

Mobility Data Science: Perspectives and Challenges*

MOHAMED MOKBEL, MAHMOUD SAKR, LI XIONG, ANDREAS ZÄIJFLE, JUSSARA ALMEIDA, TAYLOR ANDERSON, WALID G. AREF, GENNADY ANDRIENKO, NATALIA ANDRIENKO, YANG CAO, SANJAY CHAWLA, REYNOLD CHENG, PANOS CHRYSANTHIS, XIQI FEI, GABRIEL GHINITA, ANITA GRASER, DIMITRIOS GUNOPULOS, CHRISTIAN S. JENSEN, JOON-SEOK KIM, KYOUNG-SOOK KIM, PEER KRÖGER, JOHN KRUMM, JOHANNES LAUER, AMR MAGDY, MARIO A. NASCIMENTO, SIVA RAVADA, MATTHIAS RENZ, DIMITRIS SACHARIDIS, FLORA SALIM, MOHAMED SARWAT, MAXIME SCHOEMANS, CYRUS SHAHABI, BETTINA SPECKMANN, EGEMEN TANIN, XU TENG, YANNIS THEODORIDIS, KRISTIAN TORP, GOCE TRAJCEVSKI, MARC VAN KREVELD, CAROLA WENK, MARTIN WERNER, RAYMOND WONG, SONG WU, JIANQIU XU, MOUSTAFA YOUSSEF, DEMETRIS ZEINALIPOUR, MENGXUAN ZHANG, and ESTEBAN ZIMÁNYI

Mobility data captures the locations of moving objects such as humans, animals, and cars. With the availability of GPS-equipped mobile devices and other inexpensive location-tracking technologies, mobility data is collected ubiquitously. In recent years, the use of mobility data has demonstrated significant impact in various domains including traffic management, urban planning, and health sciences. In this paper, we present the domain of mobility data science. Towards a unified approach to mobility data science, we present a pipeline having the following components: mobility data collection, cleaning, analysis, management, and privacy. For each of these components, we explain how mobility data science differs from general data science, we survey the current state of the art, and describe open challenges for the research community in the coming years.

ACM Reference Format:

Mohamed Mokbel, Mahmoud Sakr, Li Xiong, Andreas ZÄijfle, Jussara Almeida, Taylor Anderson, Walid G. Aref, Gennady Andrienko, Natalia Andrienko, Yang Cao, Sanjay Chawla, Reynold Cheng, Panos Chrysanthos, Xiqi Fei, Gabriel Ghinita, Anita Graser, Dimitrios Gunopulos, Christian S. Jensen, Joon-Seok Kim, Kyoung-Sook Kim, Peer Kröger, John Krumm, Johannes Lauer, Amr Magdy, Mario A. Nascimento, Siva Ravada, Matthias Renz, Dimitris Sacharidis, Flora Salim, Mohamed Sarwat, Maxime Schoemans, Cyrus Shahabi, Bettina Speckmann, Egemen Tanin, Xu Teng, Yannis Theodoridis, Kristian Torp, Goce Trajcevski, Marc van Kreveld, Carola Wenk, Martin Werner, Raymond Wong, Song Wu, Jianqiu Xu, Moustafa Youssef, Demetris Zeinalipour, Mengxuan Zhang, and Esteban Zimányi. 2018. Mobility Data Science: Perspectives and Challenges. *ACM Trans. Graph.* 37, 4, Article 111 (August 2018), 30 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

*This is a community publication. The authors of this article met in Dagstuhl for Seminar #22021 on Mobility Data Science [153]. The first four authors co-organized the Dagstuhl seminar leading to this article and coordinated the creation of this manuscript. All other authors contributed equally to this research. The seminar was held in the week of January 9 - 14, 2022. It had 47 participants specialized in different topics: data management, mobility analysis, geography, privacy, urban computing, systems, simulation, indoors, visualization, information integration, and theory. Due to COVID-19, the seminar took place in hybrid mode, with 8 onsite, and 39 remote participants. Despite the challenge of different time zones of the participants, all sessions were attended by at least 37 participants. It was an excellent opportunity to start the discussion about next-decade opportunities and challenges for mobility data science.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

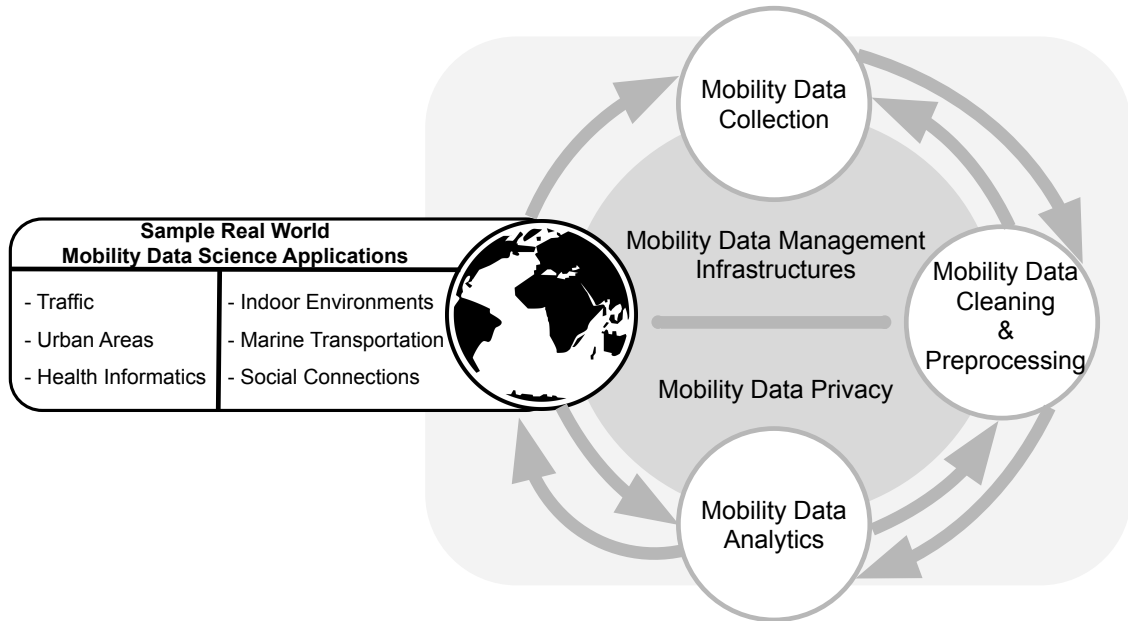


Fig. 1. The Mobility Data Science Pipeline

1 INTRODUCTION

The volume of mobility data being collected has been steadily increasing since the advent of affordable personal location-enabled mobile devices. Examples of mobility data continuously generated and collected in huge volumes include: (a) individual sporadic locations obtained from mobile app data and location-based social networks, (b) individual pedestrians, biking, or driving trajectories constrained by underlying side walks, biking trails, and road network, respectively, (c) indoor individual or asset tracking data obtained from RFID and bluetooth devices, (d) athletes movement data in various sports obtained from wearable devices, (e) public transportation, taxis, ride sharing, and delivery logistics trajectories obtained by location-tracking devices and specially designed app services, (f) aircraft and vessels trajectories moving in an unconstrained environment (i.e., no underlying road network) obtained by air and sea traffic monitoring services, and (g) animal tracking data moving freely in the space obtained from physically tagged and remotely sensed animals. Generally speaking, for each moving object, mobility data is typically available in the form of a sequence of (*location*, *timestamp*) pairs. The *location* attribute could be as simple as a point, represented by either *latitude* and *longitude* coordinates or as relative coordinates with respect to the underlying space. The *location* attribute could also be an area, which can represent the mobility of objects with spatial extents, e.g., flocks or group movement.

The ability of understanding and analyzing mobility data is crucial for various widely used important sectors and applications. In transportation and traffic management, analyzing traffic data through vehicle mobility helps in predicting accidents [159], traffic congestions [260], and better route planning [51]. In ride sharing and delivery logistics application, analyzing trip mobility data help in data-driven eco route planning, which results in huge cost and energy savings [97]. In location-based services, analyzing people movements around the city significantly helps in trip planning activities [219], finding popular tourists sites and restaurants [119], and data-driven routing and querying [220]. In indoor navigation, understanding how people move indoors helps in understanding the traffic for various stores inside a mall, which is needed in various market research studies [115]. In urban planning, driving data can significantly

help in building highly accurate, reliable, and annotated maps [160] as well as deciding on good location for various facilities, e.g., restaurants, retail stores, and clinics [208]. In social computing, analyzing how people move in cities and regions helps in understanding the demand for infrastructure and energy as a means of reducing inequalities [202]. In disaster response, analyzing crowd movement helps in preparing for natural disasters through rescuing and evacuation efforts [106]. In health informatics, connected wearables can monitor and analyze the movement of elderly people, allowing for timely, and potentially life-saving, interventions [135]. In pandemic prevention, the ability of privacy-preserving individual tracking allows for contact tracing, which was deemed as cornerstone in limiting pandemic spread [156, 279].

Despite the common goal of acquiring, managing, and generating insights from mobility data, the mobility data science community is largely fragmented, developing solutions in silos. It stems from a range of disciplines with expertise in moving object data storage and management [100], geographic information science [89], spatiotemporal data mining [212], human mobility modelling [27], ubiquitous computing, computational geometry and more. The sheer volumes of mobility data along with the immense need of mobility data analysis in various applications call for employing a complete Data Science pipeline [192] over mobility data. This includes the whole pipeline of Data Science applications, starting from the data storage and management infrastructure and going through data collection, data cleaning and preprocessing, and data analysis. Unfortunately, this is not straightforward as current Data Science systems, tools, and algorithms are not directly applicable to mobility data. This is mainly due to the fact that these systems, tools, and algorithms, are designed in a generic way to support any data type, and hence they do not lend themselves to the distinguishing characteristics of mobility data. Examples of such characteristics include the spatial and temporal dimensions of the data, the rate of updates, and the privacy requirements. In particular, mobility data is always spatial, where nearby objects are more related to each other. This is unlike traditional data, where the concept of nearby and locality is not taken into account. Also, similar to time series data, mobility data is temporal, where one object may have hundreds of updates to its location, and all updates are related to each other (e.g., one trajectory). This is again unlike traditional data, where temporal updates of a single object are not frequent and older updates would be of less importance. Similar to streaming data, mobility data has a high frequency of updates, which is not supported in typical data science applications. Finally, mobility data is more sensitive to privacy. While privacy-preserving in traditional data can be achieved by removing (quasi-)identifier attributes, in mobility data, locations by themselves are considered private information that can reveal not only the users' identities, but also their behavior, life style, medical conditions, and work places.

Motivated by ubiquity and sheer volumes of mobility data, the importance of mobility applications, and the lack of support from current data science pipelines, this paper presents a pipeline for Mobility Data Science. We define *Mobility Data Science* as an interdisciplinary field that uses scientific methods, processes, algorithms and systems to extract or extrapolate knowledge and insights from potentially noisy, structured and unstructured mobility data, and apply knowledge from mobility data across a broad range of application domains. While currently, the community of developers, practitioners, and researchers, dealing with mobility data use off-the-shelf data science techniques and systems to collect, clean, manage, and analyze their mobility data, we firmly believe that this ends up to sub-bar performance. We urge such community to build its own mobility data science pipeline to better serve its own purpose. This paper makes the case for the need for a mobility data science pipeline along with the challenges that need to be addressed to realize it.

2 MOBILITY DATA COLLECTION

The abundance availability of real data is a cornerstone to any data science application, and mobility data science applications are of no exception. However, it is much easier to collect and find tons of data for data science applications than it is the case for mobility data science. In particular, for data science applications, well-established research in anonymizing personal data allows wide data sharing. This is to the extent that governments have released various datasets for public (e.g., Data.gov). In addition, companies already collect their own inventory data that does not include any personal identifiers, and hence it is suitable to be fed to data science applications. On the other side, data-driven mobility data science research has been in a constant struggle with the need for available mobility data. A main reason is that non-aggregated individual human location data is considered personal identifiable information as it may lead to tracing an individual’s identity. For example, it has been shown that only a few spatial locations are sufficient to uniquely identify individuals even among a large population of people [204]. As a result, most datasets are collected in aggregated form, which hinders the deployment of various mobility data science applications. This sections discuss current efforts and challenges of mobility data collection.

2.1 Efforts in Mobility Data Collection

Before the wide availability of personal digital devices, human mobility data collection was expensive and therefore datasets were very sparse. With the advent of personal location-enabled devices, many people’s movements have started leaving digital traces that are being collected either by industry as a means of providing location-based services [198] or by governmental entries as a means of data analysis, e.g., traffic-related studies [234]. However, this did not result in a similar explosion of publicly available mobility data, mainly due to privacy and data sharing concerns.

Current efforts in releasing public non-aggregated mobility data is mainly limited to small datasets, small regions, while removing locations that can lead to one’s whereabouts. This mostly include trips obtained from taxis, ride sharing services, or public transportation. Some of these datasets include detailed trajectory data for the following cities (ordered alphabetically): (1) *Athens* [28]. 500K trajectories collected over 5 days in downtown Athens, Greece, (2) *Beijing 1* [273]. 17+K trajectories with 26 Million GPS points over three years in Beijing, China, (3) *Beijing 2* [261]. 10+K trajectories with 15 Million GPS points over one week in Beijing, China, (4) *Rio* [70]. 12+K Buses with detailed trajectories of 118+ Million GPS points over 30 days in Rio de Janeiro, Brazil, (5) *Rome* [41]. 320 Taxis with detailed trajectories of 21+ Million GPS points over 30 days in Rome, Italy, (6) *San Francisco 1* [180]. 536 Taxis with detailed trajectories of 11+ Million GPS points over 30 days in San Francisco, CA, USA, (7) *San Francisco 2* [1]. 20+K detailed trajectories with 5+ Million GPS points in San Francisco, CA, USA. (8) *Shenzhen* [240]. 664 Taxis with detailed trajectories of 1.1+ Millions GPS points over one day in Shenzhen, China, (9) *Singapore* [108]. 84K trajectories with 80+ Million GPS points over a month in Singapore. Other datasets only include the origin and destination of each trajectory. Examples include the following cities: (1) *Austin* [194]. 1.5 Million trips for a period of 10 months in Austin, TX, USA, (2) *Guangdong* [258]. 2.5 Million trips over one day in Guangdong Province, China, (3) *NYC* [169]. 1.5 Million Taxi trips over a period of 6 months in New York City, NY, USA, (4) *Porto* [181]. 426K Taxi trips over three months in Porto, Portugal.

Other than trip and trajectory road network data, there are tons of available biking data across the world, including tens of millions trips in Bay Area [140], Boston [139], Chicago [71], Columbus [62]. London [136], Los Angeles [38], Madrid [142], Minneapolis [165], New York City [60], Philadelphia [112], Portland [39], and Washington D.C. [44]. There are also available public marine traffic that include detailed vessel trajectories (e.g. [177]), sport data sets for basketball and soccer that include a variety of events took place in major leagues within one season [175], and indoor

data about the behaviors of nearly 30 Year-10 students and their teachers collected over four weeks in Australia, with spatial reference (associations to rooms) and highly granular wearable data [84].

