higher education levels and higher salaries. Moreover, we could observe only users who shared their professional profiles on Twitter, which may further biased our training set. In terms of occupational-salary classification, the data in [40] was collected in 2010 thus may not contain more recent professions. These biases may induce limits in the representativeness of our training data and thus in the predictions' precision. However, results based on this method of SES annotation performed best in our measurements, indicating that professions are among the most predictive features of SES, as has been reported in [11].

(iv) *Annotated Home Locations.* The remote sensing annotation was done by experts and their evaluation was based on visual inspection and biased by some unavoidable subjectivity. Although their annotations were cross-referenced and found to be consistent, they still contained biases, like over-representative middle classes, which somewhat undermined the prediction task based on this dataset.

(v) *Different Sets of Users.* Our methodologies rely on nonentirely overlapping user sets when they turn to SES inference using occupational data, census data, or remotely sensed values as proxy for individual socioeconomic status. The obtained results are undoubtedly linked to the set of individuals used in each dataset, which may affect the discriminative analysis on the advantages that each proxy provides for the inference task. On the other hand, due to the same collection filters and preprocessing conditions, users in these subsets may be considered similar enough to be able to compare the performance provided by the different methods.

Despite these shortcomings, using all three datasets, we were able to infer SES with performances close to earlier reported results, which were based on more thoroughly annotated datasets. Our results and our approach of using open, crawlable, or remotely sensed data highlight the potential of the proposed methodologies.

7. Conclusions

In this work we proposed a novel methodology for the inference of the SES of Twitter users. We built our models combining information obtained from numerous sources, including Twitter, census data, LinkedIn, and Google Maps. We developed precise methods of home location inference from geolocation, novel annotation of remotely sensed images of living environments, and effective combination of datasets collected from multiple sources. In terms of novelty, we demonstrated that, within the French Twitter space, the utilization of words in different topic categories, identified via advanced semantic analysis of tweets, can discriminate

between people of different income and that the mobility patterns and such predictability of users' whereabouts are strongly dependent on SES of people. Furthermore, we showed that among the candidate socioeconomic proxies chosen, the best results were obtained using occupational data. More importantly, we presented a proof-of-concept that our methods are competitive in terms of SES inference when compared to other methods relying on domain specific information. We can identify several future directions and applications of our work. First, further development of data annotation of remotely sensed information is a promising direction. Note that after training, our model requires as input only information that can be collected exclusively from Twitter, without relying on other data sources. This holds a large potential in terms of SES inference of larger sets of Twitter users, which in turn opens the door for studies to address population level correlations of SES with language, space, time, or social network. As such, our methodology has the merit not only of addressing open scientific questions, but also of contributing to the development of new applications in recommendation systems, in predicting customer behavior, or in online social services.

Data Availability

In order to uphold the strict privacy laws in France as well as the agreement signed with our data provider GNIP, full disclosure of the original dataset is not possible. The GitHub repository containing the data collection and preprocessing pipelines is available at <https://github.com/jaklevab/TWITTERSES>.

Conflicts of Interest

The authors declare no conflicts of interest.

Acknowledgments

We thank J-Ph. Magué, J-P. Chevrot, D. Seddah, D. Carnino, and E. De La Clergerie for constructive discussions and for their advice on data management and analysis. We are grateful to J. Altnéder and M. Hunyadi for their contributions as expert architects for data annotation. The manuscript was presented at the 2018 IEEE 18th International Conference on Data Mining, IWSC'18 2nd International Workshop on Social Computing (Singapore, 17th November, 2018). This work was supported by the SoSweet ANR project (ANR-15-CE38-0011), the MOTIf Stic-AmSud project (18-STIC-07), and the ACADEMICS project financed by IDEX LYON.

References

- [1] V. Mayer-Schönberger and K. Cukier, *Big Data: A Revolution That Transforms How We Work, Live, and Think*, John Murray, 2012.
- [2] D. Lazer, A. Pentland, L. Adamic et al., "Life in the network: the coming age of computational social science," *Science*, vol. 323, no. 5915, p. 721, 2009.

