Graceful Degradation for Real-time Visualization of Streaming Geospatial Data

João Rafael[1], João Moreira[1], Daniel Mendes[2], Mário Alves[1], and Daniel Gonçalves[1]

[1]INESC-ID Lisboa, Instituto Superior Técnico, University of Lisbon
[2]INESC TEC, Faculty of Engineering, University of Porto

Abstract

The availability of devices that can record locations and are connected to the Internet creates a huge amount of geospatial data that are continuously streamed. The informative visualization of such data is a challenging problem, given their sheer volume, and the real-time nature of the incoming stream. A simple approach like plotting all datapoints would generate visual noise, and not scale well. To tackle this problem, we have developed a visualization technique based on graceful degradation along three overlaid time periods (ongoing, recent, and history), each with a different visual idiom. A usability test of the proposed technique showed promising results.

CCS Concepts
• ***Human-centered computing*** → ***Visualization; Visual analytics; Geographic visualization;***

1. Introduction

Access to location data has become increasingly available. Such geographic data, available in large amounts and produced in real-time, can be analyzed using Infomation Visualization. Geodata can be divided into subgroups such as combinations of locations and time intervals or trajectories showing where someone has been. In some cases, geodata may show useful patterns when in great numbers. For instance, urban planners could study these patterns and identify roads that need more lanes due to being overly congested [DHZ*15]. Advertisers can study the most congested streets during rush hour and focus on those streets when putting up billboards. If people can visualize this data successfully and efficiently, they can find patterns and use this understanding for a myriad goals. However, visualizing geodata starts to be a computationally demanding process if they arrive in large amounts of data and in real-time.

Aggregation can help deal with large amounts of data, and enable an overall view of multiple movements' spatiotemporal distribution. Clustering algorithms function well in loading large amounts of data, which facilitates displaying macro information without data loss [AAB*13]. There are several aggregation techniques: spatial heatmaps [CHW*18] aggregate data into continuous surfaces, cluster heatmaps [YZTZ19] represent data in discrete squares, and density heatmaps [SWvv12] aggregate data into continuous surfaces. Specifically, in data streaming, several approaches have been used, such as an M-Kernel merging technique [ZCWQ03], SOMKE [CHM12], K-means [BF08], and heatmaps [DPMO12]. In data streaming, there is little research related to visualizing spatiotemporal data in real-time. The simplest way to represent this data is using points for each object's current location and lines representing where they have been [SKG*18]. However, since GPS data often have errors, it must undergo preprocessing [GSV17].

So, existing real-time visualizations focus on representations that do not scale visually (representing points and lines with no aggregation). Big Data visualization in ways that reduce clutter and perform well are restricted to static datasets. We aim to show data evolving in real-time, differentiating recent data from older data, yet maintaining an uncluttered map view. For this, we propose an adaptation of the *graceful degradation* concept [PMG19] to real-time streaming geospatial data visualization. This gradual simplification consists of dividing time into different periods. The data gets simpler when it transitions from more recent periods into older ones and using different ways to visualize each period. We have implemented a visualization prototype based on these ideas, with real-time test data, and conducted a user study to evaluate the approach's performance towards defined data analysis tasks.

2. Time periods and their representations

Towards our goal of visualizing large amounts of data in real-time, we divide time into several periods and represent each period using different visualization techniques. As data passes from one period

DOI: 10.2312/evs.20211058

to another, it will be gracefully degraded [PMG19], aggregating trajectories. It gradually reduces the content being shown, which decreases the amount of memory being used. After consideration, we conceived three periods: *ongoing*, *recent* and *history*. These are rendered in recency order. The more recent data will be drawn on top of older data, since the more recent data must stand out (Fig. 1). We used ColorBrewer [HB03] to select effective color schemes. For the *ongoing* period, all marks were orange. For the *recent* period, we used a sequential scale from blue (most recent) to green (less recent). We had a purple scale for the cluster heatmap of the history period. For the spatial heatmap, we chose purple and grey.

2.1. Time period *Ongoing*

The *ongoing* period is the most detailed, containing trajectories as they are received. It corresponds to the ones that have not ended. If an object stops, its trajectory until that point will no longer be part of this period and be moved to the recent period. If it starts to move again, its trajectory, which begins with the point where it last stopped, will be part of the *ongoing* period. This period's data are not reduced at all. Like Gomes et al. [GSV17], we show each object's current location with a small circle, and its complete trajectory is shown using lines (Fig. 2). To make it clear that data in the *ongoing* period correspond to something currently moving, we animate lines. Rather than appearing immediately, the animation mimics a movement from one point of the line to the other.