However, there are some large-scale aggregated datasets on a coarse granularity that can help in high level analysis, but not to get insight details of mobility data. Examples of such aggregate data include origin-destination employment statistics in the USA that contains home-to-work commuting flows aggregated to census tract level [91], cell phone trace datasets capturing the locations of individuals aggregated to their nearest cell tower [237], foot traffic data of check-ins of 35 million anonymized mobile devices in USA are aggregated to census block groups [199], and a global database about aggregate indoor occupant behavior, composed of 34 datasets from 15 countries and 39 institutions, collected by occupancy sensors that measure the occupancy count of each space being monitored [73]. An additional source of human mobility data is location-based social network (LBSN) data. LBSN data captures both 1) discrete check-ins between users and locations, and 2) a social network between users. This dimension of location bridges the gap between the physical world and online social networking services [271]. However, it has been shown existing LBSN data sets are too small to broadly understand, analyze, and predict human behavior [127].

The lack of available mobility data, combined with the need to stress test various research ideas have motivated various research groups to either develop their own data simulators or develop publicly available simulators that can also be used by other researchers for benchmark datasets. However, such simulators were mainly designed to test specific aspects of research, but not meant to be representative of real mobility data. For example, various simulators were mainly designed to test new index structures for mobility data, query processing algorithms, and system infrastructure scalability for managing spatiotemporal data (e.g., [154]). Within the transportation community, more fine granularity simulators (e.g., [34]) were proposed to study traffic infrastructure, but none of them is meant to provide comprehensive mobility study.

2.2 Challenges in Mobility Data Collection

This section presents some of the challenges in mobility data collection that the community needs to address towards realizing the pipeline of mobility data science.

Challenge 1. Mobility Data Privacy. In most cases, (human) mobility data is sensitive and considered as personal identifiable information. This puts major privacy concerns on data sharing. Hence, any attempt to collect fine granularity detailed trajectory or human mobility data must first address the privacy challenge. Though the general topic of data privacy has been well studied in literature with practical solutions, such solutions are not directly applicable to the case of mobility data. In particular, mobility data gives rise to the TUL (Trajectory-User Linking) problem [86]. To protect users' actual locations, while preserving meaningful mobility information for various learning tasks, one may wish to generate realistic motions based on real-world mobility datasets [274]. Since privacy is a core problem in mobility data that does not only impact data collection, but also impacts all other components of the mobility data science pipeline, we dedicate Section 6 to discuss mobility data privacy in details.

Challenge 2. Mobility Data Bias. Mobility data collection procedures suffer from all kinds of bias. For example, mobile applications data and mobile phone network data are biased against people who do not use smart phones or use prepaid plans. Most traffic counting sensors are installed to count cars but do not count pedestrians, cyclists, wheelchairs, or similar. Cells in mobile phone networks vary widely in size. The data traces that are usually collected in cellular networks are cellular themselves. This affects rural areas with larger cells more than urban areas. Volunteered tracking data is biased towards technically savvy people. Sports tracking data is biased towards health conscious middle

and upper class. It is important to understand, measure, and mitigate data bias in mobility datasets to ensure that actions and policies that are based on mobility data science results are equitable, fair, and include vulnerable populations [207].

Challenge 3. Incentives for Data Sharing. Users need to have good incentives to share their locations. To some degree, users kind of agree to share their locations with commercial entities to get location-based services, ride sharing, cell phone coverage, delivery, among other services. However, it is understood that users would be reluctant to publicly share their mobility traces. Meanwhile, the biking community have shown a great example for sharing their biking trails. A main reason is that, in many places of the world, most of these trails are not really home-to-work commuting, but it is more of an outdoor activities. Hence, sharing biking trails helps fellow bikers in knowing the conditions of biking trails, which is a great incentive for sharing. More incentives need to be given for drivers to share their mobility traces, even for sporadic trips that do not lead to identifiable locations. Sharing could be for part of the trajectory where rewards are given back based on the sharing length and resolution. A gamification concept may be exploited to encourage more participants to share.

Challenge 4. Simulated Mobility Data. The dire need to mobility data along with the difficulty of obtaining them made it apparent that simulated synthetic data is immensely needed to enrich and train mobility data science applications. However, the challenge is to go beyond earlier attempts of simulating data for testing very specific techniques to simulating data for the general purpose of having realistic life scenarios. Empowered by modern computational capabilities that make it possible to simulate large populations, the mobility community should work with social scientists to create realistic individual-level human mobility data. Lessons have been learned from the experience of the deep learning community, by applying generative adversarial networks (GANs) for trajectory generation [264]. However, it is unclear as of yet, how to measure the realism of mobility data. If synthetic mobility data is too realistic, for example, due to training on real human trajectories, it may invade someone's privacy if, for instance, it shows where members of a given household actually visit. On the flipside, benchmark data that is too disconnected from the real-world and does not represent realistic human behavior would not allow to generalize to the real-world.

3 MOBILITY DATA CLEANING

Until the early 21st century, location data and mobility data available for geographic information science (GIS) was mainly collected, curated, standardized [79, 80], and published by authoritative sources such as the United States Geological Survey (USGS) [233]. Now, data used for mobility data science is often obtained from sources of volunteered geographic information (VGI) [218]. Such data is contributed by millions of individual users (more than ten million contributors in the case of OpenStreetMap [171]) and is rarely curated. Mobility data collected from such sources is highly uncertain due to physical limitations of sensing devices, due to obsolescence of observations, and in many cases plain incorrect due to deliberate misinformation [158]. Consequentially, our ability to unearth valuable knowledge from large sets of mobility data is often impaired by the uncertainty of the data which geography has been named the “the Achilles heel of GIS” [90].

Data cleaning and preprocessing is a milestone to all data science. In fact, it has been reported that data scientists spend more than 80% of their time in data cleaning and preparations [163]. As a result, there are huge efforts in the data science community dedicated to developing various data cleaning algorithms [57] and full-fledged systems [68]. Mobility data is of no exception in terms of its need for data cleaning and preparation procedures. But for numerous reasons, data cleaning and preparation yields unique challenges. This section discusses current efforts and challenges of mobility data cleaning.

3.1 Efforts in Mobility Data Cleaning

A recent survey [126] and data quality assessment tool [92] have discussed various sorts of errors that negatively impact data quality in spatial and mobile environments. Motivated by the inaccuracy of location tracking devices, several efforts were dedicated to address: (a) the spatial inherent inaccuracy of GPS devices and (b) the uncertainty of moving objects whereabouts between each two known locations, which is a result of low sampling rates due to bandwidth and battery limitations.

As the spatial inaccuracy indicates erroneous GPS coordinates, the efforts to identify and correct such coordinates have focused on either finding and eliminating outliers or map matching all coordinates to an underlying fixed and trusted infrastructure (e.g., road network map). For the case of map matching, existing efforts aim to match/snap all GPS traces to an underlying road network [42, 46]. Proposed techniques vary from as simple as snapping each point to its nearest road to applying Markov Chain to identify the most probable road segment that each point should be snapped to. In case there is no underlying road infrastructure (e.g., marine transportation or animal movement), outlier detection techniques are used to identify and remove erroneous points [226].

Irrespective of the collection method and device settings, there is also indispensable uncertainty in movement data caused by their discreteness. Since time is continuous, the data cannot refer to every possible instant. For any two successive instants, there is a temporal gap where the whereabouts of the moving objects are unknown. To overcome such location uncertainty, several efforts were dedicated to modeling the uncertainty of mobility data surveyed in [280].

3.2 Challenges in Mobility Data Cleaning

This section delves into some challenges linked to cleaning mobility data that the community need to tackle.

Challenge 5. Inaccuracy in the Movement Space Infrastructure. A unique challenge in mobility data is that in many cases, its reference points are the ones that are inaccurate. In particular, mobility data that represent movement on a road network may be more accurate than the road network itself. Road networks, like any other type of data, suffer from all sorts of inaccuracy, and may not be even available in many places [161]. In fact, Microsoft has recently announced that it has found more than one million kilometers of roads missing from current maps [149]. This is why there is a whole area of industrial and academic research about map inference, which aims to infer (all or missing parts) of the road network from either satellite images [29] or trajectory data [37]. However, almost all of these techniques focus on making accurate maps in terms of topology. There need to be more efforts on map inference algorithms that go beyond inferring the map topology to inferring map metadata (e.g., road speed, traffic lights, number of lanes, and turns), without which, mobility data would not be accurate as its road network reference itself is missing important data. A major step towards cleaning mobility data would be to first clean its reference map.

Challenge 6. Filling in Temporal Mobility Gaps. As mentioned earlier, there are lots of efforts dedicated to modeling the uncertainty of moving objects whereabouts between each two consecutive time instances. However, uncertainty poses different challenges to down stream functions and applications, including the need to develop new techniques for indexing, query processing, and data analysis for various uncertainty models. One way to overcome this is to try to infer the actual whereabouts of a moving objects between any two time instances with known locations. There are already several efforts to insert artificial points between each two consecutive trajectory points, with the promise that these points act as if the trajectory was collected in a very high sampling rate. This process has various names, e.g., trajectory interpolation [137, 270], trajectory completion [131], trajectory data cleaning [263], trajectory restoration [125], trajectory map matching [42], trajectory recovery [245], and trajectory imputation [77]. However, the large majority of such work rely on matching the trajectory points on the underlying road network, where the imputation becomes finding the road network shortest path between each two consecutive trajectory points. Unfortunately, this is not applicable to the case where the road network is unknown, untrusted, or inaccurate. Hence, more recent attempts try to do data-driven trajectory imputation without relying on the underlying road network [77, 81]. However, these techniques are either not scalable to city-scale trajectory datasets, or require dense historical data that derives its imputation process. There is an immense need to develop a scalable, accurate, and fine-grained imputation that almost mimics a continuous data stream of trajectory locations.

4 MOBILITY DATA ANALYTICS

Spatial data is special. Unlike non-spatial features, location attributes (e.g., longitude and latitude) rarely exhibit linear or other simple functional relationships to variables of interest. It rarely makes sense to model a variable of interest directly in relation to spatial attributes. Instead, it is distances that matter. According to Tobler's first rule of Geography, "everything is related to everything else, but closer things are more related than things that are far apart" [223]. For mobility data, proximity is further extended with time, i.e., objects that are close in space and time. In addition to this concept of spatiotemporal autocorrelation, what makes mobility data even more challenging to handle is that it is often observed from humans whose behavior can often be irrational and difficult to explain. As Nobel Prize laureate Murray Gell-Mann famously said, "Think how hard physics would be if particles could think" [173]. But unlike in physics, the "particles" of interest are often humans who can think. Data collection sensors have the capability to capture the spatiotemporal locations of the moving objects, but not their behavioral aspects. These difficulties require new paradigms, techniques, and algorithms to analyze and learn from the spatiotemporal data, and that can explain and predict the associated behavior. This section discusses current efforts and challenges of mobility data analysis.

4.1 Efforts in Mobility Data Analytics

Mobility data analytics has already gained momentum in research in the recent years. Dedicated workshops have existed in major conferences; including the ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data (BigSpatial) since 2011 [211], the Big Mobility Data Analytics (BMDA) workshop in EDBT since 2018 [178], and the ACM SIGSPATIAL International Workshop on Animal Movement Ecology and Human Mobility (HANIMOB)@SIGSPATIAL since 2021 [172]. Surveys on the status of research exist [20, 200].

Mobility data analytics encompasses various application domains and involves analyzing data from different – sources such as urban [267], maritime [61], aviation [59], animal movement [172], and indoor movement [115]. Among these different themes urban mobility stands out with a fairly large body of research including green routing [10], traffic anomaly detection [174], hot spot and hot path analysis [167], road traffic prediction [162], and travel time estimation [242]. Trajectories of moving objects have been used as means to create and continuously update the road network [160]. Public transport systems also collect ticketing data in the form of passenger check-ins, sometimes also associated with check-outs. This data has been shown very useful to transit planners in understanding passenger demand and movement patterns in daily operations as well as in the strategic long-term planning of the network [229]. Personal mobility of individuals is also a subject of analysis that includes analyses, e.g., activity recognition [50, 176], personalized routing [67], matching with ride-sharing services [19], and crowd-sourcing [179].

While a significant portion of research focuses on understanding and analyzing data through analytics, there are also important efforts dedicated to developing generic analysis tools for spatiotemporal data that are agnostic to the application domain. Efforts on generic methods for mobility data analysis include, among many others, trajectory clustering [246], trajectory similarity measures [226], outlier detection [102], transportation mode classification [40], spatiotemporal pattern detection [201], and trajectory completion [122]. However, and despite these many research efforts towards analyzing mobility data, there is lack of common data analysis tools and systems. The scientific software environment for mobility data analysis is rather fragmented. For example, [118] lists 58 packages in their review of R packages for movement and [93] reviews Python libraries for movement data analysis and visualization.

Recent years have seen a notable increase on research on deep learning for mobility data analysis [138, 252]. This brought an increased adoption of various paradigms and (adapted versions of) architectures used in other areas

where deep learning has brought improvements in tasks, e.g., clustering/classification [150], prediction [123] and recommendation [30], information propagation [276], etc. For example, Generative Adversarial Network (GAN) based architectures have been used recently to learn representations of trajectories and generate synthetic trajectories techniques [85]. Given the introduction of Transformers [235], transformed-based approaches have also been used for mobility modelling and trajectory prediction [256], given the sequential properties of mobility data. Other deep learning approaches such as contrastive learning [275] have also been exploited in mobile data settings, along with investigation of the impact/benefits of representation learning [87].

4.2 Challenges in Mobility Data Analysis

This section highlights open problems related to mobility data analysis, that needs consideration from the community.

Challenge 7. ML for Mobility Data. The state-of-the-art deep learning (DL) models, such as Transformers [235] were not developed initially for mobility data science in mind. They were derived from NLP and computer vision domains. The community needs to provide best-case practices for doing ML (and DL) for mobility data.

A major hurdle, and a research opportunity as well, is that existing ML and analytics tools, e.g., TensorFlow, or PyTorch, do not support location and mobility as base data types to reason about. So, even the basic analysis, such as clustering, classification, similarity, etc, need to be extended when mobility data is involved. These tasks, as well as higher-level analysis, can not be totally independent. Instead, common basic building blocks could have an impact on all or some of them. For example, exploring the effectiveness of embedding for mobility data analysis is a basic block that could impact different ML-based analysis tasks. This raises a challenge to build analysis primitives and common building blocks for applications that could shape a framework of ML-based mobility data analysis.

Another major hurdle is the robustness in data-driven mobility models. It is widely known that data-driven models (as in the case of ML or DL) are only as good as the data that it is used to be trained on. However, given the changes of mobility behaviors, such as the COVID-19 pandemic and the associated lockdowns, and environmental events and disasters, traditional ML based, and even recent DL based, methods are no longer robust. The models' performance deteriorate in unseen events, especially as new behaviors emerge and then persist. Recent effort includes the incorporation of 'contextual-awareness' and 'memory' in an enhanced event-aware spatiotemporal network [247] for predicting mobility in multiple modes of transportation including taxi, cycling, subway during the unprecedented events like COVID lockdowns, or snowstorms, as it emerged and up to 30 days post the event. However, more work to be done on modelling and understanding mobility behavior, that are robust to changes due to societal events.

