Joint Learning User's Activities and Profiles from GPS Data

Vincent Wenchen Zheng
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
vincentz@cse.ust.hk

Yu Zheng
Microsoft Research Asia
4F, Sigma building, No.49
Zhichun road, Haidian District
Beijing 100190, China
yuzheng@microsoft.com

Qiang Yang
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
qyang@cse.ust.hk

ABSTRACT

As the GPS-enabled mobile devices become extensively available, we are now given a chance to better understand human behaviors from a large amount of the GPS trajectories representing the mobile users' location histories. In this paper, we aim to establish a framework, which can jointly learn the user activities (what is the user doing) and profiles (what is the user's background, such as occupation, gender, age, *etc.*) from the GPS data. We will show that, learning user activities and learning user profiles can be beneficial to each other in nature, so we try to put them together and formulate a joint learning problem under a probabilistic collaborative filtering framework. In particular, for activity recognition, we manage to extract the location semantics from the raw GPS data and use it, together with the user profile, as the input; and we will output the corresponding activities of daily living. For user profile learning, we build a mobile social network among the users by modeling their similarities with the performed activities and known user backgrounds. Compared with the other work on solely learning user activities or profiles from GPS data, our approach is advantageous by exploiting the connections between the user activities and profiles for joint learning.

Categories and Subject Descriptors

H.4.3 [**Information System Application**]: Communications Applications - Information browsers.; I.5.2 [**Pattern Recognition**]: Design Methodology - Classifier design and evaluation.

General Terms

Algorithm, Design, Experimentation

Keywords

Activity Recognition, User Profile Learning, GPS

1. INTRODUCTION

ACM LBSN '09, November 3, 2009. Seattle, WA, USA (c) 2009 ACM ISBN 978-1-60558-860-5 $10.00.

As the GPS-enabled devices, such as GPS-phones, become extensively available, people are now able to know and record his/her physical location trajectories. Such location trajectories based on raw GPS coordinates and time stamps have been shown to be attractive to the mobile users [15, 16] on their travel experience sharing and also geo-related Web communities related to outdoor sports [1, 2]. By using these GPS data, a mobile user can well manage his/her life log and also build a connection with other mobile users by setting up a mobile social based on user similarities [5]. In this paper, we aim to exploit more meaning information from the raw GPS data, including the user activities and user profiles, to better understand the mobile users' behaviors. Generally speaking, for user activities, we want to infer what the user is doing, instead of only knowing where he/she is. For example, if we know that the mobile user is near a shopping area, we may be able to infer that the user is shopping at that moment. We will define a set of common daily living activities that we can recognize from the GPS data, such as working, studying, shopping, *etc.* We can also expect some user annotations of what they are/were doing when they are actually using the GPS services. However, it is not always that the users can provide the activity annotations because they may forgot what activities they performed or just reluctant to do too many annotations. For user profiles, we are referring to user backgrounds, such as his/her occupations, age, gender, *etc.* Such user information can be provided by the mobile users, when they are using some online geo-related applications. Again, this information can be expected to be incomplete, as the users may not be reluctant to always fill all their personal information.

In this paper, we aim to establish a framework, which can jointly learn the user activities and profiles from the GPS data. Compared with previous work, which tried to conduct merely the user activity recognition from the GPS data [6] or merely the user profile learning from the RFID sensor data [8], our joint learning framework is advantageous by exploiting the connections between the user activities and user profiles. We observe that, learning user activities and learning user profiles can be beneficial to each other in nature. For example, if we have known that a mobile user is a student (*i.e.* occupation as a student) and he/she is in the university area from the GPS coordinates, we may infer that he/she is quite likely to be studying in the campus at the moment. On the contrary, if we have known that a mobile user is studying in the campus, we may infer that he/she is quite likely to be a student. Similar observations may be found in gender and age inference on shopping activities (*e.g.* young ladies are more fond of shopping than gentlemen), *etc.* Therefore, we try to put the user activity recognition and user profile learning together and formulate a joint learning problem. We put forward a probabilistic collaborative filtering framework to model such a joint learning problem. In particular, for activity recognition, we manage to extract the location semantics from the raw GPS data and use it, together with the user profile, as the input; and we will output the corresponding activities of daily living. For user profile learning, we build a mobile social network among the users by modeling their similarities with the performed activities and known user backgrounds.

