

**Spatio-temporal patterns of human mobility
from geo-social networks for urban computing:
Analysis, models & applications**



Xiao Zhou

Supervisor: Prof. Cecilia Mascolo

Department of Computer Science and Technology
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

In loving memory of Grandma & Grandpa.

Declaration

I hereby declare that the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except where specified in the text. This dissertation does not exceed the prescribed limit of 60,000 words, including footnotes, tables, and equations.

Xiao Zhou
May 2020

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Cecilia Mascolo. To me, she means more than a supervisor, but one of the most important people in my life. Without her encouragement, guidance, and considerable patience, I would not be studying this degree and cannot even imagine where I would be today. Thank you so much for your unreserved support to make the impossible possible.

Subsequently, a heartfelt "Thank you" also goes to Anastasios Noulas, whose creativity has been continuously inspiring me throughout my Ph.D. I can never thank him enough for all the constructive feedback he has been giving me. I am further grateful to Desislava Hristova for supporting me with her expertise and infecting me with a passion for research in the first year when I needed it the most. I would like to give special thanks to Petko Georgiev for taking great patience to answer my stupid questions when I was learning programming skills from scratch. I have also been fortunate to have the mentorship of Xing Xie at Microsoft Research. I am humbled by the great work he does and for having the chance to be a small part of it.

My sincere thanks to those with whom I have co-authored papers over the last three and half years: Zhongxiang Zhao, Yingzi Wang, Jianxun Lian, Danyang Liu, Max Sklar, and Enhong Chen for their guidance and help during the progress of my research.

I would not have enjoyed my Ph.D. time so much if it were not for my new and old friends alike. I am so lucky to have had the chance to meet my friends at the lab: Alessandro, Lorena, David, Sandra, Andrea, Andreas, Krittika, Apinan, Dimitris, Dionysis, Jagmohan, Zhao, Young, Alex, and others. I am grateful for all the happy moments we have enjoyed together and the support you have given me. Also, I would like to express my gratitude to my closest friends at Queens' College: Jingchen, Yaoyao, and Jiaqi for always being by my side. I will never forget the times they believed in me even when I doubted myself. Last but not least, I am grateful to all the friends I made during my internships at Tencent and Microsoft Research. They have truly made my time in Beijing an experience I would love to repeat.

Finally, I owe sincere and earnest gratitude to my parents for their unconditional love and endless support throughout my life, without which any of my achievements would not have been possible.

Abstract

The availability of rich information about fine-grained user mobility in urban environments from increasingly geographically-aware social networking services and the rapid development of machine learning applications greatly facilitate the investigation of urban issues. In this setting, urban computing emerges intending to tackle a variety of challenges faced by cities nowadays and to offer promising approaches to improving our living environment. Leveraging massive amounts of data from geo-social networks with unprecedented richness, we show how to devise novel algorithmic techniques to reveal underlying urban mobility patterns for better policy-making and more efficient mobile applications in this dissertation.

Building upon the foundation of existing research efforts in urban computing field and basic machine learning techniques, in this dissertation, we propose a general framework of urban computing with geo-social network data and develop novel algorithms tailored for three urban computing tasks. We begin by exploring how the transition data recording human movements between urban venues from geo-social networks can be aggregated and utilised to detect spatio-temporal changes of local graphs in urban areas. We further explore how this can be used as a proxy to track and predict socio-economic deprivation changes as government financial effort is put in developing areas by supervised machine learning methods. We then study how to extract latent patterns from collective user-venue interactions with the help of a spatio-temporal aware topic modeling approach for the benefit of urban infrastructure planning. After that, we propose a model to detect the gap between user-side demand and venue-side supply levels for certain types of services in urban environments to suggest further policymaking and investment optimisation. Finally, we address a mobility prediction task, the application aim of which is to recommend new places to explore in the city for mobile users. To this end, we develop a deep learning framework that integrates memory network and topic modeling techniques. Extensive experiments indicate that the proposed architecture can enhance the prediction performance in various recommendation scenarios with high interpretability.

All in all, the insights drawn and the techniques developed in this dissertation make a substantial step in addressing issues in cities and open the door to future possibilities in the promising urban computing area.

Table of contents

List of figures	xv
List of tables	xvii
1 Introduction	1
1.1 The rise of urban computing	2
1.2 Digital data for urban computing	3
1.3 Current research efforts & limitations	5
1.4 Thesis & substantiation	7
1.5 Contributions & chapter outline	8
1.6 List of publications	10
2 Urban computing with GSN data	13
2.1 Urban computing framework with GSN data	14
2.2 Urban data preprocessing	15
2.2.1 Data cleaning	15
2.2.2 Data integration	15
2.3 Machine learning for urban computing	16
2.3.1 Classification & prediction	16
2.3.2 Cluster analysis of mobility patterns	19
2.3.3 Topic modeling & urban semantics	22
2.4 Mining graphs & place networks	22
2.5 Metric learning-based venue recommendation	23
2.6 Challenges & opportunities	26
2.6.1 Possible limitations in GSN data	26
2.6.2 Knowledge discovery from heterogeneous data sources	27
2.6.3 Innovative techniques tailored for urban computing	27
2.6.4 Effective urban computing platforms for diverse applications	27

2.7	Present dissertation & future outlook	28
3	Cultural investment & urban socio-economic development	31
3.1	Introduction	31
3.2	Dataset	34
3.2.1	Socio-economic data	34
3.2.2	Cultural expenditure data	35
3.2.3	Foursquare data	36
3.2.4	Spatial unit of data	37
3.3	Metrics	38
3.3.1	Geographic features	38
3.3.2	Network metrics	40
3.3.3	Growth rate	40
3.3.4	Cultural advantage metrics	41
3.4	Preliminary analysis	44
3.5	Hypotheses	47
3.6	Evaluation & results	47
3.6.1	[H1] Network and local properties	48
3.6.2	[H2] Growth of network and local properties	49
3.6.3	[H3] Prediction model and overall evaluation	51
3.6.4	[H4] Individual features evaluation	53
3.7	Discussion & conclusions	55
4	Cultural patterns extraction & modern urban planning	57
4.1	Introduction	57
4.2	Related work	59
4.3	Approach at a glance	60
4.4	Datasets	61
4.5	Cultural patterns extraction	63
4.5.1	Temporal factors extraction	63
4.5.2	Temporal latent Dirichlet allocation	65
4.5.3	Evaluation of TLDA	68
4.6	Refined urban cultural planning	69
4.6.1	Demand range determination	70
4.6.2	Demand-supply interaction model	72
4.7	Experiments	73
4.7.1	Data preprocessing	73

4.7.2	Cultural patterns extraction for Beijing	73
4.7.3	Refined cultural planning for Beijing	77
4.8	Discussion & conclusions	81
5	Topic-enhanced memory network for POI recommendation	83
5.1	Introduction	83
5.2	Related work	85
5.3	Problem formulation	86
5.4	Topic-enhanced memory networks	87
5.4.1	General framework	88
5.4.2	Memory network	88
5.4.3	Temporal latent Dirichlet allocation	91
5.4.4	Fusion of MN & TLDA	91
5.4.5	Geographical modeling	92
5.4.6	Joint training	93
5.5	Experiments	94
5.5.1	Experimental setup	94
5.5.2	Experimental results	97
5.5.3	Detailed analysis of the proposed model	101
5.6	Discussion & conclusions	105
6	Reflections & outlook	107
6.1	Thesis summary & contributions	107
6.2	Directions for future research	110
6.2.1	Possible future work of the studies involved	110
6.2.2	Future research directions in urban computing	111
	References	113

List of figures

1.1	Typical digital data sources and research topics covered in urban computing.	3
2.1	General framework of urban computing with GSN data.	14
2.2	Geometric comparisons of MF, CML, and LRML recommendation algorithms. MF-based methods calculate the dot product of user and item vectors; CML minimises the Euclidean distance between positive user-item pair, and; LRML learns a relational vector for each user-item pair.	26
3.1	Spatial distribution of Foursquare venues in London. Cultural venues by parent categories are coloured, and other categories of Foursquare venues are in grey.	37
3.2	Map of IMD rank change from 2010 to 2015 for London wards and sub regions.	39
3.3	Initial IMD score, CEA, and CVA of London boroughs in 2010. Each London borough is represented by a circle, the colour and size of which indicate the IMD score of the borough. Yellow means relatively more deprived and purple indicates less deprived. The larger the circle, the higher level of deprivation the borough.	45
3.4	Culture expenditure and Foursquare features changes of London boroughs. Yellow circles represent relatively more deprived boroughs and purple circles stand for well-off ones.	45
3.5	Spatial distribution map of ward groups in London. NA denotes the ward's relative IMD ranking remains unchanged between 2010 and 2015.	46
3.6	Means plot for average clustering coefficient in independent one-way ANOVA analysis. The red dashed line denotes the average level of all the wards. . .	48
3.7	Means plots for variables with statistically significant effects between cultural advantaged and disadvantaged groups in independent one-way ANOVA analysis. Dashed lines denote the average feature values.	49

3.8	Means plots for variables with statistically significant effects in factorial repeated measures ANOVA analysis.	50
3.9	Distribution of the IMD change for London wards.	52
3.10	Evaluation for supervised prediction methods on different ward sets.	53
3.11	Relative importance evaluation for each feature in random forest classification.	54
3.12	Evaluation for supervised prediction methods on different feature sets.	54
4.1	Description of main phases and algorithms utilised in the research.	61
4.2	Temporal cultural check-in distribution of four months in Beijing. The x-axis labels and bars above coloured purple represent weekend days. Grey bars represent rest days (public holidays or weekends).	64
4.3	Graphical model for TLDA.	67
4.4	TCV scores for different number of patterns.	74
4.5	Probabilities of venue categories over patterns.	75
4.6	Cosine similarity between venue categories.	75
4.7	Comparison between TLDA and LDA in CV. Only the coherent value of top venue categories Since the LDA model does not contain the temporal part extended by TLDA, only the venue category clustering is evaluated.	76
4.8	Temporal characteristics of cultural patterns in Beijing.	77
4.9	Results of DSI model. For the first two rows, the darker the color, the higher the demand or supply level. For the DSR, red and blue represent high and low ratios, respectively.	78
4.10	Correlation between DSR and travel distance.	80
4.11	Number of venues in each pattern per million people in Beijing Districts.	81
5.1	The overall architecture of the TEMN. It consists of three modules represented by different colours. A user, a visited POI, a negative sampling POI, and the user's visiting history are taken as input for POI recommendation.	88
5.2	Evaluation of Top- N POI recommendation where N ranges from 1 to 10 on the two WeChat datasets in two scenarios.	100
5.3	Attention weights over memory slices for patterns.	102
5.4	Average geographical influence of POIs and geographical preference of users for patterns.	103
5.5	Effect of weighting parameters in the loss function.	104
5.6	Effect of weighting parameter in score function.	105

List of tables

3.1	Comparison table of time scales for datasets.	37
3.2	Network properties at each snapshot. Number of nodes $ V $, number of edges $ E $, average clustering coefficient $\langle C \rangle$, and average degree $\langle k \rangle$	38
3.3	Description of the variables used in the analyses. P represents preliminary analysis.	43
3.4	Groups of London wards in ANOVA analyses.	46
4.1	Notation and description of parameters in TLDA.	66
5.1	Statistics of the evaluation WeChat datasets.	95
5.2	Performance comparison of different methods in three recommendation scenarios. Best performance is in boldface. We use "*" to mark the best performance from baselines for each comparison and report the improvement ratio of our model over the best baseline performance for each scenario in parentheses.	98
5.3	Pattern and top categories.	102

Chapter 1

Introduction

The past few decades have witnessed rapid urbanisation and the rise of megacities globally (Daggitt et al., 2016; Nations, 2018; Zhang et al., 2017). According to a recently released United Nations report¹, the proportion of the world's population living in urban areas increased from 30% in 1950 to 55% in 2018 and is projected to 68% by 2050. In 1950, there were only two metropolises of 10 million people or above: New York and Tokyo (Sattler and Brandes, 2015). Until 1975, two more cities, Shanghai and Mexico City joined the megacity club. By 2018, however, the number had rocketed to 33, leading to one in eight people worldwide living in a megacity nowadays (Nations, 2018). As centres of civilisation and economic life, metropolises have created enormous economic benefits and opportunities to modern society, but must now cope with the burden of population concentration and face unprecedented challenges, including but not limited to air pollution, excessive energy consumption, traffic congestion, and a wide range of socio-economic issues due to unbalanced development (Altomare et al., 2017; Psyllidis et al., 2015; Weng et al., 2018; Zhang et al., 2017; Zheng et al., 2014a). These side effects attached to the rapid progress of urbanisation have significantly deteriorated the quality of life of urban dwellers (Galbrun et al., 2016) and inhibited the sustainable development of cities (Nations, 2018).

To fundamentally solve the problems haunting cities, a deep understanding of the underlying mechanisms of how cities function is essential. With a particular strength to see into city dynamics, the study of urban mobility offers us unparalleled opportunities to get to the root of these urban issues and come up with corresponding solutions. More specifically, investigating the mobility trends of city dwellers has the potential to tackle various urban challenges by allowing us to detect individual tastes, uncover hidden patterns of human activities, measure the city's pulse, and eventually find possible solutions to transform cities into more efficient, greener and smarter places (Paldino et al., 2015). However, given the

¹<https://esa.un.org/unpd/wup/Publications>

intrinsic complexity of urban settings, as well as the lack of reliable data sources, fine-grained spatio-temporal urban mobility characterisation and prediction seemed nearly impossible years ago. For a long time, urban studies have relied heavily on conventional survey methods and individual interviews, which are costly, time-consuming, and limited in geographical and temporal scope (Zhang et al., 2017). Fortunately, with the advent of the information age, the increasing complexity of the contemporary urban setting is nowadays coupled with massive amounts of digital data and advanced computational techniques, which opens up new horizons for human mobility studies and fosters the recent emergence of urban computing. As a highly interdisciplinary field of research, urban computing links up computational science with a wide range of city-related subjects. It empowers digital data, with the help of ubiquitous computing technologies, to handle urban issues and to significantly improve the quality of life in urban environments (Zheng et al., 2014a).

In the following sections, we will discuss in more detail urban computing (Section 1.1), digital data sources for urban computing (Section 1.2), and current research efforts as well as limitations in this field (Section 1.3), before presenting the thesis (Section 1.4) and contributions (Section 1.5) of this dissertation.

1.1 The rise of urban computing

The term "urban computing" was initially coined by Eric Paulos at UbiComp 2004 (Paulos et al., 2004), where he spearheaded this new direction and proposed that urban computing research mainly focused on addressing issues and concerns embedded within the urban living. Subsequently, it was defined in a more precise manner by Zheng et al. (Zheng et al., 2014a) as a process of acquiring, integrating, and analysing large-scale data derived from heterogeneous sources including human beings, vehicles, and sensors in the context of urban environments, to overcome various challenges cities face.

The emergence of urban computing captures a pivotal moment when ongoing rapid urbanisation comes along with the increasing proliferation of mobile devices and the widespread influence of wireless technologies in urban areas (Foth, 2008). Fuelled by large amounts of data generated from sensor networks, mobile devices as well as social media platforms, urban computing has brought a revolution to urban studies by offering fresh perspectives on the nature of urban phenomena, striving to solve serious urban problems bedevilling contemporary cities, and predicting the future development of urban areas in novel and compelling ways (Kindberg et al., 2007; Psyllidis et al., 2015; Zheng et al., 2014a). Even though still at the early stages, urban computing has risen rapidly in recent years and naturally intersected with a vast array of disciplines, including urban planning, economics, transportation, sociology, and

environmental science (Silva et al., 2019; Zheng et al., 2014a). Through unlocking the power of digital data generated in the urban context and taking advantage of computer-mediated approaches, urban computing has shown its huge potential in addressing complex real-world issues in various aspects of urban life. Figure 1.1 presents some representative digital data sources commonly exploited for urban computing tasks and major areas benefiting from urban computing techniques. In the following sections, more details about these data sources and application areas of urban computing will be introduced.

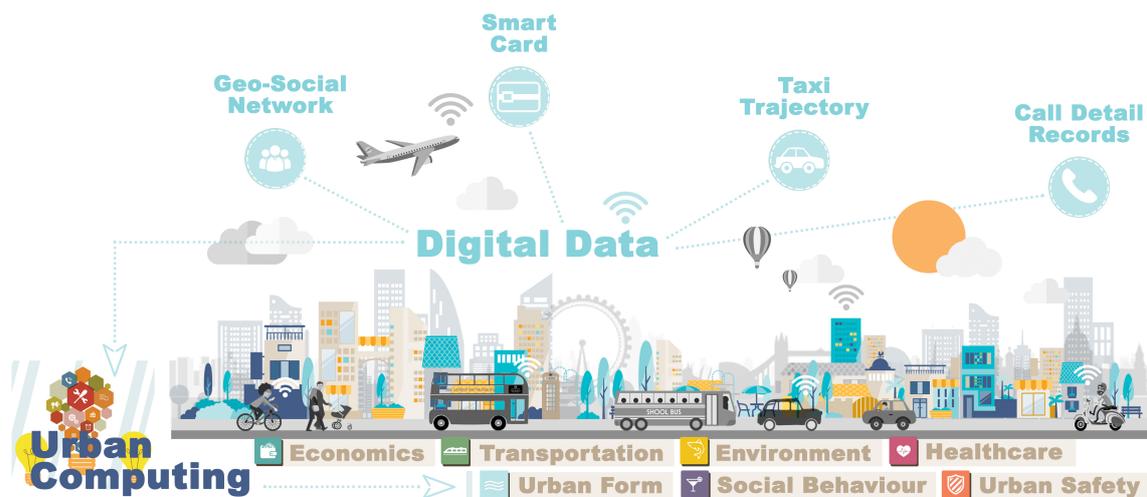


Fig. 1.1 Typical digital data sources and research topics covered in urban computing.

1.2 Digital data for urban computing

For dozens of years, human mobility profiling in cities has traditionally relied on the knowledge discovered from direct investigations (e.g., surveys and face-to-face interviews) or official data (e.g., census and statistics) (Capineri, 2016). Until recently, the rapid advances in information and communication technologies have made citizens act as sensors that produce digital footprints in the urban environment (Haklay et al., 2008). In this context, smart card data, taxi trajectory data, mobile phone call detail record (CDR), and geo-social network (GSN) data emerged as new types of data, providing us with valuable information about human movements in cities, which is unreachable in the past. More specifically, the implementation of contactless smart card systems for electronic ticketing in urban public transport, such as underground, light rail, and bus, allows citizens to swipe a card to pay for the ticket while entering (leaving) a station or getting on (off) a bus. Such ticketing systems generate massive transition records, each of which consists of the check-in station, check-out station,

and timestamps (Zheng et al., 2014a). Also, since mobile GPS devices are broadly installed in cars nowadays, taxis running in cities have generated large amounts of taxi trajectory data. Usually, anonymous taxi identity, timestamp, and geographic coordinates are recorded, enabling the mobility tracking of taxis. Another digital data source contributing to urban studies is CDR, which is created by telephone exchanges and collected by mobile network operators. Each time a mobile phone user makes a phone call, a record is created, providing information about the starting time, its duration, and the caller's and receiver's estimated locations (Zheng et al., 2014a). In addition to the digital data sources mentioned above, geo-social networks (GSNs) provide another particular form of urban mobility data. The research projects presented in this dissertation mainly rely on data from GSNs, which will be introduced and highlighted next.

The increasing availability of location-acquisition and wireless communication technologies (e.g., GPS and Wi-Fi) empowers online social networks to add a location dimension, fostering various geo-social networking services (Armenatzoglou et al., 2013; Bao et al., 2012; Bawa-Cavia, 2011). These GSNs not only allow users to maintain cyber links with one another but also provide platforms for them to share their real-life experiences at specific locations in the physical world in a variety of ways (Ferrari et al., 2011; Lee and Sumiya, 2010; Zhang et al., 2012). For instance, one can easily upload geotagged photos to Flickr², comment on an event happening at a place through Twitter³, check into and rate a restaurant visited on Foursquare⁴, advertise an exact location to her friends in real-time with the help of WeChat⁵, or share jogging routes through Joyrun⁶ (Bahir and Peled, 2013; Gao et al., 2013; Scellato et al., 2010). The provision of these GSN services has encouraged millions of registered users to share their movement behaviours and activity participation in urban spaces (Noulas et al., 2011; Zhi et al., 2016). As a result, massive amounts of user-generated data regarding urban venues emerge on social media sites (Hasan and Ukkusuri, 2015; Hasan et al., 2013; Shen and Karimi, 2016; Silva et al., 2019). To get a taste of how large scale the user base is, consider that Foursquare currently attracts more than 50 million monthly active users creating over 9 million check-ins in a single day on its Swarm app (Foursquare, 2019) and WeChat with 1.08 billion monthly active users in 2018 (WeChat, 2019). Besides, as other types of digital traces introduced previously, GSN data also offer spatio-temporal information associated with users' behaviours on geo-social networking sites. Much like smart card data, taxi trajectory, and CDR, GSN data keep temporal information regarding

²<https://www.flickr.com>

³<https://twitter.com>

⁴<https://foursquare.com>

⁵<https://www.wechat.com>

⁶<https://www.thejoyrun.com>

an event in the form of timestamps. While in the spatial dimension, GSN data provide rich geo-location information (e.g., check-ins at venues, geotagged photos, and comments on places) presented in the forms of precise latitude-longitude coordinates, the name of a nearby landmark, or a symbolic representation like home, office, or shopping mall (Capineri, 2016; Lee and Sumiya, 2010; Zheng et al., 2014a; Zheng and Zhou, 2011). From these additional location-related contents, we are thus able to infer the purpose of a user's visit to a place (Hasan et al., 2013), investigate the lifestyles of individuals (Hasan and Ukkusuri, 2015), and better understand the semantics of urban areas (Silva et al., 2019).

In a nutshell, the recent rise of digital technologies has inspired individuals to share their location preferences in multiple ways. This ubiquitous user-generated information regarding urban spaces has led to a new generation of 'human knowledge' and the development of a better understanding of individual and collective human movements in cities (Shen and Karimi, 2016; Steiger et al., 2016; Wyly, 2014). Various types of digital data introduced above that reflect urban mobility in different layers have significantly contributed to urban computing studies, as will be discussed in the next section.

1.3 Current research efforts & limitations

With the emergence of the Big Data era, urban computing has become one of the subjects that benefit from the increasing availability of large-scale digital data. Through the lens of a wide range of urban mobility data sources, changes in urban places and behaviours of citizens have become unprecedentedly observable at very fine-grained spatial and temporal scales, offering us precious opportunities to measure the urban pulse in nearly real-time (Hristova et al., 2016). This recent trend has promoted the development of urban computing and led to fruitful research in the area. This section will give a glimpse into the current research in urban computing and some limitations that stimulate this dissertation.

Generally speaking, much research in the field of urban computing has been devoted to the resolution of urban issues and the development of intelligent services in cities, including but not limited to the following topics:

Economics. From an economic perspective, Karamshuk et al. (2013) exploited Foursquare data to show the power of geographical features and user mobility signals in the popularity prediction of retail stores. Taking advantage of the categorical information of Foursquare venues, Quercia and Saez (2014) studied the relationship between the presence of different types of physical venues and the deprivation levels in London neighbourhoods. To examine whether there is an association between art and the economic status of urban areas, Seresinhe et al. (2016) utilised geotagged photos on Flickr and found that when there was a higher

proportion of art pictures in the neighbourhoods, the increase in residential property prices was also larger.

Transportation. To understand the city traffic dynamics, An et al. (2016), for instance, developed an approach to measuring the evolution of recurrent urban congestions using mobility data gathered from a GPS-equipped vehicle. From an individual-level perspective, Yuan et al. (2011) built a system that offered personalised driving routes according to traffic conditions and drivers' driving habits learnt from taxi trajectories. To encourage public transport usage, Lathia and Capra (2011) studied the extent that financial incentives implemented by transport authority correlated with changes in users' travel behaviours using smart card data.

Environment. Besides traffic jams, rapid urban development has also brought about a series of environmental issues. Focusing on noise pollution in urban areas, Zheng et al. (2014b) proposed a model to infer the category and level of noise for each region in New York City during a certain period of time by exploiting noise complaint data and GSN data. Contributing to air pollution control, Zheng et al. (2015) presented a system that was capable of forecasting fine-grained air quality using data from air quality monitoring stations.

Healthcare. The wellbeing and physical health of residents are also hot topics in the field of urban computing. For instance, to investigate the relationships between beautiful surroundings and happiness, Seresinhe et al. (2019) drew on individual happiness data from a smartphone app and scenic ratings of geotagged photos. The researchers found that people tended to be happier in more scenic places, regardless of whether they were natural or built-up areas. While focusing on the monitoring of physical health, Lima et al. (2015) presented a model that described how diseases circulated as people moving across urban regions by exploiting cellular network data.

Urban form. From the view of spatial structure, Daggitt et al. (2016) utilised Foursquare data to analyse the urban growth patterns in 100 major cities worldwide. They revealed the existence of a strong spatial correlation that nearby pairs of cities are more likely to share similar growth profiles than remote ones. Bawa-Cavia (2011) also employed large-scale Foursquare data and proposed comparative measures of polycentricity and fragmentation to discuss the spatial structure of three megacities: London, New York, and Paris.

Social behaviour. By considering the social ties of individuals, De Domenico et al. (2013) showed that it is possible to exploit friends' CDR data to forecast user movement. They also discussed how human movement correlation linked to social interactions by investigating the co-location of individuals and the number of phone calls between them. Leveraging Foursquare check-ins, Hristova et al. (2016) defined several metrics to measure the social

diversity of places and captured the mobility patterns of their visitors. The authors further discovered that the features extracted can be used as signals for gentrification in urban areas.

Urban safety. Given that cities are the places where people congregate and plenty of events happen, crowd management in public venues is crucial to urban safety. In this line of research, Botta et al. (2019) exploited Instagram data to estimate the size of crowds attracted by football matches. Crime is another major problem that threatens public safety. In order to infer the crime rate at neighbourhood level, Wang et al. (2016) drew on datasets of taxi flows and Foursquare venues in Chicago and showed that incorporating new features from digital data could significantly enhance the crime rate prediction than using traditional features alone.

As urban computing is highly interdisciplinary and comprehensive, only some striking findings in representative publications are introduced in this section. After reviewing the existing literature in the domain of urban computing, we find that although growing numbers of high-quality publications have been offered and shed light on multiple urban challenges, there are still some missing puzzle pieces. For instance, the potential of network science in giving insight into socio-economic dynamics along with the large-scale human mobility traces made available by GSN applications has so far been largely untapped in culture-led urban regeneration studies. Also, the power of tracking millions of interactions of mobile users with urban venues across space and time in GSNs to devise fine-grained indicators of urban areas where there is a lack of certain types of resources has not been well explored. Besides, for urban computing research, we advocate the adoption of more advanced algorithms, the design of novel algorithms tailed for urban settings, and the development of more sophisticated applications. In response to this, we discuss the building of personalised recommender systems that can capture comprehensive user preferences for urban venue recommendation with high predictive accuracy and interpretability in this dissertation. Generally speaking, this thesis advances the existing body of knowledge in urban computing by broadening its research scope, devising a series of innovative techniques, and developing more sophisticated mobile applications in urban environments.

1.4 Thesis & substantiation

As discussed in the previous sections, the ubiquitous usage of GSNs has generated a large amount of urban mobility data with unprecedentedly high spatio-temporal resolution. Consequently, GSNs have facilitated the recent rise of urban computing and provided unprecedented opportunities to foster algorithmic approaches to tackling complex urban issues in compelling ways. In such a context, this dissertation seeks to advance the field of urban computing by leveraging large-scale GSNs data and proposing innovative techniques to offer new insights

into human mobility patterns in time and space. It leads to a better understanding of urban development and more intelligent mobile applications.

The central thesis of this dissertation is that *the exploration of geo-social network data with various analytical techniques can advance our understanding of the nature of urban phenomena and the development of intelligent urban computing applications.*

We substantiate this statement with three closely related research themes. First, we explore how the collective mobility patterns and local network features can be used as a proxy to track and predict socio-economic deprivation changes as the government putting financial effort in developing areas. Next, we extract patterns of urban cultural interactions to detect the real-world demand and supply levels of various cultural resources for further cultural infrastructure planning. Then from a user's perspective, a personalised recommender system is developed to recommend new places for the user to explore in the city based on her preferences detected from check-in records left on geo-social networking sites. The next section will detail the contributions of this dissertation and show how they link to the research themes raised here and their corresponding chapters.

1.5 Contributions & chapter outline

This dissertation contributes to the field of urban computing through the development and assessment of novel approaches for human mobility modeling to tackle complex urban development challenges. We shall start with an introduction of the framework for urban computing and commonly used techniques in Chapter 2, before offering three main contributions that map to the later chapters in the rest of the dissertation as follows:

Contribution 1: Evaluation of investment in culture-led regeneration

In Chapter 3, we study how to leverage GSN data to assess the impact of government-led investment on socio-economic indicators in cities. Statistical analysis, network analysis, and supervised learning algorithms are employed to track the impact and measure the effectiveness of cultural investment in small urban areas. More specifically, taking advantage of nearly 4 million transition records for three years in London from a popular location-based social network service, Foursquare, we study how the socio-economic impact of government cultural expenditure in London neighbourhoods can be detected and predicted. Our analysis shows that network indicators such as the average clustering coefficient or centrality can be exploited to estimate the likelihood of local growth in response to cultural investment. We subsequently integrate these features in supervised learning models to infer socio-economic deprivation changes for London's neighbourhoods. This research presents how geo-social

and mobile services can be used as a proxy to track and predict socio-economic deprivation changes as government financial effort is put in developing urban areas and thus gives evidence and suggestions for further policy-making and investment optimisation.

Contribution 2: Optimal resource allocation for urban development

In Chapter 4, we propose a data-driven framework for optimal allocation of cultural establishments and related resources across urban regions to reduce financial costs in terms of planning and improve quality of life in the city, more generally. We make use of a large longitudinal dataset of user location check-ins in the city of Beijing from the online social network WeChat. We exploit rich spatio-temporal representations on user activity at cultural venues and use a novel extended version of the traditional latent Dirichlet allocation model (Blei et al., 2003) that incorporates temporal information to identify latent patterns of urban cultural interactions. To evaluate the performance of our proposed model and to determine the optimal hyperparameter settings, an evaluation approach for topic modeling is also devised. Then using the characteristic typologies of mobile user cultural activities emitted by the model, we determine the levels of demand for different types of cultural resources across urban areas. We then compare those with the corresponding levels of supply as driven by the presence and spatial reach of cultural venues in local areas to obtain high-resolution maps that indicate urban regions with a lack of cultural resources, and thus give suggestions for further urban cultural planning and investment optimisation.

Contribution 3: New place recommendation in the city

In Chapter 5, we detail the design and evaluation of a novel point-of-interest (POI) recommender system, that is able to capture complicated user preferences and fine-grained user-POI relationship. Technically, we propose a novel topic-enhanced memory network, a deep architecture to integrate the topic model and memory network capitalising on the strengths of both the global structure of latent patterns and local neighbourhood-based features in a non-linear fashion. Specifically, temporal latent Dirichlet allocation is employed as the topic module to capture user intrinsic preference considering temporal effects, while the memory-augmented neural network is utilised to learn complex user-POI interactions more explicitly and dynamically. We further incorporate a geographical module to exploit user-specific spatial preference and POI-specific spatial influence to enhance recommendations. The proposed unified hybrid model is widely applicable to various POI recommendation scenarios. Extensive experiments on real-world WeChat datasets demonstrate the effectiveness of the proposed model, that it offers improvement ratios of 3.25% and 29.95% for the context-aware and sequential recommendation, respectively. Moreover, qualitative visualisation of the attention weights

and topic modeling provide insight into the proposed model’s recommendation process and results.

Finally, we review the contributions of our research and draw conclusions, as well as outlining how future research may build on the work we have described in this dissertation in Chapter 6.

1.6 List of publications

During my Ph.D. studies, I have engaged in several fruitful collaborations with both academic and industry co-authors, which have yielded publications mainly focusing on the application of machine learning and deep learning algorithms, as well as the development of innovative techniques in the areas of spatio-temporal modeling of human mobility, local development prediction, and personalised recommender systems. Some of the publications are directly related to this dissertation. In particular, Chapter 3 draws from a study (Zhou et al., 2017) published in the journal *Royal Society Open Science*, Chapter 4 is based on an accepted paper (Zhou et al., 2018b) at the Conference on Knowledge Discovery and Data Mining (KDD), and Chapter 5 builds on a recently published paper at KDD 2019 (Zhou et al., 2019). Beyond that, some other works that are not fully covered by this dissertation are also listed below.

Works related to this dissertation

- **Xiao Zhou**, Desislava Hristova, Anastasios Noulas, Cecilia Mascolo. Detecting socio-economic impact of cultural investment through geo-social network analysis. In *Proceedings of the 11th International AAAI Conference on Web and Social Media (ICWSM 2017)*.
- **Xiao Zhou**, Desislava Hristova, Anastasios Noulas, Cecilia Mascolo, Max Sklar. Cultural investment and urban socio-economic development: a geo-social network approach. *Royal Society Open Science*, 4(9):170413, 2017.
- **Xiao Zhou**, Anastasios Noulas, Cecilia Mascolo, Zhongxiang Zhao. Discovering latent patterns of urban cultural interactions in WeChat for modern city planning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2018)*.
- **Xiao Zhou**, Cecilia Mascolo, Zhongxiang Zhao. Topic-enhanced memory networks for personalised point-of-interest recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2019)*.

Other works during PhD study

- **Xiao Zhou**, Desislava Hristova, Anastasios Noulas, Cecilia Mascolo. Evaluating the impact of the 2012 Olympic Games policy on the regeneration of east London using spatio-temporal Big Data. In *Proceedings of International Conference on Frontiers of Data Science for Government: Ideas, Practices, and Projections* (Data for Policy 2016).
- Yingzi Wang, **Xiao Zhou**, Anastasios Noulas, Cecilia Mascolo, Xing Xie, Enhong Chen. Predicting the spatio-temporal evolution of chronic diseases in population with human mobility data. In *Proceedings of 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence* (IJCAI 2018).
- **Xiao Zhou**, Danyang Liu, Jianxun Lian, Xing Xie. Collaborative metric learning with memory network for multi-relational recommender systems. In *Proceedings of 28th International Joint Conference on Artificial Intelligence* (IJCAI 2019).

I acknowledge with gratitude the help and advice of my co-authors in producing the papers listed above and point out that while I made use of their expertise and suggestions, the results and technical contributions outlined in this dissertation are my work, conducted towards the exploration of the thesis in the following chapters.

Chapter 2

Urban computing with GSN data

The previous chapter introduces urban computing, as a highly data-driven and multidisciplinary domain, aimed at solving real-world problems in urban settings by invoking the power of digital information. This chapter will offer a more detailed description of this process and an introduction to relevant techniques to implement it. Here we exemplify the most commonly used techniques in the domain of urban computing and their application scenarios accordingly. We also introduce the background knowledge of some basic versions of the algorithms utilised in this dissertation and advocate that new algorithms tailored for urban computing scenarios should be developed. Some opportunities and key limitations of existing approaches in urban computing are also outlined at the end of this chapter. As GSN information for urban computing constitutes a relatively recent field of investigation and shows some distinctive characteristics compared with other types of digital data, this dissertation focuses more on urban computing with GSN data.

