

Semantic Trajectory Compression

Falko Schmid¹, Kai-Florian Richter¹, and Patrick Laube²

¹ Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition,
University of Bremen, P.O. Box 330 440, 28334 Bremen, Germany

{schmid,richter}@sfbtr8.uni-bremen.de

² Department of Geomatics, The University of Melbourne, VIC 3010, Australia
plaube@unimelb.edu.au

Abstract. In the light of rapidly growing repositories capturing the movement trajectories of people in spacetime, the need for trajectory compression becomes obvious. This paper argues for *semantic trajectory compression* (STC) as a means of substantially compressing the movement trajectories in an urban environment with acceptable information loss. STC exploits that human urban movement and its large-scale use (LBS, navigation) is embedded in some geographic context, typically defined by transportation networks. STC achieves its compression rate by replacing raw, highly redundant position information from, for example, GPS sensors with a semantic representation of the trajectory consisting of a sequence of *events*. The paper explains the underlying principles of STC and presents an example use case.

Keywords: Trajectories, Moving Objects, Semantic Description, Data Compression.

1 Motivation

Trajectories, the representation of movement by means of positioning fixes, usually contain data which can be considered as highly redundant information. Movement often happens along network infrastructure, such as streets or railway tracks, and the significant behavior patterns, such as stops, are performed along it as well. Especially in dense urban environments there are not many alternatives for reaching a certain destination other than to move along available network links. The representation and storage of trajectories by means of lists of fixes also pose questions about knowledge gain and further processing; trajectories are only meaningful when their spatial context is considered. Relating trajectories to their spatial context at an early stage will lead to improved means of analyzing them with standardized methods in a later stage.

Figure 1b is a visualization of a stream of raw positional data (Figure 1a) produced by a tracking system (e.g. GPS). Figure 1c depicts the same movement embedded in its geographical context. An object moved through a system of streets to reach its destination. The actual information contained in trajectories is the sequence of implicitly encoded spatio-temporal events, i.e., a single datum is usually not of interest, but rather the *significant* information with respect to




Fig. 1. Problem overview: (a) Raw positional data; (b) trajectory in two-dimensional space, moving from origin to destination; (c) trajectory embedded in geographic context, the semantically annotated *map* features a train line #6 with two stations; tram lines #3 and #5 with several stops; major streets (Upper case) and minor streets (lower case); (d) minimalist representation of the same trajectory, used for the semantic compression: origin, street *g*, street *B*, tram line #3, straight, destination

the movement. Significance depends on the application context, but is always based on the spatial course and the determination of events (usually stops).

Having a look at the trajectory, it becomes obvious that its course can be described by referring to elements of the network without losing relevant information. The course can be expressed by the streets and tram tracks it moves along (street *g*, street *B*, tram line #3, and straight along the streets *D*, *G*, *H*). This is a minimalist representation of the movement (Figure 1d). Instead of using a large number of point coordinates, the movement can be described with elements of the transportation network, annotated with the behaviorally significant elements (in this case origin, stops, and destination).

Networks as a constraining basis for movement reduce the dimensionality of space and, thus, allow efficient indexing structures for moving objects [1]. Most work focused on the geometry of the underlying networks. However, the database community has acknowledged semantics—the meaningful annotation of moves with labels from the embedding environment—as being paramount for the interpretation and analysis of raw trajectory data [2, 3]. Whereas exploiting semantics is a young branch in spatial database research, in spatial cognition the semantics of movement has long been exploited for designing better wayfinding instructions [4, 5]. Going a step beyond utilizing spatial infrastructure as a suitable representation, a semantic representation of the trajectory can be implemented that focuses on qualitative change in course and events without losing the conceptual information of the movement data.

2 Semantics in Trajectories

The majority of systems tracking the movement of individuals produce lists of time-stamped position samples, so-called *fixes* in the form of tuples (x, y, t) . Even though this is a discrete approximation of the movement behavior, it is widely accepted to model the respective movement as a sequence of fixes, connected with straight line segments. In its most simple form a *trajectory* is a 2-dimensional polygonal line connecting the fixes of a moving individual. For example, in Figure 1, an individual has moved from location *ori* to destination *dest*, starting at 00:00 arriving at 59:40. Figure 1a illustrates a trajectory's raw data, Figure 1b the respective trajectory in two-dimensional space. Note the raw data column only illustrates a subset of the plotted fixes.