2. RELATED WORK

2.1 Activity Recognition

Activity recognition aims to infer a user's behaviors from the observations such as sensor data (including GPS data), and has various applications including medical care, logistics service, security control [10], *etc.* Existing learning-based activity recognition algorithms cover HMM [9], abstract Hidden Markov Model [4], Dynamic Bayesian Networks [9], Conditional Random Fields and its variants [6], *etc.* A major drawback of these algorithms is that, most of these algorithms are usually in a supervised learning setting, which requires a lot of labeled data to train the recognition model. Therefore, there is some other work trying to exploit some common-sense knowledge from the Web [13, 14] and/or utilize the unlabeled data [12] to reduce the human labeling efforts. However, these algorithms do not exploit the user profiles as auxiliary information source to help the activity recognition. In our approach, we will conduct the activity recognition from the GPS data in a semi-supervised learning setting to reduce the user efforts, and also utilize the information from user profiles to boost the activity recognition performance.

2.2 User Profile Learning

User profile learning aims to infer the user backgrounds based on available sensor data for each user and/or mobile social network among the users. In [8], the authors established an RFID sensor network in an office building and collected the sensor data labeled with user profile attributes such as smoking or not (binary classification), for each user. After that, they used these labeled data to train a support vector machine for classifying the attributes. The drawback for this work is that, it is a supervised learning framework and requires the target users to input complete profile information. However, in practice, the user profile information is usually incomplete. Such incomplete information problem not only happens in mobile user profile learning but also in recommendation systems, which use the opinions of a community to help individuals of the community identify the content of interest over a set of items [3]. In recommender systems, a well-known technique is Collaborative Filtering (CF), which predicts the rating of an item for a user. The main idea of collaborative filtering is that, similar users will have similar ratings over similar items. In GPS data mining, there has been work in learning user similarities based on the location histories. For example, in [5], the authors proposed to take into account both the sequence property of people's movement behaviors and the hierarchy property of geographic spaces, and formulate a hierarchical-graph-based similarity measurement to measure the user similarities. However, such work neither exactly models the profile learning nor taking user activities into consideration in measuring the similarities.

3. JOINT LEARNING OF USER ACTIVITIES AND PROFILES

3.1 An Overview of the Learning Framework

In this paper, we consider the joint user activity and profile learning problem. To help readers understand our work's main idea, we give an overview of our approach in Figure 1.

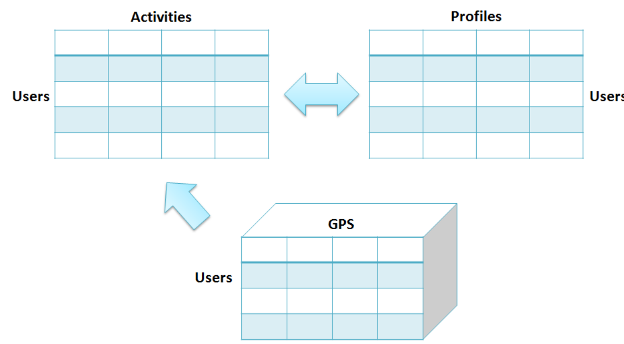

Figure 1: An overview of the learning framework.