Chapter outline. In Section 2.1, an overall framework of urban computing with GSN data is presented. Next, some techniques involved in each phase of the framework will be introduced accordingly in the following sections. More specifically, in Section 2.2, we discuss essential data preprocessing techniques and highlight the significance of integrating multiple urban data sources to tackle complex urban issues. Section 2.3 introduces a variety of typical machine learning algorithms used to model human mobility patterns for urban computing. Section 2.4 shows a graphing approach that maps users' movements between urban places. It further presents concrete examples of how this method can be used to explore city dynamics. Section 2.5 gives a background on the venue recommendation problem in GSN services and discusses its relationship to human mobility in urban environments. In Section 2.6, we broadly discuss main challenges and future work for research in the domain of urban computing leveraging GSN data. Finally, we summarise the chapter in Section 2.7.

2.1 Urban computing framework with GSN data

Urban computing leveraging GSN data focuses on using advanced modeling techniques to extract valuable knowledge for deeper insights into urban issues and to support better services in cities in various fields. Figure 2.1 presents a general framework of urban computing with GSN data, which is composed of three main components: (i) data preprocessing, (ii) data analytics, and (iii) application development.

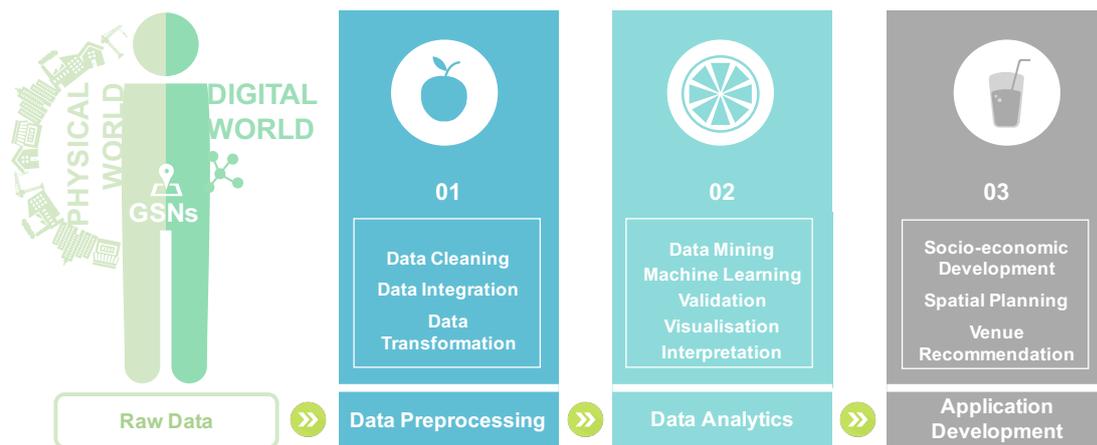


Fig. 2.1 General framework of urban computing with GSN data.

In the data preprocessing step, raw GSN data are firstly cleaned, integrated, and consolidated into forms appropriate for mining in the following data analytics stage. Intelligent approaches are then applied to extract meaningful patterns from the structured data. Typical techniques adopted in this process include statistical methods, machine learning algorithms, and information retrieval models. Appropriate modeling, validation, and visualisation approaches are selected according to data forms and research goals. Based on the knowledge gained from data analytics, related services and applications are eventually developed.

Following such a general framework, different urban computing tasks can be fulfilled by mining different GSN data with various preprocessing and analytical techniques. Below we will discuss some of the techniques commonly used in each step in greater detail. Furthermore, the necessary background knowledge of basic algorithms employed in this dissertation will be provided as well. Some of the studies are discussed to present a more concrete illustration of the components of the urban computing framework with GSN data.

2.2 Urban data preprocessing

Different GSN application sites create data that have various types, versatile structures, and broad semantic meanings. Typical forms of GSN data include sequence data (e.g., historical visiting records), graph data (e.g., social networks), and categorical data (e.g., categorical information of venues). Given the diversity of data from GSNs and the goals of urban computing, how to handle complex GSN data with rich spatio-temporal and semantic information properly has become a challenging task.

2.2.1 Data cleaning

Ensuring the quality of the data is essential for the ongoing success of the urban computing project. The first step to deal with raw GSN data is always data cleaning. In the context of data science, data cleaning refers to the processes of filtering and modifying the data so that it is easier to explore, understand, and model. Even though the specific steps and techniques for data cleaning may vary from dataset to dataset, its general objective is to make the datasets consistent and free of errors that could be problematic for data analysis. Some common tasks covered by data cleaning are fixing structural errors, dealing with missing data, and filtering irrelevant observations. Some statistical analysis and visualisation techniques, such as standard deviation, normalising, and histogram, can be employed to assist in cleaning large-scale GSN data.

Through effective data cleaning, processed GSN data should be well organised and stored by an indexing structure that can simultaneously incorporate spatio-temporal and semantic information for efficient data analytics later. In some cases, GSN data can also be transformed into place networks for graph-based analysis.

2.2.2 Data integration

Even though GSN data can reveal powerful insights into human mobility patterns in urban areas, in consideration of the complex urban issues, sometimes it is necessary to harness a diversity of data sources in a single urban computing task. As data is becoming increasingly available, various open datasets released by governments and private companies have begun to play an important role in the domain of urban computing. Data.gov¹ and Data.gov.uk² are two well-known government open data sites that offer a wide range of public sector data, ranging from demographic statistics to datasets from various government departments, such as trade

¹<https://www.data.gov>

²<https://ckan.publishing.service.gov.uk>

data in goods and services, health status, traffic statistics, and crime figures. According to which area we aim to contribute to, different urban computing tasks can be fulfilled by combining information from GSNs with other relevant data sources.

The mining of diverse urban data sources often leads to more fruitful findings and the development of more sophisticated applications due to the mutual enhancement and consolidation of such multiple sources. On the other hand, it also makes urban computing tasks more challenging. Data integration has, therefore, become a crucial phase to combine heterogeneous data sources at different levels and lay the foundations for data analytics.

2.3 Machine learning for urban computing

After data cleaning and integration, the processed data is ready for analysis. Next, the selection of appropriate methods is particularly vital, which depends not only on data forms but also on research purposes. For instance, according to whether the class label information is given or not, supervised or unsupervised machine learning methods can be employed. Moreover, with different purposes, urban computing tasks can be generally classified into descriptive tasks and predictive tasks. Here descriptive models characterise properties of the data to reveal certain types of urban phenomena, while predictive tasks perform induction on the current data to make predictions for future decision making. Besides the above considerations, it should also be noticed that even though rich GSN data provide fertile ground for urban computing, some basic machine learning techniques are not up to the complex urban computing tasks. Extended versions or novel algorithms tailored to urban issues should be developed. In this section, some commonly used machine learning algorithms for urban computing will be introduced, before specific examples are given to illustrate how they can help solve real-world urban problems.

2.3.1 Classification & prediction

In machine learning, classification is a supervised learning approach aiming to find a model that can predict the class label of new observation based on the analysis of training data for which the class labels are known (Han et al., 2011). A classification task can be fulfilled by various types of classification algorithms, such as decision tree, random forest, logistic regression, and naïve Bayes. Next, we describe some commonly used classification algorithms before introducing how they can be adopted for prediction tasks in urban computing.

Typical classification algorithms

Decision tree. The first classification technique we would like to introduce is the decision tree. As one of the most commonly used supervised learning algorithms, the decision tree utilises a flowchart-like tree structure that can automatically choose the splitting variables and split points to achieve the best fit (Hastie et al., 2009). In a decision tree, the topmost node is known as the root node; each non-leaf node in the tree denotes a test on a certain attribute; each branch represents a possible decision, and each leaf node corresponds to a class label. Given a tuple, for which the associated class label is unknown, the attribute values of the tuple are tested against the decision tree. Then a path can be traced from the root to a leaf node that holds the class prediction for the tuple. There are multiple reasons why decision tree models are popular. As the feature space partition is fully described by a single tree, a key advantage of the decision tree is its interpretability. Also, decision tree models can handle mixed discrete and continuous inputs, deal with missing values, and scale well to large datasets. Besides, they are insensitive to monotone transformations of the input and relatively robust to outliers (Hastie et al., 2009; Murphy, 2012). Despite the aforementioned advantages, decision trees are high variance estimators that a slight change to the input data can have a big impact on the structure of the tree, making them unstable (Murphy, 2012).

Random forest. One way to reduce the variance of an estimate is to average together many estimates (Murphy, 2012). Building upon such an idea, Breiman (2001) proposed the random forest algorithm that can reduce the variance of individual decision trees by building a collection of de-correlated trees on random subsets of the training data and averaging the predictions of the trees (Hastie et al., 2009). More specifically, to generate individual decision trees, random forest algorithm adopts a random selection of attributes at each non-leaf node to make the decision. This mechanism leads to the trees generated having different feature subsets and reduces the variance of the individual decision trees. When used for classification, the committee of trees each casts a vote for the predicted class so that a random forest can conduct the final prediction using the majority vote. Random forest generally has high prediction accuracy with little tuning required (Hastie et al., 2009). It also offers a variable importance plot as an output to visualise the relative importance of features for the classification (see Figure 3.11 for an example). A primary disadvantage of random forest classifiers is that the results obtained are not easily interpretable (Murphy, 2012).

Naïve Bayes. Naïve Bayes is a statistical classifier based on Bayes' theorem. It can be employed to predict the probability that a given tuple belongs to a particular class. The most distinguishing characteristic of the naïve Bayes algorithm is that it assumes the features

are conditionally independent given the class label (Murphy, 2012). By making such an assumption, the computations involved are greatly simplified, which also results in the algorithm “naïve”. Simple as naïve Bayes is, it usually does surprisingly well and outperforms far more sophisticated alternatives (Domingos and Pazzani, 1997; Hastie et al., 2009). One possible reason for this is that the naïve Bayes model is simple enough and thus relatively immune to overfitting (Murphy, 2012). Furthermore, when it is applied to large databases, the naïve Bayes classifier has also exhibited high accuracy and speed (Han et al., 2011).

Logistic regression. Logistic regression is one of the most popular algorithms used for binary classification (classification problems with two possible outcomes). It is known as a linear classifier that arises from the desire to model the posterior probabilities of the classes scaled between zero and one, while at the same time ensuring that they sum to one (Hastie et al., 2009). Mathematically, the resulting posterior class probabilities are typically given by a softmax transformation acting on a linear function of the input variables (Bishop, 2006). Logistic regression performs well when the dataset is linearly separable and provides some hints about understanding the role of the input variables in explaining the outcome (Hastie et al., 2009). However, a simple logistic regression algorithm can be prone to overfitting, especially when the number of features is large.

Applications of classification in urban computing

To infer venue categories, Falcone et al. (2014) proposed a supervised learning framework based on spatio-temporal features extracted from geotagged tweets. They defined the problem as a binary classification task and employed six classifiers to predict whether a venue belongs to a certain category or not. To reveal how GSN data can be utilised to predict socio-economic deprivation of urban areas, Venerandi et al. (2015) explored a Foursquare dataset and proposed a metric to build classifiers. The experimental results they presented show that naïve Bayes offered better performance than decision trees and logistic regression. In addition, the exploration of GSN data can also play a role in the studies of public health. In this line of research, De Choudhury et al. (2016) focused on diet-related health issues and utilised posts shared on Instagram to predict whether an urban area has limited access to healthy and affordable food. In another study on public health, Gomide et al. (2011) showed that data from Twitter is helpful in disease surveillance and monitoring in urban areas. More specifically, they built a classifier to model and forecast dengue epidemics based on the knowledge extracted from Twitter messages mentioning Dengue. Event detection is another application of classification techniques in urban computing leveraging GSN data. Thanks to the real-time nature of information offered by GSNs, event identification and

tracking in urban areas become possible. This characteristic is particularly beneficial to disaster prevention and reduction in cities. For instance, Matsu et al. (2010) investigated how microblogging data from Twitter can be used to detect earthquakes. To enable this, they devised a classifier of tweets taking advantage of features such as the keywords, number of words in a tweet, and its content. They further established a model that can estimate the centre of an earthquake and developed an earthquake reporting system in Japan.

2.3.2 Cluster analysis of mobility patterns

Unlike classification that handles data with class labels, cluster analysis, as a canonical example of unsupervised learning, deals with unlabelled data. The goal of clustering is to group a set of data objects based on the principle that maximising the intra-group similarity and minimising the inter-group similarity (Han et al., 2011; Murphy, 2012). Although the analysis objects for clustering can come in many forms, in urban computing, clustering algorithms are typically employed to discover clusters of geo-locations spatially in a two-dimensional space. In this subsection, some clustering techniques commonly utilised in the domain of urban computing are introduced before examples are given to illustrate how these techniques can provide clustering solutions to issues confronting cities.

Typical clustering algorithms

One of the simplest and most fundamental cluster analyses is partitioning clustering, which decomposes the data objects into several exclusive clusters. This type of clustering method requires a user-specified number of clusters and adopts an iterative relocation technique to optimise the partitioning by moving the objects from one cluster to another (Poole and Mackworth, 2010). The k-means algorithm (Hartigan and Wong, 1979), one of the most well-known partitioning clustering methods, will be introduced first.

k-means. As its name implies, the k-means clustering defines the average of all data points within a cluster as its centroid. This algorithm starts by randomly selecting k data objects, representing initial cluster centroids. Then each remaining data point is assigned to the closest cluster according to the Euclidean distance between the data point and the cluster centroid. Next, the k-means algorithm iteratively optimises the positions of the centroids by computing the new mean of the data points assigned to each cluster, before reassigning the objects using the updated cluster centres. These two phases of re-assigning data points to clusters and re-computing the cluster centroids are repeated until there is no change in the assignment of the data points (Bishop, 2006). Although the k-means algorithm works fine on spherical-shaped clusters, it is not suitable to detect clusters in non-convex shapes or clusters

with varying sizes. Besides, as squared Euclidean distance is chosen as the dissimilarity measure in k-means clustering, data variables are limited to the quantitative type (Hastie et al., 2009; Poole and Mackworth, 2010). Moreover, this mechanism also makes the algorithm sensitive to outliers and noise points (Bishop, 2006; Han et al., 2011).

By clustering data points based on distance, most partitioning clustering algorithms encounter difficulty when dealing with arbitrary-shaped clusters. To discover clusters of non-spherical shape, alternatively, density-based clustering methods can be utilised. The main strategy behind such methods is to grow a given cluster by examining whether the density in its neighbourhood is larger than a certain threshold. Besides enabling us to find clusters of arbitrary shapes, such a method can also be employed to filter out noise points. Two representative density-based clustering methods, namely, DBSCAN and OPTICS, will be introduced in the following paragraphs before their applications in urban computing.

DBSCAN. The density-based spatial clustering of applications with noise (DBSCAN) (Ester et al., 1996) is a well-known density-based clustering algorithm that is capable of identifying arbitrary-shaped clusters and is robust against noise and outliers (Oyelade et al., 2016). The key task of DBSCAN is to discover core data objects with high-density neighbourhoods. These core objects and their neighbourhoods collectively shape dense regions as clusters in DBSCAN. To this end, two input parameters are required. Firstly, for a given data point, the radius considered to define its neighbourhood should be specified. Then a density threshold is needed to determine whether a neighbourhood is dense or not. Since the neighbourhood for each data point is defined as space within a fixed radius, the neighbourhood density can be measured simply by counting the number of data points it covers. Thus in DBSCAN clustering, a data object is considered as a core object if its neighbourhood of a given radius contains at least a certain number of data objects.

While DBSCAN is a powerful tool for separating high-density clusters from low-density regions, it has troubles with clusters of varying densities. Moreover, although DBSCAN liberates us from specifying the number of clusters as opposed to k-means, it leaves the users responsible for setting appropriate density parameters to form acceptable clusters. In practice, however, it is generally rather difficult to determine these parameters empirically since real-world data are often high-dimensional with skewed distributions. In such cases, a global parameter setting in DBSCAN would hardly capture the intrinsic clustering structures well (Han et al., 2011).

OPTICS. Ordering points to identify the clustering structure (OPTICS) (Ankerst et al., 1999) can be seen as an extension of the DBSCAN method that is less sensitive to parameter settings. More specifically, to overcome the limitation of using one set of global parameters in the

DBSCAN clustering analysis, the OPTICS algorithm reveals the intrinsic clustering structure instead of producing a clustering of the dataset explicitly. To do so, OPTICS extracts two extra pieces of information, namely core distance and reachability distance from the datasets and outputs a cluster ordering. This special clustering ordering produced is a list of all the data points under analysis where objects in a denser cluster are presented closer to each other. It can be used to derive some basic clustering information, including cluster centres and a density-based clustering structure. Furthermore, such an ingenious design of cluster ordering in OPTICS allows it to offer a visualisation of the clustering using a reachability plot.

Applications of clustering in urban computing

Various types of clustering techniques have been applied to problems in urban computing domain. In this direction, Le Falher et al. (2015) focused on how to quantify the similarity of neighbourhoods in European and American cities taking advantage of the k-means algorithm. In this regard, the authors utilised rich information extracted from Foursquare check-ins and clustered the venues in each city according to the temporal distribution of associated check-ins in a day. In another paper using the k-means algorithm, Frias-Martinez et al. (2012) revealed clustering patterns of twitter behaviours to determine land uses in urban areas. Their results demonstrate that geo-located tweets could serve as a powerful data source to characterise urban areas, reaching competitive effects with costly conventional approaches. In a similar direction, Karina et al. (2015) proposed a framework to discover the functional use of urban areas. Instead of employing the k-means method, the authors used the DBSCAN algorithm to cluster and label urban areas by exploiting Foursquare check-in data. Zhang and Pelechris (2014) investigated the phenomenon of homophily implied in human visiting patterns. They used GSN data collected from Gowalla and revealed that peer influence tends to play a role when friends are in proximity to each other. Additionally, depending on the types of venues to visit, the degree of influence may vary greatly. In this work, the DBSCAN clustering algorithm is applied to the historical check-ins of each user to infer her home location. Contributing to the area of event detection, Arcaini et al. (2016) extended the DBSCAN to uncover spatio-temporal aperiodic and periodic characteristics of events occurring at certain locations by using geo-tagged Twitter messages. Also by exploring geo-tagged tweets, Falcone et al. (2014) proposed an approach to identify venue categories in the city. Their promising experimental results illustrate that categorical information is crucial to a better understanding of human mobility patterns. And the spatial clustering technique used to group locations in this paper is OPTICS.

2.3.3 Topic modeling & urban semantics

To obtain a deeper understanding of urban issues, taking advantage of the rich semantic information contained in GSN data, and integrating new techniques from other domains have become a trend. For instance, text data and categorical information related to urban places in GSNs can be analysed to learn city semantics by fusing traditional data mining approaches with information retrieval and natural language processing techniques. A typical technique applicable to such cases is topic modeling, a tool commonly used for discovering hidden semantic structure in text data. Through topic modeling techniques, major topics involved in a collection of documents can be discovered, which substantially enhances the power of GSN data for urban semantics studies. The following paragraph will discuss some primary studies on urban semantics investigation with GSN data and topic modeling.

In this direction, Joseph et al. (2012) utilised the topic modeling technique of Latent Dirichlet Allocation (LDA) (Blei et al., 2003) to group users based on their Foursquare check-ins. To apply LDA to user check-in behaviours modeling, the authors regard each user as a document. Then each check-in record in a user's visiting history can be viewed as a word in the document. By doing this, the authors obtained several latent topics that each can be depicted by a set of venues and identified groups of users accordingly. Also with the LDA technique, Long et al. (2012) explored the intrinsic relations existing between urban places through analysing sequential check-in data collected from Foursquare. Their basic assumption is that venues appearing together frequently in check-in sequences may represent the topic of urban areas. Unlike the above papers that view venues as words in documents, some other studies adopting topic modeling approaches more conventionally, analysing text data posted on GSNs. For instance, Schwartz et al. (2013) showed that topics derived from text messages on Twitter mapped to U.S. counties using LDA can improve the prediction of life satisfaction of users. Similarly, Paul and Dredze (2011) applied topic modeling to health-related tweets and showed the broad applicability of Twitter data for public health studies.

2.4 Mining graphs & place networks

As non-linear data structures containing rich information, graphs have been successfully applied in a wide range of domains, such as bioinformatics, computer vision, and text retrieval. A network is a collection of connected objects. The objects correspond to mathematical abstractions called nodes, which are connected by a set of edges, representing relationships between these objects (Han et al., 2011). By building an urban transition graph to map user movements between venues, graph mining techniques can play a role in city dynamics

studies. More specifically, the nodes in this type of particular graph are the city venues, and an edge exists between a pair of venues if a user created a check-in at one place right after visiting the other. If necessary, it can be built as a directed weighted graph, where the weight of an edge is the total number of movements observed from one node to the other. Next, some works will be introduced to exemplify the applications of graph mining in urban computing.

Through exploring check-ins from Brightkite using network analysis methods, Jiang and Miao (2015) investigated the underlying evolving mechanisms of cities. Their study showed that GSN data can serve as a proxy to track and understand the changes in cities at much finer spatio-temporal scales than conventional data sources. Hristova et al. (2016) proposed a method to capture the social diversity of urban locations from a network perspective. To accomplish this goal, the authors utilised both a spatial network of Foursquare venues and social network information from Twitter to gain insights into the social roles of urban places. Instead of uncovering the social meanings carried by urban places, Ferreira et al. (2015) looked at the temporal attributes contained in the transition graph. The authors studied the mobility of tourists using their check-in records on Foursquare and proposed an approach to detect the time when sights in the city are popular. A fascinating implication of their research is that it can benefit venue recommendation services, the topic of which will be discussed in the following section.

2.5 Metric learning-based venue recommendation

Based on urban mobility knowledge gained through GSN data analysis, human movements across venues are somehow predictable, leading to a classical recommendation problem, point-of-interest (POI) recommendation for mobile users. With the ability to identify a set of venues that a given user would be more likely to enjoy from millions of places in the city, POI recommender systems have played an increasingly vital role in modern society. Concretely speaking, they enrich urban life for individuals and promote business objectives for online platforms. As efficient information filtering tools, recommender systems have gone through several generations of algorithms in the past few decades. Depending on different types of objects recommended, researchers also have developed variants of classical recommendation algorithms accordingly. For instance, when recommending new urban venues to users, spatio-temporal information in the physical world should be fully considered in the design of algorithms for more accurate recommendations. In Chapter 5, we propose a deep learning framework that captures fine-grained user preferences for personalised POI recommendation. As our model proposed is highly inspired by the recent advances in metric

learning for collaborative filtering, we will introduce the evolution of relevant techniques as background information in this section.

Collaborative filtering (CF) is a widely used class of recommendation algorithms that recommend items to a target user by compiling preferences from other users in the system (Ekstrand et al., 2011; Linden et al., 2003; Resnick et al., 1994). The underlying assumption behind CF is that if the target user u_1 has the same opinion as another user u_2 on an item v_1 , user u_1 's preference on another item v_2 would be more likely to be the same as that of user u_2 than a randomly chosen user. Over the past decade, the matrix factorization (MF) technique (Koren et al., 2009) has become one of the most well-known CF algorithms and has spurred a large number of variations (Agarwal and Chen, 2010; He et al., 2016; Hu et al., 2008; Koren, 2008; Koren et al., 2009; Rendle et al., 2009). The original MF models the explicit user feedback (i.e., ratings) by mapping users U and items V to a latent factor space, so that unobserved user-item interactions can be predicted through the dot product of their latent factor vectors. To formalise it, let p denote the latent factor vector of user $u \in \mathcal{U}$, and let q denote the latent factor vector of item $v \in \mathcal{V}$. MF method scores user-item pairs with $s(u, v) = p \cdot q$. In doing so, MF can generate relatively accurate recommendations for users by capturing the prominent characteristics of users and items (Koren and Bell, 2015). However, the adoption of the dot product in the MF approach makes it difficult to depict the user-item relationship in a more fine-grained way or capture user-user and item-item similarities. This is essentially due to the fact that the dot product does not satisfy the triangle inequality, resulting in the positive relationships between user-item pairs hardly be propagated to user-user or item-item pairs in the recommender system (Hsieh et al., 2017).

Considering the limitations of the MF approach, Hsieh et al. (2017) applied metric learning techniques to CF, named collaborative metric learning (CML) algorithm for recommender systems. The fundamental distinction between CML and conventional MF is that CML is a metric learning approach that involves the measurement of the distance between each pair of user and item, while MF is not. In metric learning, given a set of objects among which pairs of objects are similar or dissimilar, the core idea is to learn a distance metric that assigns smaller distances to relatively more similar pairs, and larger distances to dissimilar ones. Following such a pattern, CML learns the user vector p of a user u and the item vector q of item v by calculating and minimising the Euclidean distance between them if a positive interaction exists through the scoring function: $s(u, v) = \|p - q\|_2^2$. Mathematically, such a distance metric is required to satisfy the triangle inequality condition (Kulis et al., 2013; Xing et al., 2003) that for three objects, the sum of any two pairwise distances should be no less than the remaining pairwise distance. As a metric learning scheme, CML obeys the triangle inequality, which makes it possible to propagate the observed user-item similarity

information to the unknown user-user and item-item relationships. Thus, CML can learn the underlying user-user and item-item similarities simultaneously (Hsieh et al., 2017; Tay et al., 2018). Compared to MF-based models, the CML algorithm achieved highly competitive performance, proving the feasibility and effectiveness of metric learning techniques for recommender systems (Chen et al., 2012; Garcia-Duran et al., 2018; He et al., 2017b; Kang et al., 2018).

Despite its improvements over the MF method, CML is not without flaws. One of its weaknesses is that the scoring function of CML is geometrically restrictive (Tay et al., 2018). To be specific, for each given user-item pair that has positive interactions, the objective function of CML tries to fit the user-item pair into the same point in vector space by optimising for $\|p - q\|_2 \approx 0$. Such a feature of CML may pose problems for more fine-grained recommendations, particularly when the dataset is large. Considering that popular GSNs usually have millions of users and items as introduced in Section 1.2, to get a good fit for each user and all her interacted items onto a single point in vector space can be very challenging. Apart from this, CML is an ill-posed algebraic system (Tikhonov and Arsenin, 1977) from a more theoretically grounded perspective, which further limits the geometric flexibility of the method. Relevant proof of this can be found in Tay et al. (2018).

In an attempt to solve the problems with CML, Tay et al. (2018) proposed Latent Relational Metric Learning (LRML) for recommender systems. Unlike the CML algorithm that measures user-item Euclidean distance in vector space and tries to fit positive user-item pairs into the same point, LRML learns a latent relation vector r specific to each user-item pair with an aim to describe the interaction between user u and item i via $p + r \approx q$. This mechanism took inspiration from the recent advances in word embeddings (Mikolov et al., 2013) and knowledge graphs (Bordes et al., 2013). In this way, it significantly alleviates the geometric inflexibility of previous metric learning approaches and allows LRML to scale to a large number of interactions with better performance. Besides, the adoption of a memory network in the architecture of LRML enables the recommender system to capture fine-grained user preferences. Even though this technique has not been applied to POI recommender systems yet, we believe it a promising way to enhance personalised venue recommendations in cities and propose the topic-enhanced memory network (TEMN) for POI recommendation based on this technique in Chapter 5. There, more details of the LRML algorithm and how it serves as part of our TEMN model will be introduced as well.

To display the relations between the three recommendation techniques introduced above more clearly, a geometric comparison of MF, CML, and LRML is shown in Figure 2.2.

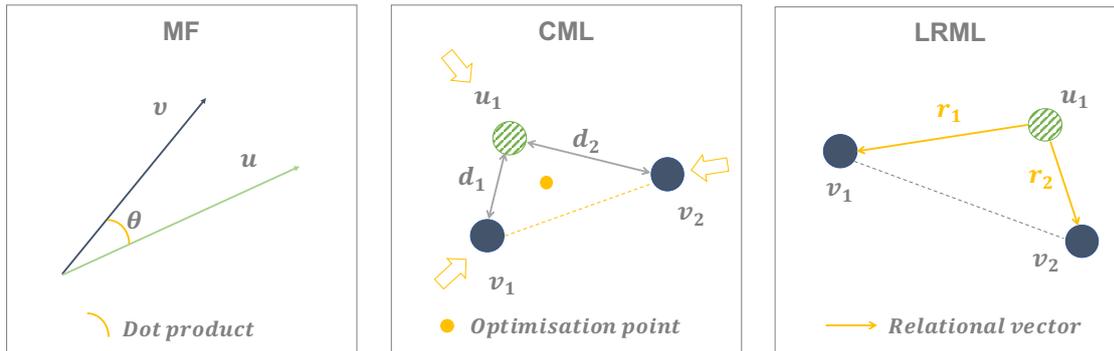


Fig. 2.2 Geometric comparisons of MF, CML, and LRML recommendation algorithms. MF-based methods calculate the dot product of user and item vectors; CML minimises the Euclidean distance between positive user-item pair, and; LRML learns a relational vector for each user-item pair.

2.6 Challenges & opportunities

The extensive overview of relevant techniques in the field of urban computing leveraging GSNs presented in previous sections places us in the right spot to discuss key challenges and opportunities in this certain area. Generally speaking, these challenges and opportunities are embodied in multiple aspects of the urban computing framework, as will be discussed below.

2.6.1 Possible limitations in GSN data

Even though GSN is a powerful data source to solve urban problems, some limitations remain, which deserve particular attention in practice. First, GSN users are generally biased towards individuals who are younger, city dwellers, and smartphone owners. That is to say, GSN data generated by relatively older and poorer residents in more deprived areas are limited, leading to those people and areas under-represented. In such cases, what we observe through the lens of GSN data may only reflect the behaviour of a fraction of the population. Apart from the bias introduced by the user structure, the contents shared in GSNs may also be biased. For instance, users might prefer not to post all their movements information in the physical world online due to privacy and security concerns. Or users might tend to amplify check-ins at trendy venues to impress others. Therefore, when exploring innovative solutions for urban challenges leveraging GSN data, whether the dataset is suitable for the research topic in a certain area needs to be seriously considered. Necessary observation and statistical analysis of the data need to be taken with care.

2.6.2 Knowledge discovery from heterogeneous data sources

Due to the complex nature of urban issues, to come up with effective solutions to deal with them usually relies on information from heterogeneous sources. Depending on the specific urban computing task, relevant data can be analysed together with GSN data to have a more thorough knowledge of the urban problem and thus provide deeper insights. It is conceivable that mining rich interconnected information networks may enable us to disclose more meaningful patterns than what we can obtain from a single data source. However, to discover knowledge from various data sources with different structures and semantics also poses challenges to multiple aspects of urban computing, including data fusion techniques, high efficiency on large-scale datasets, and cross-disciplinary knowledge.

2.6.3 Innovative techniques tailored for urban computing

As we have seen in previous sections, some classical data analysis techniques, including classification methods, clustering analysis, topic modeling, and network analysis, have been adopted to explore GSN data with success and revolutionised urban studies. Despite the significant advances, the study of data analysis techniques for urban computing is still in its infancy, that most of the studies in this research area applied vanilla versions of the algorithms in other domains to urban computing tasks directly. However, unlike purely digital goods or real-world services, GSNs link the physical and virtual worlds by offering users a way to share their life experiences with city places via the "check-in" function. Therefore, user behaviours reflected in GSN data can be affected by both online and real-life factors. Considering the distinguishing characteristics of GSN data, novel techniques tailored for urban computing leveraging GSNs are eagerly expected.

2.6.4 Effective urban computing platforms for diverse applications

Urban computing is a dynamic and fast-expanding field that covers a wide range of applications, from social economics to spatial planning, to traffic control, and even to health monitoring. As the study of urban computing is so naturally multidisciplinary, it requires researchers to be equipped with extensive knowledge and creative thinking to construct domain-dedicated systems for in-depth mining of GSN data. Apart from the diversity of application domains, urban computing tasks usually need to handle large-scale datasets and sometimes even real-time streaming data. In such cases, advanced technologies are required to guarantee efficient and steady urban computing services. Another point worth noting at the application level of urban computing is that we cannot expect every urban planner, urban

policymaker, and even mobile user to know algorithms and technical details. User-friendly urban computing platforms allowing people to perform data analysis and visualise results with high interpretability are needed. The construction of such efficient urban computing platforms for diverse applications remains a challenging area of research.

2.7 Present dissertation & future outlook

The availability of vast amounts of GSN data, thanks to the convergence of social networks and geographic information, provides us an unprecedented opportunity for urban studies in the big data era. With that, this chapter has reviewed the typical framework of urban computing leveraging GSN data as well as the most common algorithms used in this area. We also reviewed recent efforts in the literature to exemplify the main research trends and classical techniques commonly used in the domain of urban computing with GSN data. Last but not least, we have presented the limitations of GSN data and discussed current challenges as well as opportunities when exploring this particular type of data for urban computing tasks, arguing that the research area has not been fully explored. Given the highlighted limitations of data sourced from GSNs, and the increasing demand for novel techniques tailored for urban computing scenarios, some principles lie at the core of our research throughout this dissertation. These principles are believed as the key to the success of urban computing leveraging GSN data and will be explained in detail below.

This dissertation is a step towards gaining a deeper understanding of the human mobility patterns for urban computing. As pointed in Section 2.6.1, some possible biases in GSN data should be considered in this research area so that we can tackle appropriate urban challenges using appropriate datasets. Given that GSN data are generally biased towards more developed regions, two global cities, London and Beijing, are chosen to study in this dissertation. Besides, necessary data preprocessing and primary analysis are taken on data sourced from two popular GSN sites, Foursquare, and WeChat to ensure it is representative of the population. In addition, we devise novel approaches tailored to urban computing tasks with an aim to optimise urban development policies and services primarily relying on GSN data. More specifically, in Chapter 3, statistical analysis, network analysis, and supervised learning approaches are utilised to track and predict socio-economic changes as government financial effort is put in developing urban areas. We also present how GSN data can be integrated with government financial data and socio-economic indicators, and employ four classifiers for the prediction. Different from Chapter 3 that mainly uses supervised machine learning techniques, Chapter 4 focuses more on unsupervised learning algorithms. In this part of the research, we propose a novel density-based clustering algorithm and a

topic modeling method for optimal allocation of cultural establishments and related resources across urban regions. Both of the two proposed algorithms take the distinct characteristics of GSN data and urban human mobility into account. To be more specific, the clustering algorithm devised enables personalised density threshold learning for each user; and the proposed topic modeling approach extends the vanilla LDA by considering spatio-temporal effects of human mobility in the city. Further, we build a hybrid model that combines supervised and unsupervised learning and capitalises on the advances in both memory networks and topic modeling for new venues recommendation in Chapter 5. It is an end-to-end deep learning framework that integrates neighbourhood-based and global preferences of users and incorporates various types of contextual information available in GSNs. Besides quantitative improvements, by incorporating neural attention mechanisms and topic modeling, the interpretability of the POI recommendation is also significantly promoted.

Chapter 3

Cultural investment & urban socio-economic development

In this chapter, we show how to exploit collective transition behaviours in GSNs together with official datasets to detect and predict socio-economic changes in urban areas. The research presented in this chapter provides a case for integrating heterogeneous urban datasets to explore effective data-driven frameworks for local governments to optimise decision-making processes and financial allocations for urban development.

