

Mining urban lifestyles: urban computing, human behavior and recommender systems

Sharon Xu¹, Riccardo Di Clemente^{2,3}, and Marta C. González^{4,5*}

¹Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²Centre for Advanced Spatial Analysis, University College London, London WC1E 6BT, GBR

³Department of Computer Science, University of Exeter, Exeter EX4 4PY, GBR

⁴Department of City and Regional Planning, UC Berkeley, 406 Wurster Hall, Berkeley, 94720, USA

⁵Lawrence Berkeley National Laboratory, Cyclotron Road, Berkeley, 94720, USA

*martag@berkeley.edu

This is a postprint of the 5th chapter published in the book *Big Data Recommender Systems - Volume 2: Application Paradigms* and is subject to Institution of Engineering and Technology Copyright. The copy of record is available at the IET Digital Library DOI: [10.1049/PBPC035G_ch5](https://doi.org/10.1049/PBPC035G_ch5)

In the last decade, the digital age has sharply redefined the way we study human behavior. With the advancement of data storage and sensing technologies, electronic records now encompass a diverse spectrum of human activity, ranging from location data^{1,2}, phone^{3,4} and email communication⁵ to Twitter activity⁶ and open-source contributions on Wikipedia and OpenStreetMap^{7,8}. In particular, the study of the shopping and mobility patterns of individual consumers has the potential to give deeper insight into the lifestyles and infrastructure of the region. Credit card records (CCRs) provide detailed insight into purchase behavior and have been found to have inherent regularity in consumer shopping patterns⁹; call detail records (CDRs) present new opportunities to understand human mobility¹⁰, analyze wealth¹¹, and model social network dynamics¹².

Regarding the analysis of CDR data, there exists a wide body of work characterizing human mobility patterns. As a notable example,¹⁰ describes the temporal and spatial regularity of human trajectories, showing that each individual can be described by a time independent travel distance and a high probability of returning to a small number of locations. Further, the authors are able to model individual travel patterns using a single spatial probability distribution. There has also been work at the intersection of similar datasets, such as the inference of friendships from mobile phone data¹³, or the analysis such data in relation to metrics on spending behavior such as diversity, engagement, and loyalty¹⁴. Recent work¹⁵ uses the Jaccard distance as a similarity measure on motifs among spending categories, then applies community detection algorithms to find clusters of users. These studies propose models for either mobility or spending behavior, but not in conjunction.

The only known paper that incorporates both aspects¹⁶ frames its analysis only on an aggregate scale of city regions. However, the coupled collaborative filtering methods (also known as collective matrix factorization) used in¹⁶ have been successfully applied in a variety of urban computing applications for data fusion and prediction¹⁷⁻¹⁹, from location-based activity recommendations^{20,21} to travel speed estimation on road segments²². Recent work includes methods that use Laplacian regularization²³ to leverage social network information, and use geometric deep learning matrix completion methods to model nonlinearities²⁴.

In this chapter, we jointly model the lifestyles of individuals, a more challenging problem with higher variability when compared to the aggregated behavior of city regions. Using collective matrix factorization, we propose a unified dual view of lifestyles. Understanding these lifestyles will not only inform commercial opportunities, but also help policymakers and nonprofit organizations understand the characteristics and needs of the entire region, as well as of the individuals within that region. The applications of this range from targeted advertisements and promotions to the diffusion of digital financial services among low-income groups.

Mining Shopping and Mobility Patterns

Location and transactional data offer valuable perspectives on the lifestyles of each user. For example, we may expect the shopping purchases of middle-aged parents to include groceries and fuel, while their mobility patterns may center around localities near home and work locations, in addition to points of interest such as supermarket, laundry, and so on. We use mobility information to aid in the prediction of shopping behavior, connecting the two views using collective matrix factorization²⁵. In

this way, we discover representative patterns relating shopping and mobility, characterizing behavior for a richer understanding into urban lifestyles and improved prediction of behavior.

The high granularity of such digital records allows modeling at the level of the individual, providing a new framework in which to relate movement and spending. However, in using CDR data for data on individuals, we must deal with issues of sparsity and lack of contextual information on the user's activities. In proposing this dual view of lifestyles, our contributions can be summarised as follows:

Prediction of Shopping Behavior with Data Sparsity

There are many individuals for which we have no CDR data. To deal with this data sparsity issue, we construct a framework that uses mobility patterns as supplementary information in the prediction of shopping behavior. We connect the two perspectives on lifestyles using collective matrix factorization (collective matrix factorization). In comparison to modeling only shopping behavior, we find that incorporating mobility information in the prediction of shopping lifestyles leads to a significant reduction in root mean square error (RMSE).

