

# VISTA: A visual analytics platform for semantic annotation of trajectories

Amílcar Soares  
Institute for Big Data Analytics  
Halifax, Canada  
amilcar.soares@dal.ca

Jordan Rose  
Institute for Big Data Analytics  
Halifax, Canada  
jrose@dal.ca

Mohammad Etemad  
Institute for Big Data Analytics  
Halifax, Canada  
etemad@dal.ca

Chiara Renso  
ISTI-CNR  
Pisa, Italy  
chiara.renso@isti.cnr.it

Stan Matwin  
Institute for Big Data Analytics  
Halifax, Canada  
stan@dal.ca

## ABSTRACT

Most of the trajectory datasets only record the spatio-temporal position of the moving object, thus lacking semantics and this is due to the fact that this information mainly depends on the domain expert labeling, a time-consuming and complex process. This paper is a contribution in facilitating and supporting the manual annotation of trajectory data thanks to a visual-analytics-based platform named VISTA. VISTA is designed to assist the user in the trajectory annotation process in a multi-role user environment. A session manager creates a tagging session selecting the trajectory data and the semantic contextual information. The VISTA platform also supports the creation of several features that will assist the tagging users in identifying the trajectory segments that will be annotated. A distinctive feature of VISTA is the visual analytics functionalities that support the users in exploring and processing the trajectory data, the associated features and the semantic information for a proper comprehension of how to properly label trajectories.

## 1 INTRODUCTION

The increasing access to positioning devices technologies such as smart-phones, GPS-enabled cameras and sensors, resulted in vast volumes of mobility data collected, stored and available for analysis. Such mobility data are typically modeled as streams of spatio-temporal points, called trajectories. There is a growing research interest in analysis methods for semantically enriched trajectories [7, 8]. However, only a few datasets containing semantically labeled trajectory data are available. This is primarily because the semantic labels tend to indicate a specific behavior whose identification depends on the humans' interpretation: in fact, the understanding of the moving object behavior depends on multiple factors and, in most of the cases, it cannot be automatically inferred. The process of manual annotation may be, therefore, complex and time-consuming, even by domain experts who might be uncertain in the correct data interpretation or they might disagree. Indeed, the labeling process is complex due to a number of factors: (1) the annotator user needs to have an immediate view and understanding, not only of the spatio-temporal trajectory details and the numerical features deriving from them but also of the contextual semantic information; (2) different annotators might have different roles and interpretations on how to label trajectories (e.g., "is the ship fishing or not?"); (3) not

only the whole trajectory might express a single behavior, very often the parts of a trajectory have to be individually labeled since different behavior may co-exist during the same journey (e.g., a vessel is first fishing then traveling to harbor). How to correctly identify the switch points from one behavior to the other is another critical issue to be solved. We believe that these issues might be alleviated by a proper trajectory annotation platform that effectively assists the user in coping with all these aspects during the semantic labeling process.

A promising approach is to build a platform supporting the annotation by exploiting concepts from Visual Analytics. The idea of Visual Analytics is to create systems that enable analytical reasoning about complex problems with the goal of making the data processing and the inferred information clear and evident for the analysis [5]. This is accomplished by integrating data processing methods with visualizations designed to assist users in making efficient decisions. We can observe that the human mind generally conveys information more effectively through visualization than only relying on textual and numerical data. Combining visualization with user interactions enables a system to explore the data from different perspectives, thus linking and combining distinct information pieces to derive new insights from the data.

Pioneering works on semantic trajectories annotations [6, 10] used a predefined set of rules established by domain experts to automatically assign the semantic labels to the whole trajectories or to segments of trajectories. Although these methods work well for many scenarios, they are not suitable in case there are no clear criteria for identifying the object behavior and/or the segment to be labeled. In this case, we need the human intervention to establish the most appropriate label for each trajectory (or trajectory segment) based on both numerical and semantic features. Recently, some works proposed methods to simplify the manual annotation task like in [9] where a web interface has been proposed to upload personal trajectories and annotate each segment with the activity performed by the user. However, in this case, all trajectory segments need to be annotated directly by the traveling user, and this becomes unfeasible when the number of trajectory segments is very high and/or the annotator does not represent the entity which performed the movement like in the case of vehicles, vessels or animals. To cope with this problem, other methods propose machine learning approaches (e.g. semi-supervised or active learning) to automatically classify trajectories into semantic labels by starting from a small set of manually annotated traces [3, 4]. All these approaches must rely on an accurately labeled dataset to reach a good classification accuracy. Therefore, there is a strong need for supporting the

manual annotation process which can lead to good quality annotated datasets and therefore reliable analysis findings. To the best of our knowledge, no approach combines visual analytics with trajectory processing functions to assist the user in the process of manually tagging trajectory data.