Chapter outline. After first describing the datasets (Section 3.2) and metrics (Section 3.3) that are used, we run a preliminary analysis on London boroughs to get a general view and lay a basis for further investigation in Section 3.4. We then outline four hypotheses in Section 3.5, underpinning our analyses grounded on the preliminary analysis and existing literature. These hypotheses derive from two key concerns: the relationships between urban socio-economic development, cultural investment, and geo-social network features; the feasibility of predicting socio-economic development through cultural expenditure, geo-social network, and geographic features. We then examine the rationality of our hypotheses using ANOVA analysis and a supervised learning classification framework in Section 3.6. Finally, in Section 3.7, we conclude with a discussion, including the contributions and limitations of our findings.

3.1 Introduction

In 1997, the striking 'Bilbao miracle' created by Guggenheim Museum not only provided Bilbao, a depressed northern Spanish port town, with dramatic socio-economic growth but also demonstrated that cities can blossom with cultural investment (González, 2011; Jonathan, 2007). Even though the ability of culture to promote local regeneration has received

general acceptance, large-scale evaluation and prediction of its impact are still not widely practised. The potential of network science in offering insight on deprivation dynamics (Eagle et al., 2010) along with the millions of human mobility traces made available by location-based applications has so far been largely untapped in culture-led regeneration studies. In this chapter, we propose a new fusion of techniques using geo-social network data from Foursquare¹ to quantify the effect of cultural investment on the urban regeneration process and predict its outcome in London's neighbourhoods.

Culture-led urban regeneration, as one of the main branches of urban regeneration, has received increasing attention globally in recent decades and been applied by many governments as a boost to revitalising depressed urban areas. Historically, it was in the 1970s when the culture was first used as a catalyst to accelerate urban regeneration and by the late 1980s when the term 'culture-led regeneration' started to emerge in literature (Jonathan, 2007). Ever since then, the significant role that culture can play in urban regeneration has been widely discussed by researchers. For instance, Keddie (2014) pointed out culture's effectiveness in reducing the deprivation level and promoting 'social mixing' for urban areas; Jonathan (2007) stated that culture can be utilised by cities to improve the existing environment, attract tourism, increase employment, and reinforce civic pride. In addition to the benefits mentioned above, another 'by-product' of culture-led regeneration is creating the city branding, which is thought to be particularly attractive to those international metropolises with an expectation to make the city an alluring base so as to promote its functional role in the global economy (Sassen, 2013; Webber, 2007). Realising the positive effects that might be brought, a growing number of cities have begun to put more effort and allocate more financial resources to culture to promote urban regeneration. London, the city we choose to study in this research, is no exception. Despite local government budgets experiencing considerable pressure as a result of the central government funding cuts in recent years, the local authorities in London remain significant supporters of arts and culture. In 2013/14 for example, the spending of London boroughs on arts and culture was £220.5 million, representing around 3 per cent of the total local authority spending in the city, in comparison with 2.2 per cent nationally (Councils, 2014).

However, how to measure the socio-economic impact of culture-related policy and expenditure is still an open question. Conventionally, the investigation of socio-economic deprivation for urban areas has largely relied on government statistics, with data generally obtained through the traditional survey. It is usually costly to implement and takes a few years to carry out each time. With an aim to overcome this limitation, researchers have recently started to mine low-cost, real-time, and fine-grained new data sources for socio-economic

¹<https://foursquare.com>

deprivation study. For instance, Eagle et al. (2010) discovered a high correlation between call network diversity and urban area deprivation using call records data; Louail et al. (2017) showed the application of bank card transaction data on socio-economic inequalities study in cities; Quercia et al. (2012) found the topic of tweets and the deprivation level of urban areas are correlated; Smith et al. (2013) used Oyster Card data to identify areas of high deprivation level in London; Quercia and Saez (2014) explored the relationship between the presence of certain Foursquare venues with social deprivation; Venerandi et al. (2015) used Foursquare and OpenStreetMap datasets to explore the correlation between urban elements and deprivation of neighbourhoods; And Hristova et al. (2016) discussed the relationship between the prosperity of people and urban places, and distinguished between different categories and urban geographies using Foursquare and Twitter data. In line with this stream of research, we take advantage of the geo-social network data from Foursquare and show its success in capturing and predicting the socio-economic change related to government cultural expenditure which has vital implications for culture-led urban regeneration in neighbourhoods.

More specifically, we utilise the spatial network of Foursquare venues formed by the trajectories of users to track the changes of urban areas. The reason why such kind of geo-social network data is used is that cities are complex systems where the effects of regeneration often grow from the bottom up. Different from traditional analyses that have often adopted top-down methodologies that largely ignore local features, network science provides opportunities to link behaviours of individuals together in a spatio-temporal framework and thus enables significant insights into local physical and social transitions (Batty, 2013). This makes it possible to observe and understand the ever-changing dynamics of culture-led urban regeneration at fine grain as government financial effort is put into developing urban areas. In recent years, some researchers have begun to take advantage of Foursquare data to understand cities. Karamshuk et al. (2013) focused on the problem of optimal retail store placement and explored how the popularity of three retail store chains in New York is shaped; Georgiev et al. (2014) extracted indicators of the spatial positioning of retailers as well as the mobility trends of users to model the economic impact of the Olympic Games on local businesses in London; And Noulas et al. (2015) investigated the topological properties of the urban place networks created by mobility data across a large set of 100 cities globally and applied supervised learning algorithms to predict new links between venues. To the best of our knowledge, this is the first work in which collective transition data from geo-social network is used in culture-led regeneration studies. The main contributions of this chapter include:

- We propose an innovative approach to giving insights on underlying relationships between socio-economic status, cultural investment and geo-social network properties using a fusion

of techniques, including network analysis, statistical analysis, and supervised machine learning algorithms.

- We demonstrate how datasets from government and geo-social networks with different spatial and temporal granularities can be analysed jointly and produce the inference of local socio-economic change for small urban areas.
- We define new metrics on cultural investment and cultural features in geo-social networks to measure the priority level of culture for urban areas and show how the differences in these metrics reflect on the network properties of local areas.
- Applying traditional network metrics to the geo-social graph of transitions between venues on Foursquare, we show that areas with high cultural investment and deprivation level experience significant growth in the following years.
- We prove it feasible to adopt geographic, cultural expenditure, and geo-social network features to predict the binary socio-economic deprivation change for small urban areas with high predictive performance. Our evaluation shows the effectiveness of our prediction models with AUC values of up to 0.85, which is much higher than random guessing.
- We evaluate the predictive capability for different classifiers and features, with Naive Bayes and random forests being the classification methods that give the best performance. In terms of the predictive features that work best, geo-social network features as a whole are the most powerful predictors.

Overall, our findings open new directions for the detection and prediction of socio-economic conditions in the urban environment through collective transition behaviours in geo-social networks.

3.2 Dataset

In this section, we describe datasets used for the study and present their basic properties. In total, there are three major data sources, which are: socio-economic data, cultural expenditure data, and Foursquare data.

3.2.1 Socio-economic data

The dataset used to evaluate socio-economic status for neighbourhoods is the English Indices of Deprivation, an official measure of relative deprivation for small areas (Lower Super

Output Areas (LSOA²) in England calculated by the Department for Communities and Local Government (DCLG). It is organised across seven sub-domains (Health Deprivation and Disability; Employment Deprivation; Income Deprivation; Education, Skills and Training Deprivation; Crime; Barriers to Housing and Services; and Living Environment Deprivation), offers deprivation scores for each and produces the Index of Multiple Deprivation (IMD), which reflects the overall deprivation level. This IMD is the index we particularly focus on in the study to assess the socio-economic status of London areas. It has been calculated since the 1970s and is updated every 3-5 years. The latest version of this index is the IMD 2015³, which updates the previous version of the IMD 2010⁴. In this part of the research, we employ these two successive versions of the IMD to track changes of relative deprivation levels, and thus to understand the socio-economic changes of London areas through a comparative analysis. In the published data, each LSOA in England is given an IMD score and is ranked from the most deprived to the least deprived, allowing users to be aware of how much more or less deprived an area compared to another. A range of summary measures is also available for users to describe deprivation for higher-level geographies. It is worth noting that the IMD scores are not directly comparable between years. However, it is possible to compare IMD ranking changes for neighbourhoods between 2010 and 2015 versions to get a view of whether an area became relatively worse or better in terms of the socio-economic condition during the period, and how large the change was. A more detailed explanation will be given in Subsection 3.3.1.

3.2.2 Cultural expenditure data

The cultural expenditure data utilised in this work is the local authority revenue expenditure and financing derived from DCLG⁵. This dataset is based on returns from all the 444 local authorities in England, showing how they spend their money in each financial year. It provides information about the local authority revenue spending on various service areas, one of which is 'cultural and related services'. This specific category of cultural expenditure can be further divided into five sub-areas: culture and heritage, recreation and sport, open spaces, tourism, and library service. In this study, revenue spending data for financial years 2010/11, 2011/12, and 2012/13 of all these cultural sub-areas are used.

²<https://data.london.gov.uk/dataset/statistical-gis-boundary-files-london>

³<https://www.gov.uk/government/statistics/english-indices-of-deprivation-2015>

⁴<https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010>

⁵<https://www.gov.uk/government/collections/local-authority-revenue-expenditure-and-financing>

3.2.3 Foursquare data

Alongside the two official datasets from the government introduced above, we also employ user mobility records and venue information of London through a three-year-long dataset from Foursquare. This location-based social network dataset contains 'transitions' (successive pairs of check-ins created by users) occurring within London from January 2011 to December 2013. For each transition, venue IDs and timestamps of both origin and destination are recorded. In addition, information of Foursquare venues including the geographic coordinates, category, creation time and the total number of users that have check-in(s) at each venue is also available. In total, there are 3,992,664 transitions generated between 17,804 venues in London during the study period. In Figure 3.1, we map these Foursquare venues by parent categories with cultural venues coloured. Here, cultural venues are defined and selected as urban places of arts, media, sports, libraries, museums, parks, play, countryside, built heritage, tourism, and creative industries, following the line set by the Office of the Deputy Prime Minister in Regeneration through Culture, Sport and Tourism⁶. As we can see from the figure, cultural venues tend to be situated in Inner London than outer suburbs. It is also noticeable that the density of Foursquare cultural venues, in general, is higher for West London than the east of the city.

The Foursquare dataset can be represented as a spatial network of venues connected by transition flows of users for a certain period. It is a directed graph where nodes represent start and end venues, while edges correspond to transitions. In the graph, two nodes are linked if at least one trip exists between two venues during the time. The weight of an edge is proportional to the number of transitions made by all users between the two venues. Since the timestamps of each transition are obtainable, we can study how links are created and how the values of graph features are changing over time. As one would expect, investments need time to attain an observable effect. The result of cultural expenditure, no matter on the organisation of a music festival, the renovation of a library, or the construction of a new art gallery may take days, months or years to be observable. Here, we assume the impact of cultural expenditure from local authorities can be observed after 9 months on average through Foursquare. The reason why 9 months is chosen as the delay time is because the newly opened cultural venues become popular after this certain period of time averagely according to their venue visitation signatures on Foursquare over time. Based on this assumption, expenditure and geo-social network data are compared according to time scales in Table 3.1, where we look at annual snapshots of the data for different years. Formally, we define our yearly dataset as a directed graph $G_t = (V_t, E_t)$ for $t = 1, 2, 3$, which indicates three snapshots in time of the dataset. The set of nodes $V_t = \{v_1, v_2, v_3, \dots, v_{N_t}\}$ is composed of N_t Foursquare

⁶<http://www.communities.gov.uk/archived/publications/localgovernment/regenerationthroughculture>

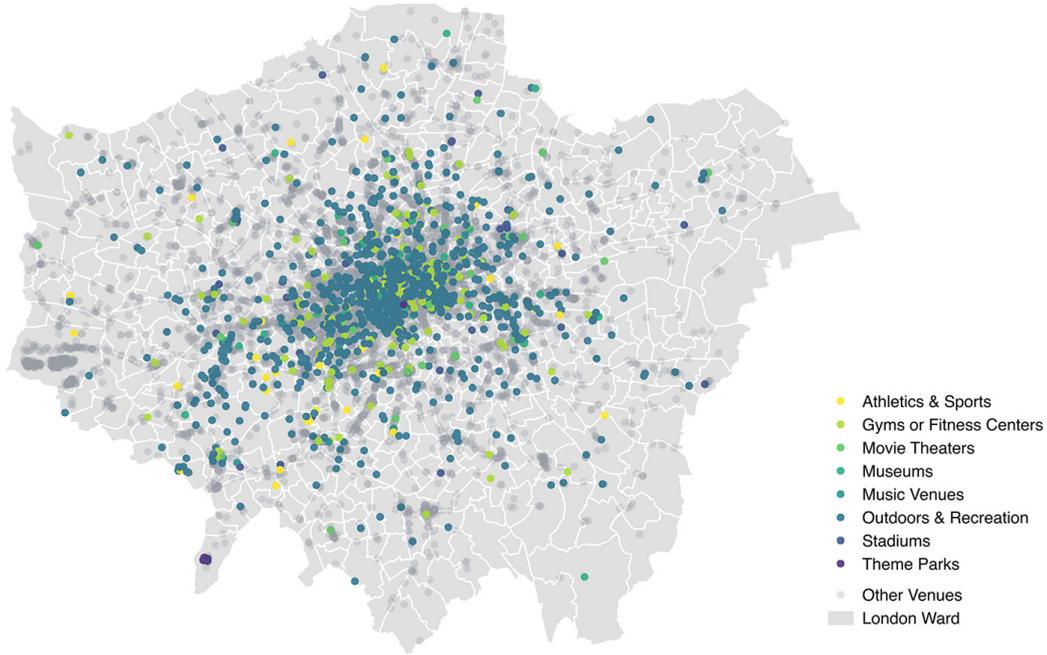


Fig. 3.1 Spatial distribution of Foursquare venues in London. Cultural venues by parent categories are coloured, and other categories of Foursquare venues are in grey.

venues and the set of edges $E_t \subseteq V_t \times V_t$ is composed of pairs of venues that have at least one transition generated between each other during time t . An edge $(v_{ot}, v_{dt}) \in E_t$ is called a transition edge between v_{ot} and v_{dt} , where v_{ot} represents the origin of the transition and v_{dt} represents the destination. The network properties for each year are shown in Table 3.2.

Expenditure dataset	GSN dataset
Financial year 2010/11 (Apr 2010 - Mar 2011)	Jan 2011 - Dec 2011
Financial year 2011/12 (Apr 2011 - Mar 2012)	Jan 2012 - Dec 2012
Financial year 2012/13 (Apr 2012 - Mar 2013)	Jan 2013 - Dec 2013

Table 3.1 Comparison table of time scales for datasets.

3.2.4 Spatial unit of data

Since the first two datasets are provided by the government at different geographic levels, here, we introduce the spatial units used in this research and demonstrate how they are applied in our further investigation. This clarification is necessary since understanding how geo-referenced data is aggregated spatially is key to linking the datasets used in the chapter.

t	Duration	 V 	 E 	$\langle C \rangle$	$\langle k \rangle$
1	January 2011 - December 2011	15832	469229	0.221	59
2	January 2012 - December 2012	16189	715113	0.228	70
3	January 2013 - December 2013	17684	742017	0.240	84

Table 3.2 Network properties at each snapshot. Number of nodes $|V|$, number of edges $|E|$, average clustering coefficient $\langle C \rangle$, and average degree $\langle k \rangle$.

The two geographic levels of London areas used in this research are borough and ward⁷. In the research, we perform our exploratory analysis on the aggregate borough level to have a first look and then improve on geographic granularity by using wards as smaller localities for our statistical evaluation and prediction. Here, the reason why the ward is selected as the geographic unit for the prediction is that it provides a relatively sufficient training set size for supervised learning models, which is unreachable for the borough (there are 32 boroughs and 625 wards in London). As mentioned above, the IMD data is available at several geographic levels, including both borough and ward, by consulting its official document. However, the DCLG only provides expenditure data initially at the borough level. We obtain the data for wards by dividing the cultural expenditure of a borough by the number of wards it includes, assuming resources are spent proportionally. As for Foursquare data, check-ins and venues are described at the level of latitude and longitude coordinates. This fine-grained spatial representation of activity at the Foursquare data layer allows for standard geographic aggregation methods to be applied in order to attain representations at the ward and borough levels, making, therefore, the linkage of the three datasets possible.

3.3 Metrics

Leveraging on the described data, we introduce some metrics highlighting the advantage that certain neighbourhoods have in terms of cultural expenditure and the properties of geo-social networks for local areas. We also compute a number of geographic features based on the neighbourhood's location and Foursquare venue information.

3.3.1 Geographic features

In Figure 3.2, we map the IMD change for London wards from 2010 to 2015 with sub-region information provided. Here, we first rank the wards according to their IMD scores from the most deprived to the least in 2010 and 2015, respectively. Then, we subtract the ranking

⁷<https://data.london.gov.uk/dataset/statistical-gis-boundary-files-london>

in 2010 from that in 2015 for each ward to look at the change, based on which, wards are coloured blue or red, according to whether they became relatively more deprived or less deprived from 2010 to 2015 compared to other London neighbourhoods. Please note that when we say less deprived or more deprived in the rest of this chapter, it means the IMD ranking of an area in 2015 became relatively lower or higher compared to that in 2010. The larger the change, the darker the shade. From the result, we can see that the ward that experienced the largest improvement locates in East London with an increase of 212 in the IMD ranking. On the other hand, the rank of a ward in Central London dropped most significantly by a number of 190. We can also observe from Figure 3.2 that areas with similar IMD ranking changes tend to be spatially clustered. And neighbourhoods showing larger improvements in terms of the overall deprivation level are more likely to be in East London. These findings suggest that geographic factors can influence the deprivation level change in urban areas. In this case, we involve a number of geographic features as metrics, including the sub-region a ward belongs to, its area size, and how far it is from the city centre.

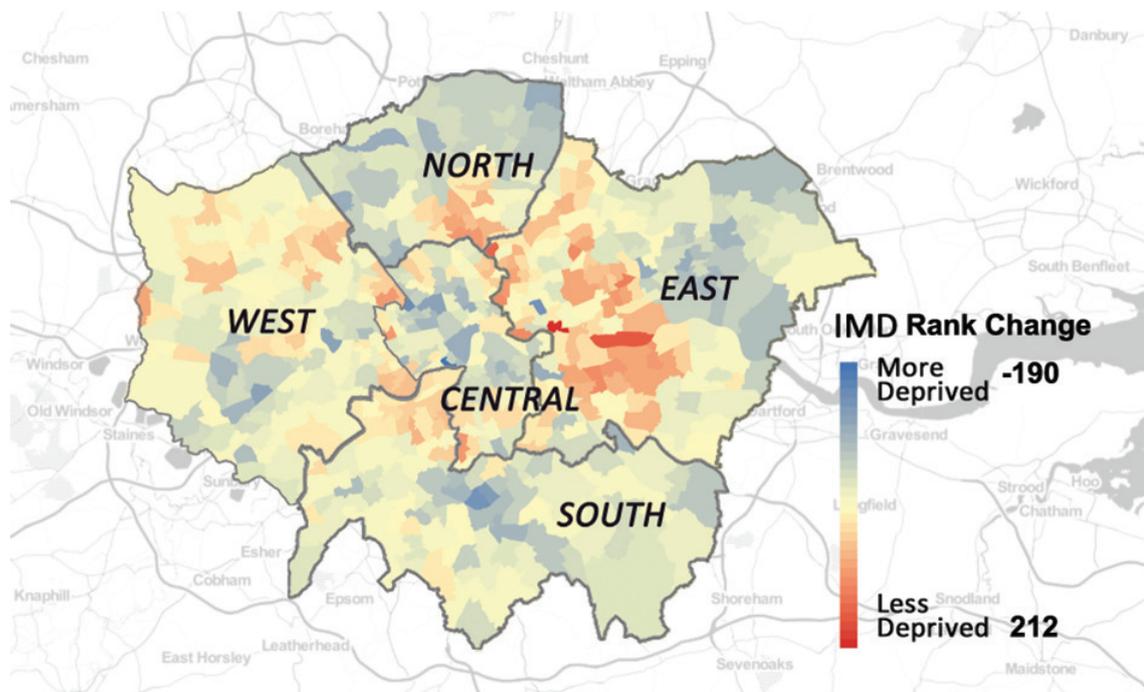


Fig. 3.2 Map of IMD rank change from 2010 to 2015 for London wards and sub regions.

In addition, some metrics about Foursquare venues with their geographic properties considered are also given, including the number of venues created (*VC*) and venue created density (*VCD*). Different from the concept of the node we described in the previous section, *VC* is defined as the total number of venues emerging in an area during a certain period and estimated on the basis of creation time information obtained from Foursquare venue profiles.

To divide the VC by *area size*, we get VCD , which represents the average number of new venues created in an urban area per square kilometre.

3.3.2 Network metrics

The network measures applied in this research are in-degree centrality (IC), out-degree centrality (OC), in-degree/out-degree ratio (IOR), and average clustering coefficient (ACC). In-degree centrality of an area i represents how many in-flow transitions the nodes of area i receive from nodes of other areas. In contrast, out-degree centrality measures how many out-flow transitions start from i , but flow to other areas. We also introduce a metric called IOR , which indicates the ratio of in-flow transitions over out-flow transitions. If the IOR of an area is high, it means the area is more likely to be an attractive place to visit for people from other places. For area i , the IOR can be calculated by:

$$IOR_i = \frac{IC_i}{OC_i} \quad (3.1)$$

The local clustering coefficient captures the degree to which the neighbours of a given node are connected with each other. For a node i with degree k_i , the local clustering coefficient (Miller and Blair, 2009) is defined as:

$$C_i = \frac{L_i}{k_i(k_i - 1)} \quad (3.2)$$

where L_i represents the number of edges between the k_i neighbours of node i . Then, the average clustering coefficient, which reflects the overall level of clustering in an area is measured as the average of the local clustering coefficients of all the nodes within it.

3.3.3 Growth rate

We also introduce growth rate metrics for some features to present changes in urban areas. Take the growth rate of nodes (GRN) as an example, we define this metric to reveal the temporal change of nodes in geo-social network graphs. If the number of nodes we observe in a network snapshot during a year $(t - 1)$ is N_{t-1} , and a number of N_t in the subsequent period t , we calculate the GRN of graph for t as:

$$GRN_t = \frac{N_t}{N_{t-1}} \quad (3.3)$$

In a similar way, other growth rate measures listed in Table 3.3 can be obtained.

3.3.4 Cultural advantage metrics

In order to measure how the cultural level of a neighbourhood is higher or lower than the average city, we introduce two cultural advantage metrics which rely instead on the concept of location quotients in economic geography. Location quotients capture regional industry specifics by comparing an area's business composition to that of a larger geographic context (i.e., state or nation) and can be calculated by the following formula:

$$LQ_i^j = \frac{q_i^j}{\sum_{j \in J} q_i^j} \cdot \left(\frac{\sum_{i \in I} q_i^j}{\sum_{i \in I} \sum_{j \in J} q_i^j} \right)^{-1} \quad (3.4)$$

In the equation, LQ_i^j represents location quotients for each industry j and for each region i ; q_i^j denotes the gross output of industry j in region i ; I and J are the sets of regions and industries, respectively. Here, the values of location quotients vary by region due to its industry makeup and can be interpreted by comparing with 1. If the value is greater than 1, it signals that the concentration of a certain industry in a particular region is higher than the average level and a value less than 1 indicates the industry is relatively scarce in that region (Miller and Blair, 2009; Schütz, 2017).

Cultural expenditure advantage

Inspired by industry location quotient, we define a metric called cultural expenditure advantage (CEA) to evaluate the priority of cultural expenditure for a neighbourhood in the city. This metric reflects the extent to which a local authority spends more on culture than the city average level. The CEA for area i in the city can be represented as:

$$CEA_i = \frac{CE_i}{TE_i} \cdot \left(\frac{\sum_{i \in I} CE_i}{\sum_{i \in I} TE_i} \right)^{-1} \quad (3.5)$$

where CE_i is the amount of cultural expenditure of neighbourhood i ; TE_i is the amount of total expenditure of i ; and I is the set of neighbourhoods in the city. Through comparing CEA with 1, whether the cultural expenditure of an area is higher than the city average can be evaluated.

Cultural venue advantage

Similar to the CEA defined previously, a metric of cultural venue advantage (CVA) is given to reflect the extent to which a neighbourhood has more cultural venues than the city average. Here, cultural venues include 8 major categories of Foursquare culture-related places presented in Figure 3.1. The CVA for neighbourhood i can be defined as:

$$CVA_i = \frac{CV_i}{TV_i} \cdot \left(\frac{\sum_{i \in I} CV_i}{\sum_{i \in I} TV_i} \right)^{-1} \quad (3.6)$$

where CV_i is the number of cultural venues in i and TV_i is the total number of venues in i .

A summary of all the metrics used in the following analyses and tests are listed in Table 3.3 with their categories, descriptions, and where they are applied in this chapter provided.

Category	Metric	Description	Application
Initial IMD	<i>Initial IMD</i>	Rank of IMD at the beginning	P [H3]
	<i>Sub Region</i>	Sub-region of London where a ward locates	[H3]
	<i>Area</i>	Size of a ward (km^2)	[H3]
	<i>Distance</i>	Distance from the centre of London to spatial centre of a ward (km)	[H3]
Geographic	<i>VC</i>	Number of venues created in an area	P [H1] [H2]
	<i>VCD</i>	Number of venues created in an area per km^2	[H1] [H2]
	<i>CVA</i>	Extent to which an area provides more cultural venues than city average	P
	<i>GRVC</i>	Growth rate of venues created number	[H3]
	<i>N</i>	Number of nodes for an area	[H1] [H2]
	<i>IC</i>	Number of in-flow transitions an area receives from other areas	P [H1] [H2]
Network	<i>OC</i>	Number of out-flow transitions an area receives from other areas	P [H1] [H2]
	<i>IOR</i>	Ratio of number of in-flow transitions over out-flow transitions	[H1] [H2]
	<i>ACC</i>	Degree to which nodes in a ward tend to clustering together	[H1] [H2]
	<i>GRN</i>	Growth rate of number of nodes	[H3]
	<i>GRI</i>	Growth rate of number of in-flow transitions	[H3]
	<i>GRO</i>	Growth rate of number of on-flow transitions	[H3]
	<i>GRIOR</i>	Growth rate of ratio of in-flow transitions over out-flow transitions	[H3]
	<i>GRACC</i>	Growth rate of average clustering coefficient	[H3]
	<i>CE</i>	Expenditure on cultural and related services	P
	<i>CEA</i>	Extent to which an area spends more on culture than city average	P [H3]
Cultural Expenditure	<i>CEOP</i>	Expenditure on open spaces per capita	[H3]
	<i>CECH</i>	Expenditure on culture and heritable per capita	[H3]
	<i>CELS</i>	Expenditure on library services per capita	[H3]
	<i>CERS</i>	Expenditure on recreation and sport per capita	[H3]
	<i>CET</i>	Expenditure on tourism per capita	[H3]

Table 3.3 Description of the variables used in the analyses. P represents preliminary analysis.

3.4 Preliminary analysis

The exploration of relationships between the IMD, cultural expenditure, and network features is the foundation of our prediction task for the socio-economic deprivation change. Before discussing it in more depth on a finer spatial scale, we run a preliminary analysis on London boroughs in this section to visualise the relationship patterns and provide evidence for further investigation.

In this part of analysis, we investigate how areas with different deprivation levels spent their money on the culture at the start of our observation period, how they adjusted their priorities in the following years, and how their network graphs changed accordingly. Firstly, Figure 3.3 is created to reveal the initial relationship between the IMD score and cultural advantage metrics (*CEA* and *CVA*) in 2010. In this figure, the colour bar on the right presents the IMD score of London boroughs in 2010, where yellow means relatively more deprived and purple indicates less deprived. The IMD score is represented by the circle size, where the larger the circle, the higher level of deprivation the area. In addition, we partition the figure at $CEA = 1$ and $CVA = 1$. Through these two axes across the 1, Figure 3.3 is split into quadrants, allowing us to see where boroughs with different deprivation levels are centralised. As we can observe from this plot, yellow circles cluster in the middle/lower part, suggesting more deprived boroughs spent relatively less on culture-related services and showed average cultural venue advantage at the beginning of the study.

Then, in Figure 3.4, we discuss how London boroughs spent their money on culture in the next two years, and how their network and local properties changed. It can be observed that yellow circles, which represent more deprived boroughs, show at the upper right of the charts. In contrast, purple circles, which stand for well-off areas, present in the lower left part of the plots. These signals suggest that compared to prosperous areas, more deprived neighbourhoods in London spent more money on culture, had a larger number of venues created, and showed higher in-degree centrality and out-degree centrality from 2011 to 2013. Additionally, these trends were more obvious among most and least deprived boroughs compared to average ones.

Through the above analysis, we find that the initial deprivation level and cultural expenditure strategy may influence the local network graph of urban areas. Furthermore, the boost for region development in more deprived areas would probably occur after investing more in cultural and related services there.

In the following sections, this question will be discussed in greater depth by improving spatial and temporal granularity, and involving statistical and machine learning techniques. To prepare for the analysis, we distinguish London's neighbourhoods on the basis of cultural expenditure priority and the deprivation level. Specifically, wards are firstly grouped into

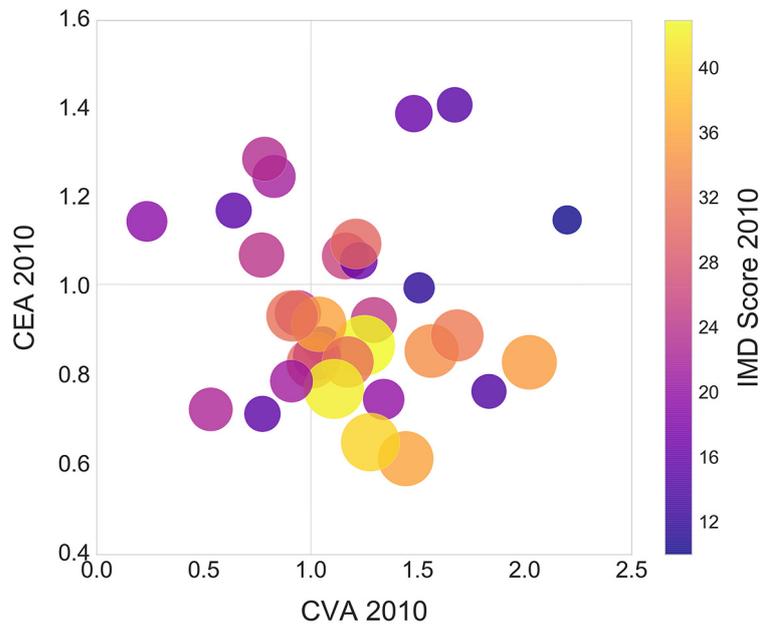


Fig. 3.3 Initial IMD score, CEA, and CVA of London boroughs in 2010. Each London borough is represented by a circle, the colour and size of which indicate the IMD score of the borough. Yellow means relatively more deprived and purple indicates less deprived. The larger the circle, the higher level of deprivation the borough.

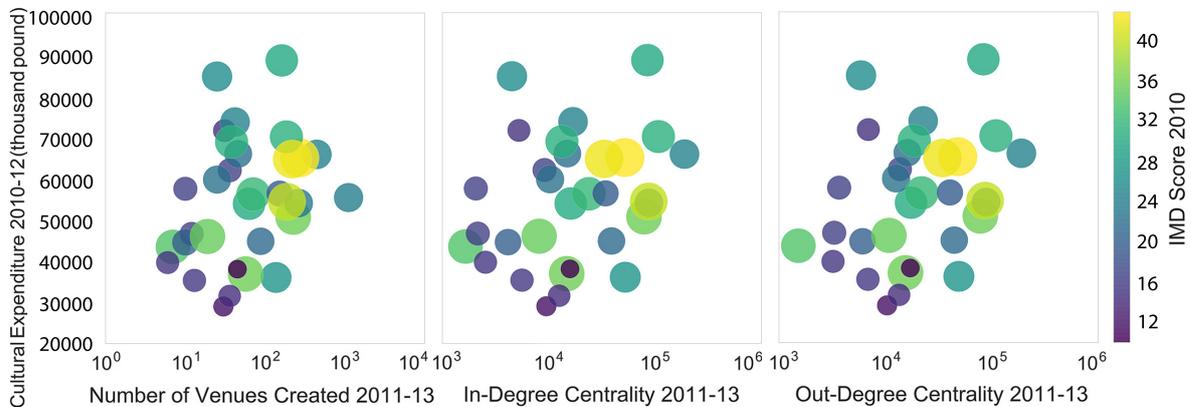


Fig. 3.4 Culture expenditure and Foursquare features changes of London boroughs. Yellow circles represent relatively more deprived boroughs and purple circles stand for well-off ones.

two categories, more deprived and less deprived, according to whether their IMD 2010 deprivation level is higher or lower than the city average. Then, the two groups are further classified according to their cultural spending priorities. If the *CEA* of a ward is more than 1, it is clustered into the more advantaged groups; otherwise, it is put into the less advantaged groups. On the basis of these rules, we can identify four groups of cases outlined in Table 3.4 and mapped in Figure 3.5. The two largest groups are Group 2 and Group 1, which

include 192 and 160 wards, respectively. Again, at the ward level, it indicates that less deprived areas tend to be more cultural advantaged, and more deprived areas tend to be less cultural advantaged. We can also observe that the majority of wards in Group 2, which are more deprived and less advantaged in cultural spending, are located in East London. The relationship patterns between deprivation level, cultural expenditure, and network properties discovered in this preliminary analysis part lead to our hypotheses in the next section. And the four distinct groups of the neighbourhoods will be referred to in our following hypotheses evaluation.

	Group 1	Group 2	Group 3	Group 4
Initial IMD	Less deprived	More deprived	More deprived	Less deprived
CEA	More advantaged	Less advantaged	More advantaged	Less advantaged
Number	160	192	88	114

Table 3.4 Groups of London wards in ANOVA analyses.

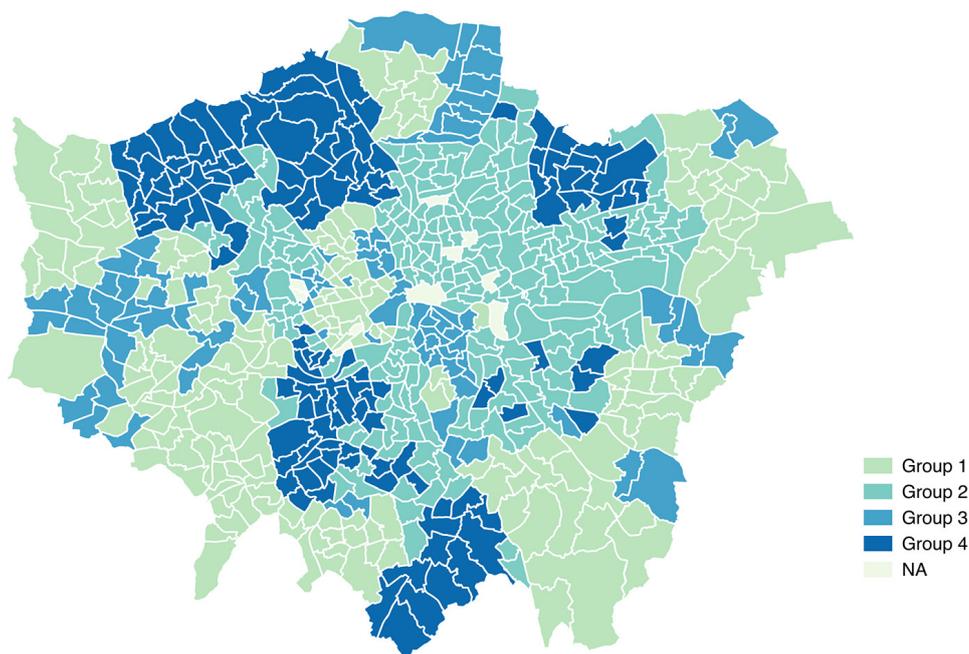


Fig. 3.5 Spatial distribution map of ward groups in London. NA denotes the ward's relative IMD ranking remains unchanged between 2010 and 2015.

