

Wise Sliding Window Segmentation: A classification-aided approach for trajectory segmentation

Mohammad Etemad¹, Zahra Etemad³, Amílcar Soares¹, Vania Bogorny⁴, Stan Matwin^{1,2}, and Luis Torgo¹

¹ Institute for Big Data Analytics, Dalhousie University, Halifax,

² Institute for Computer Science, Polish Academy of Sciences, Warsaw

³ Bu-Ali Sina University, Hamedan, Iran

⁴ PPGCC - Universidade Federal de Santa Catarina (UFSC), Brazil

Abstract. Large amounts of mobility data are being generated from many different sources, and several data mining methods have been proposed for this data. One of the most critical steps for trajectory data mining is segmentation. This task can be seen as a pre-processing step in which a trajectory is divided into several meaningful consecutive subsequences. This process is necessary because trajectory patterns may not hold in the entire trajectory but on trajectory parts. In this work we propose a supervised trajectory segmentation algorithm, called Wise Sliding Window Segmentation (WS-II). It processes the trajectory coordinates to find behavioral changes in space and time, generating an error signal that is further used to train a binary classifier for segmenting trajectory data. This algorithm is flexible and can be used in different domains. We evaluate our method over three real datasets from different domains (meteorology, fishing, and individuals movements), and compare it with four other trajectory segmentation algorithms: OWS, GRASP-UTS, CB-SMoT, and SPD. We observed that the proposed algorithm achieves the highest performance for all datasets with statistically significant differences in terms of the harmonic mean of purity and coverage.

Keywords: Trajectory Segmentation · Spatio-temporal Segmentation · Trajectory Partition · Supervised Trajectory Segmentation

1 Introduction

An essential task for mobility data mining is trajectory segmentation. Different to classical data mining, in trajectory data mining, the attributes/features are extracted from subtrajectory parts. The partitioning is necessary because a mobility pattern, in general, does not hold for the entire trajectory, but for subtrajectory parts. Therefore, the segmentation process becomes one of the most critical pre-processing steps for trajectory data mining.

Trajectory segmentation is the process of splitting a given trajectory into several homogeneous segments regarding some criteria. This task plays a pivotal

role in trajectory mining since it affects the features of each segment, as the features may depend on the size of the trajectory segment, independently of the application domain, such as fishing detection [15], animal behavior [8, 7], tourism [6], traffic dynamics [5, 7, 16, 13], vessel movement patterns [2] etc.

A trajectory is a sequence of points located in space and time, and different criteria can be used to split trajectories. There are several approaches that can be used for trajectory segmentation such as CB-SMoT [12], SPD [17], WK-Means [10], GRASP-UTS [15], TRACLUS [9], OWS [4], etc. Different to previous approaches where no training step is performed, we propose in this paper a supervised strategy to segment trajectory data. To the best of our knowledge this is the first approach that actually learns partitioning positions (i.e., the last trajectory point of a segment) from trajectory data characteristics for a given application domain. The main advantage of this supervised strategy is that the transitioning characteristics of a behavior change can be learned from the training data. The model built to forecast partitioning positions is further used to segment trajectories. After that, a majority vote strategy decides the proper location to place a partitioning position.

In summary, the main contributions of this work include: (i) a method for producing training data from partitioning positions on a labeled trajectory; (ii) a method to decide when a partitioning position occurs in a trajectory; and (iii) an empirical study comparing WS-II and several baselines for segmentation.

This paper is organized as follows. Section 2 shows the definitions necessary to describe our trajectory segmentation method. In Section 3, the related works are described. In Section 4, we propose WS-II with details. In Section 5, we applied the proposed method and other trajectory segmentation algorithms on three datasets and reported their performance results. Finally, we conclude our work in Section 6.

2 Definitions

In this section we present the basic concepts related to trajectories and used throughout this paper.

The trajectory of a moving object o can be described by a time ordered sequence of locations the object has visited. We call these locations, *trajectory points*.

Trajectory Point A *trajectory point*, l_i^o , is the location of object o at time i , and is defined as,

$$l_i^o = \langle x_i^o, y_i^o \rangle \quad (1)$$

where x_i^o is the longitude of the location which varies from 0° to $\pm 180^\circ$, while y_i^o is the latitude which varies from 0° to $\pm 90^\circ$.

