Estimating Disaggregated Employment Size from Points-of-Interest and Census Data: From Mining the Web to Model Implementation and Visualization

Filipe Rodrigues, Ana Alves, Evgheni Polisciuc
Department of Informatics Engineering
University of Coimbra
Portugal
Email: {fmpr,ana}@dei.uc.pt,
evgheni@student.dei.uc.pt

Shan Jiang, Joseph Ferreira
Massachusetts Institute of Technology
Boston, USA
Email: {shanjiang,jf}@mit.edu

Francisco C. Pereira
Singapore-MIT Alliance
for Research and Technology
Singapore
Email: camara@smart.mit.edu

***Abstract*—The global spread of internet access and the ubiquity of internet capable devices has lead to an increased online presence on the behalf of companies and businesses, namely in collaborative platforms called local directories, where Points-of-Interest (POIs) are usually classified with a set of categories and tags. Such information can be extremely useful, especially if aggregated under a common (shared) taxonomy. This article proposes a complete framework for the urban planning task of disaggregated employment size estimation based on collaborative online POI data, collected using web mining techniques. In order to make the analysis possible, we present a machine learning approach to automatically classify POIs to a common taxonomy - the North American Industry Classification System. This hierarchical taxonomy is applied in many areas, particularly in urban planning, since it allows for a proper analysis of the data at different levels of detail, depending on the practical application at hand. The classified POIs are then used to estimate disaggregated employment size, at a finer level than previously possible, using a maximum likelihood estimator. We empirically show that the automatically-classified online POIs are competitive with proprietary gold-standard POI data. This fact is then supported through a set of new visualizations that allow us to understand the spatial distribution of the classification error and its relation with employment size error.**

***Keywords*—machine learning, spatial analysis, points-of-interest, urban planning, GIS.**

I. Introduction

With the increasing number of mobile devices and social networks in the latest years, the amount of geo-referenced information available on the Web is growing at an astonishing rate. Capture devices such as camera-phones and GPS-enabled cameras can automatically associate geographic data with images, which is significantly increasing the number of geo-referenced data available online. Social networks also have an important role. They are a great medium where users can share information they collect with their mobile devices. As a consequence, the amount of online descriptive information about places has reached reasonable dimensions for many cities in the world.

A point of interest, or POI for short, is a specific point location that someone may find useful or interesting. POIs can be used in navigation, characterization of a place, sociological studies, city dynamics analysis, geo-reference of texts, etc [1]. Such a simple information structure can be used and enriched such that context-aware systems behave more intelligently.

In spite of their importance, the production of POIs is scattered across a myriad of different websites, systems and devices, thus making it extremely difficult to obtain an exhaustive database of such wealthy information. There are hundreds, if not thousands, of POI directories in the Web like Yahoo![1], Manta[2] and Yellow Pages[3], each one using its own taxonomy of categories or tags. Therefore, it is essential to unify these different sources by mapping them to a common taxonomy. Otherwise, their application as a whole becomes impractical.

In this article, we propose the use of machine learning techniques to automatically classify POIs from different sources to a standard taxonomy such as NAICS [2] (U.S., Canada and Mexico) or ISIC[4] (United Nations), thereby allowing a proper analysis and visualization of the POI data, especially when the latter comes from different sources. A good example is land use analysis, a central pillar in urban planning. If the POIs do not share a common taxonomy then we are not able to determine, for instance, how many POIs of universities exist in a given area, since a POI source can classify them as "schools" while others classify them as "higher education." This makes the whole analysis unreliable. In the particular case that we explore in this article, we are interested in classifying POIs according to the North American Industry Classification System (NAICS). The NAICS is the standard used by Federal statistical agencies in classifying business establishments for the purpose of collecting, analyzing, and publishing statistical data related to the U.S. business economy [2]. NAICS was developed under the auspices of the Office of Management and Budget (OMB), and was adopted in 1997 to replace the old Standard Industrial Classification (SIC) system.

NAICS is a two to six-digit hierarchical classification code system, offering five levels of detail. Each digit in the code is part of a series of progressively narrower categories, and more digits in the code signify greater classification detail.

[1] http://local.yahoo.com
[2] http://www.manta.com
[3] http://www.yellowpages.com
[4] http://unstats.un.org/unsd/cr/registry/isic-4.asp

The first two digits designate the economic sector, the third digit designates the sub-sector, the fourth digit designates the industry group, the fifth digit designates the NAICS industry, and the sixth digit designates the national industry. A complete and valid NAICS code contains six digits [3]. By having different levels of detail, NAICS codes allow us to perform analysis at different granularities depending on the practical application at hand.

Figure 1 shows part of the NAICS hierarchy for the retail economic sector.

- **44-45 - Retail Trade**
  - **441 - Motor Vehicle and Parts Dealers**
    - **4411 - Automobile Dealers**
      - **44111 - New Car Dealers**
        - **441110 - New Car Dealers**
  - **(...)**
  - **451 - Sporting Goods, Hobby, Musical Instrument, and Book Stores**
    - **4511 - Sporting Goods, Hobby, Musical Instrument Stores**
      - **45111 - Sporting Goods Stores**
        - **451110 - Sporting Goods Stores**
      - **45112 – Hobby, Toy, and Game Stores**

Fig. 1. Example of the NAICS hierarchy for the retail economic sector.

In order to make learning possible, a POI matching technique is proposed, allowing the establishment of golden dataset from which various typical machine learning models are then estimated. After comparing several classification methods, we apply the results to the urban modeling task of estimating employment size at a disaggregated level. This task is traditionally made at a coarser level (Traffic Analysis Zone, Census Tract or Block Group level) than what could be now possible, through the use of POI data.

Finally, we explore innovative visualization techniques to (1) understand the POI distribution across space, (2) identify spatial areas of POI classification error, and (3) relate the latter with the accuracy of employment size estimation model.

In summary, the main contributions of this article are:

- A POI matching algorithm;
- A machine learning approach to automatically classify POIs to standard classification system (NAICS);
- A model to estimate employment at a disaggregated level;
- A collection of visualization techniques that allow a deeper understanding of the geographical data and the models developed.

The remainder of this article is organized as follows. Section II presents previous related studies. Section III explains our data analysis and modeling methodology, from data preparation to model generation and validation. Section IV shows the obtained results. In Section V we describe the application of this methodology to the field of urban planning. Section VI presents a collection of visualization techniques of the data used and the models produced. We finish the article with conclusions and future work.

II. State of the Art

The applications of machine learning algorithms in classification tasks are vast and cover diverse areas that range from Speech Recognition to Medicine, including forecasting in Economics and Environmental Engineering or Road Traffic Prediction. In urban planning, land-use/land-cover information has long been recognized as a very important material [4]. However, as Fresco [5] claimed, accurate data on actual land-use cannot be easily found at both global/continental and national/regional scales.

In order to cope with these problems, automatic approaches to classify land use are being developed using distinct techniques.

A common approach to infer land-use/land-cover is to use satellite imagery. However, while these approaches have already proven to get good results, they are more suited to land-cover inference which is considered somehow different from land-use by many authors. Campbell [6], for example, considers land-cover to be concrete whereas land-use is abstract. That is, land-cover can be mapped directly from images, while land-use requires land-cover and additional information on how the land is used. Danoedoro [7] tries to improve land-use classification via satellite imagery by combining spectral classification, image segmentation and visual interpretation. Although he showed that satellite imagery could be used for generating socio-economic function of land-use at 83.63% accuracy, he is the first to recognize that applying such techniques to highly populated areas would be problematic.

Li et al. [8] use data mining techniques to discover knowledge from GIS databases and remote sensing image data that could be used for land use classification. In the field of remote sensing, Bayes classification (or maximum likelihood classification) is most widely used and, for most multi-spectral remote sensing data, the Bayes method classifies the coarse classes correctly, such as water, residential area, green patches, etc. But usually more detailed classification is required in land use classification. In order to subdivide some of the classes, Li et al. proposes the use of inductive learning techniques, particularly the C5.0 algorithm. By using these techniques they were able to get an overall accuracy of 89%. Comparing their final result with the result produced only by Bayes classification, the overall accuracy increased 11%.

An alternative to satellite imagery is the POI data. Using a large commercial POI database, Santos and Moreira [9] create and classify location contexts using decision trees. They identify clusters by means of a density-based clustering algorithm (Shared Nearest Neighbor algorithm) which allow them to define areas (or regions) through the application of a concave hull algorithm they developed to the POIs within each cluster. Finally, making use of the C5.0 algorithm, they classify a given location according to such characteristics as the number of POIs in a cluster, the size of the area of the cluster and the categories of the POIs within the cluster.

