An Algorithmic Approach to Quantifying GPS Trajectory Error

1st Matthew Plaudis
Department of Computer Science
University of Victoria
matthewplaudis@uvic.ca

2nd Muhammad Azam
Department of Computer Science
University of Victoria
mwrazam@uvic.ca

3rd Derek Jacoby
Department of Computer Science
University of Victoria
derekja@gmail.com

4th Marc-Antoine Drouin
Digital Technologies
National Research Council of Canada
Marc-Antoine.Drouin@nrc-cnrc.gc.ca

5th Yvonne Coady
Department of Computer Science
University of Victoria
ycoady@gmail.com

Abstract

The alignment of aerial and satellite imagery with ground sensor data is an ongoing research challenge. In dense urban environments, part of this challenge is induced by the positioning error of Global Positioning System (GPS). Despite the potential for error, many studies use GPS in order to infer road networks because GPS data is inexpensive and can be acquired quickly. Major transit organizations are freely providing data on the real-time position of their buses as well as ground truth route trajectories. This work exploits geospatial open data to construct a database of historical GPS from bus roads. Using this database, the GPS error map along main arteries of major cities can be reconstructed. The extraction of error maps is highly relevant for the planning and the joint exploitation of airborne and ground-based imagery. In this work, we use bus routes in downtown Victoria, BC, Canada and Adelaide, Australia to demonstrate the extraction GPS error maps.

1. Introduction

The increasing applications and more widespread use of Global Positioning Systems (GPS) have provided researchers with the opportunity to analyze travel patterns and observe real-time status updates of moving vehicles [9]. GPS receivers have been largely integrated into commonly used items such as cars and mobile phones, and can record a large amount of trajectory data [3]. Using the General Transit Feed Specification (GTFS), major transit organizations are freely providing GPS data from which we can infer vehicle trajectories [1]. Intuitively we can define a trajectory as a sequence of connected spatial points sampled from a continuously moving object, where each point has an associated time stamp and position [3].

GTFS data provides a ground truth transit route against which we can compare trajectories constructed from recorded GPS points. Since buses repeat the same route thousands of times every year, we can freely accumulate a large amount of data in order to detect and quantify route deviations. Further, open-source geospatial data allows the production of GPS error maps that can be used in multiple applications including the fusion of airborne imagery with ground-based sensors. As a case study, we fuse GTFS data for buses in Victoria, BC, Canada, as well as Adelaide, Australia, with high resolution satellite imagery in order to investigate trends in GPS errors.

Our proposed work shares research interests with literature that investigates algorithmic approaches to map inference using error prone GPS data and satellite imagery. The main challenge with automatic map inference stems from the inconsistencies in the data used by methods that attempt to generate a road network. Due to the low precision of GPS devices, recorded points are usually sampled inaccurately [3]. GPS errors mainly consist of two types:

- **Measurement Error**: In a GPS, the location of a device is calculated by its relative distance to multiple satellites. As the number and the stability of satellite connections oscillate from time to time, the sampled trajectory point usually lands to an arbitrary place nearby the actual location [3]. This can be observed in area B of Figure 1.

- **Sampling Error**: Although an object is moving continuously, its position can only be sampled periodically by the GPS device. Hence, the frequency of GPS position sampling, which is called the sampling rate, is an influential factor to the accuracy of GPS trajectory representation [3]. This can be observed in area A of Figure 1.

Figure 1. An illustration of instances of potential GPS trajectories (red) and their corresponding ground truth routes (green) of vehicles in downtown Victoria, BC, Canada. **Area A**: Shows an example of sampling error. Due to the combination of roadway speeds of $\geq$ 50km/h and turns at corners, the trajectories created using sequential GPS points go through buildings and over corners. **Area B**: An example of measurement error. The vehicle in this area is travelling on a road with taller buildings compared to the surrounding areas. As a result, the calculated distance and location of a vehicle may be skewed due to occlusion. **Area C**: An example of a GPS recording with little to no sampling or measurement error. There is less potential for both occlusion (travelling in a more suburban area with shorter buildings) and sampling errors (the trajectory direction does not change and road speeds are slower, meaning more closely sampled GPS points).

