Trajectory Bayesian Indexing : The Airport Ground Traffic Case

Cynthia Delauney[1], Nicolas Baskiotis[1] and Vincent Guigue[1]

***Abstract*— In this paper, we propose a new approach of indexing trajectories to efficiently distinguish abnormal behaviors from normal ones. After a discretization step, trajectories are considered as sets of triplets (location, velocity, direction). Those triplets are seen as words and a multinomial modeling is learned to estimate the probability of each word. The originality of our work consists in computing the likelihood of all measures and aggregating them by trajectories and spatial cells. The achieved representation is light and offers new opportunities to request normal or abnormal behaviors. The interest of our approach is demonstrated on a plane trajectory dataset provided by Paris-Charles de Gaulle airport. Several experiments are carried out to promote the proposed likelihood descriptors; in particular, experiments show how to extract easily relevant specific trajectories. A t-SNE diagram is also presented to achieve an overall discriminative representation of the whole dataset.**

Keywords: Trajectory indexing, clustering, Bayesian characterization

I. Introduction

This article aims at proposing a new approach to tackle efficiently trajectory datasets. In our modern connected world, more and more spatiotemporal data are collected: people get equipped with smartphones, smartwatchs, and soon, smartcars. Therefore, exploiting those datasets efficiently and bringing added value to the users has become a leading issue. Namely, we aim at handling and browsing existing traces, then at understanding them. Multiple challenges are already identified, ranging from sequence categorization [1] or traffic modeling [2] to anomaly detection [3] and safety improvement [4].

This paper concerns a specific task: monitoring airport planes when they drive on the tarmac, specifically when they can be seen as vehicles. Airport authorities have two main concerns: safety aspects and performance optimization. Fortunately, incidents are rare, but each of those situations have to be understood in depth to prevent new occurrences. Thus, a main interest lies in browsing all situations close to a given incident, which requires a relevant indexing of the trajectories. Another classical use-case arises when a safety patrol notices an anomaly in the field; in this case, it is mandatory to check what happened in the past hours on a particular point of the map. This corresponds to an IR (Information Retrieval) problem which requires efficient indexing. A last use-case concerns performance issues and an efficient approach to understand reasons of a traffic jam. In a nutshell, we propose in this work a tool to mine efficiently a big dataset of trajectories which are complex and structured objects.

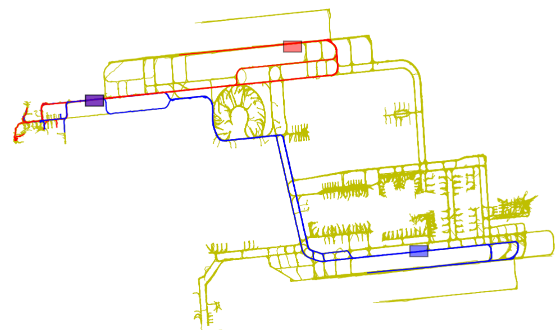

Fig. 1. Two sets of trajectories associated to standard spatial queries; All trajectories cross the northwestern tile, then red ones are continuing towards the northern runway while blue ones proceed to the southern runway.

Trajectory indexing is a well known task where contributions are numerous [5], [6], [7]. However our concern is two folds: firstly, computing similarities between trajectories or between trajectories and queries; then, being able to distinguish normal and abnormal situations efficiently. In that sense, the philosophy of our work is close to [8] who proposed to discover spatial, speed and directional pattern in the series. The first step of the proposed approach is to discretize the space and to estimate direction and speed of planes for each radar echo of the dataset. For efficiency reasons, we discard chaining information and rely only on the computed direction at each point of the trajectory. The extracted triplets (location, velocity, direction) are considered as words and exploit the bag of words formalism. An example of word based query is provided in Fig 1. The originality of our approach consists in addressing the question of normality: we aim at discovering what distinguishes a likely movement from another less common move. We would like to consider the likelihood as a central point of the indexing process: probabilistic models are often used inside categorization algorithms [9], however, the idea here is to use them as trajectory descriptors. We learn a multinomial modeling to estimate the probability of each word. Afterwards, the likelihood of all measurements is computed and aggregated by trajectories and spatial cell. Using such a Bayesian indexing strategy enables us to face efficiently a new kind of semantic request, for

[1] Sorbonne Universités, UPMC Univ. Paris 6, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris. `firstname.name@lip6.fr`

instance: *finding all trajectories crossing this area with a weak likelihood*.

