City Research Online

City, University of London Institutional Repository

Citation: Andrienko, N., Andrienko, G., Garcia, J. M. C. & Scarlatti, D. (2019). Analysis of Flight Variability: a Systematic Approach. IEEE Transactions on Visualization and Computer Graphics, 25(1), pp. 54-64. doi: 10.1109/TVCG.2018.2864811

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/20334/>

Link to published version: <https://doi.org/10.1109/TVCG.2018.2864811>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Analysis of Flight Variability: a Systematic Approach

Natalia Andrienko, Gennady Andrienko, Jose Manuel Cordero Garcia, and David Scarlatti

Fig. 1. A: Planned flight trajectories are represented in an artificial space with polar coordinates: movement direction (angle) vs. distance from the cruise phase start (radius). B: A density map summarizes the whole trajectories. C: The density map summarizes the segments that were substituted by shorter paths in the real flights. The inset on the bottom right shows a filtering window around a density hot spot. D: The trajectories crossing the hot spot in the artificial space are shown on a geographic map with 5% opacity.

Abstract—In movement data analysis, there exists a problem of comparing multiple trajectories of moving objects to common or distinct reference trajectories. We introduce a general conceptual framework for comparative analysis of trajectories and an analytical procedure, which consists of (1) finding corresponding points in pairs of trajectories, (2) computation of pairwise difference measures, and (3) interactive visual analysis of the distributions of the differences with respect to space, time, set of moving objects, trajectory structures, and spatio-temporal context. We propose a combination of visualisation, interaction, and data transformation techniques supporting the analysis and demonstrate the use of our approach for solving a challenging problem from the aviation domain.

1 INTRODUCTION

Comparison of trajectories is a common task in movement analysis. Thus, there may be a need to compare actual trajectories of moving entities to expected ones (such as typical, predicted, or planned), or trajectories of movement under different conditions. Comparison tasks may involve multiple and even very many trajectories, and it may be necessary to compare each trajectory to a certain reference trajectory. There may be a common reference trajectory for a set of trajectories (e.g., a central trajectory of a cluster of similar trajectories [8, 14]), or each trajectory may have its individual reference. For example, flights in aviation are generally conducted according to previously created plans. For each flight, there is a planned trajectory and an actual trajectory, which may deviate from the plan. The comparison task involves detecting, measuring, and analysing the deviations of actual flights from the respective plans.

Trajectories can differ in various aspects: route geometry, times of trip start and end, times of reaching corresponding intermediate positions, and characteristics of the movement along the route, such as speed and acceleration. For 3D movement, relevant characteristics include the altitudes or depths of the positions along the route.

While a need may arise in detailed examination of specific differences between two trajectories (e.g., in a case of an incident or an abnormality), in applications dealing with large amounts of movement

data comparison tasks mainly aim at revealing and understanding overall patterns of trajectory variation with respect to space, time, variety of moving objects, and different aspects of the movement context, such as weather and occurring events.

The research contribution of our paper consists of two major components. First, we present a *conceptual framework* for comparative analysis of trajectories. Its essential component is a method to *quantify differences* between trajectories at a high level of detail. Second, we propose a *generic analytical procedure* for comparative analyses involving a large number of trajectories. The procedure consists of

1. pairwise *point matching* between trajectories,
2. creation of *difference data* by computing various difference measures for the matched pairs of points,
3. when appropriate, division of trajectories into *structural parts* based on essential distinctions, e.g., in the character, purpose, or context of the movement, and
4. *analysis of the distribution* of the difference data with respect to space, time, set of moving objects, trajectory structure, movement characteristics, and movement context.

The procedure is supported by the following techniques:

- *creation of artificial spaces* according to analysis foci and visualisation of the distributions of difference data in these spaces,
- dynamic, filter-sensitive spatial and spatio-temporal aggregation of difference data in geographic and artificial spaces,
- techniques supporting comparison of aggregates that summarise different subsets of data.

We demonstrate the effectiveness of the analytical procedure and technique combination by applying them to challenging problems requiring comparative analysis of real-world massive movement data from the aviation domain and, as a supplement, ground transportation [16].

The remainder of the paper has the following structure. Based on an overview of related work (section 2), we present the conceptual foundations for comparative analyses of trajectories (section 3), introduce the analysis procedure and supporting techniques (section 4), apply them for aviation data analysis (section 5), and discuss the overall approach (section 6).