Adding Contextual Information to Location Data

We transform mobility data using external data sources to better relate CCR to CDR Data. Although CCRs provide high granularity at the level of the individual user, spatial granularity can range from a radius of 200 - 1000 meters, and there is no contextual information for the user's activities within that region. Thus, there has been little previous work leveraging CCR data for prediction with CDR data.

Multi-Perspective Lifestyles

We describe the mappings between shopping and mobility patterns, connecting the two views to provide a novel understanding of consumer behavior in urban regions.

Data

The primary datasets used in this chapter consist of two sets of anonymised data for residents in Mexico throughout five months in 2015:

- Call detail records (CDRs). CDRs are produced with each telephone exchange, These location records give the nearest cellular tower at the time of a placed call. There are 1192 cell towers throughout Mexico City – as users tend to visit a small subset of these towers, this mobility data is extremely sparse. In a count matrix denoting user visits to towers, 98% of entries indicate zero visits.
- Credit card records (CCRS). CCRs are recorded with each purchase and denote the purchase category, or Merchant Category Code (MCC), of the transaction as well as the amount spent. Each month, we have on the order of 10 million financial transactions and 200 million location records.

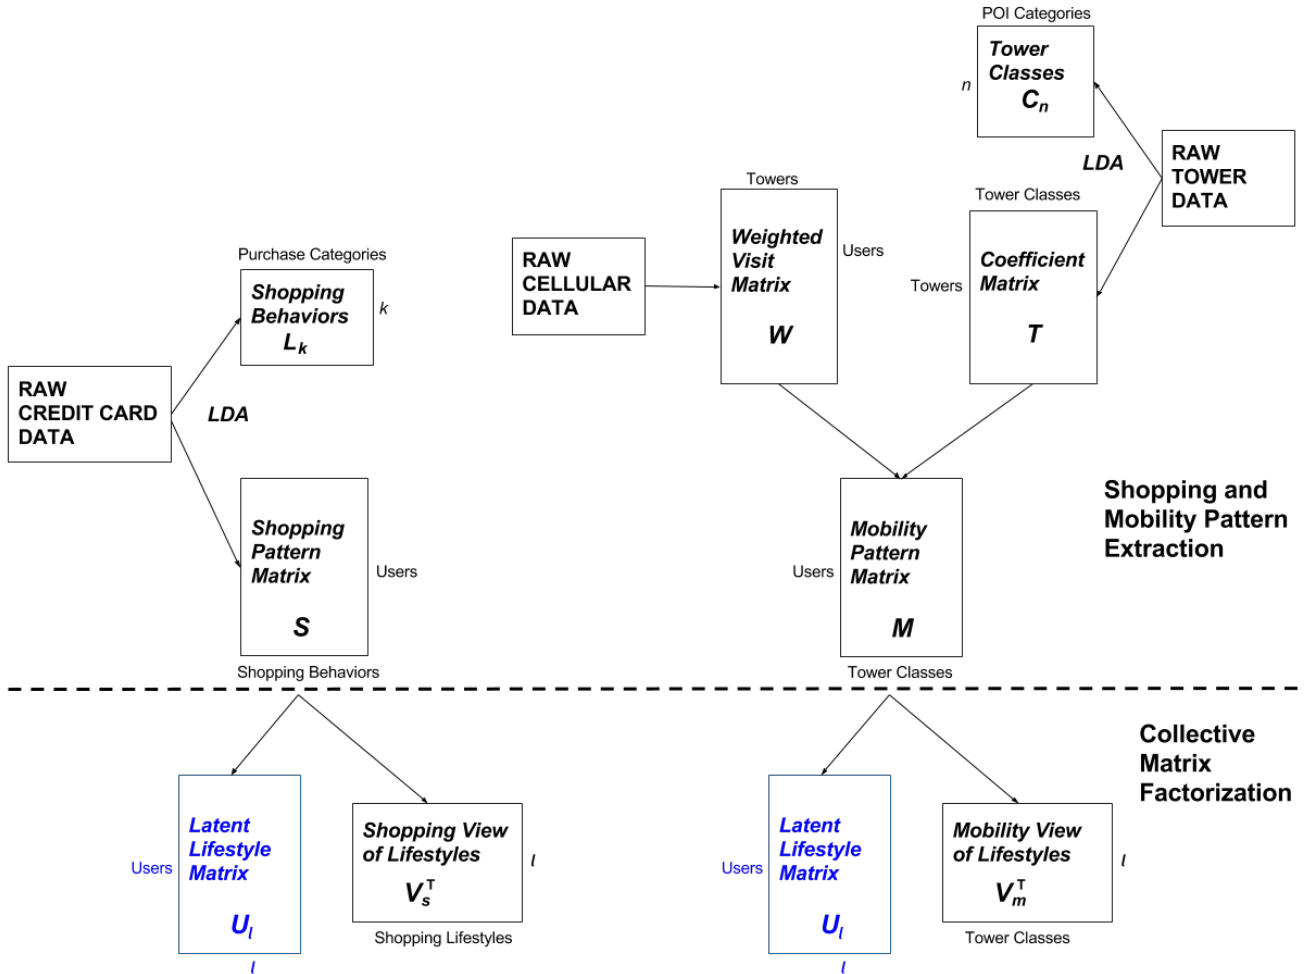
Discovering Shopping Patterns

Our spending habits reflect our lifestyles, capturing an essential aspect of our behavior. Within the computational social science community, the question remains whether pervasive trends exist among disparate groups at urban scale¹⁵. In this chapter, we use latent Dirichlet allocation (LDA)²⁶ to identify topics (behavioral patterns) among individuals, representing each individual's spending lifestyle as a finite mixture of an underlying set of behaviors. Each behavioral pattern, in turn, is modeled as a mixture of a set of words (Merchant Category Codes, or MCCs). These topics are determined by co-occurrences of words within a document. For example, in an article database, we may uncover a topic containing the words "data", "processing", "computer", and so on because these words frequently appear in an article together.

By putting a Dirichlet prior on the per-user behavior distribution and per-behavior MCC distribution, LDA controls the sparsity of the number of topics per document (the number of behaviors per individual), as well as the number of words per topic (the number of MCCs per behavioral pattern). In this way, each individual is represented by a small number of behaviors, and each behavior involves making a small set of purchase categories with high frequency.