After a brief discussion on related works (section II), section III describes our modeling in depth. Section IV shows the relevance of our approach by analyzing the results associated to classical and more original queries. The categorization ability of our trajectory description is also illustrated using t-SNE [10] to propose an overall view of the dataset.

II. Related Work

In order to handle trajectories efficiently, we investigate two bibliography axis: light trajectory representations on the one hand, to be able to compute requests on real world dataset, and similarity measures between trajectories on the other hand, which is critical for both trajectory mining and categorization applications. Those topics are closely related since similarity largely depends on the representation choice. Generally speaking, comparing two structured objects of different sizes, for instance two randomly picked trajectories, is a challenging task.

First of all, lets consider the indexing task which is crucial to perform relevant request in a big trajectory dataset. Theoretical foundations relied on an efficient space discretization as R-Tree [5]. Then, several adaptations have been proposed like TB-Tree [11] or 3D-R-tree [12]. [11] distinguished two types of spatiotemporal queries : coordinate-based queries which were local and can be enriched with a timing or contextual precisions and trajectory-based queries which involved the whole topology. Both are required in our work.

It rapidly became clear that the noise in the location measurement was a big issue. One solution was to cluster trajectories in order to introduce a meaning in the groups. Early spatial data mining algorithm like DBSCAN [13] or STING [14] focused on static localized entities. DBSCAN, as many other proposal from the 90', was not really dedicated to spatial data mining, it was simply applied on 2D datasets that could be geographical coordinates. STING was more interesting from our applicative point of view as it proposed a global solution including a geographical description based hierarchical discrete structure, a fast clustering algorithm and some request examples. Several metrics were really dedicated to measure sequence similarities like Dynamic Time Warping (DTW) [15] but they could hardly scale up. Edit Distance on Real sequences (EDR) [16] was applied on large datasets, but it required a huge computing grid. More recent approaches dealt with sequences and proposed specific focus on time aspects in TF-OPTICS algorithm [17] or on speed and directional patterns in addition to locations [8].

Another way of handling and understanding trajectories was to consider them as a set of meaningful subparts. Thus, the challenge became two folds: first, segmenting trajectories, then matching the parts. From this perspective, Longest Common SubSequence algorithm (LCSS) [18] focused on an hybrid criterion assuming that the matching score was directly linked to the segmentation. Map-matching algorithms [19] could be seen as a trajectory similarity measurement tool in the particular case where an expert provided a reliable map. Then, computations were focused on the mapping task and required a Viterbi-like dynamic programming algorithm. The matching step became trivial as all segments belonged to a discrete finite set (that is often of limited size). In the general case, more attention was given to the space dividing process [20]. TRACLUS algorithm [21] insisted on the interest of learning segmentation and matching at the same time to bring to light common patterns shared by many trajectories. [22] shared the same philosophy but introduced a distance between full trajectories as a global learning criterion.

In order to take into account chaining in trajectories, several strategies arose in the literature. Indexing task was tackle by referencing successive position in a discrete geographical grid [6]. Regarding trajectories extracted from video flow, Markovian models were often used to estimate the transition probabilities [23]. Modeling transitions was also a way to switch to predictive applications [24].