In order to use POI data for the classification of places and land-use analysis, POI classification is an essential task. Griffin et al. [10] use decision trees to classify GPS-derived POIs. However, they refer to POIs as "personal" locations to a given individual (i.e., home, work, restaurant, etc.). The main goal of their approach is then to automatically classify trips. In their approach, they start by determining clusters of trip-stops (i.e., stops that took more than 5 minutes) using a density-based clustering algorithm (Dbscan). Then, they make use of the C4.5 algorithm to classify the generated clusters as being "home", "work", "restaurant", etc., based on the time of the day and the length of the stay. However, to our best knowledge, no previous approaches have been made to classify POIs to a classification system such as NAICS. The latter is widely used for industry classification and has already been used, for instance, to classify Web Sites through machine learning techniques [11].

Spatial analysis has long been a topic of interest for researchers, who seek a comprehensive understanding on how the city behaves in different perspectives and its impact in the economy. Methods for analyzing spatial (and space-time) data have already been well developed by statisticians [12] and econometricians [13].

Visualization of geo-referenced data is one effort in such understanding. Instead of simply presenting information on a map, visualization facilitates the recognition of patterns in data by preprocessing and applying statistical filtering (e.g., average, deviation, clustering) over large datasets. Keim et al. [14] were pioneers in using visual approaches to explore heterogeneous and noisy large amount of spatial data. The authors showed how visualizations offered a qualitative overview of the data and allowed unexpectedly detected phenomena to be pointed out and explored using further quantitative analysis. Later, Costa and Venturini [15] improved these methodologies in order to give the possibility to interact with such artifacts. The authors applied them to a large POI database and showed with linear computation time it would be possible to present and interact with up to one million spatial points.

Currid et al. [16] try to understand the importance of agglomeration economies as a backbone to urban and regional growth, by identifying clusters of several "advanced" service sectors (professional, management, media, finance, art and culture, engineering and high technology) and comparing them in the top ten populous metropolitan areas in the U.S. They concluded that there are three spatial typologies of growth in the advanced services within U.S. urban regions. These typologies allowed them to understand qualities of place in general and of places specifically that drive the agglomeration of advanced services.

On a particular case study of the biotech industry in the U.S., Sambidi and Harrison[17] also analyze factors affecting site-selection of industries, testing the hypothesis of spatial agglomeration economies in that industry and confirm it using spatial econometrics. In the same topic, Arbia[18] classifies the spatial processes of individual firms into a birth process (new firms) and a growth process (existing firms) and proposes a model of economic activities on a continuous space also with the purpose of studying the geographical concentration of economic activities and analyze the economic behavior of individual firms.

III. Approach

In this section, we describe our approach, particularly what are the sources of our POI data, how we generate the training data, what methods we use for classification and how we perform validation.

A. POI Sources

Our data consists of a large set of POIs extracted from Yahoo! through their public API, another set acquired to Dun & Bradstreet (D&B) [19], a consultancy company that specializes in commercial information and insight for businesses, and a third one from InfoUSA[5] provided by the Harvard Center for Geographic Analysis (ESRI Business Analyst Data). In the first data set (from Yahoo!), the database is essentially built from user contributions. In the other two, the data acquisition process is semi-automatic and involves integration of official and corporate databases, statistical analysis and manual evaluation [19]. The POIs from D&B and InfoUSA have a NAICS code assigned (2007 version), which is not present in Yahoo!. However, each POI from Yahoo! is assigned, in average, roughly two arbitrary categories from the Yahoo! categories set. These categories are specified by the user, when adding a new POI, through a textfield and can be rather disparate since Yahoo! forces no restrictions over them. Considering that every POI source provides either some categories or tags associated with their POIs, we take advantage of this information to classify them to NAICS, where a single unifying code is assigned to each POI.

We have 156364 POIs from Yahoo!, 29402 from D&B and 196612 from InfoUSA for the area of Boston, Massachusetts. We also used 331118 POIs from Yahoo! and 16852 from D&B for the New York city area to see how our previously trained model would perform in a different city. We estimate that the Yahoo!'s categories taxonomy has more than 1300 distinct categories distributed along a 3-level hierarchy. On the other hand, NAICS has a total of 2332 distinct codes distributed along their 6-level hierarchy (1175 only in the sixth level).

Given its nature, the growth of the Yahoo! database (or any other user content platform) is considerably faster than D&B and InfoUSA, and the POI categorization follows less strict guidelines, which in some cases, as mentioned before, may become subjective. Our hypothesis is that there is considerable coherence between Yahoo! categories and NAICS codes, such that a model can be learned that automatically classifies incoming Yahoo! POIs.

In order to generate training data for the machine learning algorithms we use a *POI Matching* algorithm.

B. POI Matching and Data Preparation

When we are comparing POIs from different sources, it is important to have a way to identify similar POIs in order to correlate both databases. This requires a way to identify similarities based, not only in proximity, but also in name likeness. Our matching algorithm compares POIs according to their name, Web Site and distance. It makes use of the JaroWinklerTFIDF class from the SecondString project [20] to identify close names, ignoring misspelling errors and some abbreviations. Taking this into account, two POIs will be considered similar by our algorithm if they fit into one of the following groups:

[5] www.infousa.com

TABLE I. SOME STATISTICS OF DATASETS A AND B

	Dataset A	Dataset B
NAICS source	D&B	InfoUSA
Total POIs	7289	44634
Distinct NAICS	504	689
Distinct Yahoo! categories	802	1109
Distinct Yahoo! category combinations	569	1002
Category combinations that appear only once	136	92
Categories that appear only once	181	107
NAICS that appear only once	115	96

- The distance between the two POIs is less than 80 meters, the name similarity is above 0.70 and one or both POIs do not have website information.
- The distance between the two POIs is less than 80 meters, the name similarity is above 0.70 and the website similarity is higher than 0.60.
- The distance between the two POIs is less than 80 meters, the name similarity is above 0.60 and the website similarity is higher than 0.95.

We set the similarity thresholds to high values in order to get only high confidence matches. By manually validating a random subset of the POI matches identified (6 sets of 50 random POIs assigned to 6 volunteers), we concluded that the percentage of correct similarities identified was above 98% ($\sigma = 1.79$). Differently to validations later mentioned in this article, this is an extremely objective one, not demanding external participants or a very large sample[6].

After matching Yahoo! POIs to D&B and InfoUSA, we built two different geographic databases, where each POI contains a set of categories from Yahoo! and a NAICS classification provided by D&B and InfoUSA respectively. From this point on, we shall refer to the initial dataset, which results from POI matches between Yahoo! and D&B, as dataset A, and to the dataset resultant from the POI matching between Yahoo! and InfoUSA as dataset B. The later is six times larger than the former, due to larger coverage of InfoUSA in Boston.

Table I shows some statistic details of both datasets used.

The dataset A contains 7289 POIs for Boston and Cambridge and 2415 for New York. In comparison with the original databases, these are much smaller sets due to a very conservative POI matching approach (string similarity of at least 80%, max distance of 80 meters). However the POI quantities are high enough to build statistically valid models. We performed a detailed analysis of this data and identified 569 different category combinations which included only 802 distinct categories from the full set (of over 1300). From D&B, our data covers 504 distinct six-digit NAICS codes. However, the 2007 NAICS taxonomy has a total of 1175 six-level categories, meaning that our sample data only covers some of the most common NAICS codes, which only represents about 43% of the total number of NAICS categories. Nevertheless, the remaining ones are more exotic in our context and hence less significant for posterior analyses.

[6]Using the central limit theorem, the standard error of the mean should be near 0.73. Assuming an underestimation bias for n=6 of 5% (by the [21]), accuracy keeps very high, yielding a 95% confidence interval of [96.5%, 98.7%]

TABLE II. MOST COMMON NAICS IN THE DATASET A

NAICS code	Description	Occurrences
423730	Warm Air Heating and Air-Conditioning Equipment and Supplies Merchant Wholesalers	707
446130	Optical Goods Stores	200
314999	All Other Miscellaneous Textile Product Mills	193
493120	Refrigerated Warehousing and Storage	136
332997	Industrial Pattern Manufacturing	123

TABLE III. MOST COMMON YAHOO! CATEGORIES IN THE DATASET A

Yahoo! category	Occurrences
Salons	157
All Law Firms	129
Government	116
Trade Organizations	115
Architecture	86

Figure 2 shows the distribution of POIs along the different NAICS codes for dataset A. As we can see in the chart, the distribution is far from being uniform, which further complicates the classification task for NAICS codes with few training examples.