As a generative model, LDA allows us to calculate the probabilities (assignments to shopping behaviors) of previously unseen users. We train the model on 40% of the users, and generate the matrix S for the remaining 60%. In so doing, we set up the prediction of lifestyles for unseen users, assessing the LDA model itself in addition to the relation of shopping with mobility patterns. We experiment with the choice of number of behaviors to learn, as well as adding a categorical variable describing amount spent to each MCC. To maximise interpretability, we choose five topics while using MCCs as input only.

Figure 1. Our framework



In Fig. 2 we plot the twenty most highly weighted MCCs of the five shopping behaviors. The first shopping behavior describes credit card usage that is centered on food-related purchases such as Grocery Stores, Misc. Food Stores and Restaurants. The second shopping behavior seems to be associated primarily with business purchases, with spending within MCCs such as Fax Services and Financial Institutions. The third shopping behavior is dominated by relative “luxuries” such as purchases in the Cable and Department Store categories, and is characterised by a relatively high proportion of Air Travel and Hotel Lodging MCCs. The fourth shopping behavior contains primarily purchases in Computer Network Services and Service Stations (gas stations). The third and fourth shopping behavior describe a slightly wealthier portion of the population, as only 35% of Mexicans owned a computer in 2010²⁷, and only 44.2% own a car²⁸. Lastly, the fifth shopping behavior captures purchase primarily for toll fees and subscription services.

Mobility Pattern Extraction

Extracting Cellular Tower Location Types

Within the CDR data, each tower is the site for a corresponding cell within the Voronoi diagram; i.e., it is the closest tower to any point within this cell. We define a “visit” to a cellular tower as a call placed within its corresponding cell. In order to relate cellular towers to spending behavior, for each tower we crawl Google’s API for points of interest within a certain radius. To determine this radius, we use Delaunay triangulation, a widely used method in computational geometry. Delaunay triangulation gives the dual graph to the Voronoi diagram, maximizing the minimum angle among all the triangles within the triangulation, and connecting the sites in a nearest-neighbor fashion²⁹. For each tower, we set the crawling radius to be half the average distance from the site to its neighbors.