As shown in Figure 1, in general, we take three parts of information as inputs:

- The users' GPS trajectories, as shown in the "Users-GPS" tensor (*i.e.* a 3-D matrix) of Figure 1. At the i^{th} row, and j^{th} column of the GPS trajectory tensor, we have a vector of the raw GPS coordinates at time j for user i. Since such raw GPS data may not be meaningful enough for further activity recognition and profile learning, we will apply some post-processing on this trajectory matrix. In this way, we can extract a set of interesting locations with meaningful semantics as [16]. Further, we extract the semantics features for each location for activity recognition and profile learning. We will ignore the details here, and we would note that, finally we will get a post-processed trajectory tensor, where each user i at some time slice j (the j^{th} time period instead of one time point) stays at some location represented with a semantics feature vector.

- The users' activity sequences, as shown in the "Users-Activities" matrix of Figure 1. At the i^{th} row and j^{th} column of the activity sequence matrix, we are supposed to have the activity at some time slice j for user i. We would note that, in practice, the mobile users may only provide very few activity annotations for their GPS trajectories, so the activity sequence matrix can be incomplete and even quite sparse.

- The users' profiles, as shown in the "Users-Profiles" matrix of Figure 1. At the i^{th} row and j^{th} column of the profile matrix, we are supposed to have the j^{th} attribute for user i. However, again, we would note that, in practice, the mobile users may only provide very few activity annotations for their GPS trajectories, so the activity sequence matrix can be incomplete and even quite sparse.

We aim to output the missing activities and missing entries of the user profile matrix together. To achieve this, we utilize the inherent connection between the user activities and the profiles. On one hand, we will infer what the user is doing based on his/her current location (represented as the GPS data) and his/her profile (represented as the profile attributes such as occupation, age, gender, *etc.*). Hence, in Figure 1, we can learn the user activities ("Users-Activities" matrix) from both the GPS data and the profile data ("Users-Profiles" matrix). On the other hand, we will learn the user profiles based on user similarities, which are reflected by the user activities. In other words, if two mobile users are more similar to each other, they are very likely to conduct more similar activities; and vice versa. For example, if two users are both students (*e.g.* similar in profile attributes of occupation and age), then they may share similar activities such as studying. A possible for such an activity-based profile learning strategy over other geospatial-based user similarity mining algorithms [5] is that, it can semantically understand the user similarities, other than requiring two users to be similar only when they are physically close. We encode such knowledge in Figure 1 by learning the user profiles ("Users-Profiles" matrix) from the user activities ("Users-Activities" matrix). Since the user activities and profiles are beneficial to each other in nature, we will learn them together in a unified probabilistic collaborative filtering framework as discussed in the next section.

3.2 Probabilistic Collaborative Filtering

We propose a probabilistic learning framework based on collaborative filtering to jointly learn the user activities and profiles. In particular, we apply collective matrix factorization [7, 11] to enforce the knowledge share among the user activities and the user profiles. We factorize both the (post-processed) "Users-Activities" matrix and "Users-Profiles" matrix at the same time, and require them to share a same low-dimensional user latent feature matrix. To capture the sequential information of the user activity trajectories, we design a sequential classifier taking as inputs the user location trajectories and profile information. To help readers understand the idea of our probabilistic collaborative filtering framework, we show a graphical model in Figure 2. Let us denote the "Users-Profiles" matrix as R, where R_{ij} is the value of the user i's profile attribute j. Similarly, we denote the "Users-Activities" matrix as A, where A_{ik} is the value of the user i's distribution over some activity k, representing how often the user performs that activity. Such a "Users-Activities" matrix A is extracted from the user activity sequences Y, where each Y_i is the user i's activity sequence consisting of a series of activities corresponding to the GPS trajectories. I^A is an identity matrix to convert the varying-length user activity sequences into a fixed-length user's activity distributions over a set of predefined activities. X is the (post-processed) "Users-GPS" matrix, where each X_i is a series of interesting locations extracted from the raw GPS data. Here, each interesting location is represented as a feature vector of its location semantics such as the distribution over a set of location functionalities (*e.g.* restaurants, shopping-mall).