Challenge 8. Progressing from Next Location Prediction to Movement Behavior Understanding. Due to the wide availability of aggregated check-in and foot-traffic data, many researchers focus on the problem of location prediction, e.g., [255]. Leveraging predictions such as "User X will visit Coffee Shop A next" or " 32 ± 4 users will visit Coffee Shop A in the next hour" has some direct applications. It could be useful for providing information about parking "parking at location X appears to be a problem today, so consider ...", for battery charging opportunities, or for providing information about collective transportation status "Metro station X that you are expected to visit is closed for repairs, so instead ...". One could provide a new transportation schedule and departure time in response to problems at an anticipated future location of a user, just like airlines at times update your itinerary in case of issues. Earlier work has been based on data mining techniques to detect periodic behavior, e.g., [36, 76, 117]. Beyond predicting locations, if we understand the underlying behavior, at the individual-, group-, or population-scales, that leads to these predictions, we could understand *why* one coffee shop chain has increasing visitor rates (e.g., due to a movement

towards organic coffee sold by the coffee shop). Through inferring from the data about such behaviors, only then we can take corresponding actions not only to predict locations, but also to prescribe actions (e.g., offering more organic coffee) to improve visitor rates. This understanding of (human) behavior will broadly affect applications using mobility data. Traditional spatiotemporal data science allows for predictive analytics to predict the future. In contrast, mobility data science enables prescriptive analytics by understanding the underlying human behavior to devise actions and policies that aim to achieve desirable targets.

An open problem for understanding mobility behaviour data is the lack of labels or human annotation to provide insights on the actual observations. There are several other tricks that have been proposed, including cross-domain data fusion as well as developing interpretability mechanisms for machine learning or deep learning models. When geographical information is fused together with contextual features and social behaviours, not only location prediction can be improved, but also insights can be provided about the underlying visitor behavior [255], even if no human-labelled data are provided about the mobility behaviors.

Therefore, explainability of AI and machine learning models that have underpinned many of such predictive behavior models remain an open challenge, especially since deep learning models are black boxes. One such approach for deep-learning-based models is disentangled representation learning, and a recent work [268] shows that the disentanglement of latent spatiotemporal factors can assist the explainability of how the underlying latent factors learned by deep learning models are correlated. It can also be used for dimensionality reduction, and assist in few-shot learning cases.

Challenge 9. Visual Analytics. Visualization and exploratory analysis of mobility data has long been a hot topic in visual analytics [15]. More recently, the trend turned to combining visualization with modeling and simulation to support decision making [124]. This kind of research is by necessity application-oriented, while much less is done on developing more general ideas and approaches.

One general research problem that has only been slightly touched in visual analytics but not systematically addressed is human involvement in real-time analysis of big mobility data. Is it possible to define realistic scenarios for involving human intelligence in big data analytics taking into account the cognitive limitations of human analysts with regard to the amount of information that can be perceived, speed of processing, and time required for analytical reasoning and contributing to the analysis process? Also how to combine computational methods of analysis, such as ML, with human expert knowledge and reasoning? The involvement of human intelligence is limited to thoughtful data preparation, feature selection, parameter setting, and so on. It would be great to find ways to make more direct and effective use of human-possessed concepts and, particularly, knowledge of causal relationships. Hence, a grand research challenge for visual mobility analytics is to develop approaches to understanding and modeling mobility behaviors from low-level movement data, such as trajectories of moving entities.

The following research problem is how to analyze behaviors after they have been extracted from elementary movement data and represented by appropriate data structures. A conceptual framework should be developed to enable defining the types of conceivable patterns of movement behavior. This will provide orientation for developing visualization techniques facilitating visual discovery of behavioral patterns, as well as algorithmic methods for detection of specified types of patterns. These techniques and methods should be incorporated in systems and workflows for analyzing the contexts in which various patterns take place and developing models for describing and predicting mobility behaviors depending on the context.

5 MOBILITY DATA MANAGEMENT INFRASTRUCTURE

Classical data management systems have been designed for generic data types, where spatial and temporal data can be supported as new additional types. Yet, the core functionality of the data management engine does not acknowledge the spatial and temporal properties of mobility data. For example, mobility data calls for storing and querying locations of objects that evolve over time. The evolution can be in the location, the extent, and/or the properties of the object. The evolution can happen in discrete steps, e.g., check-ins, or in a continuous form. Thus, it is desired that the data management platform is able to represent the history, the current location, and possibly the near future of the moving object. Another example is classical index structures that are built with the assumption that the read workload is significantly higher than the write workload, and hence the index structure does not change often. Meanwhile, mobility data exhibits a different workload where the write workload (e.g., object location update) is significantly higher than the read workload, which makes all classical index structures simply not applicable to mobility data. A third example is that simple queries over mobility data, e.g., nearest neighbor search can be supported by classical data management systems by finding the distance between the user location and all other objects, sorting all objects based on that distance, and getting the closest one. This cumbersome approach is mainly due to the lack of having a specialized nearest-neighbor operator. Should we have one, that operator can seamlessly integrate with the query executor and optimize of a data management engine to efficiently support a pretty important query in most data mobility applications. Finally, a last example is that classical methods for scaling up data management in distributed environments rely on data distribution, mostly based on the data keys. This does not work well in scaling up mobility data as it is always desired to distribute mobility data in a way that spatially and temporally nearby objects are grouped together in the same cluster or computing node. This sections discusses current efforts and challenges of mobility data management.

5.1 Efforts in Mobility Data Management

There has already been extensive research in all layers of mobility data management infrastructure. On the data modeling aspect, early models based on the constraint databases model aim to support simple moving objects (i.e., points), e.g., [94]. More complex data types (e.g., moving regions) have been supported by later models based on abstract data types, e.g., [101] that is still being used in recent systems, e.g., [278]. More recent efforts have been introduced to capture the semantics of trajectories of moving objects. Other models were also proposed to capture specialized modes of movement, including indoor environments, e.g., [114], network constrained, e.g., [99], fuzzy trajectories, e.g., [227], and detecting periodic moving patterns, e.g., [33, 36, 76, 117]. In terms of indexing, tens of index structures have been proposed to support efficient indexing, storage, and retrieval for spatiotemporal data as either historical data, current locations, or continuously updated locations, e.g., [144, 147, 155, 164]. This forms the infrastructure support for various spatiotemporal query processing techniques for various query operators over moving objects, including spatiotemporal range queries [157], spatiotemporal nearest-neighbor queries, e.g., [11–13, 216, 254], reverse nearest neighbor queries [35], skyline queries [109], and scalable spatial and spatiotemporal joins, e.g., [249, 253].

In terms of academic full-fledged systems, the SECONDO system has been introduced in the early 2000 as a comprehensive testbed for distributed moving object databases covering all aspects of data modeling, indexing, and querying [98]. More recently, MobilityDB, implemented on PostGIS, has been introduced as a scalable system with a wider functionality on moving object databases [230, 278]. In terms of Big Data systems, ST-Hadoop [8], SUMMIT [7] and HadoopTrajectory [22] systems extend the Hadoop system to support spatial-temporal data, and trajectories, respectively, while other systems, e.g., [66, 145, 146], extend the Twitter Storm distributed data streaming system to

support streamed location data. TrajSpark [266], Dita [209], and TrajMesa [128] extend the Spark system to support various index structures and query operations over trajectory data. SharkDB [244] extends in-memory column-oriented storage engines to support trajectories. In the open-source community and in industry, PostGIS [182] supports very basic trajectory functions, Oracle spatial supports streaming point data to capture real time mobility [170], while Microsoft Azure [25] supports storing trajectory data in Azure table and utilizing Azure Redis for indexing. Distributed-MobilityDB [23] integrates in the one hand the trajectory data management of MobilityDB with a distributed PostgreSQL database to provide a distributed moving object database.

5.2 Challenges in Mobility Data Management Infrastructure

Though there is already a lot of work in various components of mobility data management infrastructure, there is an apparent lack of integrated systems that offer comprehensive functionality to end users, encapsulated in full-fledged systems that support mobility data science. Hence, the challenges in this section mainly focus on the system building aspect.

Challenge 10. Building Systems with Mobility Data in Mind. Location data has almost always been supported in data systems as an afterthought problem. Many systems, e.g., Postgres, Storm, Spark, and Hadoop, have not been originally designed with location data support in mind. What typically happens is that spatial data types get augmented into tuple-oriented systems to support the location data type. For example, a restaurant tuple that describes various attributes of a restaurant is augmented with the latitude and longitude of the location attribute of the restaurant to support location services. Spatial indexes are provided to speedup the access to these attributes, and some accompanying spatial operators are provided to operate on the location attributes to provide location services, e.g., range or k-nearest-neighbor searches. While this approach works to some extent, systems coming out of such approach end up with sub par performance for spatial data, and hence for mobility data. Given the myriad of applications that rely on mobility data, it is important that systems are extended with native support for locations and mobility data. So, mobility data types and operations should be integrated in the core of these systems, and not to be considered as an afterthought problem. This can go through all kinds of systems, starting from database management systems that need to be spatially- and temporally-aware to support mobility data to scalable big data and NoSQL systems, where injecting spatial- and temporal-awareness into their core functionality will inherit their scalability to support scalable mobility data science.

Challenge 11. Location Data as First-class Citizens. Having locations as core of mobility data calls for treating location data as a first-class citizen in a location data system that at the same time can be extended to support other data types [16]. These location data systems can be presented as Location+X systems, e.g., as in [16], where the data types \hat{X} can be keywords (e.g., to support spatial-keywords and tweets), graphs (e.g., to support road-network data), relational data (e.g., to support descriptions of spatial data objects), click streams (e.g., to support check-in data), document data (e.g., to support points of interest and documents that describe them), annotated trajectories (e.g., location + time + textual annotations), among others. In many location services, more than one data type X may need to be supported, e.g., a graph data type combined with a document or keyword data types, which calls for a multi-model-like data system. This gives rise to an eco-system where location is at the core with some form of an extensible multi-model data system that supports the multitude of data types \hat{X} . However, current multi-model data system technology is lacking in several aspects. First, they do not support data streaming that is a cornerstone in mobility data due to the online streamed locations of moving objects. Second, we do not want to fall into the trap of adopting existing multi-model technologies that may affect location being a first-class citizen. However, the need for

supporting multi-models in one seamlessly integrated location+X system remains a necessity. In addition to supporting location data via a native location+X engine, an ecosystem for mobility data would also include many important utilities to facilitate a broad spectrum of location service applications. From the input data side, to help navigate the vast amounts of available location datasets, and discover the right data sets for a given task, a location dataset lake infrastructure and location dataset discovery, cleaning, and integration facilities are needed. From the presentation side, a comprehensive visualization suite is envisioned to support visualizations for combinations of spatial and temporal data analytics on top of location data.

Challenge 12. Streaming, Batch, and Hybrid Workloads. Motivated by the application needs, mobility data management need to support both batch and real-time data through all systems layers from digesting the data to analyzing and visualizing it. For example, a common requirement is to visualize the positions of a fleet of vehicles in real time, which only requires access to the most recent positions of the vehicles. Yet, at the same time, there is a need to perform batch analytics on the full trajectory of these vehicles (e.g., to assess whether the trajectories exhibit some unexpected behavior). Generally speaking, the need to have both real-time and historical data has led to the development of the data warehouse domain, where operational databases cover the real-time Online Transaction processing (OLTP) while data warehouses cover the historical Online Analytical Processing (OLAP). Since having two different systems for the two kinds of workloads is very costly, a new approach referred to as Hybrid Transactional and Analytical Processing (HTAP) has been recently proposed. However, mobility data exhibits pretty different workloads from other data, where streaming data is kind of dominant in terms of objects continuously streaming their new locations. Meanwhile, historical data is not of less importance and are continuously appended. While some efforts have been spend in the direction of write-optimized indexing for location data, e.g., as in [213], more research efforts need to be spent to adopt the concepts behind HTAP systems to support the nature of mobility data.

6 MOBILITY DATA PRIVACY

As we discussed in Challenge 1, mobility data privacy is a core problem in the mobility data science pipeline. Studies have shown that location data could reveal sensitive personal information such as home and workplace, religious and sexual inclinations [185]. As localization technology advances and extremely fine-grained location tracking is being enabled, it may even reveal products of interest in the stores we have visited, doctors we saw at a hospital, book shelves of interest in a library, artifacts observed in a museum, and generally anything that might publicize our preferences, beliefs and habits. Recent survey has shown that 78% smartphone users among 180 participants believe that Apps accessing their location pose privacy threats [47].

While there are many privacy-preserving data collection and data analysis techniques developed for personal data, mobility data introduces unique challenges due to 1) spatiotemporal correlations in the mobility data which often results in increased privacy cost due to privacy composition for correlated data or downgraded utility for downstream applications, 2) complex location semantics (e.g., corresponding POIs of locations) and mobility behaviors (e.g., regular vs. one-time visit of a location) which existing privacy definitions may not be able to capture, and 3) diverse and emerging application scenarios such as contact tracing using mobility data for which existing privacy algorithms designed for aggregate data analytics are not suitable. In this section, we briefly review existing privacy notions and techniques developed for location and mobility data and discuss several open challenges.

6.1 Efforts in Mobility Data Privacy

We categorize existing techniques in mobility data privacy into two main settings corresponding to our data pipeline: 1) local setting (data collection stage), and 2) central setting (data analysis stage). In the local setting, the mobility service provider that collects mobility data is assumed to be untrusted, hence each mobile user or entity can apply privacy-preserving mechanisms before the data is collected by the service provider. In the central or global setting, the mobility service provider is assumed to be trusted and collects the raw mobility data. The provider can apply privacy-preserving mechanisms for statistical analysis, and share aggregated data, machine learning models trained from the data, or synthetic data mimicking the original data to untrusted third parties.

Local Setting. In recent years, local differential privacy (LDP), the local variant of differential privacy, [63, 95] has become the de facto standard for preserving privacy at data collection stage. Each user can perturb her raw data using an LDP mechanism before uploading it to an untrusted server. Most existing mechanisms are designed to ensure utility for aggregate queries or analytics (e.g., frequency or density estimation) and requires the aggregation of the perturbed values from a large group of users, while the individual perturbed value may not provide much utility. Several works applied existing LDP schemes to location data but the utility is poor [120, 269]. Other works relaxed LDP to personalized LDP [52]. Recent works developed improved LDP mechanisms for location data with better utility [241].

In addition to support aggregate data analytics, location based services (LBS) including range queries, spatial crowdsourcing, and the emerging contact tracing for pandemic control, require the precision of the perturbed locations themselves. Geo-indistinguishability (GeoInd) [14] relaxes LDP for location data which requires the locations to be indistinguishable only within a radius and the indistinguishability is scaled by their distances, providing better privacy utility tradeoff for LBS. Later works extended GeoInd to account for temporal correlations between consecutive locations of mobile users [251] and protection of customizable spatiotemporal activities instead of raw locations or trajectories [43]. Other works applied the GeoInd mechanisms and variants for privacy-enhanced spatial crowdsourcing and contact tracing [64, 222]. Besides statistical privacy techniques, Private Information Retrieval (PIR) and secure multiparty computation (MPC) techniques have also been developed to allow LBS queries such as range queries and contact tracing without revealing individual locations [6, 56, 88, 188], but are generally more computationally expensive and need to be designed for each different query.

