Succinct Trit-array Trie for Scalable Trajectory Similarity Search

Shunsuke Kanda
RIKEN AIP
shunsuke.kanda@riken.jp

Keisuke Fujii
Nagoya University and RIKEN AIP
fujii@i.nagoya-u.ac.jp

Koh Takeuchi
Kyoto University and RIKEN AIP
takeuchi@i.kyoto-u.ac.jp

Yasuo Tabei
RIKEN AIP
yasuo.tabei@riken.jp

ABSTRACT

Massive datasets of spatial trajectories representing the mobility of a diversity of moving objects are ubiquitous in research and industry. Similarity search of a large collection of trajectories is indispensable for turning these datasets into knowledge. Locality sensitive hashing (LSH) is a powerful technique for fast similarity searches. Recent methods employ LSH and attempt to realize an efficient similarity search of trajectories; however, those methods are inefficient in terms of search time and memory when applied to massive datasets. To address this problem, we present the *trajectory-indexing succinct trit-array trie* (tSTAT), which is a scalable method leveraging LSH for trajectory similarity searches. tSTAT quickly performs the search on a tree data structure called *trie*. We also present two novel techniques that enable to dramatically enhance the memory efficiency of tSTAT. One is a node reduction technique that substantially omits redundant trie nodes while maintaining the time performance. The other is a space-efficient representation that leverages the idea behind *succinct data structures* (i.e., a compressed data structure supporting fast data operations). We experimentally test tSTAT on its ability to retrieve similar trajectories for a query from large collections of trajectories and show that tSTAT performs superiorly in comparison to state-of-the-art similarity search methods.

CCS CONCEPTS

• **Information systems** → **Data mining**; **Data structures**; *Geographic information systems*.

KEYWORDS

Trajectory data mining, scalable similarity search, Fréchet distance, succinct data structures

ACM Reference Format:

Shunsuke Kanda, Koh Takeuchi, Keisuke Fujii, and Yasuo Tabei. 2020. Succinct Trit-array Trie for Scalable Trajectory Similarity Search. In *28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '20)*, November 3–6, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3397536.3422210>

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

SIGSPATIAL '20, November 3–6, 2020, Seattle, WA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8019-5/20/11...\$15.00

<https://doi.org/10.1145/3397536.3422210>

1 INTRODUCTION

With advances in location-acquisition technology and mobile computing, spatial trajectory data representing the mobility of a diversity of moving objects such as people, vehicles, and animals are ubiquitous in research and industry [58]. For example, a city monitoring system with the global positioning system (GPS) enables us to record a huge number of complex trajectories from vehicles [3]. A camera-based tracking system such as SportVU also enables us to collect a large number of fast and dynamic movements of sports players precisely [1]. There is, therefore, a strong and growing demand for developing new powerful methods to make the best use of huge collections of trajectories toward data mining applications such as prediction of a taxi demand in a city in the near future and development of an effective game strategy in sports.

Searching for similar trajectories in a large collection for a given query is an example of effective use of huge collections of trajectories, and it has a wide variety of applications from route recommendation [34, 45] to sports play retrieval and clustering [24, 43]. Several distance measures for trajectories such as Fréchet distance, Hausdorff distance, dynamic time warping (DTW), and longest common subsequence (LCSS) [49] have been proposed thus far. Fréchet distance is the *defacto* standard measure [4] to evaluate the similarity of trajectories. Fréchet distance can be intuitively explained as the length of the shortest leash enabling a person to walk a dog such that they move along each of their own trajectories at their own speed without going back. Fréchet distance has several advantages compared to other distance measures. For example, Fréchet distance captures the directions of trajectories in contrast to Hausdorff distance, and it satisfies the distance conditions in contrast to DTW and LCSS. Owing to those advantages, Fréchet distance has been successfully applied in various applications including detection of commuting patterns [11], handwriting recognition [48], protein structure alignment [53], and many other applications [13, 30, 59]. In addition, an algorithm competition focusing on similarity searches of trajectories using Fréchet distance was held in ACM SIGSPATIAL Cup 2017 [52]. While these applications of Fréchet distance show its practicality, similarity search of trajectories using Fréchet distance is computationally demanding because the computation time of the Fréchet distance between a pair of trajectories is quadratic to their length [9], which limits large-scale applications using Fréchet distance in practice.

Locality sensitive hashing (LSH) for Fréchet distance has been proposed for scalable similarity searches of trajectories using Fréchet distance. The first LSH method proposed by Indyk [25] approximates Fréchet distance via a product metric. Recently, more efficient

LSH methods with tighter approximation bounds have been proposed [6, 14, 15], which map trajectories into non-negative integer vectors (called *sketches*) such that Fréchet distance is preserved as the Hamming distance among sketches. LSH for Fréchet distance has been successfully applied to trajectory clustering [41, 42].

FRESH [14] is the state-of-the-art method applying LSH for Fréchet distance to similarity searches of trajectories in practice. *FRESH* uses an inverted index implemented by the hash table data structure, which stores values associated with each key. For a collection of sketches of fixed length L , *FRESH* builds L hash tables whose key is an integer at each position of sketches and value is the set of sketch identifiers with the integer. Given a query sketch of the same length L , *FRESH* retrieves the hash table with the integer at each position of the query and computes a set of sketch identifiers with the integer at each position, resulting in L sets of sketch identifiers in total for L positions of the query. The Hamming distances between each sketch in the set and the query are computed by sequentially counting the number of appearances of each integer in these sets. However, this approach suffers from performance degradation caused by the sequential counting of integers if large sets of identifiers are computed. In addition, *FRESH* consumes a large amount of memory for indexing large collections of trajectories because of the memory inefficiency of hash tables, which limits large-scale applications in practice. Thus, an important open challenge is to develop a memory-efficient data structure for fast similarity searches of sketches for trajectories.

Trie [19] is an ordered labeled tree data structure for a set of strings and supports various string operations such as string search and prefix search with a wide variety of applications in string processing. Examples are string dictionaries [28], n -gram language models [38, 39], and range query filtering [56]. A typical pointer-based representation of trie consumes a large amount of memory. Thus, recent researches have focused on space-efficient representations of trie (e.g., [28, 38, 39, 56]). To date, trie has been applied only to the limited application domains listed above. However, as we will see, there remains great potential for a wide variety of applications related to trajectories.

Our Contribution. We present a novel method called *trajectory-indexing succinct trit-array trie (tSTAT)* that efficiently performs a similarity search for trajectories with LSH for Fréchet distance. As in *FRESH*, multiple data structures according to the sketch length are built for efficient query retrievals. However, in contrast to *FRESH*, *tSTAT* enables faster similarity searches by effectively partitioning sketches into several blocks by the *multi-index approach* [23] and building the trie data structure for each block. *tSTAT*'s query retrieval is performed by an efficient trie traversal that bypasses the sequential counting of integers for computing Hamming distance. While preserving fast similarity searches, *tSTAT* successfully reduces the memory usage by presenting two novel techniques of *node reduction* and a space-efficient representation of trie by leveraging *succinct data structures* (i.e., compressed representations of data structures while supporting various data operations in the compressed format) [27, 37].