III. Bayesian modeling

Lets denote our raw trajectory dataset $\mathcal{D} = \{T_1, T_2, \ldots, T_N\}$, each trajectory T_k being a set of time-stamped 2 dimensional positions ($t \in \mathbb{R}, \ell \in \mathbb{R}^2$) sampled at a regular frequency.

$$T_k = \{\mathbf{c}, (t_1, \ell_1, \ldots, t_{|T_k|}, \ell_{|T_k|})\} \tag{1}$$

Each trajectory is associated to contextual information $\mathbf{c}$ giving the type of the aircraft, the name of the company, the nature of the move (take off, landing...), the configuration of the airport. We will exploit those pieces of information as a supervision to analyze our results in the experimental section. In our approach, we choose to discard time information and to consider only the computed direction and speed to model the motion as it is done in certain task of [8]. However, we model the joint distribution between location, velocity and direction instead of considering the marginal laws. Our hypothesis is the following: at a point of the trajectory, if we know the triplet (location, velocity, direction), the next position becomes obvious and we can assume that the transition is already described.

A. Discretization & Bag of Words

After preliminary experiments and expert feedbacks, direction is discretized in 8 categories (using classical compass intervals) and velocity in 6 intervals $[0, 3[$, $[3, 12[$, $[12, 20[$, $[20, 30[$, $[30, 60[$, $[60, +[$ in ms^{-1}. A 30×30 spatial grid is used and areas with no activities are discarded: $S = 339$ spatial tiles are remaining which leads to $Z = 6 \times 8 \times 339 = 16272$ discrete cells. A first discrete trajectory representation is obtained by considering each one of those cells as a word (cf Fig 2): $T_k = \{\mathbf{c}^\dagger, \mathbf{w}\}$ where $\mathbf{w} \in \mathbb{N}^Z$ is a fix-size integer vector counting the number of occurrences of each word, i.e., the number of measurements of the trajectory assigned to each cell of the discrete space. Consistency from one trajectory to another is guaranteed by the constant sampling frequency over the whole dataset[1]; thus, each measure point corresponds the same information weight. We slightly modified the contextual information vector $\mathbf{c}^\dagger$ by adding the departure time of the trajectory and its duration. As it is classically done in text mining, we switch to a frequency representation to improve the comparability of different trajectories. Let $w_i^f = \frac{w_i}{\sum_j w_j} \in \mathbb{R}_+$ denote the frequency of word i in the trajectory. Thus, each frequency vector $\mathbf{w}^f$ sums to one. We sometimes add a (k) exponent to the word when there is an ambiguity about the trajectory the word belongs to. As said earlier, a word corresponds to a triplet location, velocity, direction (ℓ, v, d) and the notation $\{j|\ell \in w_j\}$ is used to address the set of word indexes corresponding to the spatial tile ℓ.

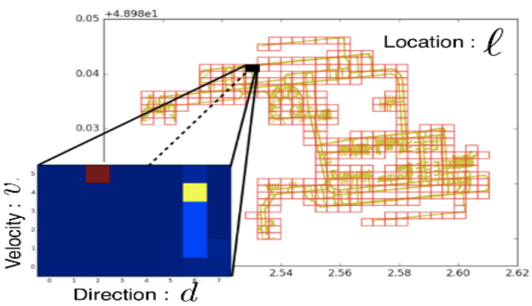

Fig. 2. Discrete joint distribution over location, velocity and direction (ℓ, v, d). High values (red) correspond to common observations while low values (blue) denotes unlikely situations. The model includes $Z = 16272$ cells.

B. Likelihood descriptors

In order to enrich the trajectory T_k, we propose to consider likelihood descriptors in the following manner. Mapping all trajectories to the Z dimensional space leads to estimate the joint distribution of location, velocity and direction, as illustrated on Fig 2. Hence, a multinomial model $\mathcal{M}$ is learned and the parameters are gathered in a vector:

$$\Theta = \begin{bmatrix} \vdots \\ \theta_i = p(w_i|\ell) \\ \vdots \end{bmatrix}, p(w_i|\ell) = \frac{\sum_k w_i^{(k)}}{\sum_k \sum_{\{j|\ell \in w_j\}} w_j^{(k)}} \quad (2)$$

So as to facilitate interpretation and prevent numerical troubles with low real values, we actually compute the conditional distribution with respect to the location ℓ.

[1] A 3Hz sampling frequency is considered in the following.