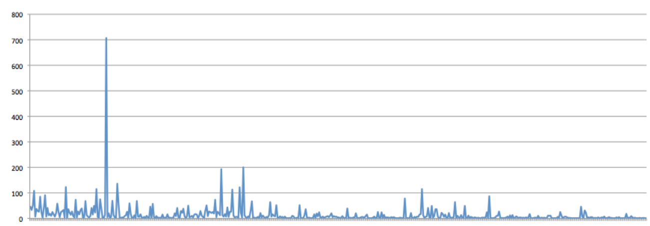

Fig. 2. Distribution of the POIs in dataset A along the different NAICS codes

Further analyses on the coherence between NAICS and Yahoo! show that only in 80.2% of the POIs in dataset A the correspondent NAICS was consistent with the most common one for that given set of categories, which means that about one fifth of the POIs are incoherent with the rest of the sample. This fact highlights the problem of allowing users to add arbitrary categories to their POIs without restrictions. For different NAICS levels, particularly for two-digit and four-digit NAICS, the same analyses showed, as expected, a higher level of coherency. For the two and four-digit NAICS, 87.1% and 83.4% of the POIs, respectively. Therefore, by having the same set of Yahoo! categories mapping to different NAICS codes in different occasions, it is not expectable that we obtain a perfect model that classifies correctly all test cases. In order to understand the impact of these inconsistencies in the results, we also modified the POI dataset so that the NAICS code of a given POI would match the NAICS codes of the other POIs with the same category set, assigning to each POI the most common NAICS code for that given category set in the dataset. The results of this experiment are also presented in Section IV.

Tables II and III show, respectively, the five most common NAICS and Yahoo! categories we identified in dataset A.

Regarding dataset B, we identified 689 distinct NAICS codes and 1109 distinct categories of the more than 1300 that we found in Yahoo!. The latter are in larger number

TABLE IV. BRIEF DESCRIPTION OF THE ALGORITHMS TESTED

Implementation	Description
ID3	Unpruned decision tree based on the ID3 algorithm.
C4.5	Pruned or unpruned C4.5 decision tree.
C4.5graft	Grafted C4.5 decision tree.
RandomForest	Forest of random trees.
RandomTree	Tree with K randomly chosen attributes at each node. Performs no pruning. Also has an option to allow estimation of class probabilities based on a hold-out set (backfitting).
JRip	Propositional rule learner, Repeated Incremental Pruning to Produce Error Reduction (RIPPER), as proposed by W. Cohen as an optimized version of IREP.
IBk	K-nearest neighbors classifier. Can select appropriate value of K based on cross-validation. Can also do distance weighting.
IB1	1 - nearest-neighbor classifier. Simplification of IBk.
K*	K* is an instance-based classifier. The class of a test instance is determined from the class of similar training instances . It uses an entropy-based distance function.
BayesNet	Bayesian Network
NaiveBayes	Naive Bayes model

than the ones from dataset A (only 802) and therefore dataset B provides a better coverage of the source taxonomy. The number of distinct category combinations almost doubled when compared to dataset A, which leads to more diversity in the training data and more accurate classifiers.

Figure 3 shows the distribution of POIs from dataset B along the different NAICS codes. Similarly to the distribution of dataset A, it is an irregular distribution.

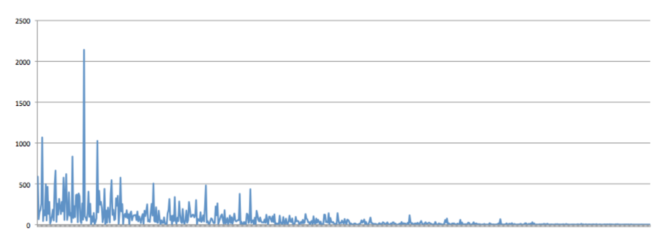

Fig. 3. Distribution of the POIs in dataset B along the different NAICS code

C. Flat Classification

The "flat classification" task corresponds to directly assigning a NAICS code to a POI given its "bag" of Yahoo! categories. It is "flat" because the inherent hierarchy of NAICS is not taken into account in the classification model. Each NAICS code is simply seen as an isolated string "tag" that is assigned to a POI.

We experimented various machine learning algorithms for this particular classification task. Table IV provides a brief description of the algorithms we tested. It is not the scope of this article to describe any of the algorithms in detail. The interested reader is redirected to dedicated literature [22], [23].

In our experiments, we built classifiers for different NAICS levels (i.e., NAICS categories with different granularities), particularly two, four and six-digit NAICS codes. This choice is typical in urban planning depending on the study at hand (e.g., level 2 allows to analyze economic sectors, while level 6 goes to the level of the establishment specificities).

For validation purposes we use ten-fold cross-validation [23]. We also performed validation with an external test set (data from a another city, New York) to understand the dependency of the generated models on the study area.

D. Hierarchical Classification

In this approach, we take advantage of the hierarchical structure of NAICS, thus the overall classifier is itself a hierarchy of classifiers. In this hierarchy, each classifier decides what classifier to use next, narrowing down the NAICS code possibilities on each step, until a final 6-digit code (or 4-digit code, depending on the goal) is achieved. Figure 4 depicts one possible hierarchy.

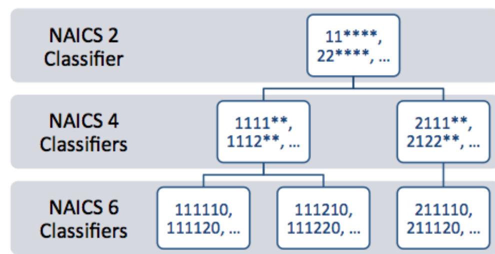

Fig. 4. A possible hierarchy of classifiers

By looking at the hierarchy above, we can see that it has 3 levels (2, 4 and 6-digit NAICS). The first level always consists of a single classifier that decides which NAICS economic sector (2-digit code) the POI belongs to. Taking the sector into account, the algorithm then decides which classifier to use next at the second level. After that, the same process repeats until a leaf node is achieved in the tree structure of the hierarchy of classifiers. To provide an example consider a POI that has the following NAICS code: 111110. According to Figure 4 the top-level classifier will decide that it belongs to sector 11 ("Agriculture, Forestry, Fishing and Hunting") and the left-most level 2 classifier will be used next. Then, this classifier will determine that the 4-digit NAICS code of the POI is 1111 ("Oilseed and Grain Farming") and, based on this decision, the left-most classifier in the third level of the Figure 4 will be used, and will supposedly classify the POI with the NAICS code 111110 ("Soybean Farming"). Of course, along this top-down process a mistake can be made by one classifier. In this case, the error would propagate downwards and there would be no way to recover from it, and hence the final NAICS code would be wrong.

Our hypotheses is that by using a hierarchy of classifiers, the classification task will be divided into several classification models, each one less complex, more accurate and dealing with a simpler problem. If we consider, for example, the ID3 algorithm, the entropy values for the different features will be computed according to a smaller class subset, and therefore the selection of the next feature to use (which is based on the entropy calculation) will be different and the resulting tree will also be different. Hopefully, the generated classifier will be more suited to that particular classification (like deciding for a POI if it belongs to the subcategory 531, 532, etc, knowing that it belongs to NAICS sector 53).

In our experiments, we use three different hierarchies of classifiers, two with 2 levels:

- NAICS 2 and NAICS 4
- NAICS 2 and NAICS 6

and other one with 3 levels:

- NAICS 2, NAICS 4 and NAICS 6

As we did for the flat classification, we also tried to test different types of machine learning algorithms: bayesian networks, tree-based learners, instance-based learners, rule-based learners. Neural networks were not possible to test due to their computational demands, both in processing power and memory.

For the hierarchical approaches we also perform ten-fold cross-validation, but the data splitting between training/testing is more prone to biased results than with standard flat classification. As in normal ten-fold cross-validation, we also start by leaving 10% of the data out for test and use the remaining 90% for training, repeating this process ten times. However, each classifier in a given level only receives the part of those 90% of training data that respects to it. For instance, a level two classifier for deciding which sub-category of NAICS sector 53 a given POI belongs to would only be trained with POIs that belong to that NAICS sector. Hence, the only classifier that receives all the training data (90%) would be the top-level classifier (i.e., the one that decides which NAICS sector a POI belong to). After the training phase, the hierarchy is tested with the 10% of the data left out. This process is repeated ten times, and the average accuracy over the ten iterations is determined.

