The Purpose of Motion: Learning Activities from Individual Mobility Networks

Salvatore Rinzivillo*, Lorenzo Gabrielli*, Mirco Nanni*, Luca Pappalardo*†‡, Dino Pedreschi† and Fosca Giannotti*
*Institute of Information Science and Technology (ISTI), National Research Council (CNR), Pisa, Italy
†Department of Computer Science, University of Pisa, Italy
‡Budapest University of Technology and Economics (BME), Budapest, Hungary
Email: name.surname@isti.cnr.it

***Abstract*—The large availability of mobility data allows us to investigate complex phenomena about human movement. However this adundance of data comes with few information about the *purpose* of movement. In this work we address the issue of activity recognition by introducing Activity-Based Cascading (ABC) classification. Such approach departs completely from probabilistic approaches for two main reasons. First, it exploits a set of structural features extracted from the Individual Mobility Network (IMN), a model able to capture the salient aspects of individual mobility. Second, it uses a cascading classification as a way to tackle the highly skewed frequency of activity classes. We show that our approach outperforms existing state-of-the-art probabilistic methods. Since it reaches high precision, ABC classification represents a very reliable *semantic amplifier* for Big Data.**

I. Introduction

Human mobility is driven by people's daily activities, such as going to work or school, shopping, transporting kids, and so on. The digital mobility traces collected through a variety of technologies, from navigation devices to smart phones, allow us to understand people's movements in great detail [1], [2]. However, they generally fail to capture the *purpose* of such movements, i.e. the kind of activity behind each travel. This deficiency is a hard obstacle to the deployment of Big Data in many domains such as urban planning, traffic management, intelligent transportation systems, socio-demographic simulation and nowcasting, and emergency management [3]. For all such applications, information on *why* people move is crucial. The availability of society-wide data about mobility and activity of people would be a driver for a better comprehension of our complex society, and for smarter knowledge services for the individual and the collective sphere.

It is not a surprise, hence, that several researchers tackled the problem of *activity recognition*, i.e. how to infer the kind of activity associated to a travel on the only basis of the observed mobility patterns [4]. Show me how you move, I'll tell what you do. The rationale behind such research follows a two steps method. First, use a small training mobility dataset annotated with activity information, obtained for instance by surveying some volunteers, to learn a classifier that maps mobility-related features into the different kinds of activities. Second, apply the classifier to unlabeled Big Data, to obtain large-scale mobility data annotated with activity information — the activity classifier acts as a *semantic amplifier* for Big Data. It goes without saying that this second step can be successful only if the predictive accuracy of the classifier is extremely high. Unfortunately, none of the methods for activity recognition proposed so far reach adequate performance for semantic amplification. Some existing methods, such as Conditional Random Fields (CRF) [5], obtain very accurate learners that, given a past history of activity-labelled movements of individuals, predict the activity associated to future unlabeled trips of the *same* individuals. As we discuss in this paper, such learners exhibit poor performance when used to predict the activity associated to the movements of *other* individuals, whose data were *not* used in the learning process; which is the situation we face in semantic amplification of unlabeled Big Data.

In this paper, we describe a model for activity recognition targeted explicitly at the semantic amplification of big mobility data, called Activity-Based Cascading (ABC) classification. ABC departs completely from probabilistic approaches for two main reasons: (i) it exploits a set of structural features extracted from the Individual Mobility Network (IMN), a model able to capture the salient aspects of individual mobility; (ii) it uses a cascading classification as a way to tackle the skewed frequency of activity classes (home and work are generally very frequent compared to shopping and leisure). We use a dataset of approximately 7,000 activity-annotated trips obtained from GPS receivers on board of private cars, and show how ABC classification reaches high precision (up to 0.98) and outperforms both state-of-the-art probabilistic methods (CRF) and Decision Tree classifiers. In summary, the novel contributions of the paper are the following:

- Summarization of individual mobility by introducing a novel graph-based representation: Individual Mobility Networks (Section III);
- Selection of predictive features extracted from the IMNs to improve decision tree classification (Section IV-A);
- Enhancing Cascade Classification with label propagation through successive steps (Section IV-B).
- Comparison with state-of-the-art methods to assess validity of the approach (Section V);

II. Related Works

The task of activity recognition consists in assigning a label to a movement according to its relevant characteristics.

The vast literature on the subject may be organized according to the type of movement observed. Many works focus on the movement of individuals to recognize gestures, indoor activities, physical activity levels, surveillance, outlier and intrusion detection [6], [7]. We are mainly interested in those works considering the movement as the physical change in position of the individual, thus leaving a geographical place to reach another one. We can identify two large groups of inference methods: supervised and unsupervised.