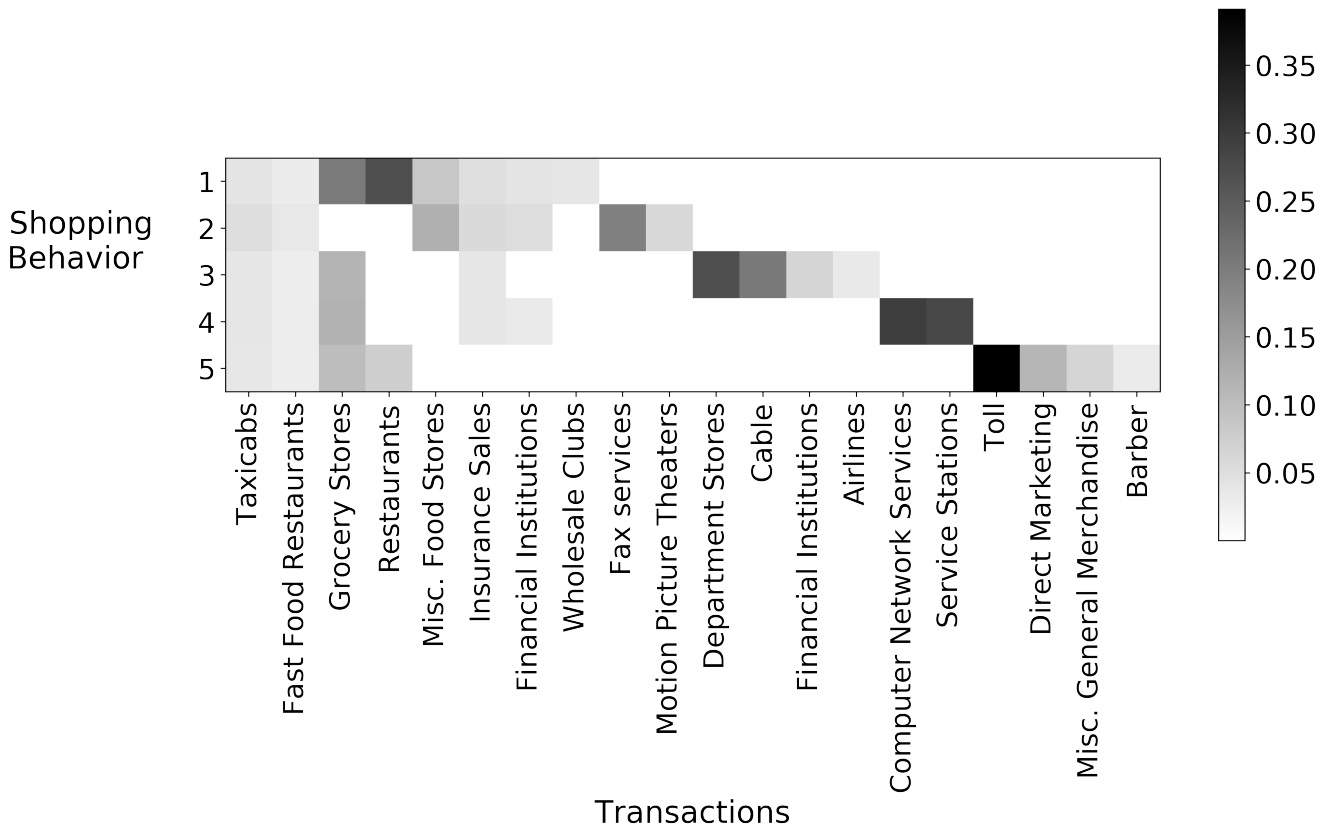


Figure 2. The top weighted purchase categories of the five shopping behaviors learned from LDA.

Treating each of the Voronoi cells as a document and the POI categories as words, we use latent Dirichlet allocation to discover underlying tower “classes” that will be more informative of shopping behavior. We remove from the vocabulary any POI categories that occur with over 25% frequency. These removed categories are uninformative classifications such as “point of interest” and “establishment”. For purposes of interpretability, we learn the LDA model with twenty classes on the 1192 towers.

In Fig. 3, we show a subset of tower classes highly weighted within our final lifestyles (see Section 6), and the corresponding points of interest with the highest probabilities. From the sample topics in Fig. 4, we see each tower class puts specific emphasis on related points of interest, such as “hospital” and “doctor”, “car rental” and “car repair”, or “book store” and “library”. In this way, we cluster the towers in terms of nearby POI categories, obtaining contextual information more directly related to shopping.

Baseline Methods

Before introducing our model, we present the results of several baseline methods, illustrating the challenges of incorporating CDR data into the prediction of shopping patterns.

Regression on Average Amount Spent

Using the columns of the per tower count matrix W directly as features, we use regression with L1 regularization to predict the average amount spent by the user per week. As we increase regularization we increase the test R-squared, but due to a combination of sparsity and lack of signal achieve a maximum test R-squared of 0 as the coefficients shrink to 0.

Classification of Primary Shopping Behavior

For each user, we take as our outcome the highest weighted shopping behavior from the topic proportions learned from LDA. This is the user’s primary behavior. Again using the columns of W as our features, we employ a range of classifiers including

Figure 3. The top weighted POI categories of a subset of tower classes learned from LDA.