We experimentally test *tSTAT* on its ability to retrieve similar trajectories for a query from large collections of real-world trajectories. The performance comparison with state-of-the-art similarity

search methods for trajectories demonstrates that *tSTAT* performs superiorly with respect to time and memory efficiency.

2 PROBLEM SETTING

This section introduces discrete Fréchet distance [17] as a similarity measure for trajectories and then presents the similarity search problem on the Fréchet distance. Although *tSTAT* is presented with discrete Fréchet distance, it is also applicable to continuous Fréchet distance [4], as presented in [14]. In the remainder of this paper Fréchet distance is used for discrete Fréchet distance.

2.1 Fréchet Distance

A trajectory P of length m is a sequence of m coordinates $P = (p_1, p_2, \dots, p_m)$, where $p_i \in \mathbb{R}^d$ is a coordinate with d dimensions. Let $P = (p_1, p_2, \dots, p_{m_1})$ and $Q = (q_1, q_2, \dots, q_{m_2})$ be two trajectories of lengths m_1 and m_2 , respectively. A *traversal* τ is a sequence of coordinate pairs $(p_i, q_j) \in P \times Q$ such that (i) τ begins with (p_1, q_1) and ends with (p_{m_1}, q_{m_2}) , and (ii) $(p_i, q_j) \in \tau$ can be followed by one of (p_{i+1}, q_j) , (p_i, q_{j+1}) , or (p_{i+1}, q_{j+1}) . For the set of all possible traversals $\mathcal{T}$ for P and Q , the (discrete) Fréchet distance between P and Q is defined as

$$\text{Frec}(P, Q) = \min_{\tau \in \mathcal{T}} \max_{(p_i, q_j) \in \tau} \|p_i - q_j\|_2,$$

where $\|\cdot\|_2$ is the L_2 norm. The computation time for $\text{Frec}(P, Q)$ is $O(m_1 m_2)$ by the dynamic programming [17].

2.2 Similarity Search for Trajectories

We formulate the similarity search problem for trajectories with LSH for Fréchet distance [14]. Given a collection of n trajectories $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$, LSH for Fréchet distance projects each P_i in $\mathcal{P}$ into sketch S_i of fixed length L , i.e., $S_i = (s_{i1}, s_{i2}, \dots, s_{iL})$ for each $s_{ij} \in \{0, 1, \dots, \sigma - 1\}$, where σ is the alphabet size. That is, it produces a collection of n sketches $\mathcal{S} = \{S_1, S_2, \dots, S_n\}$. Given a query sketch T projected from a query trajectory Q and of the same length L , the task of similarity search is to retrieve from the collection $\mathcal{S}$ all the identifiers of sketches $\mathcal{I} = \{i_1, i_2, \dots, i_k\}$ such that the Hamming distance between S_{i_j} and T is no more than a threshold K , i.e., $\text{Ham}(S_{i_j}, T) \leq K$ for each $i_j \in \mathcal{I}$, where $\text{Ham}(\cdot, \cdot)$ denotes the Hamming distance between two sketches (i.e., the number of positions at which the corresponding integers between two sketches are different). We shall call the problem for computing $\mathcal{I}$ the *Hamming distance problem*.

$\mathcal{I}$ is an approximation set of solutions that can include two types of errors consisting of false positive and false negative: (i) false positive is identifier $i \in \mathcal{I}$ such that the Hamming distance between S_i and T is no more than K (i.e., $\text{Ham}(S_i, T) \leq K$), but the Fréchet distance between P_i and Q is more than threshold R (i.e., $\text{Frec}(P_i, Q) > R$) and (ii) false negative is identifier $i \notin \mathcal{I}$ such that the Hamming distance between S_i and T is more than K (i.e., $\text{Ham}(S_i, T) > K$), but the Fréchet distance between P_i and Q is no more than R (i.e., $\text{Frec}(P_i, Q) \leq R$).

False positives in $\mathcal{I}$ can be filtered out by computing the Fréchet distance between P_i for each $i \in \mathcal{I}$ and Q , resulting in solution set $\mathcal{I}' \subseteq \mathcal{I}$. False negatives for $\mathcal{I}$ can be reduced by setting K to be large, but such a setting can increase the search time.

Table 1: Summary of similarity search methods, where n is the number of trajectories (or sketches) in a collection, L is the fixed length of sketches, σ is the alphabet size, B is the number of blocks in the multi-index ($B \leq L$), K is the Hamming distance threshold, and N_{in} is the number of internal nodes in tSTAT. C_{fresh} , C_{hms} and C_{tstat} are the verification times for candidates obtained in FRESH, HmSearch and tSTAT, respectively. The search times are obtained assuming that sketches are uniformly distributed.

Method	Data structure	Memory usage (bits)	Search time
FRESH [14]	Hash table	$O(Ln \log n + Ln \log \sigma)$	$O(L \cdot \max(1, n/\sigma)) + C_{\text{fresh}}$
HmSearch [57]	Hash table	$O(Ln \log n + Ln \log \sigma)$	$O(L \cdot \max(1, Ln/\sigma^{L-2L/K+1})) + C_{\text{hms}}$
tSTAT (this study)	STAT	$O(\sigma N_{\text{in}} + Bn \log n + Ln \log \sigma)$	$O(B(L/B)^{K/B+2}) + C_{\text{tstat}}$

3 LITERATURE REVIEW

Several efficient data structures have been proposed for similarity search over sketches in the research areas of trajectory data analysis and data mining. We briefly review the state-of-the-art methods, which are also summarized in Table 1.

3.1 Related Work on Trajectory Data Analysis

For car trajectories, a large number of similarity search methods have been proposed and most of them employ map matching that projects such trajectories into a sequence of roads [31, 47, 51, 55]. While the applicability of the map matching is limited to car trajectories, there are important similarity searches for several other types of trajectories.

There are several existing methods for trying to compute exact solutions for trajectory similarity searches. In ACM SIGSPATIAL Cup 2017 [52], similarity search methods for Fréchet distance were developed [7, 12, 16]. Recently, Bringmann et al. [10] improved the winning algorithm [7]. In another line of research, several similarity search methods for distributed query processing were proposed [46, 54]. Those methods commonly build spatial indexes (e.g., R-trees or KD-trees) using spatial points and minimum bounding rectangles of trajectories and solve similarity searches with spatial indexes. However, owing to the computational demand of Fréchet distance, designing an efficient data structure for its fast similarity searches is difficult. Thus, recent researches have focused on approximate similarity searches using LSH for Fréchet distance because the LSH enables us to solve the trajectory similarity search problem as the Hamming distance problem for sketches.

The first data structure for approximate nearest neighbor searches for trajectories was proposed by Indyk [26]. While the search time is $O(m^{O(1)} \log n)$ for the number of trajectories n and the maximum length of trajectory m , it consumes a large amount of memory in $O(|\mathcal{A}| \sqrt{m} (m \sqrt{m} n)^2)$ words for the size of domain $|\mathcal{A}|$ on which the trajectories are defined, which limits practical applications of the data structure. Later, improved data structures [6, 15] were presented, and they take $O(n \log n + nm)$ memory words while performing a similarity search in $O(m \log n)$ time.