In practice, $\mathcal{M}$ enables us to compute the likelihood of each trajectory point. Assuming a (naive) independence between measures, a trajectory likelihood can be computed. However it is well known that such a computation will greatly depends on the trajectory length and thus prevent any relevant comparison between sequences. As a consequence, we choose to compute a normalized local likelihood $\mathcal{L}$ for each spatial cell ℓ crossed by the trajectory k:

$$\bar{\mathcal{L}}_\ell^k = \frac{\sum_{\{i|\ell \in w_i\}} w_i^{(k)} \theta_i}{\sum_{\{i|\ell \in w_i\}} w_i^{(k)}} \times \frac{1}{\max_{\{i|\ell \in w_i\}} \theta_i} \quad (3)$$

Likelihoods of positions are averaged inside each area ℓ and divided by the maximum parameter of the region. This latter operation corresponds to a rough normalization regarding the entropy of the spatial cell: some cells are concentrated namely they contain few high probabilities and a lot of zeros whereas others, with higher entropy, contain many low probabilities. Our concern is to be able to compare the likelihoods of two trajectories crossing two different cells. The full representation of the trajectory is thus:

$$T_k = \{\mathbf{c}^\dagger, \mathbf{w}^f \in \mathbb{R}_+^Z, \bar{\mathcal{L}} \in \mathbb{R}_+^S\} \quad (4)$$

C. Queries

As in [11], two types of queries are investigated: content-based queries and sample-based queries. For the formers, our full representation is exploited. A particular context can be selected (take off from a given runway for instance); a spatial query consists then in selecting the words corresponding to an area (and optionally to a specified speed and/or direction). Finally, the query can specify if a standard scenario or unusual situations at a given location is sought.

The second use-case corresponds to sample-based queries. In this situation, a k-nearest neighbors search around the requested trajectory is perfomed using an Euclidian distance to rank the neighbors. The coefficients $\mathbf{w}^f$ are always exploited which enables us to return a group of trajectories that present close shapes. But on top of that, one can chose to activate $\bar{\mathcal{L}}$ descriptors to select trajectories presenting the same normality or abnormality at different points.

IV. Experiments

This section presents a proof of concept around trajectory mining with our new light indexing technique. All experiments have been conducted on 15 days dataset provided by Paris-Charles de Gaulle airport, including 5 716 plane trajectories produced by the Secondary Surveillance RADAR (SSR) which also collects plane transponders' informations.

First, basic experiments are considered to validate our representation based on the joint space *location, velocity,* *direction*. Then the interest of the likelihood descriptors is illustrated in another series of experiments. In order to promote our approach we also carry out a preliminary experiment to extract abnormal long trajectories, considered as being late.

A. Preliminary experiment on long trajectories

We investigate a simple model to extract trajectories that are longer than usual. We adopt a two-stage approach: first, trajectory start points are clustered according to their positions; then, we compute the distribution of taxiing duration and the trajectories corresponding to the highest decile are extracted as illustrated on Fig 3. By checking where long trajectories appear in our query answers will allow to provide qualitative analysis of our work. On top of that, we will focus on the likelihood of such trajectories to try to bring to light congestion precursors.

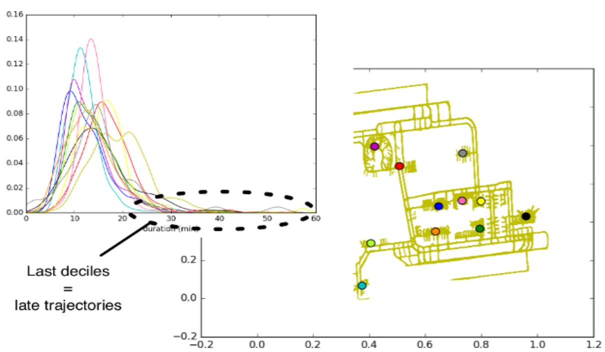

Fig. 3. Start point categorization and taxiing time distribution for each cluster. We consider the trajectories corresponding to the highest decile as being late for each of the 12 clusters.