Among the supervised approaches, some methods try to infer the mode of transportation [8], [9], the activity performed in a specific location [10], [11] and a combination of the two [4]. The methods that deals with the transportation mode try to infer if the individual is moving by foot, by car, by bike or by public transportation. This annotation exploits several features of the movements, speed, acceleration and, when available, other context data like accelerometer measurements. The learning approaches are based on discriminative methods, like decision trees [9] and conditional random fields (CRF) [4], [12]; or on generative methods, like Hidden Markov Models (HMM). In our context we are not interested in the transportation mode, since our focus is the prediction of the activity at destination. When considering the activity from movement we can distinguish two main approaches: *sequence learning approaches* consider the activity of an individual as a sequence in a fixed temporal period (usually one day) and try to predict the labels for the whole sequence [4]; *episode learning approaches* try to label each single movement episode independently from the others [9].

Unsupervised methods are mainly based on clustering techniques [13], [14], [15] or dimensionality reduction [13], [16]. In [13] the authors analyze an activity-based travel survey conducted in the Chicago metropolitan area with the aim of exploring the daily activity structure of people. They describe how the considered population can be clustered into eight (weekdays) or seven (weekends) representative groups according to the activities performed by the individuals.

III. Individual Mobility Networks

An *Individual Mobility Network* (IMN) describes the individual mobility of an individual through a graph representation of her locations and movements, grasping the relevant properties of individual mobility and removing unnecessary details.

Definition 3.1: An **Individual Mobility Network** (IMN) of an individual u is a directed graph $G_u = (V, E)$, where V is the set of nodes and E is the set of edges. On nodes and edges the following functions are defined:

- $\omega : E \rightarrow \mathbf{N}$ returns the weight of an edge (i.e. the number of travels performed by u on that edge);
- $\tau : V \rightarrow \mathbf{N}$ returns the time spent by the individual in a given location;
- $p_e : E \times T \rightarrow [0, 1]$ estimates the probability $p_e(e, t)$ of observing an individual u moving on edge e at time t;
- $p_l : V \times T \rightarrow [0, 1]$ estimates the probability $p_l(v, t)$ of observing an individual u at location v at time t.

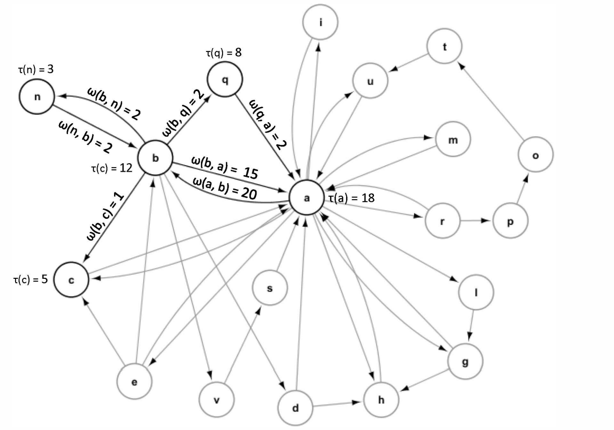

Fig. 1. The IMN extracted from the mobility of an individual. Edges represent the existence of a route between the locations. The function $\omega(e)$ indicates the number of trips performed on the edge e, while $\tau(x)$ the total time spent in a location x.

Nodes represent locations and edges represent movements between locations. We attach to both nodes and edges statistical information by means of structural annotations: edges provide information about the frequency of movements through the ω function; nodes provide an estimation of the time spent in each location through the τ function. To clarify the concept of IMN, let us consider the network in Figure 1. It describes the IMN extracted from the mobility of an individual who visited 19 distinct locations. Location a has been visited a total of 18 time units (days in the example), since $\tau(a) = 18$. The edge $e = (a, b)$ has weight $\omega(e) = \omega(a, b) = 20$, indicating that the individual moved twenty times from location a to location b.

The IMN of an individual is an abstraction of her mobility behavior. A location is an abstract entity without any reference to the geographic space. It can be interpreted as a subjective point of interest, a place around which the mobility of that individual gravitates. This allows the modeling of locations that are meaningful only for that individual, like her home or work place, etc. Accordingly, given the IMNs of two distinct individuals we are not able to determine whether they have visited the same location. This limitation, on the other hand, allow to hide the actual places visited by the individual, providing a protection layer of sensitive information.

The computation of a IMN starts from the ordered sequence of an individual's trajectories. It works in a "streaming" fashion: every time the individual performs a new trip, her IMN is updated according to the new trajectory. The origin and destination points of the new trajectory are mapped to locations in the IMN. The locations are obtained by aggregating all the origin and destination points of the past trajectories within a spatial threshold δ.

