Merging Roadmaps Using Graph Distance Measures

Erfan Hosseini Sereshgi* and Carola Wenk*

1 Introduction

Merging roadmaps is a process that can arise in a variety of situations. When new map companies or ride apps are founded, they often do not possess a road map of their own. Most use freely available maps from OpenStreetMap (openstreetmap.org) as base maps and over time improve their networks by using different sources like GPS data, satellite images or simply a roadmap from another data provider. This requires comparing a new road network to a base roadmap, and later adding it to the base roadmap. Some parts of the new road network might already be present in the base roadmap, some might be missing, and some might be partially present or have different geometries. For example, the exact locations of the centerline representations of streets often differ, as do locations of street intersections and highway ramp connections.

In this paper we consider a roadmap to be a geometric graph in the plane, i.e., the graph is given together with a planar embedding which maps vertices to points and edges to line segments. We note that some of the approaches we discuss may also apply to immersed graphs which allow transversal edge intersections and hence are able to model bridges. Fairly recently, several methods have been proposed to compare such geometric graphs, including edit distances and Fréchet-based distances; see [2] and [5] for surveys. Some but not all of these distance measures return a correspondence between the two graphs that indicate which parts of graph G correspond to graph H. Naturally, such correspondence will be beneficial for merging graphs.

Indeed, in this paper we study one particular method for comparing geometric graphs, the graph sampling method [3, 4], and we present a method for using the correspondence computed by this distance to merge two input roadmaps from the same area.

2 Merging Roadmaps using Graph Sampling

Being the most popular method for evaluating reconstructed maps, the graph sampling method has been used for almost a decade. Graph sampling was first introduced by Biagioni and Eriksson [3, 4] to evaluate their results and compare their approach with the state of the art map construction algorithms. Even though it has been wrongly called a "distance" on some occasions, as we explained in our previous work [1] graph sampling is a statistical method which discretizes 2D immersed graphs into sample points and matches the samples instead. The essential idea is to first compute a set of point samples on each map, and then to match pairs of samples—one from each map—via a one-to-one matching. For deciding whether two samples can be matched, different criteria, e.g., based on distance or orientation, can be used. Samples are placed on edges of the graph starting from a random point on each connected component following a certain interval between each consecutive pair of samples. After choosing a matching distance threshold, these samples are matched using weighted maximum matching and counted for the evaluation. The matching distance thresholds only allows the samples within that distance to be matched to each other. Fig. 2 shows an overview of this method where samples are matched based on their distance from one another and the angle between their corresponding edges.

An important attribute that we want to preserve while merging roadmaps is the integrity of the edges. The usual graph sampling method places samples on the graphs at a fixed distance (typically 5 meters), but not actually on intersections. We slightly change this placement to fit our purpose better. We sample each

*{shosseinisereshgi,cwenk}@tulane.edu;Department of Computer Science, Tulane University, USA. Partially supported by National Science Foundation grant CCF 2107434.

edge independently and always place samples on both endpoints of an edge, and then sample the remainder of the edge using a fixed interval. While this means that the length of the intervals can be inconsistent among edges, we make sure that interval lengths on a single edge differ by at most 1 meter. And by construction, the samples on the endpoints are common between adjacent edges. Our goal is to make minimal adjustments in the base map while adding as much detail as possible. Fig. 1 Shows the types of edges that our algorithm considers as candidates for addition:

1. Edges with matched endpoints only: We add a replica of these edges by connecting the corresponding matched samples on the base map
2. Edges with only one matched endpoint: We add a replica of these edges by connecting the corresponding matched sample from the base map to the unmatched endpoint.
3. Edges with no matched samples: We do not add such edges unless they evolve into type 1 or 2