- Natalia Andrienko and Gennady Andrienko are with Fraunhofer IAIS and City University London. E-mail: gennady.andrienko@iais.fraunhofer.de.
- Jose Manuel Cordero Garcia is with CRIDA (Reference Center for Research, Development and Innovation in ATM). E-mail: jmcordero@e-crida.enaire.es
- David Scarlatti is with Boeing Research & Development Europe. E-mail: david.scarlatti@boeing.com.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

2 RELEVANT WORK

There exist many visualization [24] and visual analytics [7] methods targeting at transportation data and problems. However, these data and problems are so complex, varied, and multifaceted that the existing methods are not sufficient to meet the needs of the transportation domain. Here we discuss three major components of the existing approaches: visualization, interactive filtering, and data transformations. We then review the existing approaches to point-wise matching of trajectories. Next, we describe specifics of the aviation domain and position our research against the state of the art.

2.1 Visualization of trajectories

Trajectories are traditionally visualized as lines on a map and in a space-time cube [34]. Trajectories of movement in 3D space require representation of the altitudes or depths, e.g., by varying the appearance of line segments or by using 3D space displays [23,53]. Visual displays showing multiple trajectories often suffer from severe overplotting and visual clutter. To handle this problem, movement data can be represented in an aggregated way using discrete [4,13] or continuous [38,54] aggregation. Other approaches include edge bundling [25,33] and schematic representation [31], which improve display readability but introduce distortions and undesired artefacts. Aggregated or schematic representations of trajectories hide a lot of details. More detailed exploration is possible for a small number of selected trajectories. Interactive selection that temporarily hides a part of data is often called *filtering*.

2.2 Interactive filtering of movement data

Movement data involve different aspects [6], including spatial positions and path geometries, positions in time, and various attributes characterising the movement, such as the speed and direction, the moving objects, such as their weight or load, and the movement context, such as weather parameters or density of surrounding traffic. Many such attributes can be derived from the spatial positions of the trajectory points alone or in combination with data describing the spatial, temporal, or spatio-temporal context of these positions [11]. Subsets of movement data can be selected based on any of these aspects using interactive tools designed for spatial, temporal, and attribute-based filtering [6]. A prominent example of interactive filtering by direct manipulation is FromDaDy [32], where the user selects trajectories by drawing shapes in displays presenting various 2D projections of flight trajectories. Filtering operations can be applied sequentially to results of previous operations. Temporal filtering can be based on linear [1] or cyclic [27,28] models of time or on selection of time intervals satisfying interactively specified query conditions [15].

2.3 Transformation of time and space

An important tool for movement data analysis is data transformations [6], which may affect the spatial and temporal references. Transformations of time references include replacement of absolute times by relative positions within temporal cycles (annual, weekly, or daily) or with respect to the start and/or end times of trajectories [5]. Temporal references can also be replaced by chronological ordering, as in a trajectory wall display [49] and in a matrix of tram rides along a given route with rows corresponding to trips and columns to consecutive stops [42,55]. Transformations of spatial references include replacing absolute spatial positions by relative positions within a group of jointly moving objects [14] and by positions in an abstract semantic space consisting of location categories rather than specific locations [17,37].

Furthermore, artificial spaces can not only be obtained by transforming the spatial positions of trajectory points to another coordinate system but also constructed based on any attributes associated with trajectory points. This idea has not yet been applied to spatial trajectories, but it was used for representing changes of dynamic attributes of objects that do not necessarily move or are not spatial by nature. For example, the development of Swiss cantons was represented by trajectories in a space formed by two attributes [39]. A similar representation was used for stock market data [47]. An artificial space for representing changes and long-term evolution can also be constructed from multiple attributes using dimensionality reduction techniques [18,51].

Artificial spaces can be treated similarly to usual (physical) spaces, which means that data transformed to artificial spaces can be visualised on maps using cartographic visualisation techniques [36]. One can also apply methods for spatial aggregation, e.g., build density maps [37].