Given the potential for inaccurately sampled points, it is nontrivial to decide whether two points are sampled from the same or nearby road segments [5]. There is also evidence to suggest that one cannot infer a road network starting with a low-accuracy initial graph and noisy GPS data [4]. Given the complex nature of the road network inference problem, current approaches employ heuristic methods that, when applied to more broad examples, the quality of the resulting road network is no longer guaranteed [4].

In this paper we compare the advantages and disadvantages of current methods for automating map inference. We also explore the feasibility of quantifying the error of a GPS recording relative to it's ground truth position and what implications that has in areas such as UAV flight planning: we can optimize the flight path of UAV's based on the characterization of ground-level targets. Rather than treating GPS signals as merely a noisy dataset, we propose to characterize the systematic nature of the errors in that data as a spatially-based reliability map of GPS signals over ground positions within a region. This two dimensional reliability map pushes our concept of route similarity into a purely area-based conception of curve similarity and we will compare that with other measures of route similarity.

2. Related Approaches

Current approaches attempt to implement algorithms or learning based methods that rely on a ground truth road network. Cao and Krumm [5] create an attraction mechanism in order to group GPS points with other nearby trajectories, shown in Figure 2 [5]. This mechanism employs certain heuristics in order to consider road direction and multi-lane roads. By finding the direction orthogonal to the candidate segment, an attractive force pulls that segment towards others with the same directional movement and repels it from those travelling the opposite way. This works to remove outlying points that are the result of common GPS noise. Although this method is able to differentiate opposing directions of traffic, it does not perform well in intersections when vehicles from several directions overlap [5]. The method was also confined to a street network around Microsoft Campus in Redmond, WA, USA, and has not been tested on larger roadways with over/underpasses.

Bastani et al. [6] uses a Convolutional Neural Network (CNN) called RoadTracer to infer a road network using aerial imagery and a partially constructed graph of the underlying road network. The image and the partial road network G are used as inputs by the CNN, which in turn out-

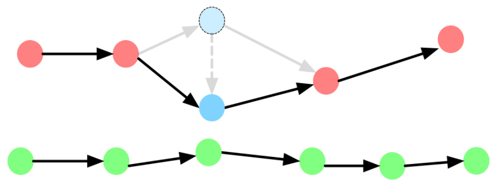

Figure 2. Figure based on the work proposed by Cao and Krumm [5] which demonstrates the attention mechanism the authors developed. Cao and Krumm demonstrate a potential solution for dealing with GPS points (red vertices) in a candidate trajectory that deviate from the most common trajectory (green vertices), due to measurement or sampling error. The original position of the point (faint blue vertex) is corrected to a position (darker blue vertex) nearer to the most common trajectory. This process is repeated for each of the GPS points in the candidate trajectory.

puts a movement decision in G. The learner compares the output against a completed road network G^* in order to adjust it's prediction. A notable finding was that RoadTracer performs much better with respect to occlusion by buildings and shadows in the Chicago and Boston regions in comparison to traditional methods that process each pixel in an image as either being 'road' or 'non-road'. Despite it's improved output when compared to previous methods, because of occlusion by tall buildings, shadows, and overpasses, accurately inferring roads from aerial imagery alone in dense urban areas and complex intersections is challenging [7].

Building on this, S. He et al. [7] proposed an alternative method of road network inference that uses both aerial imagery and GPS data. The algorithm first infers partial map correctly and then merges the constructed map portion to another map inferred by other high recall methods. The method starts from an initial map node and grows until covering the entire map region. It can also infer complex regions as it fully utilises the long-term travel information from the trajectories. The method proposed in this paper infers less travelled road networks with a high degree of accuracy that has not been previously seen. However, as the authors noted the two-step approach of RoadRunner can be computationally expensive and depending on the application this can be a hindrance.

With respect to methods that have used transit GPS data, Raymond and Imamichi [8] employ a map-matching algorithm to predict which GPS trajectories belong to which predefined route (shown in Figure 3). This is accomplished by creating sequential segments using coinciding GPS points and comparing these segments with bus routes using a cosine similarity function. Although this method showed accurate results, it is computationally expensive and relies on an accurate base map in order to map-match [9].

In May 2021 Lyu et al. [9] also investigated bus route identification. Using Fréchet distance, which captures the largest mismatched distance between two curves, a route R_i is matched with a candidate path Q_j that has the minimum maximum Fréchet distance (i.e. dissimilarity). One drawback to the solution proposed in this paper is a data preprocessing step, that may be impractical in a real time, automatic inference setting.

