

INFERENCE OF MOBILITY PATTERNS VIA SPECTRAL GRAPH WAVELETS

Xiaowen Dong[†], Antonio Ortega[‡], Pascal Frossard[†] and Pierre Vandergheynst[†]

[†] Signal Processing Laboratories, EPFL, Switzerland
{xiaowen.dong, pascal.frossard, pierre.vandergheynst}@epfl.ch

[‡] Ming Hsieh Department of Electrical Engineering, University of Southern California, USA
antonio.ortega@sipi.usc.edu

ABSTRACT

Modern data processing tasks frequently involve structured data, for example signals defined on the vertex set of a weighted graph. In this paper, we address the problem of inference of mobility patterns from data defined on geographical graphs based on spatially localized events. Specifically, we propose a model-based approach where we build a signal model for each of the expected mobility patterns. We then analyze the characteristics of the signal models by studying their spectral representations using wavelets defined on graphs, which enables us to build efficient classifier in the spectral domain. Experiments on data gathered from photo-taking events in Flickr show that we can efficiently infer mobility patterns using only coarse aggregated information, which is certainly interesting in terms of privacy protection.

Index Terms— Signals on graphs, spectral graph wavelets, mobility patterns, classification, Flickr

1. INTRODUCTION

Signals defined on the vertex set of graphs have recently attracted growing interests in the signal processing community. In this setting, the vertices of the graph represent entities and the edge weights reflect the pairwise relationships between them, while the data associated with the vertices contain information of interest that is related to the entities. It is often the case that various patterns are present in the signals on graphs, which are usually of great practical interests. In this paper, we focus on the problem of inference of mobility patterns from geo-tagged data based on spatially localized events. The aggregation of such events lead to signals defined on geographical graphs. Our goal is to infer, based only on such aggregated information, the mobility patterns within geographical graphs.

Since the signal that we consider resides in a structured and irregular space, namely graphs, traditional techniques cannot be applied directly. Our approach is to study the signal characteristics in the graph spectral domain. More specifically, we propose a model-based inference approach, where we design synthetic models tailored for different classes of signals in the local regions of the graph, which naturally correspond to expected mobility patterns. With the proposed models, we simulate event-based individual mobility traces, by aggregating which we generate synthetic signals on graphs. We then study their characteristics in the graph spectral domain via the Spectral Graph Wavelet Transform [1]. This allows us to learn efficient classifiers in the spectral domain, and infer mobility by matching local data features.

This work has been partly supported by the EDGAR project funded by Hasler Foundation, Switzerland.

We show an example based on data gathered from photo-taking events in Flickr. We specifically consider, the aggregated number of distinct Flickr users who have taken geo-tagged photos at different locations in the central area of big cities, as the signal defined on a geographical graph. Results show that our framework can efficiently infer the mobility patterns of Flickr users within these big cities.

The present paper is mainly related to two branches of work in the literature. First, graph-based methods have become popular in classification tasks. Example works include [2, 3] where the authors studied semi-supervised learning problems on graphs for classification. Second, a few recent works focused on mining Flickr data with application to inference of user mobility patterns based on empirical individual traces [4, 5, 6, 7, 8, 9]. In addition, our approach is also analogue to the texture classification approach in [10] in the sense that both perform classification with features computed from wavelet transforms of signals.

The contribution of our paper is two-fold: First, we propose a novel model-based inference approach for signals on graphs via its spectral representation, which is different from most of the traditional graph-based classification approaches like label propagation. Second, we achieve detailed and meaningful inference of mobility patterns in geographical graphs from Flickr data without looking into empirical traces, which better preserves user privacy.

2. FRAMEWORK

2.1. Spectral representations of signals on graphs

Efficient representation of signals on graphs have recently attracted increasing interest in the signal processing community. Among several works that have proposed wavelet transforms on graphs, the recently developed Spectral Graph Wavelet Transform (SGWT) [1] is of particular interest in this paper. The SGWT carries over many properties of the classical wavelet transform, especially the ability of decomposing the signals into different subbands according to the spectral characteristics of the signals on the graph. We first give a very brief introduction of the construction of SGWT.