2.4 Matching of trajectories

Our approach to comparing trajectories bases on establishing pairwise matches between the points of the trajectories to be compared in a way compliant with the following *requirements*: (1) points must be matched based solely on their spatial proximity without involving temporal and speed constraints; (2) it must not be assumed that each point of a trajectory necessarily has a match in the other trajectory; (3) each point may receive at most one match. Point matching is involved in algorithms assessing the degree of similarity of two trajectories for the purposes of search and clustering [20,40,43,50], the best known being Fréchet distance [3] and Dynamic Time Warping [44]. Most of these algorithms do not comply with the requirements 2 and 3, and those based on sequence alignment [52] comply with (2) but not with (3). Sankararaman et al. [46] provide an illustration of results of several algorithms demonstrating violations of our requirements. The map matching algorithms [41] aiming to match trajectories to predefined lines, such as streets, do not fulfil requirements 1 and 2.

2.5 Visual analytics for air traffic domain

Visual analytics approaches have been proposed for various specific problems in air traffic analysis, such as detection of holding patterns and problematic movement events [45], assessment of conflict probabilities regarding the traffic density [2], extraction and analysis of wind parameters [29], and detection of violations of rules imposed for decreasing noise in populated areas [22]. Sophisticated domain-specific analyses can be done by applying clustering to interactively selected relevant parts of trajectories [8]. Still, there are many analysis problems that have not yet been addressed in the visual analytics research. Due to the complexity and various specifics of the aviation domain, it is important to do research in collaboration with domain experts [8,30].

2.6 Positioning of our work

While much research has been done on analysis of movement data, the problem of comparative analysis of trajectories has not been addressed yet. We systematically considered this problem and developed a general conceptual framework, where we defined the possible aspects in which trajectories can differ and figured out the objectives of analysing these differences for a large set of trajectories rather than for individual trajectories. Such comparative analyses aim at discovering and understanding patterns that exist among the differences with regard to their spatial, temporal, and spatio-temporal distributions, the distribution within the trajectories, and relationships to movement characteristics and spatio-temporal context.

Based on this conceptualisation of the problem, we developed a generic analytical procedure for comparative analysis of trajectories. Both the conceptualisation and the procedure are novel research contributions. They are complemented with a proposed suite of tools that can support the realisation of the approach. This includes a novel algorithm for matching trajectories with the aim to measure their differences. The other tools (data transformations, visualisations, and interaction techniques) were created by adapting and developing ideas from the previous research, namely, artificial spaces and dynamic aggregation.

3 PROBLEM CONCEPTUALISATION

3.1 Levels of analysis tasks

According to Bertin [19], tasks (questions) in data analysis can be differentiated according to the reading levels: elementary (referring to individual data items), intermediate (referring to groups of data items), and overall (referring to a whole data set). A unifying term *synoptic* was introduced [12] to refer to the intermediate and overall levels, both of which involve abstraction. Synoptic tasks deal with multiple data items, which are considered together, and require abstraction for deriving some general statements concerning all these items.

The distinction between elementary and synoptic levels is relevant, in particular, to comparison tasks [12]. In application to trajectories, elementary comparison tasks consist of detection and examination of similarities and differences between individual trajectories at a high level of detail, i.e., with attending to their points and segments. A smallest comparison task is pairwise comparison, i.e., comparison of two trajectories. An elementary comparison task involving more than two trajectories can be decomposed into several subtasks of pairwise comparison. From two trajectories, one can be treated as a *reference* to which the other trajectory is compared. For example, a planned trajectory can be a reference to an actual trajectory, or an actual trajectory can be a reference to a trajectory generated by a predictive model.

Synoptic comparison tasks involve joint consideration of similarities and differences of multiple trajectories from their references; the latter may be either common or distinct. In analysing large collections of trajectories, synoptic tasks are of primary importance while elementary tasks need to be performed occasionally, in particular, for examination of trajectories with large differences from their references.

3.2 Difference measures

Synoptic comparison tasks can be performed using computationally derived *difference measures*, which include distance in space, difference in time, and differences in movement characteristics, such as speed, direction, and, in case of 3D trajectories, vertical position (altitude or depth). Differences between the followed routes can be represented by the spatial distances between corresponding points of two trajectories and by the differences in the path lengths of corresponding parts of the trajectories. Temporal differences are measured as differences between the absolute and/or relative times of corresponding points. The relative times may be defined with respect to the trip starts and/or ends. Differences in relative times are indicative of differences in movement speed, which can also be computed explicitly.

Derived difference measures can be organised in data records attached to components (i.e., points and segments) of trajectories. For uniformity, it can be assumed that difference data are attached to points, with data referring to segments being attached to the starting points of the segments. Each record consists of values of one or more difference measures. Since each point of a trajectory has its specific position in space (coordinates) and time (time stamp), difference data records become associated with these spatial and temporal positions, i.e., difference data are spatio-temporal by their nature.