FRESH [14] is a practical data structure for approximate similarity search of trajectories, and it uses L inverted indexes with hash tables. The search time of FRESH is $O(L \cdot \max(1, n/\sigma)) + C_{\text{fresh}}$ assuming that the hash table has a uniform distribution, where C_{fresh} is the verification time for candidates. The memory usage is $O(Ln \log n + Ln \log \sigma)$ bits.

3.2 Related Work on Data Mining

In data mining, recent similarity search methods for sketches use the *multi-index* approach [23] as follows. The approach partitions sketches in $\mathcal{S}$ into several blocks of short sketches and builds multiple inverted indexes from the short sketches in each block. The similarity search consists of two steps: filtering and verification. The filtering step roughly obtains candidates of similar sketches by querying the inverted indexes. The verification step removes false positives from these candidates by computing the Hamming distances.

Several data structures with the multi-index approach were presented specifically for similarity searches on binary sketches (i.e., $\sigma = 2$) [21, 36, 40]. These methods are not directly applicable to similarity searches on integer sketches (i.e., $\sigma > 2$). HmSearch [57] is an efficient similarity search with the multi-index approach for integer sketches. It generates sketches similar to the ones in $\mathcal{S}$ and adds them to $\mathcal{S}$, which enables fast similarity searches with hash tables. The search time is $O(L \cdot \max(1, Ln/\sigma^{L-2L/K+1})) + C_{\text{hms}}$ assuming that the sketches in $\mathcal{S}$ are uniformly distributed, where C_{hms} is the verification time for candidates. HmSearch needs a large memory consumption of $O(Ln \log n + Ln \log \sigma)$ bits because of the amplification of the sketches in $\mathcal{S}$. Although several similarity search methods applicable to integer sketches have been proposed (e.g., [32, 33, 35]), HmSearch performs the best [57].

Despite the importance of a scalable similarity search for massive trajectories, no previous work achieving fast search and memory efficiency in trajectory data analysis and data mining exists. As reviewed in this section, the current best possible methods are (i) to use L inverted indexes with hash tables or (ii) to use the multi-index approach with hash tables. However, those methods scale poorly due to the large computation cost of the search algorithm and memory inefficiency of the data structure.

4 TRAJECTORY-INDEXING SUCCINCT TRIT-ARRAY TRIE

tSTAT proposes to use an efficient approach called multi-index [23] for solving the similarity searches in the Hamming space and leverages a trie [19] for implementing the multi-index approach. tSTAT efficiently computes Hamming distance by traversing tries. A straightforward implementation of tries using pointers is not memory efficient for indexing a large collection of sketches. Instead, we present a novel node reduction technique and memory-efficient representation of trie called *succinct trit-array trie* (STAT).

	S^1	S^2	$\text{Ham}(S_i, T)$
S_1	1 2 1 2	1 1 2 3	2
S_2	2 2 1 1	1 1 2 2	5
S_3	2 3 0 3	1 3 0 3	7
S_4	1 2 1 2	0 0 0 1	2
S_5	1 0 1 0	1 0 2 3	3
S_6	2 3 0 3	1 0 2 3	5
T	T^1 1 2 1 2	T^2 0 0 2 3	

Figure 1: Collection of six sketches S (respectively, query T) of $L = 8$ and $\sigma = 4$, partitioned into two blocks S^1 and S^2 (respectively, T^1 and T^2). Each integer in S mismatched to the corresponding integer in T are colored in red.

4.1 Multi-Index Approach

The multi-index approach builds multiple indexes from collection S for solving the similarity search problem of sketches on the Hamming space. Each sketch $S_i \in S$ is partitioned into B disjoint blocks $S_i^1, S_i^2, \dots, S_i^B$ of fixed length L/B . For convenience, we assume that L is divisible by B . The j -th block of S for $j = 1, 2, \dots, B$ is denoted by $S^j = \{S_1^j, S_2^j, \dots, S_n^j\}$. Figure 1 shows an example collection of six sketches S , and S is partitioned into two blocks S^1 and S^2 .

For the similarity search for query T , we first partition T into B blocks $T^1, T^2, \dots, T^B$ in the same manner. We then assign smaller Hamming distance thresholds $K^1, K^2, \dots, K^B$ to each block, which is detailed later. We obtain the candidate sets of sketch ids $C^j = \{i \mid \text{Ham}(S_i^j, T^j) \leq K^j\}$ for each block $j = 1, 2, \dots, B$ and compute the union of the candidate sets as $C = C^1 \cup C^2 \cup \dots \cup C^B$. Finally, we remove false positives by computing $\text{Ham}(S_i, T)$ for each $i \in C$, resulting in solution set $\mathcal{I}$.

The threshold assignment requires K^j such that each solution in $\mathcal{I}$ is included in C^j for any $j = 1, 2, \dots, B$. Such assignment is satisfied by setting K^j such that $K^1 + K^2 + \dots + K^B = K - B + 1$, which is ensured by the general pigeonhole principle [40].

For the similarity search with $K = 3$ in Figure 1, we assign $K^1 = 1$ and $K^2 = 1$ such that $K^1 + K^2 = K - B + 1 = 2$ based on the general pigeonhole principle. Then, $C^1 = \{1, 4\}$ and $C^2 = \{5, 6\}$ are obtained. Since $\mathcal{I} = \{1, 4, 5\}$ is a subset of $C = \{1, 4, 5, 6\}$, we can obtain $\mathcal{I}$ by verifying each element in C .

4.2 Trie Implementing Multi-Index

Trie $\mathcal{X}^j$ is an edge-labeled tree indexing sketches in S^j (see Figure 2). Each node in $\mathcal{X}^j$ is associated with the common prefix of a subset of sketches in S^j , and the root (i.e., node without a parent) is not associated with any prefix. Each leaf (i.e., node without any children) is associated with input sketches of the same integers and has the list of their sketch ids. All outgoing edges of a node are labeled with distinct integers. For a node u in $\mathcal{X}^j$, the number of edges from the root to u is the level of u . Figure 2 shows an example of $\mathcal{X}^1$ and $\mathcal{X}^2$ built from S^1 and S^2 in Figure 1, respectively.

Searching similar sketches for query sketch T^j and threshold K^j is performed by traversing trie $\mathcal{X}^j$ in a depth-first manner. We start to traverse nodes from the root at level 0. At each level ℓ , we compare the ℓ -th integer of T^j with labels associated with outgoing

Figure 2: Tries $\mathcal{X}^1$ and $\mathcal{X}^2$ built from S^1 and S^2 of Figure 1, respectively. In each traversed node u , dist_u to T of Figure 1 is denoted by red/blue numbers when $K^1 = 1$ and $K^2 = 1$.

edges from nodes at level ℓ . For each node u at level ℓ , we compute the Hamming distance dist_u between the sequence of edge labels from the root to u and the subsequence from the first position to the ℓ -th position in T^j . If dist_u becomes more than K^j at each reached node u , we can safely stop traversing down to all the descendants under node u . The candidate set C^j can be obtained by accessing the list of sketch ids at each reached leaf v , which means $\text{dist}_v \leq K^j$.