Definition 3.2: Let P_u be the set of origin and destination points of an individual u, and δ a distance threshold. A **location** L of u is the aggregation of points in $P_L = \{p \in P_u | \forall p_i, p_j \in P_L : d(p_i, p_j) \leq \delta\}$, where $d(p_i, p_j)$ is the Euclidean distance between p_i and p_j.

When a new trajectory is produced, we search for an existing location, i.e. a place in space within a spatial threshold distance

δ from the current point. If such location exists the point is associated with such location; otherwise a new location is created. Since each update is performed in constant time, the aggregation procedure is very efficient.

IV. Individual Activity Recognition

The task of activity recognition consists in assigning an activity to an individual according to her mobility habits. In this paper we consider the specific issue of detecting the activity related to the destination of a trip, by looking at the characteristics of movements and visited locations. We tackle the problem into two steps: (i) we extracted a set of features from the IMN of each individual; (ii) we learn a classifier on the set of extracted features. To learn the classifier, we consider an extended version of the IMNs where the attributes of each trip include a label which represents the activity performed at the destination node. In the learning phase, two are the critical choices to analyze: the set of features to consider (Section IV-A) and the learning algorithm to adopt (Section IV-B).

A. *Individual Mobility features*

We consider two distinct sets of features. The first set, named *trip features*, regards the characteristics of each single trip performed by an individual. As a second set of features, we introduce a new set of attributes, named *network features*, which regard the topological structure of each IMN.

1) Network features: We identify four classes of network features to capture the salient aspects of a IMN: centrality, predictability, hubbiness and volume.

Centrality measures the degree of connectivity of a location. Given a location, we want to state whether it is at the *center* of the individual's mobility. For example, in the network of Figure 1 location n cannot be considered central, since it is connected to just two edges and visited just two times from the same location b. On the contrary, location a is a central place the individual visits starting from different other locations, presumably her home place. The *clustering coefficient* [17] of a node captures these aspects by estimating the probability that its neighbors are connected to each other. In other terms, it measures the probability that a triad of locations forms a triangle with their edges. The value of the clustering coefficient varies from zero to one: it is when all the triads of nodes are not closed, it is one when all the neighboring triads are connected, representing a dense mobility around the considered locations. Another feature of centrality is the *average path length* [17] of a location x: the average number of edges to traverse to reach the location x from any other node in the network. A location that is at the periphery of the network is likely to have a high average path length. In the example of Figure 1, location o is located at the border of the individual mobility network, having a high average path length. Location a is at the center of the network, resulting in a quite low average path length.

Predictability measures the degree of uncertainty related to the individual's locations. In details, given a location x we want to measure the accuracy of predicting the next location visited from x. Consider for example the central location a in Figure 1, from which the individual starts to visit uniformly many locations. Without other background information it is hard to estimate the probability of the next location. On the contrary, from location n the user visits always the same location b. To model these cases we adopt the concept of Shannon *entropy* [18] to measure the distribution of ingoing and outgoing trips from a node. Formally, the entropy of a location x is given by the formula:

$$E(x) = -\frac{\sum_{y \in V} p(x,y) \log p(x,y)}{\log(N)}$$

where N is the number of locations of the individual, $p(x,y)$ represents the probability of observing a travel from x to y, i.e. $p(x,y) = \omega(x,y)/\sum_{e \in E} \omega(e)$. The entropy is zero when all the outgoing trips are concentrated in a single edge, while it is one when the trips are distributed uniformly over all the available edges.

Hubbiness measures the relevance of a location in the IMN. We can capture the relevance of a node x by counting the different origin locations that have a trip ending in x and, analogously, the number of distinct destination locations for the trips leaving x. These two aspects can be summarized by the *in-degree* and the *out-degree* [17], i.e. the number of incoming and outgoing edges of a location. The relevance of a node (edge) in a graph is also proportional to number of path traversing that node (edge). Such property is synthesized by edge- and node-*betweenness* [17]. This measure computes, for each node (edge), the number of minimum paths connecting any two pairs of nodes in the network that pass through the considered node (edge). The value of betweenness is low when few paths traverse that portion of the network, making the node (edge) marginal. It is high when the number of passing paths is high, making the node (edge) relevant for the network.

Volume describes quantitatively the amount of mobility observed for each node (edge). In particular we exploit the *weight* function ω of each edge to quantitatively measure its volume. The *flow per location* is measured by summing the weights of incoming (outgoing) edges. Similarly, we state the relevance of a location by considering the value of the function τ representing the time spent in that location.