Consider an undirected and weighted graph $\mathcal{G}(V, E, w)$, where V represents the set of N vertices, E represents the set of edges, and $w : E \rightarrow \mathbb{R}^+$ associates each edge with a positive weight. Denote L as the unnormalized Laplacian matrix of the graph. Notice that L is real and symmetric, therefore it has a complete set of orthonormal eigenvectors: $L\chi_l = \lambda_l\chi_l, l = 0, \dots, N-1$. For a bandpass filter g that satisfies the admissibility conditions, the spectral graph wavelets defined at scale $\{t_j\}_{j=2}^J$ can be written as: $\Psi_{t_j} = \chi g(t_j\Lambda)\chi^*$, where Λ is the diagonal matrix that contains the eigenvalues along diagonal, and J is the total number of scales. For a discrete set

of scales $\{t_j\}_{j=1}^J$ ¹, the spectral graph wavelets actually constitute a frame in \mathbb{R}^N . The SGWT is then computed by taking the inner product between the signal and every atom in the wavelet frame. By construction, the SGWT decomposes a signal into J subbands that are computed from the (discretized) spectrum of the graph, where each subband represents a certain frequency in the graph spectrum domain. Therefore, wavelet coefficients in different subbands can be considered as spectral representations of the signal.

2.2. Model-based inference of mobility patterns on geographical graphs

The signal that we consider is closely related to the geographical graph and spatially localized events. Since we can expect different mobility patterns in the context of application at hand, we propose to create synthetic signal models that correspond to these patterns, and generate synthetic signals by simulations with the models on synthetic geographical graphs. We then analyze the characteristics of the resulting signals in the spectral domain. This leads to different spectral signatures for signals from different models, based on which we can build efficient classifier in the spectral domain.

Differences in spectral signatures of different signal models mainly lie in the shape of the spectrum in different subbands, namely the magnitude of the wavelet coefficients and how they vary from one vertex to another. This allows us to build feature vectors from spectral characteristics. In this paper, we build spectral graph wavelets with three scales, which correspond to low, middle and high frequency subbands, respectively. For each training signal, we then propose a 6-dimensional spectral feature vector as follows:

$$f(i) = \begin{cases} \max(C_{\text{low}})/\text{maxcoeff}, & i = 1, \\ \{\max(C_{\text{low}}) - \min(C_{\text{low}})\}/\text{maxcoeff}, & i = 2, \\ \max(C_{\text{mid}})/\text{maxcoeff}, & i = 3, \\ \{\max(C_{\text{mid}}) - \min(C_{\text{mid}})\}/\text{maxcoeff}, & i = 4, \\ \max(C_{\text{hig}})/\text{maxcoeff}, & i = 5, \\ \{\max(C_{\text{hig}}) - \min(C_{\text{hig}})\}/\text{maxcoeff}, & i = 6, \end{cases} \quad (1)$$

where C_{low} , C_{mid} , C_{hig} denote the magnitudes of wavelet coefficients in low, middle and high frequency subbands, respectively, maxcoeff denotes the maximum magnitude of wavelet coefficients in all subbands, and $\max(\cdot)$ and $\min(\cdot)$ denote the operators taking the maximum and minimum values.

We can then train a classifier on the synthetic signals. We apply this classifier to testing data to infer model-based patterns, where the testing data is segmented into local graph patches before classification. The classification algorithm is summarized in Algorithm 1.

Within the context of photo-taking events in Flickr, we describe in detail the construction of models and the classification of mobility patterns in the next Section.

3. MODELS AND INFERENCE OF MOBILITY PATTERNS

3.1. Mobility patterns and the models

The four mobility patterns are illustrated in Fig. 1. Notice that the signal that we consider is a mixture of two sources of information, namely user mobility in Flickr and the photo-taking events. Therefore, the proposed patterns include models for both mobility and photo-taking. We describe the models as follows².