IV. RESULTS

Table V shows the results obtained using different machine learning algorithms for different NAICS levels (two, four and six-digit codes) for dataset A. There are some missing results in the cases where the algorithm took over 72 hours to run. We can see that the tree-based (e.g., ID3, RandomForest) and instance-based learning approaches (e.g., IBk, K*) are the ones that perform better in this classification task, especially the latter. Notice that only 80.2% of data is classified in a totally non-ambiguous way. The most successful algorithm is IBk (with k=1), which essentially finds the similar test case and assigns the same NAICS code. The difference in accuracy between tree-based and instance based approaches is too small to conclude which one outperforms the other, however we could expect that instance based models bring better results since the distribution of the different Yahoo! categories is relatively even among examples of the same NAICS code (implying no clear "dominance" of some categories over others). Understandably, the Naive Bayes algorithm performs badly because the assumption that different Yahoo! categories for the same NAICS classification are independently distributed is obviously false (e.g., "Doctors & Clinics, Laboratories, Medical Laboratories" are correlated). Such assumption is not fully necessary in Bayesian Networks, which actually brings better results. Unfortunately, we could not find a model search algorithm that performs in acceptable time (less than 72 hours)

TABLE V. RESULTS OBTAINED FOR THE DIFFERENT MACHINE LEARNING ALGORITHMS WITH POIS FROM DATASET A FOR THE BOSTON AREA

Algorithm	NAICS2(kappa)	NAICS4(kappa)	NAICS6(kappa)
ID3	85.495 (0.842)	77.955 (0.776)	74.015 (0.737)
C4.5	84.241 (0.828)	77.630 (0.772)	73.071 (0.727)
Random Forest	86.174 (0.849)	79.298 (0.789)	74.753 (0.744)
Random Tree	85.303 (0.840)	77.763 (0.774)	74.192 (0.739)
JRip	81.334 (0.795)	74.340 (0.737)	69.264 (0.686)
IB1	82.736 (0.812)	74.266 (0.738)	68.644 (0.683)
IBk	86.646 (0.854)	79.475 (0.791)	75.343 (0.750)
K*	85.702 (0.844)	79.726 (0.794)	75.387 (0.751)
BayesNet	80.950 (0.790)	56.721 (0.554)	45.064 (0.438)
NaiveBayes	74.399 (0.715)	40.446 (0.382)	30.264 (0.283)

TABLE VI. RESULTS OBTAINED FOR THE DIFFERENT MACHINE LEARNING ALGORITHMS USING A RE-CLASSIFIED VERSION OF DATASET A

Algorithm	NAICS2	NAICS4	NAICS6
ID3	92.975	89.728	88.680
RandomForest	93.609	90.805	89.846
IBk	94.170	91.189	89.979

and produces a more accurate model. We used Simulated Annealing and Hill Climbing.

As expected, we obtained better results classifying POIs to the two-level NAICS than for the six-level NAICS, since the noise due to ambiguous classifications in the POI dataset is smaller (we now have 87.1% of non-ambiguous cases).

In Table VI, we can see the results obtained by changing the POI dataset A so that the NAICS codes of POIs where ambiguities arise are grouped together in the same "super-category", eliminating the inconsistencies.

By comparing the results in Table VI with the results in Table V, we realize that the NAICS labeling inconsistencies in the POI data have a major negative effect in the performance of the machine learning algorithms, reducing the accuracy in more than 16% in some cases for the six-level NAICS codes. This also gives indications for future versions of NAICS, where some categories may become aggregated according to these "super-categories".

It would be expectable to obtain accuracies closer to 100% for the results in Table VI. However, that does not happen since 115 of the 514 NAICS codes covered by our dataset A only occur once. Therefore, when we split the dataset to perform the ten-fold cross-validation, a significative number of the test cases will have NAICS codes that the algorithm was not trained for, causing it to incorrectly classify them.

Table VII shows the results we obtained by training the machine learning approaches with dataset A from Boston and Cambridge and testing them with New York POI data. As we can see in the results, if we apply the generated model to a different city, it still performs well, even though the accuracy drops a small amount in some cases. This is understandable since even the Yahoo! taxonomy differs slightly from city to city.

Table VIII shows the results obtained for the different machine learning algorithms using dataset B.

TABLE VII. RESULTS OBTAINED FOR THE DIFFERENT MACHINE LEARNING ALGORITHMS USING POI DATA FROM BOSTON FOR TRAINING AND POI DATA FROM NEW YORK FOR TESTING

Algorithm	NAICS2	NAICS4	NAICS6
ID3	85.061	75.586	70.209
RandomForest	85.488	76.867	71.318
IBk	85.360	76.909	71.276

TABLE VIII. RESULTS OBTAINED FOR THE DIFFERENT MACHINE LEARNING ALGORITHMS WITH POIS FROM DATASET B FOR THE BOSTON AREA

Algorithm	NAICS2(kappa)	NAICS4(kappa)	NAICS6(kappa)
ID3	90.567 (0.897)	85.459 (0.852)	82.091 (0.819)
C4.5	90.113 (0.800)	85.085 (0.849)	81.831 (0.816)
RandomForest	90.758 (0.899)	85.710 (0.855)	82.436 (0.823)
RandomTree	90.500 (0.896)	85.275 (0.851)	81.818 (0.817)
JRip	85.748 (0.844)	80.998 (0.807)	78.495 (0.780)
IB1	87.224 (0.861)	81.495 (0.812)	76.826 (0.766)
IBk	91.024 (0.902)	85.974 (0.858)	82.553 (0.824)
K*	90.227 (0.893)	85.849 (0.856)	82.522 (0.824)
BayesNet	88.961 (0.880)	77.964 (0.776)	67.877 (0.675)
NaiveBayes	87.910 (0.868)	70.250 (0.696)	56.052 (0.554)

TABLE IX. COMPARISON BETWEEN THE RESULTS FOR DATASET B USING FLAT CLASSIFICATION (4-DIGIT NAICS) AND HIERARCHICAL CLASSIFICATION WITH 2 LEVELS (NAICS 2 AND 4)

	Flat classification	Hierarchical classification	
Algorithm	accuracy	Level1 acc.	Level2 acc.
ID3	85.459	90.659	85.620
C4.5	85.085	90.172	84.901
RandomForest	85.710	90.959	85.969
RandomTree	85.275	90.509	85.315
JRip	80.998	85.806	80.440
IB1	81.495	87.637	81.126
IBk	85.974	91.080	86.097
K*	85.849	90.305	85.244
BayesNet	77.964	88.002	74.243
NaiveBayes	70.250	30.688	20.091

By analyzing the results from Table VIII we can see that the results have significantly improved over dataset A, which shows the importance of the training data in the performance of the machine learning algorithms.

Finally, Tables IX to XI show the results obtained using the different hierarchical classification schemes for various types of machine learning algorithms.

Our intuition was that hierarchical classification would perform generally better than standard flat classification. However, only in some algorithms the results improved. Therefore, we will not argue that hierarchical classification of POIs into NAICS is always a better solution. In fact, as shown before by comparing the datasets A and B, the quality and the dimensions of the dataset seems to have a much bigger impact on the results than whether we apply hierarchical or flat classification.

Another interesting fact in the results from the hierarchical classification is that the accuracies vary considerably with the hierarchy type used. For instance, when classifying POIs with 6-digit NAICS codes, we can see that using a two-level hierarchy the RandomForest algorithm improved over the flat classification, while using a three-level hierarchy it became worse (although the differences in accuracy are small).