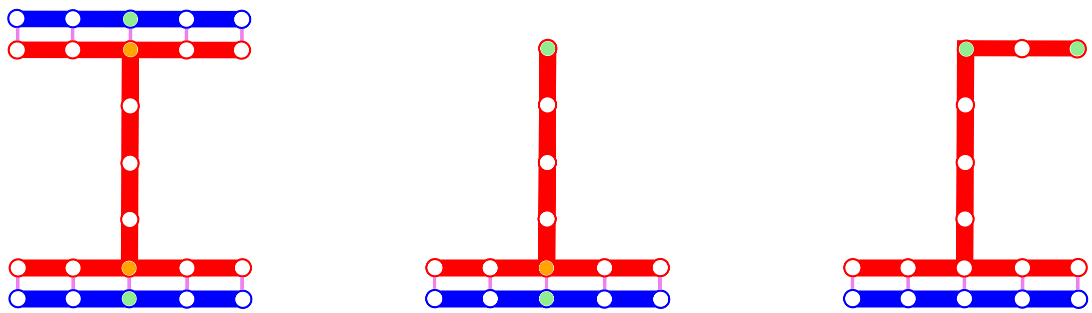

(a) Both endpoints are matched (b) Only one matched endpoint (c) No matched samples

Figure 1: Types of edges that are considered for addition. Blue is the base map and red is the supplementary map. The circles are the samples. Matched samples are connected via a violet line segment. Orange points are the matched samples on the candidate edge. Green points show the endpoints of the new edge that will be added to the base map.

As we add type 1 and 2 edges, we update the matching status of samples on the supplementary map so type 3 edges that previously had no matched samples can have a new matched endpoint because an adjacent edge was recently added to the base map. This process is necessary to avoid short isolated edges and completely isolated connected components from appearing on the base map. Note that roadmpas are singly connected components by definition, so even though edges with no matched samples usually do not appear in the results, an isolated edge/connected component can exist in a reconstructed map. Such edges are often caused by noise in GPS datasets or the map construction algorithms that were used.

In general, merging a roadmap into another roadmap is relatively simpler because the edges are more refined and precise. However since these maps have a lot in common one has to take extra care with how they choose the additional edges. This task can get quite challenging in areas with complex networks such as near the highways and intersections with more than four incoming/outgoing roads. Small angle turns like roundabouts can also become an issue when they are not correctly presented in the base map since they are made of a set of short edges.

The final step is to integrate these new edges into the base map. This means that the new edges must be reachable from other edges. Even though the samples on the endpoints of the added edges were matched, these samples do not necessarily fall on a vertex in the base map. They can be in the middle of an edge, isolated from other vertices. We overcome this issue by dividing such edges in two parts so that the matched sample would become a vertex connecting the two.

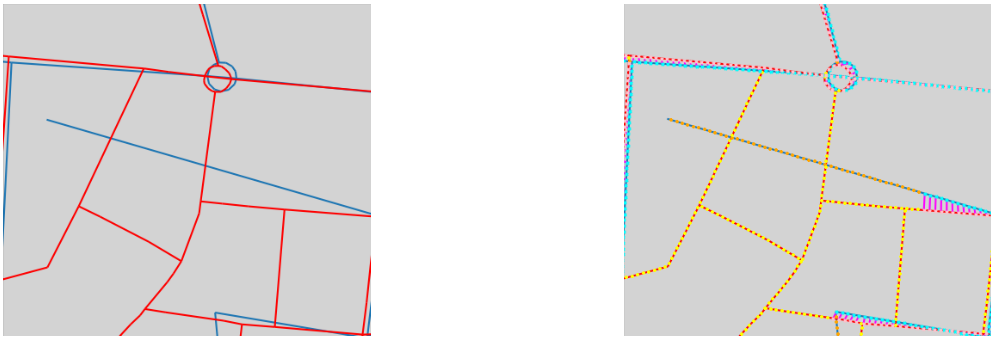

Figure 2: Two roadmaps in blue and red and their graph sampling results: matched points are connected by pink lines, unmatched points are yellow or orange.