The challenge of a visual analytics trajectory annotation system is to provide efficient and effective support for the user interaction, helping the user in focussing on each specific annotation task, highlighting the features values for each trajectory point or segment and properly visually combine contextual knowledge. Specifically, the annotators need to use their domain knowledge for thinking, creating associations, and generating insights from the trajectory data. On the other hand, the system also needs capabilities for processing and aggregating trajectory data. Deciding where to set the segmentation point in a trajectory is very challenging because it will directly affect the values of the segments' features and therefore the overall labeling of the trajectory dataset. We propose VISTA as an interactive visual user interface integrating spatio-temporal processing capabilities to play an essential role in semantic trajectory annotation. With VISTA, we provide full support to the manual trajectory annotation by tailoring a visual analytics platform that guides the user in this process. We strongly believe that VISTA can significantly contribute to the scientific community in supporting the domain experts in the annotation process and produce more reliable semantically labeled trajectory datasets available for analysis.

## 2 SYSTEM ARCHITECTURE

The architecture of VISTA has been designed to provide solutions to the three issues introduced in the previous section, namely the immediate understanding, the different users' role and the support for the segment switch points identification.

A distinctive feature of VISTA is the possibility to handle a multi-user annotation process where the users might have different roles. VISTA supports both the creation of an annotation session (e.g., upload trajectories, contextual information, feature creation, etc) by the *session manager* and the annotation process itself by the *annotator* (e.g., user analyses a trajectory, provide partitioning positions, compare the segments, annotates a trajectory, etc).

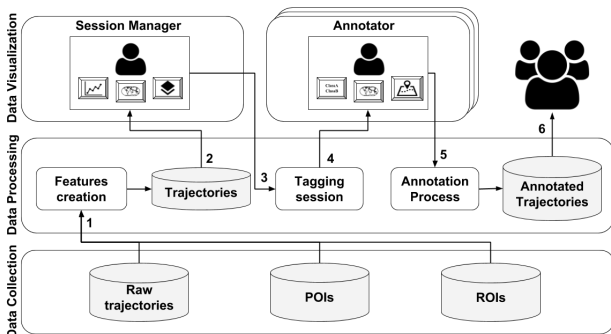


Figure 1: VISTA architecture and workflow.

The architecture and workflow overview is illustrated in Figure 1. As we can see, it is organized into three main components: (i) the data collection to handle raw trajectory and contextual geographical information like Point of Interests (POI) and Region Of Interests (ROI), (ii) the Data Processing that deal with trajectory and annotation data, and (iii) the Data Visualization to

interact with the users in the processes of both creating a tagging session, detect the switch points and annotating trajectories. The workflow of VISTA to perform trajectory annotation includes six steps depicted as the numbers of the arrows of Figure 1. In the first step, the session manager is requested to set the stage for the annotation process, namely upload raw trajectories and the POIs and ROIs that are relevant to the studied domain. In the following step (Step 2), the data processing engine automatically creates numerical features related to each trajectory point, called *point features* like the distance traveled, the estimated speed, the bearing, the bearing rate, and the acceleration. The engine also calculates the shortest distances to a POIs, and detects when the trajectory points intersect with ROIs. The shapely library<sup>1</sup> was used to calculate the shortest distance and intersections. At this point, each trajectory uploaded by the user is stored as a document in a MongoDB<sup>2</sup> collection. In step 3, VISTA displays the trajectories associated to these new features to the session manager whose tasks are: (1) the data exploration for the selection of the features that are relevant for the tagging session, (2) the creation of the annotation classes (i.e. labels) that must be used in a tagging session, and (3) the invitation to the annotator users to participate in a tagging session. With a tagging session created, it is now possible for the invited annotators to start the tagging process of trajectory data. In step 4, the annotators explore the trajectory data and create the trajectory segments that must reflect the annotation classes available for tagging. Several visualization functionalities are available for the process of tagging and are detailed in Section 3. After going through all trajectories and tagging their segments with the labels (Step 5), a dataset with annotated trajectories will be available for all the users to download (Step 6).

As we discussed above, a crucial step in annotating trajectory data is to determine the parts of the trajectory where the moving object's behavior changes. Detecting such switches in the object movement behavior is challenging for an annotating user since there is a need to explore a possibly high number of features to precisely determine where and when the object behavior changed. This is done through a process usually called *segmentation*. Segmenting a trajectory means to find the partitioning points that are used to create trajectory segments, or sub-trajectories, characterized by the fact that each segment has uniform behavior respect to some criteria [2]. Once these segments have been identified, additional numerical features like average, median, standard deviations, and percentiles may be created to better characterize the behavior of the moving object in that segment (the so-called *segment features*).