B. Validation of joint space representation

Our first experiments are illustrated on Fig 1, they correspond to 2 standard spatial queries based on words associated to 2 couples of geographical tiles. We check that we obtain relevant results. Disabling one criterion is equivalent to summing different columns of the representations. For instance Fig 4 illustrates the differences between 2 query answers on the same region, respectively with and without the direction criterion.

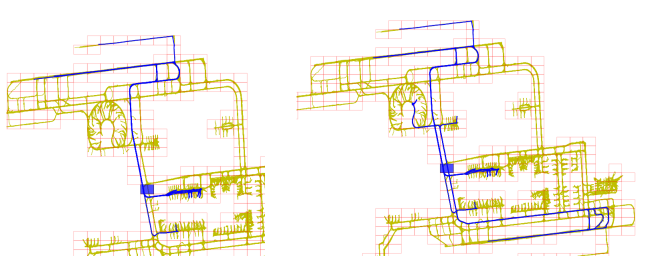

Fig. 4. Result of 2 queries on the same region with (left) or without (right) direction information

Then we move to one of the sample-queries which aim at extracting the nearest neighbors of a targeted situation. Fig 5 illustrates a 3 nearest neighbors search. At first sight, it seems that all trajectories are identical; however zooming on the departure point shows that all trajectories are close but separate. It also exhibits the noise in position measurement, especially at low speed conditions. Spatial discretization is a way to tackle this problem efficiently.

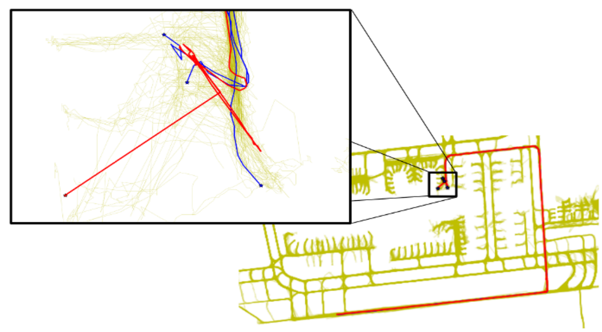

Fig. 5. The red sample is used as a query. 3 nearest neighbors are plotted in blue. The zoom show the noise in the original trajectories, especially in the low speed context, for instance, when planes leave their parking.

Another experiment based on t-SNE [10] is proposed to achieve a global vision of the trajectory dataset. In order to reduce time computation, the spatial cells are aggregated by groups of 9; hence, a 10×10 grid is used for this experiment. The dimension of the dataset is reduced to 2; this leads to a picture where each point corresponds to a trajectory. Trajectories form several separate groups but they are hard to interpret. As a consequence, we conduct a second series of experiments: we focus on take off and use a color code to distinguish the 4 airport runways (Fig 6). We note that groups are nearly perfectly homogeneous, only few outliers remain and it will be interesting to analyze them with experts.

Finally, we exploit the results from section IV-A to map long trajectories on the figure (black circled points): most of those particular trajectories are gathered (especially in the yellow, magenta and green cases). In depth analyses of the magenta cluster, presented in Fig 7, show that the two lobes of the magenta main group correspond to different situations: the left lobe (blue trajectories) matches with northern parking and associated path while the right lobe (red trajectories) matches with southern parkings. Even better, 2 magenta outliers correspond, in fact, to long trajectories. Taking a closer look at this situation, we understand that outliers wait (a long time) at a singular point of the runway entrance (Fig 8): in our representation $\mathbf{w}^f$, those regions corresponds to the highest values. In conclusion our discrete representation seems well adapted to perform an efficient trajectory categorization.