In Figure 2, dist_u for each node u and query T^1 or T^2 is represented in each node. When $K^1 = 1$ and $K^2 = 1$, we stop traversing at each node including a red number. We obtain candidates $C^1 = \{1, 4\}$ and $C^2 = \{5, 6\}$, associated with the reached leaves.

The algorithm can prune unnecessary portions of the search space depending on K^j . The number of traversed nodes in $\mathcal{X}^j$ is bounded by $O((L/B)^{K^j+2})$ when assuming the complete σ -ary trie [5]. When we assign K^j in a round robin manner, it holds $K^j \leq K/B$; thus, the algorithm takes $O((L/B)^{K/B+2} + |C^j|)$ time for $\mathcal{X}^j$.

When we represent trie $\mathcal{X}^j$ by using pointers straightforwardly, the scalability becomes a critical issue since the memory usage is $O(N^j \log(\sigma N^j))$ bits. In the remaining subsections, we present two techniques of node reduction and STAT data structure for compactly representing $\mathcal{X}^j$.

4.3 Node Reduction

A crucial observation for candidate set C is that it allows false positives for solution set $\mathcal{I}$ for the Hamming distance problem. This observation enables us to make tries more memory efficient by removing their redundant nodes.

The *weight* of node u is defined as the total number of sketch ids associated with leaves of the subtree with root u and satisfies an *antimonotonicity*, i.e., the weight of node u is no less than the one of u 's child. The weight of the root in trie $\mathcal{X}^j$ is n .

Our node reduction is a depth-first traversal leveraging the antimonotonicity. The algorithm starts from the root and visits each node. If the weight of node u is no more than hyper-parameter λ , the algorithm eliminates the subtree with root u from the tree. Then, u becomes a new leaf of the tree and has the sketch ids of previous leaves associated with the subtree with root u .

Figure 3 shows an example of node reductions with $\lambda = 1$ and $\lambda = 2$ for trie $\mathcal{X}^2$ in Figure 2. The trie with 16 nodes in Figure 2b is reduced to the tries with 11 nodes in Figure 3a and 6 nodes in

Figure 3: Trie $\mathcal{X}^2$ after node reduction in Figure 2.

Figure 3b by the node reduction algorithm with $\lambda = 1$ and $\lambda = 2$, respectively.

λ is a trade-off parameter that can control the balance between the tree traversal time and the verification time for removing false positives. The larger λ becomes, the smaller the number of nodes becomes and the larger the number of false positives becomes, resulting in a smaller tree traversal time and a larger verification time. This trade-off is investigated with various λ in Section 5.

4.4 Succinct Rank and Select Data Structures

STAT is an integer array representation of trie leveraging succinct Rank and Select data structures [27]. Given an integer array A of length M , the data structures support the following operations in compressed space:

- $\text{Rank}_c(A, i)$ returns the number of occurrences of integer c between positions 0 and $i - 1$ on A .
- $\text{Select}_c(A, i)$ returns the position in A of the i -th occurrence of integer c ; if i exceeds the number of c 's in A , it always returns M .

For $A = (0, 1, 0, 2, 1, 1, 2, 0)$, $\text{Rank}_0(A, 3) = 2$ and $\text{Select}_2(A, 1) = 6$.

For a bit array A consisting of integers from $\{0, 1\}$, A can be stored in $M + o(M)$ bits of space while supporting the operations in $O(1)$ time [27, 50]. For a trit array A consisting of integers from $\{0, 1, 2\}$, A can be stored in $M \log_2 3 + o(M)$ bits of space while supporting the operations in $O(1)$ time [18, 37].

4.5 STAT Data Structure

A basic idea behind the STAT data structure is to represent a trie using three arrays each of which represents children of a node, the number of sketch ids associated with leaves, and the sketch ids at each level. While the array representing children of a node consists of integers from $\{0, 1, 2\}$ and is indexed by the succinct rank and select data structure for trits, the array representing the number of sketch ids associated with leaves consists of integers from $\{0, 1\}$ and is indexed by the succinct rank and select data structure for bit arrays. Detailed explanation of the STAT data structure is subsequently presented.

STAT consists of three arrays H_ℓ , G_ℓ , and V_ℓ for compactly representing nodes at level ℓ of trie $\mathcal{X}^j$ (see Figure 4). To define the arrays, we introduce the order of internal nodes and leaves to each level of the trie as follows. We traverse $\mathcal{X}^j$ in a breadth-first order. The i -th internal node at a level is the i -th encountered node for internal nodes at the level. Similarly, the i -th leaf at a level is the

Figure 4: STAT for trie $\mathcal{X}^2$ in Figure 3a. The orders of internal nodes and leaves at each level are denoted by yellow and green numbers, respectively.

i -th encountered node for leaves at the level. Ids starting at zero are assigned to nodes. For example, level 2 of $\mathcal{X}^2$ in Figure 4 consists of the zeroth internal node, the first internal node, and the zeroth leaf from top to bottom.

H_ℓ is a trit (i.e., ternary digit) array and represents children of internal nodes at level ℓ . The children of an internal node is represented in a trit array of length σ , and H_ℓ is constructed by concatenating such trit arrays in the order of internal nodes. That is, the subarray of H_ℓ between positions $i\sigma$ and $(i+1)\sigma - 1$ corresponds to the i -th internal node at level ℓ . The data structure is as follows:

- $H_\ell[i\sigma + c] = 0$ if the i -th internal node at level ℓ does not have a child indicated with edge label c in the trie,
- $H_\ell[i\sigma + c] = 1$ if the i -th internal node at level ℓ has a child indicated with edge label c as an internal node, and
- $H_\ell[i\sigma + c] = 2$ if the i -th internal node at level ℓ has a child indicated with edge label c as a leaf.

H_ℓ is implemented by the Rank data structure for trits. Let us consider the zeroth internal node at level 1 in Figure 4. $H_1[0] = 1$ and $H_1[1] = 1$ because it has two children as internal nodes indicated with edge labels 0 and 1; $H_1[2] = 0$ because it does not have a child indicated with edge label 2; and $H_1[3] = 2$ because it has a child as a leaf indicated with edge label 3.

G_ℓ is a bit array representing the numbers of sketch ids associated with leaves at level ℓ in unary encoding. That is, for g sketch ids associated with a leaf, the number g is encoded into $(g - 1)$ 0s following 1. G_ℓ is constructed by concatenating the unary codes from the zeroth leaf at level ℓ and is implemented by the Select data structure for bits [27]. In Figure 4, $G_4 = (1, 0, 1, 1)$ because the zeroth, first, and second leaves at level 4 have two, one, and one sketch ids, respectively.