Raw Trajectory A *raw trajectory*, or simply *trajectory*, is a time-ordered sequence of trajectory points of some moving object o ,

$$\tau^o = \langle l_0^o, l_1^o, \dots, l_n^o \rangle \quad (2)$$

Segment or Subtrajectory is a set of consecutive trajectory points belonging to a raw trajectory $\tau^o = \langle l_0^o, l_1^o, \dots, l_n^o \rangle$,

$$s^o = \langle l_j^o, \dots, l_k^o \rangle, \quad j \geq 0, \quad k \leq n \text{ and } s^o \subset \tau^o \quad (3)$$

The process of generating segments from a trajectory is called *Trajectory Segmentation (TS)*. The most common way of defining TS involves splitting a raw trajectory into a set of non-overlapping segments. More formally:

Trajectory Segmentation Given a raw trajectory $\tau^o = \langle l_0^o, l_1^o, \dots, l_n^o \rangle$, we define a sequence of segments $S = \langle s_0^o, \dots, s_k^o \rangle$, such that

$$\forall_{s_i^o, s_{i+1}^o \in S} s_i^o = \langle l_p^o, \dots, l_{p+t}^o \rangle, \quad s_{i+1}^o = \langle l_{p+t+1}^o, \dots, l_{p+t+u}^o \rangle \quad (4)$$

and

$$s_0^o = \langle l_0^o, \dots, l_i^o \rangle, \quad s_k^o = \langle l_j^o, \dots, l_n^o \rangle \quad (5)$$

Equation 6 shows the input and output of the trajectory segmentation process, where τ is a raw trajectory which contains n trajectory points, and S is the set of all segments generated from τ using *TS*.

$$TS : \tau \longrightarrow S, \quad |\tau| = n + 1, \quad |S| = k + 1 \quad (6)$$

In this notation, $n + 1$ is the number of trajectory points and $k + 1$ is the number of segments resulting from applying *TS* to the trajectory.

We call a trajectory point at the end of each segment as *partitioning position*. This means that the result of applying *TS* to a trajectory, S contains k partitioning positions.

Problem Definition Given a raw trajectory τ^o , we would like to generate a sequence of segments $S = \langle s_0^o, \dots, s_k^o \rangle$ so that each s_i^o satisfies a certain homogeneity criteria for a given application domain. To evaluate the performance of the generated S , we rely on the knowledge of an expert user to provide a set of semantic tuples $sl_i = (sid, label)$ where *sid* identifies a segment s_i of a trajectory, generated by the expert user, and *label* is a semantic label attached by the expert to this segment, such as for instance, a transportation mode or status of fishing or non-fishing.

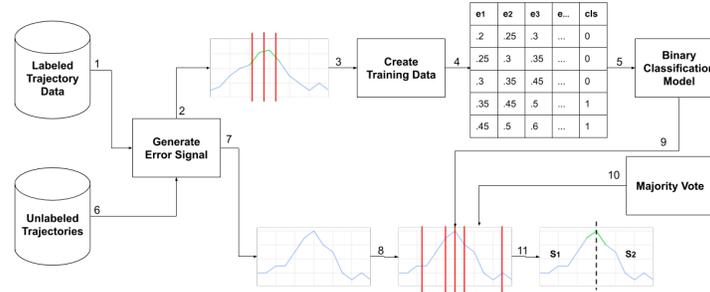


Fig. 1: A high level view of wise sliding window segmentation algorithm.

3 Related works

In this section, we give an overview of several methods for trajectory segmentation. Warped K-Means (WK-Means), which is a general-purpose segmentation algorithm based on K-Means [11], is introduced in [10]. It modifies the K-Means algorithm by minimizing a quadratic error (cost function) while imposing a sequential constraint in the segmentation step. Since WK-Means imposes a hard sequential constraint, segments can be updated while new samples arrive without affecting too much the previous clustering configuration [10]. This algorithm receives the number of segments to be found on the data (k). Having such input parameter is the main limitation of using it in domains where the number of segments is not pre-defined or is dynamic.

The Stay Point Detection (SPD) [17] is a simple algorithm that follows the idea that between each two-movements, there is a stop. SPD applies a distance threshold (θ_d) and a time threshold θ_t so that a moving object which spends more than θ_t time in the neighborhood of θ_d belongs to a stay point. Hence, each stay point identifies a segment, and the trajectory points between two stay points are generated in another segment.