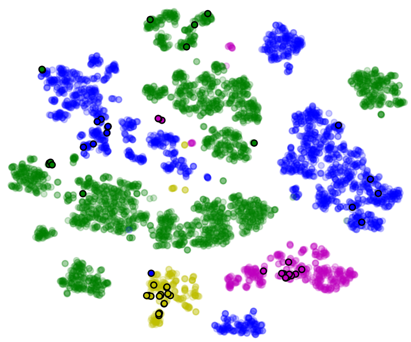

Fig. 6. Visualization with t-sne algorithm of take-off trajectories. Colors correspond to the 4 airport runways. Circled points depict long trajectories extracted in section IV-A

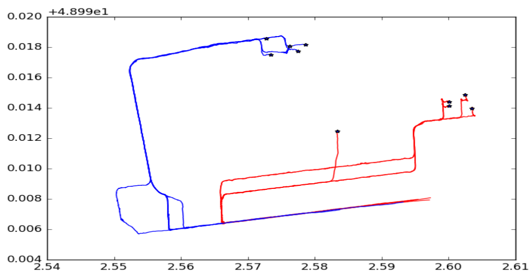

Fig. 7. In depth analyse of the magenta cluster (Fig 6): trajectories from the left (blue) and right (red) part of the main group.

C. Likelihood descriptors

The previous series of experiments (section IV-B) highlight an important requirement: recognizing what is normal or not. As explain in section III-B, a multinomial distribution is learned and specific descriptors are built (one by trajectory and by area): what appends regularly is considered as normal and unusual situations as abnormal. The asset of this rich representation (eq. 4) is to provide a new query opportunity, like *what trajectories are abnormal in area ℓ?* For instance, Fig 9 focuses on abnormal behavior at the northern take off runway waiting point: not surprisingly, all resulting trajectories correspond to long ones (according to experiment IV-A).

Two trajectories can also be compared from the spatial likelihood point of view as shown on Fig 10. According to $\mathbf{w}^f$ descriptors, those trajectories are very close, but adding likelihood information make them distinct; indeed, the first situation presents a strong likelihood drop when the plane drive by terminal 1 (the big round parking structure). In this way, our original description can distinguish situations that would have been merged by a classical spatial query.

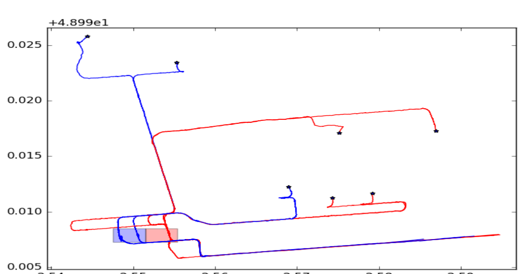

Fig. 8. Late trajectories of the magenta cluster (Fig 6): main group (blue) *vs* outliers (red). The difference comes from the highlighted tiles. Waiting point at the runway entrance is not the same.

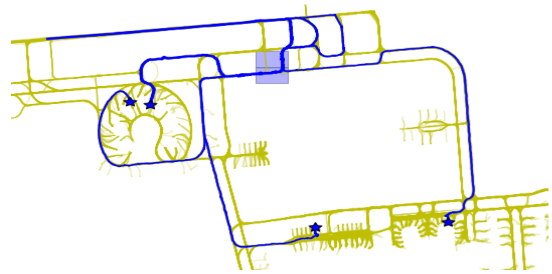

Fig. 9. Results of a query regarding trajectories with low likelihood at the waiting point of the northern take off runway.

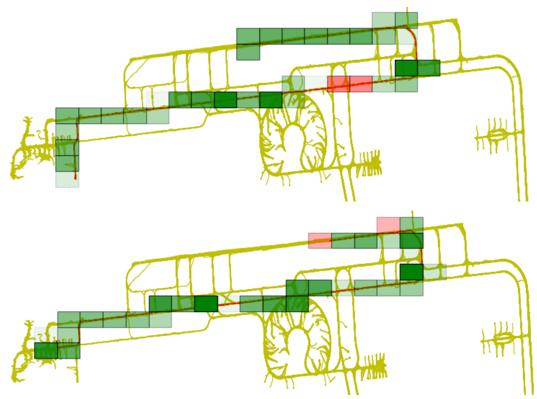

Fig. 10. Spatial likelihood of the two trajectories that are very close according to $\mathbf{w}^f$. The speed profile of those trajectories is analyzed on Figure 11(c). Likelihood above 0.55 are colored in red.