Apart from local difference measures referring to points and segments, one can derive general difference measures referring to the whole trajectories. These measures include differences in the trip origins and destinations, start and end times, trip durations, path lengths, path curvatures, average speeds, and other general features of the trips.

3.3 Context data

Since movement is much affected by the *context* [6, 10], it may be necessary to analyse the relationships between the computed differences and the context. Movement context consists of various spatial, temporal, and spatio-temporal objects, events, phenomena, and processes. *Global context*, which is common for all moving entities at a given time, can be distinguished from *local context*, which differs among locations or combinations of locations and times and, hence, is specific to each point in a trajectory. Aspects of local context that are relevant to the analysis goals can be represented by attributes attached to points of trajectories [11]. Thus, for a trajectory point, attributes can represent

- presence or absence, number, or density of spatial objects or events in the spatial, temporal, or spatio-temporal neighbourhood of the point;
- attribute values or statistical summaries of attribute values associated with the spatial positions and/or times in the neighbourhood of the point.

Such local context attributes can be derived by joining trajectory data with data describing the relevant context based on the spatial and/or temporal references specified in the context data. Aspects of global context can be represented by time-dependent attributes characterizing properties of time moments (e.g., day or night, week day or weekend),

presence of events (such as a strike or a public holiday), and/or the overall traffic situation on the territory under study (e.g., low or intense traffic). Unlike local context attributes, global context attributes are not associated with specific trajectories or positions in trajectories.

3.4 Structural parts of trajectories

Trajectories may consist of heterogeneous parts differing in essential features. Thus, travelling people may use different transportation modes, wheeled vehicles use different road categories, movements of ball game players with the ball differ from running without the ball, and flight trajectories include take-off, ascent, cruise, descent, and landing. For valid analysis, it may be appropriate to distinguish structural parts of trajectories, or movement phases, based on essential differences in movement character, purpose, context, or other aspects pertinent to the analysis goals. Analysts may strive to study how differences relate to trajectory structures, or may separately consider differences for distinct structural parts, or may focus only on certain parts of trajectories.

3.5 Foci in synoptic comparison tasks

Synoptic comparison tasks aim at discovering and studying general patterns of differences with respect to the space, time, set of moving entities, global and local contexts, as well as with respect to the internal make-up of the trajectories, including the geometry of the followed path, life time from the start to the end, movement characteristics along the path, and, possibly, essentially different structural phases of the movement. Hence, there are varieties of synoptic comparison tasks focussing on one or several of the following aspects:

- spatial distribution of differences;
- temporal distribution of differences;
- statistical distribution of differences over a set of moving entities;
- internal distribution of differences within the trajectories in relation to their spatial and temporal features and structure;
- relationships of differences to movement characteristics;
- relationships of differences to local and global contexts.

As stated in section 3.1, synoptic tasks include overall and intermediate tasks. Overall comparison tasks are applied to the entire set of trajectories, and intermediate tasks are applied to subsets. Subsets may be selected based on space (e.g., trips within or crossing specific areas or trips with specific origins and/or destinations), time (e.g., trips within a chosen time period, or trips that occurred at certain times in a time cycle), categories or characteristics of moving objects (e.g., types of vehicles or flight operators), and/or properties of the trips, such as path length, duration, curvature, average or maximal speed, etc.

3.6 Comparison tasks in flight data analysis

In aviation, flights are generally conducted according to plans, which are created in advance and agreed between flight operators and flight managers. In reality, actual flights may deviate from the plans for a variety of reasons, such as weather or traffic conditions. Deviations from plans occur frequently and are not considered problematic. Still, better compliance to flight plans is desired, and specialists are interested in analysing deviations for revealing general patterns and interdependencies that may suggest directions to improving flight planning.

As mentioned in section 3.4, flight trajectories, either planned or actual, consist of several structural phases. In analysing flights and their differences from the plans, it is important to distinguish these phases and, possibly, select only what is relevant to specific analysis goals.

4 APPROACH

The proposed analytical procedure is schematically shown in Fig. 2. Comparative analysis of trajectories is based on computing difference measures (section 3.2), which are attached to the trajectories (general measures) and their points (local measures). The measures characterize pairwise differences between each trajectory and some reference trajectory. While obtaining general difference measures is straightforward, deriving local difference measures requires finding for each trajectory point the corresponding point in the reference trajectory, which is not trivial. An algorithm for point matching is proposed in section 4.1.1.