¹In the construction of SGWT, the scale t_1 corresponds to a low pass filter h which covers low frequencies in the spectrum.

²The parameters are fixed in Section 3.1 for simplicity. In the experimental results in Section 4 we allow the parameters to vary.

Algorithm 1 Model-based inference of mobility patterns

1: Input:

\mathcal{G} : A weighted and undirected graph of N vertices

F : A signal defined on the vertex set of \mathcal{G}

2: For each vertex j in \mathcal{G} , generate a subgraph \mathcal{G}_j that contains the k -hop neighbors of j . Generate the truncated signal F_j on \mathcal{G}_j . Apply SGWT to F_j and compute the wavelet coefficients in three subbands.

3: Compute the testing features $\{f_j\}_{j=1}^N$ using Eq. (1).

4: Apply the trained classifier to $\{f_j\}_{j=1}^N$ to classify the k -hop neighborhood of vertex j into one of the expected patterns.

5: Output: Predicted classes of patterns for the neighborhoods of each vertex in \mathcal{G} .

3.1.1. Models for user mobility

We consider a synthetic geographical area, namely, an eight nearest neighbor grid of size 10 by 10 as shown in Fig. 2(a). The parameters in the models for user mobility include (i) the starting point, (ii) the direction for the next move, and (iii) the probability of staying at the current position (denoted by p_{stay}). Notice that each step is denoted by an integer in the set of positions $\text{move}_{\text{pos}} = \{1, 2, 3, 4, 5, 6, 7, 8\}$, which represents the eight possible moving directions on the grid as shown in Fig. 2(b). Furthermore, for simplicity we set $p_{\text{stay}} = 0.1$ in all the models.

The first mobility pattern is a “random” pattern for areas where there is no particular point of interest, and people just follow naturally their paths without specific targets. The user mobility is then modeled by a synthetic trajectory, where a user first moves from a boundary point to the interior of the grid, and the t -th move is computed as follows:

$$\text{move}_t = \text{move}_{t-1} + \lfloor m \rfloor + 8n, \quad (2)$$

where $\lfloor m \rfloor$ is the largest integer no greater than a normally distributed random number m , and n is an off-set parameter to ensure that $\text{move}_t \in \text{move}_{\text{pos}}$.

The second mobility pattern is called “spreading”, where people emerge from one certain position in the area (usually public transport stations), spread to the points of interest nearby, and then leave. The user mobility is modeled by a synthetic trajectory where a user starts at the certain position and makes the first move towards the points of interest; the t -th move is computed as follows:

$$\text{move}_t = \text{move}_1 + \lfloor m \rfloor + 8n, \quad (3)$$

where move_1 is the first move, and m and n are defined as before. This enforces straighter trajectories since in this case people move with more clear purposes.

The third pattern is a “straight” pattern for areas in which there is a point of particular interest (denoted by $\text{POI}_{\text{straight}}$), for example a landmark. People enter the area with clear purposes, move straight onto $\text{POI}_{\text{straight}}$, and then leave. We model the user mobility based on the distance between the current position of the user and $\text{POI}_{\text{straight}}$. Specifically, a user starts at a boundary point with the t -th move computed as:

$$\text{move}_t = \begin{cases} x \in \text{move}_{\text{pos}} \text{ such that } d_t < d_{t-1}, \\ \quad \text{before reaching } \text{POI}_{\text{straight}}, \\ x \in \text{move}_{\text{pos}} \text{ such that } d_t > d_{t-1}, \\ \quad \text{after reaching } \text{POI}_{\text{straight}}, \end{cases} \quad (4)$$

where d_{t-1} and d_t are the distances from the current and the next position to $\text{POI}_{\text{straight}}$, respectively.




Fig. 1. Four mobility patterns. The red curve simulates user movement and the blue circle represents a photo-taking event.




Fig. 2. (a) An eight nearest neighbor grid of size 10 by 10. (b) Eight possible directions to move on an eight nearest neighbor grid.