We propose to push the analysis further and we study the velocity likelihood profiles of close spatial trajectories. In Fig 11, 2 areas are considered and specific queries are built to bring to light low likelihood trajectories respectively associated to low velocity and excessive speed. The evolution of the velocity with respect to time is plotted, bolding the measures corresponding to the requested region: we note that velocities behave as expected, being respectively lower or higher than mainstream signals when the likelihood drops. Thus, the proposed indexing strategy allows to build rich semantic queries: the concept of *excessive speed* is automatically learned without any prior, it depends on the usual behavior in a particular region.

V. Conclusion

We propose in the article a new way of indexing trajectories to bring to light efficiently what is normal or not. This indexing is based on the discretization of the physical space and the estimation of direction and speed of planes in each discrete spatial cell for each radar echo

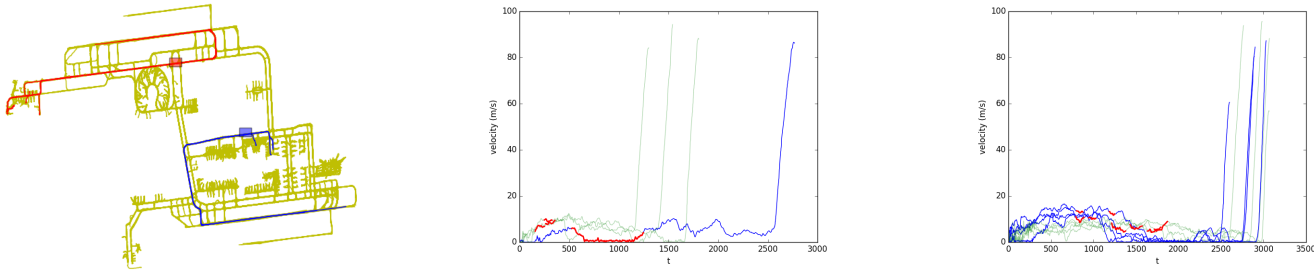

(a) Two trajectory sets corresponding to 4 queries: (1) red region, (2) red region AND abnormally low velocity (3) blue region, (4) blue region AND abnormally high velocity.

(b) Speed wrt to time for the red trajectories of (a). Normal taxiing = green, abnormal = blue. Red bold section of the curves correspond to the highlighted area: **low velocity** is associated to low likelihood and long trajectories.

(c) Speed wrt to time for the blue trajectories of (a). Normal taxiing = green, abnormal = blue. Red bold section of the curves correspond to the highlighted area: **excessive speed** is associated to low likelihood.

Fig. 11. Velocity profiles of normal *vs* abnormal trajectories.

of the dataset; those measures are considered as words corresponding to triplets (location, velocity, direction); finally, a multinomial probabilistic model is learned by counting words appearing in trajectories. The originality of our work consists in computing the likelihood of all measures and aggregating them by spatial cell for every trajectories. Thus, a light representation is obtained offering new opportunities to request normal or abnormal behaviors.

We demonstrate the interest of our approach on various use cases, including a t-SNE diagram that proposes an overall discriminative representation of the whole dataset.

Acknowledgment This work was partially founded by the FUI project AWACS supported by the French Public Investment Bank and the "Conseil Régional d'Ile-de-France". The authors gratefully acknowledge the support of these institutions, and all partners of the project: Safety Line, Aéroports de Paris, CEFA Aviation and the IFSTTAR-GRETTIA Laboratory."

References

[1] F. Jiang, Y. Wu, and A. K. Katsaggelos, "Abnormal event detection based on trajectory clustering by 2-depth greedy search," in *IEEE ICASSP*, pp. 2129–2132, 2008.

[2] M. A. Taylor, J. E. Woolley, and R. Zito, "Integration of the global positioning system and geographical information systems for traffic congestion studies," *Transportation Research : Emerging Technologies*, vol. 8, no. 1, pp. 257–285, 2000.

[3] Y. Bu, L. Chen, A. W.-C. Fu, and D. Liu, "Efficient anomaly monitoring over moving object trajectory streams," in *ACM KDD*, pp. 159–168, 2009.