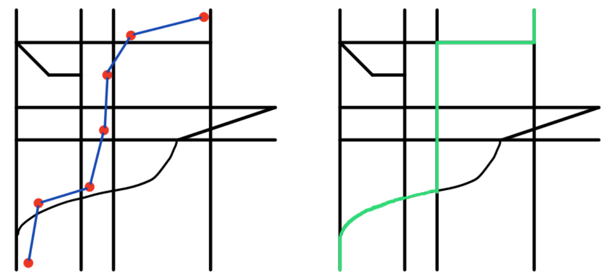

Figure 3. Figure based on the work proposed by Raymond and Imamichi [8] demonstrating the results of their map matching algorithm. The figure on the left shows the original GPS trajectories (blue) created using sequential GPS points (red), overlaid on a road network (black). Both the GPS points and their trajectories are not following a road network. On the right is the vehicle path (green) that is the result of the authors applying their map matching solution.

Common among all of the current methods is the recognition that both aerial imagery and GPS data is inherently noisy. Each method that was proposed attempts to implement heuristics in order to infer a road network or a vehicle's trajectory despite the fact that the data is imperfect. Although there is extensive research into how we can infer an accurate road map based on inaccurate data, currently there is a lack of research on how GPS quality affects performance [3].

3. Current Work & Area of Interest

Our current research focuses on a method to quantify the degree of inaccuracy between a recorded GPS point and it's ground truth. Using data collected from the GTFS for buses in Victoria, British Columbia, and Adelaide, Australia, we can compare the GPS trajectory from individual trips with ground truth routes provided by BC Transit.

Explicitly, a trajectory T_r is a sequence of spatial points $T_r : p_1 \to p_2 \to ... \to p_n$ sampled from a continuously moving object [3]. Each point p_i is chronologically ordered and consists of a position $< x_i, y_i > \in R^2$ at a timestamp t_i. if T_r consists of n points, we can say that T_r has $n-1$ seg-

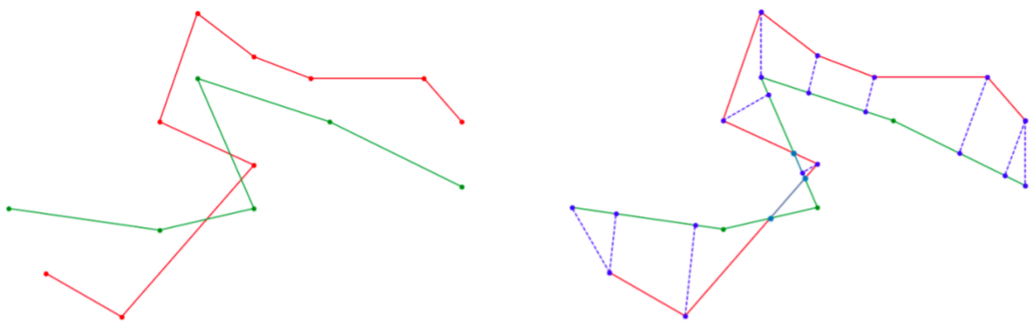

Figure 4. An example $\in R^2$ of the iterative projection construction implemented in Algorithm 1. The algorithm takes as input the vertex and edge set of both the GPS path (red) as well as the reference path (green). The algorithm processes each GPS vertex from right to left and adds $\hat{j}$ (blue) according to the outlined heuristics. $\hat{j}$ expresses a similar relationship between the two paths compared to the Fréchet distance described in Lyu et al [9]

ments, where $s_1 = p_1 \to p_2$, for example. We can also define sets $Q = \{q_1, q_2, ..., q_n\}$ and $R = \{r_1, r_2, ..., r_m\}$ as a collection of recorded trips along specified bus routes and a collection of possible bus routes, respectively. Each $r_i \in R$ and $q_j \in Q$ are trajectories by the definition given above. In order to quantify error, we employ a heuristic-based algorithm (given below in the form of a naive optimization as Algorithm 1) that calculates the vertical component of the projection of a GPS point p_i onto the most nearby segment s_i of route $r_i \in R$, denoted $\hat{j}_{proj_{s_i} p_i}$.