2.2. Time period *Recent*

The *recent* period covers an adjustable period (e.g., stopped trajectories to one hour ago or stopped trajectories to one day ago). Instead of keeping a trajectory's complete data, lines are simplified with trajectory-bundling. Each bundled line has information on how many lines are bundled in it. We encode this attribute with both width and color, with a color scale that ranges from blue to green (Fig. 3). The brighter green helps draw the user's eye immediately to streets with higher traffic. At the same time, width is used to make segments with low total thinner to not clutter the screen.

Figure 1: *Visualization with all periods. Orange circles and lines represent the ongoing period, green and blue lines the recent, and purple/grey areas the spatial heatmap encoding the history period.*

Figure 2: *Representation for the ongoing period. The circles represent the moving object, and the lines their current trajectory.*

In *recent* and *ongoing* periods, the size of lines and circles adapt to the map's zoom. Therefore, the lines are never too thin or too wide for the current zoom level. Besides, this adaptation will avoid occlusion problems. If the map is zoomed out and lines have a high stroke, they will cover each other. On the other hand, they will not convey enough detail with a zoom-in and a low stroke.

2.3. Time period *History*

The *history* period is the final accumulator for the data. Every line that is older than the recent period's limit is represented here. Rather than storing them like in other periods, we fit each one into a matrix grid encompassing the entire dataset's boundaries. For this period, we developed two representations: a **cluster heatmap**, a gird where each square's saturation increases for higher values; and a **spatial heatmap** of continuous surfaces created from interpolating discrete points (Fig. 4). This allowed us to explore a discrete approach and a continuous approach, respectively. In both representations, different resolution grids are used according to the map zoom (Fig. 5). When zoomed out, the visualization will show a lower-resolution version, while zoomed in shows a higher-resolution one. This helps the grid supplement the map's overall detail.

3. User Evaluation

We developed a functional prototype to find if users understood the difference in real-time between each data representation. Namely, if they were able to identify trends, and if they were able to understand how trends changed over time in streaming data. We also

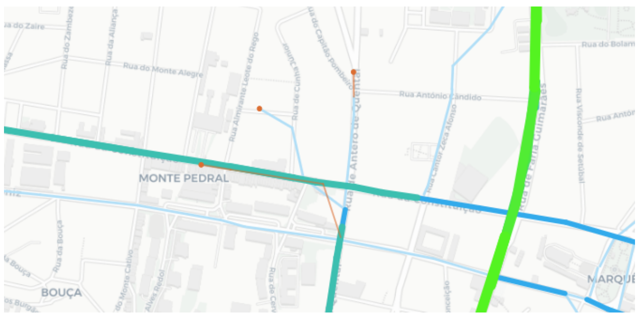

Figure 3: *Representation for the recent period. Blue lines represent newer data and green lines represent older data.*

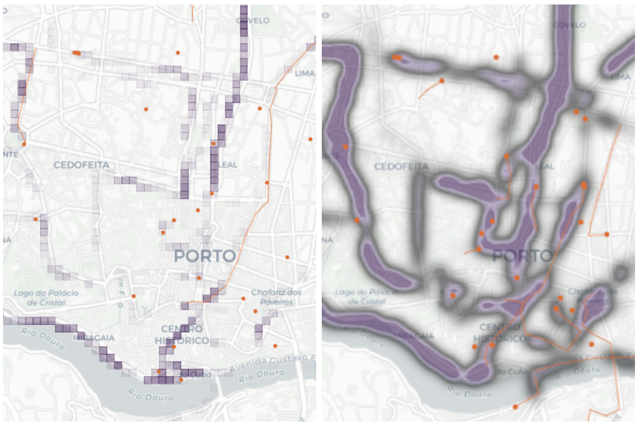

Figure 4: *Cluster heatmap (left) and spatial heatmap (right) representations for the history period. Higher color saturation corresponds to higher traffic.*

wanted to compare the two representations created for the *history* period. We used a dataset from a taxi company operating in Porto, Portugal [MGF*13]. The tests were performed remotely, sharing the screen and letting participants control the prototype remotely.