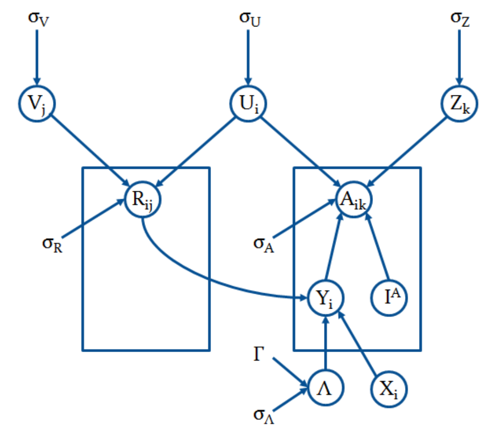

Figure 2: A graphical model for the joint learning framework.

Our probabilistic collaborative filtering framework can be depicted by two components:

- **User profile learning**. We are seeking a matrix factorization of the "Users-Profiles" matrix R as a product of two low-dimensional latent feature representations U and V. Here, U is the user feature matrix with each U_i denoted as a low-dimensional feature vector of user i. Similarly, V is the profile feature matrix with each V_j denoted as a low-dimensional feature vector of profile attribute j. Mathematically, we formulate such a matrix factorization as a conditional probability:

$$P(R|U,V,\sigma_R^2) = \prod_{i,j}\left[\mathcal{N}\left(R_{ij}|U_iV_j^T,\sigma_R^2\right)\right]^{I_{ij}^R}, \quad (1)$$

  where $\mathcal{N}(x|\mu,\sigma)$ is the probability density function of the Gaussian distribution with mean μ and variance σ, and I_{ij}^R is an indicator function that is equal to 1 if user i has provided the profile attribute j's value. As the user profile can be related to the user activities, we try to capture this relation by forcing the user feature matrix U to be shared with the "Users-Activities" matrix A's factorization. In addition, we also place the zero-mean Gaussian priors [7] on the user and profile feature matrices:

$$\begin{aligned} P(U|\sigma_U^2) &= \prod_i \mathcal{N}\left(U_i|0,\sigma_U^2 I\right) \\ P(V|\sigma_V^2) &= \prod_j \mathcal{N}\left(V_j|0,\sigma_V^2 I\right). \end{aligned} \quad (2)$$

- **User activity recognition**. Given the user activity sequences Y, we are seeking a matrix factorization of the "Users-Activities" matrix A as a product of two low-dimensional latent feature representations U and Z. Here, U is the user feature matrix shared with the "Users-Profiles" matrix factorization; Z is the activity feature matrix with each Z_k denoted as a low-dimensional feature vector of activity k. Similarly, we formulate such a matrix factorization as a conditional probability:

$$P(A|U,Z,Y,I^A,\sigma_A^2) = \prod_{i,k}\mathcal{N}\left(A = Y_i I_k^A|U_i Z_k^T,\sigma_A^2\right), \quad (3)$$

and similarly, we place the zero-mean Gaussian priors on the activity feature matrix:

$$P(Z|\sigma_Z^2) = \prod_k N\left(Z_k|0, \sigma_Z^2 I\right). \quad (4)$$

Note that the user activities can be inferred from both the user location information and the sequential information, so we model the conditional probability for the activity sequences in a parametric way:

$$P(Y|X, R, \Lambda) = \prod_{i,t} P(Y_{i,t}|Y_{i,t-1}, X_{i,t}, R_i, \Lambda), \quad (5)$$

where Λ is the model parameter for inferring the activities. The probability of $P(Y_{i,t}|Y_{i,t-1}, X_{i,t}, R_i, \Lambda)$ can be defined in various ways. For example, one may set it as multinomial logit model. We also place some Gaussian priors on the model parameter Λ:

$$P(\Lambda|\Gamma, \sigma_\Lambda^2) = \prod_y \mathcal{N}\left(\Lambda_y|\Gamma_y, \sigma_\Lambda^2 I\right), \quad (6)$$

where Γ is the prior.

