

Towards ML-enhanced Trajectory Query Processing

Guang Yang

The rapid growth in location-based services has generated massive movement data. These movement data, or trajectories, are crucial in numerous scientific research and commercial applications. The increased query needs and data volume have put modern databases under pressure. They face two significant challenges in terms of scalability and flexibility.

When it comes to scalability, low-latency services (e.g., Uber ride find and Google Maps route search) demand real-time query processing whilst processing multiple queries concurrently. On the other hand, flexibility is key for analytics in modern applications where systems have to process complex queries with constraints. The most challenging part of the problem is that modern databases cannot deal with these two challenges simultaneously. Therefore, addressing these two challenges together in large-scale trajectory data management necessitates research on optimized data structures, particularly in index structures, which are crucial to query performance.

In our work, we propose a modular approach using hybrid indexing to tackle the two challenges from two levels: (1) Global scalability point of view (across machines) and (2) Local efficiency point of view (within each machine). Together, these two levels form a comprehensive view of the modular approach to enable efficient complex trajectory data processing. Specifically, we introduce the concept of machine learning (ML)-enhanced indexes or learned indexes into trajectory indexing. However, directly applying learned indexes introduces challenges since they cannot preserve the continuity of trajectories or cope with the dynamic nature of trajectories (requiring continuous updates).

From our insight into trajectory update patterns, where spatial updates shift the data locations within a predefined region while the temporal dimension continuously grows and expires, we propose to use strategies at global and local levels. At the global level, We propose to optimize traditional indexing techniques to improve updates in spatial dimensions. This allows the data structure to be more robust to data distribution shifts across multiple machines and prevent constant retraining of the learned models. At the local level, we optimize for efficient updates and queries in the rapidly evolving temporal dimension, which exhibits a characteristics of a streaming workload. We designed learned indexes exploits the data distribution’s predictability to reduce the indexes’ latency and storage size.

We present a comprehensive view of our design approach by showcasing solutions to trajectory problems with different requirements. Specifically, at the global level, we developed a distributed approach for efficiently answering distance joins while preserving trajectory continuity and a framework for in-network range queries leveraging sensor and data distribution to ensure privacy. We then optimize each machine by transforming the temporal component of trajectories as a stream indexing problem. Two updatable and parallelizable learned indexes were developed to handle the temporal streams under various query workloads and update characteristics.