As described in Related Approaches, Lyu et al. [9] uses the concept of Fréchet distance to measure the similarity between GPS trajectories and predefined bus routes in order to implement a map matching algorithm. Map matching is the process of comparing each route $r_i \in R$ with a candidate GPS trajectory made up of several segments. If the segments of a route and a GPS trajectory within the same spatiotemporal window are similar, their Fréchet distance will be small, meaning there is a spatial similarity between them.

Comparatively, we find $\hat{j}_{proj_{s_i} p_i}$ in Algorithm 1 in order to segment individual polygons between r_i and q_j. The context in which our research finds a similar measure to Fréchet distance differs from that of Lyu et al, as we do not use the magnitude of $\hat{j}_{proj_{s_i} q_i}$ in order to characterize the similarity between a GPS trajectory and a ground truth trajectory. Further, we are focused on quantifying the spatial severity of GPS errors associated with certain regions as opposed to finding the most similar route in comparison to a candidate trip as Lyu et al. have done.

Algorithm 1: Iterative Projection Construction

input : vertex set, edge set of the graphical interpretation of a recorded bus trip $q_j = \{p_1, ..., p_n\} \in Q$ of $n-1$ segments and bus route $r_i = \{p_1, ..., p_m\} \in R$ of $m-1$ segments.
output: Simple polygonal segmentation of the area between r_i and q_j

```
for each sequential GPS vertex p_i ∈ q_j do
    for Each segment s_i ∈ r_i do
        if (s_0 and p_0) or (s_{m-1} and p_n) then
            draw ĵ between the first sequential points of
              the first segments or the last sequential
              points of the last segments
        end
        if |proj_{s_i} p_i| < 0 then
            Check if the projection intersects with
              another segment
            if intersecting with other segment then
                move onto next sequential segment in r_i
            end
            else
                Connect p_i to endpoint of s_i with ĵ
            end
        end
        if ∃ valid proj_{s_i} p_i then
            Draw ĵ from p_i onto s_i
        end
        if |proj_{s_i} p_i| > |s_i| then
            Move to next sequential segment in r_i
        end
    end
end
```

Figure 4 shows the results of applying Algorithm 1 to r_i (green) and q_j (red) in a closed system $\in R^2$. By finding $\hat{j}_{proj_{s_i} p_i}$ for each recorded point $p_i \in q_j$, we can identify simple polygonal structures that are formed between the given trip q_j and it's proposed route r_i. The area of these polygons can be calculated using basic triangulation techniques and can express the error between q_j and r_i. Larger areas correspond to lower accuracy GPS recordings

Figure 5. **Left**: GTFS transit data (pink) of route 14 in Victoria, BC, Canada, overlaid on 50cm resolution SkySat imagery from Planet Labs. Points travelling north to south are on a wider roadway and have less immediate occlusions. As a result there is little to no measurement error and all points are confined to the correct lane. Comparatively points travelling from west to east are on a narrow roadway. There are more immediate occlusions which result in higher measurement errors (more severe spread of points across the width of the roadway). **Right**: GTFS transit data (green) of route 157 in Adelaide, Australia, overlaid on 0.5m resolution Pleiades-1A imagery from Satellite Imaging Corporation. There are natural obstructions lining the sides and the median of the roadway that can contribute to measurement errors in GPS recordings.

and may be correlated with higher levels of occlusion.

More specifically, we are interesting in comparing the degree of error from occlusion in city centres with large buildings (Area B in figure 1) to the error observed in surrounding suburban neighborhoods (Area C in Figure 1). By generating a sort of 'heat map' of areas with high rates of error, Unmanned Aerial Vehicles (UAVs) can plan a flight path that that accounts for the identified occlusions in order to more accurately track ground-based targets.

In order to demonstrate the potential GPS error due to occlusions, we plotted several thousand recorded GPS points along commonly travelled bus routes in both Victoria and Adelaide using high resolution satellite imagery (50cm resolution SkySat imagery from Planet Labs and 0.5m resolution Pleiades-1A imagery from Satellite Imaging Corporation, respectively) shown in Figure 5.