TABLE X. COMPARISON BETWEEN THE RESULTS FOR DATASET B USING FLAT CLASSIFICATION (6-DIGIT NAICS) AND HIERARCHICAL CLASSIFICATION WITH 2 LEVELS (NAICS 2 AND 6)

	Flat classification	Hierarchical classification	
Algorithm	accuracy	Level1 acc.	Level2 acc.
ID3	82.091	90.659	82.100
C4.5	81.831	90.173	81.484
RandomForest	82.436	90.959	82.477
RandomTree	81.818	90.509	81.654
JRip	78.495	85.806	76.398
IB1	76.826	87.637	76.826
IBk	82.553	91.080	82.551
K*	82.522	90.305	81.661
BayesNet	67.877	89.059	69.336
NaiveBayes	56.052	88.002	59.885

TABLE XI. COMPARISON BETWEEN THE RESULTS FOR DATASET B USING FLAT CLASSIFICATION (6-DIGIT NAICS) AND HIERARCHICAL CLASSIFICATION WITH 3 LEVELS (NAICS 2, 4 AND 6)

	Flat classification	Hierarchical classification		
Algorithm	accuracy	Level1 acc.	Level2 acc.	Level3 acc.
ID3	82.091	90.659	85.620	82.111
C4.5	81.831	90.172	84.901	81.341
Random Forest	82.436	90.959	85.969	82.398
Random Tree	81.818	90.509	85.315	81.694
JRip	78.495	85.806	80.440	76.889
IB1	76.826	87.637	81.126	76.826
IBk	82.553	91.080	86.097	82.539
K*	82.522	90.305	85.244	81.486
BayesNet	67.877	-	-	-
NaiveBayes	56.052	-	-	-

V. AN APPLICATION IN URBAN PLANNING

In this section, we describe a practical application of Yahoo! POIs classified to NAICS using a non-hierarchical approach with the k-nearest neighbor classifier (see Section III-C for more details).

In the field of urban planning, urban simulation models have evolved significantly in the past several decades. For instance, the travel demand modeling approach has been evolving from the traditional Four-Step Model (FSM) to the Activity-Based Model (ABM) [24]. Consequently, requirements for disaggregated data increase greatly, ranging from population data, employment data, to travel survey data. The employment data (on the travel destination side) is usually obtained from proprietary sources, which adds another layer of barriers to widely applying the Activity-Based Modeling approach, let alone the expensive travel-survey data acquisition. In order to study this issue, researchers are trying to develop new methods of estimating disaggregated employment size and location by category.

In our case, we intend to develop a set of new methods and demonstrate their applications for estimating activities, incorporating them into travel demand and urban simulation models. This will be beneficial for cities that lack detailed survey data for building Activity-Based Models but wish to test the sensitivity of travel behavior to policy changes such as Intelligent Transportation Systems (ITS) implementations that are likely to alter activity patterns. An important step to achieve these goals is to obtain a disaggregated employment distribution by POIs of an area. For the case of Cambridge, MA, we have official data at the Block Group (BG) level (obtained from the U.S. Census Transportation Planning Package 2000), which essentially describes the total size of employees by economic sector at that spatial resolution. We need to distribute these totals into Block or Parcel level.

For demonstration purposes we only use POIs from the "Retail Trade" sector of the NAICS taxonomy, i.e., categories whose code starts by 44 or 45. Figures 5 and 6 show the aggregated retail employment density at the Block Group level and distribution of our POI data from Yahoo! at the Census Block level for Cambridge, respectively.

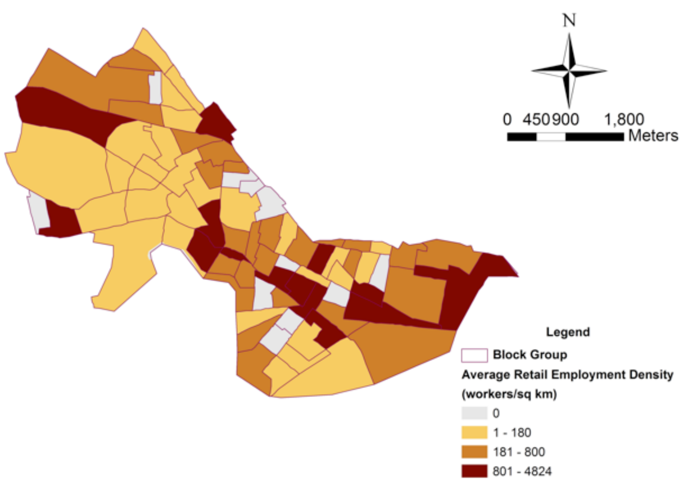

Fig. 5. Aggregated retail employment density at the Block Group level (pl/sq km= employed people per square kilometer).

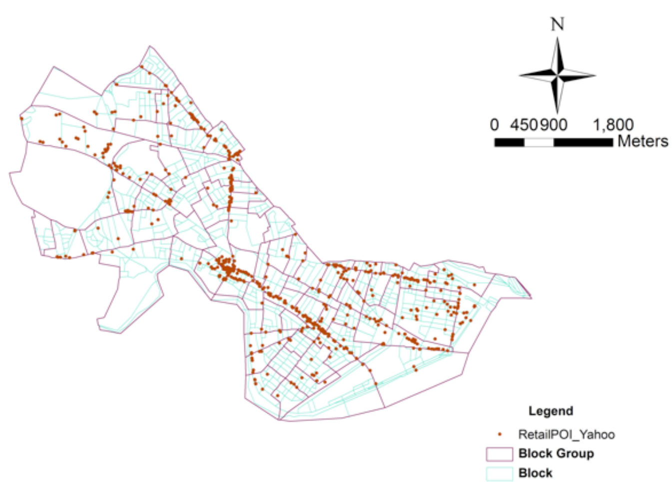

Fig. 6. Cambridge retail POI distributions from Yahoo!

By using the business establishment survey data (from InfoUSA, 2007) which is believed to be close to the population, we are able to obtain a benchmark estimate of employment size by category at the Census Block level for the study areas. This will function as a ground truth to test our algorithm. Notice however that the dates for each of the databases are quite distinct (2000 for Census, 2007 for InfoUSA and 2010 for Yahoo!) therefore some error is expected to happen.

We employ the local maximum likelihood estimation (MLE) method as described below to derive the disaggregated destination estimation at Block level.

1) We calculate the total number of POIs (destinations) by category c in each Block b.
2) We assume that the employment size at destination d in Block Group g of category c is proportional to some function f of its associated block area $a_{d,c,g}$, which means the effective area of the destination d in Block Group g of category c. The form of function f will be explored based on the empirical data, and we also allow the possibility that $f(a_{d,c,g}) = a_{d,c,g}$ which is the natural benchmark case. Mathematically, assume that for employment category c, there are $n_{c,g}$ destinations at Block Group g. For $d = 1, 2, \dots, n_{c,g}$, let random variable $e_{d,c,g}$ be the employment size of category c at destination d in Block Group g.
3) We assume that $e_{d,c,g}(d = 1, 2, \dots, n_{c,g})$ are i.i.d.$(f(a_{d,c,g}) \cdot \alpha_{c,g}, \sigma^2_{c,g})$, where $\alpha_{c,g}$ is the employment size of category c per unit of effective area at Block Group g; $\alpha_{c,g}$ and $\sigma_{c,g}$ are positive constants independent of d. $E(e_{d,c,g}) = f(a_{d,c,g}) \cdot \alpha_{c,g}$ and $Var(e_{d,c,g}) = \sigma^2_{c,g}$. We then estimate $\alpha_{c,g}$ by employing the maximum likelihood method locally at Block Group g for employment category c. Thus we obtain an estimate of employment size $e_{d,c,g}$ of category c at destination d in Block Group g.
4) Finally, we sum up the employment size in category c in Census Block b in Census Block Group g.

By employing the same local maximum likelihood method described above and using the business establishment survey data (e.g., ESRI Business Analysis package) which is believed to be close to the population POIs, we obtain a benchmark estimate of employment size by category at the Block level for the study area, $E^*_{b,c,g}$. By using the derived POI information (obtained from the machine learning algorithm), we obtain an estimate of employment size by category c at Block b for the study area, $\hat{E}_{b,c,g}$.