V_ℓ is an array of sketch ids associated with leaves at level ℓ and is constructed by concatenating the sketch ids from the zeroth leaf at level ℓ . In Figure 4, $V_4 = (5, 6, 2, 1)$ because the zeroth, first,

and second leaves at level 4 have sketch ids (5, 6), (2), and (1), respectively.

We present operations for computing children for a given internal node and sketch ids associated with a given leaf. Those operations are necessary for computing Hamming distances for a query on trie X^j . Given the i -th internal node at level ℓ and integer c , there is not a child indicated with edge label c if $H_\ell[i\sigma + c] = 0$; If $H_\ell[i\sigma + c] = 1$, the child is the i' -th internal node at level $\ell + 1$, where $i' = \text{Rank}_1(H_\ell, i\sigma + c)$; If $H_\ell[i\sigma + c] = 2$, the child is the i' -th leaf node at level $\ell + 1$, where $i' = \text{Rank}_2(H_\ell, i\sigma + c)$. Given the i -th leaf node at level ℓ , the associated sketch ids are the elements between s and e in V_ℓ , where $s = \text{Select}_1(G_\ell, i)$ and $e = \text{Select}_1(G_\ell, i + 1) - 1$.

In Figure 4, for the zeroth internal node at level 1, the child indicated with edge label 1 is the first internal node at level 2 because $H_1[1] = 1$ and $\text{Rank}_1(H_1, 1) = 1$; the child indicated with edge label 3 is the zeroth leaf at level 2 because $H_1[3] = 2$ and $\text{Rank}_2(H_1, 3) = 0$. For the zeroth leaf at level 4, the associated sketch ids are elements $V_4[0] = 5$ and $V_4[1] = 6$ because $\text{Select}_1(G_4, 1) = 0$ and $\text{Select}_1(G_4, 1) - 1 = 1$.

Analysis for Memory Efficiency. Let N_{in}^j denote the number of internal nodes in trie X^j . The total length of H_ℓ is σN_{in}^j , and the total memory usage of H_ℓ is $\sigma N_{\text{in}}^j \log_2 3 + o(\sigma N_{\text{in}}^j) = O(\sigma N_{\text{in}}^j)$ bits. The total length of G_ℓ (or V_ℓ) is n . The total memory usage of G_ℓ and V_ℓ is $n + o(n) + n \log_2 n = O(n \log n)$ bits.

The *information-theoretic lower bound (ITLB)* is the defacto standard criteria for investigating the memory efficiency of data structures and is defined as the minimum memory usage for representing tries. We analyze the memory efficiency of STAT by comparing the memory usage of STAT with that of ITLB for tries.

We ignore G_ℓ and V_ℓ in the comparison because the arrays are independent from trie data structures. The memory usage of STAT for a trie with N_{in}^j internal nodes is $\sigma N_{\text{in}}^j \log_2 3 + o(\sigma N_{\text{in}}^j)$ bits. ITLB for a trie with N^j nodes is $N^j(\sigma \log_2 \sigma - (\sigma - 1) \log_2(\sigma - 1))$ bits [8]. Roughly, STAT becomes smaller than ITLB if

$$\frac{\sigma \log_2 3}{\sigma \log_2 \sigma - (\sigma - 1) \log_2(\sigma - 1)} < \frac{N^j}{N_{\text{in}}^j}.$$

For example, when $\sigma = 256$, STAT becomes smaller if $43 < N^j/N_{\text{in}}^j$. The comparison shows that STAT is efficient for representing tries with large N^j/N_{in}^j .

N^j/N_{in}^j is related to the average number of children for each internal node, i.e., the *load factor* of H_ℓ . Since nodes at a deeper level in a trie have fewer children, STAT is more effective for representing shallow tries. Thus, applying node reduction to a trie for eliminating trie nodes at deep levels enhances the space efficiency of STAT, as demonstrated in Section 5.

4.6 Complexity Analysis

The space complexity of tSTAT is derived as follows. Let $N_{\text{in}} = \sum_{j=1}^B N_{\text{in}}^j$. B STATs are represented in $O(\sigma N_{\text{in}} + Bn \log n)$ bits of space, where the left term is for H_ℓ and the right term is for G_ℓ and V_ℓ . In addition, we need to store collection $\mathcal{S}$ in $O(Ln \log \sigma)$ bits of space for verifying C .

Table 2: Statistics of datasets. Min, Max, Mean, and Median regarding trajectory length are presented.

Dataset	Number	Min	Max	Mean	Median
Taxi	1,704,769	1	3,881	48.9	41
NBA	3,288,264	1	898	85.3	73
OSM	19,113,525	1	2,000	13.3	7

All the methods in Table 1 require $O(Ln \log \sigma)$ bits of space to store $\mathcal{S}$. In addition to the space, tSTAT uses $O(\sigma N_{\text{in}} + Bn \log n)$ bits of space although FRESH and HmSearch use $O(Ln \log n)$ bits of space. The factor of $O(n \log n)$ is obviously large for a massive collection with large n . tSTAT can relax the large factor to B/L . Although tSTAT needs $O(\sigma N_{\text{in}})$ bits of space, N_{in} can be reduced by the node reduction approach.

The time complexity of tSTAT is derived by simply assuming $\lambda = 0$. The algorithm consists of traversing nodes and verifying candidates of C . The traversal time is $O(B(L/B)^{K/B+2})$ as presented in Section 4.2. The verification time C_{tstat} contains the times of removing duplicated candidates in $C^1, C^2, \dots, C^B$ and verifying candidates in C .

5 EXPERIMENTS

We evaluated the performance of tSTAT through real-world trajectory datasets. We used three datasets with $d = 2$, as shown in Table 2. *Taxi* is 1.7 million trajectories of 442 taxis driving in the city of Porto for one year [3]. The GPS locations (i.e., latitude and longitude) were recorded every 15 seconds with mobile data terminals installed in the taxis. Each trajectory represents a taxi trip taken by a passenger. *NBA* is 3.3 million trajectories of 636 games in the 2015/16 NBA seasons [1]. Ten player locations were captured every 40 milliseconds by SportVU player tracking system. Each trajectory is segmented by stationary points (i.e., small moving distances). *OSM* is 19 million trajectories of various moving objects including cars, bikes and humans traced by GPS in Western United States, and it is created from OpenStreetMap project [2]. OSM is the largest dataset of 19 million trajectories among three datasets and enables us to evaluate the performance of similarity search methods on a huge dataset.

For each dataset, we sampled 1,000 trajectories as queries and excluded them from the collection. We selected Fréchet distance thresholds R such that 1, 10, and 100 solutions are found on average per query, respectively, which resulted in $R = 567, 2720, 7263$ for Taxi, $R = 0.15, 0.26, 0.45$ for NBA, and $R = 899, 2506, 7824$ for OSM. LSH for Fréchet distance has two parameters of hash functions δ and k [14]. Following the original paper [14], those parameters were set to $\delta = 8dR$ and $k = 1$.