In the data-fused image of Victoria on the left of Figure 5, recorded points that are travelling from north to south are on a large, 4-lane road in the downtown core of the city. These points are confined to the two lanes travelling in the same direction and do not appear to be suffering from measurement error. However, points that are travelling from west to east are doing so along a more narrow roadway with more immediate occlusions. As a result, there is a more severe spread of data across the width of the roadway. Recorded points are not adhering to a specific lane and appear to be located on sidewalks and at the edges of buildings. Although we have not yet statistically quantified the severity of the spread of data on an occluded roadway, by plotting GTFS transit data we can see that measurement error is present.

In the data-fused image of Adelaide on the right of Figure 5, other obstructions are potentially the cause of GPS error: heavy tree coverage is lining the sides of the road as well as the median. Points span the width of the road and cross over the median in certain locations. As Bastani et al. [6] noted, with regards to inferring road networks from satellite imagery alone, occlusions such as tall buildings, trees, and shadows can make such a task very difficult. The use of GPS trajectories can be useful in areas where roads are not visible from aerial imagery, however this may not be the case when the GPS data is error prone as well.

3.1. Experimental Results

We applied linear regression techniques in order to quantify the observed GPS error in Figure 5 and compared these results to straight (no turns or deviations in bearing), non-occluded segments of the same bus routes. For each observed GPS recording, an equirectangular projection was made using the latidunial and longitudinal values for both the noisy and clean segments in each city. Since the areas of investigation are relatively small, we are not concerned with any distortion related to the method of projection [11]. In Table 1 we report calculations of error observed when linear regression was applied, and show the regression line fit to each respective segment in Figure 6. The values in Table 1 are given as arbitrary units that are consistent among themselves. We do not report the results as a measure of real world distance because we are comparing the residual of each GPS point from it's respective regression line, and not an actual road.

The Mean Absolute Error can be interpreted as the average of the absolute residual distance, where the residual is the vertical distance between the data point and the regression line. Since each segment is a straight road segment, each GPS point is expected to stay near the line of best fit. In the noisy segments from both Victoria and Adelaide it is clear that there are several data points with a large residual value. The Root Mean Square Error (RMSE) is the standard deviation of the residuals. The standard deviation calculates how much the data points are spread around the regression line. Again, noisy road segments have approximately a factor of 10 greater spread than their respective clean segments in comparison, as noted in Table 1.

3.2. Relevance to Algorithm & Impact of Research

Our statistical analysis used the concept of residuals in order to quantify both the average distance of GPS points from the regression line as well as the spread of points around the regression line. The idea was to demonstrate the fact that points recorded in occluded areas, on average, have a greater deviation from the line of best fit on which they should be recorded. Using Algorithm 1 given at the beginning of this section, we can treat the regression line as the road itself since we have chosen straight road segments. We could further assume that the ground truth bus route would not deviate from the road itself, meaning any points forming the ground truth trajectory would appear directly on the regression line. From this we can perform Iterative Projection Construction (Figure 4) and quantify the area between the observed GPS trajectory and the ground truth trajectory. This will give a more accurate indication as to where deviations within a segment occur and can be applied to an entire curved trajectory.

4. Discussion & Future Work

As mentioned previously, we are currently investigating the effect of certain metropolitan areas on the severity of sampling and measurement errors in GPS data. It should be noted that buildings in Victoria and Adelaide are not tall when compared with other city centres such as Vancouver, Toronto, or New York City. We believe that the level of error will increase as the average height of buildings that border a roadway also increase, but more research is required.

We have collected real-time data on the vehicle positions of BC Transit buses on routes around Victoria, BC, Canada, and Adelaide, Australia. Using high resolution satellite imagery of downtown Victoria and Adelaide, we can plot GPS data that aligns temporally with the moment the satellite image is captured. In doing so we can observe the visual location of buses within the image and compare that with their recorded GPS location. We believe that this will confirm our hypothesis that areas in Victoria with taller buildings and faster roads will create a larger disparity between a vehicle and it's GPS location when compared to suburban areas. The ability to generate error maps of expected GPS data will inform a larger project to develop mission planning tools to perform aerial data acquisition.

We are also exploring the effect of more specific inputs on predicting error-prone regions. Depending on the severity of obstructions in a given location, a GPS receiver may be communicating with fewer satellites when compared to optimal conditions. We would like to be able to collect data on the number of satellites visible by a vehicle at recorded locations in order to better express the likelihood of error for certain geographical regions in a city. This work will support our overall goal of expressing GPS error in the form of a spatially-based reliability map using open source data.