Fig. 2. Proposed workflow for comparative analysis of trajectories.

Fig. 3. Examples of point-wise matching of two trajectories (top and centre) and trajectories of a cluster to its central trajectory (bottom).

Examining differences between two trajectories (elementary comparison) is not too difficult. It can be supported by visualizing the trajectories on a map and/or in a space-time cube [34], whereas the temporal variation of their dynamic attributes can be shown in a time graph [35]. Previously derived difference measures can be explicitly encoded using suitable visual variables [26].

In our research, we are mostly concerned with synoptic comparison tasks (section 3.1), which focus on the distributions of difference data over space, time, set of moving objects, in relation to trajectory structure, and in relation to local and global contexts (section 3.5). We propose a combination of supporting techniques, which includes creation of artificial spaces according to analysis foci (section 4.2), dynamic, filter-sensitive aggregation of difference data (section 4.3), and operations enabling comparison of aggregates. These techniques are generic, i.e., applicable to trajectories of various kinds of moving objects. We tested their effectiveness by applying to data from the aviation domain. Analysis of flight trajectories requires distinguishing between flight phases. To perform our analysis properly, we developed a method for dividing a flight trajectory into phases (section 4.1.2). In

Fig. 4. Schematic representation of the basic idea of the point matching algorithm: p_1 and q_1 are previously matched points, (p_2, q_2) is the current candidate pair, the double lines represent matches.

general, division of trajectories into structural parts is done by defining and applying domain-specific rules based on values of relevant attributes (e.g. for travel mode detection [21]).

4.1 Algorithmic foundations

4.1.1 Matching points of two trajectories

The goal is to match points along the paths from the origin to the destination based on their spatial proximity, as shown in Fig.3, so that the requirements 1-3 (section 2.4) are fulfilled. The basic idea of the Algorithm 1 is schematically represented in Fig.4, where p_i and q_j are two last matched points from trajectories P and Q, respectively, and d_{00} is the distance between them. For the following points p_{i+1} and q_{j+1} , the distances to the points $\{q_j, q_{j+1}, q_{j+2}\}$ and $\{p_i, p_{i+1}, p_{i+2}\}$, respectively, are measured and compared. Depending on which of the five distances is the smallest and whether it is also smaller than d_{00} or the distance d_{22} between p_{i+2} and q_{j+2} , either p_{i+1} is matched with one of $\{q_j, q_{j+1}, q_{j+2}\}$ or q_{j+1} is matched with one of $\{p_i, p_{i+1}, p_{i+2}\}$. When p_{i+1} is to be matched with q_j , or q_{j+1} with p_i , either p_i or q_j loses its previous match and can be matched instead to the point q_{j-1} or p_{i-1} , respectively, if the latter is free (Fig.4, bottom left).

Algorithm 1 Matching points of two trajectories

```

1: procedure MATCHPOINTSOFTRAJECTORIESPANDQ ( $P, Q$ )
2:    $M \leftarrow (1, 1)$  ▷ set of matching pairs of  $P$  and  $Q$ 
3:    $(i, j) \leftarrow (1, 1)$  ▷ running indices over  $P$  and  $Q$ 
4:   while  $i + 1 \leq P.length \vee j + 1 \leq Q.length$  do
5:     for  $k = 0, 1, 2$  do
6:       for  $n = 0, 1, 2$  do
7:         if  $i + k \leq P.length \wedge j + n \leq Q.length$  then
8:            $d_{kn} \leftarrow distance(p_{i+k}, q_{j+n})$ 
9:         else
10:           $d_{kn} \leftarrow \infty$ 
11:       if  $d_{01} < d_{00} \wedge d_{01} \leq \min(d_{10}, d_{11}, d_{12}, d_{21})$  then
12:          $M.removeLast()$ 
13:          $M.append(i, j + 1)$ 
14:       else
15:         if  $d_{10} < d_{00} \wedge d_{10} \leq \min(d_{01}, d_{11}, d_{12}, d_{21})$  then
16:            $M.removeLast()$ 
17:            $M.append(i + 1, j)$ 
18:       else
19:         if  $d_{12} < d_{22} \wedge d_{12} \leq \min(d_{01}, d_{10}, d_{11}, d_{21})$  then
20:            $M.append(i + 1, j + 2)$ 
21:       else
22:         if  $d_{21} < d_{22} \wedge d_{21} \leq \min(d_{01}, d_{10}, d_{11}, d_{12})$  then
23:            $M.append(i + 2, j + 1)$ 
24:       else
25:          $M.append(i + 1, j + 1)$ 
26:        $(i, j) \leftarrow M.last()$ 
27:   return  $M$  ▷ set of matching pairs of  $P$  and  $Q$ 