We conducted all experiments on one core of quad-core Intel Xeon CPU E5-2680 v2 clocked at 2.8 GHz in a machine with 256 GB of RAM, running the 64-bit version of CentOS 6.10 based on Linux 2.6. We implemented all data structures in C++17 and compiled the source code with g++ (version 7.3.0) in optimization mode -O3. We implemented succinct Rank and Select data structures on bits using the *succinct data structure library* [20]. As far as we know, there is no any available implementation on trits; therefore, we

Figure 5: Precision-recall curve for varying Hamming distance threshold K under fixed L and σ on each dataset.

developed a practical implementation of the data structures for trits. For a fast computation of Hamming distance on integers, we used an efficient computation algorithm exploiting a bit-parallelism technique [57]. The implementation details are presented in Appendix A. The source code used in our experiments is available at <https://github.com/kampersanda/frechetsimsearch>.

5.1 Performance of LSH

We analyzed the performance of LSH for Fréchet distance while varying L and σ . All the combinations of $L = 32, 64, 128$ and $\sigma = 2^1, 2^8, 2^{32}$ were tested. Setting σ to $2^1, 2^8$ and 2^{32} means that the resulting sketches are binary, byte, and four-byte strings, respectively. For each combination of L and σ , we computed recalls and precisions by varying Hamming distance threshold K from 0 to $L/4$ with step width $L/32$.

Figure 5 shows recall-precision curves on Taxi and NBA. Each recall value (or precision value) was averaged per query. Overall, the larger Hamming distance thresholds were used, the larger the recall values became and the smaller the precision values became on each dataset. Under each fixed L , the recall and precision values for $\sigma = 2^8$ were better than those for $\sigma = 2^1$ on each dataset while the recall and precision values for $\sigma = 2^8$ were almost the same as those for $\sigma = 2^{32}$, which showed $\sigma = 2^8$ was reasonable for achieving high recall and precision values of similarity searches using LSH. Under each fixed σ , $L = 64$ achieved reasonable recall and precision values for similarity searches using LSH. In addition, $L = 64$ enables us to efficiently implement the fast algorithm for computing Hamming distance [57], which is applicable to not only tSTAT but also HmSearch and linear search (see Appendix A.2). Thus, the following experiments were performed using $\sigma = 2^8$ and $L = 64$.

Figure 6: Results of node reduction on NBA ($R = 0.45$).

5.2 Efficiency of Node Reduction and STAT Data Structure

We evaluated the efficiency of node reduction in Section 4.3 and STAT data structure in Section 4.5 in terms of the memory efficiency and search time of tSTAT for solving the Hamming distance problem. We measured (a) the number of internal nodes N_{in} , (b) the memory usage, (c) the number of candidates $|C|$, and (d) the average search time per query by testing parameters $\lambda = 0, 2, 8, 32, 128, 512$. Setting $\lambda = 0$ results in no node being eliminated by the node reduction, i.e., the original trie is represented by the STAT data structure. All the combinations of block numbers $B = 8, 16$ and Hamming distance thresholds $K = 2, 4, 6, \dots, 12, 14$ were tested. Figure 6 shows the experimental results on NBA for $R = 0.45$. The results with other parameter settings are presented in the full version of this paper [29].

Figures 6a and 6b show the number of internal nodes N_{in} and the memory usage of tSTAT, respectively. To observe the memory efficiency of tSTAT, Figure 6b also shows the ITLB estimated from the number of nodes $N = \sum_{j=1}^B N^j$. As λ grew, the number of internal nodes N_{in} and the memory usage were dramatically reduced. As shown in Section 4.5, since the value N/N_{in} increased, the memory usage of tSTAT approached the ITLB. The reduction

of the number of internal nodes and the memory usage converged at around $\lambda = 32$.

Figures 6c and 6d show the number of candidates $|C|$ and the search time, respectively. Since the search time was affected by the number of candidates, both figures showed a similar tendency. For $B = 8$ and $K \geq 8$, as λ grew, the number of candidates and search time also increased. However, the effect of λ for the number of candidates and search time also converged at around $\lambda = 32$. Those results showed that setting $\lambda = 8$ or 32 was beneficial for improving the search and memory efficiencies of tSTAT. For $\lambda = 8$, $R = 0.45$, and $B = 8$, tSTAT achieved an average search time of 5 milliseconds and a memory usage of 0.36 GiB. The next subsection shows the efficiency of tSTAT with $\lambda = 8$ and tSTAT can achieve a memory efficiency and fast search time.

5.3 Efficiency of tSTAT

We compared tSTAT with HmSearch [57], FRESH [14], and linear search (LS) in the computation of I' . LS is a strawman baseline that computes the Hamming distance between a sketch in $\mathcal{S}$ and a query one-by-one in $O(Ln)$ time and $O(Ln \log \sigma)$ bits of space. Regarding exact solutions (without LSH), one possible baseline would be the winner of ACM SIGSPATIAL Cup 2017 [7, 10]. We did not include the winner in our competitors, since the experimental comparison between the winner and FRESH has been done in [14] and FRESH often outperformed the winner.

We fixed $L = 64$. The parameters of LSH for Fréchet distance were set to $\sigma = 2^8$, $\delta = 8dR$, $k = 1$, and $K = 2, 4, 6, \dots, 12, 14$, which were used in tSTAT, LS, and HmSearch. Following the original paper [14], FRESH was tested with the following parameter settings: $\sigma = 2^{32}$, $\delta = 4dR$, $k = 1, 2, 4$, and $K = 31, 35, 39, \dots, 59, 63$. For a fair comparison between different methods, we measured the time performance and the memory usage for varying recall for solution set I' . The implementations of HmSearch and FRESH used in these experiments are downloadable from <https://github.com/kampersanda/hmsearch> and <https://github.com/Cecca/FRESH>, respectively. Figure 7 shows experimental results. A part of the experimental results are presented in Appendix B.

Figures 7a–7c show the search time for varying the recall up to 100% for each method. The result showed tSTAT was the fastest for most results and demonstrated the efficiency of our search algorithm on tries with the multi-index approach. FRESH was slow because the implementation takes $O(n)$ time for C_{fresh} . For over 90% recalls, tSTAT was at most 60×, 34×, and 1600× faster than FRESH on Taxi, NBA, and OSM, respectively. For over 90% recalls, tSTAT was at most 3.4×, 12×, and 2.3× faster than HmSearch on Taxi, NBA, and OSM, respectively.

We also evaluated the scalability of the methods by varying the collection size n . We randomly sampled 20% to 100% of the trajectories from each dataset. Figure 7d shows the fastest search times with more than 90% recall for each method on NBA. tSTAT was much faster than LS and FRESH for a wide range of collection sizes n and Fréchet distance thresholds R . While HmSearch was also fast for small R , the time efficiency was decreased for large R . This is because, as described in Section 3.2, HmSearch stores and retrieves many sketches additionally generated, resulting in amplifying the

number of candidate sketches that need to be verified (i.e., the cost C_{hms} in Table 1) for large R .

Figures 7e–7g show the memory usage for varying the recall up to 100% for each method. Various memories consumed by HmSearch were observed for different recalls because different data structures in HmSearch are built according to Hamming distance threshold K . Those results demonstrated a high memory efficiency of tSTAT. In fact, while tSTAT consumed only 2.4–2.7 GiB of memory on OSM, HmSearch and FRESH consumed 27–65 GiB and 50–65 GiB of memories, respectively.