- [3] K. D. V. Liere and R. E. Dunlap, "The social bases of environmental concern: A review of hypotheses, explanations and empirical evidence," *Public Opinion Quarterly*, vol. 44, no. 2, pp. 181–197, 1980.
- [4] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: homophily in social networks," *Annual Review of Sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [5] Y. Leo, E. Fleury, J. I. Alvarez-Hamelin, C. Sarraute, and M. Karsai, "Socioeconomic correlations and stratification in social-communication networks," *Journal of the Royal Society Interface*, vol. 13, no. 125, Article ID 20160598, 2016.
- [6] J. L. Brown-Iannuzzi, K. B. Lundberg, and S. McKee, "The politics of socioeconomic status: how socioeconomic status may influence political attitudes and engagement," *Current Opinion in Psychology*, vol. 18, pp. 11–14, 2017.
- [7] J. L. Abitbol, M. Karsai, J. Magué, J. Chevrot, and E. Fleury, "Socioeconomic dependencies of linguistic patterns in twitter: a multivariate analysis," in *Proceedings of the World Wide Web Conference (TheWebConf '18)*, pp. 1125–1134, Lyon, France, April 2018.
- [8] M. Kosinski, D. Stillwell, and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 110, no. 15, pp. 5802–5805, 2013.
- [9] Y. Leo, M. Karsai, C. Sarraute, and E. Fleury, "Correlations and dynamics of consumption patterns in social-economic networks," *Social Network Analysis and Mining*, vol. 8, no. 1, p. 9, 2018.
- [10] T. Piketty, "Capital in the 21st century," 2014.
- [11] D. Preoțiuc-Pietro, S. Volkova, V. Lampos, Y. Bachrach, N. Aletras, and L. A. Braunstein, "Studying user income through language, behaviour and affect in social media," *PLoS ONE*, vol. 10, no. 9, Article ID e0138717, pp. 1–17, 2015.
- [12] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010.
- [13] J. Levy Abitbol, M. Karsai, and E. Fleury, "Location, occupation, and semantics based socioeconomic status inference on twitter," in *Proceedings of the 18th International Conference on Data Mining (IWSM '18) and 2nd International Workshop on Social Computing (ICDMW '18)*, pp. 1192–1199, November 2018.
- [14] J. L. Abitbol, <https://github.com/jaklevab/TWITTERSES>, 2019.
- [15] B. P. Chamberlain, C. Humby, and M. Deisenroth, "Detecting the age of twitter users, 2016," <https://arxiv.org/abs/1601.04621>.
- [16] T. Hu, H. Xiao, J. Luo, and T. T. Nguyen, "What the language you tweet says about your occupation," *Tenth International AAAI Conference on Web and Social Media*, 2017.
- [17] V. Lampos, N. Aletras, J. K. Geyti, B. Zou, and I. J. Cox, "Inferring the socioeconomic status of social media users based on behaviour and language," in *Advances in Information Retrieval, Lecture Notes in Computer Science*, pp. 689–695, Springer International Publishing, 2016.
- [18] D. Preoțiuc-Pietro, V. Lampos, and N. Aletras, "An analysis of the user occupational class through Twitter content," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*, pp. 1754–1764, Beijing, China, July 2015.
- [19] S. Volkova, G. Coppersmith, and B. Van Durme, "Inferring user political preferences from streaming communications," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL '14)*, pp. 186–196, June 2014.
- [20] H. A. Schwartz, J. C. Eichstaedt, M. L. Kern et al., "Personality, gender, and age in the language of social media: the open-vocabulary approach," *PLoS ONE*, vol. 8, no. 9, Article ID e73791, pp. 1–16, 2013.
- [21] S. Luo, F. Morone, C. Sarraute, M. Travizano, and H. A. Makse, "Inferring personal economic status from social network location," *Nature Communications*, vol. 8, Article ID 15227, 2017.
- [22] Twitter Open API, 2018, <https://developer.twitter.com/en/docs.html>.
- [23] A. Culotta, N. K. Ravi, and J. Cutler, "Predicting the demographics of Twitter users from website traffic data," in *Proceedings of the AAAI Conference on Artificial Intelligence*, January 2015.
- [24] J. Eisenstein, B. O'Connor, N. A. Smith, E. P. Xing, and R. C. Berwick, "Diffusion of lexical change in social media," *PLoS ONE*, vol. 9, no. 11, Article ID e113114, pp. 1–13, 2014.
- [25] A. Llorente, M. Garcia-Herranz, M. Cebrian, E. Moro, and Y. Moreno, "Social media fingerprints of unemployment," *PLoS ONE*, vol. 10, no. 5, pp. 1–13, 2015.
- [26] N. Aletras and B. P. Chamberlain, "Predicting twitter user socioeconomic attributes with network and language information," *Proceedings of the 29th on Hypertext and Social Media*, 2018.
- [27] M. C. González, C. A. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns," *Nature*, vol. 453, pp. 779–782, 2008.
- [28] R. Jurdak, K. Zhao, J. Liu, M. AbouJaoude, M. Cameron, and D. Newth, "Understanding human mobility from Twitter," *PLoS ONE*, vol. 10, no. 7, 2015.
- [29] B. Bernstein, "Language and social class," *The British Journal of Sociology*, vol. 11, no. 3, pp. 271–276, 1960.
- [30] R. Compton, D. Jurgens, and D. Allen, "Geotagging one hundred million Twitter accounts with total variation minimization," *IEEE International Conference on Big Data*, 2014.
- [31] T. Hu, J. Luo, H. Kautz, and A. Sadilek, "Home location inference from sparse and noisy data: models and applications," *Frontiers of Information Technology & Electronic Engineering*, vol. 17, no. 5, pp. 389–402, 2016.
- [32] Gini Index World Bank, 2010, <https://data.worldbank.org/indicator/SI.POV.GINI?locations=FR>.
- [33] INSEE, Revenus, pauvreté et niveau de vie en 2014, 2017, <https://www.insee.fr/fr/statistiques/3288151/>.
- [34] V. Pareto, "Manual of political economy," 1971.
- [35] P. Saunders, "Social class and stratification," *Routledge*, 2006.
- [36] Y. Xu, A. Belyi, I. Bojic, and C. Ratti, "Human mobility and socioeconomic status: Analysis of Singapore and Boston," *Computers, Environment and Urban Systems*, vol. 72, pp. 51–67, 2018.
- [37] LinkedIn, 2018.
- [38] LinkedInHelper, 2016, <https://linkedhelper.com/>.
- [39] P. Manzanares-Lopez, J. P. Muñoz-Gea, and J. Malgosa-Sanahuja, "Analysis of linkedin privacy settings: are they sufficient, insufficient or just unknown?" in *Proceedings of the 10th International Conference on Web Information Systems and Technologies (WEBIST '14)*, vol. 1, pp. 285–293, April 2014.
- [40] INSEE, "Les salaires dans le secteur privé et les entreprises publiques," 2010, <https://www.insee.fr/fr/statistiques/2122237/>.
- [41] Sequence Matcher Python Library, 2017.
- [42] T. Gebru, J. Krause, Y. Wang et al., "Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 114, no. 50, pp. 13108–13113, 2017.