Then the mean squared error (MSE), weighted mean squared error (WMSE), and the relative weighted mean squared error (RWMSE) can be calculated to evaluate the goodness of fit of the model (see Equations 1, 2, 3, and 4).

$$\mathrm{MSE}(\hat{E}_{b,c,g}, E^*_{b,c,g}) = \sum_{b,c,g} (\hat{E}_{b,c,g} - E^*_{b,c,g})^2 \tag{1}$$

$$\mathrm{WMSE}(\hat{E}_{b,c,g}, E^*_{b,c,g}) = \sum_{b,c,g} w_{b,c,g} (\hat{E}_{b,c,g} - E^*_{b,c,g})^2 \tag{2}$$

$$\mathrm{RWMSE}(\hat{E}_{b,c,g}, E^*_{b,c,g}) = \frac{\sum_{b,c,g} w_{b,c,g} (\hat{E}_{b,c,g} - E^*_{b,c,g})^2}{\sum_{b,c,g} w_{b,c,g} (\bar{E}_{b,c,g} - E^*_{b,c,g})^2} \tag{3}$$

$$\bar{E}_{b,c,g} = \frac{w'_{b,g} \sum_q E^*_{q,c,g}}{\sum_q w'_{q,g}} \tag{4}$$

Weights $\{w_{b,c,g}\}$ are normalized to reflect the proportion of each Census Block in the whole map. In Equation 2, when we take the weight $w_{b,c,g} = 1$ for any subscripts b, c, and g, the corresponding WMSE becomes MSE. In Equation 4, $w'_{b,g} =$ area of Block b in Block Group g, and $\bar{E}_{b,c,g}$ is the estimated employment size in Block b of category c, using the traditional disaggregation approach, assuming that the employment is uniformly distributed across blocks in each Block Group g.

If RWMSE is less than 1, it means that the quality of the derived POIs is reliable, so is the new method; the smaller the RWMSE, the more accurate is the method. If WMSE or RWMSE equals to 0, it means that the derived POIs from the Internet match exactly with the trusted proprietary POIs (treated as the population POIs). However, if RWMSE is greater than 1, it means that the derived POIs cannot well reflect the distribution of the population POIs.

Figures 7 and 8 show the estimation results of the disaggregated retail employment density at Block level in Cambridge, MA, by using POIs from infoUSA and Yahoo! respectively. By comparing the estimation results, we find that the disaggregated employment estimations by using the POIs captured from the Internet using Yahoo! and those obtained from the proprietary source (infoUSA 2007) are very close.

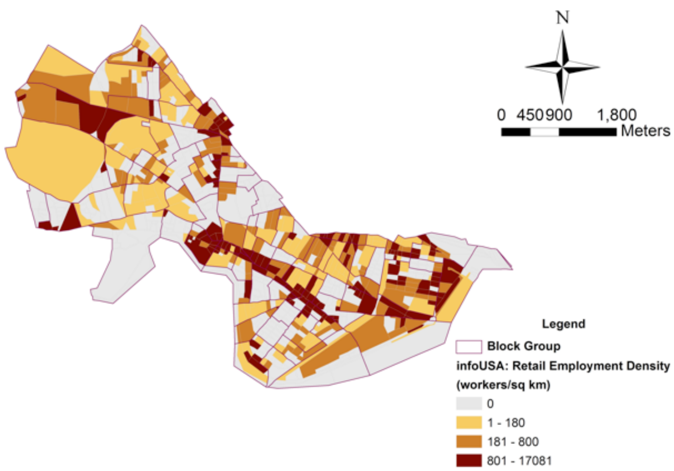

Fig. 7. Disaggregated retail employment densities at the Block level, in Cambridge, MA, by using POIs from infoUSA

Employing Equation 3, the disaggregated employment estimation at the Block level using Yahoo! POI gives RMSE = 0.312. The RMSE is significantly smaller than 1, which means that using the extracted Yahoo! online POIs to estimate the disaggregated employment sizes at the Block level has reduced the mean squared error by around 69% compared to the traditional average disaggregation approach.

VI. Model performance analysis with visual techniques

In this section, we explore visualization techniques to further assess the performance of the developed models. Furthermore, by making use of these visualization techniques we are able to identify possible ways to improve the models' performance and establish future work directions.

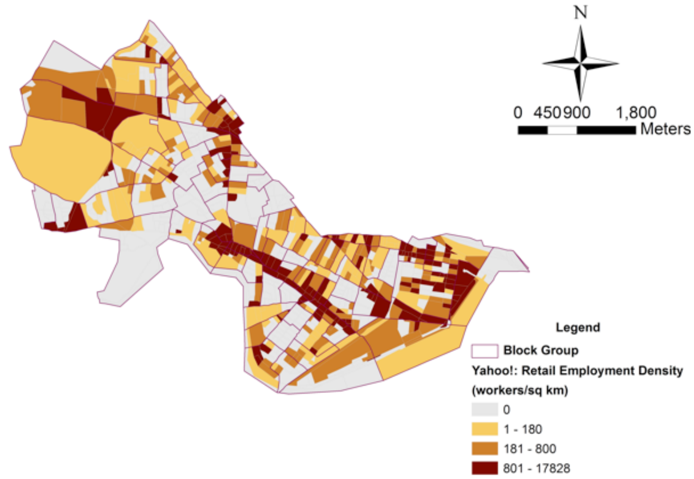

Fig. 8. Disaggregated retail employment densities at the Block level, in Cambridge, MA, by using POIs from Yahoo!

We start by visualizing the errors of the learned POI classifier. Due to the hierarchical structure of the NAICS taxonomy, there are 6 types of error depending on the level at which classifier fails first. Therefore, we define the error level $\xi(p)$ of some POI p as the first index, from right to left, where the predicted NAICS code is different from the true one. So, for example, if the classifier mislabels a POI from NAICS 921130 as 921120, we say it made a level 2 error ($\xi(p) = 2$). On there other hand, if it mislabels it as 238320 we say it made a level 6 error ($\xi(p) = 6$). Correctly classified POIs are said to have an error level of 0. Hence, the higher the error level, the worst its consequences are for practical applications that use the classified POI data. In particular, errors of levels 5 and 6 are especially bad, since they mean the classifier was unable to correctly identify the economic sector of the POI, thus invalidating analysis even at the coarsest level of granularity.

With this in mind, we developed a colour-based representation scheme, where different colours represent different error levels. Figures 9 and 10 show this visualization method applied to the Boston metropolitan area and Cambridge respectively. Figure 10 shows that the majority of the POIs are correctly classified, and that most of the errors that occur correspond to higher level errors (i.e., level 5 and 6). These findings are coherent with the results presented in Table VIII. However, we can now observe that these errors are not uniformly distributed across space. Contrarily, we can see that, regardless of the error level, the higher the POI density is in a given region, the more likely it is to have high quantities of misclassified POIs. Although intuitive, notice that this observation is of vital importance for practical uses of the classified POI data.

Despite the possibility to visualize the spatial distribution of the error of the POI classifier, the visualization method described above has some limitation when it comes to comparing different error level. This led us to the idea of using small multiples - a technique commonly used to compare various versions of data sets with the same structure in order to show shifts in relationship [25]. In other words, the different layers were separated and afterwards composed in sequenced order. The resulting artefact can be seen in Figure 11. Apart from the