## 3 DEMONSTRATION

The objective of our demonstration is to involve the user in the VISTA tagging experience by providing a trajectory dataset to be annotated together with semantic contextual information. We have selected two datasets: (1) 10 trajectories of vessels that should be annotated as "fishing" and "not fishing" activity; (2) 20 trajectories of people that should be annotated with transportation means as "walking", "bike", "train", "bus", and "car". During the demo the authors will guide attendees through the whole process where the user will experience both roles, the *session manager* and the *annotator*, whose tasks and relative interfaces are detailed in the next sections.

<sup>1</sup><https://pypi.org/project/Shapely/>

<sup>2</sup><https://www.mongodb.com/>, version 3.6.2

### 3.1 Session Manager

The role of this user is setting the stage for the actual annotation process. This is done through four screens. In the first screen, the user is requested to create and give a name to a new tagging session or select a session that was previously created. In the second screen, the user is requested to upload raw trajectories in delimited separated file format (e.g., CSV, TSV, etc) and map the file fields to the columns representing the raw trajectory data: *trajectory\_id*, *time*, *latitude*, *longitude*. In the third screen, the user is asked to upload POIs and/or ROIs that are relevant to the domain. Then, VISTA executes a process in the background to create the following point features for each raw trajectory point: time difference from the previous point, distance traveled from the previous point, speed, acceleration, bearing, jerk, bearing rate, and rate of bearing rate. The relative computation formulas and details can be found in [1]. The platform also will create other two point features using POIs and ROIs: (1) for every POIs layer, the platform will calculate the shortest distance to a POI on the particular layer; (2) for ROIs layers, the platform will verify if a trajectory point intersects a ROI in the specific layer. Finally, in the fourth screen, the user will create the labels that must be used by annotators and invite annotators in the trajectory tagging session by providing their emails. This sequence of actions represents steps 1, 2 and 3 of Figure 1. After this step, the annotators will receive a notification with the invitation to enter the tagging session.

### 3.2 Annotator

When a user receives an invitation for a tagging session, he/she can start to tag trajectories. The process of annotating trajectories is iterative and interactive, where a single trajectory is presented to the user and further explored in each iteration with the system. We recall that the two-fold objective of a tagging session is (1) to identify the segments of the trajectory with a uniform behavior, or, in other words, identify the change points and (2) actually tag the segment with the appropriate label. In VISTA, the annotator has access a dashboard providing tools to understand the behavior of the moving object, depicted in Figure 2. We observe that the main screen is divided into two interactive panels: (i) a map on the left and (ii) summary statistics on the right. The map panel visualization needs to be effective in showing the actual movement of the trajectory with point and trajectory features. For this reason, we implemented two visualization solutions to support the annotator in understanding the movement data: first, the actual moving object movement can be played, dynamically showing how the moving object performed its movement in a particular region. Second, the line colors are displayed in saturation grades of red, reflecting the value of the point feature selected on the right side of the screen; a low value is colored with a light intensity of red, and higher values are colored with a more intense red. There is an automatic interaction between the two panels: the user can select a segment and/or point features in the right panel and the relative parts in the map on the left are highlighted. Conversely, the annotator can select a trajectory point in the map, and the corresponding segment and point feature are highlighted in the right panel. On the bottom of the map, the red color legend is presented to the annotator user to have an immediate perception of what is happening between the points sequence. Inside the map and in the top left, VISTA provides some typical Geographical Information System (GIS) visualization options, such as zoom-in and zoom-out, display or

hide POIs and ROIs, and change the colors of the annotation classes (e.g., fishing or not-fishing).

In more detail, the right panel provides the following options and statistics: (i) On the top, it displays a summary computed from the trajectory data with its total number of points, total distance travelled in meters and the average sampling rate between the trajectory points; (ii) at the bottom, the user may choose to visualize the data as a line chart that shows the values of the point features following the temporal order or a scatter plot where the user can try to find correlations between two point features.

