Research Statement

XU Yixin
School of Computing and Information Systems, Singapore Management University
Tel: (65) 6826 4809; Email: yixinxu@smu.edu.sg
01 July 2023

Background

The prevalence of mobile devices enables humans, vehicles, and almost all objects in the city to report and update their locations and status in real-time. The collected spatial-temporal data, with both spatial and temporal attributes, has great potential to be discovered and drives the growth of vital applications in our daily lives, such as real-time traffic management and location-based services. Nevertheless, processing and mining the increasingly available spatial-temporal data presents challenges due to its massive size and the continuous movement of objects.

My research addresses the problem of efficiently and effectively extracting useful information from the collected spatial-temporal data and applying it to real-world applications. My goal is to design algorithms to better store, index, and update the spatial-temporal data and therefore achieve efficient processing with minimal memory overhead.

Research Areas

Ride-sharing dispatching algorithm

Ride-sharing has emerged as a prevalent transportation mode to provide timely and convenient rides to passengers. A fundamental problem of ride-sharing is determining how to dispatch vehicles to passengers in a short time, which is challenging due to the large number of involved participants, highly dynamic scenarios, and expensive road network distance computation.

I studied efficient and scalable algorithms to enable fast and high-quality dispatching in ride-sharing. I first proposed an efficient pruning algorithm to solve a fundamental problem in the dispatching process, i.e., quickly filtering out infeasible vehicles for passengers. The key insight is that the area that a vehicle may possibly visit is restricted within a set of geometric objects (ellipses or circles). Computing and indexing these geometric objects is quick and simple. Thus, it can largely reduce the computation time and updating costs for moving vehicles. The experiments on real-world data show that our algorithm prunes an order of magnitude more vehicles and reduces the update time by two to three orders of magnitude compared to state-of-the-art algorithms.

Furthermore, I studied multi-hop ride-sharing, which significantly enhances the dispatching quality and flexibility by enabling transfers between vehicles. We first proposed an algorithm that finds optimal multi-hop matches, which was demonstrated to be two orders of magnitude faster than the state-of-the-art algorithms in our experiments. To further enhance the efficiency, we employed approximation strategies such as using deep learning to predict the possible travel areas of vehicles.

The approximation strategies achieved another order of magnitude speed-up while maintaining a comparable level of matching quality in the system.

Spatial-temporal data query processing

The spatial-temporal data queries are crucial to location-based services. The goal of such queries is to find objects that satisfy specific spatial or temporal relationships. Examples of these queries include finding the nearest ATM or finding supermarkets within 1 km.

I studied an important spatial-temporal data query named all nearest neighbor query (ANN). This query aims to identify the nearest neighbor (e.g., car park) of each query object (e.g., cars) over a road network. Instead of applying an individual nearest neighbor query on each query object, which can be expensive and involve redundant computation, we proposed an algorithm that finds the nearest neighbors with only one traversal over the road network. The proposed algorithm outperforms the state-of-the-art in terms of query cost and memory consumption by orders of magnitude.

Selected Publications and Outputs

Yixin Xu, Jianzhong Qi, Renata Borovica-Gajic, Lars Kulik. GeoPrune: Efficiently Matching Trips in Ridesharing Through Geometric Properties, *International Conference on Scientific and Statistical Database Management* (SSDBM), pp. 1-12, 2020.

Yixin Xu, Lars Kulik, Renata Borovica-Gajic, Abdullah Aldwyish, Jianzhong Qi. Highly Efficient and Scalable Multi-hop Ride-sharing, *International Conference on Advances in Geographic Information Systems* (SIGSPATIAL), pp. 215-226, 2020.

Yixin Xu, Jianzhong Qi, Renata Borovica-Gajic, Lars Kulik. Finding All Nearest Neighbors with a Single Graph Traversal, *International Conference on Database Systems for Advanced Applications* (DASFAA), pp. 221-238, 2018.