[4] K. Takeda, C. Miyajima, T. Suzuki, K. Kurumida, Y. Kuroyanagi, H. Ishikawa, P. Angkititrakul, R. Terashima, T. Wakita, M. Oikawa, *et al.*, "Improving driving behavior by allowing drivers to browse their own recorded driving data," in *IEEE ITSC*, pp. 44–49, 2011.

[5] A. Guttman, *R-trees: a dynamic index structure for spatial searching*, vol. 14. ACM, 1984.

[6] V. P. Chakka, A. C. Everspaugh, and J. M. Patel, "Indexing large trajectory data sets with seti," *Ann Arbor*, vol. 1001, no. 48109-2122, p. 12, 2003.

[7] Y. Zheng and X. Zhou, *Computing with spatial trajectories*. Springer Science & Business Media, 2011.

[8] N. Pelekis, I. Kopanakis, G. Marketos, I. Ntoutsi, G. Andrienko, and Y. Theodoridis, "Similarity search in trajectory databases," in *IEEE Temporal Representation and Reasoning*, pp. 129–140, 2007.

[9] T. Hofmann, "Probabilistic latent semantic indexing," in *ACM SIGIR*, pp. 50–57, 1999.

[10] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, no. 85, pp. 2579–2605, 2008.

[11] D. Pfoser, C. S. Jensen, Y. Theodoridis, *et al.*, "Novel approaches to the indexing of moving object trajectories," in *VLDB*, pp. 395–406, 2000.

[12] Y. Theoderidis, M. Vazirgiannis, and T. Sellis, "Spatio-temporal indexing for large multimedia applications," in *IEEE IC on Multimedia Computing and Systems*, pp. 441–448, 1996.

[13] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise.," in *KDD*, vol. 96, pp. 226–231, 1996.

[14] W. Wang, J. Yang, R. Muntz, *et al.*, "Sting: A statistical information grid approach to spatial data mining," in *VLDB*, vol. 97, pp. 186–195, 1997.

[15] D. Bernad, "Finding patterns in time series: a dynamic programming approach," *AAAI Advances in knowledge discovery and data mining*, 1996.

[16] L. Chen and R. Ng, "On the marriage of lp-norms and edit distance," in *VLDB*, pp. 792–803, 2004.

[17] M. Nanni and D. Pedreschi, "Time-focused clustering of trajectories of moving objects," *Journal of Intelligent Information Systems*, vol. 27, no. 3, pp. 267–289, 2006.

[18] M. Vlachos, G. Kollios, and D. Gunopulos, "Discovering similar multidimensional trajectories," in *IEEE IC on Data Engineering*, pp. 673–684, 2002.

[19] M. A. Quddus, W. Y. Ochieng, and R. B. Noland, "Current map-matching algorithms for transport applications: State-of-the art and future research directions," *Transportation Research : Emerging Technologies*, vol. 15, no. 5, pp. 312–328, 2007.

[20] M. Hadjieleftheriou, G. Kollios, V. J. Tsotras, and D. Gunopulos, "Efficient indexing of spatiotemporal objects," in *Advances in Database Technology*, pp. 251–268, Springer, 2002.

[21] J.-G. Lee, J. Han, and K.-Y. Whang, "Trajectory clustering: a partition-and-group framework," in *ACM IC on Management of Data*, pp. 593–604, 2007.

[22] A. Anagnostopoulos, M. Vlachos, M. Hadjieleftheriou, E. Keogh, and P. S. Yu, "Global distance-based segmentation of trajectories," in *ACM KDD*, pp. 34–43, 2006.

[23] N. Saunier and T. Sayed, "Clustering vehicle trajectories with hidden markov models application to automated traffic safety analysis.," in *IJCNN*, vol. 2006, pp. 4132–4138, 2006.

[24] W. Groves, E. Nunes, and M. L. Gini, "Predicting globally and locally: A comparison of methods for vehicle trajectory prediction.," in *UDM@ IJCAI*, p. 5, 2013.