2) Trip features: In the literature, the features used for the learning phase vary according to the application scenario. However, we can identify a subset of features that are shared by many approaches [12], [9], [19]. We start from this set and select the following features: length, duration, time interval (four distinct daily time intervals: 00-06, 06-12, 12-18, 18-00), and average speed.

It is worth noting that each family of features combines both topological properties of the graph and mobility specific measures. Moreover, each feature can be computed directly from the IMN without referring to the original trajectories used to learn the model. Table I summarizes the features used for the learning phase. From each IMN we derive a training set where each row describes a single trip of the individual: some of the attributes of the row depend on the specific trip; the remaining features depend on the topology of the edge. Network features are the same for all the trips belonging to the same network edge. Since each edge refers to a pair of locations, the features of the nodes are computed for the outgoing node and the ingoing node and stored in distinct attributes of the table, using the suffix *-from* or *-to* to distinguish between them.

TABLE I. FEATURES SUMMARIZING THE MAIN ASPECTS OF INDIVIDUAL MOBILITY.

trip features	**network features**	
length duration time interval average speed	*centrality*	clustering coefficient average path length
	predictability	entropy
	hubbiness	degree betweenness
	volume	edge weight flow per location

TABLE II. SOME EXAMPLES OF FEATURES EXTRACTED FROM THE IMN IN FIGURE 1.

From	**To**	**ccFrom**	**ccTo**	**weight**	...	**duration**
a	b	0.12	0.22	20	...	11min
a	b	0.12	0.22	20	...	8min
...	...	...	...	...	...	...
b	n	0.22	0.0	2	...	3min
b	n	0.22	0.0	2	...	2min
...	...	...	...	...	...	...

Table II displays some examples extracted from the IMN of Figure 1. Given an edge, the number of examples equals the value of the edge's weight. Edge $e = (a, b)$, for instance, has weight $\omega(a, b) = 20$ and is represented by 20 rows in Table II. The trip features (such as length and duration) are inferred from the distributions related to nodes and edges.

B. Activity-Based Cascade Classification

We propose a meta-classification process based a step-wise approach, which has some similarities with nested cascade classification [20]. For this reason we call it *Activity-Based Cascade (ABC)* classification. The idea is to reduce the original multi-class classification problem to several binary classification problems, each trying to discriminate between one class and all the others. The classifiers are learnt in cascade according to their ascending frequency, as for rule-based classifiers [21]. At each step of the cascade chain, the classifier tries to distinguish the instances of its class among those ones rejected in the previous steps. For instance, if "going home" is the first class in the cascade, the first binary classifier will discriminate between "going home" and "all other classes". Then, the second classifier will discriminate for another class *without* the instances of "going home". This elimination simplifies the job of discriminating the instances of the second type, belonging to more difficult classes.

We can further improve the method by using the classification results obtained at each step as source for later classifications. Indeed, the first classifier of the cascade allows to recognize all the trips/locations of its class (for instance, all the trips of type "going home"). We can use the information from the previous steps as contextual information by computing some measures for all other trips/locations. For instance, we can compute the number of network hops between any trip/location and its closest "going home" trip/location. The result is that the number of attributes describing the instances grows at each step of the cascading process, and the later classes will benefit from a larger set of context information.

Clearly, the same step-wise process is adopted when we apply the classifier to a test set of unlabelled IMNs, in order to perform the semantic amplification. The first step of the cascade is applied to all the trips of the IMNs to classify those ones belonging to the first class (in our example above, "going home"). Then, all the remaining trips are enriched with a context feature, e.g. their network distance from the closest trip labeled as "going home". The process continues with the next step in the cascade. At the end of the cascade process, if some instances of the test set are classified as "all other classes", they remain *unclassified*, meaning that the overall classifier was not able to properly classify it.

1) ABC Evaluation Criteria: To evaluate the quality of each classifier, we split the annotated dataset into a training set (used for the learning phase) and a test set (used for measuring the perfomances). Following the approach presented in [12], the split is done according to two strategies: leave-one-week-out (WO) and leave-one-user-out (UO). The first strategy extracts the trajectories of one week of each individual, using them as test set. The second strategy extracts a single individual from the dataset, using her whole mobility as test set. The two quality measures drive the classification models towards two distinct goals. The WO strategy emphasizes what is learned from the past to predict the future. This approach is usually suited for recommendation-like services, where the previous examples are used to classify the new ones. However, it has a tendency to produce overfitted models that are hardly adaptable to new individuals. The UO strategy, instead, learns from the crowd to predict the activities of new individuals. It is more robust to overfitting since it generalizes more the learned model. In our experiments we used a generalization of this two approaches by splitting the annotated dataset into two parts corresponding to 60% and 40% of total size. Since the available dataset covers a period of one month, we generalize the WO strategy in a *Temporal Split* (TS) where three weeks are used for training and the last week is used as test set. The corresponding splitting strategy for UO is called *User Split* (US), where the 60% of the individuals compose the training set, and the remaining 40% compose the test set.