An extension of DB-SCAN [3], CB-SMoT detects stops and moves segments in a trajectory[12]. The original definitions of a ϵ -neighborhood and minimum points in DB-SCAN are altered so that CB-SMoT utilizes spatial and temporal aspects of trajectories. CB-SMoT works based on the trajectory speed, and the stop points are consecutive trajectory points where the moving object has a lower speed.

TRACCLUS [9] detects dense regions with the same line segment characteristics. This clustering algorithm has two steps: (i) partitioning of the trajectory to line segments; and (ii) clustering these lines. A cost function based on the Minimum Description Length (MDL) principle is applied in the first step to split a trajectory into its line segments. It considers three trajectory segment's attributes: (i) parallel distance, (ii) perpendicular distance, and (iii) angular distance. Clustering line segments using DB-SCAN is run in the next step [9].

GRASP-UTS is an unsupervised trajectory segmentation algorithm that benefits from the Minimum Description Length (MDL) principle to build the most homogeneous segments. First, GRASP-UTS generates random landmarks. Then, it builds homogeneous segments by swapping the trajectory points across temporally-ordered segments and adjusting the landmarks based on its cost function’s value [15]. GRASP-UTS can apply additional features on top of the raw trajectories to perform trajectory segmentation.

The OWS (Octal Window Segmentation) algorithm is based on computing the error signal generated by measuring the deviation of a middle point of an octal window [4]. The intuition behind OWS is that when a moving object changes its behavior, this shift may be detected using only its geolocation over time [4]. OWS uses interpolation methods to find the estimated position of the moving object, i.e., where it is supposed to be if its behavior does not change. Then, OWS compares the real position of the moving object with the estimated one, creating an error signal. With such a procedure, it is possible to determine where the moving object changed its behavior and to use this information to create segments.

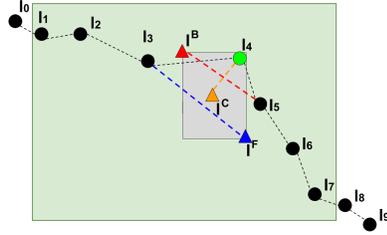
In this work, we extend the idea of OWS by using a configurable sliding window for interpolating points and a supervised strategy for deciding where partitioning positions should be placed. Unlike all previous segmentation algorithms, WS-II is supervised. This means that WS-II is able to learn the variations in the error signal generated by interpolation techniques which characterize partitioning positions over consecutive segments, avoiding in this way the decision of choosing an error threshold value (i.e., an epsilon value in the OWS) that relies on the characteristics of the domain where trajectories were collected.

4 The proposed method

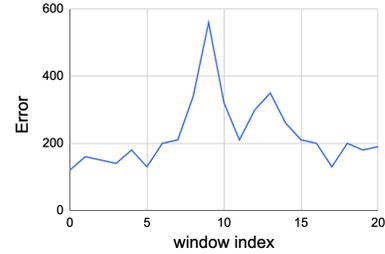
Figure 1 shows an overview of the Wise Sliding Window Segmentation (WS-II) method, which has four core procedures: Generate Error Signal, Create Training Data, Binary Classification Model, and Majority Vote. First, the WS-II creates the error signal from the labeled dataset, which is detailed in Section 4.1. The second step is to generate the training data using the error signal, by sliding a window over its values and adding the presence or absence of a partitioning position. This part is detailed in Section 4.2. The third step is to train a binary classifier to recognize the partitioning positions over the sequence of error signals. This part is detailed in Section 4.3. Finally, unlabeled trajectories can then be segmented based on the model learned in the previous step and using the majority vote, as detailed in Section 4.4.

4.1 Generating the error signal

The first step of our proposal is similar to the OWS algorithm [4], which creates a sliding window over a trajectory to compute a signal error between trajectory



(a) Fig. 2a Example of error calculation where seven trajectory points (e.g., l_1 to l_7) are selected as the *current sliding window* (e.g., green box).



(b) Fig. 2b Example of error signal generated in meters for trajectory data.

points. For each sliding window, the error is generated by calculating the deviation of the interpolated midpoint of the window from the actual midpoint. This process is repeated by sliding the window by one point forward, so receiving a new trajectory point, it adds the newer point to the window set and removes the oldest point from the set. An example of this process is shown in Figure 2a.