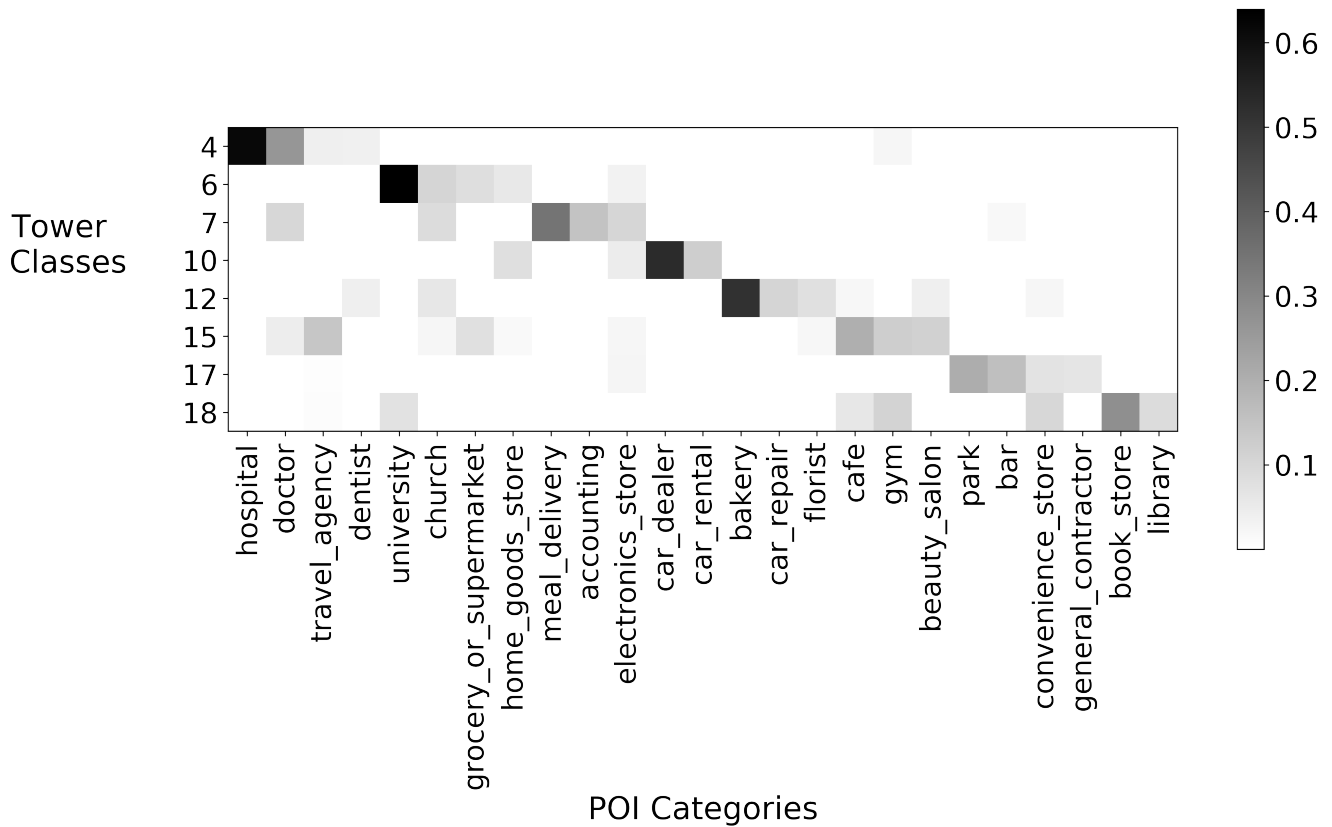


Figure 4. Sample topics from learned from LDA, treating each tower as a document and each POI as a word.

```

0.617"hospital" + 0.264"doctor"
+ 0.041"travel_agency"
+ 0.038"dentist" + 0.023"gym"

0.569"clothing_store"
+ 0.090"department_store"
+ 0.048"shopping_mall"

0.765"lodging" + 0.059"bar"
+ 0.052"museum" + 0.018"travel_agency"

0.247"atm" + 0.176"bank" + 0.118"police"
+ 0.088"post_office" + 0.079"city_hall"
+ 0.071"local_government_office"

```

SVM and AdaBoost to predict primary behavior. We find that the best classifier achieves only 21.6% accuracy, when already 21.9% of users fall into a single class.