Global Setting. Many works have applied differential privacy (DP) for computing and publishing aggregate mobility data. Compared to DP algorithms for tabular data, they typically exploit the hierarchical structure of locations and sequential patterns of trajectories to improve utility [2, 49, 151, 186, 206]. Some works also utilized the DP aggregates for task assignment in spatial crowdsourcing [221]. In practice, mobility data providers have started sharing aggregated mobility datasets with DP, esp. in response to the pandemic, such as Meta’s population density maps and Movement Range maps, Google’s COVID-19 Community Mobility Reports, and SafeGraph’s Patterns [24]. Other works have applied DP for training machine learning models using mobility data, for example, for location prediction [5]. Another line of work attempts to generate synthetic trajectories or mobility data based on raw trajectories with formal DP guarantees [104, 243]. From the privacy attack side, recent works demonstrated the possibility of membership inference attacks on aggregate location data and linking attacks, and the defense power of DP against some of these attacks, reinforcing the need for ensuring rigorous privacy even for seemingly anonymous aggregate mobility data and machine learning models trained from mobility data [116, 184].

6.2 Challenges in Mobility Data Privacy

This section highlights open problems related to mobility data privacy, that needs consideration from the community.

Challenge 12. Threat Models and Privacy Definitions. The first challenge for mobility data privacy is the need to understand the threat models and adopt or define proper criteria by which to enforce privacy. We need to define first what needs to be protected (i.e., the sensitive information). This may vary for different mobile users and applications. It may be the exact location coordinates of a user at a given time (most existing efforts focus on this). It may also be the association of a user with a sensitive place, co-location of two users (while it's okay for the users to reveal the exact location coordinates), or spatiotemporal activities of a user (e.g., stay at a place, or a trajectory). When defining privacy models and designing subsequent privacy mechanisms, there will (almost always) be attacks based on side channel information exploitation. While privacy notions like DP typically assumes the worst case which also means sacrificed utility, relaxed versions may be needed given specific threat models to enhance the privacy and utility tradeoff.

Besides developing rigorous privacy enhancing mechanisms, it is equally important to understand the privacy risks and the empirical defense power of the privacy enhancing technology (PETs). While there have been some work on privacy attacks on aggregate mobility data [183], more work is needed to understand what sensitive information may be revealed and reconstructed from mobility data based models, e.g., if membership inference attacks or feature reconstruction attacks [82, 214] can be carried out, and potentially build benchmark attacks which can be used to audit the privacy risk of mobility data science systems and privacy mechanisms.

Challenge 13. Privacy and Utility Tradeoff and Other Factors. When designing privacy mechanisms for mobility data collection and analysis, it is important to consider the utility of the privacy protected data for the downstream applications. For LBS (as typical in the local setting), the utility needs to be measured by the precision or accuracy of range queries for POI search, or contact detection for contact tracing (instead of how accurate the perturbed location is from the original location for which most algorithms following geoInd are focused on). Hybrid methods that combine DP and cryptographic techniques may be needed esp. for critical applications like contact tracing and public health [56]. For aggregate data analytics and machine learning applications using mobility data (in both local and global setting), the utility need to be measured by the accuracy of the statistics (e.g., frequency or density estimation for which most existing work focus on), the trained model, or the fidelity of the synthetic data. As a result, the algorithms need to be designed to optimize the corresponding utility and many remain an open challenge. For example, existing methods for DP trajectory synthesization are mainly based on statistical models or low-order Markov models and perform well on some utility metrics [104, 243]. While there are more powerful generative adversarial network (GAN) based models or diffusion models for generating more realistic synthetic trajectories [138, 277], ensuring formal DP for these models would result in deteriorated utility due to the complexity of the models. Designing methods for optimal privacy utility tradeoff remains an open challenge.

In addition to the privacy and utility tradeoff, privacy enhancing technology may exacerbate bias in the data or learning algorithms. Mobility data may have inherent bias as we discussed in Challenge 2. Data analysis algorithms may also have unfair performance for groups that are underrepresented in training data. It has been demonstrated that learning with DP could exacerbate such unfairness, i.e. underrepresented groups suffer from worse privacy/utility trade-offs [21]. Research is needed to understand such impact on mobility data and design privacy algorithms to optimize privacy utility tradeoff while ensuring the fairness.

Challenge 14. Explainability and Societal Education. Another important challenge of mobility data privacy is to improve the explainability of privacy definitions and mechanisms and communicate them to the stakeholders including mobile users (data contributors), mobility service providers, and data analysts. This is a general challenge for privacy enhancing technology, but more so for mobility data given the complex semantics of location information and diverse applications as we mentioned. DP-compliant algorithms and location privacy models (such as Geo-Ind) as described earlier use privacy parameters to control the trade-off between privacy guarantee and the utility of the private outputs. However, there is a significant gap between the theory and practice of DP: we lack principles and guidelines for choosing privacy parameters when collecting or processing mobility data using DP techniques in the real world. While the technology companies have employed DP in releasing the mobility datasets as we discussed earlier, the choice of the privacy parameter and the associated noise and uncertainty are often not precisely specified or uniform across companies. This makes it difficult for the downstream applications to quantify the uncertainty of the analysis result.

The parameter ϵ of DP is mathematically defined but not well-aligned with the stakeholders' interests. Even for the same ϵ , the privacy guarantees could be different based on the different variants of DP and algorithms at hand. In addition, the ϵ is not always linked to a specific privacy risk for the users (such as "the probability that an attacker can correctly infer my data") or a precise utility level for data analysts (such as "the accuracy of the DP-ML model"). To promote the adoption of mobility data privacy technology such as those based on DP, we should establish principles, design guidelines, and provide tools for explaining DP's protection and limitation from stakeholders' practical interests. For example, we can help data contributors understand the privacy risk (such as membership inference attacks or reconstruction attacks) under different privacy parameters given a concrete DP algorithm; we can also design efficient methods to visualize how data analyzers' utility metrics (such as MSE or model accuracy) may change along with different privacy parameters for specific mobility applications.

7 MOBILITY DATA SCIENCE APPLICATIONS

Mobility data science used to be limited in the domain of transportation but recent technological inventions have created an abundance of mobility data, resulting in applications in many other domains of interest for society. Such applications leverage mobility data to understand, explain, and predict where moving entities such as humans, animals, or infectious diseases go, why they go where they go, and where they will go next. This section outlines broad applications of mobility data science to illustrate the recent landscape of mobility data science.

7.1 Traffic

Traffic is a problem of global scale, as recognized by transportation science over a decade ago. Drivers in the United States spend 6.9 billion of driving-hours stuck in traffic and waste more than 11 billion liters of fuel per year according to INRIX [113]. Measured per-capita, people in Russia and Thailand spend even more time in traffic, while Brazil, South Africa, the UK, and Germany are only slightly behind the United States. Leveraging mobility data science and understanding the underlying behavior of human participants concomitantly with different transportation modes, can enable more effective solutions to multiple problems at the heart of improving traffic management. Two main lines of research focus on: (1) traffic monitoring at an aggregate level, e.g., to help city administration, and (2) services that road users are getting. Existing work towards traffic monitoring include monitoring congestion [129], assessing the safety of roads and intersections [143], traffic prediction [132], evacuation routing [265], optimizing the public transportation schedules [193]. Efforts on the services provided to road users include routing queries that balance

the traffic across roads [69], help find drivers finding nearest facilities [121], personalized routing [130], eco-routing for minimizing greenhouse emissions [134], and enabling multi-modal trip planning [225]. But there are many open opportunities and challenges in using mobility data to improve traffic conditions. One example is devising accurate models for the dynamic scheduling of public transportation. Another example is the context-aware optimization of traffic signals – e.g., incorporating the impact of additional flux of pedestrians in bus/train stations, to minimize the stop-and-go impacts for vehicles. A challenge of using mobility data science in the transportation domain is monitoring and reduction of emissions. Being able to quantify emissions (e.g., from transportation) is essential to accountability and reduction of emissions. Using data on emissions collected from in-situ sensors but also sensed remotely through earth observation (satellite) data will allow us to better understand the effects of e-mobility, better collective transportation, and infrastructure improvements.

7.2 Urban Areas

In 2018, 55% of the world’s population (4.2 billion people) resided in urban areas, and this proportion is projected to increase to 68% by 2050 [232]. Urban areas are a focal point for mobility application as they introduce a variety of mobility modalities such as electrical vehicles [236] and bicycles and scooters with respective sharing programs [133]. But by understanding how, where, and why people move in cities, outer suburban and regional areas, the demand for infrastructure and energy can be better understood [272]. Improving this understanding helps reduce urban inequalities in cities [166] such as access to high quality food [238] and healthcare [96]. Mobility data also helps improve urban safety by improving crime prediction [83] and helping to recommend safe routes [205].

A specific Urban mobility data science supports urban areas is through data-driven map construction [3] and updating of existing maps to account for blocked or new road segments [48] which is paramount in autonomous driving applications [141].

The real-time monitoring of urban mobility could result in *situational awareness*, initially a term coined in defense applications, involving *perception* of the environmental states using the surrounding data, *comprehension* of the ingested data to understand the emerging situations, and *projection* of future states and/or events that require predictive analytics. Mobility data provides critical components and insights into situational awareness in cities. When achieved, this applies not only to enabling robust critical infrastructures in cities but also to protecting them from harm, e.g., forest fires, earthquakes, and terrorist attacks. Many researchers use mobility data as input to enable situational awareness in cities as well as in airports [210].

7.3 Health Informatics

The spread of infectious diseases is a highly complex spatiotemporal process that is strongly tied to human mobility [107] and human behavior [75]. Many recent works have used human mobility data for data-driven epidemic forecasting as surveyed in [197]. A specific example of leveraging mobility data for public health is contact tracing, which refers to the process of tracking persons who may have come into spatial contact with an infected person, and subsequently collecting further information about these contacts [152]. The feature-rich interaction, processing and localization/communication modalities of smartphone devices, have brought these to battle on the technological forefront and have curbed the fast spread of pandemics, like COVID-19. To this date, the community has proposed a wide range of contact tracing approaches, including opportunistic [187] and participatory approaches [65] approaches as well as privacy-sensitive [262], decentralized [228], proximity-based (e.g., BLE, sound) [189], and location-based approaches (e.g., Wi-Fi, GPS) [65] approaches. However, a wide range of challenges remain unanswered, including methodologies to

improve the penetration and adoption rates, alleviate privacy or expectation skepticism [32], ubiquitous availability on low-end terminals as well as technological/psychological adoption barriers [31], achieving cross-country interoperability with standard formations beyond recommendations, scalability/reliability and accuracy verification of engaged spatial technologies as well as lessons about effectiveness from real large-scale deployments.

Another specific health application for mobility data is elderly health monitoring. GPS-enabled smart-watch technology can be used to monitor the movement of elderly users [217]. In particular, if the monitored user is showing early signs of dementia, her/his trajectories could show an abrupt change from her/his movement history [224]. For instance, a user who normally walks in a park then goes to a restaurant is found to only stay in the park for a substantial amount of time. Indoor sensors installed in the room can also be used to track whether an elderly person or a patient falls from the bed. Trajectory outlier analysis methods, together with gerontology knowledge, can be very useful for this kind of applications.

7.4 Indoor Environments

Indoor mobility data management has been described as a new frontier in data management [115]. But in addition to data management, large-scale indoor localization data also raises challenges in data collection, data analysis and data privacy. Indoor data collection is an open research problem due to the non-existence of the indoor equivalent of GPS: a system that can provide the user location in any building worldwide. This is particularly important in applications related to emergency management and infectious disease contact tracing. Systems have been developed over the years to address this problem based on different data sources including WiFi signal strength and time of arrival [257], cellular signal [196], Ultra-wideband [9], ultrasonic [111], magnetic tracking [215], inertial sensors [103], among others. These novel data sources enable new applications in indoor navigation, contact tracing, indoor analytics, and evacuation management.

Indoor data analytics allows to improve understanding of indoor behavior which has multiple benefits and applications, including for crowd management [4], retail and POI recommendation systems [191], and for optimizing energy use and improving sustainability in the long term [202]. For example, by utilizing WiFi logs, Ren et al. [190] find strong correlations between behaviors and user demography (e.g., age, gender and visitor types), indicating that both indoor mobility behavior, in conjunction with online behavior, can be used to predict the underlying demography of the visitors.

Occupancy behaviors are also highly linked with building management systems and controls [45]. By having a more accurate energy use estimation using indoor spatial and mobility data, in addition to historical energy consumption data, the performance of the buildings can be better optimized, towards achieving a more sustainable operations [72]. The responsible use of mobility behavior analytics, including indoor and outdoor mobility behaviors, strongly points towards the increased capacity for improving sustainable operations of buildings [202], enabling net zero goals to be achieved.

7.5 Marine Transportation

According to UNCTAD, Over 80% of the volume of international trade in goods is carried by sea, and the percentage is even higher for most developing countries [231]. Estimates say that the global shipping activity emits 3% of the global emissions worldwide in 2022 [110]. These significant numbers, as well as the availability of large-scale ship trajectory data obtained from the automatic identification system (AIS) [18] motivated a lot of research efforts on mobility data

analysis for maritime. The stakeholders who seek benefit of such analyses include the maritime authorities, environment officers, ship owners, port and canal managers, and the transport and logistic sector.

One major challenge is to ensure the safety at sea, which splits down to the technical challenges of identifying positional anomalies [195], locating dark vessels (vessels that switch off their AIS devices) [148], and cleaning location and identity spoofing [74]. Additionally, an essential aspect is the detection of fishing activities to ensure the sustainable fishing practices [58]. Since vessels do not have fixed routes in the sea, research has also investigated the density of ship routes [250].

Multi-criteria routing using multiple optimization criteria including estimated time of arrival, fuel consumption, safety, and comfort has been increasingly recognised as an important path planning problem [105]. An optimization of ship routes could effectively lead to significant reductions of GHG emissions and contribute to the actions against anthropogenic global warming. The influence of ocean currents, waves, and wind on the course and speed of ships have been known for centuries. Used optimally, ocean currents lead to more efficient paths between two given ports. Ship route computation approaches that exploit the potentials of wind, wave and weather models aiming at minimize fuel consumption have been addressed by the marine science, maritime engineering and transportation community [78].

Since green mobility is currently gaining huge attention, CO₂ emission aware ship routing is expecting to get an enormous impact in economy, politics and society and provides very promising opportunities for the spatial and spatiotemporal database and mobility community. Marine transportation becomes particularly important in the scope of climate change (e.g., the advent of hydrogen/battery/fossil/atom hybrid vessels) as well as digitization for new infrastructure-free localization technologies on-board.