V. EXPERIMENTS AND EVALUATIONS

In our experiments we used a large GPS dataset of about 150k vehicles moving in Tuscany during May 2011. A portion of the movements have been annotated by volunteers to reconstruct their activities during May 2011. In particular, volunteers annotated 6,953 distinct trips performed by 65 vehicles. The annotation was performed by using 13 distinct activities: going home, working, daily shopping, shopping, social activities, leisure, services, education and training, bring and get, touring, other. We introduced the activity type *none* for the activities that did not receive any tagging. Such *none* activities, which correspond to about the 5% of the training set, have been removed before the learning phase. Figure 2 shows the distribution of activities for all the movements in the training set. Clearly, "going home" and "working" activities have the highest frequencies in the dataset: this highlights the fact that the majority of movements are performed for systematic routines. A large part of the activities, however, is annotated with the "other" label generally due to the presence of activities that are difficult to be categorized by the volunteer. For example, the "bring and get" activity is rarely annotated and often avoided, since it includes short stops. We computed the IMN of each individual in the dataset through the incremental procedure described in Section III, using $\delta = 200$m as spatial threshold for computing the network's locations.

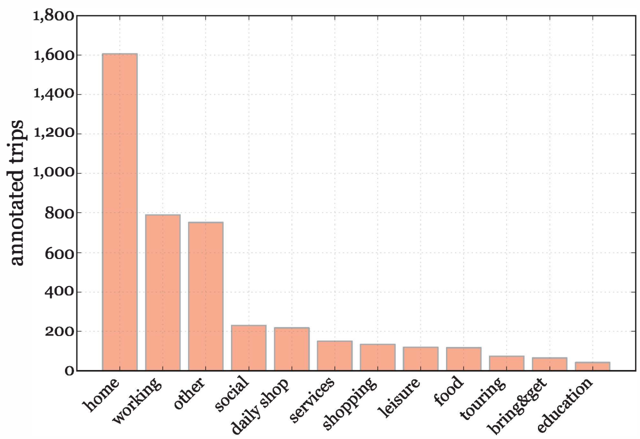

Fig. 2. Distribution of activities performed by the users in the training set.

A. Competing classification approaches

The task of classifying individual activities by movements is usually approached with two main strategies. One group of methods [12], [9], based mainly on probabilistic approaches, consider the problem of tagging a whole sequence of activities of an individual. They start from the assumption that an individual performs similar schedules in different days. The second group of methods, based on decision tree and Hidden Markov Models, try to tag each single episode. Since a IMN hides the specific sequence of movements performed by the individual, we cannot use directly sequence-based learning methods. However, we will show that the topological features of a IMN are capable of subsuming the probabilistic dependencies among the different activities of an individual. To this aim, we compare the performance of the ABC classifier introduced in Section IV-B with two methods which represents the two families of approaches described above: Conditional Random Fields (CRF) and Random Forest (RF).

a) Conditional Random Fields (CRF): Conditional Random Fields (CRF) [5] are probabilistic graphical models to label sequence data. They are based on undirected graphical structures, where each link represents the conditional probability distribution over hidden states. Unlike Hidden Markov Models, there is no assumption of dependency in the structure. In [12] the authors show an instance of CRF for the activity recognition problem where the nodes represent sequences of observations. A trajectory is associated with a node and all its properties are attached to the corresponding vertex, while hidden states represent the activities to be predicted. We used the network structure as presented in [12] with the implementation of the library `CRF++`[1]. Since the method depends on the sequential order of the movement episodes, we implemented this approach by considering directly the original trajectories of the vehicles. The original article considers only trip features. We extended the implementation by including also the features derived from the corresponding IMN for each set of trajectories of the corresponding vehicle. The efficiency of the learning phase is determined by the complexity of the model and, hence, by the number of features considered.

b) Random Forest: The learning methods based on decision trees provide an efficient solution for classification problems. However, the learning phase and the classification step are usually limited to a single episode, with limited knowledge of contextual episodes (i.e. previous or successive movements). To overcome this problem, we extend the set of features for each movement to classify by coding the contextual information in a set of numeric attributes. This approach simplifies the search space compared with that structured in a CRF and focuses only on relevant properties of each movement. The procedure to extract these features from a IMN is described in Section IV-A. Given the large set of features and the multi-class scenario, we adopted a Random Forest classifier, as provided by the `scikit-learn` Python package.[2]