Characterizing Mobility Patterns

From the Voronoi diagram of the p cell tower locations, we construct a matrix $W \in \mathbb{R}^{n \times p}$ where each entry w_{ij} is the number of days individual i visited tower j throughout five months. We weight these counts using TF-IDF, a common method for text representation³⁰. Using TF-IDF, we offset the tower counts by the frequency of the tower in the data, so that a user's visit to an uncommonly visited tower is assigned a higher weight. We now have a matrix $W \in \mathbb{R}^{n \times p}$ that characterises users in terms of tower visits, and a matrix $C_m \in \mathbb{R}^{p \times d}$, where d is the chosen number of tower classes. We define our mobility pattern matrix as $M = WT$, achieving a significant dimensionality reduction with $M \in \mathbb{R}^{n \times d}$. In this manner, we obtain a representation of mobility more closely related to shopping behavior, as users are now characterised by their visits to tower classes defined by POI categories.

Predicting Shopping Behavior

For many users, we have access to data on mobility patterns (M) but not shopping patterns (S). In this section, we describe our methodology for incorporating mobility information in addition to shopping information for the matrix completion problem of predicting the shopping behavior of unseen users.

Collective Matrix Factorization

We denote S as the matrix of behavior proportions obtained from latent Dirichlet allocation, and M as the matrix of weighted visit frequencies to the different tower classes. Modeling each user's shopping and mobility behavior as two views of the same lifestyle, we assume that S and M are generated from a matrix U_l containing the latent lifestyle information of each user.

$$S \approx U_l V_s^T$$

$$M \approx U_l V_m^T$$

Traditionally, the objective function under this model is represented as

$$\mathcal{L}(U_l, V_s, V_m) = \|S - U_l V_s^T\|^2 + \|M - U_l V_m^T\|^2 + \lambda_1 \|U_l\|^2 + \lambda_2 \|V_s\|^2 + \lambda_3 \|V_m\|^2$$

In this chapter, we use group-wise sparse collective matrix factorization³¹, which puts group-sparse priors following $\mathcal{N}(0, \sigma_k^2)$ on the columns of matrices V_s and V_m , where the columns are the groups indexed by k and σ_k^2 is small. This allows the matrix to learn private factors for the relation between latent lifestyles (U_l) and the shopping aspect (V_s), and correspondingly between latent lifestyles (U_l) and the mobility aspect (V_m). More specifically, if the k^{th} column of V_m is null, then the k^{th} factor impacts only the shopping pattern matrix S .

Results

Prediction

In our problem, credit card data is unknown for many users, but we would like to use mobility information to predict their shopping behavior; i.e., S contains many empty rows. Thus, to test the performance within this setting, we remove rows from the shopping behavior S to predict the shopping behavior of users for which we have no credit card information. We use 10-fold cross validation and compare our collective matrix factorization predictions with the actual values. We use the popular metric root mean square error (RMSE) to evaluate our model.

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{i,j} (S_{i,j} - \hat{S}_{i,j})^2}$$

Using cross-validation to determine the rank (number of lifestyles), we find that the inclusion of mobility data leads to a 1.3% decrease in RMSE and obtain a test error of 21.6%.

Dual Lifestyles

Using collective matrix factorization, we also obtain both the dual shopping and mobility views of these latent lifestyles, in V_s and V_m respectively.

Lifestyle 1 is connected with wealthier shopping behavior typical common to urban white collars. The top weighted shopping patterns indicate spending on cable, air travel, hotels and at department stores as well as gas stations and computer network services (Fig. 2: behaviors 3 and 4, respectively). This suggests that people who can afford to spend on relative luxuries tend to

have vehicles and thus higher mobility, visiting a wider range of tower classes. The mobility patterns of this lifestyle focus on areas with points of interest such as universities, accounting, electronics, bakeries and car repair (Fig. 3: tower classes 6, 7, 12, 17 and 20).

Lifestyle 2 is extremely food-oriented, with high weight on shopping behavior 1. Mobility patterns suggest visits to cafes, gyms and convenience stores.

Lifestyle 3 primarily captures the transportation aspect of lifestyles. Top weighted mobility patterns indicate visits to areas with car rental and car repair (tower classes 10 and 12), while shopping patterns include gas stations in behavior 4 and food in behavior 1.

Discussion

In this study, we relate the shopping and mobility patterns of consumers on an individual level for the first time. Viewing these as aspects of the same underlying lifestyle, we set up a framework to incorporate CDR data in the prediction of shopping patterns for unseen users. We achieve a significant increase in prediction and recover interesting relationships between shopping and mobility.