3.5 Hypotheses

In order to explore the role that geo-social network data can play in culture-led urban regeneration study, we firstly reveal the underlying relationships between socio-economic status, local cultural expenditure, and geo-social network graph. As a base condition, we expect that neighbourhoods with different socio-economic conditions and amounts of cultural expenditure will show different network properties in the geo-social graph. Therefore, we hypothesise that:

[H1] Areas with high cultural investment and deprivation level *have significantly different networks and local properties* from areas with low cultural investment and deprivation level.

This assertion lays the foundation of further investigation into the nature of culture-led urban regeneration, where based on existing case studies from literature (González, 2011) and our preliminary analysis, we expect that cultural investment in more deprived areas results in growth. Specifically:

[H2] Areas with high cultural investment and deprivation level experience *significant growth with respect to network and local properties* from areas with low cultural investment and deprivation level.

Additionally, network studies have shown great potential in reflecting socio-economic conditions in communication networks and the prediction of deprivation (Eagle et al., 2010). Based on this existing knowledge, we propose that network features such as the centrality and clustering coefficient, combined with geographic and cultural expenditure factors are able to predict local socio-economic changes. We put forward the following two hypotheses to this end:

[H3] Network features of areas together with cultural expenditure and geographic features are powerful signals in predicting socio-economic change.

[H4] Network features of areas are better predictors of improvement than expenditure and geographic features.

3.6 Evaluation & results

In the following hypotheses evaluation, London's neighbourhoods will be studied at ward level and grouped according to Table 3.4 before two types of ANOVA analyses are run. We aim to get a deeper understanding of the relationship patterns we found at the borough level, and test how they vary between ward groups and different periods. We evaluate the last two hypotheses using a prediction framework, which allows us to reason about the predictive power of different features described in Section 3.3.

3.6.1 [H1] Network and local properties

To test the [H1], we employ independent one-way ANOVA to examine whether a statistically significant difference is found in terms of a set of network and local features between different ward groups. Each feature value is the average for the three years from 2011 to 2013.

From the output results, we can see that there are significant effects of groups on six features at the $p < .05$ level. The feature that distinguishes less and more deprived groups is the average clustering coefficient with a statistically significant main effect of $F(3, 550) = 4.15, p = .006$. This shows that neighbourhoods from different socio-economic status groups presented significantly different clustering patterns in their network graphs. Furthermore, through taking a comparison between groups in Figure 3.6, we find that less deprived wards (Group 1 and Group 4) have higher means of average clustering coefficient than more deprived ones (Group 2 and Group 3), which illustrates that venues in less deprived areas are more likely to cluster together. Group 1, which represents less deprived and more cultural spending advantaged neighbourhoods, is the only group that exceeds the average of all wards (red dashed line in Figure 3.6).

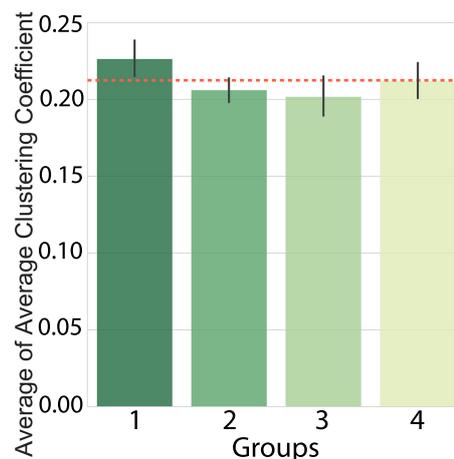


Fig. 3.6 Means plot for average clustering coefficient in independent one-way ANOVA analysis. The red dashed line denotes the average level of all the wards.

The other five features that show statistically significant effects reveal differences between cultural advantaged and disadvantaged groups. The means of these five factors for different groups are plotted in Figure 3.7, from which we can find areas that gave a higher priority to culture (Group 1 and Group 3) had larger venue created number, node number, in-degree centrality, out-degree centrality, and venue created density on average. Moreover, Group 3, which was more deprived but invested a relatively larger amount of money in culture from the financial year 2010/11 to 2012/13, had the highest means in most cases. This result indicates that with more effort being put into culture, the stimulation of local business and

the enhancement of vitality for urban areas may appear. Additionally, a more striking effect can probably be seen in more deprived areas.

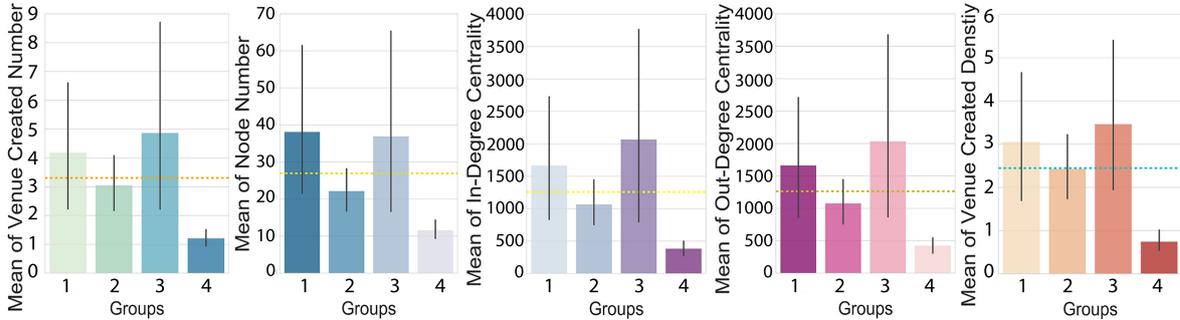


Fig. 3.7 Means plots for variables with statistically significant effects between cultural advantaged and disadvantaged groups in independent one-way ANOVA analysis. Dashed lines denote the average feature values.

Through the one-way ANOVA analysis, we find that urban areas of different socio-economic status and cultural investment priorities vary in terms of local network features. Furthermore, areas with high cultural investment and deprivation level *show significantly different values in network and local properties from those areas with low cultural investment and deprivation level*, which suggests that [H1] is true.

3.6.2 [H2] Growth of network and local properties

After discussing the differences between groups on average, we test whether different groups of areas experienced significantly different growth patterns with respect to network and local properties in this subsection. Technically, we examine whether there are statistically significant differences between years and whether interaction effects exist between our two main factors, group and time, by factorial repeated measures ANOVA analysis.

We present the means plots of five dependent variables that show statistically significant effects in Figure 3.8. From these plots, it can also be found that Group 1 and Group 3, which gave a high priority to culture had dramatic advantages in terms of almost all the features during the three years. Then, to consider the groups separately, Group 3 (more deprived and more advantaged in cultural investment) is still the one that had the highest means in general, while Group 4 (less deprived and less advantaged) is the lowest with respect to all the features. These results demonstrate that significant differences not only exist between ward groups but also show between different time points. It also reveals that the advantage of culture-supporting areas in various local and network properties is a dynamic and continuous process rather than an occasional phenomenon shown in a single year.

Moreover, significant interaction effects between group and year are found in three features: venue created density ($p = .008$), in-degree centrality ($p = .038$), and out-degree centrality ($p = .037$). We then run pairwise comparisons between different years and find that each pair of time points are observed to have statistical significance on venue created density, which reinforces that significant changes exist between years on this feature. As for the two centrality metrics, Group 1 differed significantly from Group 2 and Group 3, suggesting that less deprived neighbourhoods that spent more money on culture experienced significantly different changes with regard to in-flow and out-flow transitions compared to more deprived neighbourhoods.

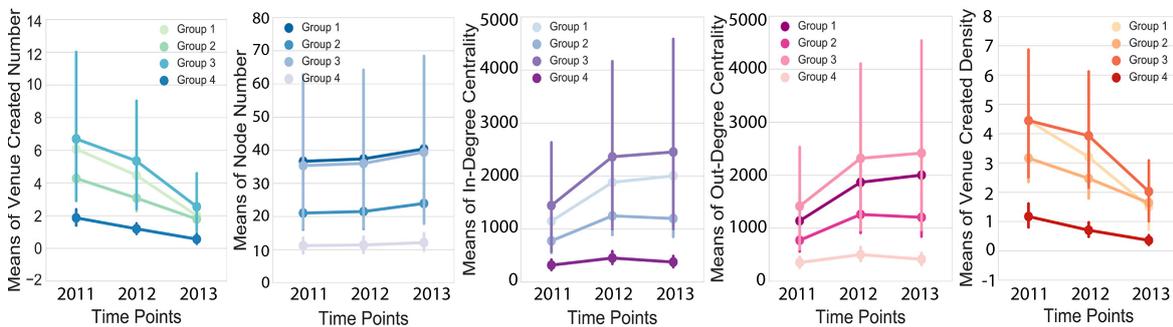


Fig. 3.8 Means plots for variables with statistically significant effects in factorial repeated measures ANOVA analysis.

With the help of factorial repeated measures ANOVA, in this subsection, we detect how the growth of network and local properties varied between groups and time points. We confirm [H2] by finding that areas with high cultural investment and deprivation level experience significant growth with respect to network and local properties from areas with lower cultural investment and deprivation level. The interaction effect between group and time discovered can be further studied to explore whether there was a major culture-related policy or investment taking effect in a certain group at a certain time point.

In summary, the ANOVA analysis results for wards presented in the evaluation of [H1] and [H2] confirm the trend that we observed in the preliminary analysis for boroughs. Significant differences show in geo-social network variables between groups, time points, as well as their interaction effects, suggesting that urban areas with different socio-economic situations and cultural investment attitudes present different network graph patterns. Generally, when more investment are made in culture, the local network of an area grows in several ways. Additionally, this effect is more evident in those more deprived areas.

Building on the findings from preliminary and ANOVA analyses, we introduce a supervised learning framework that exploits the prediction features displayed in Table 3.3 to predict the relative IMD changes for London wards next. We assess whether our prediction

features from three main categories of geographic, cultural expenditure, and network can be combined to build prediction models with good performance as a response to [H3]. Followed by the overall evaluation of classification models, our focus turns to explore the predictive power of different feature classes, especially the network feature set to test [H4].

3.6.3 [H3] Prediction model and overall evaluation

In this subsection, we establish prediction models for the IMD change and discuss the performance of various methods on different neighbourhood sets.

The target feature of our prediction is the binary IMD change in 2015 compared to the initial condition in 2010, which can either become relatively more deprived or less deprived. We propose a supervised learning approach to tackle this binary classification problem: for each ward, we collect its initial IMD rank in 2010 and some basic geographic features; We calculate the average local cultural expenditure of five kinds and the average *CEA* during financial years from 2010/11 to 2012/13; Also, we compute local and network metrics for the beginning and end snapshots of 2011 and 2013, respectively, before calculating multiple growth rate features. The full list of prediction features we adopt to discriminate areas that are more or less likely became better or worse in terms of the socio-economic condition has been presented in Table 3.3. After all these values are collected and calculated, we train classifiers and use a stratified 10-fold cross validation as the evaluation approach. The supervised learning methods we implement are classification tree, random forest, logistic regression, and Naive Bayes using the algorithms in library Scikit-learn⁸. After models are trained and established, we employ the AUC value, the average accuracy, and the average precision as measures to evaluate the prediction performance for our classifiers. Precision is the fraction of positive predictions that are correct. Accuracy represents the proportion of the total number of predictions that are correct. AUC, which is not as intuitive as the previous two measures, stands for area under ROC (receiver operating characteristic) curve and is commonly used as a measure of the overall quality of binary classification models.

Instead of only focusing on the whole set of London's wards, we also classify them according to how large their IMD rank changes are from 2010 to 2015, so as to discuss whether the prediction effectiveness varies when looking at areas with different IMD changes. In Figure 3.9, we present the IMD change distribution of wards in London. As we can see from Figure 3.9a, the IMD change of the whole ward set is normally distributed. Even though there are 625 London wards in total, 385 wards that have data available for all the prediction features are involved. The number of wards for each subset is also presented in Figure 3.9b

⁸<http://scikit-learn.org/stable/index.html>

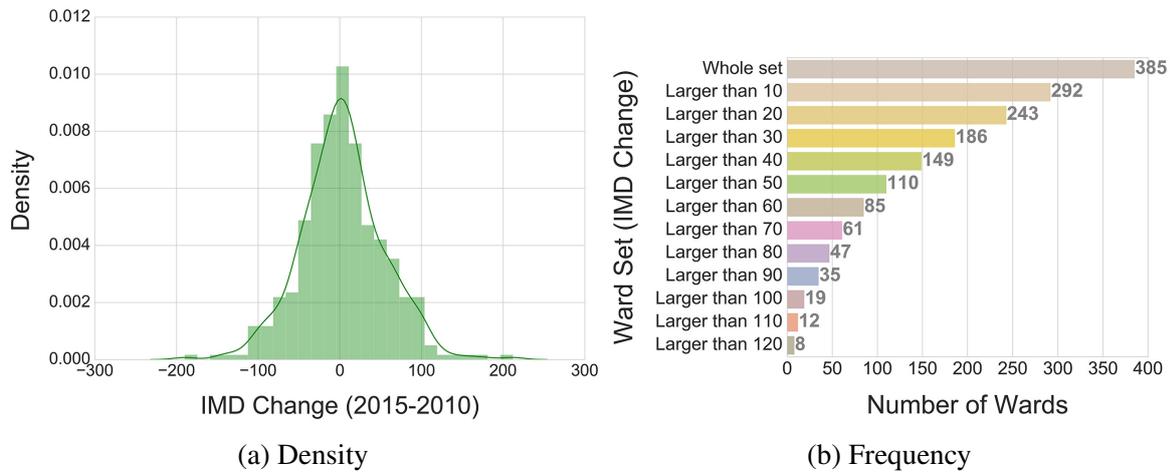


Fig. 3.9 Distribution of the IMD change for London wards.

on the right. Due to the consideration of sufficient sample size for our supervised learning models, besides the whole ward set, four subsets are also chosen to run the test, which are wards with the IMD rank change larger than 10, larger than 20, larger than 30, and larger than 40, respectively.

Overall, our prediction results shown in Figure 3.10 reveal that the inclusion of network, cultural expenditure, and geographic features offers high prediction performance by giving AUC scores over 0.7 for almost all the classifiers. And the best prediction performance shows when we look at wards with the IMD rank change larger than 40 using Naive Bayes, that the AUC reaches 0.865. As the AUC value lies between 0.5 to 1, where a random classifier has an AUC of 0.5 and a perfect classifier's AUC is equal to 1 (Trevor et al., 2009), the predictive performance of our models is much better than random guessing. With regard to the accuracy and precision measures, the scores are also higher than 0.7 in general.

In addition, for different ward sets, we observe that it shows a rising tendency for all the classifiers in terms of evaluation measures in Figure 3.10. This finding suggests that better prediction results can be achieved from wards that have larger IMD changes. The reason for this is probably that neighbourhoods which experienced larger IMD changes showed more evident changes in local and network properties, making their socio-economic changes easier to be predicted.

When comparing the performance between different classifiers, we can see that Naive Bayes and random forest outperform the other two methods with high values in terms of all the three metrics. Followed by Naive Bayes and random forest, logistic regression performs slightly worse, whereas classification tree presents the lowest values. While we have not explored exhaustively initialisation parameters of the four classifiers, what is important with

regards to the goals of the present work, is that their performance evolves steadily with respect to the feature exploration we are demonstrating next.

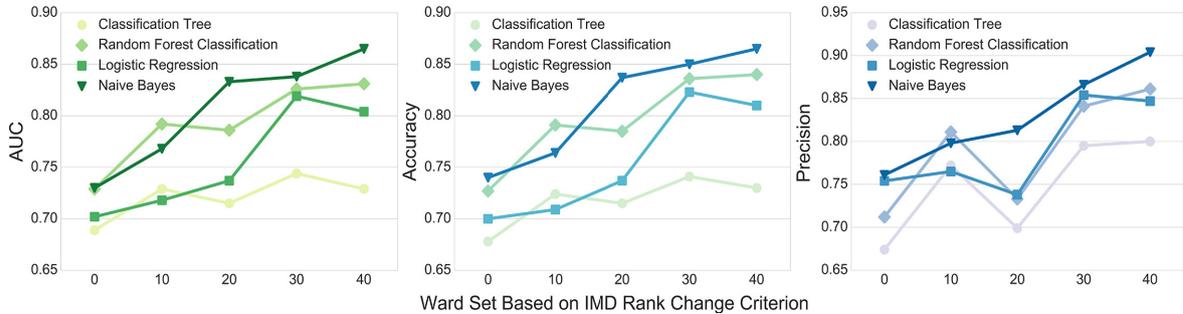


Fig. 3.10 Evaluation for supervised prediction methods on different ward sets.

The high performance of our models presented by the overall evaluation result shows the feasibility and superiority of utilising network features together with cultural expenditure and geographic features to predict the IMD change for urban areas [H3]. This desired outcome once again proves that crowdsourced data can play a significant role in socio-economic deprivation prediction instead of expensive census data.

3.6.4 [H4] Individual features evaluation

After evaluating the overall performance of prediction models, we investigate the predictive power of individual features in this part of analysis.

Firstly, we study the predictive power of each individual feature in the prediction model. We take random forest classifier as an example and compute relative importance for each feature, which is determined in terms of the Gini index (Breiman, 2017). In Figure 3.11, we can see that *CEA* (cultural expenditure advantage) plays the most significant role, being the only feature with an importance score over 0.1. It is followed by features *GRIOR* (growth rate of the ratio of in-degree centrality over out-degree centrality) and *CEOP* (expenditure on open spaces per capita) with values around 0.09. While for the last two features, *Sub region* (sub-region of London where a ward locates) and *GRVC* (growth rate of venues created number), the importance scores are less than 0.03.

Next, in order to understand to which extent different feature classes are contributing to the prediction, we test what prediction performance can be achieved by removing one feature class with respect to the full model. The prediction results of these new models with two feature classes considered in each are shown in Figure 3.12. From the figure, we can see that geographic features make the smallest contribution to prediction models, as the reduction of



Fig. 3.11 Relative importance evaluation for each feature in random forest classification.

prediction effectiveness is least when they are removed. In contrast, network features as a whole are the strongest signal to predict the IMD change [H4].

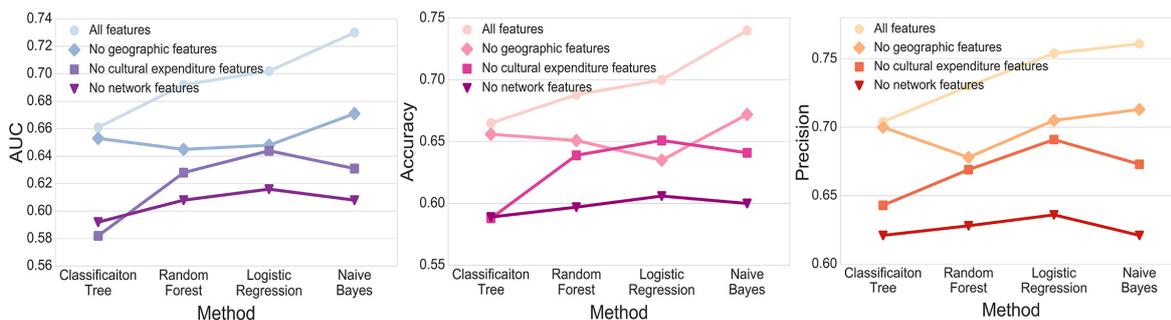


Fig. 3.12 Evaluation for supervised prediction methods on different feature sets.

In conclusion, we present the feasibility and effectiveness of using geo-social network, cultural expenditure, and geographic features to infer whether the socio-economic status gets relatively better or worth for small urban areas in the evaluation of [H3] and [H4]. The results of our prediction models are favourable with AUC scores higher than 0.7 in general. Moreover, the prediction performance sees an improvement when we focus on neighbourhoods that have larger IMD changes, and employ Naive Bayes and random forest classifiers. To evaluate the prediction contribution of features separately, network is the category that plays the most significant role in the prediction and cultural expenditure advantage is the most powerful individual feature.

3.7 Discussion & conclusions

In this chapter, we have investigated the socio-economic impact of cultural expenditure on London neighbourhoods, visible through the lens of location-based mobile data. Finding evidence of the regenerative phenomena in GSNs along with such investment in local areas, we take a step further by trying to predict socio-economic impact based on geographic, network and cultural expenditure features. Overall, we have put forth evidence of the potential of using geo-social data for detecting and predicting the impact of culture-led regeneration strategies. This has a number of significant implications for location-based mobile systems, local governments and policymakers alike.

Firstly, we have explored the relationship between socio-economic condition, cultural investment, and geo-social network graph, finding that regions that spent more on culture experienced an improvement of local development, especially for more deprived neighbourhoods. This observation verifies the effectiveness of implementing cultural strategies in urban regeneration projects and illustrates that culture-led regeneration policies are more suitable for underprivileged areas. On the basis of this finding, we suggest governments and policymakers taking socio-economic conditions into consideration as an important factor when implementing cultural strategies to promote local development. Furthermore, our research has presented geo-social network data's ability to enrich or even replace traditional census-based deprivation statistics by proposing a supervised learning framework for predicting the outcome of cultural investment in London neighbourhoods. Although at present we perform our evaluation on annual snapshots, which nevertheless improves on current 5-year government census statistics, our future work will involve higher temporal resolution in order to explore these effects even further. Moreover, integrating predictive growth modelling in current urban planning systems could significantly help the government decide on the amount of cultural expenditure and along with geo-social network metrics, predict the impact of such investments.

One notable limitation of our research is that Foursquare data, as well as other social media data in general, is shown to be biased towards more central than peripheral parts of the city and often omits significant portions of the population who might be more deprived. Also, the Foursquare venues cannot represent the entire set of urban places exactly as how they present in the city, but they undeniably provide us an inspirational view to understand cities in fine-grained spatio-temporal contexts. In our case, Foursquare data make it possible to observe culture-led regeneration policies taking place and detect their effects. Furthermore, it has been shown to have the interesting potential of uncovering gentrification processes in the city where more affluent residents tracked by Foursquare might replace the local deprived population (Hristova et al., 2016). Additionally, cultural activities have been widely

associated with benefits related to health, well-being, and prosperity (Guetzkow et al., 2002). One possible future application of our research is to the location-based application domain where recommendation systems can make use of geo-social and public data to help direct users to areas of 'cultural buzz' (Currid and Williams, 2009). As a whole, our work aims to shed light on the practicality of such future applications and invites further research into the exploration of culture-led urban development using digital traces.

Chapter 4

Cultural patterns extraction & modern urban planning

Different from the previous chapter that explores ways to assist urban governments in smart financial allocation, this chapter discusses how GSN data can be utilised and how new algorithms tailored for urban settings can be devised to optimise urban facilities planning.

Chapter outline. In this chapter, we exploit large-scale spatio-temporal footprints of users in the online social network WeChat to obtain patterns of urban cultural interactions in the city of Beijing. We first utilise a latent Dirichlet allocation (LDA) (Blei et al., 2003) based method that takes as input check-ins at venues over time to identify clusters of mobile users with similar cultural profile patterns in Section 4.5. Having identified the geographic spread of the check-in activity of each user cluster, we then estimate the primary locations of users in a cluster, in terms of home or work, and evaluate the degree of accessibility to cultural venue resources for a user in Section 4.6. Next, we empirically demonstrate how the supply-demand balance of cultural services in a city can be highly skewed and pinpoint in cartographic terms the areas in the urban territory where supply could improve through appropriate investment in Section 4.7.

4.1 Introduction

The opportunity to enjoy cultural and entertainment activities is an essential element of urban life. Cultural spending is a considerable fraction of the annual budget of a city, as it is a key catalyst of social life and an important quality of life indicator in urban environments. Typically, in megacities like Beijing, London or New York, the financial resources allocated to support cultural events and related urban development (e.g. museum construction and

maintenance) is in the order of tens of millions of dollars annually (local government, 2015). Furthermore, in today's inter-connected world, the opportunity to experience a diverse set of cultural activities is amplified; citizens equipped with mobile devices can utilise a wide range of mobile applications and services that inform them on on-going social and cultural events, as well as on the best areas in the city to explore culture. In this setting, urban culture explorers are also mobile users who emit digital traces of where and when they are participating in cultural events. The new window onto the cultural life of a city, opened by the availability of novel sources of mobile user data, paves the way to the development of new monitoring technologies of urban cultural life. This can power evidence-based cultural policy design and optimise urban planning decisions. As an example, by tracking the interactions of mobile users with cultural venues across space and time, fine-grained indicators of areas in the city where there is lack, or excessive supply, of cultural resources can be devised. In more detail, we make the following contributions:

- **Cultural patterns extraction from check-in data:** We obtain raw representations of time-stamped check-in data at WeChat venues and use those as input to an extended version of the standard LDA model that takes into account temporal information (TLDA) on when venues of certain cultural categories are visited. We evaluate the performance of the model with a novel metric of coherence between top cultural venue categories observed in a cluster of users and the time periods of activity that are characteristic of each pattern. Using the metric we optimise the TLDA model and identify the presence of six latent cultural patterns in Beijing, each of which bears characteristic spatio-temporal manifestations of user activity at cultural venues. The description of the TLDA model and the corresponding data representations are described in Section 4.5.
- **Determining cultural demand patterns of users spatially:** Having obtained the latent cultural patterns through the TLDA model, we then exploit the frequencies of user check-ins across space to determine the *levels of demand* of cultural activities in different areas of the city. In this context, we present POPTICS, a user-personalised version of the OPTICS algorithm (Ankerst et al., 1999), used here to identify clusters of user activity hotspots. These primary locations of user activity are the means to quantify demand levels for cultural resources spatially. Overall, the output of this process corresponds to a set of heat maps depicting the intensity levels of user activity geographically for each of the cultural pattern emitted by the TLDA model. We consider such intensity levels to reflect user-driven demand of cultural resources geographically and present our results in Section 4.6.
- **Identification of areas that lack cultural offering:** In addition to obtaining spatial descriptions of the demand levels for each cultural pattern observed in the city, we determine

the *supply levels* of cultural resources using the spatial distribution of cultural venues and users' check-ins belonging to each TLDA pattern as input. For each region in the city, we obtain a demand-supply ratio (DSR), high values of which are indicative of an area lacking cultural venues, whereas low values suggest oversupply of cultural establishments in a region. In Section 4.7 we generate precision maps of such supply and demand patterns for each cultural pattern and demonstrate how users who live in high-DSR neighbourhoods but adhere to a specific cultural pattern travel longer distances in the city to access the resources they are interested in. The latter is an indication of how lack of resources in an area translates to larger travel distance for its residents.

In summary, the methodology put forward in the present chapter paves the way to novel data-driven urban cultural planning schemes that dynamically adapt to the profiles of residents active in city regions. Such schemes exploit the rich characteristics of new digital data sources and have the potential to improve planning decisions that currently tend to rely solely on residential population density and are agnostic to both personalised user preferences and fine-grained temporal signatures of user behaviour.

4.2 Related work

The flourishing of location technology services (Bauer et al., 2012; Hasan and Ukkusuri, 2014) and the distinct advantage of topic modeling (Blei et al., 2003) in uncovering latent patterns have encouraged researchers to employ them as data sources and methods, respectively, in large-scale urban human mobility studies. From the perspective of the users, Jiang et al. (2015) proposed a topic-based model to recommend personalized points of interest (POIs) for tourists. Yuan et al. (2013) leveraged topic modeling to learn lifestyles for individuals based on their digital footprints and social links. Kurashima et al. (2013) established a geo-topic model for location recommendation with the consideration of the activity area of users. Yin et al. (2013) proposed a recommender system for both venues and events according to personal preferences. From an urban planning perspective, Bauer et al. (2012) analysed the comments published alongside Foursquare check-ins to detect the topics in neighbourhoods. Yuan et al. (2012) utilised human mobility flows between POIs to identify land use functions for urban regions.

Even though topic modeling has been shown as an effective unsupervised approach to discovering latent mobility patterns in existing literature, it still has some limitations. Firstly, urban activities are time-sensitive that citizens show strong time preferences for different activities (Bauer et al., 2012; Hasan and Ukkusuri, 2014). However, these temporal behavioral patterns in urban daily life can hardly be captured since classical topic modeling, such as LDA,

does not have time factors integrated (Chen, 2017). Some existing works (Amoualian et al., 2016; Blei and Lafferty, 2006; Wang et al., 2012) considering temporal dynamic information in topic modeling cannot characterise the stable visiting time preferences of users and group them accordingly. Secondly, evaluation methods are missing in the application of topic modeling to human mobility study. The number of topics was either selected without testing (Bauer et al., 2012; Hasan and Ukkusuri, 2015; Yuan et al., 2013) or by using traditional evaluation methods in text mining cases directly (Hu et al., 2013). Thirdly, few published research findings obtained through topic modeling can be used in practical urban planning. And finally, current publications focus on mining the patterns of all types of urban activities in general, without looking at particular subgroups of people or urban venues.

Our research is different from existing works in the following aspects: i) We integrate time parameter with a standard topic model to gain insights into temporal behavioural patterns along with the cultural activities of people. ii) We devise a novel evaluation method for the temporal topic model to value its performance quantitatively. iii) We detect the heterogeneity on various cultural demand and supply levels in the city environment. To the best of our knowledge, this is the first methodological work for urban cultural patterns mining using large-scale location services data in practical urban cultural planning.

4.3 Approach at a glance

The research framework outlined in Figure 4.1 includes two main phases: cultural patterns extraction, followed by the spatial distribution of cultural patterns. These two phases as a whole provide an integrated approach to optimising the cultural resources allocation in cities and refining the urban cultural planning scheme by using GSN data as input. To this aim, we utilise three models tailored to the characteristics of spatio-temporal social network data.

Cultural Patterns Extraction. Cultural pattern extraction begins with raw check-in data processing. We build the tuple (u, v, t) to denote each qualified check-in record, indicating user u visited venue category v at time t . The whole set of check-ins formatted in this way is stored in a data cube (U, V, T) displaying the check-in history of users in the city before being sent to the TLDA model as input data. At this stage, temporal coherence value (TCV) measurement is designed to evaluate the performance of TLDA and choose the optimal input parameter K , which denotes the number of latent patterns Z . The TLDA model applied in this research enables us to essentially group U , T , and V into K cultural patterns and discover the associations between users, time, and various cultural activities.

Spatial Distribution of Cultural Patterns. Building on the output of TLDA, we further devise the POPTICS algorithm and the demand-supply interaction (DSI) model to get a

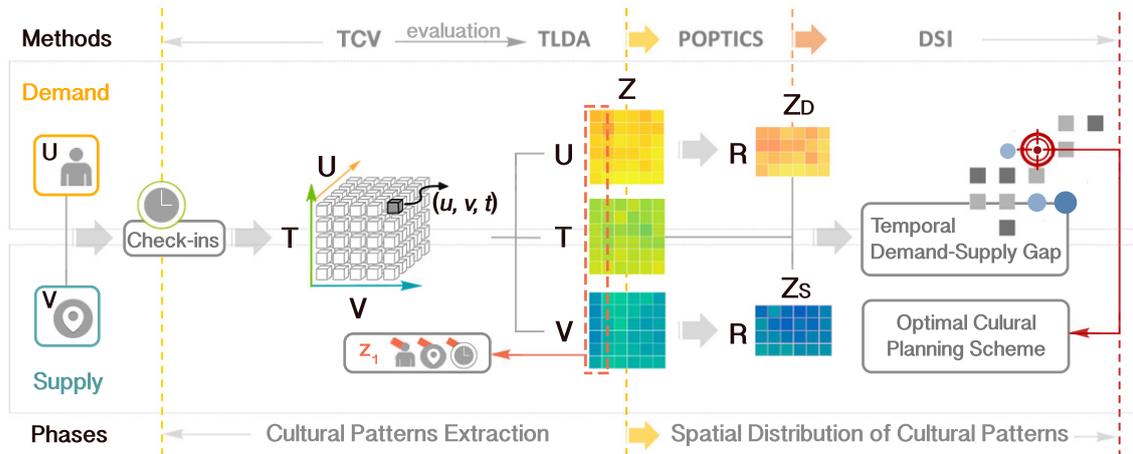


Fig. 4.1 Description of main phases and algorithms utilised in the research.

general picture of how cultural check-ins and venues belonging to different patterns distribute spatially in the city, so as to uncover the gap between user-side demand and venue-side supply for different cultural patterns. Our primary assumption here is that the cultural demand of users and supply capability of cultural venues are heterogeneous across regions that by mapping the demand-supply ratio, we can infer whether a type of cultural resources is insufficient or over-supplied relatively in the city.

4.4 Datasets

WeChat¹ is a mobile social application launched by Tencent in January 2011 and has currently become the most popular mobile instant messaging application dominating the Chinese market. According to the latest data report published by the WeChat team at the Tencent Global Partners Conference (WeChat, 2017), by the end of September 2017, it had 902 million average daily logged-in users sending 38 billion messages in total every day. WeChat's 'Moments' function, which is an equivalent of Facebook's timeline feature, allows users to share their status or anything of interest via photos, text, videos, or web links with their contacts. When a user posts on Moments, a timestamp will be generated automatically. Additionally, the user has the option to share their current place from a list of pre-selected locations nearby in WeChat. This real-time location-based service provided by 'Moments' depicts the routine lives of users at a fine-grained spatio-temporal scale, and provides a precious dataset that enables us to discover urban activity patterns.

¹<http://www.wechat.com/en>

Advantages of WeChat Dataset. For our analysis, the WeChat dataset possesses a number of natural advantages:

- *High population coverage levels in cities.* According to a survey by Tencent in September 2015, 93% of the population in the first-tier cities in China² were WeChat users. As for the selected case study in this research, Beijing has 21,136,081 monthly active users³ according to our statistics, making up 97.4% of the residents⁴ in September 2017. This high popularity makes the observation of users' mobility patterns through the lens of WeChat a reliable proxy to the real-world mobility of Beijing residents.
- *Wide age distribution of users.* Compared with many other social media services, WeChat penetration among middle-aged and senior users is relatively high. Although people born in the 80s and 90s are still the major groups, the monthly active WeChat users between 55-70 years old in September 2017 were approximately 50 million. In other words, WeChat is a representative social media data source reflecting a wide range of age groups in the general population.
- *Private circle visibility.* With an in-group design, the social circle of a user on WeChat is mainly comprised of relatives, friends, and colleagues who have a close relationship with him in life. Additionally, WeChat empowers the user with the right to control over exactly who has access to every single post on his 'Moments'. This powerful feature creates a secure and private environment that encourages WeChat users to communicate freely and share their check-ins.