7.6 Social Connections

Location-based social networks (LBSNs) bridge the gap between the physical world and online social networking services [271]. LBSN data capture both human mobility (in the form of check-ins to discrete points of interest) and a social network between individual humans. Combining mobility data and social networks, LBSN data finds many applications. A first application found in the literature was on modeling and describing human mobility patterns (e.g., [55, 168]), analyzing these patterns (e.g., [54]), and explaining why individual user choose locations and how social ties affect this choice (e.g., [239]). Another application is that of location recommendation, which leverages check-ins of users and their ratings in the user-location network to recommend new locations to users [26]. A closely related application area is that of location prediction (e.g., [53]), which predicts the future check-ins of users. Another active research field in LBSNs analysis is friend recommendation or social link prediction (e.g., [203]), which suggests new friends to users based on similar interests at similar locations, while also having similar social connections. Other research topics concerning LBSNs include efficient query processing (e.g., [17]) finding user communities (e.g., [259]), and estimating the social influence of users (e.g., [248]).

This plethora of applications and research shows how mobility data in connection with social network data can be used to understand the social fabric that ties us together. A potential future application is using human mobility data to reinforce this social fabric by recommending social events and meetings to groups of people to help people find new friends, collaborators, sports mates, teachers, mentors, and family members.

8 CONCLUSIONS

This paper presented the current state of mobility data science pipeline in addressing the specific challenges of mobility data. A main question that this paper answered is how mobility data science is different from data science. The

space and time dimensions in mobility data call for different methods of data acquisition, management, analysis, and privacy preservation which are not addressed by the common data science tools. Accordingly we surveyed the main problems that are currently being researched, we identified major research questions for the coming years, and described applications that lead to broader impacts of mobility data science. Co-authored by a diversity of academics and industry professionals, this paper also conferred a community effort to sketch the boundary of mobility data science as an interdisciplinary field and bring together a dedicated research community around the identified research challenges.

9 ACKNOWLEDGEMENTS

Walid G. Aref acknowledges the support of the National Science Foundation under Grant Number IIS-1910216.

REFERENCES

- [1] ACM SIGSPATIAL CUP 2017. <http://sigspatial2017.sigspatial.org/giscup2017/download>.
- [2] G. Acs and C. Castelluccia. A case study: Privacy preserving release of spatio-temporal density in paris. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1679–1688, 2014.
- [3] M. Ahmed, S. Karagiorgou, D. Pfoser, C. Wenk, M. Ahmed, S. Karagiorgou, D. Pfoser, and C. Wenk. *Map construction algorithms*. Springer, 2015.
- [4] T. Ahmed, T. B. Pedersen, and H. Lu. Finding dense locations in indoor tracking data. In *2014 IEEE 15th International Conference on Mobile Data Management*, volume 1, pages 189–194. IEEE, 2014.
- [5] R. Ahuja, G. Ghinita, and C. Shahabi. Differentially-private next-location prediction with neural networks. In A. Bonifati, Y. Zhou, M. A. V. Salles, A. Böhm, D. Olteanu, G. H. L. Fletcher, A. Khan, and B. Yang, editors, *Proceedings of the 23rd International Conference on Extending Database Technology, EDBT 2020, Copenhagen, Denmark, March 30 - April 02, 2020*, pages 121–132. OpenProceedings.org, 2020.
- [6] W. Al Amiri, M. Baza, K. Banawan, M. Mahmoud, W. Alasmay, and K. Akkaya. Privacy-preserving smart parking system using blockchain and private information retrieval. In *2019 international conference on smart applications, communications and networking (SmartNets)*, pages 1–6. IEEE, 2019.
- [7] L. Alarabi and M. F. Mokbel. A demonstration of summit: a scalable data management framework for massive trajectory. In *IEEE International Conference on Mobile Data Management, MDM*, pages 226–227, June 2020.
- [8] L. Alarabi, M. F. Mokbel, and M. Musleh. ST-Hadoop: a MapReduce Framework for Spatio-temporal Data. *GeoInformatica*, 22(4):785–813, 2018.
- [9] A. Alarifi, A. Al-Salman, M. Alsaleh, A. Alnafessah, S. Al-Hadhrani, M. A. Al-Ammar, and H. S. Al-Khalifa. Ultra wideband indoor positioning technologies: Analysis and recent advances. *Sensors*, 16(5):707, 2016.
- [10] A. M. Almeida, J. L. Leite, J. A. Macedo, and J. C. Machado. Gps2gr: Optimized urban green routes based on GPS trajectories. In *Proceedings of the 8th ACM SIGSPATIAL Workshop on GeoStreaming*, pages 39–48, 2017.
- [11] A. M. Aly, W. G. Aref, and M. Ouzzani. Spatial queries with two knn predicates. *Proc. VLDB Endow.*, 5(11):1100–1111, 2012.
- [12] A. M. Aly, W. G. Aref, and M. Ouzzani. Cost estimation of spatial k-nearest-neighbor operators. In *Proceedings of the 18th International Conference on Extending Database Technology, EDBT 2015, Brussels, Belgium, March 23-27, 2015*, pages 457–468. OpenProceedings.org, 2015.
- [13] A. M. Aly, W. G. Aref, and M. Ouzzani. Spatial queries with k-nearest-neighbor and relational predicates. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems, Bellevue, WA, USA, November 3-6, 2015*, pages 28:1–28:10. ACM, 2015.
- [14] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914, 2013.
- [15] G. Andrienko, N. Andrienko, P. Bak, D. Keim, and S. Wrobel. *Visual Analytics of Movement*. Springer, 2013.
- [16] W. G. Aref, Y. Rayhan, L. Zhou, and A. Daghistani. ILX: intelligent "location+x" data systems (vision paper). *CoRR*, abs/2206.09520, 2022.
- [17] N. Armentzoglou, S. Papadopoulos, and D. Papadias. A general framework for geo-social query processing. *Proc. of the VLDB Endowment*, 6(10):913–924, 2013.
- [18] A. Artikis and D. Zissis. *Guide to Maritime Informatics*. Springer, 2021.
- [19] M. Asghari, D. Deng, C. Shahabi, U. Demiryurek, and Y. Li. Price-aware real-time ride-sharing at scale: an auction-based approach. In *Proceedings of the 24th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 1–10, 2016.
- [20] G. Atluri, A. Karpatne, and V. Kumar. Spatio-temporal data mining: A survey of problems and methods. *ACM Comput. Surv.*, 51(4), aug 2018.
- [21] E. Bagdasaryan, O. Poursaeed, and V. Shmatikov. Differential privacy has disparate impact on model accuracy. *Advances in neural information processing systems*, 32, 2019.
- [22] M. Bakli, M. Sakr, and T. H. A. Soliman. Hadooptrajectory: a hadoop spatiotemporal data processing extension. *Journal of Geographical Systems*, pages 1–25, 2019.
- [23] M. Bakli, M. Sakr, and E. Zimányi. Distributed spatiotemporal trajectory query processing in sql. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, pages 87–98, 2020.
- [24] S. Balsari, C. Buckee, J. Chan, and A. Schroeder. The user of human mobility data in public health emergencies. In *CrisisReady*, 2022.

- [25] J. Bao, R. Li, X. Yi, and Y. Zheng. Managing massive trajectories on the cloud. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPACIAL '16, New York, NY, USA, 2016. Association for Computing Machinery.
- [26] J. Bao, Y. Zheng, D. Wilkie, and M. Mokbel. Recommendations in location-based social networks: a survey. *Geoinformatica*, 19(3):525–565, 2015.
- [27] H. Barbosa, M. Barthelemy, G. Ghoshal, C. R. James, M. Lenormand, T. Louail, R. Menezes, J. J. Ramasco, F. Simini, and M. Tomasini. Human mobility: Models and applications. *Physics Reports*, 734:1–74, 2018. Human mobility: Models and applications.
- [28] E. Barmounakis and N. Geroliminis. On the New Era of Urban Traffic Monitoring with Massive Drone Data: The pNEUMA Large-scale Field Experiment. *Transportation Research Part C: Emerging Technologies*, 111:50–71, 2020.
- [29] F. Bastani, S. He, S. Abbar, M. Alizadeh, H. Balakrishnan, S. Chawla, S. Madden, and D. J. DeWitt. RoadTracer: Automatic Extraction of Road Networks From Aerial Images. In *CVPR*, pages 4720–4728, Salt Lake City, UT, USA, June 2018.
- [30] Z. Batmaz, A. Yürekli, A. Bilge, and C. Kaleli. A review on deep learning for recommender systems: challenges and remedies. *Artif. Intell. Rev.*, 52(1):1–37, 2019.
- [31] L. Bedogni, F. Montori, and F. Salim. Location contact tracing: Penetration, privacy, position, and performance. *Digital Government: Research and Practice*, 3(3):1–13, 2022.
- [32] L. Bedogni, S. K. Rumi, and F. D. Salim. Modelling memory for individual re-identification in decentralised mobile contact tracing applications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(1):1–21, 2021.
- [33] T. Behr, V. T. de Almeida, and R. H. Güting. Representation of periodic moving objects in databases. In *Proceedings of the 14th Annual ACM International Symposium on Advances in Geographic Information Systems*, GIS '06, page 43, New York, NY, USA, 2006. Association for Computing Machinery.
- [34] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz. SUMO – Simulation of Urban MObility: An Overview. In *Proceedings of SIMUL 2011, The Third International Conference on Advances in System Simulation*, 2011.
- [35] R. Benetis, C. S. Jensen, G. Karčiauskas, and S. Šaltenis. Nearest and reverse nearest neighbor queries for moving objects. *The VLDB Journal*, 15:229–249, 2006.
- [36] C. Berberidis, I. P. Vlahavas, W. G. Aref, M. J. Atallah, and A. K. Elmagarmid. On the discovery of weak periodicities in large time series. In *Principles of Data Mining and Knowledge Discovery, 6th European Conference, PKDD 2002, Helsinki, Finland, August 19-23, 2002, Proceedings*, volume 2431 of *Lecture Notes in Computer Science*, pages 51–61. Springer, 2002.
- [37] J. Biagioni and J. Eriksson. Inferring Road Maps from Global Positioning System Traces: Survey and Comparative Evaluation. *Transportation Research Record: Journal of the Transportation Research Board*, 2291(1):61–71, 2012.
- [38] Bike Share Metro Data. <https://bikeshare.metro.net/about/data/>.
- [39] Bike Town System Data. <https://www.biketownpdx.com/system-data>.
- [40] F. Biljecki, H. Ledoux, and P. Van Oosterom. Transportation mode-based segmentation and classification of movement trajectories. *International Journal of Geographical Information Science*, 27(2):385–407, 2013.
- [41] L. Bracciale, M. Bonola, P. Loreti, G. Bianchi, R. Amici, and A. Rabuffi. CRAWDAD dataset roma/taxi (v. 2014-07-17). Downloaded from <https://crawdad.org/roma/taxi/20140717>, July 2014.
- [42] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On Map-Matching Vehicle Tracking Data. In *VLDB*, pages 853–864, Trondheim, Norway, Aug. 2005.
- [43] Y. Cao, Y. Xiao, L. Xiong, L. Bai, and M. Yoshikawa. Protecting spatiotemporal event privacy in continuous location-based services. *IEEE Transactions on Knowledge and Data Engineering*, 33(8):3141–3154, 2019.
- [44] Capital Bike Share System Data. <https://www.capitalbikeshare.com/system-data>.
- [45] S. Carlucci, M. De Simone, S. K. Firth, M. B. Kjærgaard, R. Markovic, M. S. Rahaman, M. K. Annaqeeb, S. Biandrate, A. Das, J. W. Dziedzic, et al. Modeling occupant behavior in buildings. *Building and Environment*, 174:106768, 2020.
- [46] P. Chao, Y. Xu, W. Hua, and X. Zhou. A Survey on Map-Matching Algorithms. In *Australasian Database Conference, ADC*, pages 121–133, Melbourne, Australia, Feb. 2020.
- [47] K. Chatzikokolakis, E. ElSalamouny, C. Palamidessi, P. Anna, et al. Methods for location privacy: A comparative overview. *Foundations and Trends® in Privacy and Security*, 1(4):199–257, 2017.
- [48] C. Chen, C. Lu, Q. Huang, Q. Yang, D. Gunopulos, and L. J. Guibas. City-scale map creation and updating using GPS collections. In B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, and R. Rastogi, editors, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 1465–1474. ACM, 2016.
- [49] J. Chen and R. Cheng. Efficient evaluation of imprecise location-dependent queries. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 586–595. IEEE, 2006.
- [50] K. Chen, D. Zhang, L. Yao, B. Guo, Z. Yu, and Y. Liu. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Comput. Surv.*, 54(4), may 2021.
- [51] L. Chen, S. Shang, C. S. Jensen, B. Yao, Z. Zhang, and L. Shao. Effective and efficient reuse of past travel behavior for route recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 488–498, 2019.
- [52] R. Chen, H. Li, A. K. Qin, S. P. Kasiviswanathan, and H. Jin. Private spatial data aggregation in the local setting. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*, pages 289–300. IEEE, 2016.