Figure 7h shows the construction time for varying the recall up to 100% for each method on OSM. We measured the execution time for producing sketches from input trajectories and building the index from the sketches. tSTAT took 5.0–6.1 minutes in total while LS took 3.3 minutes. On the other hand, the time needed for only building STATs from sketches was only 1.7–2.8 minutes, which demonstrated fast algorithms for constructing STATs. HmSearch and FRESH took 32–57 minutes and 18–51 minutes, respectively, which were much larger than the construction times of tSTAT.

Considering the time and memory efficiency of tSTAT, the result demonstrated the feasibility of finding similar trajectories from large collections for a query, which would be beneficial to analyze trajectories that are similar to a query of interest.

Figure 8 shows an example of querying NBA trajectories using tSTAT with $R = 0.45$. Similar movements of NBA players for a short movement of Rajon Rondo in the match between Sacramento Kings and Oklahoma City Thunder on December 6, 2015 were retrieved from a large collection of 3.3 million trajectories in the NBA dataset. tSTAT successfully found similar movements of NBA players such as a movement of Tim Duncan in the match between San Antonio Spurs and Toronto Raptors on December 9, 2015. The result demonstrated an effectiveness of tSTAT for discovering similar movements of NBA players, which would be beneficial to analyze various movements in sports data analyses [24, 43, 44].

6 CONCLUSION

We presented tSTAT, a novel method for fast and memory-efficient trajectory similarity search under Fréchet distance. Experimental results on real-world large datasets demonstrated that tSTAT was faster and more memory-efficient than state-of-the-art similarity search methods.

Our method is applicable to other distance measures such as continuous Fréchet distance and dynamic time warping distance on trajectories, and several methods were presented for those distance measures [14, 15]. Thus, One important future work is to develop similarity searches for other distance measures by extending the idea behind tSTAT. This would be beneficial for users analyzing massive trajectories in research and industry.

ACKNOWLEDGMENTS

KF's work was supported by JSPS KAKENHI (Grant Numbers 19H04941 and 20H04075). We would like to thank Yoichi Sasaki for introducing succinct data structures on trits to us. We also would like to thank Giulio Ermanno Pibiri for useful comments.

Figure 7: Performance evaluation of tSTAT, LS, HmSearch, and FRESH. Each of charts (a)–(c) shows average search times per query in milliseconds (ms) for varying recalls. Chart (d) shows average search times per query in ms for varying the number of trajectories, and it demonstrates scalabilities of similarity search methods. Each of charts (e)–(g) shows memory usages in GiB for varying recalls. Chart (h) shows construction times in minutes for varying recalls. The charts except (d) are plotted in the logarithmic scale.

REFERENCES

- [1] NBA Advanced Stats. <https://stats.nba.com>.
- [2] OpenStreetMap Data Extracts. <http://download.geofabrik.de/index.html>.
- [3] Taxi Service Trajectory (TST) Prediction Challenge 2015. <http://www.geolink.pt/ecmlpkdd2015-challenge/index.html>.
- [4] Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.*, 5(01n02):75–91, 1995.
- [5] Abdullah N Arslan and Ömer Eğecioğlu. Dictionary look-up within small edit distance. In *COCOON*, pages 127–136, 2002.
- [6] Maria Astefanoaei, Paul Cesaretti, Panagiota Katsikouli, Mayank Goswami, and Rik Sarkar. Multi-resolution sketches and locality sensitive hashing for fast trajectory processing. In *SIGSPATIAL*, pages 279–288, 2018.
- [7] Julian Baldus and Karl Bringmann. A fast implementation of near neighbors queries for Fréchet distance (GIS Cup). In *SIGSPATIAL*, page 99, 2017.
- [8] David Benoit, Erik D. Demaine, J. Ian Munro, Rajeev Raman, Venkatesh Raman, and S. Srinivasa Rao. Representing trees of higher degree. *Algorithmica*, 43(4):275–292, 2005.
- [9] Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *FOCS*, pages 661–670, 2014.
- [10] Karl Bringmann, Marvin Künnemann, and André Nusser. Walking the dog fast in practice: Algorithm engineering of the Fréchet distance. In *SoCG*, volume 129, pages 17:1–17:21, 2019.
- [11] Kevin Buchin, Maike Buchin, Joachim Gudmundsson, Maarten Löffler, and Jun Luo. Detecting commuting patterns by clustering subtrajectories. *Int. J. Comput. Geom. Appl.*, 21(03):253–282, 2011.
- [12] Kevin Buchin, Yago Diez, Tom van Diggelen, and Wouter Meulemans. Efficient trajectory queries under the Fréchet distance (GIS Cup). In *SIGSPATIAL*, page 99, 2017.

Figure 8: Example of querying NBA trajectories using tSTAT with $R = 0.45$.