Moments Check-in Data. The main dataset employed for this study comes from the anonymized logs of complete WeChat Moments posting activities. We collected all check-in records with venue information provided in Beijing during the four months of October 2016, January, April, and July 2017. In total, there were 56,239,429 check-ins created by 9,517,175 users at 2,428,182 venues. For each check-in, information about user ID, timestamp, coordinates, and POI category is provided. Through the lens of this dataset, we are thus able to observe who visited where at what time for what purpose. We obtained IRB approval at the University of Cambridge to work on the data for this research project.

²Beijing, Guangzhou, Shanghai and Shenzhen

³users who have logged into WeChat within the month

⁴The permanent population of Beijing is 21.907 million by the end of 2017 according to Beijing Municipal Bureau of Statistics. <http://www.bjstats.gov.cn/tjsj/yjdsj/rk/2017>

4.5 Cultural patterns extraction

In this section, we present how urban cultural patterns can be extracted from Moments check-in records. Before discussing the issue in more depth, we first state the meaning of culture and culture-related terms in this work.

Cultural Venues. Cultural venues are defined as urban places of arts, media, sports, libraries, museums, parks, play, countryside, built heritage, tourism, and creative industries, following the line set by the Office of the Deputy Prime Minister in Regeneration through Culture, Sport and Tourism (ODPM, 1999; Zhou et al., 2017). Based on this definition, 37 categories of WeChat cultural venues are selected in this research.

Cultural Check-ins. The check-in activities taking place at the cultural venues are called cultural check-ins.

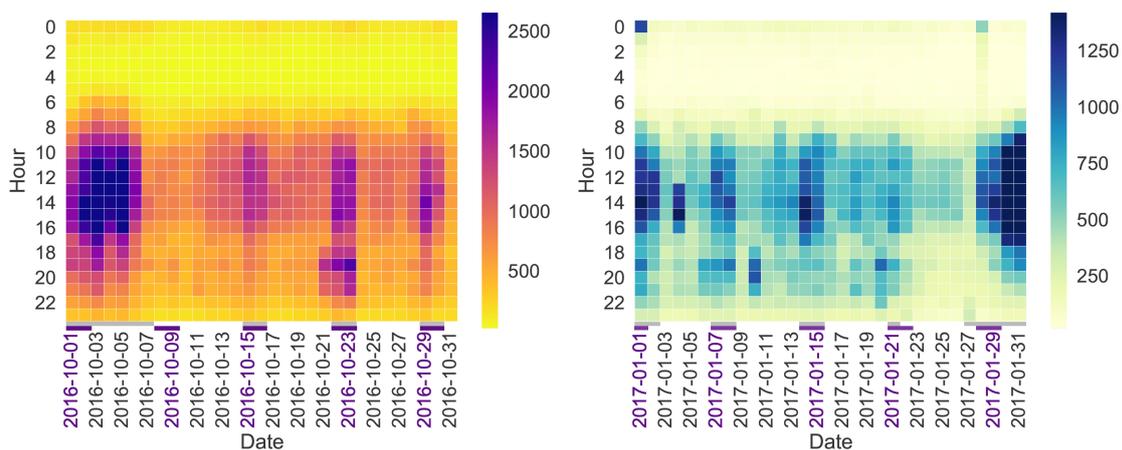
Cultural Fans. To classify individual cultural patterns, we only focus on users who have a certain minimum level of cultural check-ins during the observation time and call them cultural fans.

In the following subsections, we illustrate the necessity of considering temporal factors in urban cultural patterns mining, before introducing the TLDA model which integrates temporal characteristics with a particular subgroup of cultural activities, followed by the introduction of a novel evaluation method for the TLDA.

4.5.1 Temporal factors extraction

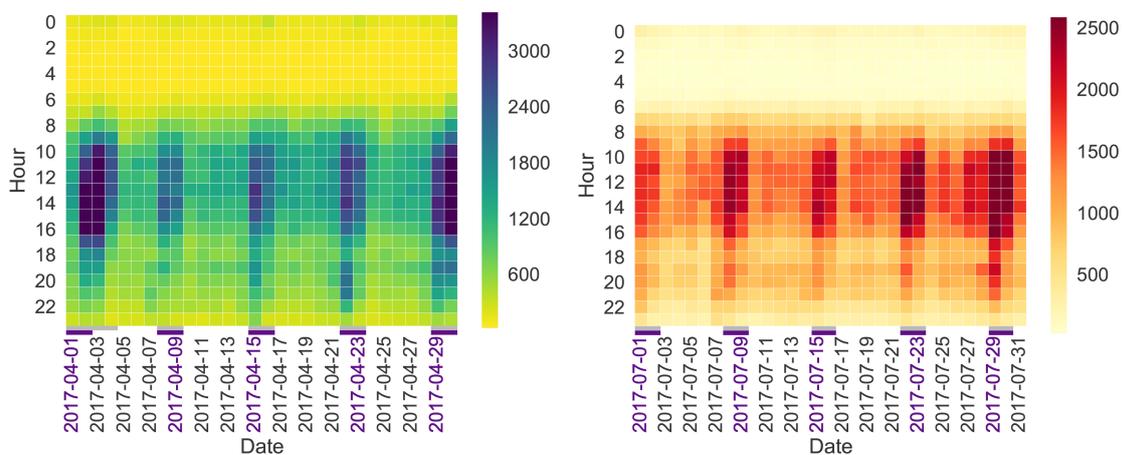
In Figure 4.2 we show the temporal cultural check-in distribution of the four selected months in Beijing. The purpose of creating these heat maps in a calendar format is to present the cultural check-in frequency by date and hour corresponding to four seasons chronologically. In each subfigure, the date is plotted along the horizontal axis with hour appearing on the vertical axis to unveil cultural visiting patterns associated with temporal factors, which will later be explored in further depth.

It can be seen from Figure 4.2 that hourly and weekly cultural visiting patterns are both significant in general. For all the seasons, the least likely hours for cultural check-in creation is during the night, from 0 to 6am. After that, the hourly frequency of cultural check-ins increases gradually and stays at a relatively high level during the daytime. On a daily basis, two peak periods can be recognized, among which, the highest one lasts for around seven hours from 10am to 4pm while a lower peak appears between 7pm to 9pm. As for weekly patterns, we can find that the check-in frequency is significantly higher for weekends (purple x-axis labels and bars) than weekdays.



(a) Cultural check-ins in October 2016.

(b) Cultural check-ins in January 2017.



(c) Cultural check-ins in April 2017.

(d) Cultural check-ins in July 2017.

Fig. 4.2 Temporal cultural check-in distribution of four months in Beijing. The x-axis labels and bars above coloured purple represent weekend days. Grey bars represent rest days (public holidays or weekends).

It is also noticeable that hourly and weekly patterns are more evident in April and July, while less regular in October and January. This observation can be explained by comparing our calendar heat maps with the public holidays in China. At the beginning of October 2016, we can see a dramatically high cultural check-in frequency for six continuous days, when people have one week off for the National Day. It is the longest holiday after the Chinese New Year and is also called 'Golden Week' for people to reunite with families and take trips. Then the second graph follows and presents the situation in the 'New Year' month. It can be discovered that both the first cells corresponding to the New Year's Day (01/01/2017) and the Spring Festival (28/01/2017) have distinctively higher values compared to the rest. In addition, during the week before the Chinese New Year, the number of cultural check-ins is much smaller than that in other weeks. Moreover, on 27/01/2017, the day before the Lunar New Year, the check-in frequency is particularly low, forming a sharp contrast with the following Spring Festival week. This is because typically, Chinese people prefer to stay at home with families before the New Year's Eve, waiting for the coming new year, but would like to hang out with friends in the next few days.

4.5.2 Temporal latent Dirichlet allocation

The classical LDA (Blei et al., 2003) is a hierarchical Bayesian model that has been shown as an effective unsupervised learning method in discovering structural daily routines (Farrahi and Gatica-Perez, 2011; Hasan and Ukkusuri, 2014; Huynh et al., 2008; Sun and Yin, 2017). However, the original LDA approach is built based on the 'bag-of-words' assumption (Hasan and Ukkusuri, 2014), which means that it only considers the number of times each word appears in a document, without involving any temporal consideration (Chen, 2017). According to the observations made in the previous subsection, the periodicity of cultural check-ins in different levels of temporal granularity is so obvious that we choose to explicitly incorporate it to the LDA model and thus utilise temporal latent Dirichlet allocation (TLDA) Chen (2017). The TLDA is an extended version of LDA that integrates time factors into the original model, so as to uncover multiple associations between users, urban activities, and their corresponding temporal characteristics. Specifically, the TLDA model is an unsupervised machine learning algorithm that characterises each user as a mixture of patterns with both venue and temporal preferences considered. In the TLDA model, users' check-in data containing timestamp and POI information are taken as the input. Here the set of a user's historical check-ins as a whole is viewed as a document, where POIs and time slots associated with the check-ins are regarded as words analogously. As a generative probabilistic model in nature, TLDA models the underlying generative process whereby the user's check-in set is created to infer her probability distribution over several latent patterns.

Besides TLDA, several other studies have also attempted to add a temporal dimension to the LDA model. For example, Blei and Lafferty (2006) developed a dynamic topic model to capture the evolution of topics from a sequential collection of documents. Similarly, to model the topic transitions in temporally sequenced documents on social media platforms, Wang et al. (2012) proposed TM-LDA. Considering the topic dependency between consecutive documents, Amoualian et al. (2016) developed another extended version of LDA, called streaming-LDA. The main reason TLDA is adopted in this study rather than other extended versions is that our aim here is to group users by capturing their stable temporal preferences, while other models focus more on the dynamic changes of topics. Furthermore, in our case, each user's check-in records as a whole are seen as a document. There is no sequential relationship between documents, leading to other extended versions of LDA introduced above inappropriate to the task.

The graphical model representation of TLDA is shown in Figure 4.3, where the lower part integrated by red arrows is the addition of TLDA. In the figure, circles represent parameters and the meaning of which are described in Table 4.1. In general, the framework of the TLDA model consists of four hierarchical layers, including a user layer, a time layer, a venue category layer, and a cultural latent pattern layer. The cultural pattern layer is the key layer that links the other three. Like its predecessor, TLDA is also a generative model, the goal of which is to find the best set of latent variables (cultural patterns) that can explain the observed data (cultural check-ins by users) (Hasan and Ukkusuri, 2014). To generate a cultural venue category, the pattern distribution of the corresponding user is sampled from a prior Dirichlet distribution parameterized by α , $\theta_u \sim Dir(\alpha)$. In a similar way, the pattern distribution of time is sampled from a prior Dirichlet distribution parameterized by γ . Based on these, the pattern assignment Z_{ut} of the venue category is drawn from a multinomial distribution $Z_{ut} \sim Multi(\theta_u, \phi_t)$.

Notation	Description
α	Dirichlet prior over the pattern-user distributions
β	Dirichlet prior over the venue-pattern distributions
γ	Dirichlet prior over the pattern-time distributions
θ_u	pattern distribution of user u
ϕ_z	venue distribution of pattern z
ϕ_t	pattern distribution of time t
z_{ut}	pattern of venue of user u at time t
v_{ut}	venue of user u 's check-in at time t

Table 4.1 Notation and description of parameters in TLDA.

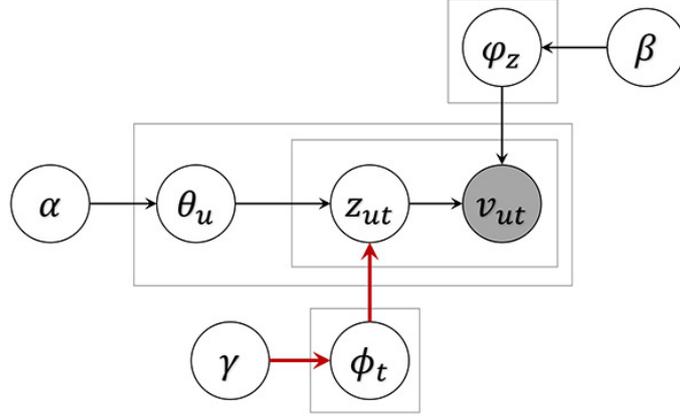


Fig. 4.3 Graphical model for TLDA.

$$P(z_{ut}|\alpha, \gamma) = \int_{\theta_u} P(z_u|\theta_u)P(\theta_u|\alpha) \int_{\phi_t} P(z_t|\phi_t)P(\phi_t|\gamma) \quad (4.1)$$

Then, the venue category is generated by sampling $V_{ut} \sim \text{Multi}(\varphi_z)$. φ_z specifies the venue distribution of pattern z , which is drawn from a prior Dirichlet distribution parameterized by β .

$$P(v_{ut}|z_{ut}) = \int_{\varphi_z} P(v_{ut}|\varphi_z)P(\varphi_z|\beta) \quad (4.2)$$

After that, we estimate the maximum likelihood of v_{ut} of u at time t by integrating θ_u , ϕ_t and φ_z , as shown in the following equation.

$$P(v_{ut}|\alpha, \gamma, \beta) = \int_{\theta_u} \int_{\phi_t} \int_{\varphi_z} P(v_{ut}, z_{ut}, \theta_u, \phi_t, \varphi_z|\alpha, \gamma, \beta) \quad (4.3)$$

In the last step, we use the Gibbs sampling algorithm (Griffiths, 2002) to estimate the probability distributions of pattern-user, venue-pattern, and pattern-time as follows:

$$P\left(z_{ut}(v) = k | \overrightarrow{z_{-ut}}, \overrightarrow{V_{-ut}}\right) \propto \frac{n_{u,-ut}^k + \alpha_k}{\sum_{k=1}^K (n_{u,-ut}^k + \alpha_k)} \frac{n_{t,-ut}^k + \gamma_t}{\sum_{k=1}^K (n_{t,-ut}^k + \gamma_t)} \frac{n_{k,-ut}^v + \beta_v}{\sum_{v=1}^V (n_{k,-ut}^v + \beta_v)} \quad (4.4)$$

where $\overrightarrow{V_{-ut}}$ is the probability distribution of venue categories over K patterns except the current one. $n_{u,-ut}^k$ is the total number of venue categories from user u 's check-in history assigned to pattern k , not including the current one. $n_{t,-ut}^k$ is the number of time slices assigned to pattern k , not including the current one. And $n_{k,-ut}^v$ is the number of venues assigned to pattern k except the current one.

4.5.3 Evaluation of TLDA

TLDA is an unsupervised learning method that requires a pre-specified number of patterns K . As the temporal layer is integrated into the TLDA, conventional LDA evaluation approaches cannot be used directly to the extended model. To handle this obstacle, we design temporal coherence value (TCV) to evaluate the performance of the TLDA model and to find the optimum number of cultural patterns in the city. Inspired by the coherence value (CV) measurement (Röder et al., 2015), the TCV introduced in this chapter is used to measure the coherence level between top cultural venue categories and time periods within each pattern, before averaging them to evaluate the overall performance of the TLDA model.

How the TCV works step-by-step is depicted in Algorithm 1. To run the model, we need three inputs: from the output of TLDA TV , we obtain 1) top venue categories V^* and 2) top time periods T^* in each cultural pattern; 3) all the check-in activities SW . It should be stated that to eliminate the influence of variance between users' check-in frequencies, the input SW used in TCV is constructed by a sliding window that moves over the original check-ins of all the users.

Algorithm 1 Temporal Coherence Value Calculation

Input: $VT([V_1^*, V_2^*, \dots, V_K^*], [T_1^*, T_2^*, \dots, T_K^*]), SW$

Output: \bar{m}

```

1: initialize  $S = set()$ 
2: for each  $(V^*, T^*)$  in  $VT$  do:
3:   for each  $v^*$  in  $V^*$  do:
4:      $S_{set}^{one} = \{(v^*, V^*, T^*) | v^* \in V^*\}$ 
5:      $S \leftarrow S + S_{set}^{one}$ 
6: for each  $S_{set}^{one}$  in  $S$  do:
7:   for each  $t_j^*$  in  $T^*$  do:
8:      $\vec{w}(j) = NPMI(v^*, t_j^*)^\tau$ 
9:      $\vec{W}(j) = \sum_{i=1}^I NPMI(v_i^*, t_j^*)^\tau$ 
10:     $m_q = \cos(\vec{w}_q, \vec{W}_q)$ 
11: return  $\bar{m} = \frac{\sum_{q=1}^Q m_q}{Q}$ 

```

In the algorithm, we firstly define a segmentation S_{set}^{one} for each top venue category v^* in each pattern as equation 4.5 shows. Here S is used to denote the list of S_{set}^{one} . The total number of S_{set}^{one} in S is denoted by Q .

$$S_{set}^{one} = \{(v^*, V^*, T^*) | v^* \in V^*\} \quad (4.5)$$

For each S_{set}^{one} , we calculate the *normalised pointwise mutual information* (NPMI) (Bouma, 2009) for v^*-T^* vector and V^*-T^* vector, respectively. The j th element of the time vector t_j^* and venue category v^* has the NPMI:

$$\vec{w}(j) = NPMI(v^*, t_j^*)^\tau = \left(\frac{\log \frac{P(v^*, t_j^*) + \varepsilon}{P(v^*) \cdot P(t_j^*)}}{-\log(P(v^*, t_j^*) + \varepsilon)} \right)^\tau \quad (4.6)$$

where $P(v^*, t_j^*)$ is the probability of the co-occurrence of v^* and t_j^* , while $P(v^*)$ and $P(t_j^*)$ mean the probabilities of v^* and t_j^* ($t_j^* \in T^*$), respectively. ε is added to avoid logarithm of zero, and an increase of τ gives higher NPMI values more weight. After calculating the NPMI value for each venue category according to the above formula, we aggregate them to obtain the j th element of the time vector of V^* by the following equation.

$$\vec{W}(j) = \sum_{i=1}^I NPMI(v_i^*, t_j^*)^\tau \quad (4.7)$$

where v_i^* represents the i th venue category in V^* .

Cosine similarity is then calculated between pairs of context vectors \vec{w}_q and \vec{W}_q to obtain the coherence score m_q for each S_{set}^{one} by formula 4.8 before we average over all top venue categories in patterns to get the final TCV \bar{m} for the model through equation 4.9.

$$m_q = \cos(\vec{w}_q, \vec{W}_q) \quad (4.8)$$

$$\bar{m} = \frac{\sum_{q=1}^Q m_q}{Q} \quad (4.9)$$

The higher the TCV score, the better the clustering result of the TLDA model.

4.6 Refined urban cultural planning

Different from current cultural planning frameworks, which mainly consider the population of urban areas when allocating cultural resources, we propose a refined urban cultural planning scheme based on the results of the TLDA. Our core viewpoint here is that urban regions are heterogeneous in terms of cultural demand and supply capability, which should not be treated uniformly. By employing the TLDA model, we can group users according to their cultural tastes, and cluster cultural venues based on their similarities derived from human mobility behaviours. Then, after aggregating the users and cultural facilities into urban regions, we can get an idea about how the cultural demand and supply are distributed spatially in the

city for different cultural patterns. Moreover, through learning the supply-demand balance across regions, we can detect the areas where particular cultural services are needed, and we are thus able to provide city government with a priority list when the financial budget is compiled for culture-related planning.

4.6.1 Demand range determination

In this part, we explore a way to determine the main activity range of individual users as reflected in their historical check-ins. More specifically, we aim to detect valid visits for the user, map the active ranges of areas which he visits frequently, and finally, determine the centre and radius for these active ranges. Among existing clustering methods, we find OPTICS (Ankerst et al., 1999) a suitable approach for our problem. OPTICS is an algorithm for finding meaningful density-based clusters in spatial data (Kriegel et al., 2011). This method requires two parameters as input: the maximum radius to consider, and the least number of points to form a cluster. As check-in frequencies of users can vary greatly, setting a common number of minimum points for all users is inadequate. Considering this limitation, we propose a modified version as Algorithm 2 shows, named POPTICS that defines a different threshold for each user separately.

In POPTICS, we collect all the N locations of check-ins $L_u = [l_1, l_2, \dots, l_N]$ for each user u . Here places with more than one check-in records are counted repeatedly, as they are more important in the user's life and should be given higher weights. An input parameter η is set to denote the percentage of a user's total check-in locations L_u being considered in the calculation of core distances. Here ηL_u is varied for different users. The core distance for location l_i is defined as the Euclidean distance between l_i and the ηL_u -th nearest point to it, as shown in the following function.

$$CD(i) = \min_{\eta L_u} Dist(i, j) (j = 1, 2, 3, \dots, N) \quad (4.10)$$

After calculating the core distance, for location l_o , we define the reachability distance from l_o to l_i as:

$$RD(o, i) = \max(CD(o), Dist(o, i)) \quad (4.11)$$

According to the reachability distances, an ordered list of locations is generated. Then, to find meaningful cluster(s) of locations and detect outliers, a threshold of maximum reachability distance, rd_{th} is set according to the score derived from formula 4.12. The lower the score, the better the chosen rd_{th} is. The group of all valid points is denoted by RD^* as equation 4.13 shows.

Algorithm 2 POPTICS

Input: $L_u = [l_1, l_2, \dots, l_N], \eta$ **Output:** cluster groups of locations $GL = [L_1^*, L_2^*, \dots, L_r^*]$
cluster of locations $L_i^* = [l_{i1}, l_{i2}, \dots, l_{ig}]$

```

1: initialize  $CD=list()$ ,  $RD=list(maxdis)$ ,  $RD(0)=0$ ,
2:  $seeds=1, 2, \dots, N$ ,  $ind=1$ ,  $order=list()$ ,  $GL=list()$ ,  $tmp\_L=list()$ 
3: for  $i = 1, 2, 3, \dots, N$  do:
4:    $CD \leftarrow CD + \min_{\eta} Dist(i, j)$ 
5: while  $seeds! = \{\}$  do:
6:    $seeds.move(ind)$ 
7:    $order \leftarrow order + ind$ 
8:   for each  $ii$  in  $seeds$  do:
9:      $cur\_rd \leftarrow \max(CD(ind), Dist(ind, ii))$ 
10:     $RD(ii) \leftarrow \min(RD(ii), cur\_rd)$ 
11:    $ind \leftarrow \{min - index(RD_{ii}) | ii \in seeds\}$ 
12:  $rd_{th} \leftarrow \min std(RD^*) \frac{N}{len(RD^*)}$ 
13: for  $ii$  in  $order$  do:
14:   if  $RD(ii) < rd_{th}$  then:
15:      $tmp\_L \leftarrow tmp\_L + l_{ii}$ 
16:   else:
17:      $GL \leftarrow GL + tmp\_L$ 
18:      $tmp\_L.clear()$ 
19:  $GL \leftarrow GL + tmp\_L$ 

```

$$score(rd_{th}) = std(RD^*) \frac{N}{len(RD^*)} \quad (4.12)$$

$$RD^* = \{rd_i | rd_i \in RD \text{ and } rd_i < rd_{th}\} \quad (4.13)$$

4.6.2 Demand-supply interaction model

In this part, we display the demand-supply interaction model (DSI). Through the TLDA model, each user u is labelled as a member of a particular cultural pattern z . Through POPTICS, the active centre μ and radius r of each user are determined. Also, the sub active range of locations belonging to pattern z can be drawn, which has the same centre μ and a smaller pattern radius r_{u_z} . These results allow us to link users with the urban areas and thus give us an indication of the demand levels of different cultural types in urban regions. The assumption here is that for a certain user u from a cultural group z , his demand for this type of cultural service is highest at the active centre μ , and decays as the distance increasing until r_{u_z} . The attenuation pattern is depicted by a Gaussian function. For a point x within user u 's pattern range r_{u_z} in the city, the demand influence it gets from u can be obtained by:

$$d_{u_z}(x) = Norm(x, \mu, r_{u_z}) = \frac{1}{\sqrt{2\pi r_{u_z}^2}} \exp\left(-\frac{(x - \mu)^2}{2r_{u_z}^2}\right) \quad (4.14)$$

The total demand in terms of pattern z for area x is the aggregation of influences from all users in pattern z .

$$D_z(x) = \sum_{u_z} d_{u_z}(x) \quad (4.15)$$

Next, we turn our focus to the supply of patterns. For each venue category v in pattern z , we calculate the supply capability of v spatially. If a user u once created check-in(s) at venue v , then the centre of the user μ is covered by the service range of venue v . We find all the users who had check-ins at v , calculate the distances between their centres with v . The average of the distances is set as the standard deviation for the attenuation distribution, and denoted by σ . Based on this assumption, the supply capability of cultural pattern z contributed by venue v in area x can be obtained through:

$$s_{v_z}(x) = Norm(x, v_z, \sigma_{v_z}) = \frac{1}{\sqrt{2\pi \sigma_{v_z}^2}} \exp\left(-\frac{(x - v_z)^2}{2\sigma_{v_z}^2}\right) \quad (4.16)$$

The total supply level of area x in the city in terms of pattern z can be achieved by the following equation.

$$S_z(x) = \sum_{v_z} s_{v_z}(x) \quad (4.17)$$

We then define a metric called demand-supply ratio (DSR) to capture the desirability level of a certain type of cultural service z in urban areas as:

$$DSR_z(x) = \frac{D_z(x)}{S_z(x)} \quad (4.18)$$

The higher the DSR is, the greater the need for particular cultural facilities, and the higher the priority of the area in the proposed urban cultural planning scheme.

4.7 Experiments

Until now, we have provided a holistic framework for urban cultural studies from extracting spatio-temporal cultural patterns to refining cultural planning for the city. Next, we will employ WeChat Moments data described in Section 4.4 and use Beijing as a case to present how these models can be applied jointly in practice.

4.7.1 Data preprocessing

We firstly filter cultural fans based on users with at least 20 check-ins at cultural venues during the observation time. After this procedure, our dataset shows that there are 1,082 cultural venues grouped in 37 categories, and 324,809 cultural check-ins created by 18,234 cultural fans in Beijing during the selected four months. Besides venue categorical information, we also represent a temporal label for each cultural check-in with three levels of identifiers, including month of the year, day of the week, and hour of the day. Following this form of expression, a user's check-in history can be represented as (User3, ((Concert hall, JulFri20), (Golf, OctSun10), (Yoga, AprFri18))), for example. A collection of all the cultural fans' check-ins constitutes the whole corpus, which is the input data for our analysis.

4.7.2 Cultural patterns extraction for Beijing

We first run the TLDA model with the optimum number of patterns K given by TCV. We adopt 7 numbers from 3 to 9 as candidates, run the TLDA for 100 iterations each, and get their respective average TCV scores as shown in Figure 4.4. We can see that 6 gives the

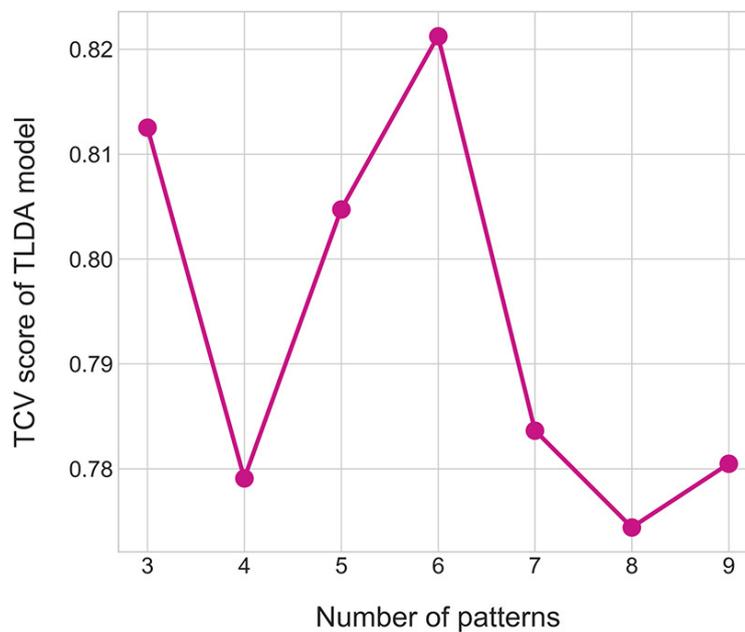


Fig. 4.4 TCV scores for different number of patterns.

best performance, suggesting that cultural behaviours in Beijing should be classified into six groups based on their categorical and temporal characteristics. We then select 6 as the value of K and run the TLDA model again to extract cultural patterns for the city. The main output resulting from this process are three matrices: pattern-user matrix, pattern-time matrix, and venue-pattern matrix. These matrices provide us the probability distribution of users, time periods, and venue categories over 6 patterns, respectively. These outputs as a whole tell us which group of people prefer to do what type of cultural activities at what time.

Figure 4.5 presents the probabilities of cultural venue categories over patterns. The colour of a cell indicates the probability a category belongs to a certain cultural pattern. As we can observe, 37 cultural venue categories can be clustered separately into six cultural groups based on which pattern has the highest probability. The name of the category is printed in black as a top venue if its probability is higher than 0.1. Otherwise, it is considered not a typical category for the pattern and is coloured in grey.

We also compute the cosine similarity between each pair of venue categories and show the results in Figure 4.6. As it can be seen in the figure, the clustering result of venues is desirable in the sense that all the within-group similarities are higher than 0.9, while most of the inter-group similarities are lower than 0.1.

To further evaluate the clustering performance of the model, we compare the CV scores (Röder et al., 2015) of TLDA with that of LDA. Since the LDA model does not contain the temporal part extended by TLDA, only the venue category clustering is evaluated. Again, we

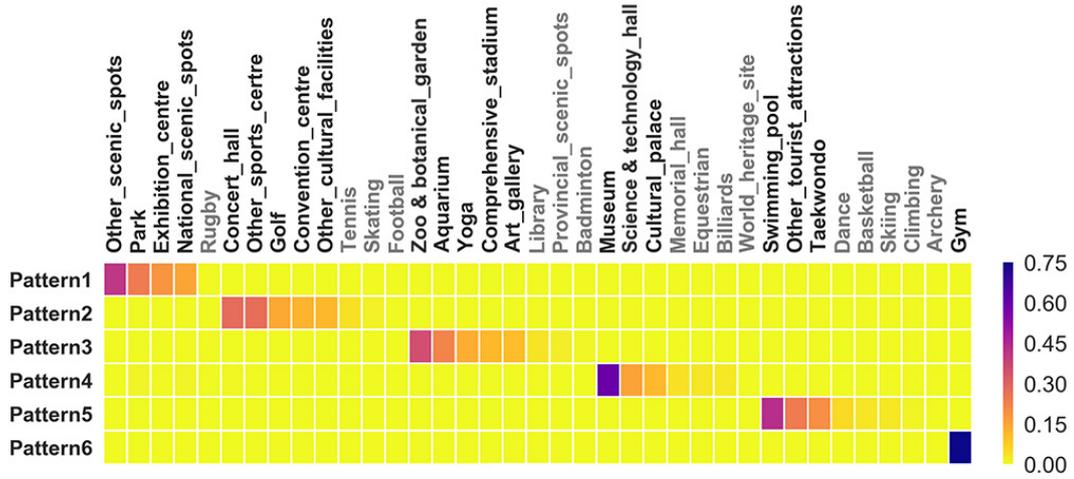


Fig. 4.5 Probabilities of venue categories over patterns.

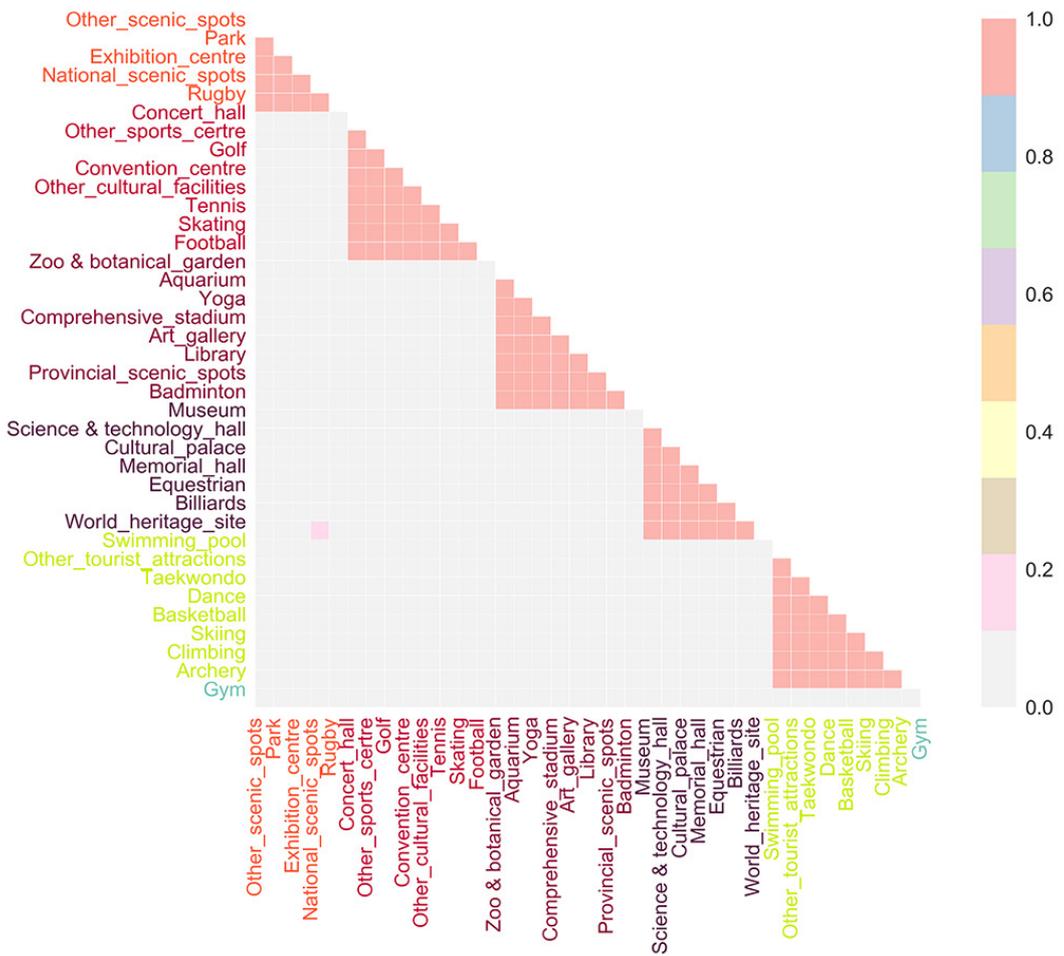


Fig. 4.6 Cosine similarity between venue categories.