There are many directions for future work. In terms of modeling formulation, it would be interesting to introduce a temporal dimension into the task of shopping prediction, as human behavior and needs vary over time. There is also the opportunity to include social regularization in the collective matrix factorization formulation, constraining each user to be similar to his or her neighborhood. In addition, stronger prediction methods may be achieved by modeling nonlinear relationships using geometric deep learning methods described by²⁴.

Acknowledgements

We thank Grandata for supplying the data. (For contractual and privacy reasons, the raw data cannot be provided.)

References

1. Song, C., Qu, Z., Blumm, N. & Barabási, A. L. Limits of predictability in human mobility. *Science* **327**, 1018–1021, DOI: [10.1126/science.1177170](https://doi.org/10.1126/science.1177170) (2010).
2. Hasan, S., Schneider, C. M., Ukkusuri, S. V. & González, M. C. Spatiotemporal Patterns of Urban Human Mobility. *J. Stat. Phys.* **151**, 304–318, DOI: [10.1007/s10955-012-0645-0](https://doi.org/10.1007/s10955-012-0645-0) (2013).
3. Jo, H.-H., Karsai, M., Kertész, J. & Kaski, K. Circadian pattern and burstiness in mobile phone communication. *New J. Phys.* **14**, 013055xw, DOI: [10.1088/1367-2630/14/1/013055](https://doi.org/10.1088/1367-2630/14/1/013055) (2012).
4. Aledavood, T. *et al.* Daily rhythms in mobile telephone communication. *PLoS ONE* **10**, e0138098, DOI: [10.1371/journal.pone.0138098](https://doi.org/10.1371/journal.pone.0138098) (2015).
5. Malmgren, R. D., Stouffer, D. B., Motter, A. E. & Amaral, L. A. A Poissonian explanation for heavy tails in e-mail communication. *Proc. Natl. Acad. Sci. United States Am.* **105**, 18153–18158, DOI: [10.1073/pnas.0800332105](https://doi.org/10.1073/pnas.0800332105) (2008).
6. ten Thij, M., Bhulai, S. & Kampstra, P. Circadian patterns in twitter. *Data Anal. 2014* 12–17 (2014).
7. Yasserli, T., Sumi, R. & Kertész, J. Circadian patterns of wikipedia editorial activity: A demographic analysis. *PLoS ONE* **7**, e30091, DOI: [10.1371/journal.pone.0030091](https://doi.org/10.1371/journal.pone.0030091) (2012).
8. Yasserli, T., Quattrone, G. & Mashhadi, A. Temporal analysis of activity patterns of editors in collaborative mapping project of openstreetmap. In *Proceedings of the 9th International Symposium on Open Collaboration, WikiSym + OpenSym 2013*, 13, DOI: [10.1145/2491055.2491068](https://doi.org/10.1145/2491055.2491068) (ACM, 2013).
9. Krumme, C., Llorente, A., Cebrian, M., Pentland, A. & Moro, E. The predictability of consumer visitation patterns. *Sci. Reports* **3**, DOI: [10.1038/srep01645](https://doi.org/10.1038/srep01645) (2013).
10. González, M. C., Hidalgo, C. A. & Barabási, A. L. Understanding individual human mobility patterns (Nature (2008) 453, (779-782)). *Nature* **458**, 238, DOI: [10.1038/nature07850](https://doi.org/10.1038/nature07850) (2009).
11. Blumenstock, J., Cadamuro, G. & On, R. Predicting poverty and wealth from mobile phone metadata. *Science* **350**, 1073–1076, DOI: [10.1126/science.aac4420](https://doi.org/10.1126/science.aac4420) (2015).