B. Accuracy evaluation

How to compare the performances of the two learning approaches presented in Section IV-B? On one hand, the CRF classifier uses the raw trajectories of the training set. On the other hand, the Random Forest classifier is based on the IMNs extracted from the raw trajectories. The set of classes to predict, conversely, is fixed for both the methods. We compare the two approaches by means of two distinct accuracy methods, the Time Split and the User split methods introduced in Section IV-B1; using only trip features and both trip and network features.

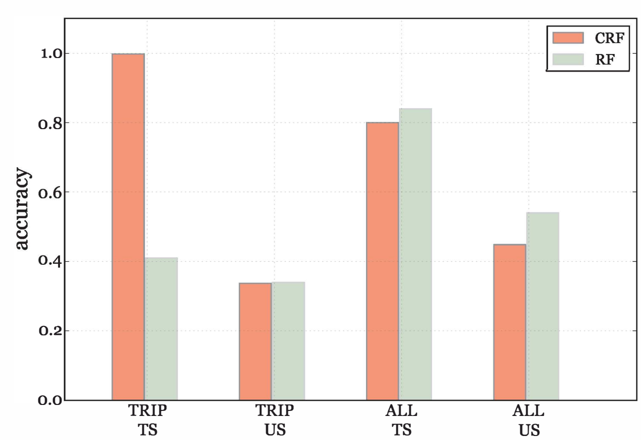

Fig. 3. Comparison between the accuracy of Conditional Random Fields (CRF) and the accuracy of Random Forest classifier (RF), with Time Split (TS) and User Split (US), and two sets of features, trip features and trip and network features.

Figure 3 shows the accuracy of the two methods. Both methods have better performances when using the Time Split (TS) accuracy: they are more robust when we learn the model from the past and apply it to the future. Using the TS accuracy metric, the CRF classifier outperforms the Random Forest, although the extended set of features (the ones including the

[1] https://code.google.com/p/crfpp/

[2] http://scikit-learn.org/

network features) negatively influences the CRF results. In the RF, on the contrary, the extended network features are effective and contribute to increase the accuracy of the classification. The Time Split accuracy measure, however, has a marginal significance for our problem. Our goal is to perform semantic amplification: extend the annotation performed on a small portion of the population to a larger set of moving individuals. Therefore, in our context the User Split (US) accuracy measure is the crucial scenario. Moving to the US accuracy, both models produce a significant drop in the accuracy. In particular the CRF method shows a significant difference w.r.t. the TS scenario, suggesting a low efficacy in generalizing the model to new behaviors. The RF method suffers from an analogous drop in accuracy, but reach better overall performance than CRF on the extended set of features.

The experiments show that the proposed extended feature set are effective for both methods. They improve prediction accuracy for both the TS and the US accuracy measures. When considering the US accuracy, however, their contribution is still more valuable, since they are capable of representing general attributes of individual mobility. Thus, the proposed RF approach is more effective from two points of view: (i) we get a very compact and aggregated view over individual mobility; (ii) we are capable to better generalize mobility features to predict individuals' activities.

C. Results

In this section we compare the results of the RF classifier and the ABC classifier introduced in Section IV-B, using the US accuracy measure. As noted in Section IV-B, the ABC method might leave some output instances unclassified. In this context the accuracy (which would simply consider unclassified instances as errors) would not be a sufficiently informative measure. For this reason, we adopt the standard precision and recall measures, in order to better highlight the trade-off between quality of the classification versus completeness. Table III shows, for each class, the precision and the recall achieved by the two classifiers. Results for RF shows that some classes, e.g. "education", "bring/get" and "touring", have low precision and recall, presumably due the low frequency of such classes in the annotated dataset (see Figure 2).

The Activity-Based Cascade (ABC) classification presents several advantages with respect to the RF. First, it mitigates the differences in the support of some activities, such as the "bring and get" activity. Moreover, it produces specialized decision tree models, which help in the understanding of each activity. The ABC classification is presented with two different modes: (i) a step-wise cascading classification; (ii) a propagation of the classes assigned in the previous steps to the current one. In both cases, the accuracy of the ABC approach outperforms the RF approach presented above and provides better results (around 10% of increment of accuracy) with label propagation. The comparative results are shown in Table III. The test set used to evaluate the performance contains $2,594$ tuples. The table reports the support for each activity, i.e. the number of samples for which the ABC classifier succeeds in assigning a label. It is clear that the ABC classifier is not able to classify all the instances in the test set. In the last column of Table III we report the number of tuples labelled by the classifier. At the end of the cascading classification, the samples that are not classified by the last decision tree remain unclassified. This is due to the selectiveness of the decision tree in the cascading chain, which is emphasized in the label propagation mode, since the additional features contribute to better distinguish the current class from the remaining ones. As a consequence, ABC classifies a reduced number of examples, but with a higher confidence.