The fourth pattern is named “bi-directional” suitable for areas like streets and rivers. In these places people usually move with clear directionality and seldom make a detour. We assume that people move only within a “bi-directional” area, for example a grid of size 10 by 3 as the red area shown in Fig. 2(a). A user starts from a random point within the left boundary, and stop after reaching the right boundary. In this specific case, at each step, the t -th move is computed as:

$$\text{move}_t = x \in \{6, 7, 8\}. \quad (5)$$

3.1.2. Models for photo-taking events

In the models for photo-taking events, we consider the probability of photo-taking actions at different positions (denoted by p_{photo}), which is chosen in accordance with the models for user mobility that we have proposed before. In the “random” pattern, since we do not assume any particular point of interest, we set $p_{\text{photo}} = 0.1$. For “spreading”, it is usually the case that points of interest lie in certain directions with some distance from the position people emerge, and people start taking photos after reaching these points of interest. Therefore, we define p_{photo} as follows:

$$p_{\text{photo}} = \begin{cases} 0.3, & s \geq 2 \text{ and } s \leq 4, \\ 0.1, & s > 4, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where s is the number of steps the user has made. Next, in the “straight” pattern, people go straight to $\text{POI}_{\text{straight}}$ to take photos. We set p_{photo} higher at and around $\text{POI}_{\text{straight}}$:

$$p_{\text{photo}} = \begin{cases} 0.5, & \text{at } \text{POI}_{\text{straight}}, \\ 0.3, & \text{at the neighbors of } \text{POI}_{\text{straight}}, \\ 0.1, & \text{otherwise.} \end{cases} \quad (7)$$

Finally, for “bi-directional”, people tend to take photos at certain fixed points of interest (denoted by $\text{POI}_{\text{bi-directional}}$) within the “bi-directional” area. Therefore, we set p_{photo} higher at $\text{POI}_{\text{bi-directional}}$:

$$p_{\text{photo}} = \begin{cases} 0.3, & \text{at } \text{POI}_{\text{bi-directional}}, \\ 0.1, & \text{otherwise.} \end{cases} \quad (8)$$

The models for user mobility and models for photo-taking events together constitute the proposed mobility patterns.

3.2. Classification of mobility patterns based on spectral signatures

Based on the mobility patterns, we have simulated numerous synthetic runs of individual mobility traces with photo-taking events. By aggregating such individual traces, we get the aggregated number of distinct users that have taken photos at different positions in the grid, which is viewed as a signal defined on the vertex set of the grid graph with unit edge weights. The signal is then decomposed into three spectral subbands using the three-scale spectral graph wavelets. As we can imagine, for different patterns, the spectrum in different subbands behave in different ways. Specifically, in “random” we expect a flat spectrum and in “spreading” we expect a smoothly decaying spectrum near the position people emerge. The spectrum of “straight” shows a sharp peak at $\text{POI}_{\text{straight}}$, while the one for “bi-directional” is a combination of flat spectrum and less sharp peaks at $\text{POI}_{\text{bi-directional}}$. All these signatures are represented by the spectral feature vectors computed in Eq. (1). We then train a classifier with synthetically generated training signals using the LIBSVM library [11], and apply it to testing signal from Flickr data using Algorithm 1 (with the choice of $k = 3$).

4. EXPERIMENTAL RESULTS

4.1. Results on synthetic data

In this experiment, we compare two classifiers, one is trained using the proposed spectral features, the other using features computed from the spatial variations of the signals. More specifically, for each synthetic signal and graph, we look at the ratios of the mean and maximum differences in signal values along all the edges to the maximum difference in signal values (not necessarily along edges) to build a 2-dimensional feature vector. Both classifiers are trained using the LIBSVM library with default parameters. Due to space limit, we leave the discussions on dimensions of feature vectors and parameter selection in LIBSVM for the long version of the paper.