After we have formulated the conditional probabilities for both activity recognition and user profile learning, we can calculate the corresponding posterior probabilities and thus have a final optimization function[1] as

$$maxP(U, V, Z, Y, \Lambda|X, R, I^A, \Gamma, \sigma_U^2, \sigma_V^2, \sigma_Z^2, \sigma_\Lambda^2, \sigma_R^2). \quad (7)$$

Finally, given the GPS data X, user profiles R and partially labeled activity sequences Y, we can optimize the Equation (7) and output the missing entries of the user profiles and missing activities for the GPS trajectories, completing our joint learning task.

4. CONCLUSION

In this paper, we have show how to go a step beyond GPS data mining to learn the user activities and profiles from the raw GPS data. We exploit the inherent connections between the user activities and user profiles, and provide a possible solution to jointly learn them together to better understand the mobile users' behaviors. In particular, we formulate the joint learning problem as a probabilistic collaborative filtering problem, which takes as input the partially labeled activity sequences based on the GPS trajectories and incomplete user profile matrices. Finally, we aim to output the missing activity labels for the activity sequences and the missing entries for the user profile matrices.

Acknowledgement

This work is supported by Microsoft Research Asia Internet Services Theme Research Program.

5. REFERENCES

[1] Bikely: http://www.bikely.com/.

[2] Gps track route exchange forum: http://www.gpsxchange.com/.

[3] G. Adomavicius and A. Tuzhhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. In *IEEE Transaction on Knowledge and Data Engineering*, pages 734–749.

[4] H. H. Bui. A general model for online probabilistic plan recognition. In *Proc. of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 1309–1318, 2003.

[5] Q. Li, Y. Zheng, X. Xie, Y. Chen, W. Liu, and W.-Y. Ma. Mining user similarity based on location history. In *Proc. of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems*, pages 1–10, 2008.

[6] L. Liao, D. Fox, and H. A. Kautz. Extracting places and activities from gps traces using hierarchical conditional random fields. *International Journal of Robotics Research*, 26(1):119–134, 2007.

[7] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *Proc. of the 17th ACM conference on Information and knowledge management*, pages 931–940, 2008.

[8] Y. Matsuo, N. Okazaki, K. Izumi, Y. Nakamura, T. Nishimura, and K. Hasida. Inferring long-term user property based on users location history. In *Proc. of the 20th International Joint Conference on Artificial Intelligence*, 2007.

[9] D. J. Patterson, D. Fox, H. A. Kautz, and M. Philipose. Fine-grained activity recognition by aggregating abstract object usage. In *Proc. of the Ninth IEEE International Symposium on Wearable Computers*, pages 44–51, 2005.

[10] M. E. Pollack, L. E. Brown, D. Colbry, C. E. McCarthy, C. Orosz, B. Peintner, S. Ramakrishnan, and I. Tsamardinos. Autominder: an intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems*, 44(3-4):273–282, 2003.

[11] A. P. Singh and G. J. Gordon. Relational learning via collective matrix factorization. In *Proc. of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 650–658, 2008.

[12] S. Wang, W. Pentney, and A. maria Popescu. Common sense based joint training of human activity recognizers. In *Proc. of the 20th International Joint Conference on Artificial Intelligence*, pages 2237–2242, 2007.

[13] D. Wyatt, M. Philipose, and T. Choudhury. Unsupervised activity recognition using automatically mined common sense. In *Proc. of the 20th National Conference on Artificial Intelligence*, pages 21–27, 2005.

[14] V. W. Zheng, D. H. Hu, and Q. Yang. Cross-domain activity recognition. In *Proc. of the 11th International Conference on Ubiquitous Computing*, 2009.

[15] Y. Zheng, L. Liu, L. Wang, and X. Xie. Learning transportation mode from raw gps data for geographic applications on the web. In *Proc. of the 17th international conference on World Wide Web*, pages 247–256, 2008.

[16] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma. Mining interesting locations and travel sequences from gps trajectories. In *Proc. of the 18th International World Wide Web Conference*, 2009.

[1] Details are skipped here.