3.1. Method

We defined ten tasks according to two well-known types: analyze and search [TUoBC14]. Each participant was asked to complete all of them: **1.** "Follow a specific vehicle and point out when its trajectory has transitioned to the recent period"; **2.** "Indicate the street that was most-used recently"; **3.** "Indicate the area that has seen the most traffic overall (not just recently)"; **4.** "Explain whether, recently, there has been more traffic outside or inside the city"; **5.** "After viewing the prototype at two different times in the simulation, explain whether trends in traffic have changed between these two times or not"; **6.** "Indicate an area that was used in the past, but not recently"; **7.** "Go to 'Ponte da Arrábida' and indicate how many vehicles passed through it recently"; **8.** "Indicate one vehicle that does not follow the trends displayed by the visualization"; **9.** "A scenario will be shown. This scenario will have a gradual but significant change. When it happens, indicate it"; **10.** "Explore the visualization and make any observations".

Figure 5: *Different grid resolution in the history period using the cluster heatmap. As the user zooms out (left to right), the grid will decrease in resolution, thus the squares will increase in size in relation to the map.*

All tasks were performed using the functional prototype in real-time with streaming data, with the tree periods enabled all the time. We divided tasks into three sets with incremental difficulty. Their order in each set followed a Latin square distribution. They started with tasks 1 to 3, another order of tasks 4 to 8, then task 9, and finally task 10. Task 10 is a freeform task where the user can make any observations about anything they desire. Thus, this task did not have any correct answer. The users were asked to voice their thoughts about whatever parts of the map or trends they noticed until they felt they had no further observations to present. Tasks 3 and 6, which focus specifically on the *history* period, were performed twice to compare both representations' clarity and understandability. To avoid the order bias, each user experienced the representations in an alternated order.

For each task, we measured the elapsed time, the success or failure in completing it, and user classified its easiness in a Likert scale rating from 1 (very difficult) to 5 (very easy). After all tasks were performed, participants were asked to fill out a questionnaire with two sections. The first was about the System Usability Scale (SUS), and the second about the Raw NASA Task Load Index (TLX) tests.

3.2. Participants

A total of 21 people participated in the tests, resulting in 42 tests for tasks 3 and 6 (considering we did these tasks for both representations of the *history* period). Of these participants, 13 were male and 8 female, and the vast majority (18 participants) were between 18 and 25 years old. As for visual problems, all participants had normal or corrected to normal vision, and none were colorblind.

3.3. Task Performance

Firstly, we checked if there was a significant difference between the *history* period representations in tasks 3 and 6. We evaluated two sets of variables. In the first, we tested the elapsed time (Fig. 6), success, and each user's reported easiness (Tab. 1). In the second, we evaluated the rates given by participants for the representation, the transition between periods, and how easily they could distinguish between periods. Tasks 3/C and 6/C correspond to Tasks 3 and 6 with the cluster heatmap representation of the *history* period, and 3/S and 6/S to the same tasks with the spatial heatmap.

To analyze each variable in both sets, we performed the Friedman test. In case this test indicated significant differences, we performed the posthoc Wilcoxon Signed Rank test with the Bonferroni correction. To evaluate success, we used McNemar's test. In the first set, after comparing the two different representations, the only significant difference was found in the easiness of task 3 ($p = 0.046$). The **spatial heatmap** representation was the most successful overall. In the second set, again only in task 3, there was a significant difference in all variables (with p-values of 0.003, 0.018, and 0.013, respectively). Once again, the **spatial heatmap** representation was the most successful overall.

In our second statistical analysis, we checked if there was a significant difference between all ten tasks. This time, we evaluated the first set of variables (time, success, and easiness). All statistical tests were the same as with our first analysis. In both time and easiness, there was a significant difference across tasks ($p < 0.0005$). In

Task	T1	T2	T3/S	T3/C	T4	T5	T6/S	T6/C	T7	T8	T9
Success Rate	100%	100%	86%	71%	86%	100%	100%	100%	100%	100%	86%
Easiness	5 (0)	5 (0)	4 (0.25)	4 (1)	5 (0)	5 (0.25)	5 (0.25)	5 (1)	5 (0)	5 (0)	4 (2)

Table 1: *Success rate and reported easiness (median and interquartile range; 1 - very difficult, 5 - very easy) for each task.*

completion time, task 1 was significantly different from all others. It took the shortest time to perform. On the other side of the spectrum, task 10 was also significantly different, taking the longest time. Regarding significant differences in the remaining tasks, only **task 3** (both representations) and **task 9** stand out with high completion times. Then, in the easiness variable, the Friedman test revealed significant differences ($p < 0.0005$). The posthoc showed them to be between two groups of tasks: the first with tasks 1, 2, 7, and 8, and the second with task 3 (both representations) and task 9, which had higher easiness values. Therefore, **task 3** (both) and **task 9** were considered to be more challenging.