ABC classification uses a fixed order of the classes in the cascading chain: in our experiments we sorted in ascending order the classes according to the their support, following the strategy in [21]. To check the significance of such sorting with respect to the outcome of the classification, we performed a series of classification processes by considering different shuffling of the class order (Figure 4). It is clear that the ordering have little and statistically non-significant influence on the accuracy.

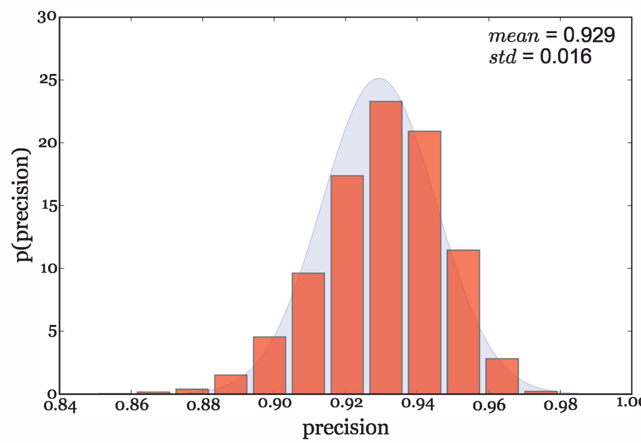

Fig. 4. Distribution of precision for 10k executions of the ABC classifier, varying class order for progressive classification.

VI. Conclusions

In this paper we described a methodology to recognize individuals' activities from the observation of their movements. The approach is based on the definition of an Individual Mobility Network (IMN), a graph-based model capable of grasping the salient properties of individual mobility. The set of annotated movements are enriched with features extracted from the IMN, and used to train a classifier for mapping mobility properties to semantic activities. The analytical process we proposed, Activity-Based Cascading (ABC) classification, departs completely from previous approaches and outperforms the state-of-the-art probabilistic methods. The contribution of each single feature extracted from the IMN can be helpful to further improve the quality of the classification, and it is left for future work.

Since the features of IMNs are abstracted from real geography, we can exploit this methodology to transfer the model to a different territory. From the experiments we have evidence that the approach is valid at a regional level, since the outcomes of the annotation within the same region provide reasonable results. As a preliminary proof, we focused on two cities in the region of Tuscany, restricting the training set to Pisa and the test set to Florence. The results are promising: high frequency activities (going home and working) show high precision

TABLE III. PRECISION AND RECALL OF THE RF CLASSIFIER (LEFT) AND THE ABC CLASSIFIER (RIGHT) WITH EXTENDED FEATURES, LABEL PROPAGATION, AND USER SPLIT (US) VALIDATION.

RF classifier			activity	ABC classifier			
precision	recall	f1-score		precision	recall	f1-score	support
0.91	0.93	0.92	Going Home	0.99	0.99	0.99	930
0.00	0.00	0.00	Bring and get	1.00	0.68	0.81	22
0.00	0.00	0.00	Education	1.00	0.25	0.40	4
0.15	0.14	0.14	Daily shopping	0.97	0.85	0.91	68
0.50	0.37	0.43	Working	0.93	0.96	0.94	258
0.34	0.6	0.43	Other	0.90	0.98	0.94	384
0.09	0.05	0.06	Shopping	0.87	0.85	0.86	39
0.13	0.09	0.11	Leisure	0.84	0.84	0.84	51
0.24	0.11	0.15	Services	0.83	0.71	0.77	42
0.04	0.01	0.02	Touring	0.77	0.83	0.80	12
0.06	0.06	0.06	Food	0.76	0.49	0.60	53
0.08	0.04	0.05	Social activities	0.65	0.69	0.67	49
0,54	**0,54**	**0,54**	avg / total	**0.94**	**0.94**	**0.94**	1912

(respectively 90% and 80%). Activities with lower support, instead, have lower accuracy because of the low number of samples.

ACKNOWLEDGMENT

This work has been partially funded by the European Union under the FP7-ICT Program: Project DataSim n. FP7-ICT-270833, and Project PETRA n. 609042 (FP7-SMARTCITIES-2013); and by the MIUR and MISE under the Industria 2015 program: Project MOTUS grating degree n. 0000089 - application code MS01_00015.