- [53] C. Cheng, H. Yang, M. R. Lyu, and I. King. Where you like to go next: Successive point-of-interest recommendation. In *IJCAI*, volume 13, pages 2605–2611, 2013.
- [54] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui. Exploring millions of footprints in location sharing services. *ICWSM*, 2011:81–88, 2011.
- [55] E. Cho, S. A. Myers, and J. Leskovec. 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, 2011.
- [56] H. Cho, D. Ippolito, and Y. W. Yu. Contact tracing mobile apps for COVID-19: Privacy considerations and related trade-offs. *arXiv preprint arXiv:2003.11511*, 2020.
- [57] X. Chu, I. F. Ilyas, S. Krishnan, and J. Wang. Data cleaning: Overview and emerging challenges. In *Proceedings of the 2016 international conference on management of data*, pages 2201–2206, 2016.
- [58] B. Chuaysi and S. Kiattisin. Fishing vessels behavior identification for combating iuu fishing: Enable traceability at sea. *Wireless Personal Communications*, 115:2971–2993, 2020.
- [59] S.-H. Chung, H.-L. Ma, M. Hansen, and T.-M. Choi. Data science and analytics in aviation, 2020.
- [60] CitiBike System Data. <https://ride.citibikenyc.com/system-data>.
- [61] C. Claramunt, C. Ray, L. Salmon, E. Camossi, M. Hadzagic, A. Joussetme, G. Andrienko, N. Andrienko, Y. Theodoridis, and G. Vouros. Maritime data integration and analysis: recent progress and research challenges. *Advances in Database Technology-EDBT*, 2017:192–197, 2017.
- [62] CoGo Bike Share System Data. <https://www.cogobikeshare.com/system-data>.
- [63] G. Cormode, S. Jha, T. Kulkarni, N. Li, D. Srivastava, and T. Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1655–1658, 2018.
- [64] Y. Da, R. Ahuja, L. Xiong, and C. Shahabi. React: real-time contact tracing and risk monitoring via privacy-enhanced mobile tracking. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pages 2729–2732. IEEE, 2021.
- [65] Y. Da, R. Ahuja, L. Xiong, and C. Shahabi. REACT: real-time contact tracing and risk monitoring via privacy-enhanced mobile tracking. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021*, pages 2729–2732. IEEE, 2021.
- [66] A. Daghistani, W. G. Aref, A. Ghafour, and A. R. Mahmood. SWARM: adaptive load balancing in distributed streaming systems for big spatial data. *ACM Trans. Spatial Algorithms Syst.*, 7(3):14:1–14:43, 2021.
- [67] J. Dai, B. Yang, C. Guo, and Z. Ding. Personalized route recommendation using big trajectory data. In *2015 IEEE 31st International Conference on Data Engineering*, pages 543–554, 2015.
- [68] M. Dallachiesa, A. Ebad, A. Eldawy, A. Elmagarmid, I. F. Ilyas, M. Ouzzani, and N. Tang. Nadeef: a commodity data cleaning system. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 541–552, 2013.
- [69] A. M. De Souza, R. S. Yokoyama, G. Maia, A. Loureiro, and L. Villas. Real-time path planning to prevent traffic jam through an intelligent transportation system. In *2016 IEEE symposium on computers and communication (ISCC)*, pages 726–731. IEEE, 2016.
- [70] D. Dias and L. H. M. K. Costa. CRAWDAD dataset coppe-ufjr/riobuses (v. 2018-03-19). Downloaded from <https://crawdad.org/coppe-ufjr/RioBuses/20180319>, Mar. 2018.
- [71] Divvy Bikes System Data. <https://www.divvybikes.com/system-data>.
- [72] B. Dong, Y. Liu, H. Fontenot, M. Ouf, M. Osman, A. Chong, S. Qin, F. Salim, H. Xue, D. Yan, et al. Occupant behavior modeling methods for resilient building design, operation and policy at urban scale: A review. *Applied Energy*, 293:116856, 2021.
- [73] B. Dong, Y. Liu, W. Mu, Z. Jiang, P. Pandey, T. Hong, B. Olesen, T. Lawrence, Z. O’Neil, et al. A global building occupant behavior database. *Scientific data*, 9(1):369, 2022.
- [74] E. d’Afflisio, P. Braca, and P. Willett. Malicious ais spoofing and abnormal stealth deviations: A comprehensive statistical framework for maritime anomaly detection. *IEEE Transactions on Aerospace and Electronic Systems*, 57(4):2093–2108, 2021.
- [75] J. Elarde, J.-S. Kim, H. Kavak, A. Züfle, and T. Anderson. Change of human mobility during COVID-19: A united states case study. *PLoS one*, 16(11):e0259031, 2021.
- [76] M. G. Elfeky, W. G. Aref, and A. K. Elmagarmid. Using convolution to mine obscure periodic patterns in one pass. In *Advances in Database Technology - EDBT 2004, 9th International Conference on Extending Database Technology, Heraklion, Crete, Greece, March 14-18, 2004, Proceedings*, volume 2992 of *Lecture Notes in Computer Science*, pages 605–620. Springer, 2004.
- [77] M. M. Elshrif, K. Isufaj, and M. F. Mokbel. Network-less trajectory imputation. In *SIGSPATIAL*, pages 8:1–8:10, Nov. 2022.
- [78] M.-C. Fang and Y.-H. Lin. The optimization of ship weather-routing algorithm based on the composite influence of multi-dynamic elements (ii): Optimized routings. *Applied Ocean Research*, 50:130–140, 2015.
- [79] Federal Geographic Data Committee. Geospatial Metadata Standards and Guidelines. <https://www.fgdc.gov/metadata/geospatial-metadata-standard>.
- [80] R. G. Fegeas, J. L. Cascio, and R. A. Lazar. An Overview of FIPS 173, The Spatial Data Transfer Standard. *Cartography and Geographic Information Systems*, 19(5):278–293, 1992.
- [81] J. Feng, Y. Li, C. Zhang, F. Sun, F. Meng, A. Guo, and D. Jin. DeepMove: Predicting Human Mobility with Attentional Recurrent Networks. In *WWW*, pages 1459–1468, Apr. 2018.
- [82] M. Fredrikson, S. Jha, and T. Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1322–1333, 2015.
- [83] K. Fu, Z. Chen, and C.-T. Lu. Streetnet: preference learning with convolutional neural network on urban crime perception. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 269–278, 2018.

- [84] N. Gao, M. Marschall, J. Burry, S. Watkins, and F. D. Salim. Understanding occupants' behaviour, engagement, emotion, and comfort indoors with heterogeneous sensors and wearables. *Scientific Data*, 9(1):261, 2022.
- [85] N. Gao, H. Xue, W. Shao, S. Zhao, K. K. Qin, A. Prabowo, M. S. Rahaman, and F. D. Salim. Generative adversarial networks for spatio-temporal data: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(2):1–25, 2022.
- [86] Q. Gao, F. Zhou, K. Zhang, G. Trajcevski, X. Luo, and F. Zhang. Identifying human mobility via trajectory embeddings. In *IJCAI*, volume 17, pages 1689–1695, 2017.
- [87] Q. Gao, F. Zhou, T. Zhong, G. Trajcevski, X. Yang, and T. Li. Contextual spatio-temporal graph representation learning for reinforced human mobility mining. *Inf. Sci.*, 606:230–249, 2022.
- [88] G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. Private queries in location based services: anonymizers are not necessary. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 121–132, 2008.
- [89] M. Goodchild. Twenty years of progress: Giscience in 2010. *Journal of Spatial Information Science*, 1:3–20, 2010.
- [90] M. F. Goodchild. Uncertainty: The Achilles Heel of GIS. *Geo Info Systems*, 8(11):50–52, 1998.
- [91] M. Graham, M. Kutzbach, and B. McKenzie. Design comparison of lodes and acs commuting data products. Technical report, US Census Bureau, Center for Economic Studies, 2014.
- [92] A. Graser. An exploratory data analysis protocol for identifying problems in continuous movement data. *Journal of Location Based Services*, 15(2):89–117, 2021.
- [93] A. Graser. The State of Trajectory Visualization in Notebook Environments. *GI Forum*, 1:73–91, 2023.
- [94] S. Grumbach, P. Rigaux, and L. Segoufin. The dedale system for complex spatial queries. *SIGMOD Rec.*, 27(2):213–224, jun 1998.
- [95] X. Gu, M. Li, L. Xiong, and Y. Cao. Providing input-discriminative protection for local differential privacy. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 505–516. IEEE, 2020.
- [96] M. F. Guagliardo. Spatial accessibility of primary care: concepts, methods and challenges. *International journal of health geographics*, 3(1):1–13, 2004.
- [97] C. Guo, B. Yang, O. Andersen, C. S. Jensen, and K. Torp. Ecomark 2.0: empowering eco-routing with vehicular environmental models and actual vehicle fuel consumption data. *GeoInformatica*, 19:567–599, 2015.
- [98] R. H. Güting, T. Behr, and C. Düntgen. Secondo: A platform for moving objects database research and for publishing and integrating research implementations. *IEEE Data Eng. Bull.*, 33:56–63, 2010.
- [99] R. H. Güting, V. T. De Almeida, and Z. Ding. Modeling and querying moving objects in networks. *The VLDB Journal*, 15:165–190, 2006.
- [100] R. H. Güting and M. Schneider. *Moving objects databases*. Elsevier, 2005.
- [101] R. Gajting, M. Bathlen, M. Erwig, C. Jensen, N. Lorentzos, M. Schneider, and M. Vazirgiannis. A foundation for representing and querying moving objects. *ACM Transactions on Database Systems (TODS)*, 25:1–42, 03 2000.
- [102] X. Han, R. Cheng, C. Ma, and T. Grubenmann. Deeptea: effective and efficient online time-dependent trajectory outlier detection. *Proceedings of the VLDB Endowment*, 15(7):1493–1505, 2022.
- [103] R. Harle. A survey of indoor inertial positioning systems for pedestrians. *IEEE Communications Surveys & Tutorials*, 15(3):1281–1293, 2013.
- [104] X. He, G. Cormode, A. Machanavajjhala, C. Procopiuc, and D. Srivastava. Dpt: differentially private trajectory synthesis using hierarchical reference systems. *Proceedings of the VLDB Endowment*, 8(11):1154–1165, 2015.
- [105] J. Hinnenthal and G. Clauss. Robust Pareto-optimum routing of ships utilising deterministic and ensemble weather forecasts. *Ships and Offshore Structures*, 5(2):105–114, Aug. 2010.
- [106] B. Hong, B. J. Bonczak, A. Gupta, and C. E. Kontokosta. Measuring inequality in community resilience to natural disasters using large-scale mobility data. *Nature communications*, 12(1):1870, 2021.
- [107] X. Hou, S. Gao, Q. Li, Y. Kang, N. Chen, K. Chen, J. Rao, J. S. Ellenberg, and J. A. Patz. Intracounty modeling of COVID-19 infection with human mobility: Assessing spatial heterogeneity with business traffic, age, and race. *Proceedings of the National Academy of Sciences*, 118(24):e2020524118, 2021.
- [108] X. Huang, Y. Yin, S. Lim, G. Wang, B. Hu, J. Varadarajan, S. Zheng, A. Bulusu, and R. Zimmermann. Grab-posisi: An extensive real-life GPS trajectory dataset in southeast asia. In *Proceedings of the ACM SIGSPATIAL International Workshop on Prediction of Human Mobility, PredictGIS 2019*, pages 1–10, Chicago, IL, USA, Nov. 2019.
- [109] Z. Huang, H. Lu, B. C. Ooi, and A. K. H. Tung. Continuous Skyline Queries for Moving Objects. *IEEE Transactions on Knowledge and Data Engineering*, 18(12):1645–1658, 2006.
- [110] Iea (2023), CO2 emissions in 2022, iea, paris. <https://www.iea.org/reports/co2-emissions-in-2022> (Accessed 05/08/2023), License: CC BY 4.0.
- [111] F. Ijaz, H. K. Yang, A. W. Ahmad, and C. Lee. Indoor positioning: A review of indoor ultrasonic positioning systems. In *2013 15th International Conference on Advanced Communications Technology (ICACT)*, pages 1146–1150. IEEE, 2013.
- [112] Indego Trip Data. <https://www.rideindego.com/about/data/>.
- [113] INRIX. 2022 global traffic scorecard: <https://inrix.com/scorecard/> (accessed 01/11/2023).
- [114] C. S. Jensen, H. Lu, and B. Yang. Graph model based indoor tracking. In *2009 Tenth International Conference on Mobile Data Management: Systems, Services and Middleware*, pages 122–131, 2009.
- [115] C. S. Jensen, H. Lu, and B. Yang. Indoor-a new data management frontier. *IEEE Data Eng. Bull.*, 33(2):12–17, 2010.

- [116] F. Jin, W. Hua, J. Xu, and X. Zhou. Moving object linking based on historical trace. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1058–1069. IEEE, 2019.
- [117] T. Jindal, P. Giridhar, L. Tang, J. Li, and J. Han. Spatiotemporal periodical pattern mining in traffic data. In *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing, UrbComp@KDD 2013, Chicago, Illinois, USA, August 11, 2013*, pages 11:1–11:8. ACM, 2013.
- [118] R. Joo, M. E. Boone, T. A. Clay, S. C. Patrick, S. Clusella, R. Trullas, and M. Basille. Navigating through the R packages for movement. *Journal of Animal Ecology*, 89(1):248–267, Jan. 2020.
- [119] G. Jossé, K. A. Schmid, A. Züfle, G. Skoumas, M. Schubert, and D. Pfoser. Turismo: A user-preference tourist trip search engine. In *Advances in Spatial and Temporal Databases: 14th International Symposium, SSTD 2015, Hong Kong, China, August 26–28, 2015. Proceedings 14*, pages 514–519. Springer, 2015.
- [120] J. W. Kim and B. Jang. Workload-aware indoor positioning data collection via local differential privacy. *IEEE communications letters*, 23(8):1352–1356, 2019.
- [121] M. R. Kolahdouzan and C. Shahabi. Voronoi-based K nearest neighbor search for spatial network databases. In M. A. Nascimento, M. T. Özsu, D. Kossmann, R. J. Miller, J. A. Blakeley, and K. B. Schiefer, editors, *(e)Proceedings of the Thirtieth International Conference on Very Large Data Bases, VLDB 2004, Toronto, Canada, August 31 - September 3 2004*, pages 840–851. Morgan Kaufmann, 2004.
- [122] J. Krumm. Maximum entropy bridgelets for trajectory completion. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pages 1–8, 2022.
- [123] P. Lara-Benitez, M. Carranza-García, and J. C. Riquelme. An experimental review on deep learning architectures for time series forecasting. *Int. J. Neural Syst.*, 31(3):2130001:1–2130001:28, 2021.
- [124] C. Lee, Y. Kim, S. Jin, D. Kim, R. Maciejewski, D. Ebert, and S. Ko. A visual analytics system for exploring, monitoring, and forecasting road traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3133–3146, 2020.
- [125] B. Li, Z. Cai, M. Kang, S. Su, S. Zhang, L. Jiang, and Y. Ge. A Trajectory Restoration Algorithm for Low-sampling-rate Floating Car Data and Complex Urban Road Networks. *International Journal of GIS*, 35(4):717–740, 2021.
- [126] H. Li, H. Lu, C. S. Jensen, B. Tang, and M. A. Cheema. Spatial Data Quality in the Internet of Things: Management, Exploitation, and Prospects. *ACM Computing Surveys*, 55(3):57:1–57:41, 2023.
- [127] M. Li, R. Westerholt, H. Fan, and A. Zipf. Assessing spatiotemporal predictability of lbsn: a case study of three foursquare datasets. *GeoInformatica*, pages 1–21, 2016.
- [128] R. Li, H. He, R. Wang, S. Ruan, Y. Sui, J. Bao, and Y. Zheng. Trajmesa: A distributed nosql storage engine for big trajectory data. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 2002–2005. IEEE, 2020.
- [129] X. Li, J. Han, J.-G. Lee, and H. Gonzalez. Traffic density-based discovery of hot routes in road networks. In *Advances in Spatial and Temporal Databases: 10th International Symposium, SSTD 2007, Boston, MA, USA, July 16–18, 2007. Proceedings 10*, pages 441–459. Springer, 2007.
- [130] Y. Li, D. Gunopulos, C. Lu, and L. J. Guibas. Personalized travel time prediction using a small number of probe vehicles. *ACM Trans. Spatial Algorithms Syst.*, 5(1):4:1–4:27, 2019.
- [131] Y. Li, Y. Li, D. Gunopulos, and L. J. Guibas. Knowledge-based Trajectory Completion from Sparse GPS Samples. In *SIGSPATIAL*, pages 33:1–33:10, Oct. 2016.
- [132] Y. Li, R. Yu, C. Shahabi, and Y. Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [133] Y. Li, Y. Zheng, H. Zhang, and L. Chen. Traffic prediction in a bike-sharing system. In *Proceedings of the 23rd SIGSPATIAL international conference on advances in geographic information systems*, pages 1–10, 2015.
- [134] C. Lin, K. L. Choy, G. T. Ho, S. H. Chung, and H. Lam. Survey of green vehicle routing problem: past and future trends. *Expert systems with applications*, 41(4):1118–1138, 2014.
- [135] Z. Lin, S. Lyu, H. Cao, F. Xu, Y. Wei, H. Samet, and Y. Li. Healthwalks: Sensing fine-grained individual health condition via mobility data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4):1–26, 2020.
- [136] London Cycling Data. <https://cycling.data.tfl.gov.uk/>.
- [137] J. A. Long. Kinematic Interpolation of Movement Data. *International Journal of Geographical Information Science*, 30(5):854–868, 2016.
- [138] M. Luca, G. Barlacchi, B. Lepri, and L. Pappalardo. A survey on deep learning for human mobility. *ACM Comput. Surv.*, 55(2):7:1–7:44, 2023.
- [139] BlueBikes System Data. <https://www.bluebikes.com/system-data>.
- [140] Lyft Bikes Bay Wheels Trip Data. <https://www.lyft.com/bikes/bay-wheels/system-data>.
- [141] J. Macfarlane and M. Stroila. Addressing the uncertainties in autonomous driving. *SIGSPATIAL Special*, 8(2):35–40, 2016.
- [142] Madrid Open Data. [https://opendata.emtmadrid.es/Datos-estaticos/Datos-generales-\(1\)](https://opendata.emtmadrid.es/Datos-estaticos/Datos-generales-(1)).
- [143] H. Maeda, Y. Sekimoto, and T. Seto. Lightweight road manager: smartphone-based automatic determination of road damage status by deep neural network. In *Proceedings of the 5th ACM SIGSPATIAL International Workshop on Mobile Geographic Information Systems*, pages 37–45, 2016.
- [144] A. R. Mahmood, A. M. Aly, T. Kuznetsova, S. M. Basalamah, and W. G. Aref. Disk-based indexing of recent trajectories. *ACM Trans. Spatial Algorithms Syst.*, 4(3):7:1–7:27, 2018.
- [145] A. R. Mahmood, A. M. Aly, T. Qadah, E. K. Rezig, A. Daghistani, A. Madkour, A. S. Abdelhamid, M. S. Hassan, W. G. Aref, and S. M. Basalamah. Tornado: A distributed spatio-textual stream processing system. *Proc. VLDB Endow.*, 8(12):2020–2023, 2015.