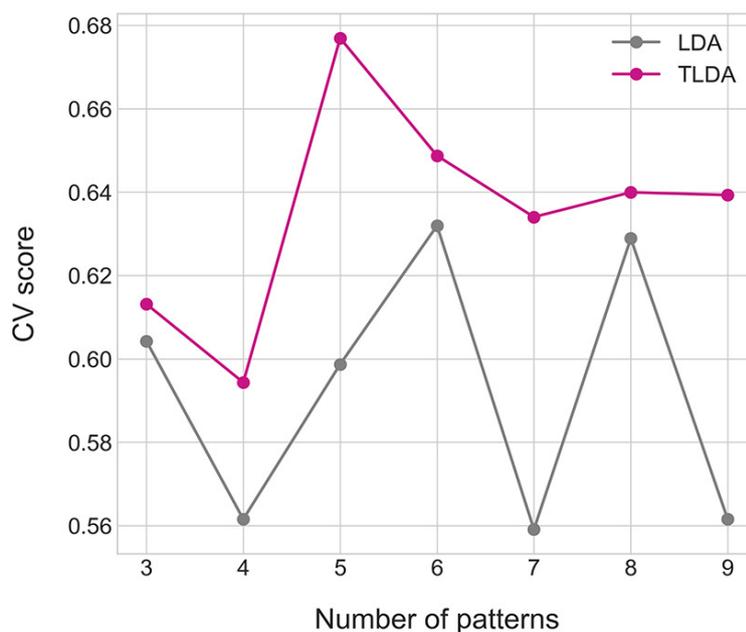


Fig. 4.7 Comparison between TLDA and LDA in CV. Only the coherent value of top venue categories Since the LDA model does not contain the temporal part extended by TLDA, only the venue category clustering is evaluated.

run the analysis iteratively with K setting as 3-9 and present the results in Figure 4.7. We can see that the TLDA model outperforms LDA in all the seven cases. This result indicates that the TLDA model not only enriches LDA by considering temporal features, it is also superior to classical LDA by generating more coherent topics.

In addition to venue category information, TLDA also provides us another point of view to learn about cultural patterns temporally. The temporal characteristics of the cultural patterns are displayed and compared in Figure 4.8. For comparison, values are shown in percentage to present the degrees to which time periods are representative of the patterns. To present the hourly patterns more clearly, we group 24 hours into five slots, which are morning (6-11am), noon (11am-14pm), afternoon (14-19pm), evening (19-24pm), and night (0-6am).

The results in Figure 4.7 and Figure 4.8 together reveal the key characteristics of the six cultural patterns detected in Beijing. We can observe that pattern one is composed of people who love travelling and prefer wide-open space. They do not like visiting scenic spots and parks during winter very much, perhaps due to harsher weather conditions. Compared to other groups, the first pattern has the highest percentage of activity in the morning and the lowest during night-time. This finding can be linked with what we can expect from real-life experience, that parks in Beijing are always full of people doing morning exercises. Music fans make up the majority of pattern 2. For this group of people, their cultural

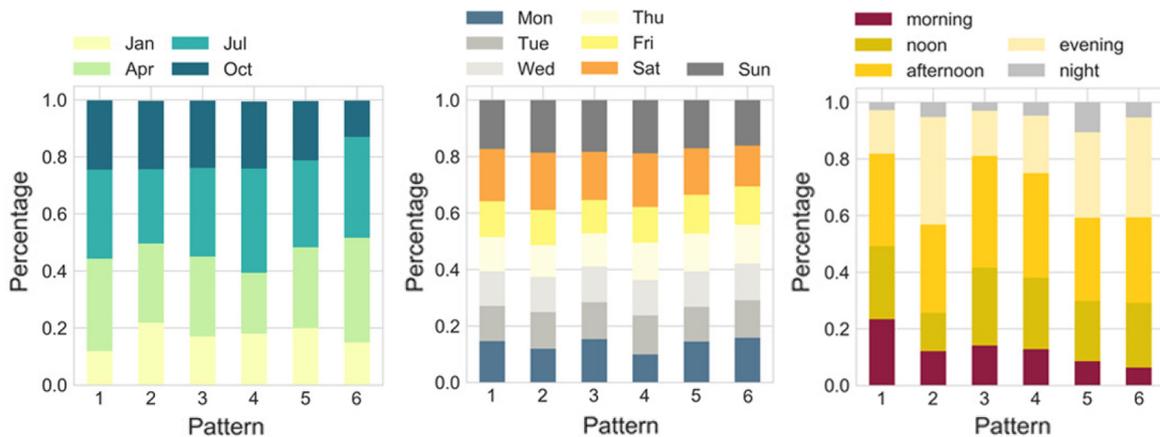


Fig. 4.8 Temporal characteristics of cultural patterns in Beijing.

visiting frequencies during the four seasons are rather balanced. However, on weekends and evenings, they present considerably a higher probability of being active when compared to their counterparts. This can be explained by the fact that concerts are usually being scheduled and attended in the evening hours. The third pattern are nature lovers who like plants and animals in particular. This group of people present hourly sensitive features as they prefer visiting cultural places during the daytime to evening or at night. Then, the fourth pattern corresponds to museum lovers. This cultural group are particularly active in summer and the afternoons. Moreover, they have the lowest percentage of activity on Monday compared to other patterns. This phenomenon can probably be explained by the fact that many museums are closed on Mondays. The fifth group are sports fans, especially swimming enthusiasts. These people have the highest percentage of night-time activity. The last group of people are gym lovers with a single top cultural venue category becoming prominent here with an extremely high percentage of 0.99. Spring and summer time is the most popular period for them to exercise. They do not like going to the gym in the morning, preferring evenings in most cases. Additionally, even though weekends make the greatest contributions to almost all the cultural patterns, they are not prominent in the case of pattern 6.

4.7.3 Refined cultural planning for Beijing

After uncovering the cultural patterns, we map how the demand and supply levels of each pattern are distributed spatially, and calculate the demand-supply ratio for urban areas. In this part of the analysis, we divide the city into 400m by 400m grids aligned with the latitude and longitude dimensions. Each cell is called an area, and the centroid of which is used to represent the cell's demand-supply balance.

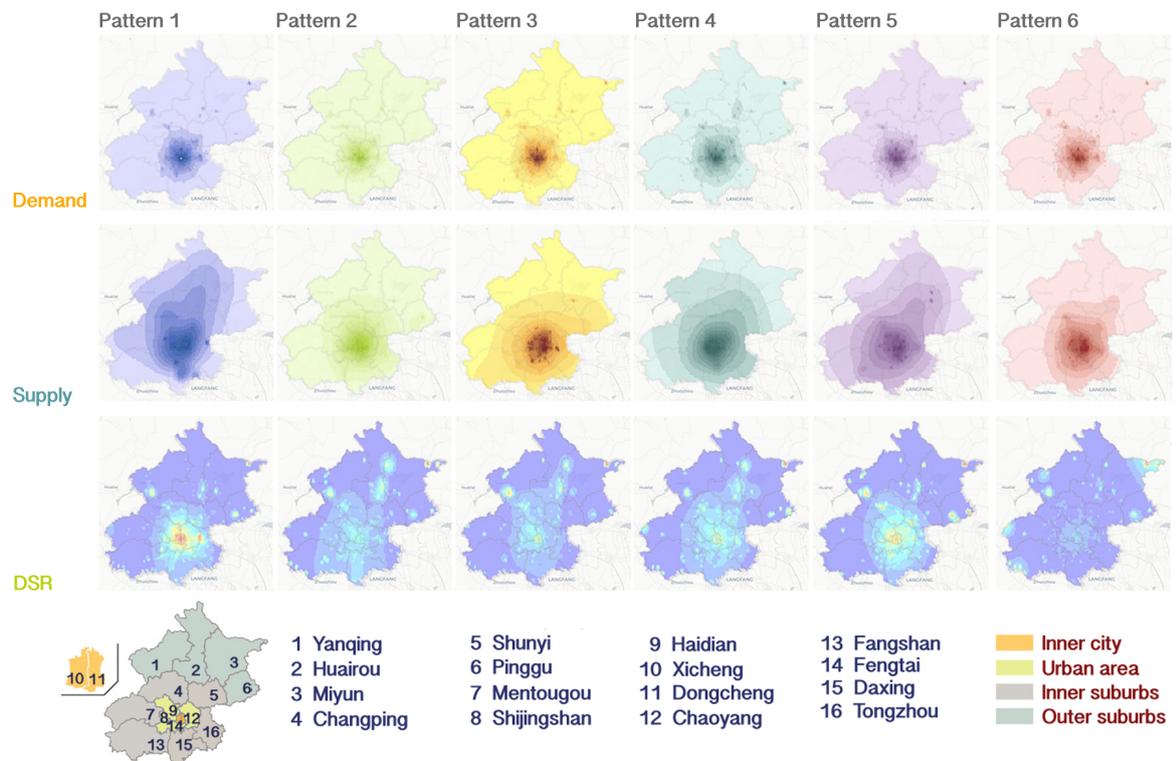


Fig. 4.9 Results of DSI model. For the first two rows, the darker the color, the higher the demand or supply level. For the DSR, red and blue represent high and low ratios, respectively.

From the demand perspective, we begin with the application of the POPTICS algorithm to find centres of activity range for six groups of cultural fans in their daily lives based on all the check-ins they created previously. Then, we collect a particular subgroup of cultural check-ins for each cultural fan according to the pattern he belongs to. We find his influential radius r_{u_c} , and calculate the demand value he contributes to his surrounding areas based on Equation 4.14. After the need of all users in a certain pattern group is aggregated by Equation 4.15, the overall demand for each cultural pattern can be obtained as the first row in Figure 4.9 suggests. With respect to the supply side, we calculate the supply capability of each area in terms of various cultural patterns according to equations 4.16 and 4.17. The supply levels categorised by patterns across the city are visualised in the middle row subfigures, followed by the final demand-supply ratio output shown in the last row in Figure 4.9. For the first two rows, the darker the colour, the higher the demand or supply level; while for the DSR, red and blue represent high and low ratios, respectively. As can be observed from the figure, the demand for different cultural patterns is distributed in a similar manner spatially. The highest demand areas cluster in the urban area between the 2nd and 5th ring road, while some hotspots shown in suburbs areas like Yanqing, Huairou, and Miyun Districts. When we look at the supply level, six patterns present a more heterogeneous behaviour. Although a general pattern can be discovered that the cultural supply capabilities show a decreasing trend from the city centre to the suburbs, this inequality is less obvious in patterns 1, 4 and 5. From our final results of DSR, we can find the inner city inside the former city walls (Xicheng and Dongcheng) is in great need of cultural services of pattern 1 and 5, like parks, swimming pools, and exhibition centres. For patterns 2 and 3, the demand-supply gaps of related cultural facilities are relatively equal within the city, while for the last pattern, the need for gym service is relatively greater in outer suburbs.

Through the demand-supply analysis above using the DSI model, we get priority lists of urban areas in terms of different types of cultural services according to the levels of need. To validate our model, we calculate the Pearson correlation between the DSR value and the average distance users need to travel for a particular kind of cultural services. The correlation coefficients for the six patterns are presented in Figure 4.10. From this figure, we can see that five patterns show high positive correlations except pattern one. The distinctive result observed in pattern one can probably be explained by its top venue categories. As an ancient city, many of the scenic spots in Beijing are historical relics, the locations of which are not decided by modern urban planning. The overall result of the correlation analysis suggests that users from areas in great needs of a type of cultural services generally have to travel longer distances to be served. It further indicates that the facilities in the users' surrounding areas are not enough to fulfil their needs, and thus provides evidence for our results that

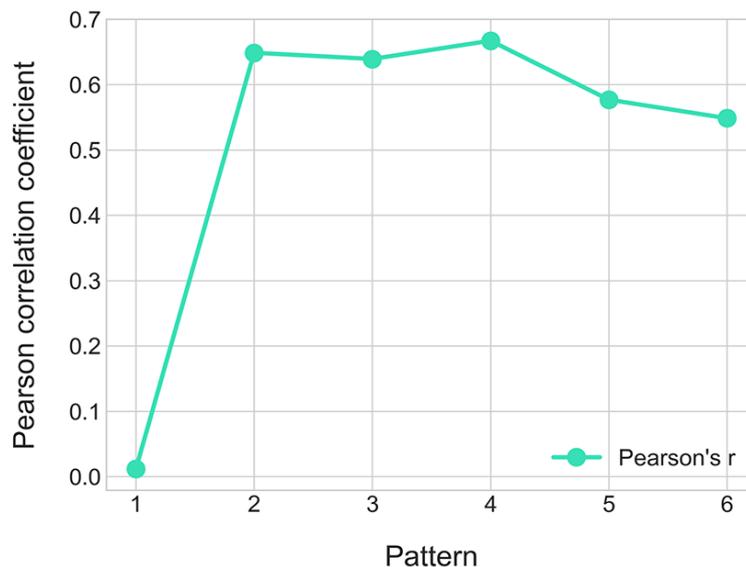


Fig. 4.10 Correlation between DSR and travel distance.

services in high priority areas detected by the models indeed are insufficient compared to their counterparts.

In the last part of the analysis, to display how the proposed approach to urban resource allocation is different from traditional ways more intuitively, we present the possible results obtained using a traditional method. Without access to GSN data and cutting-edge techniques, the primary source of information that we can rely on is government statistical data. To measure the average share of resources in different districts, we divide the number of venues assigned to each cultural pattern by the district population and present the results in Figure 4.11. Here the darker the colour is, the more sufficient the resources are. As shown in the figure, except for the first pattern, all the other cultural patterns have relatively more resources in central areas. This observation that cultural services in pattern 1 are more needed in the city centre compared to their counterparts is generally consistent with that shown by the DSR results in Figure 4.10. However, visible differences can be discovered between the results of the proposed model and that of the traditional method. Compared with the traditional one, the method devised in this chapter allows a fine-grained mapping of the demand-supply levels of certain types of facilities in small urban areas. Furthermore, this new approach leveraging large-scale GSN data enables us to consider the personalised preferences of residents rather than treat them homogeneously by utilising population density as the basis only in a more traditional style.

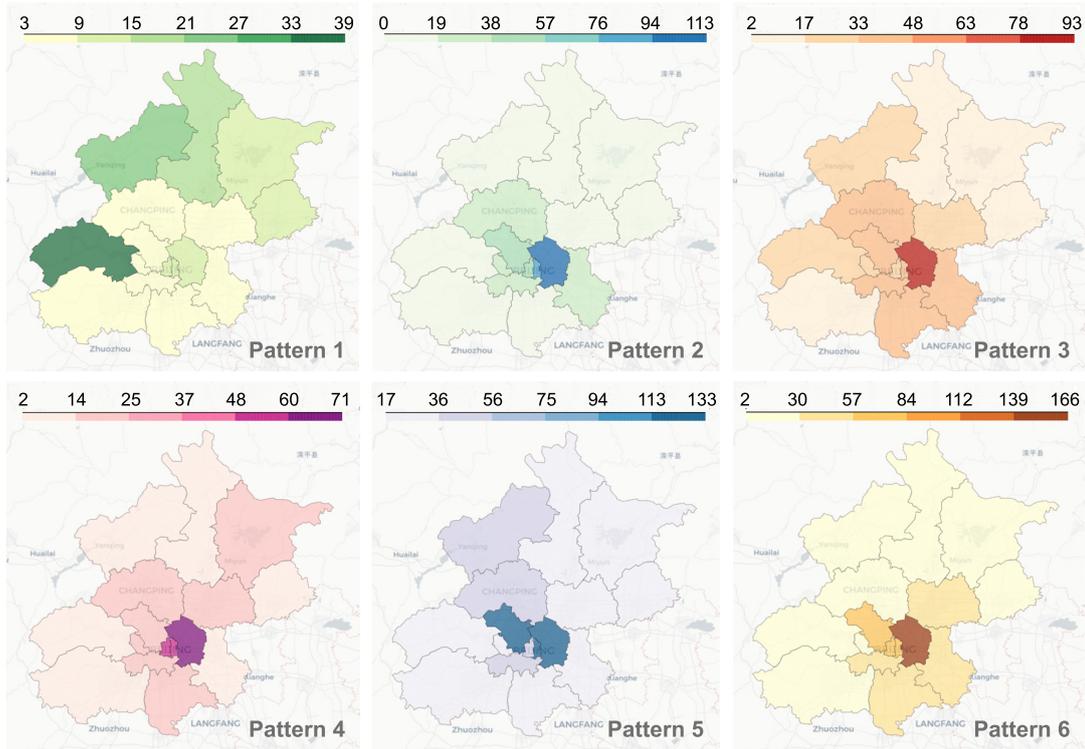


Fig. 4.11 Number of venues in each pattern per million people in Beijing Districts.

4.8 Discussion & conclusions

In this chapter, we have proposed a data-driven framework for urban cultural planning. The framework exploits a time-aware topic model to identify latent patterns of urban cultural interactions. Using then a density-based algorithm named POPTICS, we identify the primary locations of activity of mobile users and couple this with the TLDA output to generate cartographic representations indicative of the demand-supply balance for cultural resources in the city. We evaluate our approach using implicit user feedback, demonstrating how users active in areas that lack cultural establishments bear larger transportation costs to access cultural resources. Besides urban policymakers, the findings of this research can also provide suggestions to business owners on the opening hours, and citizens on neighbourhood characteristics in the city. For instance, a gym lover in Beijing may consider choosing to live in an area where the DSR of pattern 6 is relatively low, so as to enjoy better gym services around his neighbourhood. Overall, we demonstrate how the new generation of datasets emerging through modern location-based systems can provide an edge in city planning as they offer rich views on urban mobility dynamics and allow for the development of population adaptive frameworks that move beyond static representations of area-level population densities.

Chapter 5

Topic-enhanced memory network for POI recommendation

The previous two chapters have shown that collective human movement data in GSNs have the potential to support better government decision making. Next, from users' perspectives, we present how to recommend new places to individuals based on their fine-grained tastes in GSNs.

Chapter outline. In this chapter, we propose a novel topic-enhanced memory network (TEMN) (Section 5.4), a deep architecture to integrate the topic model and memory network capitalising on the strengths of both the global patterns and local neighbourhood-based features of user-POI interactions. We further incorporate a geographical module to exploit user-specific spatial preference and POI-specific spatial influence to enhance recommendations. In Section 5.5, extensive experiments on real-world WeChat datasets demonstrate the proposed unified hybrid model's effectiveness in various POI recommendation scenarios. We also present how a qualitative analysis of the attention weights and topic modeling can provide insight into the model's recommendation process and results.

5.1 Introduction

In the information explosion era, recommender systems have become increasingly important in our daily lives by improving the suggestions we receive online based on personal tastes (Zhou et al., 2018a). As a particular type of recommender systems focusing on real-world locations, point-of-interest (POI) recommendation has drawn intensive attention recently with the prevalence of smart mobile devices and the rapid growth of location-based

social networking services. Specifically, GSNs, such as Yelp¹ and Foursquare² link the physical and virtual worlds by offering users a way to share their life experiences with POIs via the "check-in" function. Through the lens of check-in data in GSNs, rich knowledge becomes available to discover users' visiting preferences for customised POI recommendation (Zhang and Chow, 2015). Such kind of services not only benefit users by providing them with suggestions of appealing venues to explore, but also facilitate targeted advertising with significant economic efficiency enhancement.

Unlike other recommender systems that push digital goods, e.g. e-books, news, and movies, POI recommendation aims at offering users preferred new venues to explore in the physical world (Wang et al., 2018), which can be largely affected by various real-life factors and thus faces more challenges. Firstly, to experience a POI, a user has to physically visit it, which is generally more costly and time-consuming than watching a movie or listening to a song online. Also, even if the user does visit the venue, she might prefer not to leave a check-in record due to privacy and security concerns (Xie et al., 2016). For these reasons, the number of POIs that a user generally interacts with is extremely small compared with the total number of POIs in the GSN, leading to a very sparse user-POI matrix that plagues a considerable number of current POI recommender systems (Yin et al., 2017). Moreover, unlike pure online interactions, users' activities in the physical world are limited by travel distance and time (Wang et al., 2018), making POI recommendation a more complex issue. Considering the distinguishing characteristics of the POI recommendation task, corresponding research efforts have been spent improving its effectiveness. The existing approaches in this area can be roughly grouped into two categories: neighbourhood-based methods and latent factor models (Ebesu et al., 2018). Neighbourhood-based approaches are good at describing local strong associations between users and POIs, but typically ignore the vast majority of user-POI interactions that are not similar enough. While for latent factor models like Matrix Factorization (MF) (Koren et al., 2009), they can easily extract the global structure of relationships between users and POIs, but have limitations in capturing their local features. Instead of establishing either a neighbourhood-based or latent factor model, we propose a unified hybrid architecture that combines the benefits of both techniques to enrich predictive capabilities by capturing both global and neighbourhood features of users and POIs. Furthermore, aiming to learn higher-order complex relations between users and POIs that cannot be obtained through simple functions, we adopt deep learning techniques.

The model proposed in this part of the research, *Topic-Enhanced Memory Networks (TEMN)*, is an end-to-end framework for personalised POI recommendation, which con-

¹www.yelp.com

²<https://foursquare.com/>

sists of three integrated components. Specifically, a memory-augmented neural network is employed to model fine-grained relations between each user-POI pair via a metric learning approach. This memory network (MN) component is linked with a topic modeling module, temporal latent Dirichlet allocation (TLDA) (Section 4.5.2) through user embedding. In this way, TLDA module can not only be used to infer the general preferences of users towards POIs and time slots, but also influence the representations of POIs and users in MN to enhance the recommendation performance. Furthermore, to capture geographical effects in POI recommendation, a geographical modeling module is also devised.

Our primary contributions can be summarised as follows:

- We propose an end-to-end deep learning framework that integrates neighbourhood-based and global preferences of users.
- We devise a more flexible architecture that incorporates multiple types of contextual information into POI recommendation and makes it applicable to various recommendation scenarios.
- We build a hybrid model that combines supervised and unsupervised learning and capitalises on the advances in both memory networks and topic modeling. Through a mutual learning mechanism, our model is also able to provide users' probability distributions over topics that are influenced by memory network.
- Comprehensive experiments on large WeChat datasets in different recommendation scenarios demonstrate the effectiveness of TEMN against competitive state-of-the-art baselines (improvement ratio of 3.25% and 29.95% for context-aware and sequential recommendation, respectively).
- Besides quantitative improvements, by incorporating neural attention mechanisms and topic modeling in TEMN, the interpretability of the POI recommendation is significantly promoted.

5.2 Related work

To alleviate the sparsity issue of user-POI interaction data and cater to the context-awareness nature of POI recommendation, auxiliary information has been incorporated into existing POI recommender systems (Wang et al., 2018). For instance, Lian et al. (2014) incorporated spatial clustering features of check-ins into a weighted matrix factorization framework. In some other works (Feng et al., 2017; Xie et al., 2016; Zhao et al., 2017), POIs clustered

within the same geographical region were restricted to similar representations compared with their distant counterparts in the recommendation. Another group of researchers focused on how temporal effects can be utilised to improve POI recommendation by considering either temporal cyclic patterns (Gao et al., 2013; Zhang et al., 2014) or sequential influence (Chen et al., 2014; Cheng et al., 2013; Zhao et al., 2017). Besides spatio-temporal information, other types of information have also been explored to facilitate POI recommendation performance, such as social influence (Cheng et al., 2012; Cho et al., 2011; Li et al., 2016; Tang et al., 2013; Ye et al., 2011; Zhang et al., 2016; Zhang and Chow, 2015), categorical information (Zhang and Chow, 2015; Zhao et al., 2015), visual content (Wang et al., 2017b), and text information (Chang et al., 2018; Wang et al., 2017a) of POIs.

From the angle of recommendation techniques, scholars have proposed some effective models. For example, MF (Koren et al., 2009) and its variants have been widely applied to POI recommendation tasks. Vanilla MF was initially used to deal with general POI recommendation, where only user-POI interactions are leveraged. When additional information like spatio-temporal information and the social relationship became available, some models derived from basic MF (e.g., IRenMF (Liu et al., 2014) were developed. For instance, GeoMF (Lian et al., 2014) and GeoIE (Wang et al., 2018)) were proposed and adapted to the location-aware recommendation. Markov chain models (e.g., LBPR (He et al., 2017a), NLPMM (Chen et al., 2014), FPMC-LR (Cheng et al., 2013), LORE (Zhang et al., 2014)), recurrent neural network (RNN) (e.g., ST-RNN (Liu et al., 2016), CARA (Manotumrukha et al., 2018), and DRCF (Liu et al., 2016)) can also be employed when temporal effects are particularly valued in sequential POI recommendation.

5.3 Problem formulation

In this section, we formulate our research problems and introduce the notations used throughout the chapter.

In a POI recommender system, let \mathcal{U} and \mathcal{V} denote the user set and the POI set, respectively. The user-POI interaction matrix $Y = \{y_{uv} \mid u \in \mathcal{U}, v \in \mathcal{V}\}$ is defined according to users' implicit feedback, where each entry y_{uv} in $Y \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{V}|}$ records whether user u has visited POI v as follows:

$$y_{uv} = \begin{cases} 1, & \text{if interaction } (u, v) \text{ is observed;} \\ 0, & \text{otherwise.} \end{cases} \quad (5.1)$$

Here a value of 1 for y_{uv} indicates that user u has visited the venue v before. Otherwise, it is assigned a value of 0, meaning that there is no interaction existing between them and the preference of user u for POI v is currently unclear. Mathematically, the main goal of the POI recommender system is to estimate the scores of the unobserved entries in matrix Y for the recommendation.

Since we aim to propose a more generic framework for POI recommendation that is sufficiently flexible to incorporate various contextual information and is widely applicable to different scenarios, we present the definitions of three POI recommendation scenarios involved in this chapter as follows:

General POI recommendation. In this recommendation scenario, only the interactions between users \mathcal{U} and POIs \mathcal{V} are taken into account. In other words, the interaction matrix Y is taken as the only input to the recommender system. The main task here is to provide each $u \in \mathcal{U}$ a list of POIs consisting of venues that the target user u has not explored before, but is potentially interested in.

Sequential POI recommendation. Sequential recommendation predicts successive POIs that are likely to be visited by a user given her check-in history. In such a scenario that sequence matters, for each user u , her past interactions with POIs are firstly sorted according to the check-in timestamps ascendingly. Instead of using the whole visiting sequence of the user directly, we divide it into some shorter segments by setting a time interval threshold ΔT . Specifically, if the time interval between two adjacent check-ins is larger than ΔT , they are not regarded as successive check-ins, since the large time interval may indicate they are irrelevant (Zhang et al., 2014). Based on this assumption, a check-in sequence set forms for each user and can be used as training data for sequential POI recommendation.

Context-aware POI recommendation. As stated in Section 5.2, multiple types of rich contextual information associated with check-ins can be used to enhance POI recommender systems. In this research, we will focus more on spatio-temporal aware POI recommendation.

5.4 Topic-enhanced memory networks

In this section, our proposed model Topic-Enhanced Memory Networks for POI recommendation (TEMN) is introduced. After presenting the general architecture of the model, we will describe its main components respectively before elaborating on how to integrate them with each other for overall training.

5.4.1 General framework

In Figure 5.1, we present a visual depiction of TEMN’s overall architecture. It can be seen that at a high level, TEMN consists of three key parts: a memory network, TLDA, and geographical modeling. The first two components are linked with each other, which allows for modeling the non-linear interactions between local features learned from the neighbourhood-based memory network and global preferences extracted from the topic model. With such an architecture, given a user u , the final prediction score for each unvisited POI v is a reflection of both the fine-grained relationship between user-POI pair $\langle u, v \rangle$ and the interest of user u in general. Technically, this capability is realised by the topic-enhanced user embedding, which is made up of two parts: 1) a representation related to the user’s memory component that encodes her visiting records (named as memory embedding), and 2) a vector employed to encode her intrinsic preference extracted through topic modeling (named as intrinsic embedding).

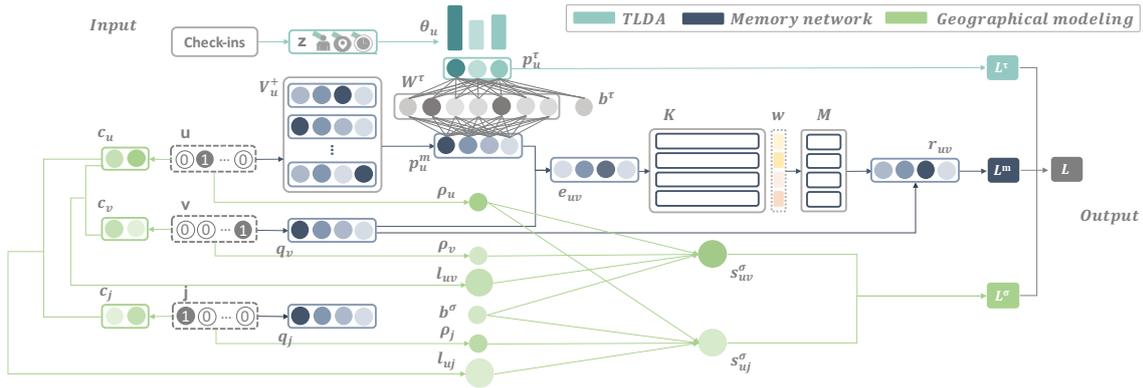


Fig. 5.1 The overall architecture of the TEMN. It consists of three modules represented by different colours. A user, a visited POI, a negative sampling POI, and the user’s visiting history are taken as input for POI recommendation.

5.4.2 Memory network

In our framework, we employ LRML (Tay et al., 2018) as a module of memory network to encode users’ previous check-in records, capture the local relationships between venues, and learn a relational vector for each user-POI pair which represents complex interactions between them. The reason that metric learning approaches, or to be more specific, LRML is utilised as a module for the recommendation in our framework has been provided in Section 2.5. Next, we will take a closer look at this memory network module in TEMN and introduces the function of each layer in the memory network in more detail.

User & POI embedding

In the memory network, POIs are initially represented by binary sparse vectors through one-hot encoding. Then they are further transformed into low-dimensional dense vectors to get POI embeddings. Formally, given a venue v , let $q_v \in \mathbb{R}^d$ denotes its embedding vector. Here d is the dimension size of the POI embedding. Different from POI embedding method, for each user u , her memory embedding $p_u^m \in \mathbb{R}^d$ is derived based on the POIs she has visited V_u^+ as:

$$p_u^m = \frac{1}{|V_u^+|} \sum q_i^m (i \in V_u^+) \quad (5.2)$$

where q_i^m is the embedding vector of POI i . Then for each given user-POI pair $\langle u, v \rangle$, we calculate their joint embedding e_{uv} through Hadamard product as below:

$$e_{uv} = p_u^m \odot q_v \quad (5.3)$$

The dimension size of the generated vector e_{uv} is the same as that of p_u^m and q_v . This joint embedding of user-POI pair will be fed into a memory-augmented neural network later to make predictions.

Key-value memory networks

The core of the memory network module consists of a memory matrix $\mathbf{M} \in \mathbb{R}^{h \times d}$, where h represents the number of memory slots and d denotes the size of each memory cell. In our model, this memory matrix stores the information related to user-POI interactions and can be read and updated adaptively by adopting key-value attention mechanisms. Technically, a key matrix $\mathbf{K} \in \mathbb{R}^{d \times h}$ is employed, where the number of key vectors is h , and the dimensionality of each key slot $k_i \in \mathbf{K}$ is d , which is the same as that in user and POI embedding. Our main aim here is to learn an attention vector w . For each $\langle u, v \rangle$ pair, when their joint embedding e_{uv} is fed into the networks, the similarity between e_{uv} and each key vector k_i will be first calculated by the dot product before being converted to a relevance probability using softmax function:

$$w_i = \text{softmax}(e_{uv}^T \cdot k_i) = \frac{\exp(e_{uv}^T \cdot k_i)}{\sum_j \exp(e_{uv}^T \cdot k_j)}, \forall j = 1, 2, \dots, h \quad (5.4)$$

where w_i represents the attention weight of the i -th element in the attention vector w , which will be further utilised to collect useful information from each memory slice $m_i \in \mathbf{M}$ to generate the relation vector r_{uv} of the given user-POI pair $\langle u, v \rangle$ through:

$$r_{uv} = \sum_i w_i m_i, \forall i = 1, 2, \dots, h \quad (5.5)$$

This obtained latent relation vector r_{uv} is the output of the memory network module.

Optimisation

The latent representation of user-POI interactions acts as a translation vector that enables the model to learn relationships between each specific user-POI pair in metric space adaptively (Tay et al., 2018). The scoring function implemented here is based on the idea of TransE (Bordes et al., 2013) model, where we compute the loss of the triple $\langle p_u^m, r_{uv}, q_v \rangle$ by the Euclidean distance and calculate the score for each user-POI pair $\langle u, v \rangle$ as:

$$s_{uv}^m = -\|p_u^m + r_{uv} - q_v\|_2^2 \quad (5.6)$$

where $\|\cdot\|_2$ is the $L2$ -norm of the vector. s_{uv}^m represents the score of user-POI pair $\langle u, v \rangle$ predicted through the memory network. The higher the score, the more likely user u would be interested in the recommended venue v .

We optimise the memory network module under the Bayesian personalised ranking (BPR) (Rendle et al., 2009) criterion, where the key assumption is that observed entries should be ranked higher than their unobserved counterparts. Given a positive instance $v \in \mathcal{V}_u^+$ visited by user u , a corrupted example $j \notin \mathcal{V}_u^+$ which the user has not interacted with yet would be selected based on a certain negative sampling rule. Then both of the POIs v and j would go through the same user and POI embedding layer as introduced, respectively. When calculating the score s_{uj}^m for the corrupted pair $\langle u, j \rangle$, the relation vector used is the same as that generated by positive pair $\langle u, v \rangle$. The reason for doing like this is that we empirically found it providing much better performance than generating another relation vector for the negative sample. Based on pairwise comparisons, the objective function of the memory network module is defined as:

$$\mathcal{L}^m = \sum_u \sum_{v \in \mathcal{V}_u^+} \sum_{j \notin \mathcal{V}_u^+} \max(0, s_{uj}^m - s_{uv}^m + \lambda^m) \quad (5.7)$$

where λ^m is the margin that separates positive and negative examples. Rectified linear unit (*ReLU*) function is adopted by us as the non-linear activation function here as it gives better experimental results than using the sigmoid function.

5.4.3 Temporal latent Dirichlet allocation

Apart from user’s memory embedding p_u^m achieved through the memory network, which captures neighbourhood-based interests of a user, we think that the intrinsic preferences of a user are also essential in creating a comprehensive user profile for personalised POI recommendation. For this reason, we employ the TLDA model in this subsection to mine the inner interests of users from a more general perspective.

As introduced in Subsection 4.5.2, the TLDA model is an unsupervised machine learning algorithm initially proposed by Chen (2017) to improve the performance of blog search engines. Until recently, we applied it to the user profiling question leveraging human mobility data for the first time (Zhou et al., 2018b). With the help of TLDA, each user can be characterised as a mixture of patterns with both spatial and temporal preferences considered. In this study, the probability distribution over latent patterns for a user inferred through TLDA serves as the user’s intrinsic embedding in the TEMN. As a detailed introduction about TLDA has been provided in Subsection 4.5.2, we will not repeat it here.

Eventually, the probability distribution over π patterns for a given user u can be estimated by applying the Gibbs sampling algorithm (Griffiths, 2002) (see Equation 4.4). It can be seen as a pattern mixture that represents each user’s intrinsic embedding denoted as p_u^τ :

$$p_u^\tau = (P(z_u^1|\theta_u), P(z_u^2|\theta_u), \dots, P(z_u^\pi|\theta_u)) \quad (5.8)$$

It is worth noting that the TLDA module not only provides us the pattern-user probability distribution, the distributions of venues and time slots associated with each pattern can also be estimated to help us understand what each pattern might refer to.

5.4.4 Fusion of MN & TLDA

In the last two subsections, the modules of the memory network and TLDA have been introduced, respectively. Essentially, they are not two separate parts, but an interrelated whole system that is capable of learning both neighbourhood-based and global characteristics of users and POIs, as well as the fine-grained relationships between specific user-POI pair. In this subsection, we will introduce how the memory network and TLDA can be fused, and mutually reinforce each other under the TEMN framework.