- 101, 2017.
- [13] Jonathan C Campbell, Jonathan Tremblay, and Clark Verbrugge. Clustering player paths. In *FDG*, 2015.
 - [14] Matteo Ceccarello, Anne Driemel, and Francesco Silvestri. FRESH: Fréchet similarity with hashing. In *WADS*, pages 254–268, 2019.
 - [15] Anne Driemel and Francesco Silvestri. Locality-sensitive hashing of curves. In *SoCG*, 2017.
 - [16] Fabian Dütsch and Jan Vahrenhold. A filter-and-refinement-algorithm for range queries based on the Fréchet distance (GIS Cup). In *SIGSPATIAL*, page 100, 2017.
 - [17] Thomas Eiter and Heikki Mannila. Computing discrete Fréchet distance. Technical report, TU Vienna, 1994.
 - [18] Johannes Fischer and Daniel Peters. GLOUDS: Representing tree-like graphs. *J. Discrete Algorithm.*, 36:39–49, 2016.
 - [19] Edward Fredkin. Trie memory. *Commun. ACM*, 3(9):490–499, 1960.
 - [20] Simon Gog, Timo Beller, Alistair Moffat, and Matthias Petri. From theory to practice: Plug and play with succinct data structures. In *SEA*, pages 326–337, 2014.
 - [21] Simon Gog and Rossano Venturini. Fast and compact Hamming distance index. In *SIGIR*, pages 285–294, 2016.
 - [22] Rodrigo González, Szymon Grabowski, Veli Mäkinen, and Gonzalo Navarro. Practical implementation of rank and select queries. In *WEA*, pages 27–38, 2005.
 - [23] Dan Greene, Michal Parnas, and Frances Yao. Multi-index hashing for information retrieval. In *FOCS*, pages 722–731, 1994.
 - [24] Joachim Gudmundsson and Michael Horton. Spatio-temporal analysis of team sports. *ACM Comput. Surv.*, 50(2), 2017.
 - [25] Piotr Indyk. Approximate nearest neighbor algorithms for Fréchet distance via product metrics. In *SCG*, pages 102–106, 2002.
 - [26] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *STOC*, pages 604–613, 1998.
 - [27] Guy Jacobson. Space-efficient static trees and graphs. In *FOCS*, pages 549–554, 1989.
 - [28] Shunsuke Kanda, Kazuhiro Morita, and Masao Fuketa. Compressed double-array tries for string dictionaries supporting fast lookup. *Knowl. Inf. Syst.*, 51(3):1023–1042, 2017.
 - [29] Shunsuke Kanda, Koh Takeuchi, Keisuke Fujii, and Yasuo Tabei. Succinct trit-array trie for scalable trajectory similarity search. *CoRR*, abs/2005.1, 2020.
 - [30] Maximilian Konzack, Thomas McKetterick, Tim Ophelders, Maike Buchin, Luca Giuggioli, Jed Long, Trisalyn Nelson, Michel A Westenberg, and Kevin Buchin. Visual analytics of delays and interaction in movement data. *Int. J. Geogr. Inf. Sci.*, 31(2):320–345, 2017.
 - [31] Benjamin Krogh, Christian S Jensen, and Kristian Torp. Efficient In-memory indexing of network-constrained trajectories. In *SIGSPATIAL*, 2016.
 - [32] Chen Li, Jiaheng Lu, and Yiming Lu. Efficient merging and filtering algorithms for approximate string searches. In *ICDE*, pages 257–266, 2008.
 - [33] Alex X Liu, Ke Shen, and Eric Torng. Large scale hamming distance query processing. In *ICDE*, pages 553–564, 2011.
 - [34] Wuman Luo, Haoyu Tan, Lei Chen, and Lionel M Ni. Finding time period-based most frequent path in big trajectory data. In *SIGMOD*, pages 713–724, 2013.
 - [35] Gurmeet Singh Manku, Arvind Jain, and Anish Das Sarma. Detecting near-duplicates for web crawling. In *WWW*, pages 141–150, 2007.
 - [36] Mohammad Norouzi, Ali Punjani, and David J Fleet. Fast exact search in Hamming space with multi-index hashing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(6):1107–1119, 2014.
 - [37] Mihai Patrascu. Succincter. In *FOCS*, pages 305–313, 2008.
 - [38] Giulio Ermanno Pibiri and Rossano Venturini. Efficient data structures for massive n-gram datasets. In *SIGIR*, pages 615–624, 2017.
 - [39] Giulio Ermanno Pibiri and Rossano Venturini. Handling massive N-gram datasets efficiently. *ACM Trans. Inf. Syst.*, 37(2), 2019.
 - [40] Jianbin Qin, Chuan Xiao, Yaoshu Wang, and Wei Wang. Generalizing the pigeon-hole principle for similarity search in Hamming space. *IEEE Trans. Knowl. Data Eng.*, 2019.
 - [41] Maede Rayatidamavandi, Yu Zhuang, and Mahshid Rahnamay-Naeini. A comparison of hash-based methods for trajectory clustering. In *DASC*, pages 107–112, 2017.
 - [42] Ivan Sanchez, Zay Maung Maung Aye, Benjamin I. P. Rubinstein, and Kotagiri Ramamohanarao. Fast trajectory clustering using hashing methods. In *IJCNN*, pages 3689–3696, 2016.
 - [43] Long Sha, Patrick Lucey, Yisong Yue, Peter Carr, Charlie Rohlf, and Iain Matthews. Chalkboarding: A new spatiotemporal query paradigm for sports play retrieval. In *IUI*, pages 336–347, 2016.
 - [44] Long Sha, Patrick Lucey, Stephan Zheng, Taehwan Kim, Yisong Yue, and Sridha Sridharan. Fine-grained retrieval of sports plays using tree-based alignment of trajectories. *CoRR*, abs/1710.0, 2017.
 - [45] Shuo Shang, Ruogu Ding, Kai Zheng, Christian S Jensen, Panos Kalnis, and Xiaofang Zhou. Personalized trajectory matching in spatial networks. *The VLDB Journal*, 23(3):449–468, 2014.
 - [46] Zeyuan Shang, Guoliang Li, and Zhifeng Bao. DITA: Distributed in-memory trajectory analytics. In *SIGMOD*, pages 725–740, 2018.
 - [47] Renchu Song, Weiwei Sun, Baihua Zheng, and Yu Zheng. PRESS: A novel framework of trajectory compression in road networks. *PVLDB*, 7(9):661–672, 2014.
 - [48] E Sriraghavendra, K Karthik, and Chiranjib Bhattacharyya. Fréchet distance based approach for searching online handwritten documents. In *ICDAR*, volume 1, pages 461–465, 2007.
 - [49] Kevin Toohey and Matt Duckham. Trajectory similarity measures. *SIGSPATIAL Special*, 7(1):43–50, 2015.
 - [50] Sebastiano Vigna. Broadword implementation of rank/select queries. In *WEA*, pages 154–168, 2008.
 - [51] Sheng Wang, Zhifeng Bao, J Shane Culpepper, Zizhe Xie, Qizhi Liu, and Xiaolin Qin. Torch: A search engine for trajectory data. In *SIGIR*, pages 535–544, 2018.
 - [52] Martin Werner and Dev Oliver. ACM SIGSPATIAL GIS Cup 2017: Range queries under Fréchet distance. *SIGSPATIAL Special*, 10(1):24–27, 2018.
 - [53] Tim Wylie and Binhai Zhu. Protein chain pair simplification under the discrete Fréchet distance. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 10(6):1372–1383, 2013.
 - [54] Dong Xie, Feifei Li, and Jeff M Phillips. Distributed trajectory similarity search. *PVLDB*, 10(11):1478–1489, 2017.
 - [55] H Yuan and G Li. Distributed In-memory trajectory similarity search and join on road network. In *ICDE*, pages 1262–1273, 2019.
 - [56] Huan Chen Zhang, Hyeontaek Lim, Viktor Leis, David G Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. SuRF: Practical range query filtering with fast succinct tries. In *SIGMOD*, pages 323–336, 2018.
 - [57] Xiaoyang Zhang, Jianbin Qin, Wei Wang, Yifang Sun, and Jiaheng Lu. HmSearch: An efficient Hamming distance query processing algorithm. In *SSDBM*, page 19, 2013.
 - [58] Yu Zheng. Trajectory data mining: an overview. *ACM Trans. Intell. Syst. Technol.*, 6(3):29, 2015.
 - [59] Haohan Zhu, Jun Luo, Hang Yin, Xiaotao Zhou, Joshua Zhexue Huang, and F Benjamin Zhan. Mining trajectory corridors using Fréchet distance and meshing grids. In *SIGKDD*, pages 228–237, 2010.

This section shows the other experimental results not presented in Section 5.3. Figure 10 shows the search time for showing the scalability. For all the datasets, the search times of tSTAT were always fast. Figure 11 shows the construction time. Also as observed in Section 5.3, the time needed for only building STATs from sketches was not significant because the construction times of tSTAT was approximate to those of LS.

Figure 10: Results of average search times per query in ms for varying the number of trajectories.

Figure 11: Results for construction times in minutes for varying recalls.