In Figure 2a, the green rectangle is a sliding window of size 7, the l^B (red triangle), and l^F (blue triangle) are the interpolated positions. l^F is generated using extrapolation on the first three points (l_1, l_2, l_3) and l^B is generated using the last three points inside the window (l_5, l_6, l_7). The green dot (l_4) is assumed to be the missing point in the sliding window, while the l^C (orange triangle) is generated as a middle point between l^B and l^F . The distance between the midpoint (l^C) and the missing point (l_4) is called the *error value* of this window. In the example of Figure 2a, the haversine distance from the estimated position l^C to the real position l_i is visible. This may indicate that the moving object's behavior has changed at position l_4 .

An example of the error signal from a trajectory is shown in Figure 2b. A raw trajectory with 26 points that forms 20 sliding window of size 7 (generating error for point index 4 to 23) is displayed in this example. The first three and last three error values are dropped. Window index is the index of the middle trajectory point in each window. Figure 2b illustrates a situation in which there are several trajectory points (e.g., around trajectory points 8 and 12) along the raw trajectory where the estimated positions were far from the real trajectory positions. These boundaries are considered as potential partitioning positions for creating trajectory segments.

4.2 Creating Training Data

The second core procedure of WS-II is to create a training dataset using the sequential error values extracted in the previous step. First, we create an array of size q of error signals that will belong to the first training sample, and we use the ground truth information (i.e., if in this particular region there was a change

	e_1	e_2	e_3	e_4	e_5	e_6	e_7	label
...
w_1	120	160	150	140	180	130	200	0
w_2	160	150	140	180	130	200	210	0
w_3	150	140	180	130	200	210	340	0
w_4	140	180	130	200	210	340	560	1
w_5	180	130	200	210	340	560	320	1
w_6	130	200	210	340	560	320	210	1
w_{10}	560	320	210	120	320	273	200	1
w_{11}	320	210	120	320	273	200	130	0

(a) Example of a training set generated by WS-II.

	e_1	e_2	e_3	e_4	e_5	e_6	e_7	b_{cls}	m_{cls}
w_{m-1}
w_m	100	150	230	160	170	320	400	0	...
w_{m+1}	150	230	160	170	320	400	390	1	...
w_{m+2}	230	160	170	320	400	390	320	0	...
w_{m+3}	160	170	320	400	390	320	380	0	0
w_{m+4}	170	320	400	390	320	380	330	0	1
w_{m+5}	320	400	390	320	380	330	490	1	0
w_{m+6}	400	390	320	380	330	490	450	1	0
w_{m+7}	390	320	380	330	490	450	230	1	...
w_{m+8}	320	380	330	490	450	230	160	0	...
w_{m+9}	380	330	490	450	230	160	190	0	...
w_{m+10}

(b) Example of the majority vote mechanism.

Fig. 3: Examples of data tables generated by our algorithm.

in the behavior) to annotate the label of this sample. If this window includes a partitioning position, it is labeled as 1 and 0 otherwise. By receiving every new trajectory point, we remove one point from the start of our window and add the new point to the end of the window. Then we create our next sample by applying the same step of labeling 1 when a partitioning position is present in the sliding window, and 0 if it is not. This procedure is repeated until all the error signals are evaluated.

To understand how the labeling process works, we show an example in Figure 3a. In this example, the training data are created for the sliding window built with seven (e_1 to e_7) trajectory points over eleven slides (i.e., w_1 to w_{11}). As can be seen in Figure 3a, from w_1 to w_3 , there was no big change in the error signal (ranging from 120 to 340 meters). In w_4 , the value of 560 characterizes a high jump in the estimated error and actually reflects a real change in the behavior of the moving object, resulting in a positive example (i.e., there is a partitioning position) in the training data. Examples from w_4 to w_{10} are labeled as positive due to the presence of partitioning position in the sliding window. From w_{11} , the samples are again labeled as negative examples due to the absence of a partitioning position in the data.

4.3 Binary Classification Model

A binary classifier is used by WS-II to categorize each error signal sample into either a partitioning position or not. The labeled trajectory data created in the previous step is used to generate training samples for this binary classifier so that it can classify signal samples into a class where a sliding window has a partitioning position (e.g., value 1) or a class when it does not have a partitioning position (e.g., value 0).

It was observed that the error signal has its minimum fluctuations far from a partitioning position, and it has its maximum fluctuations while transitioning from one segment to a new one. Therefore, detecting the area that includes

partitioning positions is an indicator that the behavior has changed. We apply the binary classifier to identify these areas over a trajectory that has the highest likelihood of containing partitioning positions. In this work, we used a Random Forest classifier [1] to benefit from its bagging power while processing long window sizes faster by limiting the number of features. However, we emphasize that any classification model can be used in this step. After forecasting these transitioning areas, we use a majority vote mechanism to decide precisely where to place a partitioning position, explained in the next section.