A *map* is a semantically annotated network of edges and nodes. A map represents the transport network of an urban environment, featuring streets, bus, tram and train lines (see Figure 1c). In a map vertices are unambiguously defined, either by IDs or by (x, y) coordinate tuples. Further, edges may have a label (street name, or bus, tram, train line). This is a $n : 1$ -relation as several edges can have the same label. Vertices of bus, tram, and train lines are stops and stations; these may be labeled with the stops' names. The labeling of edges and vertices can extend several levels. An edge may at the same time have a local street name (e.g., "Ostertorsteinweg"), be part of a national highway system (e.g., "A7"), and be part of a bus or tram line (e.g., "Tram #3").

3 Semantic Trajectory Compression

If movement happens in a transport network, as is usually the case in urban environments, trajectories can be mapped to a map representing this environment. The mapping of fixes to vertices and edges of the transport network then allows for exploiting this structure to restricting a trajectory's representation to the significant events. A network reduces the dimensionality of a two-dimensional movement space. It allows for concise positioning of a moving object through time-stamping along edges and at vertices, which both have unique identifiers. In a semantically annotated map, edges and vertices can be aggregated according to shared labels, for example their street names or the train lines. Often, several consecutive edges represent the same street and, thus, share the same label. Tram and bus lines may extend over large sections of an urban transport network. Thus, the semantic annotation of the network offers a high-level reference system for urban spaces, which is exploited in STC.

Taking this perspective, streets and tram, bus or train lines are viewed as *mobility channels* that moving objects hop on, ride for a while, and hop off again to catch another channel that brings them closer to their destination. In terms of trajectory compression, this perspective has the advantage that only little information needs to be retained for describing the movement of an individual in terms of riding such channels. For most kinds of movement storing a sequence of the identifiers of the specific channels and hop-on and hop-off times results

in a sufficient approximation of the individual’s movement through the network. At the same time, this drops a large amount of fixes, which are highly correlated and, hence, redundant. Semantic compression of trajectories makes use of principles and methods that have been previously implemented for the generation of cognitively motivated route descriptions (the GUARD process, cf. [6]). Broadly, it is based on three steps:

1. Identify the relevant events along the trajectory. Relevant events are origin and destination, as well as street intersections and public transport stops (see Figure 1c).
2. For each event, determine all possible descriptions of how movement continues from here. These descriptions are egocentric direction relations (straight, left, right, etc.; in Figure 1c straight from edge D to edge G) or changes in labels of network elements (in Figure 1c change from label street B to tram line #3) for capturing the motion continuation of an event.
3. Based on the descriptions, combine consecutive events into sequences of events. These sequences are termed (*spatial*) *chunks* [5]. The compressed trajectory consists of sequences of such spatial chunks (Figure 1d).

In decompression, the aim is to reconstruct movement through an environment. In the chosen semantic approach, decompression does not restore the original trajectory, but rather the path through the network along with inferred time-stamps. The path contains all information on changes of direction as well as places along the way; each such event point is coupled with a time-stamp stating when in the travel behavior it occurred. Note that for all reconstructed points in the decompressed trajectory, i.e., those that are not original points retained in the compressed trajectory, the time-stamp is calculated based on an assumed linear movement behavior between start and end point of a chunk. While this time estimation is a simplification resulting in information loss, it provides no limitation for the targeted applications (see Section 5).

In a nutshell, the decompression algorithm iterates through the sequence of chunks stored in the compressed trajectory. It returns a sequence of vertices that are a geometric representation of the travelled path through the network. In more detail, beginning with the start vertex of a chunk the algorithm adds geometric edges to the reconstructed path until the end vertex is reached. To this end, it uses different strategies to determine which edge is to be added; these strategies depend on the description used for chunking. Each added vertex is linked to a time-stamp, which is calculated assuming constant movement speed, i.e., representing a fraction of time corresponding to the fraction of the distance travelled between start and end vertex.

4 Example Use Case

Figure 2 shows an example use case of applying semantic trajectory compression. The geometric representation of the path contains 115 points in space-time (115 tuples of (x, y, t)). It further comprises 52 events, i.e., 52 intersections and stops

along the way. Performing compression yields the following 6 elements as result:

```
((3490254.00 5882064.00 00:00) (3490057.00 5882110.00 01:12) "Bei den Drei Pfählen")
((3490057.00 5882110.00 01:12) (3489534.00 5882241.50 04:47) "Am Hulsberg")
((3489534.00 5882241.50 04:47) (3488929.50 5882100.00 08:21) "Am Schwarzen Meer")
((3488929.50 5882100.00 08:21) (3488222.50 5882314.50 13:09) "Vor dem Steintor")
((3488222.50 5882314.50 13:09) (3487688.75 5882291.00 16:17) "Ostertorsteinweg")
((3487688.75 5882291.00 16:17) (3487544.75 5882351.00 17:21) "Am Wall")
```