12. Morse, S., Gonzalez, M. C. & Markuzon, N. Persistent cascades: Measuring fundamental communication structure in social networks. *Proc. - 2016 IEEE Int. Conf. on Big Data, Big Data 2016* 969–975, DOI: [10.1109/BigData.2016.7840695](https://doi.org/10.1109/BigData.2016.7840695) (2016).
13. Eagle, N., Pentland, A. & Lazer, D. Inferring friendship network structure by using mobile phone data. *Proc. Natl. Acad. Sci. United States Am.* **106**, 15274–15278, DOI: [10.1073/pnas.0900282106](https://doi.org/10.1073/pnas.0900282106) (2009).
14. Singh, V. K., Freeman, L., Lepri, B. & Pentland, A. S. Predicting Spending Behavior Using Socio-mobile Features. In *2013 International Conference on Social Computing, SOCIALCOM '13*, 174–179, DOI: [10.1109/SocialCom.2013.33](https://doi.org/10.1109/SocialCom.2013.33) (IEEE, Washington, DC, USA, 2013).
15. Di Clemente, R. *et al.* Sequences of purchases in credit card data reveal lifestyles in urban populations. *Nat. Commun.* **9**, DOI: [10.1038/s41467-018-05690-8](https://doi.org/10.1038/s41467-018-05690-8) (2018).
16. Hu, T., Song, R., Wang, Y., Xie, X. & Luo, J. Mining shopping patterns for divergent urban regions by incorporating mobility data. In *International Conference on Information and Knowledge Management, Proceedings*, vol. 24-28-Octo, 569–578, DOI: [10.1145/2983323.2983803](https://doi.org/10.1145/2983323.2983803) (ACM, 2016).
17. Zheng, Y. Methodologies for Cross-Domain Data Fusion: An Overview. *IEEE Transactions on Big Data* **1**, 16–34, DOI: [10.1109/tbdata.2015.2465959](https://doi.org/10.1109/tbdata.2015.2465959) (2015).
18. Zheng, Y., Capra, L., Wolfson, O. & Yang, H. Urban computing: concepts, methodologies, and applications. *ACM Transactions on Intell. Syst. Technol. (TIST)* **5**, 38, DOI: [10.1145/2629592](https://doi.org/10.1145/2629592) (2014).
19. Zheng, Y. Trajectory data mining: An overview. *ACM Transactions on Intell. Syst. Technol.* **6**, 29, DOI: [10.1145/2743025](https://doi.org/10.1145/2743025) (2015).
20. Zheng, V. W., Zheng, Y., Xie, X. & Yang, Q. Collaborative location and activity recommendations with GPS history data. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, 1029–1038, DOI: [10.1145/1772690.1772795](https://doi.org/10.1145/1772690.1772795) (ACM, 2010).
21. Zheng, V. W., Cao, B., Zheng, Y., Xie, X. & Yang, Q. Collaborative filtering meets mobile recommendation: A user-centered approach. In *Proceedings of the National Conference on Artificial Intelligence*, vol. 1, 236–241 (2010).
22. Shang, J., Zheng, Y., Tong, W., Chang, E. & Yu, Y. Inferring gas consumption and pollution emission of vehicles throughout a city. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1027–1036, DOI: [10.1145/2623330.2623653](https://doi.org/10.1145/2623330.2623653) (ACM, 2014).
23. Leng, C., Zhang, H., Cai, G., Cheng, I. & Basu, A. Graph regularized Lp smooth non-negative matrix factorization for data representation. *IEEE/CAA J. Autom. Sinica* **6**, 584–595, DOI: [10.1109/JAS.2019.1911417](https://doi.org/10.1109/JAS.2019.1911417) (2019).
24. Bronstein, M. M., Bruna, J., Lecun, Y., Szlam, A. & Vandergheynst, P. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Process. Mag.* **34**, 18–42, DOI: [10.1109/MSP.2017.2693418](https://doi.org/10.1109/MSP.2017.2693418) (2017).
25. Singh, A. P. & Gordon, G. J. Relational learning via collective matrix factorization. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '08*, 650–658, DOI: [10.1145/1401890.1401969](https://doi.org/10.1145/1401890.1401969) (ACM, New York, NY, USA, 2008).
26. Blei, D. M., Ng, A. Y. & Jordan, M. I. Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022, DOI: [10.1016/b978-0-12-411519-4.00006-9](https://doi.org/10.1016/b978-0-12-411519-4.00006-9) (2003).
27. Villagram, L. *For most Mexicans, the digital age is still out of reach*, vol. 16 (Dallas News, 2012).
28. Morales, S. E., Cosart, T. & Holben, W. E. Bacterial gene abundances as indicators of greenhouse gas emission in soils. *The ISME J.* **4**, 799–808, DOI: [10.1038/ismej.2010.8](https://doi.org/10.1038/ismej.2010.8) (2010).
29. Aurenhammer, F. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Comput. Surv. (CSUR)* **23**, 345–405, DOI: [10.1145/116873.116880](https://doi.org/10.1145/116873.116880) (1991).
30. Zhang, W., Yoshida, T. & Tang, X. A comparative study of TF*IDF, LSI and multi-words for text classification. *Expert. Syst. with Appl.* **38**, 2758–2765, DOI: [10.1016/j.eswa.2010.08.066](https://doi.org/10.1016/j.eswa.2010.08.066) (2011).
31. Klami, A., Bouchard, G. & Tripathi, A. Group-sparse embeddings in collective matrix factorization. *2nd Int. Conf. on Learn. Represent. ICLR 2014 - Conf. Track Proc.* (2014). [1312.5921](https://arxiv.org/abs/1312.5921).