- [146] A. R. Mahmood, A. Daghistani, A. M. Aly, M. Tang, S. M. Basalamah, S. Prabhakar, and W. G. Aref. Adaptive processing of spatial-keyword data over a distributed streaming cluster. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2018, Seattle, WA, USA, November 06-09, 2018*, pages 219–228. ACM, 2018.
- [147] A. R. Mahmood, S. Punni, and W. G. Aref. Spatio-temporal access methods: A survey (2010 - 2017). *GeoInformatica*, 23(1):1–36, 2019.
- [148] F. Mazzarella, M. Vespe, A. Alessandrini, D. Tarchi, G. Aulicino, and A. Vollero. A novel anomaly detection approach to identify intentional ais on-off switching. *Expert Systems with Applications*, 78:110–123, 2017.
- [149] Discover New Roads with Bing Maps. <https://blogs.bing.com/maps/2022-12/Bing-Maps-is-bringing-new-roads/>.
- [150] E. Min, X. Guo, Q. Liu, G. Zhang, J. Cui, and J. Long. A survey of clustering with deep learning: From the perspective of network architecture. *IEEE Access*, 6:39501–39514, 2018.
- [151] D. J. Mir, S. Isaacman, R. Cáceres, M. Martonosi, and R. N. Wright. Dp-where: Differentially private modeling of human mobility. In *2013 IEEE international conference on big data*, pages 580–588. IEEE, 2013.
- [152] M. Mokbel, S. Abbar, and R. Stanojevic. Contact tracing: Beyond the apps. *SIGSPATIAL Special*, 12(2):15–24, 2020.
- [153] M. Mokbel, M. Sakr, L. Xiong, A. Züfle, J. Almeida, T. Anderson, W. Aref, G. Andrienko, N. Andrienko, Y. Cao, et al. Mobility data science (dagstuhl seminar 20201). In *Dagstuhl reports*, volume 12. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [154] M. F. Mokbel, L. Alarabi, J. Bao, A. Eldawy, A. Magdy, M. Sarwat, E. Waytas, and S. Yackel. Mntg: An extensible web-based traffic generator. In *Advances in Spatial and Temporal Databases: 13th International Symposium, SSTD 2013, Munich, Germany, August 21-23, 2013. Proceedings 13*, pages 38–55. Springer, 2013.
- [155] M. F. Mokbel, T. M. Ghanem, and W. G. Aref. Spatio-temporal access methods. *IEEE Data Eng. Bull.*, 26(2):40–49, 2003.
- [156] M. F. Mokbel, L. Xiong, and D. Zeinalipour-Yazti. Introduction to the special issue on contact tracing, 2022.
- [157] M. F. Mokbel, X. Xiong, and W. G. Aref. SINA: Scalable Incremental Processing of Continuous Queries in Spatio-temporal Databases. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 623–634, June 2004.
- [158] P. Mooney, M. Minghini, et al. A review of openstreetmap data. *Mapping and the citizen sensor*, pages 37–59, 2017.
- [159] S. Moosavi, M. H. Samavatian, S. Parthasarathy, R. Teodorescu, and R. Ramnath. Accident risk prediction based on heterogeneous sparse data: New dataset and insights. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 33–42, 2019.
- [160] M. Musleh, S. Abbar, R. Stanojevic, and M. Mokbel. Qarta: an ml-based system for accurate map services. *Proceedings of the VLDB Endowment*, 14(11):2273–2282, 2021.
- [161] M. Musleh and M. F. Mokbel. RASSED: A Scalable Dashboard for Monitoring Road Network Updates in OSM. In *MDM*, pages 214–221, June 2022.
- [162] A. M. Nagy and V. Simon. Survey on traffic prediction in smart cities. *Pervasive and Mobile Computing*, 50:148–163, 2018.
- [163] New York Times. For Big-Data Scientists, 'Janitor Work' Is Key Hurdle to Insights <https://www.nytimes.com/2014/08/18/technology/for-big-data-scientists-hurdle-to-insights-is-janitor-work.html> (Accessed 05/05/2023).
- [164] L. Nguyen-Dinh, W. G. Aref, and M. F. Mokbel. Spatio-temporal access methods: Part 2 (2003 - 2010). *IEEE Data Eng. Bull.*, 33(2):46–55, 2010.
- [165] NiceRide System Data. <https://www.niceridemn.com/system-data>.
- [166] J. Nijman and Y. D. Wei. Urban inequalities in the 21st century economy. *Applied geography*, 117:102188, 2020.
- [167] P. Nikitopoulos, A.-I. Paraskevopoulos, C. Doukeridis, N. Pelekis, and Y. Theodoridis. Hot spot analysis over big trajectory data. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 761–770, 2018.
- [168] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil. An empirical study of geographic user activity patterns in Foursquare. *ICWSM*, 11:70–573, 2011.
- [169] Kaggle. New York City Taxi Trip Duration. <https://www.kaggle.com/c/nyc-taxi-trip-duration/data>.
- [170] Oracle and/or its affiliates. Stream IoT data to an autonomous database using serverless functions F35431-06, 2022.
- [171] Open street map. <http://www.openstreetmap.org>.
- [172] F. Ossi, F. Hachem, F. Cagnacci, U. Demšar, and M. L. Damiani. Hanimob 2021 workshop report: The 1st acm sigspatial workshop on animal movement ecology and human mobility. *SIGSPATIAL Special*, 13(1-3):33–36, 2022.
- [173] S. E. Page. Computational models from A to Z. *Complexity*, 5(1):35–41, 1999.
- [174] B. Pan, Y. Zheng, D. Wilkie, and C. Shahabi. Crowd sensing of traffic anomalies based on human mobility and social media. In *Proceedings of the 21st ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 344–353, 2013.
- [175] L. Pappalardo, P. Cintia, A. Rossi, E. Massucco, P. Ferragina, D. Pedreschi, and F. Giannotti. A public data set of spatio-temporal match events in soccer competitions. *Nature Scientific data*, 6(1):236, 2019.
- [176] C. Parent, S. Spaccapietra, C. Renso, G. Andrienko, N. Andrienko, V. Bogorny, M. L. Damiani, A. Gkoulalas-Divanis, J. Macedo, N. Pelekis, et al. Semantic trajectories modeling and analysis. *ACM Computing Surveys (CSUR)*, 45(4):1–32, 2013.
- [177] K. Patroumpas, E. Alevizos, A. Artikis, M. Voudas, N. Pelekis, and Y. Theodoridis. Online event recognition from moving vessel trajectories. *GeoInformatica*, 21:389–427, 2017.
- [178] N. Pelekis, C. Renso, Y. Theodoridis, and K. Zeitouni. Editor's note. *GeoInformatica*, 26(3):449, 2022.
- [179] D. Pfoser. *Crowdsourcing Geographic Information Systems*, pages 1–8. Springer New York, New York, NY, 2016.
- [180] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauer. CRAWDAD dataset epfl/mobility (v. 2009-02-24). Downloaded from <https://crawdad.org/epfl/mobility/20090224>, Feb. 2009.
- [181] Taxi Service Trajectory. Prediction Challenge. ECML PKDD 2015. <http://www.geolink.pt/ecmlpkdd2015-challenge/dataset.html>.

- [182] PostGIS Project Steering Committee and others (<https://postgis.net>). Postgis, spatial and geographic objects for postgresql, 2023.
- [183] A. Pyrgelis, C. Troncoso, and E. D. Cristofaro. Knock knock, who's there? membership inference on aggregate location data, 2018.
- [184] A. Pyrgelis, C. Troncoso, and E. De Cristofaro. Knock knock, who's there? membership inference on aggregate location data. *arXiv preprint arXiv:1708.06145*, 2017.
- [185] A. Pyrgelis, C. Troncoso, and E. De Cristofaro. What does the crowd say about you? evaluating aggregation-based location privacy. *Proceedings on privacy enhancing technologies*, 2017(4):156–176, 2017.
- [186] W. Qardaji, W. Yang, and N. Li. Differentially private grids for geospatial data. In *2013 IEEE 29th international conference on data engineering (ICDE)*, pages 757–768. IEEE, 2013.
- [187] S. Rambhatla, S. Zeighami, K. Shahabi, C. Shahabi, and Y. Liu. Toward accurate spatiotemporal COVID-19 risk scores using high-resolution real-world mobility data. *ACM Trans. Spatial Algorithms Syst.*, 8(2):1–30, 2022.
- [188] L. Reichert, S. Brack, and B. Scheuermann. Privacy-preserving contact tracing of COVID-19 patients. *Cryptology ePrint Archive*, 2020.
- [189] L. Reichert, S. Brack, and B. Scheuermann. A survey of automatic contact tracing approaches using bluetooth low energy. *ACM Transactions on Computing for Healthcare*, 2(2):1–33, 2021.
- [190] Y. Ren, M. Tomko, F. D. Salim, J. Chan, and M. Sanderson. Understanding the predictability of user demographics from cyber-physical-social behaviours in indoor retail spaces. *EPJ Data Science*, 7(1):1–21, 2018.
- [191] Y. Ren, M. Tomko, F. D. Salim, K. Ong, and M. Sanderson. Analyzing web behavior in indoor retail spaces. *Journal of the Association for Information Science and Technology*, 68(1):62–76, 2017.
- [192] E. K. Rezig, L. Cao, M. Stonebraker, G. Simonini, W. Tao, S. Madden, M. Ouzzani, N. Tang, and A. K. Elmagarmid. Data civilizer 2.0: A holistic framework for data preparation and analytics. *Proc. VLDB Endow.*, 12(12):1954–1957, aug 2019.
- [193] K. Richly, R. Teusner, A. Immer, F. Windheuser, and L. Wolf. Optimizing routes of public transportation systems by analyzing the data of taxi rides. In *Proceedings of the 1st International ACM SIGSPATIAL Workshop on Smart Cities and Urban Analytics*, pages 70–76, 2015.
- [194] Ride Austin Dataset. <https://data.world/ride-austin/ride-austin-june-6-april-13>.
- [195] M. Riveiro, G. Pallotta, and M. Vespe. Maritime anomaly detection: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(5):e1266, 2018.
- [196] H. Rizk, M. Torki, and M. Youssef. Cellindeep: Robust and accurate cellular-based indoor localization via deep learning. *IEEE Sensors Journal*, 19(6):2305–2312, 2018.
- [197] A. Rodríguez, H. Kamarthi, P. Agarwal, J. Ho, M. Patel, S. Sapre, and B. A. Prakash. Data-centric epidemic forecasting: A survey. *arXiv preprint arXiv:2207.09370*, 2022.
- [198] Safegraph. Safegraph. Places Data Curated for Accurate Geospatial Analytics. <https://www.safegraph.com/>.
- [199] SafeGraph Inc. Weekly Patterns Dataset (<https://docs.safegraph.com/docs/weekly-patterns>).
- [200] M. Sakr, C. Ray, and C. Renso. Big mobility data analytics: Recent advances and open problems. *Geoinformatica*, 26(4):541–549, oct 2022.
- [201] M. A. Sakr and R. H. Güting. Group spatiotemporal pattern queries. *Geoinformatica*, 18:699–746, 2014.
- [202] F. D. Salim, B. Dong, M. Ouf, Q. Wang, I. Pigliautile, X. Kang, T. Hong, W. Wu, Y. Liu, S. K. Rumi, M. S. Rahaman, J. An, H. Deng, W. Shao, J. Dziedzic, F. C. Sangogboye, M. B. Kjærsgaard, M. Kong, C. Fabiani, A. L. Pisello, and D. Yan. Modelling urban-scale occupant behaviour, mobility, and energy in buildings: A survey. *Building and Environment*, 183, 2020.
- [203] S. Scellato, A. Noulas, and C. Mascolo. Exploiting place features in link prediction on location-based social networks. In *ACM SIGKDD*, pages 1046–1054, 2011.
- [204] E. Seglem, A. Züfle, J. Stutzki, F. Borutta, E. Faerman, and M. Schubert. On privacy in spatio-temporal data: User identification using microblog data. In *Advances in Spatial and Temporal Databases: 15th International Symposium, SSTD 2017, Arlington, VA, USA, August 21–23, 2017, Proceedings 15*, pages 43–61. Springer, 2017.
- [205] S. Shah, F. Bao, C.-T. Lu, and I.-R. Chen. Crowdsafe: crowd sourcing of crime incidents and safe routing on mobile devices. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 521–524, 2011.
- [206] S. Shaham, G. Ghinita, R. Ahuja, J. Krumm, and C. Shahabi. Htf: Homogeneous tree framework for differentially-private release of large geospatial datasets with self-tuning structure height. *ACM Transactions on Spatial Algorithms and Systems*, 2022.
- [207] S. Shaham, G. Ghinita, and C. Shahabi. Models and mechanisms for spatial data fairness. *Proc. VLDB Endow.*, 16(2):167–179, 2022.
- [208] S. Shang, B. Yuan, K. Deng, K. Xie, and X. Zhou. Finding the most accessible locations: reverse path nearest neighbor query in road networks. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 181–190, 2011.
- [209] Z. Shang, G. Li, and Z. Bao. Dita: distributed in-memory trajectory analytics. In *Proceedings of the 2018 International Conference on Management of Data*, pages 725–740, 2018.
- [210] W. Shao, A. Prabowo, S. Zhao, S. Tan, P. Koniusz, J. Chan, X. Hei, B. Feest, and F. D. Salim. Flight delay prediction using airport situational awareness map. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 432–435, 2019.
- [211] A. Shashidharan, V. Chandola, and R. R. Vatsavai. The 9th acm sigspatial international workshop on analytics for big spatial data (bigspatial 2020) november 3, 2020. *SIGSPATIAL Special*, 12(3):15–16, 2021.
- [212] S. Shekhar, Z. Jiang, R. Y. Ali, E. Eftelioglu, X. Tang, V. M. V. Gunturi, and X. Zhou. Spatiotemporal data mining: A computational perspective. *ISPRS International Journal of Geo-Information*, 4(4):2306–2338, 2015.