The key part connecting the two modules is the user embedding. More specifically, we define our problem as a multi-class classification task and use a neural network approach to learn the topic-enhanced user embedding for the memory network. As our main aim is to map the memory embedding p_u^m of a user u to her intrinsic embedding p_u^τ through the neural network architecture, p_u^m is taken as the input data, and the intrinsic embedding p_u^τ of the user

is applied as the output where each element of the vector represents a class. Mathematically, the model is defined as:

$$O^\tau = f(W^\tau p_u^m + b^\tau) \quad (5.9)$$

where W^τ and b^τ denote the weight matrix and bias vector, respectively. $f(\cdot)$ is the activation function. To get the final predicted probabilities \hat{p}_u^τ , the output vector O^τ is converted by the softmax function as:

$$\hat{p}_u^\tau = \text{softmax}(O^\tau) \quad (5.10)$$

By introducing the softmax activation function in the output layer, the cross-entropy loss is used for optimisation:

$$\mathcal{L}^\tau = - \sum_u \sum_i p_u^\tau(i) \cdot \log(\hat{p}_u^\tau(i)), \forall i = 1, 2, \dots, \pi \quad (5.11)$$

where i represents one of the π patterns in the TLDA. The selection of activation function $f(\cdot)$ in Equation 5.9 is not necessarily restricted to a particular function. Many options like sigmoid, hyperbolic tangent, and *ReLU* can be taken in practice according to the specific application domain. Alternatively, one can also choose to apply Equation 5.9 without an activation function and convert it to a softmax regression model. Here, we choose the *ReLU* function as it yields the most favourable result.

5.4.5 Geographical modeling

Given the TEMN is expected to be competent for the context-aware recommendation, it is necessary to incorporate spatio-temporal information into the model. As introduced in the last subsection, temporal effects have been considered in the design of TLDA. In this subsection, we are going to present how geographical influence is taken into account in the TEMN.

Most existing methods assume that the geographical effect between a pair of user and POI is determined by their physical distance, which fails to capture the asymmetry between the user and POI, the high variation of geographical preferences across users, and the differences of geographical influence among POIs. In this sense, we think that to measure the geographical effect only with physical distance is not enough. Instead, we introduce the concepts of *geographical preference* and *geographical influence* which are specifically learned for each user and POI so as to depict the fine-grained geographical effect between them.

For a user u , each POI $v \in \mathcal{V}_u^+$ visited by her has certain geographic coordinates $c_v = (g_{lat}^v, g_{lon}^v)$ recorded simultaneously when the certain check-in was created. Here g_{lat}^v and g_{lon}^v represent the latitude and longitude of venue v , respectively. Through an investigation about her check-in history, the central point c_u of user u where her check-ins clustered around geographically can be determined by:

$$c_u = \left(\frac{1}{|\mathcal{V}_u^+|} \sum_i g_{lat}^i, \frac{1}{|\mathcal{V}_u^+|} \sum_i g_{lon}^i \right), \forall i \in |\mathcal{V}_u^+| \quad (5.12)$$

Then the physical distance l_{uv} between user u and venue v can be calculated through:

$$l_{uv} = \|c_u - c_v\|_2 \quad (5.13)$$

To model the geographical effect between each user-POI pair $\langle u, v \rangle$, three factors are introduced in this part: 1) the *geographical preference* ρ_u of user u , which characterises how sensitive the user is to distance; 2) the *geographical influence* ρ_v , representing the influential level of the venue reflected on physical distance; and 3) the geographical distance between them l_{uv} . Based on this, the geographical score s_{uv}^σ of user-POI pair $\langle u, v \rangle$ is calculated by:

$$s_{uv}^\sigma = \rho_u l_{uv} + \rho_v l_{uv} + b_u^\sigma = (\rho_u + \rho_v) l_{uv} + b^\sigma \quad (5.14)$$

where b^σ denotes the geographical bias. Similarly, like that in the memory network (Section 5.4.2), the objective function applied here is also based on BPR (Rendle et al., 2009). Given a venue v which has been visited by user u , and a negative venue j that she never interacted with, the pairwise loss function is defined as:

$$\mathcal{L}^\sigma = \sum_u \sum_{v \in \mathcal{V}_u^+} \sum_{j \notin \mathcal{V}_u^+} \max(0, s_{uj}^\sigma - s_{uv}^\sigma + \lambda^\sigma) \quad (5.15)$$

where λ^σ is the margin separating positive and negative examples.

5.4.6 Joint training

Our complete recommender system is a hybrid of memory network, TLDA, and geographical modeling, which is a topic-enhanced memory network with spatio-temporal contextual information considered. Until now, we have elaborated each module respectively and introduced how memory network is linked with the TLDA. In this part, we present how the TEMN framework functions as a whole.

For each pair of user-POI $\langle u, v \rangle$, we can calculate its score achieved from the memory network s_{uv}^m through Equation 5.6, and its geographical score s_{uv}^σ by Equation 5.14. Then the overall score s_{uv} can be easily obtained through the following equation:

$$s_{uv} = s_{uv}^m + \eta s_{uv}^\sigma \quad (5.16)$$

where η is a weighting parameter, which can be optimally tuned to study the effects of incorporating geographical modeling for POI recommendation under our framework. The overall objective function for TEMN is a weighted sum of all the contributions of each related individual objective as:

$$\mathcal{L} = \mathcal{L}^m + \zeta \mathcal{L}^\tau + \varepsilon \mathcal{L}^\sigma + \lambda \|\rho\|_2^2 \quad (5.17)$$

where \mathcal{L}^m , \mathcal{L}^τ , and \mathcal{L}^σ denote the loss induced by the memory network, TLDA, and geographical modeling respectively. ζ and ε are weighting parameters specified to influence to what degree the TLDA and geographical effects should be taken into account during the optimisation. ρ is the model parameter set; and λ is a parameter that controls the importance of the last term, where we regularise all the model parameters to prevent overfitting. In the training phase, mini-batch stochastic gradient descent (SGD) is employed to minimise the objective function.

5.5 Experiments

The effectiveness of the proposed model for top-N POI recommendation under various scenarios is evaluated in this section. We begin with experimental setup and then present and analyse the experimental results both quantitatively and qualitatively. We also provide detailed explanations of how we did our experiments for reproducibility. Our code is publicly available at <https://github.com/XiaoZHOU/CAM/TEMN>.

5.5.1 Experimental setup

Datasets

We conduct our experiments on the WeChat³ Moments check-in data. This study was approved by Tencent⁴ and the Computer Science and Technology Department University of Cambridge Ethics Committee. Initially, the whole dataset contains all the POI check-in

³<http://www.wechat.com/en>

⁴<https://www.tencent.com/en-us/index.html>

records created by WeChat users living in Beijing for one year spanning from September 2016 to August 2017. To ensure the quality of the data, we retain users who interacted with at least ten unique POIs. After filtering the data, we obtain an evaluation dataset, termed as *WeChat(GPR)* for the general and context-aware recommendations. For the sequential POI recommendation scenario, the dataset is further processed as introduced in Section 5.3. More specifically, we set the threshold ΔT to one day (24 hours), and keep sequences with at least five successive check-ins. This newly created dataset containing the set of successive check-ins for each user is called *WeChat(SPR)*. The statistics of the two datasets are summarised in Table 5.1. Apart from the interaction information between users and POIs, the timestamp associated with each check-in and the side information of each POI consisting of coordinates and categories (e.g., restaurant, art museum, sports centre) is also available.

Dataset	#Users	#POIs	#Chek-ins	Density
WeChat (GPR)	75,973	28,183	5,644,965	0.264%
WeChat (SPR)	28,566	13,826	509,589	0.129%

Table 5.1 Statistics of the evaluation WeChat datasets.

Evaluation protocols

In the experiments, the check-in data are transformed into binary implicit feedback as ground truth, indicating whether the user had visited the specific POI. For each dataset, we holdout the latest 15% visiting history of each user to construct the test set, and split the remaining data into training (70%) and validation (15%) sets. The validation set is used to tune hyper-parameters and the final performance comparison is conducted on the test set. Closely following the setup that in some representative papers (He et al., 2017c; Rendle et al., 2009), we randomly sample 100 POIs that are not visited by the user and rank the test POI along with the 100 negative samples for evaluation.

Two common ranking evaluation metrics, *Hit Ratio (HR)* and *Normalised Discounted Cumulative Gain (NDCG)* are adopted. *HR* measures whether the test POI shows within the top N in the ranked list and the *NDCG* takes the position of the test POI into account and penalises the score if it is ranked lower in the list.

Baselines

We compare our proposed approach against competitive baselines representing latent factor models, topic models, metric learning methods, Markov chain models, and deep learning-based models under various recommendation scenarios.

To evaluate the performance of our model in *general POI recommendation*, the following methods are considered:

- **MF** (Koren et al., 2009) Matrix factorization is one of the most popular model-based collaborative filtering approaches for the recommendation, which models the user-POI relationship using the inner product. A more detailed description of the MF method can be found in Section 2.5.
- **BPR** (Rendle et al., 2009) Bayesian personalised ranking optimises the MF model with a pairwise ranking loss. It is tailored to recommendations with implicit feedback data.
- **LDA** (Blei et al., 2003) It is an unsupervised learning approach initially applied to cluster documents and discover topics based on their contents. It is the prototype of our TLDA module that does not consider temporal factors in human mobility.
- **CML** (Hsieh et al., 2017) This is a method used to exam whether the latent relational learning of r_{uv} is necessary in our design of the memory network module. CML minimises the Euclidean distance between user and POI vectors as $\|p_u^m - q_v\|_2^2$. Please refer to Section 2.5 for more detailed information about the CML algorithm.
- **LRML** (Tay et al., 2018) This is a memory network-based approach that learns a translation vector between a pair of user and item. It acts as the memory network module in our TEMN framework without topic-enhanced effects and geographical influence.
- **TEMN(GPR)** It is a variant of TEMN, which keeps the memory network component, replaces the TLDA with LDA, and removes the geographical modeling part.

The second group of baseline approaches for *context-aware POI recommendation* that support the integration of contextual information associated with POI check-ins are:

- **GeoMF** (Lian et al., 2014) This method extends the MF by augmenting latent factors with the user’s activity region and POI’s influence area.
- **TLDA** (Zhou et al., 2018b) TLDA is an extension model of LDA that can be applied to extract users’ lifestyle patterns with temporal preferences considered. It is also employed as the topic module in our framework to enhance memory networks.
- **TEMN(CPR)** It is our complete model with the WeChat (GPR) dataset employed.

The last group of baselines are models devised for *sequential POI recommendation*, which include:

- **LORE** (Zhang et al., 2014) It considers sequential influence and geographical effects in recommendation by adopting additive Markov chain.
- **ST-RNN** (Liu et al., 2016) This is an RNN-based state-of-the-art POI recommendation model, which considers information about the time interval and distance between pair of check-ins.
- **TEMN(SPR)** Our complete model with the WeChat (SPR) data fed into, which is employed to test our model’s effectiveness in successive POI recommendation.

Parameter tuning

We implemented the TEMN model in TensorFlow⁵. For our model, all hyperparameters are tuned according to the validation set based on the NDCG metric. We randomly initialise model parameters according to the uniform distribution and optimise the model by conducting stochastic gradient descent (SGD). The learning rate of SGD and the regularisation parameter are determined by grid search in the range of $\{0.0001, 0.005, 0.001, 0.01\}$ and $\{0.0001, 0.001, 0.01, 0.1\}$, respectively. The dimensionality of user and item embeddings d is tuned amongst $\{20, 50, 75, 100\}$. The number of memory slices in M is tuned amongst $h = \{5, 10, 20, 50\}$. And the number of patterns set in TLDA is tuned among $\{3, 5, 10, 20\}$. For MN module and geographical modeling that minimise the hinge loss, the margin is tuned amongst $\{0.1, 0.2, 0.25, 0.5\}$.

Eventually, we find that the following hyperparameters work well: the learning rate is set to 0.005, the regularisation parameter is set to 0.0001, both of λ^m and λ^σ are set to 0.2. The dimension of user and item embeddings d is set to 50, and the number of memory slots h is set to 10. The number of patterns set in TLDA is 10. We set the weighting parameters $\zeta = 0.2$, $\varepsilon = 0.1$, and $\eta=0.4$ in general and context-aware recommendation; while for sequential POI recommendation, we set $\eta = 1.6$ in the overall score function, ζ , and ε to 0.1 and 0.4, respectively.

5.5.2 Experimental results

The comparison results of our proposed model and the baselines on two WeChat datasets in three POI recommendation scenarios are reported in Table 5.2. Some key observations from the experimental results are summarised as follows:

⁵<https://www.tensorflow.org/>

Type	Method	HR@5	NDCG@5	HR@10	NDCG@10
General	MF	0.50961	0.30345	0.63174	0.34028
	BPR	0.52561	0.33747	0.64313	0.37597
	LDA	0.69221	0.53960	0.80958	0.57778
	CML	0.67983	0.45613	0.78813	0.49163
	LRML	0.69594*	0.54669*	0.81199*	0.58436*
	TEMN(GPR)	0.70389(+1.141%)	0.55221(+1.010%)	0.81752(+0.682%)	0.58914(+0.817%)
Context-aware	GeoMF	0.63714	0.50435	0.70742	0.52741
	TLDA	0.71518*	0.55852*	0.83033*	0.59601*
	ST-RNN	0.60240	0.47044	0.74372	0.51621
	TEMN(CPR)	0.72876(+1.899%)	0.57666(+3.248%)	0.83398(+0.440%)	0.61053(+2.436%)
Sequential	LORE	0.53187*	0.37880	0.69148*	0.43049
	ST-RNN	0.53041	0.42867*	0.62798	0.45999*
	TEMN(SPR)	0.62105(+16.767%)	0.54847(+27.947%)	0.69769(+0.897%)	0.57304(+24.577%)

Table 5.2 Performance comparison of different methods in three recommendation scenarios. Best performance is in boldface. We use "*" to mark the best performance from baselines for each comparison and report the improvement ratio of our model over the best baseline performance for each scenario in parentheses.

(1) **Overall Comparison.** Encouragingly, it is clear that the performance of our proposed model TEMN is consistently better than all the baselines under different conditions by a relative large margin (the relative improvement ratio over the best baseline in the general, context-aware, and sequential recommendation scenario is 1.14%, 3.25%, and 27.95%, respectively). Additionally, as can be seen from both Table 5.2 and Figure 5.2, TEMN offers significantly larger improvement for NDCG than HR in general, suggesting that the proposed model gives a higher ranking to the positive test POI, and shows the feasibility and effectiveness of applying the topic-enhanced memory network to the top-N POI recommendation.

(2) **Comparison against Individual Modules.** Comparing the full TEMN model with its component modules, we find that the overall performance order is as follows: TEMN(CPR) > TLDA > LRML. This indicates that an integrated structure of memory network and topic model performs better than the individual components. To further compare the performance of the memory network module (LRML) and topic model module (TLDA), TLDA shows its superiority by giving around 2-3% improvement over LRML across all the metrics and cut-offs. We can also observe that TLDA beats all the other baselines in both general and context-aware recommendations, which highlights the importance of users' long-term interest mining and temporal effect in POI recommendation. In a nutshell, even though the performances of the individual modules, LRML and TLDA, are marginally worse than the complete model, they provide the best baseline performance in general and context-aware recommendation scenarios, respectively.

(3) **Comparison against Topic Models & MF-based Models.** To compare models with different recommendation mechanisms, neural network-based approaches (CML, LRML, and ST-RNN) and topic models (LDA and TLDA) outperform MF-based methods (MF, BPR, and GeoMF) in most cases, which indicates the usefulness of capturing more complex user-POI relationships through nonlinear methods and the global preference of users through topic modeling. The relatively poor performance of MF-based approaches suggests that adopting the dot product may not be enough to depict user-POI interactions as discussed in Section 2.5. Another finding is that LRML performs better than CML, proving the advantage of employing a translation vector to capture the relation between a user-POI pair. In addition, GeoMF, TLDA, and TEMN(CPR) all give better recommendation results than their prototypes (MF, LDA, and TEMN(GPR)), indicating the value of incorporating spatio-temporal information into POI recommender systems.

(4) **Sequential POI Recommendation.** Utilising a different dataset (WeChat (SPR)), the experimental results of models in the third group for sequential POI recommendation are not directly comparable with that in the other two groups and are thus discussed individually here.

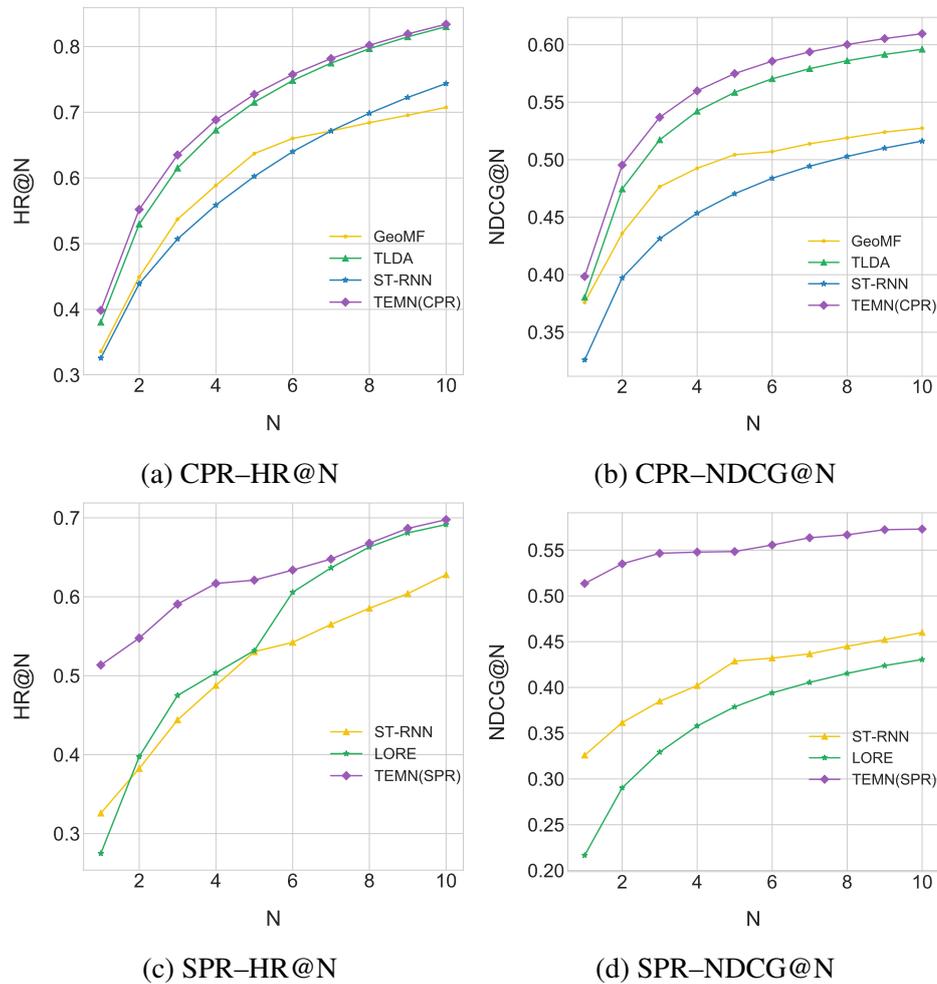


Fig. 5.2 Evaluation of Top- N POI recommendation where N ranges from 1 to 10 on the two WeChat datasets in two scenarios.

As can be seen from the table, our proposed model TEMN(STR) gives the best performance among the methods applied to the sequential POI recommendation task. This result indicates the expressive power of the memory mechanism to model more recent user-POI interactions and the effectiveness of geographical module to characterise fine-grained spatial effects specific to each user-POI pair in sequential recommendations. To compare the two baseline methods in this group, it is interesting to see that the Markov chain model (LORE) performs better for HR , while the RNN-based method (ST-RNN) shows better results for $NDCG$. This finding may be useful for practical applications: when the length of recommendation list is limited (N is small in the top- N recommendation), our TEMN model and ST-RNN are better choices for successive POI recommendations.

5.5.3 Detailed analysis of the proposed model

After discussing the experimental results quantitatively, in this subsection we present some qualitative observations to show the interpretability of our proposed model.

Explainable recommender system is a hot topic that has attracted much attention in recent years. However, due to the black-box nature of many existing recommendation algorithms, the interpretability of recommender systems is still an open research question. Take MF as an example, although MF has been shown as a successful recommendation technique that achieves a significant improvement in recommendation accuracy, it is not intuitive to understand the meaning of the latent factors of users and items in MF (Kim and Shin, 2017; Zhang and Chen, 2018). Deep learning-based recommender systems also suffer from the problem of model explainability. Such systems are often viewed as black boxes since it is not clear about the decision-making process of neural networks, making it difficult to understand why an item got a higher score of recommendation than other candidates (Zhang and Chen, 2018). To overcome this limitation, the neural attention mechanism emerged as a possible technique to improve the interpretability and has been applied to some deep learning models. LRML is one such example that leveraged the recent advancements of deep learning and interpretable attention module to deliver recommendations (Tay et al., 2018). To take this one step further, the proposed architecture of TEMN incorporates topic information and the attention mechanism in memory networks, which further enhance the interpretability of the memory-based recommender systems.

Attention visualisation

The use of attention mechanism in TEMN enables us to visualise the weighted importance of memory slices with respect to multiple patterns. To illustrate this, we plot a heatmap of the weights in Figure 5.3, where the colour scale represents the strength of the attention weights, and each row represents the average attention vector for each pattern. As can be observed, it is clear that each pattern has learned a certain type of selection rules across memory slices. Moreover, some patterns are particularly similar to each other than their counterparts. When we compared these results with the semantic meaning for each pattern as presented in Table 5.3, we can find that similar patterns like 2, 7, 8, and 9 all have residential district as their top category. In contrast, pattern 6, which is related to Olympic venues, shows a distinctive attention vector, suggesting that the visiting features of this pattern can be extremely different from other types.

Pattern	Top Categories
0	Airport, Central business district, Hub station
1	Shopping mall, Square, Commercial street
2	University, Shopping mall, Residential district
3	Bar street, Stadium
4	Historic district, Shopping mall
5	High-tech zone, University, Shopping mall
6	Olympic Park
7	Residential district, SOHO
8	Hub station, Residential district
9	University, Residential district

Table 5.3 Pattern and top categories.

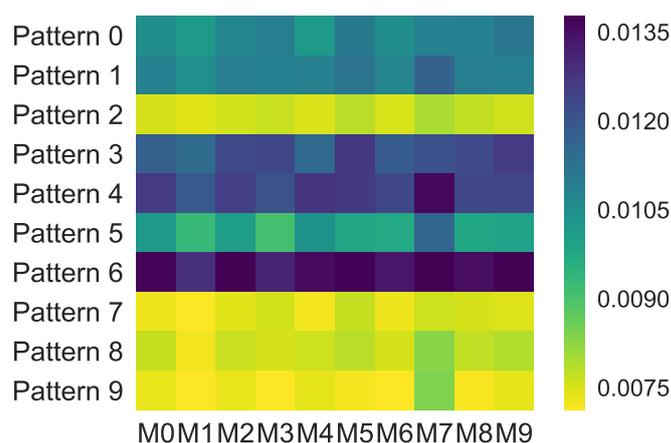


Fig. 5.3 Attention weights over memory slices for patterns.

Geographical preferences of patterns

As introduced in 5.4.5, the *geographical influence* ρ_v of a POI and *geographical preference* ρ_u of a user can be learned through our model. By calculating the average *geographical influence* of POIs and *geographical preference* of users for each pattern, we can gain insights into the geographical characteristics of different patterns. The calculated results are presented in Figure 5.4, where a smaller number means the POI or user is less sensitive to distance. It can be seen that patterns of 0 and 8 give the smallest value of ρ_v , indicating that the visits to POIs belonging to these patterns (e.g. airport, central business district, hub station) are not influenced by distance too much; while for pattern 1, which is represented by POI categories more related to local services, it is rather sensitive to distance that people are unlikely to visit

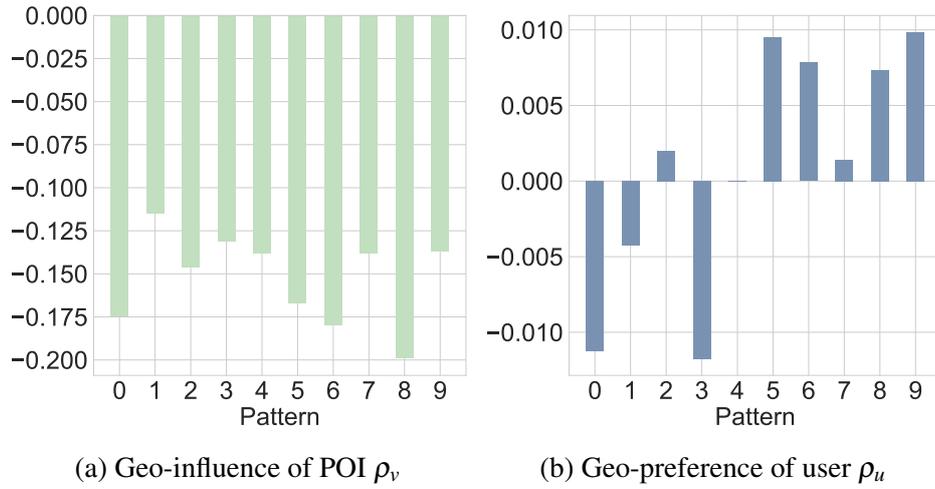


Fig. 5.4 Average geographical influence of POIs and geographical preference of users for patterns.

them from distant places. On the other hand, users from pattern 0 and 3 are relatively more willing to visit POIs far away than their counterparts by showing the smaller average ρ_u .

Effect of weighting parameters

The weighting parameters utilised in the overall loss function (Equation 5.17) and score function (Equation 5.16) are tuned by us to study the effect of incorporating the TLDA module and geographical modeling for POI recommendation. Specifically, we study the performance of our models on NDCG@5 and NDCG@10 in *context-aware POI recommendation* (CPR) and *sequential POI recommendation* (SPR) scenarios by tuning the relevant weighting parameters as shown in Figure 5.5 and Figure 5.6.

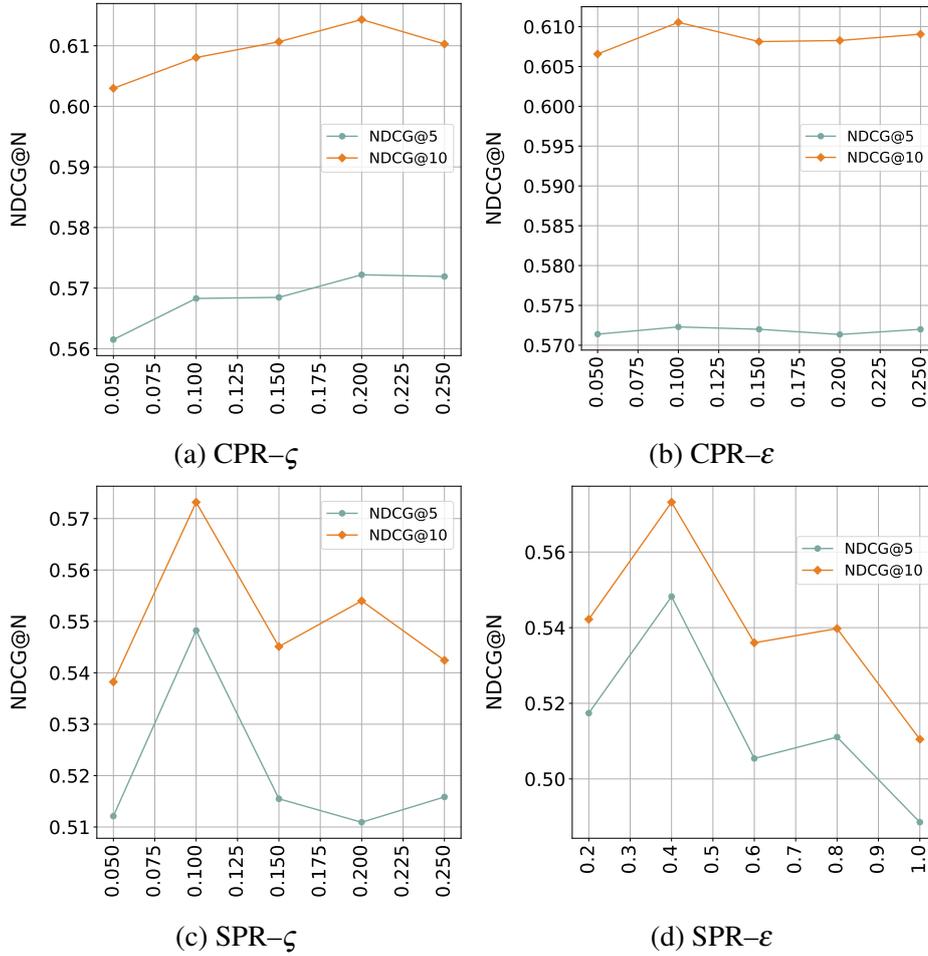


Fig. 5.5 Effect of weighting parameters in the loss function.

It can be seen from Figure 5.5a and Figure 5.5c, when the weighting parameter of TLDA component ζ is set to a small value of 0.05, models for CPR and SPR all give unfavourable results. When TLDA begins to play an increasingly important role, the performance is improved dramatically. The best results are achieved when $\zeta \approx 0.2$ for CPR and $\zeta \approx 0.1$ for SPR. However, when the weights of the TLDA module continue to rise, the performance goes down. This result indicates that integrating TLDA module with memory network indeed helps to offer better POI recommendation, however, paying too much attention to user's intrinsic preference may weaken her neighbourhood-based characteristics.

As for the influence of geographical modeling, an interesting observation is that in the CPR scenario, the performance of models is stable when ϵ is tuned; while for SPR, the curves fluctuate substantially with the increasing of ϵ . Moreover, it can be found that when the weighting parameters of the geographical module are set to relatively larger values ($\epsilon \approx 0.4$ and $\eta \approx 1.6$), the performance of our models is most satisfying. These findings suggest

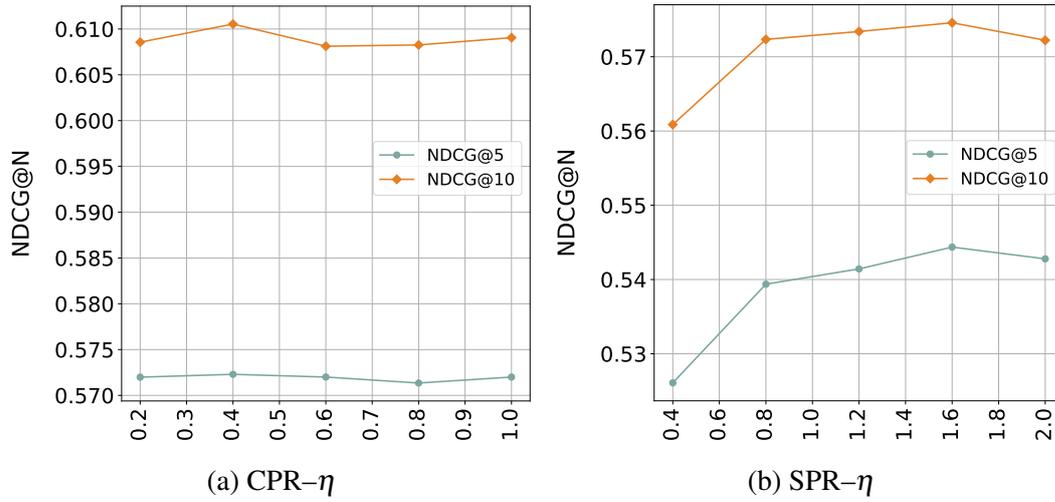


Fig. 5.6 Effect of weighting parameter in score function.

that geographical effects have a greater influence on sequential POI recommendation, where short-term preferences value more.

5.6 Discussion & conclusions

We introduce a novel end-to-end architecture, named Topic-Enhanced Memory Network (TEMN), for better POI recommendation with spatio-temporal contextual information considered. Different from existing deep learning frameworks for POI recommendation, TEMN is a unified hybrid model that leverages TLDA and external memory network with a neural attention mechanism to capture both global and fine-grained preferences of users. Comprehensive experiments under multiple configurations demonstrate the proposed architecture's significant improvements over competitive baselines. Qualitative visualisation of the attention weights and semantic investigation about patterns learned via TLDA module provide insight into the learning process and the recommendation results. In future work, we plan to extend TEMN to incorporate richer content information such as users' profiles and POIs' attributes.

Chapter 6

Reflections & outlook

The preceding chapters have presented a variety of modeling approaches for tackling some tricky challenges cities face nowadays and discussed the results of our research, concerning the thesis: *the exploration of geo-social network data with various analytical techniques can advance our understanding of the nature of urban phenomena and the development of intelligent urban computing applications*. In this chapter, we provide a summary of the major contributions presented in this dissertation that support the thesis above, followed by highlighted directions for future research.

6.1 Thesis summary & contributions

The recent emergence of massive GSN data, thanks to the proliferation of mobile services together with advances in machine learning, has stimulated the urban computing research in this dissertation. Essentially, the growing adoption of machine learning algorithms in urban studies empowers large-scale digital data to reveal human mobility patterns and promote fresh solutions to various urban problems. In this dissertation, we have taken a step forward in both the modelling of human mobility and the development of computing frameworks that could support urban mobile applications. Through this perspective, new algorithms tailored to specific urban problems have been proposed to this end, demonstrating how people-place interactions can be modelled to gain novel insights into the dynamics of cities. In more detail, the research findings we have presented and the key contributions of this dissertation are summarised below.

Revealing underlying local evolution in geo-social graphs: In Chapter 3, we proposed an innovative approach to giving insights on underlying relationships between socio-economic status, cultural investment, and geo-social network properties using a fusion of techniques,

including network analysis, statistical analysis, and supervised machine learning. Through the analyses, we find that initial deprivation level and cultural expenditure strategy in neighbourhoods may lead to different network properties in the geo-social graph. Furthermore, investing more in cultural and related services can boost local development for relatively more deprived urban areas.

Defining novel metrics to characterise urban areas leveraging multiplex interactions:

In Chapter 3, we demonstrate how datasets from government and geo-social networks with different spatial and temporal granularities can be integrated and analysed jointly to produce the inference of local socio-economic change at a fine temporal grain. Leveraging heterogeneous datasets, we define new metrics on cultural investment and features in geo-social networks to measure the priority level of culture for urban areas and show how the differences in these metrics reflect on the network properties of local areas.

Predicting socio-economic conditions through collective transitions in geo-social graphs:

Based on the preliminary analysis of local network graphs in Chapter 3, we have further proven it feasible to adopt network features, such as the centrality and clustering coefficient, combined with geographic and cultural expenditure factors to predict the socio-economic deprivation change for small urban areas with high prediction performances. Furthermore, we have also evaluated the predictive capability for different classes of features, which demonstrates that geo-social network features as a whole are more powerful predictors than geographic and expenditure signals.