When a trajectory is shown on the screen, the annotator user is requested to provide the partitioning positions to split the trajectory into the segments representing uniform behavior to be assigned to the class labels. This is the most challenging functionality to develop since the correct segmentation is fundamental for a good quality annotation process. However, identifying the switch points is particularly difficult since many different aspects have to be considered at the same time. It is, therefore, challenging to support the user in this "multi-dimensional view" where the most crucial aspects have to be considered jointly. For this reason, we have created a drag and drop tool on the map where the user can add a new partitioning point to the screen, assign to a class label and pin exactly the switch point where to split the trajectory. By adding a new partitioning position, all the statistics regarding the classes of the trajectory have to be automatically recomputed since new labeled segments are provided. This is captured by the two panels described above. First, on the map panel, the colors of the trajectory segments change according to the class provided. Second, the colors of both the scatter plots and the line chart change to reflect the new information provided by the user. We also provide statistical measures regarding the trajectory segments and their classes. Trajectories are sent to user in a sequential way, one per interaction, which the user's objective is to segment it using the drag and drop pins provided in the left side of the panel. When a trajectory is completely annotated, the user can press the Next button on the bottom of the screen to receive a new trajectory to be annotated.

### 3.3 Summarizing annotations from users

When all the users conclude the annotation process, a summarizing screen (Figure 3) is created with the objective of exploring how the different users tagged the data. In VISTA, the users can confront their tagging with other annotators by exploring how they annotated their dataset and how the results of their tagging session are similar or not to the other users. The tagging session results can be compared for both point and segment features. Figure 3 shows the main screen with two panels: (1) on the left, the user is able to analyze statistics (e.g. minimum, maximum, average, standard deviation, and percentiles) of all point features available on the platform; (2) on the right, the user may want to analyze statistics regarding the segment features. For both charts, the average behavior per user and class are plotted with the objective of understanding if the users most likely agree or disagrees regarding some feature.

## 4 CONCLUSIONS AND FUTURE WORKS

We have proposed VISTA, an interactive tool based on visual analytics principles, supporting the users in semantically annotate trajectory data. A distinctive feature of VISTA is the support for the identification of trajectory segments and the assignment to the relative semantic label. We intend to expand this

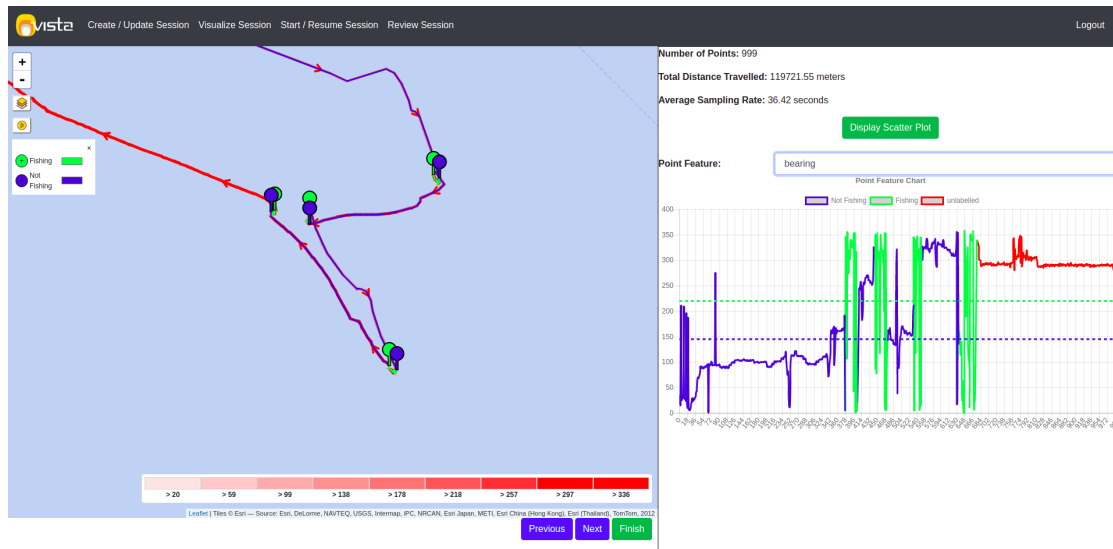


Figure 2: Elements of the annotator user dashboard with the vessels trajectories dataset.