As can be seen, STC achieves a high compression rate. Instead of the 115 original points, it ends up with only 6 items, which corresponds to a compression rate of 94.78%. Considering that each item in the compressed trajectory consists of three elements, the ratio is still 18 to 115 elements or 84.35%. Decompressing the compressed trajectory reconstructs the original path. It also keeps the time-stamps explicitly stated in the compressed trajectory. There are some differences in the geometric representation—in this case the reconstructed path contains 3 coordinates more than the original path. This can be explained with ambiguities in the underlying geographic data set that for some streets has individual representations of different lanes, resulting in different geometric representations. However, there is no visual or semantic difference between the original and the reconstructed path; all events of the original path are correctly reconstructed. Regarding time, the original time-stamps stored in the compressed trajectory are retained; all other reconstructed events are annotated with estimated time-stamps assuming linear movement within a chunk.




Fig. 2. The map shows part of the inner-city region of Bremen, Germany. a) The displayed path (the bold line) runs from right to left. The dots on the path mark all event points along the way. b) The events stored in the compressed trajectory.

5 Conclusions and Outlook

This paper presents a novel approach for compressing large volumes of trajectory data by exploiting the semantic embedding of movement in a geographical context. Inspired by network-constrained object indexing and techniques used in spatial cognition and wayfinding, the paper presents *semantic trajectory compression* (STC). STC matches the movement to the underlying map and aggregates chunks based on identical semantic descriptions. Initial experiments with a set of use case trajectories captured with volunteers in the city of Bremen serve

as a proof of concept, deliver promising results for future experiments and help to identify limitations and a road map for future work. After implementing an STC prototype, future work will focus on evaluating the STC algorithm with large and diverse trajectory data. Extensive experiments with real, recorded trajectory data shall identify possible conceptual shortcomings and reveal the runtime characteristics of the STC algorithm for various scenarios.

As a main contribution, the paper illustrates that the embedding of human movement in the geographic context of an urban street network can successfully be exploited for compressing large volumes of raw trajectory data with acceptable information loss. The reconstructed information is suited for a number of applications based on individual spatial profiles which are not built upon a fine-grained analysis of movement dynamics (e.g., ascending and descending velocity). Specifically, this holds for prior-knowledge based navigation support [7] which relies on previously visited places and traveled paths. Also, most applications within the field of Location Based Services that rather rely on a clean *model* of movement than its detailed dynamics will benefit from semantically compressed trajectories.

Acknowledgments

Support by the German Science Foundation (DFG) and Group of Eight / DAAD Australia Germany Joint Research Co-operation Scheme is acknowledged.

References

- [1] Li, X., Lin, H.: Indexing network-constrained trajectories for connectivity-based queries. *Int. Journal of Geographical Information Science* 20(3), 303–328 (2006)
- [2] Alvares, L.O., Bogorny, V., Kuijpers, B., Fernandes de Macedo, J.A., Moelans, B., Vaisman, A.: A model for enriching trajectories with semantic geographical information. In: *GIS 2007: Proc. of the 15th annual ACM international symposium on Advances in GIS*, pp. 1–8. ACM, New York (2007)
- [3] Spaccapietra, S., Parent, C., Damiani, M.L., de Macedo, J.A., Portoa, F., Vangenot, C.: A conceptual view on trajectories. *Data and Knowledge Engineering* 65(1), 126–146 (2008)
- [4] Tversky, B., Lee, P.U.: How space structures language. In: Freksa, C., Habel, C., Wender, K.F. (eds.) *Spatial Cognition 1998. LNCS (LNAI)*, vol. 1404, pp. 157–175. Springer, Heidelberg (1998)
- [5] Klippel, A., Hansen, S., Richter, K.F., Winter, S.: Urban granularities – a data structure for cognitively ergonomic route directions. *GeoInformatica* 13(2), 223–247 (2009)
- [6] Richter, K.F.: *Context-Specific Route Directions - Generation of Cognitively Motivated Wayfinding Instructions*. DisKI, vol. 314. IOS Press, Amsterdam (2008); also appeared as SFB/TR 8 Monographs Volume 3
- [7] Schmid, F.: Knowledge based wayfinding maps for small display cartography. *Journal of Location Based Services* 2(1), 57–83 (2008)