Extracting spatio-temporal patterns from digital footprints in geo-social networks:

In Chapter 4, we presented how to exploit large-scale time-stamped check-in data at urban venues in the geo-social network to obtain latent patterns of urban cultural interactions. We employ an extended version of the standard topic modeling approach that takes check-ins as input to identify clusters of mobile users and venues according to their spatio-temporal visiting profiles. To examine the performance of the new topic modeling algorithm, we propose a novel evaluation method which is capable of measuring the intra-cluster coherence values among top venues and time slots for each latent pattern.

Determining primary visiting regions for users by a personalised clustering algorithm:

In Chapter 4, we explored novel methods to determine the main activity clusters for each user, as reflected in the geographic spread of her check-in activities. Technically, we proposed a novel density-based clustering algorithm, which is capable of detecting valid data points, determining visiting centres and radii of clusters for each user. Compared with the standard

OPTICS algorithm introduced in Chapter 2, the proposed clustering approach called POPTICS can learn a personalised threshold for each user. It emancipates users from setting the global input parameters of the maximum radius to consider, and the least number of points to form a cluster.

Identifying urban areas that lack service offerings: The primary locations of user activity through POPTICS are the means to quantifying demand levels for resources spatially. Overall, the output of this process corresponds to a set of heat maps depicting the intensity levels of user activity geographically for each of the pattern emitted by the TLDA model. We consider such intensity levels to reflect the user-driven demand for resources geographically. In addition to obtaining spatial descriptions of the demand levels for each pattern observed in the city, we determine the supply levels of resources using the spatial distribution of venues and users' check-ins belonging to each pattern as input. For each region in the city, we obtain a demand-supply ratio (DSR), high values of which are indicative of an area lacking resources, whereas low values indicate oversupply of a certain type of establishments in the region. We generate precision maps of such supply and demand levels for each pattern and demonstrate how users living in high-DSR neighbourhoods but adhere to a specific pattern travel long distances in the city to access the resources.

Recommending new urban venues in geo-social networks: In Chapter 5, we propose an end-to-end deep learning framework for personalized POI recommendation, termed as TEMN. Different from the DSR model that pinpoints in cartographic terms the areas in the urban territory where supply could improve through appropriate investment for government thanks to TLDA, the TEMN focuses more on the needs of individual users. By integrating the memory network with the topic modeling module, the proposed recommender system enables us to capture both neighbourhood-based and global preferences of users, and model fine-grained relations between each user-POI pair via a metric learning approach. Moreover, as TLDA is able to represent temporal characteristics of user-place interactions, geographical modeling module can model geographical effects, TEMN is also applicable to the spatio-temporal context-aware recommendation. Technically, TEMN is a hybrid model that combines supervised and unsupervised learning and capitalises on the advances in both memory networks and topic modeling through a mutual learning mechanism. Besides quantitative improvements, by incorporating neural attention mechanisms and topic modeling in TEMN, the interpretability of the POI recommendation is significantly promoted.

To summarise, this dissertation has shown how human mobility data in GSNs can be exploited to tackle various urban issues by devising computational algorithms and frameworks. In all the three studies, only the anonymous check-in data are utilised as the GSN

data for the general applicability of the proposed methods. Meanwhile, we emphasise the importance of the integration of GSN data with other datasets and the fusion of advanced computing techniques for a more comprehensive understanding of urban problems and the development of more effective solutions. Our findings open new directions for the detection and prediction of spatio-temporal human mobility patterns in the urban environment to tackle various challenges cities face nowadays through collective transition behaviours in geo-social networks. We expect that our contributions will encourage researchers to rethink city studies and develop novel urban computing techniques that may complement or even replace conventional methods in the big data era. Ultimately, this dissertation has laid the foundation for new research directions that enable computational algorithms to model mobility patterns in city environments and intelligent urban services to citizens, which as whole making cities better places.

6.2 Directions for future research

The research I have presented so far in this dissertation touches on some of the key issues in urban development, utilising data-driven computational approaches. This chapter presents possible future work arising from the studies in this dissertation, followed by a discussion of other directions for future research in the urban computing field.

6.2.1 Possible future work of the studies involved

In this dissertation, we have investigated the socio-economic development of urban regions along with government-led cultural investment (Chapter 3), latent patterns extraction of user-venue interactions for optimal urban services allocation (Chapter 4), and topic-enhanced memory networks for personalised POI recommendation (Chapter 5). Next, we will discuss possible future work of the three research projects introduced in this dissertation, respectively.

While our findings in Chapter 3 have revealed the underlying relationship between socio-economic conditions, government financial efforts, and local geo-social graph features, the in-depth reasons remain unclear. Through multi-level statistical analyses and classification techniques, we discover some general rules that can assist in making effective investment decisions based on socio-economic deprivation levels and local network characteristics of neighbourhoods. However, it is questionable whether all the neighbourhoods followed the same rules. Some intriguing results might be found if we pay more attention to exceptions to the general patterns. In addition, the target variable to be predicted currently is whether the overall socio-economic status of a particular urban area becomes better or not. The research

question can be further specified to the prediction of how large the changes would be for various subdomains through regression methods.

A data-driven framework is proposed for the optimal allocation of establishments and related resources across urban areas in Chapter 4. The final outputs of the framework are a variety of precision heat maps indicating regions that lack certain types of urban services. To validate the results generated, we currently choose to examine whether the users living in high-DSR neighbourhoods but adhere to a specific cultural pattern travel longer distances in the city to access the resources they are interested in. Even though the overall result of the correlation analysis suggests the reasonableness of our approach, it can be further validated by surveys. Apart from this, the results of our model can not only benefit long-term infrastructure planning as presented in this dissertation but can also recommend short-term event venues by involving higher temporal resolution for analysis in future work.

Similar to that in Chapter 4, Chapter 5 also exploits WeChat Moments check-in data created in the city of Beijing, which makes use of the geographical locations, timestamps, as well as categorical information associated with each check-in record for personalised POI recommendation. WeChat and many other geo-social networking sites may offer data with additional types of information, which can be further investigated to enhance the performance of existing systems or promote novel architectures for POI recommendation.

Besides, for all the three research projects we introduced, conclusions are drawn on the basis of results in a single city each time. A direction for future research is exploring whether the algorithmic models proposed are suitable to other city environments. Furthermore, we plan to develop interactive data visualisation and analytics tools online based on the models devised in this dissertation.

6.2.2 Future research directions in urban computing

Given that urban computing covers a broad spectrum of topics on nearly every aspect of urban life, opportunities for expanding on the current body of research are numerous. Here we summarise some of the more salient points that may enlighten future work in this field.

Emerging challenges of ongoing rapid urbanisation. In the wake of developments in science and technology, we have become more capable of modeling and solving real-world problems by applying advanced algorithms to massive amounts of data. Under such a background, the emergence of urban computing takes a significant step towards bridging between the ever-growing power of machine learning algorithms and the increasing diversity of urban problems. Apart from socio-economics, infrastructure planning, and venue recommendation covered by this dissertation, there are other categories of urban computing tasks, including

transportation, environment, energy, and safety. With the rapid urbanisation, cities nowadays have to face some serious environmental issues, such as air quality, noise, and water pollution, which are not problems in the past. New urban challenges will keep emerging, which leads to the generation of new urban computing tasks.

Rural computing. The expansion of research topics and questions in urban computing not only reflects in the broadening of disciplines but also in the widening of geographic scope. As mentioned in Subsection 2.6.1, most of the urban computing studies currently focus on the more developed areas due to the limitation that smartphones and mobile technologies have not yet been ubiquitous in deprived regions. However, this situation can be expected to change in the coming decades. When the time comes, a variety of new subjects regarding rural development will pop up, which may significantly enrich the research field of urban computing.

Knowledge fusion. An urban computing system usually requires merging multiple data sources and multidimensional knowledge discovery. This is due to the fact that there are typically multiple reasons that lead to an urban phenomenon. For instance, the air quality in a certain area can be affected by various factors, including vehicle emissions, industrial dust, fuel burnings, as well as local geographical and meteorological conditions. Since the source of urban air pollution is a mixture, it is still hard to answer which factor plays a more prominent role in influencing the air quality for a specific location during the given period. Establishing models to deal with data from one source for a single issue with simple relationships is relatively easy. However, in urban settings, the fusion of knowledge does not mean directly putting together a collection of features extracted from multiple sources. Instead, it requires a deep understanding and appropriate usage of each data source in each part of the urban computing framework.



This dissertation has demonstrated that the massive amount of data generated from geo-social networks and the advances in computing technology hold great promise to revolutionise urban sciences. We hope the insights drawn and the techniques developed in this dissertation can help towards a better understanding of human urban movement and motivate new initiatives that tackle urban challenges to improve the quality of life in cities.

References

- Agarwal, D. and Chen, B.-C. (2010). flda: matrix factorization through latent dirichlet allocation. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 91–100. ACM.
- Altomare, A., Cesario, E., Comito, C., Marozzo, F., and Talia, D. (2017). Trajectory pattern mining for urban computing in the cloud. *IEEE Transactions on Parallel and Distributed Systems*, 28(2):586–599.
- Amoualian, H., Clausel, M., Gaussier, E., and Amini, M.-R. (2016). Streaming-lda: A copula-based approach to modeling topic dependencies in document streams. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 695–704.
- An, S., Yang, H., Wang, J., Cui, N., and Cui, J. (2016). Mining urban recurrent congestion evolution patterns from gps-equipped vehicle mobility data. *Information Sciences*, 373:515–526.
- Ankerst, M., Breunig, M. M., Kriegel, H.-P., and Sander, J. (1999). Optics: ordering points to identify the clustering structure. In *ACM Sigmod record*, volume 28, pages 49–60. ACM.
- Arcaini, P., Bordogna, G., Ienco, D., and Sterlacchini, S. (2016). User-driven geo-temporal density-based exploration of periodic and not periodic events reported in social networks. *Information Sciences*, 340:122–143.
- Armenatzoglou, N., Papadopoulos, S., and Papadias, D. (2013). A general framework for geo-social query processing. *Proceedings of the VLDB Endowment*, 6(10):913–924.
- Bahir, E. and Peled, A. (2013). Identifying and tracking major events using geo-social networks. *Social science computer review*, 31(4):458–470.
- Bao, J., Zheng, Y., and Mokbel, M. F. (2012). Location-based and preference-aware recommendation using sparse geo-social networking data. In *Proceedings of the 20th international conference on advances in geographic information systems*, pages 199–208. ACM.
- Batty, M. (2013). *The new science of cities*. MIT press.
- Bauer, S., Noulas, A., Séaghdha, D. O., Clark, S., and Mascolo, C. (2012). Talking places: Modelling and analysing linguistic content in foursquare. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*, pages 348–357. IEEE.

- Bawa-Cavia, A. (2011). Sensing the urban: using location-based social network data in urban analysis. In *Pervasive PURBA Workshop*, pages 1–7.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- Blei, D. M. and Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795.
- Botta, F., Moat, H. S., and Preis, T. (2019). Measuring the size of a crowd using instagram. *Environment and Planning B: Urban Analytics and City Science*, 0(0):2399808319841615.
- Bouma, G. (2009). Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL*, pages 31–40.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Breiman, L. (2017). *Classification and regression trees*. Routledge.
- Capineri, C. (2016). The nature of volunteered geographic information. *European handbook of crowdsourced geographic information*, page 15.
- Chang, B., Park, Y., Park, D., Kim, S., and Kang, J. (2018). Content-aware hierarchical point-of-interest embedding model for successive poi recommendation. In *IJCAI*, pages 3301–3307.
- Chen, L.-C. (2017). An effective lda-based time topic model to improve blog search performance. *Information Processing & Management*, 53(6):1299–1319.
- Chen, M., Liu, Y., and Yu, X. (2014). Nlpmm: A next location predictor with markov modeling. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 186–197. Springer.
- Chen, S., Moore, J. L., Turnbull, D., and Joachims, T. (2012). Playlist prediction via metric embedding. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 714–722. ACM.
- Cheng, C., Yang, H., King, I., and Lyu, M. R. (2012). Fused matrix factorization with geographical and social influence in location-based social networks. In *Aaai*, volume 12, pages 17–23.
- Cheng, C., Yang, H., Lyu, M. R., and King, I. (2013). Where you like to go next: Successive point-of-interest recommendation. In *IJCAI*, volume 13, pages 2605–2611.
- Cho, E., Myers, S. A., and Leskovec, J. (2011). Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090. ACM.

- Councils, L. (2014). London local government's support for arts and culture a brief overview.
- Currid, E. and Williams, S. (2009). The geography of buzz: art, culture and the social milieu in los angeles and new york. *Journal of Economic Geography*, 10(3):423–451.
- Daggitt, M. L., Noulas, A., Shaw, B., and Mascolo, C. (2016). Tracking urban activity growth globally with big location data. *Royal Society open science*, 3(4):150688.
- De Choudhury, M., Sharma, S., and Kiciman, E. (2016). Characterizing dietary choices, nutrition, and language in food deserts via social media. In *Proceedings of the 19th acm conference on computer-supported cooperative work & social computing*, pages 1157–1170. ACM.
- De Domenico, M., Lima, A., and Musolesi, M. (2013). Interdependence and predictability of human mobility and social interactions. *Pervasive and Mobile Computing*, 9(6):798–807.
- Domingos, P. and Pazzani, M. (1997). On the optimality of the simple bayesian classifier under zero-one loss. *Machine learning*, 29(2-3):103–130.
- Eagle, N., Macy, M., and Claxton, R. (2010). Network diversity and economic development. *Science*, 328(5981):1029–1031.
- Ebesu, T., Shen, B., and Fang, Y. (2018). Collaborative memory network for recommendation systems. *arXiv preprint arXiv:1804.10862*.
- Ekstrand, M. D., Riedl, J. T., Konstan, J. A., et al. (2011). Collaborative filtering recommender systems. *Foundations and Trends® in Human–Computer Interaction*, 4(2):81–173.
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231.
- Falcone, D., Mascolo, C., Comito, C., Talia, D., and Crowcroft, J. (2014). What is this place? inferring place categories through user patterns identification in geo-tagged tweets. In *6th International Conference on Mobile Computing, Applications and Services*, pages 10–19. IEEE.
- Farrahi, K. and Gatica-Perez, D. (2011). Discovering routines from large-scale human locations using probabilistic topic models. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(1):3.
- Feng, S., Cong, G., An, B., and Chee, Y. M. (2017). Poi2vec: Geographical latent representation for predicting future visitors. In *AAAI*, pages 102–108.
- Ferrari, L., Rosi, A., Mamei, M., and Zambonelli, F. (2011). Extracting urban patterns from location-based social networks. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, pages 9–16. ACM.
- Ferreira, A. P. G., Silva, T. H., and Loureiro, A. A. F. (2015). Beyond sights: Large scale study of tourists' behavior using foursquare data. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pages 1117–1124. IEEE.

- Foth, M. (2008). *Handbook of research on urban informatics: the practice and promise of the real-time city: the practice and promise of the real-time city*. IGI Global.
- Foursquare (2019). About us.
- Frias-Martinez, V., Soto, V., Hohwald, H., and Frias-Martinez, E. (2012). Characterizing urban landscapes using geolocated tweets. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, pages 239–248. IEEE.
- Galbrun, E., Pelechris, K., and Terzi, E. (2016). Urban navigation beyond shortest route: The case of safe paths. *Information Systems*, 57:160–171.
- Gao, H., Tang, J., Hu, X., and Liu, H. (2013). Exploring temporal effects for location recommendation on location-based social networks. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 93–100. ACM.
- Garcia-Duran, A., Gonzalez, R., Onoro-Rubio, D., Niepert, M., and Li, H. (2018). Transrev: Modeling reviews as translations from users to items. *arXiv preprint arXiv:1801.10095*.
- Georgiev, P. I., Noulas, A., and Mascolo, C. (2014). Where businesses thrive: Predicting the impact of the olympic games on local retailers through location-based services data. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- Gomide, J., Veloso, A., Meira Jr, W., Almeida, V., Benevenuto, F., Ferraz, F., and Teixeira, M. (2011). Dengue surveillance based on a computational model of spatio-temporal locality of twitter. In *Proceedings of the 3rd international web science conference*, page 3. ACM.
- González, S. (2011). Bilbao and barcelona ‘in motion’. how urban regeneration ‘models’ travel and mutate in the global flows of policy tourism. *Urban studies*, 48(7):1397–1418.
- Griffiths, T. (2002). Gibbs sampling in the generative model of latent dirichlet allocation.
- Guetzkow, J. et al. (2002). How the arts impact communities: An introduction to the literature on arts impact studies. Technical report.
- Haklay, M., Singleton, A., and Parker, C. (2008). Web mapping 2.0: The neogeography of the geoweb. *Geography Compass*, 2(6):2011–2039.
- Han, J., Pei, J., and Kamber, M. (2011). *Data mining: concepts and techniques*. Elsevier.
- Hartigan, J. A. and Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108.
- Hasan, S. and Ukkusuri, S. V. (2014). Urban activity pattern classification using topic models from online geo-location data. *Transportation Research Part C: Emerging Technologies*, 44:363–381.
- Hasan, S. and Ukkusuri, S. V. (2015). Location contexts of user check-ins to model urban geo life-style patterns. *PloS one*, 10(5):e0124819.

- Hasan, S., Zhan, X., and Ukkusuri, S. V. (2013). Understanding urban human activity and mobility patterns using large-scale location-based data from online social media. In *Proceedings of the 2nd ACM SIGKDD international workshop on urban computing*, page 6. ACM.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media.
- He, J., Li, X., and Liao, L. (2017a). Category-aware next point-of-interest recommendation via listwise bayesian personalized ranking. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 1837–1843. AAAI Press.
- He, R., Kang, W.-C., and McAuley, J. (2017b). Translation-based recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pages 161–169. ACM.
- He, X., Liao, L., Zhang, H., Nie, L., Hu, X., and Chua, T.-S. (2017c). Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 173–182. International World Wide Web Conferences Steering Committee.
- He, X., Zhang, H., Kan, M.-Y., and Chua, T.-S. (2016). Fast matrix factorization for online recommendation with implicit feedback. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 549–558. ACM.
- Hristova, D., Williams, M. J., Musolesi, M., Panzarasa, P., and Mascolo, C. (2016). Measuring urban social diversity using interconnected geo-social networks. In *Proceedings of the 25th international conference on World Wide Web*, pages 21–30. International World Wide Web Conferences Steering Committee.
- Hsieh, C.-K., Yang, L., Cui, Y., Lin, T.-Y., Belongie, S., and Estrin, D. (2017). Collaborative metric learning. In *Proceedings of the 26th International Conference on World Wide Web*, pages 193–201. International World Wide Web Conferences Steering Committee.
- Hu, B., Jamali, M., and Ester, M. (2013). Spatio-temporal topic modeling in mobile social media for location recommendation. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 1073–1078. IEEE.
- Hu, Y., Koren, Y., and Volinsky, C. (2008). Collaborative filtering for implicit feedback datasets. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*, pages 263–272. Ieee.
- Huynh, T., Fritz, M., and Schiele, B. (2008). Discovery of activity patterns using topic models. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 10–19. ACM.
- Jiang, B. and Miao, Y. (2015). The evolution of natural cities from the perspective of location-based social media. *The Professional Geographer*, 67(2):295–306.
- Jiang, S., Qian, X., Shen, J., Fu, Y., and Mei, T. (2015). Author topic model-based collaborative filtering for personalized poi recommendations. *IEEE transactions on multimedia*, 17(6):907–918.

- Jonathan, V. (2007). The emergence of culture-led regeneration: a policy concept and its discontents. *Research paper, Centre for Cultural Policy Studies, University of Warwick*.
- Joseph, K., Tan, C. H., and Carley, K. M. (2012). Beyond local, categories and friends: clustering foursquare users with latent topics. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 919–926. ACM.
- Kang, W.-C., Wan, M., and McAuley, J. (2018). Recommendation through mixtures of heterogeneous item relationships. *arXiv preprint arXiv:1808.10031*.
- Karamshuk, D., Noulas, A., Scellato, S., Nicosia, V., and Mascolo, C. (2013). Geo-spotting: mining online location-based services for optimal retail store placement. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 793–801. ACM.
- Karina, V. R. C., Quercia, D., Francesco, B., Fraternali, P., et al. (2015). Taxonomy-based discovery and annotation of functional areas in the city. In *International Conference on Web and Social Media*, pages 445–453.
- Keddie, J. (2014). *Negotiating urban change in gentrifying London: experiences of long-term residents and early gentrifiers in Bermondsey*. PhD thesis, The London School of Economics and Political Science (LSE).
- Kim, D. and Shin, H. (2017). An explanatory matrix factorization with user comments data. In *RecSys Posters*.
- Kindberg, T., Chalmers, M., and Paulos, E. (2007). Guest editors' introduction: Urban computing. *IEEE Pervasive Computing*, 6(3):18–20.
- Koren, Y. (2008). Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434. ACM.
- Koren, Y. and Bell, R. (2015). Advances in collaborative filtering. In *Recommender systems handbook*, pages 77–118. Springer.
- Koren, Y., Bell, R., and Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, (8):30–37.
- Kriegel, H.-P., Kröger, P., Sander, J., and Zimek, A. (2011). Density-based clustering. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(3):231–240.
- Kulis, B. et al. (2013). Metric learning: A survey. *Foundations and Trends® in Machine Learning*, 5(4):287–364.
- Kurashima, T., Iwata, T., Hoshide, T., Takaya, N., and Fujimura, K. (2013). Geo topic model: joint modeling of user's activity area and interests for location recommendation. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 375–384. ACM.
- Lathia, N. and Capra, L. (2011). How smart is your smartcard?: measuring travel behaviours, perceptions, and incentives. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 291–300. ACM.

- Le Falher, G., Gionis, A., and Mathioudakis, M. (2015). Where is the soho of rome? measures and algorithms for finding similar neighborhoods in cities. In *Ninth International AAAI Conference on Web and Social Media*.
- Lee, R. and Sumiya, K. (2010). Measuring geographical regularities of crowd behaviors for twitter-based geo-social event detection. In *Proceedings of the 2nd ACM SIGSPATIAL international workshop on location based social networks*, pages 1–10. ACM.
- Li, H., Ge, Y., Hong, R., and Zhu, H. (2016). Point-of-interest recommendations: Learning potential check-ins from friends. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 975–984. ACM.
- Lian, D., Zhao, C., Xie, X., Sun, G., Chen, E., and Rui, Y. (2014). Geomf: joint geographical modeling and matrix factorization for point-of-interest recommendation. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 831–840. ACM.
- Lima, A., De Domenico, M., Pejovic, V., and Musolesi, M. (2015). Disease containment strategies based on mobility and information dissemination. *Scientific reports*, 5:10650.
- Linden, G., Smith, B., and York, J. (2003). Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing*, (1):76–80.
- Liu, Q., Wu, S., Wang, L., and Tan, T. (2016). Predicting the next location: A recurrent model with spatial and temporal contexts. In *AAAI*, pages 194–200.
- Liu, Y., Wei, W., Sun, A., and Miao, C. (2014). Exploiting geographical neighborhood characteristics for location recommendation. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 739–748. ACM.
- local government, L. (2015). London local government’s support for arts and culture a brief overview.
- Long, X., Jin, L., and Joshi, J. (2012). Exploring trajectory-driven local geographic topics in foursquare. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 927–934. ACM.
- Louail, T., Lenormand, M., Arias, J. M., and Ramasco, J. J. (2017). Crowdsourcing the robin hood effect in cities. *Applied network science*, 2(1):11.
- Manotumruksa, J., Macdonald, C., and Ounis, I. (2018). A contextual attention recurrent architecture for context-aware venue recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 555–564. ACM.
- Matsu, Y., Okazaki, M., and Sakak, T. (2010). Earthquake shakes twitter users: real-time event detection by social sensors. In *WWW 2010: Proceedings of the 19th World Wide Web Conference*.

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Miller, R. E. and Blair, P. D. (2009). *Input-output analysis: foundations and extensions*. Cambridge university press.
- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Nations, U. (2018). 2018 revision of world urbanization prospects.
- Noulas, A., Scellato, S., Mascolo, C., and Pontil, M. (2011). An empirical study of geographic user activity patterns in foursquare. In *Fifth international AAAI conference on weblogs and social media*.
- Noulas, A., Shaw, B., Lambiotte, R., and Mascolo, C. (2015). Topological properties and temporal dynamics of place networks in urban environments. In *Proceedings of the 24th International Conference on World Wide Web*, pages 431–441. ACM.
- ODPM (1999). Regeneration through culture, sport and tourism.
- Oyelade, J., Isewon, I., Oladipupo, F., Aromolaran, O., Uwoghiren, E., Ameh, F., Achas, M., and Adebisi, E. (2016). Clustering algorithms: Their application to gene expression data. *Bioinformatics and Biology insights*, 10:BBI–S38316.
- Paldino, S., Bojic, I., Sobolevsky, S., Ratti, C., and González, M. C. (2015). Urban magnetism through the lens of geo-tagged photography. *EPJ Data Science*, 4(1):5.
- Paul, M. J. and Dredze, M. (2011). You are what you tweet: Analyzing twitter for public health. In *Fifth International AAAI Conference on Weblogs and Social Media*.
- Paulos, E., Anderson, K., and Townsend, A. (2004). Ubicomp in the urban frontier.
- Poole, D. L. and Mackworth, A. K. (2010). *Artificial Intelligence: foundations of computational agents*. Cambridge University Press.
- Psyllidis, A., Bozzon, A., Bocconi, S., and Bolivar, C. T. (2015). A platform for urban analytics and semantic data integration in city planning. In *International conference on computer-aided architectural design futures*, pages 21–36. Springer.
- Quercia, D. and Saez, D. (2014). Mining urban deprivation from foursquare: Implicit crowdsourcing of city land use. *IEEE Pervasive Computing*, 13(2):30–36.
- Quercia, D., Séaghdha, D. Ó., and Crowcroft, J. (2012). Talk of the city: Our tweets, our community happiness. In *Sixth International AAAI Conference on Weblogs and Social Media*.
- Rendle, S., Freudenthaler, C., Gantner, Z., and Schmidt-Thieme, L. (2009). Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 452–461. AUAI Press.

- Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, J. (1994). Grouplens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186. ACM.
- Röder, M., Both, A., and Hinneburg, A. (2015). Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*, pages 399–408. ACM.
- Sassen, S. (2013). *The global city: New york, london, tokyo*. Princeton University Press.
- Sattler, B. and Brandes, P. (2015). The megacity state: The world’s biggest cities shaping our future.
- Scellato, S., Mascolo, C., Musolesi, M., and Latora, V. (2010). Distance matters: Geo-social metrics for online social networks. In *WOSN*.
- Schütz, M. H. (2017). Australia’s regional innovation systems: inter-industry interaction in innovative activities in three australian territories. *Economic Systems Research*, 29(3):357–384.
- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Lucas, R. E., Agrawal, M., Park, G. J., Lakshminanth, S. K., Jha, S., Seligman, M. E., et al. (2013). Characterizing geographic variation in well-being using tweets. In *Seventh International AAAI Conference on Weblogs and Social Media*.
- Seresinhe, C. I., Preis, T., MacKerron, G., and Moat, H. S. (2019). Happiness is greater in more scenic locations. *Scientific reports*, 9(1):4498.
- Seresinhe, C. I., Preis, T., and Moat, H. S. (2016). Quantifying the link between art and property prices in urban neighbourhoods. *Royal Society open science*, 3(4):160146.
- Shen, Y. and Karimi, K. (2016). Urban function connectivity: Characterisation of functional urban streets with social media check-in data. *Cities*, 55:9–21.
- Silva, T. H., Viana, A. C., Benevenuto, F., Villas, L., Salles, J., Loureiro, A., and Quercia, D. (2019). Urban computing leveraging location-based social network data: A survey. *ACM Computing Surveys (CSUR)*, 52(1):17.
- Smith, C., Mashhadi, A., and Capra, L. (2013). Ubiquitous sensing for mapping poverty in developing countries. *Paper submitted to the Orange D4D Challenge*.
- Steiger, E., Westerholt, R., and Zipf, A. (2016). Research on social media feeds—a giscience perspective. *European Handbook of Crowdsourced Geographic Information*, page 237.
- Sun, L. and Yin, Y. (2017). Discovering themes and trends in transportation research using topic modeling. *Transportation Research Part C: Emerging Technologies*, 77:49–66.
- Tang, J., Hu, X., Gao, H., and Liu, H. (2013). Exploiting local and global social context for recommendation. In *IJCAI*, volume 13, pages 2712–2718.
- Tay, Y., Tuan, L. A., and Hui, S. C. (2018). Latent relational metric learning via memory-based attention for collaborative ranking. In *Proceedings of the 27th International Conference on World Wide Web*.

- Trevor, H., Robert, T., and JH, F. (2009). The elements of statistical learning: data mining, inference, and prediction.
- Venerandi, A., Quattrone, G., Capra, L., Quercia, D., and Saez-Trumper, D. (2015). Measuring urban deprivation from user generated content. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 254–264. ACM.
- Wang, H., Fu, Y., Wang, Q., Yin, H., Du, C., and Xiong, H. (2017a). A location-sentiment-aware recommender system for both home-town and out-of-town users. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1143. ACM.
- Wang, H., Kifer, D., Graif, C., and Li, Z. (2016). Crime rate inference with big data. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 635–644. ACM.
- Wang, H., Shen, H., Ouyang, W., and Cheng, X. (2018). Exploiting poi-specific geographical influence for point-of-interest recommendation. In *IJCAI*, pages 3877–3883.
- Wang, S., Wang, Y., Tang, J., Shu, K., Ranganath, S., and Liu, H. (2017b). What your images reveal: Exploiting visual contents for point-of-interest recommendation. In *Proceedings of the 26th International Conference on World Wide Web*, pages 391–400. International World Wide Web Conferences Steering Committee.
- Wang, Y., Agichtein, E., and Benzi, M. (2012). Tm-lda: efficient online modeling of latent topic transitions in social media. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 123–131.
- Webber, R. (2007). The metropolitan habitus: its manifestations, locations, and consumption profiles. *Environment and Planning A*, 39(1):182–207.
- WeChat (2017). The 2017 wechat data report.
- WeChat (2019). The 2018 wechat data report.
- Weng, X., Liu, Y., Song, H., Yao, S., and Zhang, P. (2018). Mining urban passengers’ travel patterns from incomplete data with use cases. *Computer Networks*, 134:116–126.
- Wyly, E. (2014). The new quantitative revolution. *Dialogues in Human Geography*, 4(1):26–38.
- Xie, M., Yin, H., Wang, H., Xu, F., Chen, W., and Wang, S. (2016). Learning graph-based poi embedding for location-based recommendation. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 15–24. ACM.
- Xing, E. P., Jordan, M. I., Russell, S. J., and Ng, A. Y. (2003). Distance metric learning with application to clustering with side-information. In *Advances in neural information processing systems*, pages 521–528.

- Ye, M., Liu, X., and Lee, W.-C. (2011). Exploring social influence for recommendation—a probabilistic generative model approach. *arXiv preprint arXiv:1109.0758*.
- Yin, H., Sun, Y., Cui, B., Hu, Z., and Chen, L. (2013). Lcars: a location-content-aware recommender system. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 221–229. ACM.
- Yin, H., Wang, W., Wang, H., Chen, L., and Zhou, X. (2017). Spatial-aware hierarchical collaborative deep learning for poi recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 29(11):2537–2551.
- Yuan, J., Zheng, Y., and Xie, X. (2012). Discovering regions of different functions in a city using human mobility and pois. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 186–194. ACM.
- Yuan, J., Zheng, Y., Xie, X., and Sun, G. (2011). Driving with knowledge from the physical world. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 316–324. ACM.
- Yuan, N. J., Zhang, F., Lian, D., Zheng, K., Yu, S., and Xie, X. (2013). We know how you live: exploring the spectrum of urban lifestyles. In *Proceedings of the first ACM conference on Online social networks*, pages 3–14. ACM.
- Zhang, C., Shou, L., Chen, K., Chen, G., and Bei, Y. (2012). Evaluating geo-social influence in location-based social networks. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 1442–1451. ACM.
- Zhang, C., Zhang, K., Yuan, Q., Peng, H., Zheng, Y., Hanratty, T., Wang, S., and Han, J. (2017). Regions, periods, activities: Uncovering urban dynamics via cross-modal representation learning. In *Proceedings of the 26th International Conference on World Wide Web*, pages 361–370. International World Wide Web Conferences Steering Committee.
- Zhang, C., Zhang, K., Yuan, Q., Zhang, L., Hanratty, T., and Han, J. (2016). Gmove: Group-level mobility modeling using geo-tagged social media. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1305–1314. ACM.
- Zhang, J.-D. and Chow, C.-Y. (2015). Geosoca: Exploiting geographical, social and categorical correlations for point-of-interest recommendations. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 443–452. ACM.
- Zhang, J.-D., Chow, C.-Y., and Li, Y. (2014). Lore: Exploiting sequential influence for location recommendations. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 103–112. ACM.
- Zhang, K. and Pelechris, K. (2014). Understanding spatial homophily: the case of peer influence and social selection. In *Proceedings of the 23rd international conference on World wide web*, pages 271–282. ACM.
- Zhang, Y. and Chen, X. (2018). Explainable recommendation: A survey and new perspectives. *arXiv preprint arXiv:1804.11192*.

- Zhao, K., Cong, G., Yuan, Q., and Zhu, K. Q. (2015). Sar: A sentiment-aspect-region model for user preference analysis in geo-tagged reviews. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*, pages 675–686. IEEE.
- Zhao, S., Zhao, T., King, I., and Lyu, M. R. (2017). Geo-teaser: Geo-temporal sequential embedding rank for point-of-interest recommendation. In *Proceedings of the 26th international conference on world wide web companion*, pages 153–162. International World Wide Web Conferences Steering Committee.
- Zheng, Y., Capra, L., Wolfson, O., and Yang, H. (2014a). Urban computing: concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(3):38.
- Zheng, Y., Liu, T., Wang, Y., Zhu, Y., Liu, Y., and Chang, E. (2014b). Diagnosing new york city’s noises with ubiquitous data. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 715–725. ACM.
- Zheng, Y., Yi, X., Li, M., Li, R., Shan, Z., Chang, E., and Li, T. (2015). Forecasting fine-grained air quality based on big data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2267–2276. ACM.
- Zheng, Y. and Zhou, X. (2011). *Computing with spatial trajectories*. Springer Science & Business Media.
- Zhi, Y., Li, H., Wang, D., Deng, M., Wang, S., Gao, J., Duan, Z., and Liu, Y. (2016). Latent spatio-temporal activity structures: a new approach to inferring intra-urban functional regions via social media check-in data. *Geo-spatial Information Science*, 19(2):94–105.
- Zhou, M., Ding, Z., Tang, J., and Yin, D. (2018a). Micro behaviors: A new perspective in e-commerce recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 727–735. ACM.
- Zhou, X., Hristova, D., Noulas, A., Mascolo, C., and Sklar, M. (2017). Cultural investment and urban socio-economic development: a geosocial network approach. *Royal Society open science*, 4(9):170413.
- Zhou, X., Mascolo, C., and Zhao, Z. (2019). Topic-enhanced memory networks for personalised point-of-interest recommendation. *arXiv preprint arXiv:1905.13127*.
- Zhou, X., Noulas, A., Mascolo, C., and Zhao, Z. (2018b). Discovering latent patterns of urban cultural interactions in wechat for modern city planning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1069–1078. ACM.