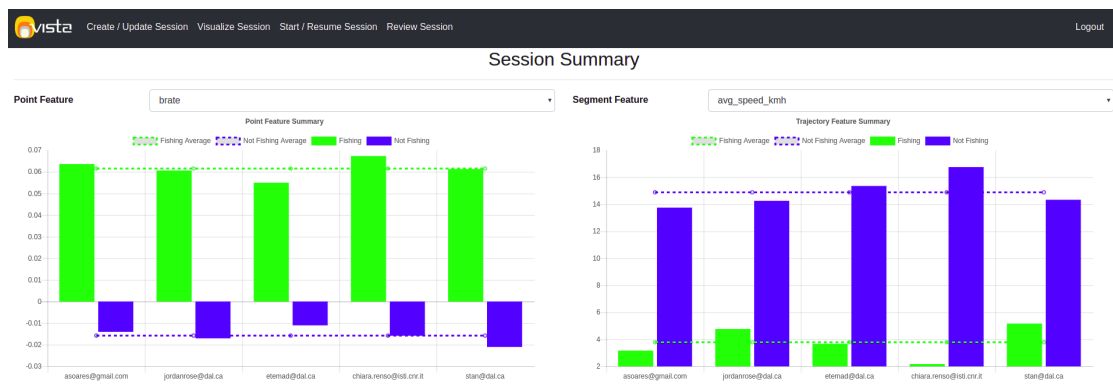


Figure 3: Summary results with the users' annotations.

platform into two directions. First, we want to create a module that automatically suggests how to segment the trajectory by learning from the previous interactions with the platform. Second, we intend to improve the results comparing how the labels have been assigned by the different users, highlighting when users agree or disagree in identifying a specific behavior of a moving object. The tool is available for testing at the URL <https://bigdata.cs.dal.ca/resources>.

## ACKNOWLEDGEMENTS

This work has been supported by the MASTER project that has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement N. 777695. The authors would also like to thank NSERC (Natural Sciences and Engineering Research Council of Canada) for financial support.

## REFERENCES

[1] Mohammad Etemad, Amílcar Soares Júnior, and Stan Matwin. 2018. Predicting Transportation Modes of GPS Trajectories using Feature Engineering and Noise Removal. In *31st Canadian Conference on Artificial Intelligence, Canadian AI 2018, Toronto, Canada, May 8–11*. Springer, 259–264.

[2] Amílcar Soares Júnior, Bruno Neiva Moreno, Valéria Cesário Times, Stan Matwin, and Lucídio dos Anjos Formiga Cabral. 2015. GRASP-UTS: an algorithm for unsupervised trajectory segmentation. *International Journal of*

*Geographical Information Science* 29, 1 (2015), 46–68.

[3] Amílcar Soares Júnior, Chiara Renso, and Stan Matwin. 2017. ANALYTIC: An Active Learning System for Trajectory Classification. *IEEE computer graphics and applications* 37, 5 (2017), 28–39.

[4] Amílcar Soares Junior, Valeria Cesario Times, Chiara Renso, Stan Matwin, and Lucidio A. F. Cabral. 2018. A Semi-Supervised Approach for the Semantic Segmentation of Trajectories. In *2018 19th IEEE International Conference on Mobile Data Management (MDM)*. 145–154.

[5] Daniel Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, Jörn Kohlhammer, and Guy Melançon. 2008. Visual analytics: Definition, process, and challenges. In *Information visualization*. Springer, 154–175.

[6] Christine Parent, Stefano Spaccapietra, Chiara Renso, Gennady Andrienko, Natalia Andrienko, Vania Bogorny, Maria Luisa Damiani, Aris Gkoulalas-Divanis, Jose Macedo, Nikos Pelekis, et al. 2013. Semantic trajectories modeling and analysis. *ACM Computing Surveys (CSUR)* 45, 4 (2013), 42.

[7] Christine Parent, Stefano Spaccapietra, Chiara Renso, Gennady L. Andrienko, Natalia V. Andrienko, Vania Bogorny, Maria Luisa Damiani, Aris Gkoulalas-Divanis, José Antônio Fernandes de Macêdo, Nikos Pelekis, Yannis Theodoridis, and Zhixian Yan. 2013. Semantic trajectories modeling and analysis. *ACM Comput. Surv.* 45, 4 (2013), 42:1–42:32. <https://doi.org/10.1145/2501654.2501656>

[8] Chiara Renso, Stefano Spaccapietra, and Esteban Zimanyi (Eds.). 2013. *Mobility Data: Modeling, Management, and Understanding*. Cambridge Press.

[9] Salvatore Rinzivillo, Fernando de Lucca Siqueira, Lorenzo Gabrielli, Chiara Renso, and Vania Bogorny. 2013. Where Have You Been Today? Annotating Trajectories with DayTag. In *Advances in Spatial and Temporal Databases - 13th International Symposium, SSTD 2013, Munich, Germany, August 21-23, 2013. Proceedings*. 467–471. [https://doi.org/10.1007/978-3-642-40235-7\\_30](https://doi.org/10.1007/978-3-642-40235-7_30)

[10] Zhixian Yan, Dipanjan Chakraborty, Christine Parent, Stefano Spaccapietra, and Karl Aberer. 2011. SeMiTri: a framework for semantic annotation of heterogeneous trajectories. In *Proceedings of EDBT*. ACM, 259–270.