	Victoria, BC		Adelaide, SA	
	Noisy ($n = 491$)	Clean ($n = 239$)	Noisy ($n = 156$)	Clean ($n = 142$)
Mean Absolute Error	2.124×10^{-5}	7.410×10^{-6}	1.120×10^{-5}	2.958×10^{-6}
Root Mean Squared Error	2.999×10^{-5}	9.134×10^{-6}	1.773×10^{-5}	3.926×10^{-6}

Table 1. Results of applying linear regression to noisy (identified occlusions) and clean (non-occluded) straight road segments (no change in bearing) on bus routes in Victoria, BC, and Adelaide, SA. Truncated to the nearest thousandth.

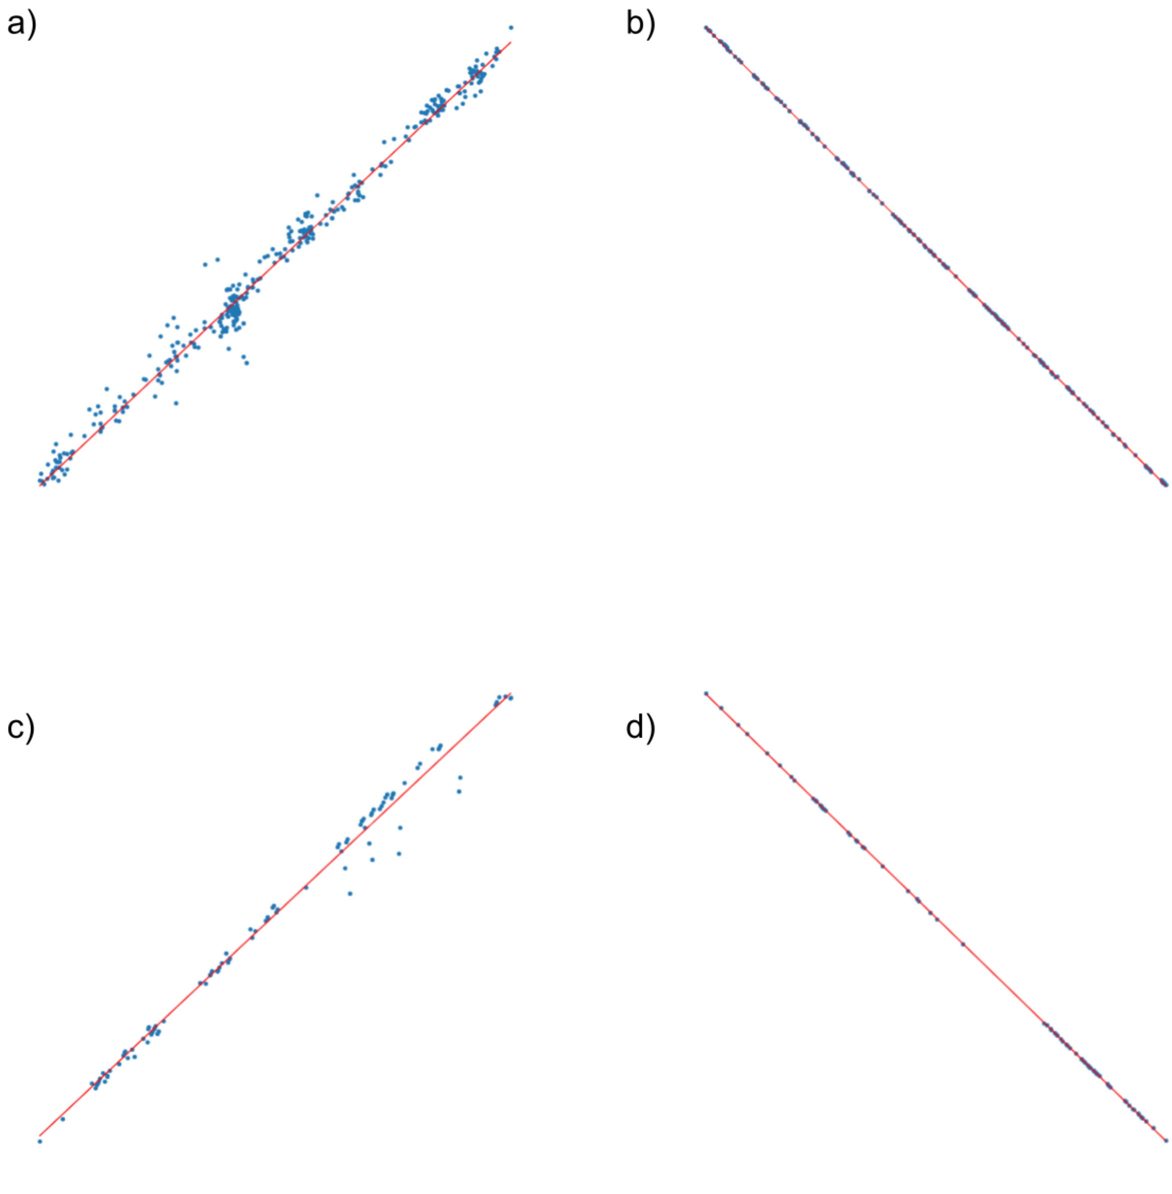

Figure 6. Regression Lines fit to segments of the 14 bus route in Victoria, BC (a) and b)), and the 157 bus route in Adelaide, SA (c) and d)). **a)** Regression of the data points observed in Figure 5 Left moving west to east. There is a greater MAE, RMSE when compared to **b)**, a straight road segment with no visible occlusions on the same bus route. Similarly, **c)** is the regression of the data points observed in Figure 5 Right. This road segment has natural obstructions that create visible errors in GPS when compared to a non-occluded segment **d)**. Again the noisy segment was found to have a greater MAE and RMSE.

References

[1] "GTFS Static Overview — Static Transit," *Google Developers.* https://developers.google.com/transit/gtfs.

[2] "Satellite Navigation with GPS." https://seos-project.eu/GPS/GPS-c01-p02.html (accessed Jul. 06, 2021).

[3] Chao, Pingfu, et al. "A Survey and Quantitative Study on Map Inference Algorithms from GPS Trajectories." *IEEE Transactions on Knowledge and Data Engineering*, 2020, pp. 1–1. *DOI.org (Crossref)*, doi:10.1109/TKDE.2020.2977034.