3.4. Usability and Task Load

Our third analysis focused on the System Usability Scale (SUS) and the NASA Task Load Index (TLX) results. On the former, having a score above 87.5 corresponds to the top-fourth of responses [BKM08]. The average SUS score for our prototype was 90.12, which matches an adjective rating of **Excellent** and only 0.78 points below the Best Imaginable rating. Regarding individual questions, the most positive ratings were in the fifth (how well integrated the system's functions were) and eighth (how complicated it was to use). The most varied answers were in the fourth (whether the user would require a technician's help to use the system).

The NASA TLX of our work was 22.58. Participants felt they performed the tasks well. The Performance sub-scale had a median of 3 (1 - Success or good performance, 20 - Failure or poor performance), and a score of 12.86. Frustration was also low (16.67) and Physical Demand was down too. The highest demand was Mental Demand, scoring 32.86. Finally, Temporal Demand and Effort scored 24.29 and 30.24, respectively.

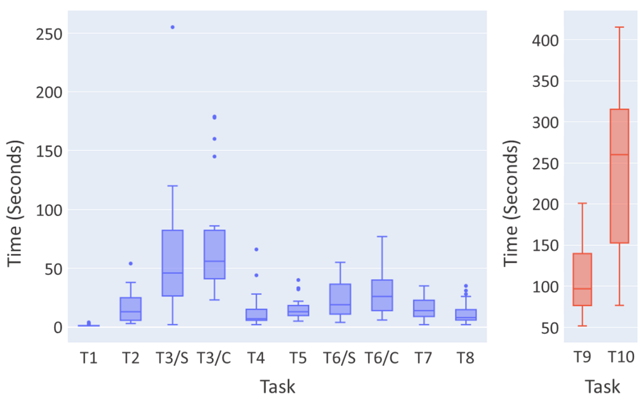

Figure 6: *Completion time for each task.*

3.5. Discussion

We can conclude that users were generally able to understand how the functional prototype worked, the differences between each period, and how trends could evolve over time. Our goal was to make a visually clear and easily understandable way to show streamed trajectory data, and participants were indeed not confused with our representations. Given that the *ongoing* period is the simplest, it is no surprise that it was very well understood. Tasks 1 and 8 focused on it and had excellent results. Participants also had no difficulty understanding the *recent* period: tasks 2, 4, 5, 6, and 7 were focused on this and had good results. The *history* period was somewhat harder to understand, with tasks 3 and 6 having lesser results. Comparing the different representations for this period, most participants preferred the spatial heatmap, as it was easier to use.

Participants had some difficulties analyzing how trends changed over time. While task 5 had excellent results, task 9 proved to have significant differences compared to most other tasks. However, the results were still satisfactory, given that task 9 is the most complex. Task 10 was focused on analysing the visualization freely. Participants' observations provided useful information regarding how they interpreted the visualization, comparing the three different periods to identify the trends being shown. Considering the SUS and NASA TLX scores, we can conclude that our approach's usability was overall excellent.

4. Conclusion and Future Work

Analyzing streaming geospatial data in real-time can be challenging due to the huge amount information that is evolving and needs to be visualized. To tackle this challenge, we resorted to the graceful degradation concept. We conceived three periods (*ongoing*, *recent*, and *history*) that aggregate and display data differently. For the *history* period, we developed two alternative visual idioms. Through a user evaluation, we found out that our approach showed very promising results. The *ongoing* and *recent* periods were both very easily understood. Still, the representations for the *history* period were slightly harder to understand. For future work, we think the graceful degradation concept could be further explored. Our modular period concept could be built upon to incorporate periods with different data simplifications or other representations. Also, more granular forms of degradation could be developed, allowing for more and configurable periods.

Acknowledgments

This work was partially supported by FCT through projects PTDC/CCI-CIF/28939/2017, UIDB/50021/2020, SFRH/BD/143496/2019 and POCI-01-0145-FEDER-030740 – PTDC/CCICOM/30740/2017.