We generate 4000 synthetic signals on the grid graph, each of which is obtained by aggregating one hundred synthetic runs of mobility traces with photo-taking events. We also allow the parameters in the mobility patterns, for example p_{photo} and the locations of the special points in the models, to vary within a certain range. Due to randomness in the synthetic signals introduced by simulations and flexible parameter selections, the trained classifiers take into account the effect of noise, and the classification performance does not depend too much on specific choices of parameters. The 4000 synthetic signals are divided into two halves for training and testing, both of which contain 500 synthetic signals for each mobility pattern. We measure quantitatively the classification accuracy by applying the trained classifiers to the testing signals and the results are shown in Table 1. As we can see, although the classifier trained using spatial features performs slightly better in the “random” pattern, the proposed spectral-feature based classifier achieves clearly higher accuracies in the “spreading” and “bi-directional” patterns, and leads to a better overall performance. Notice that the “random” pattern is the one most likely to be affected by the randomness contained in the synthetic signals, and the resulting spectrum might be misleading.

4.2. Results on real world data

We now apply the two trained classifiers to the data that we have collected from Flickr. Specifically, we consider the aggregated number of distinct Flickr users that have taken photos at different geographical locations in central London between January 2010 and June 2012

Table 1. Classification accuracy of the two classifiers.

	accuracy in “random”	accuracy in “spreading”	accuracy in “straight”	accuracy in “bi-directional”	overall accuracy
Classifier based on spectral features	80.6%	89.6%	100%	85.4%	88.9%
Classifier based on spatial features	82.6%	81%	100%	65%	82.2%




Fig. 3. Aggregated number of distinct Flickr users that have taken photos at different locations in central London between January 2010 and June 2012. The size and color of the balls indicate the number of distinct users associated with these geographical locations.

as our signal. The original locations are tiled such that each resulting location represents an area of 10 meters by 10 meters. The graph is computed by assigning an edge between two locations when the distance between them is shorter than around 30 meters, and the edge weight is set to be inversely proportional to the distance. The signal and graph are shown in Fig. 3.

The class labels obtained by applying the two classifiers are shown in Fig. 4(a) and Fig. 4(b), where label 1 (blue) represents vertices that have less than ten 3-hop neighbors, whose neighborhood is not of significant interest. Label 2 (green), 3 (yellow), 4 (orange) and 5 (red) represent vertices whose neighborhoods are classified as “random”, “spreading”, “straight” and “bi-directional” patterns respectively. As we can see in Fig. 4(a), the red vertices are mostly streets or bridges over the river, which corresponds well to the assumptions we made in the “bi-directional” pattern. The orange vertices clearly represent landmarks or locations of significant interest, whose neighborhoods fit well in the “straight” pattern. The yellow vertices stand for regions similar to the one in “spreading”, where we expect movement of users from certain positions to the points of interest nearby. Therefore, they are usually located in the surrounding areas of the orange vertices. Furthermore, it is also possible to infer directionality of user movement by looking at the locations of these yellow vertices. For example, in Fig. 5(a) we can probably infer that people move from the pier and Westminster tube station to Big Ben, while in Fig. 5(b) it is clear that people get out of Piccadilly Circus tube station and several bus stops and move towards the points of interest to take photos. Finally, the green vertices represent areas where the numbers of photographers are quite uniform, for example squares without particularly interesting landmarks. In contrast, as we can see in Fig. 4(b), the labels obtained by the spatial-feature based classifier are clearly not as meaningful, which shows that this classifier is not robust enough to be generalized to signals on real world graphs that are usually much more irregular than the grid graph.