REFERENCES

[1] M. C. González, C. A. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 7196, pp. 779–782, Jun. 2008.

[2] L. Pappalardo, S. Rinzivillo, Z. Qu, D. Pedreschi, and F. Giannotti, "Understanding the patterns of car travel," *The European Physical Journal Special Topics*, vol. 215, no. 1, pp. 61–73, 2013.

[3] M. Batty, K. Axhausen, F. Giannotti, A. Pozdnoukhov, A. Bazzani, M. Wachowicz, G. Ouzounis, and Y. Portugali, "Smart cities of the future," *The European Physical Journal Special Topics*, vol. 214, no. 1, pp. 481–518, 2012.

[4] L. Liao, D. J. Patterson, D. Fox, and H. Kautz, "Learning and inferring transportation routines," *Artif. Intell.*, vol. 171, no. 5-6, pp. 311–331, Apr. 2007.

[5] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proceedings of the Eighteenth International Conference on Machine Learning*, ser. ICML 2001. Morgan Kaufmann Publishers Inc., 2001, pp. 282–289.

[6] A. M. Ladd, K. E. Bekris, A. Rudys, G. Marceau, L. E. Kavraki, and D. S. Wallach, "Robotics-based location sensing using wireless ethernet," in *Proceedings of the 8th annual international conference on Mobile computing and networking*, ser. MobiCom '02. New York, NY, USA: ACM, 2002, pp. 227–238.

[7] P. Bahl and V. N. Padmanabhan, "Radar: An in-building rf-based user location and tracking system," in *INFOCOM*, 2000, pp. 775–784.

[8] Y. Zheng, Q. Li, Y. Chen, X. Xie, and W.-Y. Ma, "Understanding mobility based on gps data," in *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 2008, pp. 312–321.

[9] S. Reddy, M. Mun, J. Burke, D. Estrin, M. Hansen, and M. Srivastava, "Using mobile phones to determine transportation modes," *ACM Trans. Sen. Netw.*, vol. 6, no. 2, pp. 13:1–13:27, February 2010.

[10] M.-J. Lee and C.-W. Chung, "A user similarity calculation based on the location for social network services," in *Proceedings of the 16th international conference on Database systems for advanced applications - Volume Part I*, ser. DASFAA'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 38–52.

[11] V. W. Zheng, Y. Zheng, X. Xie, and Q. Yang, "Collaborative location and activity recommendations with gps history data," in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 1029–1038.

[12] L. Liao, D. Fox, and H. A. Kautz, "Extracting places and activities from gps traces using hierarchical conditional random fields," *The International Journal of Robotics Research*, vol. 26, no. 1, pp. 119–134, 2007.

[13] S. Jiang, J. Ferreira, and M. C. González, "Clustering daily patterns of human activities in the city," *Data Mining and Knowledge Discovery*, pp. 478–510, 2012.

[14] N. Eagle and A. (Sandy) Pentland, "Reality mining: sensing complex social systems," *Personal Ubiquitous Computing*, vol. 10, no. 4, pp. 255–268, Mar. 2006.

[15] L. Ferrari and M. Mamei, "Classification and prediction of whereabouts patterns from the reality mining dataset," *Pervasive and Mobile Computing*, vol. 9, no. 4, pp. 516 – 527, 2013.

[16] K. Farrahi and D. Gatica-Perez, "Discovering routines from large-scale human locations using probabilistic topic models," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 1, pp. 3:1–3:27, Jan. 2011.

[17] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, no. 2, pp. 167–256, 2003.

[18] J. K. Koerner, "Coding of an information source having ambiguous alphabet and the entropy of graphs," in *Transactions of of the 6th Prague Conference on Information Theory, etc., Academia, Prague*, 1973, pp. 411–425.

[19] T. Huynh and B. Schiele, "Analyzing features for activity recognition," in *Proceedings of the 2005 joint conference on Smart objects and ambient intelligence: innovative context-aware services: usages and technologies*, ser. sOc-EUSAI '05. New York, NY, USA: ACM, 2005, pp. 159–163.

[20] R. Verschae, J. Ruiz-del Solar, and M. Correa, "A unified learning framework for object detection and classification using nested cascades of boosted classifiers," *Machine Vision and Applications*, vol. 19, no. 2, pp. 85–103, 2008.

[21] W. W. Cohen, "Fast effective rule induction," in *In Proceedings of the Twelfth International Conference on Machine Learning*. Morgan Kaufmann, 1995, pp. 115–123.