References

[AAB*13] ANDRIENKO G., ANDRIENKO N., BAK P., KEIM D., WROBEL S.: *Visual Analytics of Movement*. Springer, Berlin, Dec 2013. doi:10.1007/978-3-642-37583-5. 1

[BF08] BARBAKH W., FYFE C.: Online clustering algorithms. *International Journal of Neural Systems 18*, 03 (2008), 185–194. 1

[BKM08] BANGOR A., KORTUM P. T., MILLER J. T.: An empirical evaluation of the system usability scale. *International Journal of Human–Computer Interaction 24*, 6 (2008), 574–594. URL: https://doi.org/10.1080/10447310802205776, doi:10.1080/10447310802205776. 4

[CHM12] CAO Y., HE H., MAN H.: Somke: Kernel density estimation over data streams by sequences of self-organizing maps. *IEEE transactions on neural networks and learning systems 23*, 8 (2012), 1254–1268. 1

[CHW*18] CHEN W., HUANG Z., WU F., ZHU M., GUAN H., MACIEJEWSKI R.: VAUD: A visual analysis approach for exploring spatio-temporal urban data. *IEEE Transactions on Visualization and Computer Graphics* (Sep. 2018). doi:10.1109/TVCG.2017.2758362. 1

[DHZ*15] DENG Z., HU Y., ZHU M., HUANG X., DU B.: A scalable and fast optics for clustering trajectory big data. *Cluster Computing* (Jun 2015). URL: https://doi.org/10.1007/s10586-014-0413-9, doi:10.1007/s10586-014-0413-9. 1

[DPMO12] DUCHOWSKI A. T., PRICE M. M., MEYER M., ORERO P.: Aggregate gaze visualization with real-time heatmaps. In *Proceedings of the symposium on eye tracking research and applications* (2012), pp. 13–20. 1

[GSV17] GOMES G. A. M., SANTOS E., VIDAL C. A.: Interactive visualization of traffic dynamics based on trajectory data. In *2017 30th SIBGRAPI Conference on Graphics, Patterns and Images* (Oct 2017). doi:10.1109/SIBGRAPI.2017.21. 1, 2

[HB03] HARROWER M., BREWER C. A.: Colorbrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal 40*, 1 (2003), 27–37. 2

[MGF*13] MOREIRA-MATIAS L., GAMA J., FERREIRA M., MENDES-MOREIRA J., DAMAS L.: Predicting taxi–passenger demand using streaming data. *IEEE Transactions on Intelligent Transportation Systems 14*, 3 (2013), 1393–1402. 3

[PMG19] PIRES G., MENDES D., GONÇALVES D.: Vismillion: A novel interactive visualization technique for real-time big data. In *2019 International Conference on Graphics and Interaction (ICGI)* (2019), IEEE, pp. 86–93. 1, 2

[SKG*18] STEPTOE M., KRÜGER R., GARCIA R., LIANG X., MACIEJEWSKI R.: A visual analytics framework for exploring theme park dynamics. *Transactions on Interactive Intelligent Systems* (Feb. 2018). URL: http://doi.acm.org/10.1145/3162076, doi:10.1145/3162076. 1

[SWvv12] SCHEEPENS R., WILLEMS N., VAN DE WETERING H., VAN WIJK J.: Interactive density maps for moving objects. *IEEE Computer Graphics and Applications* (Jan 2012). doi:10.1109/MCG.2011.88. 1

[TUoBC14] TAMARA (UNIVERSITY OF BRITISH COLUMBIA VANCOUVER C. M.: *Visualization Analysis and Design*. Taylor and Francis Inc, Nov. 2014. URL: https://www.ebook.de/de/product/22371060/tamara_university_of_british_columbia_vancouver_canada_munzner_visualization_analysis_and_design.html. 3

[YZTZ19] YANG C., ZHANG Y., TANG B., ZHU M.: Vaite: A visualization-assisted interactive big urban trajectory data exploration system. In *2019 IEEE 35th International Conference on Data Engineering* (April 2019). doi:10.1109/ICDE.2019.00233. 1

[ZCWQ03] ZHOU A., CAI Z., WEI L., QIAN W.: M-kernel merging: Towards density estimation over data streams. In *Eighth International Conference on Database Systems for Advanced Applications, 2003.(DASFAA 2003). Proceedings.* (2003), IEEE, pp. 285–292. 1