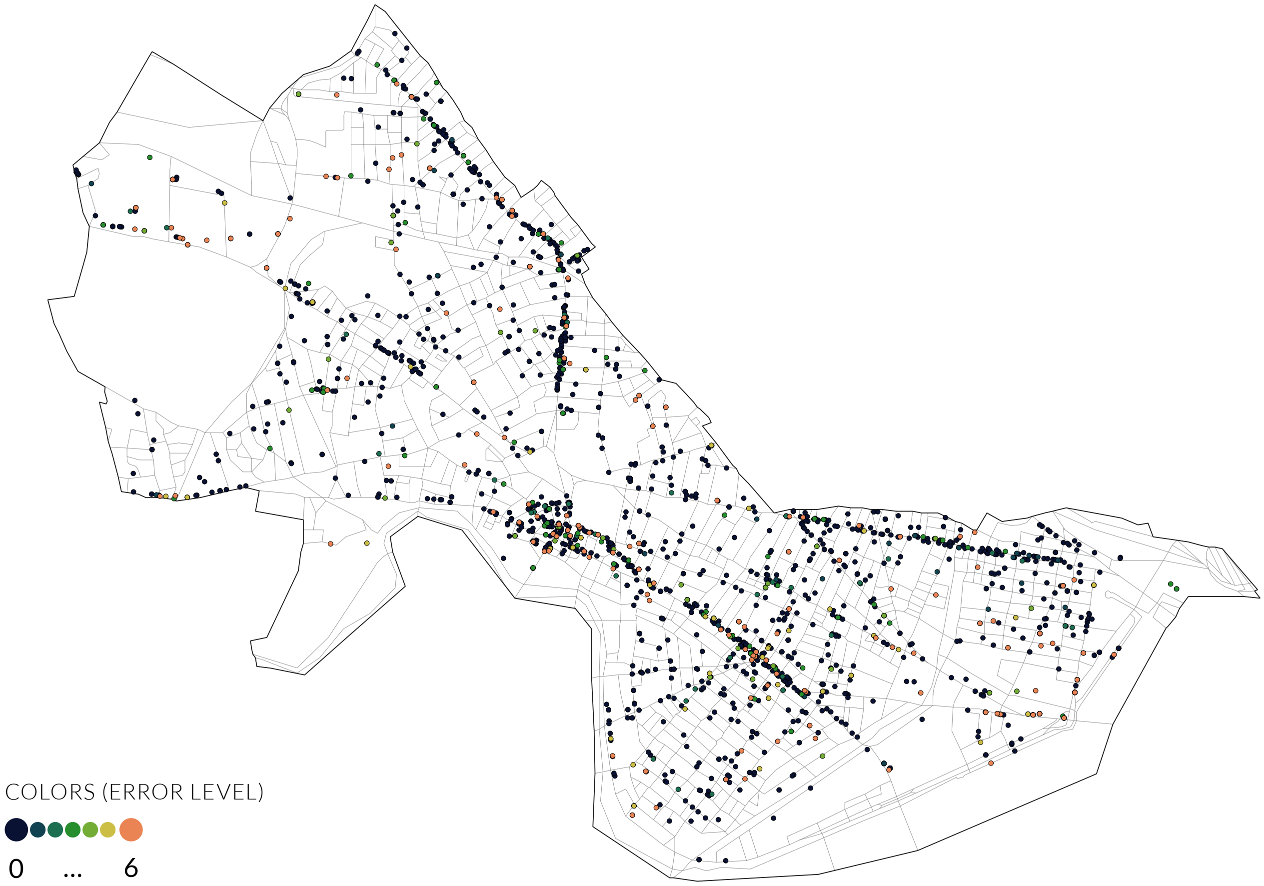

Fig. 10. Spatial distribution of the POI classification error in the area of Cambridge, MA.

fact that most of the times the POIs are correctly classified, as the figure evidences, the distribution of the misclassified POIs is not uniform even at the individual error levels. Furthermore, we can see that level 6 errors are the most frequent ones.

The small multiples visualization allowed us to visualize and compare the spatial distribution of different error levels. However, it is hard to really understand the relation between classification error and POI density from these methods. In order to better understand this relation, a new visualization method was developed. The first step consists in constructing a modular grid over the map, whose size can be manipulated, which then reflects on the resulting visual impact. Then, the average error level $\overline{\xi_j}$ of a grid cell j is simply calculated as follows:

$$\overline{\xi_j} = \frac{1}{|P_j|} \sum_{p \in P_j} \xi(p) \qquad (5)$$

where P_j denotes the set of all POIs in the grid cell j. The value of $\overline{\xi_j}$ is then used to visually represent how severe the misclassification is in that grid cell. By having a single-value representation of the misclassification error for a given region, we can now easily overlay this information with POI density. Figure 12 shows an example of this visualization method for the Boston metropolitan area. The radius of the black circles represent the values of ξ_j and the different intensities of red represent the POI density.

The visualization from Figure 12 allows us to compare POI density and average error per cell. Hence, we can see that saturated red cells with small circles in it represent optimal areas in terms of the reliability of the POI data, since the average classification error is small even with high POI densities. On the other hand, the cells with large circles and almost white background colour represent areas with many classification errors, even though the number of POIs is small. The worst case would be for a cells to have large circles and red saturated colours. That would mean that there are many incorrectly classified POIs. As Figure 12 evidences, there are no such regions in our classified data. Furthermore, we can see that the regions where misclassification is more severe, correspond to regions of low POI density. Thus, from a perspective of the “functional regions” of space, where each region is defined by the economical activities that take place there, these are the areas where the resultant analysis would be less reliable.

The visualization method described above can easily be extended to include the performance of practical application of the POI data - in this case, the employment size estimation model from Section V. Figure 13 shows the resulting visual-

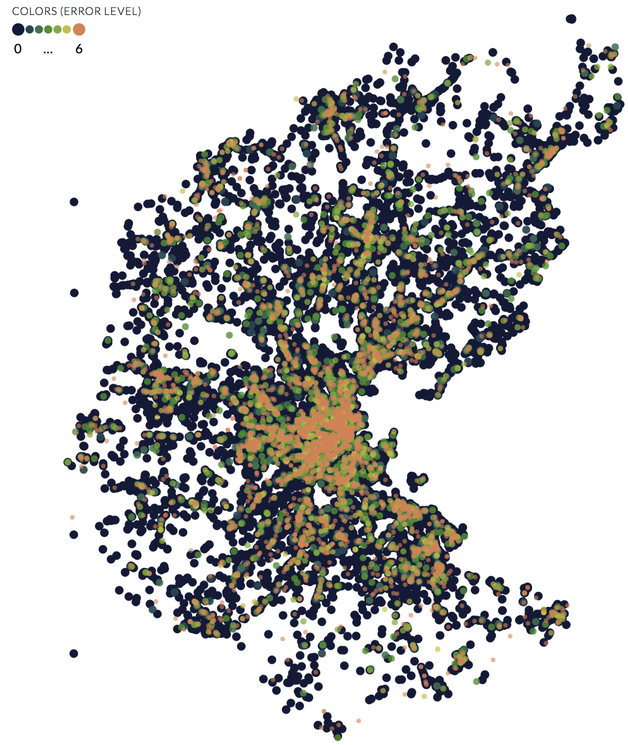

Fig. 9. Spatial distribution of the POI classification error in the Boston metropolitan area.

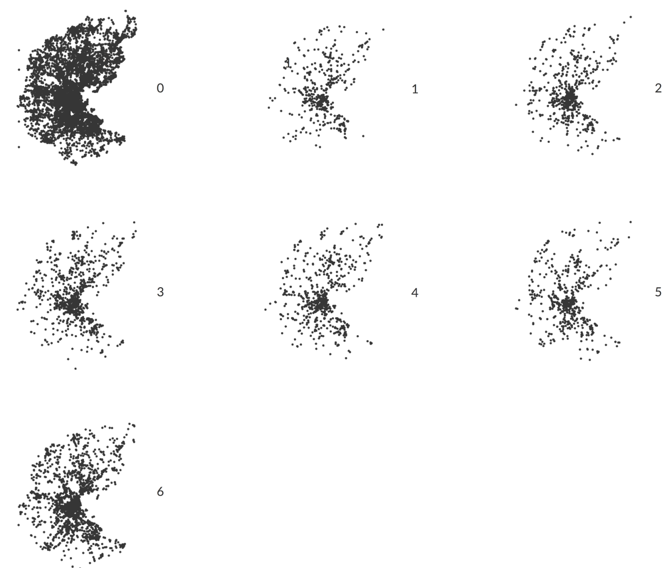

Fig. 11. Distribution of the classification error through space for different error levels ξ.

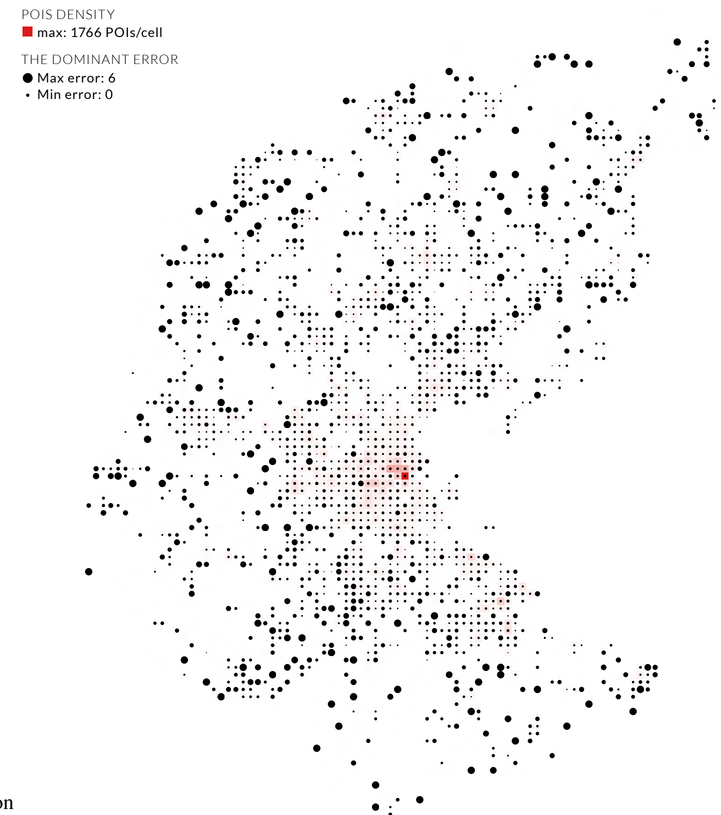

Fig. 12. Visualization of the relation between POI density and average POI classification error level $\overline{\xi_j}$ for the Boston metropolitan area. High intensity red squared represent zones high POI density. Circle radius represent the value of the average error level $\overline{\xi_j}$.

ization for the study area of Cambridge. As we can see from this figure, there is some relation between POI classification error and estimated employment size error. In fact, a closer look at this relation shows a Spearman correlation of 0.179 (p-value = 0.006). Hence, a future work direction should be on improving the classifier's performance. But, more importantly, Figure 13 suggests the hypothesis that regions with lower POI densities are more likely to have higher estimated employment size errors, which in turn favors the idea that building a more comprehensive POI database, through the aggregation of multiple online sources of data, constitutes a very important future work direction.