[4] He, Songtao, et al. "RoadRunner: Improving the Precision of Road Network Inference from GPS Trajectories." *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM,* 2018, pp. 3–12. *DOI.org (Crossref)*, doi:10.1145/3274895.3274974.

[5] Cao, Lili, and John Krumm. "From GPS Traces to a Routable Road Map." *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '09, ACM Press*, 2009, p. 3. *DOI.org (Crossref)*, doi:10.1145/1653771.1653776.

[6] Bastani, Favyen, et al. "RoadTracer: Automatic Extraction of Road Networks from Aerial Images." *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, IEEE,* 2018, pp. 4720–28. *DOI.org (Crossref)*, doi:10.1109/CVPR.2018.00496.

[7] S. He et al., "RoadRunner: improving the precision of road network inference from GPS trajectories," in *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, Seattle Washington, Nov. 2018, pp. 3–12. doi: 10.1145/3274895.3274974.

[8] Raymond, Rudy, and Takashi Imamichi. "Bus Trajectory Identification by Map-Matching." *2016 23rd International Conference on Pattern Recognition (ICPR), IEEE*, 2016, pp. 1618–23. *DOI.org (Crossref)*, doi:10.1109/ICPR.2016.7899868.

[9] Lyu, Cheng, et al. "A Partial-Fréchet-Distance-Based Framework for Bus Route Identification." *IEEE Transactions on Intelligent Transportation Systems*, 2021, pp. 1–6. *DOI.org (Crossref)*, doi:10.1109/TITS.2021.3069630.

[10] D. Kloster, "New tallest building on Island is in downtown Victoria," *Times Colonist*. https://www.timescolonist.com/real-estate/new-tallest-building-on-island-is-in-downtown-victoria-1.24159688.

[11] "Equirectangular—Help — ArcGIS for Desktop." https://desktop.arcgis.com/en/arcmap/10.3/guide-books/map-projections/equirectangular.htm (accessed Aug. 05, 2021).