- [213] J. Shin, J. Wang, and W. G. Aref. The LSM rum-tree: A log structured merge r-tree for update-intensive spatial workloads. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021*, pages 2285–2290. IEEE, 2021.
- [214] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.
- [215] Y. Shu, C. Bo, G. Shen, C. Zhao, L. Li, and F. Zhao. Magicol: Indoor localization using pervasive magnetic field and opportunistic wifi sensing. *IEEE Journal on Selected Areas in Communications*, 33(7):1443–1457, 2015.
- [216] Y. N. Silva, W. G. Aref, P. Larson, S. Pearson, and M. H. Ali. Similarity queries: their conceptual evaluation, transformations, and processing. *VLDB J.*, 22(3):395–420, 2013.
- [217] T. G. Stavropoulos, A. Papastergiou, L. Mpaltadoros, S. Nikolopoulos, and I. Kompatsiaris. Iot wearable sensors and devices in elderly care: A literature review. *Sensors*, 20(10):2826, 2020.
- [218] D. Sui, S. Elwood, and M. Goodchild. *Crowdsourcing Geographic Knowledge: Volunteered Geographic Information (VGI) in Theory and Practice*. Springer Science & Business Media, 2012.
- [219] X. Teng, G. Trajcevski, J.-S. Kim, and A. Züfle. Semantically diverse path search. In *2020 21st IEEE International Conference on Mobile Data Management (MDM)*, pages 69–78. IEEE, 2020.
- [220] X. Teng, G. Trajcevski, and A. Züfle. Semantically diverse paths with range and origin constraints. In *Proceedings of the 29th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 375–378, 2021.
- [221] H. To, G. Ghinita, and C. Shahabi. A framework for protecting worker location privacy in spatial crowdsourcing. *Proceedings of the VLDB Endowment*, 7(10):919–930, 2014.
- [222] H. To, C. Shahabi, and L. Xiong. Privacy-preserving online task assignment in spatial crowdsourcing with untrusted server. In *2018 IEEE 34th international conference on data engineering (ICDE)*, pages 833–844. IEEE, 2018.
- [223] W. R. Tobler. A computer movie simulating urban growth in the detroit region. *Economic geography*, 46(sup1):234–240, 1970.
- [224] M. I. Tolea, J. C. Morris, and J. E. Galvin. Trajectory of mobility decline by type of dementia. *Alzheimer disease and associated disorders*, 30(1):60, 2016.
- [225] D. Tomaras, V. Kalogeraki, T. Liebig, and D. Gunopulos. Crowd-based ecofriendly trip planning. In *19th IEEE International Conference on Mobile Data Management, MDM 2018, Aalborg, Denmark, June 25-28, 2018*, pages 24–33. IEEE Computer Society, 2018.
- [226] K. Toohey and M. Duckham. Trajectory similarity measures. *Sigspatial Special*, 7(1):43–50, 2015.
- [227] G. Trajcevski, O. Wolfson, K. Hinrichs, and S. Chamberlain. Managing uncertainty in moving objects databases. *ACM Trans. Database Syst.*, 29(3):463–507, sep 2004.
- [228] C. Troncoso, M. Payer, J.-P. Hubaux, M. Salathé, J. Larus, E. Bugnion, W. Lueks, T. Stadler, A. Pyrgelis, D. Antonioli, et al. Decentralized privacy-preserving proximity tracing. *arXiv preprint arXiv:2005.12273*, 2020.
- [229] R. Truong, O. Gkountouna, D. Pfoser, and A. Züfle. Towards a better understanding of public transportation traffic: A case study of the washington, dc metro. *Urban Science*, 2(3):65, 2018.
- [230] ULB Data science lab. MobilityDB, an open source geospatial trajectory data management and analysis platform (<https://mobilitydb.com>).
- [231] United Nations Conference on Trade and Development. Review of maritime transport 2022. <https://unctad.org/rmt2022>.
- [232] United Nations Department of Economic and Social Affairs. Revision of world urbanization prospects. *New York: United Nations Department of Economic and Social Affairs*, 2018.
- [233] United States Geological Survey. USGS Science Data Catalog. <https://data.usgs.gov/datacatalog/>.
- [234] U.S. Department of Transportation. Data Inventory. <https://www.transportation.gov/data>.
- [235] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [236] M. M. Vazifeh, H. Zhang, P. Santi, and C. Ratti. Optimizing the deployment of electric vehicle charging stations using pervasive mobility data. *Transportation Research Part A: Policy and Practice*, 121:75–91, 2019.
- [237] Y. Vigfusson, T. A. Karlsson, D. Onken, C. Song, A. F. Einarsson, N. Kishore, R. M. Mitchell, E. Brooks-Pollock, G. Sigmundsdottir, and L. Danon. Cell-phone traces reveal infection-associated behavioral change. *Proceedings of the National Academy of Sciences*, 118(6):e2005241118, 2021.
- [238] R. E. Walker, C. R. Keane, and J. G. Burke. Disparities and access to healthy food in the united states: A review of food deserts literature. *Health & place*, 16(5):876–884, 2010.
- [239] F. Wang, G. Wang, and S. Y. Philip. Why checkins: Exploring user motivation on location based social networks. In *Data Mining Workshop (ICDMW)*, pages 27–34. IEEE, 2014.
- [240] G. Wang, X. Chen, F. Zhang, Y. Wang, and D. Zhang. Experience: Understanding Long-Term Evolving Patterns of Shared Electric Vehicle Networks. In *Proceeding of the International Conference on Mobile Computing and Networking, MobiCom*, pages 1–12, Los Cabos, Mexico, Oct. 2019.
- [241] H. Wang, H. Hong, L. Xiong, Z. Qin, and Y. Hong. Privlbs: Local differential privacy for location-based services with staircase randomized response. In *Proceedings of the ACM Conference on Computer and Communications Security, 2022*.
- [242] H. Wang, X. Tang, Y.-H. Kuo, D. Kifer, and Z. Li. A simple baseline for travel time estimation using large-scale trip data. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–22, 2019.
- [243] H. Wang, Z. Zhang, T. Wang, S. He, M. Backes, J. Chen, and Y. Zhang. Privtrace: Differentially private trajectory synthesis by adaptive markov model. In *USENIX Security*, 2023.

- [244] H. Wang, K. Zheng, J. Xu, B. Zheng, X. Zhou, and S. Sadiq. Sharkdb: An in-memory column-oriented trajectory storage. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*, pages 1409–1418, 2014.
- [245] J. Wang, N. Wu, X. Lu, W. X. Zhao, and K. Feng. Deep Trajectory Recovery with Fine-Grained Calibration using Kalman Filter. *TKDE*, 33(3):921–934, 2021.
- [246] S. Wang, Z. Bao, J. S. Culpepper, and G. Cong. A survey on trajectory data management, analytics, and learning. *ACM Comput. Surv.*, 54(2), mar 2021.
- [247] Z. Wang, R. Jiang, H. Xue, F. D. Salim, X. Song, and R. Shibasaki. Event-aware multimodal mobility nowcasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4228–4236, 2022.
- [248] Y.-T. Wen, P.-R. Lei, W.-C. Peng, and X.-F. Zhou. Exploring social influence on location-based social networks. In *ICDM*, pages 1043–1048. IEEE, 2014.
- [249] R. T. Whitman, B. G. Marsh, M. B. Park, and E. G. Hoel. Distributed spatial and spatio-temporal join on apache spark. *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, 5(1):1–28, 2019.
- [250] L. Wu, Y. Xu, Q. Wang, F. Wang, and Z. Xu. Mapping global shipping density from ais data. *The Journal of Navigation*, 70(1):67–81, 2017.
- [251] Y. Xiao and L. Xiong. Protecting locations with differential privacy under temporal correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1298–1309, 2015.
- [252] J. Xie, Z. Song, Y. Li, Y. Zhang, H. Yu, J. Zhan, Z. Ma, Y. Qiao, J. Zhang, and J. Guo. A survey on machine learning-based mobile big data analysis: Challenges and applications. *Wirel. Commun. Mob. Comput.*, 2018:8738613:1–8738613:19, 2018.
- [253] X. Xiong, M. F. Mokbel, W. G. Aref, S. E. Hambrusch, and S. Prabhakar. Scalable spatio-temporal continuous query processing for location-aware services. In *Proceedings of the 16th International Conference on Scientific and Statistical Database Management (SSDBM 2004), 21-23 June 2004, Santorini Island, Greece*, pages 317–326. IEEE Computer Society, 2004.
- [254] J. Xu, R. H. Güting, and Y. Gao. Continuous k nearest neighbor queries over large multi-attribute trajectories: a systematic approach. *GeoInformatica*, 22:723–766, 2018.
- [255] H. Xue, F. Salim, Y. Ren, and N. Oliver. Mobtcast: Leveraging auxiliary trajectory forecasting for human mobility prediction. *Advances in Neural Information Processing Systems*, 34:30380–30391, 2021.
- [256] H. Xue and F. D. Salim. Termcast: Temporal relation modeling for effective urban flow forecasting. In *Advances in Knowledge Discovery and Data Mining: 25th Pacific-Asia Conference, PAKDD 2021*, pages 741–753. Springer, 2021.
- [257] C. Yang and H.-R. Shao. Wifi-based indoor positioning. *IEEE Communications Magazine*, 53(3):150–157, 2015.
- [258] Y. Yang, F. Zhang, and D. Zhang. SharedEdge: GPS-Free Fine-Grained Travel Time Estimation in State-Level Highway Systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(1):48:1–48:26, 2018.
- [259] H. Yin, Z. Hu, X. Zhou, H. Wang, K. Zheng, Q. V. H. Nguyen, and S. Sadiq. Discovering interpretable geo-social communities for user behavior prediction. In *ICDE*, pages 942–953. IEEE, 2016.
- [260] H. Yuan and G. Li. A survey of traffic prediction: from spatio-temporal data to intelligent transportation. *Data Science and Engineering*, 6:63–85, 2021.
- [261] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang. T-drive: driving directions based on taxi trajectories. In *Proceedings of the 18th SIGSPATIAL International conference on advances in geographic information systems*, pages 99–108, 2010.
- [262] A. Zaidi, R. Ahuja, and C. Shahabi. Differentially private occupancy monitoring from wifi access points. In *23rd IEEE International Conference on Mobile Data Management, MDM 2022, Paphos, Cyprus, June 6-9, 2022*, pages 361–366. IEEE, 2022.
- [263] A. Zhang, S. Song, J. Wang, and P. S. Yu. Time Series Data Cleaning: From Anomaly Detection to Anomaly Repairing. *PVLDB*, 10(10):1046–1057, 2017.
- [264] L. Zhang, L. Zhao, and D. Pfoser. Factorized deep generative models for end-to-end trajectory generation with spatiotemporal validity constraints. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pages 1–12, 2022.
- [265] P. Zhang, H. Zhang, and D. Guo. Evacuation shelter and route selection based on multi-objective optimization approach. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, pages 1–5, 2015.
- [266] Z. Zhang, C. Jin, J. Mao, X. Yang, and A. Zhou. Trajspark: A scalable and efficient in-memory management system for big trajectory data. In *Web and Big Data: First International Joint Conference, APWeb-WAIM 2017, Beijing, China, July 7–9, 2017, Proceedings, Part I 1*, pages 11–26. Springer, 2017.
- [267] K. Zhao, S. Tarkoma, S. Liu, and H. Vo. Urban human mobility data mining: An overview. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 1911–1920. IEEE, 2016.
- [268] S. Zhao, W. Shao, J. Chan, and F. D. Salim. Measuring disentangled generative spatio-temporal representation. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pages 522–530. SIAM, 2022.
- [269] X. Zhao, Y. Li, Y. Yuan, X. Bi, and G. Wang. Ldpart: effective location-record data publication via local differential privacy. *IEEE Access*, 7:31435–31445, 2019.
- [270] K. Zheng, Y. Zheng, X. Xie, and X. Zhou. Reducing Uncertainty of Low-Sampling-Rate Trajectories. In *ICDE*, pages 1144–1155, Apr. 2012.
- [271] Y. Zheng. Location-based social networks: Users. In *Computing with spatial trajectories*, pages 243–276. Springer, 2011.
- [272] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Urban computing: concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(3):1–55, 2014.

- [273] Y. Zheng, X. Xie, and W.-Y. Ma. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data(base) Engineering Bulletin*, June 2010.
- [274] F. Zhou, X. Liu, K. Zhang, and G. Trajcevski. Toward discriminating and synthesizing motion traces using deep probabilistic generative models. *IEEE Transactions on Neural Networks and Learning Systems*, 32(6):2401–2414, 2021.
- [275] F. Zhou, P. Wang, X. Xu, W. Tai, and G. Trajcevski. Contrastive trajectory learning for tour recommendation. *ACM Trans. Intell. Syst. Technol.*, 13(1):4:1–4:25, 2022.
- [276] F. Zhou, X. Xu, G. Trajcevski, and K. Zhang. A survey of information cascade analysis: Models, predictions, and recent advances. *ACM Comput. Surv.*, 54(2):27:1–27:36, 2022.
- [277] Y. Zhu, Y. Ye, S. Zhang, X. Zhao, and J. J. Q. Yu. Difftraj: Generating GPS trajectory with diffusion probabilistic model. In *NeurIPS*, 2023.
- [278] E. Zimányi, M. Sakr, and A. Lesuisse. Mobilitydb: A mobility database based on postgresql and postgis. *ACM Trans. Database Syst.*, 45(4), dec 2020.
- [279] A. Züfle, T. Anderson, and S. Gao. Introduction to the special issue on understanding the spread of COVID-19, part 1. *ACM Transactions on Spatial Algorithms and Systems*, 8(3):1–5, 2022.
- [280] A. Züfle, G. Trajcevski, D. Pfoser, M. Renz, M. T. Rice, T. Leslie, P. Delamater, and T. Emrich. Handling uncertainty in geo-spatial data. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pages 1467–1470. IEEE, 2017.