4.4 Majority Vote

At this step, we use the same sliding window of size q to decide if a partitioning position occurred. Since we are using a window slide point by point, each trajectory point can be part of q sliding windows, and we classify each window using the binary classifier. This means that we have q outputs that the binary classifier generates for a trajectory point belongs to the q windows. Using a majority vote mechanism for these q outputs leads us to the final decision: the trajectory point is a partitioning position if more than 50% of the sampled signals are labeled as a partitioning position.

Leveraging this feature and applying the voting technique, we can have a more robust evaluation to support if a point is a partitioning position or not. The decision to identify a trajectory point as a partitioning position is supported by q results, each of which contributes $1/q$ to the final decision. This means a misclassification of the binary classifier weights $1/q$. Although increasing q can make the algorithm more robust to noise, it will make it fail to identify segments with a length smaller than q . Furthermore, the algorithm is more robust against noisy points, which may happen in trajectory data due to device collection errors.

An example of the advantages of the majority vote mechanism are exemplified in Figure 3b, where a window with $q = 7$ was used. In Figure 3b, the column b_{cls} was forecast by the binary classifier for w_m to w_{m+9} . It is possible to see in Figure 3b that w_{m+3} is decided by evaluating the b_{cls} column values from w_m to w_{m+6} (0,1,0,0,0,1,1). The decision regarding a majority vote for w_{m+3} is equal to 0 since $|\#0| = 4$ and $|\#1| = 3$. For deciding the final value of w_{m+4} the lines from w_{m+1} to w_{m+7} are used. The evaluation of the set (1,0,0,0,1,1,1) through a majority vote ($|\#0| = 3$ and $|\#1| = 4$) results in the decision of 1 (i.e., a partitioning position occurred). As previously stated, such strategy makes WS-II robust against spatial jumps due to GPS error in the data collection process.

5 Experimental Evaluation

In this section, we evaluate the proposed method and compare it to state-of-the-art approaches. In Section 5.1, we describe the datasets. In Section 5.2, we detail the experimental setup and we report the results in Section 5.3.

5.1 Datasets

We evaluate our method on three datasets. The first is a fishing dataset containing 5190 trajectory points and 153 segments, where fishing activity labels (e.g., fishing or not-fishing) were provided by specialists and used to create trajectory segments. The second is the Atlantic hurricane dataset, which contains 1990 trajectory points and 182 segments. The Saffir-Simpson scale was used to determine the type of hurricane, and the transitions from one hurricane-level to another was used for creating trajectory segments. Finally, a subset of the Geolife dataset containing 12,955 trajectory points and 181 segments was used as a third dataset. For this dataset, we use the transportation mode as the ground truth for creating the segments. The reason that we did not use the full Geolife data set was that some of the segmentation algorithms, such as GRASP-UTS were not able to provide segments in a reasonable time. We create a sample Automatic Identification System (AIS) data to debug our algorithm and test our code and made it available to public ⁵.

5.2 Experimental Setup

In this work we measure the trajectory segmentation performance using Harmonic mean of Purity and Coverage, introduced in [4]. The use of purity and coverage for trajectory segmentation performance measurement originally is introduced in [15]. We do not use clustering measures such as completeness and homogeneity since the segmentation task is different from clustering. In trajectory segmentation, the order of the segments is essential, and adjacent segments can come from the same cluster. For example, an object moving to a shopping store and going back home characterizes two segments, that would be in the same "walk" cluster.

In each experiment, we divided the dataset into ten folds, one of which is applied to tuning/training the algorithm and the rest to testing its performance. Each fold contains different trajectories of different moving objects; therefore, we individually segment each trajectory and report the average results.

Since we divide data into ten folds, we calculate ten values for the Harmonic means. A boxplot is used to show the visual difference between these ten values for each algorithm, Figure 4. Although the boxplot can show the difference between the performance of algorithms, we perform a Mann Whitney U test (having only ten numbers, we could not prove the data follows normal distribution, so we did not use T-test) to show that the difference between the median of each set is not generated randomly.