VII. Conclusion and Future Work

In this paper, a complete framework for employment size estimation at a disaggregated level based on online collaborative POI data mined from the Web was proposed. We empirically showed that it possible to classify POI to the widely used NAICS taxonomy with several different machine learning algorithms using only the categories or tags that are commonly associated with them. We matched two different POI databases (InfoUSA and Dun & Bradstreet) to Yahoo!, in order to build two reliable training sets that have POIs with user provided bags of categories classified with NAICS codes. We tested several classification algorithms and the results show that the best approaches for this particular task are inductive based algorithms, namely instance based and tree based learning. These allow for an accuracy as high as 82% in the most complex task (classification with 6-digit NAICS codes). We also tried to perform classification in a hierarchical way, however the results did not showed many improvements

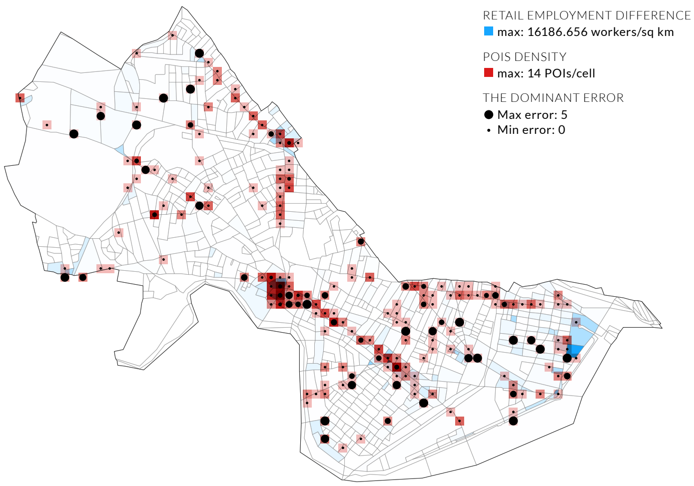

Fig. 13. Relation between POI density (red), estimated disaggregated employment size error (blue), and POI classification error (black circles) for the area of Cambridge, MA.

over the flat approaches, leading us to the conclusion that the size of the training set and its consistency/quality can have a larger impact on the results than the classification algorithm itself (except maybe for Bayesian approaches).

The classified POIs were applied to the urban planning task of employment size and location disaggregation from Block Group level to Block level and the results show encouraging quality. This strengthens the idea that well classified POI data to a convenient taxonomy like the NAICS is of great use and can have many distinct applications.

To the authors best knowledge, this is the only work that proposes an automatic approach for classifying POIs to the NAICS, and, therefore, a comparison with other works is not possible. Thus, we contribute with a novel approach to this important problem that has high impact in urban planning and space classification.

Furthermore, we propose several visualization techniques to help improve the overall quality of the proposed framework, by helping us to (1) identify regions of high classification error, (2) understand the relationship between POI classification error and POI density, (3) comprehend how POI classification error relates with estimated employment size error, (4) visualize how all these aspect distribute among space, and many other interesting aspects of the models and the data.

Future work will investigate the use of semantic enrichment of the POI data in order to help improve the POI classification accuracy. Furthermore, we also intend to explore other sources of online data, so that a more comprehensive POI dataset could be built.

REFERENCES

[1] F. Rodrigues, F. Pereira, A. Alves, S. Jiang and J. Ferreira. Automatic Classification of Points-of-Interest for Land-use Analysis. Proceedings of the Fourth International Conference on Advanced Geographic Information Systems, Applications, and Services (GEOProcessing), pp. 41–49, 2012.

[2] U. C. Bureau. North american industry classification system (naics): Introduction, February 2010. http://www.census.gov/eos/www/naics/.

[3] N. Association. Naics association: Frequently asked questions, February 2010. http://www.naics.com/faq.htm.

[4] D. T. Lindgren. *Land-use Planning and Remote Sensing*. Martinus-Nijhoff, Boston, MA, 1985.

[5] L. O. Fresco. *The Future of the Land – Mobilizing and Integrating Knowledge for Land-use Options*. John Wiley & Sons, Chichester, 1997.

[6] J. B. Campbell. *Mapping the Land – Aerial Imagery for Land use Information*. Association of American Geographers, Washington, D.C., 1983.

[7] P. Danoedoro. Extracting land-use information related to socio-economic function from quickbird imagery: A case study of semarang area, indonesia. *Map Asia 2006*, 2006.

[8] D. Li, K. Di, and D. Li. Land use classification of remote sensing image with gis data based on spatial data mining techniques. *Geo-Spatial Information Science*, 3:30–35, 2000.

[9] M. Santos and A. Moreira. Automatic classification of location contexts with decision trees. *CSMU-2006 : Proceedings of the Conference on Mobile and Ubiquitous Systems, Guimares, Portugal*, 2006.

[10] T. Griffin, T. Huang, and R. Halverson. Computerized trip classification of gps data. *Proceedings of 3rd International Conference on Cybernetics and Information Technologies, Systems and Applications (CITSA 2006)*, 2006.

[11] J. Pierre. On the automated classification of web sites. *Linkoping Electronic Articles in Computer and Information Science*, 6, 2001.

[12] R. P. Haining. *Spatial data analysis in the social and environmental sciences*. Cambridge University Press, Cambridge, 1990.

[13] L. Anselin and R. Florax. *New directions in spatial econometrics*. Springer, New York, 1995.

[14] D. Keim, C. Panse, M. Sips and S. North. Pixel based visual mining of geo-spatial data. *Computers & Graphics (CAG)*, 2004, vol. 28.

[15] D. Costa and G. Venturini. A Visual and Interactive Data Exploration Method for Large Data Sets and Clustering. *Advanced Data Mining and Applications, Lecture Notes in Computer Science*, 2007, vol. 4632, pp. 553–561.

[16] E. Currid and J. Connolly. Patterns of knowledge: The geography of advanced services and the case of art and culture. *Annals of the Association of American Geographers*, 98:414–434, 2008.

[17] P. Sambidi and W. Harrison. Spatial clustering of the u.s. biotech industry. 2006 Annual meeting, July 23-26, Long Beach, CA 21360, American Agricultural Economics Association (New Name 2008: Agricultural and Applied Economics Association), 2006.

[18] G. Arbia. Modelling the geography of economic activities on a continuous space. *Papers in Regional Science*, 80(4):411–424, 2001.

[19] D. . Bradstreet. D & b website, February 2010. http://www.dnb.com/.

[20] W. Cohen, P. Ravikumar, and S. Fienberg. A comparison of string distance metrics for name-matching tasks. *Proceedings of the IJCAI-2003 Workshop on Information Integration on the Web (IIWeb-03), Acapulco, Mexico*, pp. 73–78, 2003.

[21] J. Gurland and R. Tripathi. A simple approximation for unbiased estimation of the standard deviation. *American Statistician*, 25(4):30–32, 1971.

[22] I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques, 2nd Edition*. Morgan Kaufmann, 2005.

[23] T. M. Mitchell. *Machine Learning*. McGraw-Hill, New York, 1997.

[24] M. McNally and C. Rindt. *The Activity-Based Approach*. Handbook of Transportation Modeling. Elsevier, Amsterdam, London, 2008.

[25] E. Tufte. Visual Display of Quantitative Information. pp. 170–175, 2001.