3 Experiments

3.1 Data

In our experiments, and throughout this paper, we used reconstructed maps and roadmaps from OpenStreetMap (OSM) for the *Chicago*, *Athens*, and *Berlin* data sets that are available on `mapconstruction.org`. For *Berlin*, we have small ($16km^2$) roadmaps from TeleAtlas (TA) from 2007 and OpenStreetMap (OSM) from April 2013. Similarly, for *Athens*, we have TA maps from 2007 and OSM maps from 2010. Unfortunately, the TA maps are not publicly available. Finally, we have used the Graph Sampling Toolkit from [1] at `https://github.com/Erfanh1995/GraphSamplingToolkit` for visualizations.

3.2 Experiments on Roadmaps

We show some results of merging TA maps into OSM maps of *Athens* and *Berlin*. We used 5 meter intervals, the matching distance threshold of 15 meters, and weighted maximum matching for Graph Sampling. Despite having a few minor artifacts, the final roadmaps look very promising. As shown in Fig. 3, The input maps are quite similar to each other except for a roundabout that seems to be missing from the OSM map. This particular path which is covered with unmatched green samples in Fig. 4 is perfectly added to the base map. Furthermore, a highway ramp has been added to the left side of the roadmap however this seems to be an error since such a ramp already exists in the base map in a different shape and position. In Fig. 5, a roundabout is missing in the OSM map (blue) which has been added from TA map (red). Several missing roads are also added in the merge with no notable error or artifact.

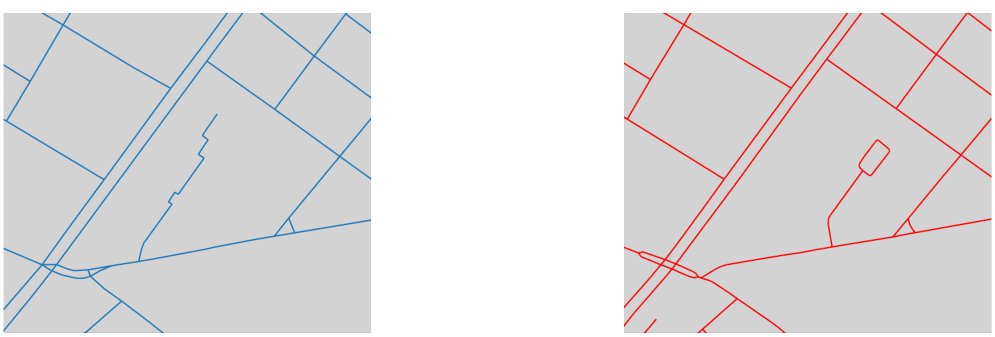

Figure 3: The base map (OSM) and the supplementary map (TA) in blue and red respectively

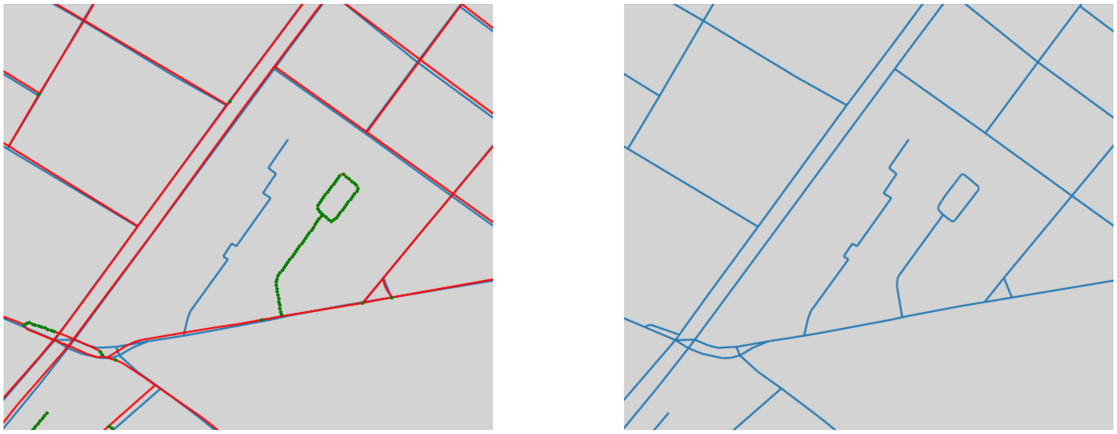

Figure 4: The base map (OSM) and the supplementary map (TA) in blue and red. Green points are unmatched points on the supplementary map. The roadmap on the right is the final merged map.