The state of the art methods that we compare to our approach requires some parameterization. The input parameter values estimation for GRASP-UTS was using a grid search with all combinations of values reported in [15]. For the SPD algorithm, we used the suggested parameters on the original paper for the subset of Geolife dataset, and for the rest of datasets we used a grid search to find the

⁵ <https://github.com/metemaad/WS-II>

best parameters. For CB-SMoT, we applied a grid search to tune parameters using the parameter tuning fold. For OWS, we have tested the four kernels (e.g., random walk, kinematic, linear, and cubic) and used the same strategy reported in [4] to find the best value of *epsilon*. We decided only to report the random walk kernel findings since it obtained good results for all datasets. For a fair comparison between OWS and WS-II, we only report the WS-II results with the random walk kernel. Details regarding the input parameter values ranges for all algorithms can be found in the following link ⁶

5.3 Results and discussion

Figure 4.a displays the results of executing different segmentation algorithms on the Fishing dataset. A Mann Whitney U test indicated that WS-II produces statistically significant higher median ($mean = 94.32, std=0.9$) harmonic mean for trajectory segmentation comparing to OWS with random walk kernel ($p_{value} = 9.133e-05, mean=89.04, std=1.03$). Therefore, the proposed method achieved better performance in comparison to other trajectory segmentation methods.

A fishing activity is characterized by several ship turns. We believe that WS-II had a better result when compared with the other algorithms because of its capability to analyze not only a single trajectory point, but a larger region (i.e., a larger sliding window size). By analyzing a larger region, WS-II will only place a partitioning position when a partitioning position actually occurred (e.g., learned from the training data). Since a single turn is not enough to characterize a fishing activity, WS-II’s strategy of analyzing a larger window is more robust in learning such behavior.

In this experiment, we compare five trajectory segmentation algorithms: CB-SMoT, SPD, GRASP-UTS, OWS with Random Walk kernel, and our proposed trajectory segmentation algorithm (WS-II) on Atlantic hurricane dataset. Figure 4.b shows that WS-II performed better than all other algorithms. A Mann Whitney U test indicated that WS-II produces statistically significant higher median ($mean=94.68, std=2.23$) harmonic mean for trajectory segmentation comparing to OWS with random walk kernel ($p_{value}=9.1e-05, mean=85.67, std=0.59$).

In this experiment, we applied all the segmentation algorithms on a subset of Geolife containing ten different users. Each user’s trajectory creates a fold and we use one fold to tune up our algorithm each time. Figure 4.c depicts our experiment results. Moreover, a Mann Whitney U test supports the claim that WS-II produces statistically significant higher median ($mean = 92.8, std=2.11$) harmonic mean for trajectory segmentation comparing to OWS with random walk kernel ($p_{value} = 0.00065, mean=88.94, std=5.06$).

In the Geolife dataset, there are two major types of movement: (1) fast movements of buses, trains, and cars; and (2) slow movements of walk and bike, which have a random nature. The selection of a random walk seems to be a reasonable decision in this dataset because as long as a moving object moves slowly, the random walk kernel seems to reproduce the random nature of the movement. On

⁶ <https://github.com/metemaad/WS-II>

the other hand, for the moving object that travels fast, the behavior of random walk is similar to a linear interpolation kernel in terms of direction because the direction variation decreases.

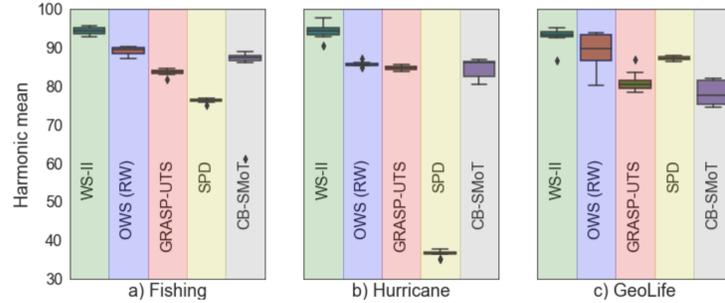


Fig. 4: Our proposed method outperforms CB-SMoT, SPD, GRASP-UTS, and OWS with Random Walk kernel on Fishing, Hurricane and Geolife dataset

6 Conclusions

In this paper we presented a supervised method for trajectory segmentation named Wise Sliding Window Segmentation (WS-II), that uses a trained model for deciding where partitioning positions should be placed. With the majority voting strategy the method becomes more robust to noise points and avoiding unnecessary partitions. The experimental results show that WS-II achieves better performance in terms of a harmonic mean of purity and coverage when compared with state-of-art trajectory segmentation algorithms in three datasets of different domains. One limitation of WS-II, which is a limitation for all learning methods, is that several domains do not have a labeled dataset where the patterns of movement behavior change can be learned. Although there are tools in the literature that encourage and assist the user in the process of labeling trajectory data [14], most trajectory datasets still do not provide any type of ground truth for validating supervised methods. As future work, we would like test how this algorithm performs with different sample sizes for training, i.e., how large the labeled data needs to be to find good results.