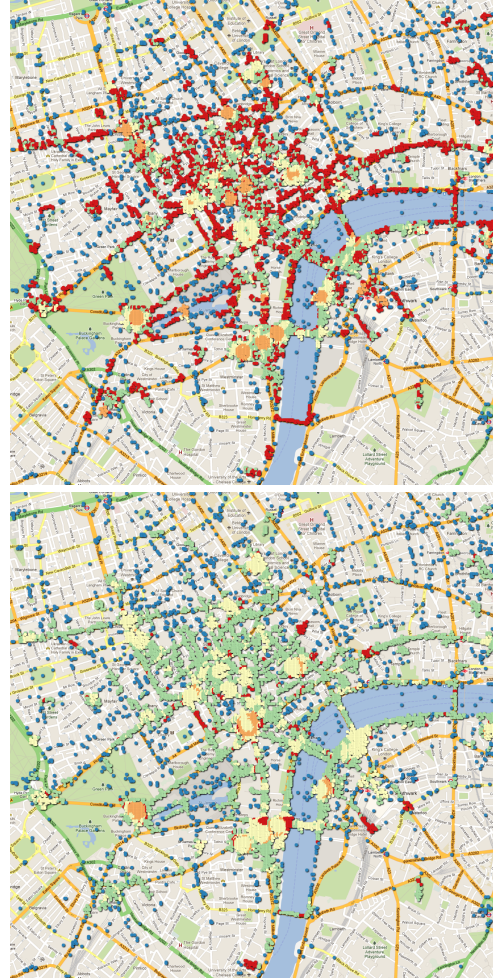


Fig. 4. Class labels of mobility patterns in London obtained by classifier trained using spectral (top) and spatial features (bottom). The colors green, yellow, orange and red represent “random”, “spreading”, “straight” and “bi-directional” patterns, respectively. Blue vertices are not of interest due to small neighborhoods.




Fig. 5. Class labels of mobility patterns near Big Ben and Westminster Bridge (left) and Piccadilly Circus tube station (right).

5. CONCLUSION

We have shown in this paper that: (i) Inference based on spectral graph models is of great potential in learning tasks with signals defined on geographical graphs, and leads to pretty robust algorithms. (ii) Only coarse aggregated information is sufficient to provide a detailed understanding of mobility patterns in geographical graphs such as the one from Flickr, which interestingly permits to understanding of people behaviors without sacrificing privacy.

6. REFERENCES

- [1] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on Graphs via Spectral Graph Theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, Mar 2011.
- [2] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf, "Learning with Local and Global Consistency," in *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- [3] X. Zhu, "Semi-Supervised Learning with Graphs," *PhD thesis, CMU-LTI-05-192, Carnegie Mellon University*, 2005.
- [4] F. Girardin, F. Dal Fiore, C. Ratti, and J. Blat, "Leveraging Explicitly Disclosed Location Information to Understand Tourist Dynamics: A Case Study," *Journal of Location Based Services*, vol. 2, no. 1, pp. 41–56, Mar 2008.
- [5] T. Kurashima, T. Iwata, G. Irie, and K. Fujimura, "Travel Route Recommendation Using Geotags in Photo Sharing Sites," in *ACM International Conference on Information and Knowledge Management*, 2010.
- [6] X. Lu, C. Wang, J.-M. Yang, Y. Pang, and L. Zhang, "Photo2Trip: Generating Travel Routes from Geo-tagged Photos for Trip Planning," in *ACM International Conference on Multimedia*, 2010.
- [7] M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu, "Automatic Construction of Travel Itineraries using Social Breadcrumbs," in *ACM Conference on Hypertext and Hypermedia*, 2010.
- [8] M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu, "Constructing Travel Itineraries from Tagged Geo-temporal Breadcrumbs," in *International Conference on World Wide Web*, 2010.
- [9] Y.-T. Zheng, Z.-J. Zha, and T.-S. Chua, "Mining Travel Patterns from Geotagged Photos," *ACM Transactions on Intelligent Systems and Technology*, vol. 3, no. 3, May 2012.
- [10] R. Porter and N. Canagarajah, "A Robust Automatic Clustering Scheme for Image Segmentation using Wavelets," *IEEE Transactions on Image Processing*, vol. 5, no. 4, pp. 662–665, Apr 1996.
- [11] C.-C. Chang and C.-J. Lin, "LIBSVM: A Library for Support Vector Machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, no. 3, Apr 2011.