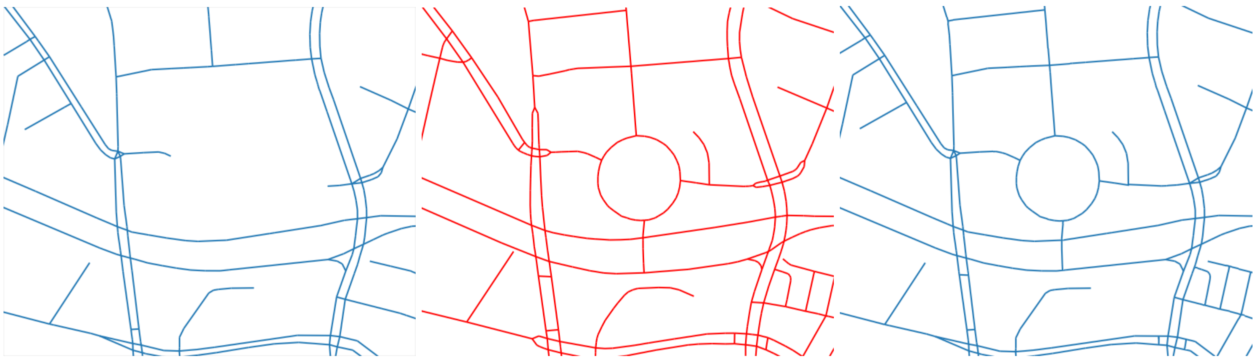

Figure 5: The OSM map in blue and the TA map in red. The blue map on the right is the result of merging TA into OSM.

4 Discussion

We presented an effective method to add details to roadmaps from other sources using the Graph Sampling method. The quality of the merged map can depend on the input maps, matching distance threshold and the type of matching algorithm that is used for Graph Sampling. Regarding evaluation, although the outputs look quite promising, more experiments are required to determine how our method can perform against other map merging algorithms. Such experiments can be difficult to conduct due to the lack of any ground-truth.

The process can be quick on small roadmaps and a bit time consuming on larger datasets, however, the majority of the runtime is needed for the Graph Sampling and can be accelerated by using the greedy matching provided in [1]. The merging can be done in linear time in regard to the number of edges, and can also be improved further to possibly add more details in uncertain situations such as having a partially matched edge.

As we mentioned earlier, correspondences from other graph distance measures can be used to merge road maps. Although most of which are more expensive than Graph Sampling, in some cases they may yield interesting or better results. Moreover, a combination of such correspondences may be used in a multi-step algorithm to capture more details and achieve more refined roads.

References

[1] Jordi Aguilar, Kevin Buchin, Maike Buchin, Erfan Hosseini Sereshgi, Rodrigo I. Silveira, and Carola Wenk. Graph sampling for map comparison. *In submission to ACM TSAS*, 2022.

[2] M. Ahmed, S. Karagiorgou, D. Pfoser, and C. Wenk. *Map Construction Algorithms*. Springer International Publishing, first edition, 2015.

[3] 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*, 2291:61–71, 2012.

[4] J. Biagioni and J. Eriksson. Map inference in the face of noise and disparity. In *Proc. 20th ACM SIGSPATIAL GIS*, pages 79–88, 2012.

[5] M. Buchin, E. Chambers, P. Fang, B.T. Fasy, E. Gasparovic, E. Munch, and C. Wenk. Distances between immersed graphs: Metric properties. *La Matematica*, 2022. Accepted.