Acknowledgments This work was financed by the Brazilian Agencies CNPq, CAPES(Project Big Data Analytics [CAPES/PRINT process number 88887.310782/2018-00]), FAPESC (Project Match co-financing of H2020 Projects - Grant 2018TR 1266); the European Union’s Horizon 2020 research and innovation programme under Grant Agreement 777695 (MASTER1); and the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

1. Breiman, L.: Random forests. *Machine Learning* **45**(1), 5–32 (Oct 2001). <https://doi.org/10.1023/A:1010933404324>
2. Carlini, E., de Lira, V.M., Soares, A., Etemad, M., Machado, B.B., Matwin, S.: Uncovering vessel movement patterns from ais data with graph evolution analysis **2020** (2020)
3. Ester, M., Kriegel, H.P., Sander, J., Xu, X., et al.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *Kdd*. vol. 96, pp. 226–231 (1996)
4. Etemad, M., Hoseyni, A., Rose, J., Matwin, S.: A trajectory segmentation algorithm based on interpolation-based change detection strategies. In: *EDBT/ICDT Workshops* (2019)
5. Etemad, M., Júnior, A.S., Matwin, S.: Predicting transportation modes of gps trajectories using feature engineering and noise removal. In: *Canadian Conference on Artificial Intelligence*. pp. 259–264. Springer (2018)
6. Feng, S., Cong, G., An, B., Chee, Y.M.: Poi2vec: Geographical latent representation for predicting future visitors. In: *Thirty-First AAAI Conference on Artificial Intelligence* (2017)
7. Júnior, A.S., Renso, C., Matwin, S.: Analytic: An active learning system for trajectory classification. *IEEE computer graphics and applications* **37**(5), 28–39 (2017)
8. Junior, A.S., Times, V.C., Renso, C., Matwin, S., Cabral, L.A.: A semi-supervised approach for the semantic segmentation of trajectories. In: *19th IEEE International Conference on Mobile Data Management (MDM)*. pp. 145–154. IEEE (2018)
9. Lee, J.G., Han, J., Whang, K.Y.: Trajectory clustering: a partition-and-group framework. In: *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. pp. 593–604. ACM (2007)
10. Leiva, L.A., Vidal, E.: Warped k-means: An algorithm to cluster sequentially-distributed data. *Information Sciences* **237**, 196–210 (2013)
11. MacQueen, J., et al.: Some methods for classification and analysis of multivariate observations. In: *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*. vol. 1, pp. 281–297. Oakland, CA, USA (1967)
12. Palma, A.T., Bogorny, V., Kuijpers, B., Alvares, L.O.: A clustering-based approach for discovering interesting places in trajectories. In: *Proceedings of the 2008 ACM symposium on Applied computing*. pp. 863–868. ACM (2008)
13. Soares, A., Dividino, R., Abreu, F., Brousseau, M., Isenor, A.W., Webb, S., Matwin, S.: Crisis: Integrating ais and ocean data streams using semantic web standards for event detection. *International Conference on Military Communications and Information Systems* (2019)
14. Soares, A., Rose, J., Etemad, M., Renso, C., Matwin, S.: Vista: A visual analytics platform for semantic annotation of trajectories. In: *EDBT*. pp. 570–573 (2019)
15. Soares Júnior, A., Moreno, B.N., Times, V.C., Matwin, S., Cabral, L.d.A.F.: Grasp-uts: an algorithm for unsupervised trajectory segmentation. *International Journal of Geographical Information Science* **29**(1), 46–68 (2015)
16. Varlamis, I., Tserpes, K., Etemad, M., Júnior, A.S., Matwin, S.: A network abstraction of multi-vessel trajectory data for detecting anomalies. In: *EDBT/ICDT Workshops* (2019)
17. Zheng, Y., Zhang, L., Ma, Z., Xie, X., Ma, W.Y.: Recommending friends and locations based on individual location history. *ACM Transactions on the Web (TWEB)* **5**(1), 5 (2011)