- [43] N. Jean, M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, vol. 353, no. 6301, pp. 790–794, 2016.
- [44] Google Maps Static API, 2018, <https://developers.google.com/maps/>.
- [45] M. Castelluccio, G. Poggi, C. Sansone, and L. Verdoliva, Land use classification in remote sensing images by convolutional neural networks, 2015, <https://arxiv.org/abs/1508.00092>.
- [46] UC Merced Land Use Dataset, 2017, <http://weegee.vision.ucmerced.edu/datasets/landuse.html>.
- [47] F. Chollet et al. Keras. <https://keras.io>, 2015, date of access: November 2018.
- [48] J. Deng, W. Dong, and R. Socher, "ImageNet: a large-scale hierarchical image database," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '09)*, pp. 248–255, Miami, FL, USA, June 2009.
- [49] J. Leskovec, A. Rajaraman, and J. Ullman, *Mining of Massive Datasets*, Cambridge University Press, 2014.
- [50] T. Mikolov, K. Chen, G. Corrado, and J. Dean, Efficient estimation of word representations in vector space, 2013 <https://arxiv.org/abs/1301.3781>.
- [51] T. Chen and C. Guestrin, "XGBoost: a scalable tree boosting system," *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [52] L. Torlay, M. Perrone-Bertolotti, E. Thomas, and M. Baciau, "Machine learning–XGBoost analysis of language networks to classify patients with epilepsy," *Brain Informatics*, vol. 4, no. 3, pp. 159–169, 2017.
- [53] J. Foster, T. Provost, and R. Kohavi, "The case against accuracy estimation for comparing induction algorithms," in *Proceedings of the 15th International Conference on Machine Learning (ICML '98)*, pp. 445–453, San Francisco, Calif, USA, 1998.