```

This is a single-pass algorithm in which three distances are calculated for each trajectory point; hence, the computational complexity is $\mathcal{O}(n + m)$, where n and m are the numbers of points in the two trajectories.

Figure 3 shows examples of results of Algorithm 1. The upper two images demonstrate matching of two trajectories coloured in red and blue. The points that got matches are represented by dots of the same colours as the trajectories, and the points without matches are represented by pink and cyan dots. Yellow lines connect the matched points. These lines are only visible where the routes diverge. As can be seen, the algorithm handles quite well unequal numbers of available points and unequal point spacing in two trajectories. The lower image demonstrates matching of a cluster of similar flight trajectories (thin blue lines) to the cluster's central trajectory (thick red line), which was constructed with ignoring the initial and final parts of the flights [8]. Please note that, although Algorithm 1 begins with matching the initial points of two trajectories (line 3), this match can later be cancelled (line 16) if a better match is found. The better match, in turn, can be replaced in the next step by a yet better one. Hence, an initial part of

one of the trajectories can finally have no match, as it happened to the initial parts of the flights in Fig.3. This is a valid result, as well as the absence of matches for the final parts of the flights, since the reference trajectory lacks both the initial and final parts.

4.1.2 Identification of flight phases

In analysing flight trajectories, it is important to distinguish three phases of each flight: ascent, cruise, and descent. The task is not trivial, because both ascent and descent can be (and most often are) step-wise, and the altitude (flight level) can also change during the cruise phase. It is not easy to determine whether a flat segment of a trajectory (i.e., where the altitude is constant) is already the cruise or still a step in the ascent phase, and whether a decrease of the altitude means the beginning of the descent or just a change of the cruise level. Flight phases can be quite easily distinguished visually from a representation of the vertical profile of a flight over time. However, it is much harder to define formal rules that could be automatically applied.

The main idea of our approach is first to select roughly a time interval when the flight is very likely to be in its cruise phase (we explain in the next paragraph how this is done). From this interval, we find the minimal altitude among the flat segments or peak points, which we treat as the minimal cruise altitude. We include peak points into consideration because the ascent in a very short flight may be immediately followed by the descent, i.e., the cruise phase may consist of a single peak point. We then scan the entire trajectory and find the first and the last flat segments where the altitude is not lower than the minimal cruise altitude. These segments are treated as the beginning and end of the cruise phase. The trajectory parts before and after them are categorised as ascent and descent, respectively. In Fig.5, this approach has been applied to a set of flights between two cities.

A tricky part of the approach is the initial selection of a time interval that is likely to contain the cruise phase but no flat segments of the ascent and descent. By interactive visual exploration of numerous flights with diverse durations, we found out that the selection needs to be done differently for long and short flights. For long flights, it usually takes up to 45 minutes to reach the cruise altitude after the take-off whereas the descent may last for up to 75 minutes. For short flights, the approximate durations of the ascent and descent are better to be estimated proportionally to the overall flight duration. We found that about 30% of the overall duration may be spent for the ascent and about 40% for the descent. Hence, the interval $[t_{start} + 45 \text{ minutes}, t_{end} - 75 \text{ minutes}]$ can be used for finding the minimal cruise altitude when the flight is long (3 hours or more), and the interval $[t_{start} + d \cdot 0.3, t_{end} - d \cdot 0.4]$, where d is the total flight duration, can be used when the flight is short (less than 3 hours). The use of these intervals has been approved by the aviation experts. Please note that the initially selected time interval generally differs from the finally identified cruise phase. In particular, the former may neither fully contain the latter nor be contained in it.

4.2 Artificial spaces

Figure 5, bottom, shows an example of a constructed artificial space. The general idea is to choose two numeric attributes associated with trajectory points and treat their value ranges as spatial dimensions, which can be arranged as Cartesian or polar coordinate systems. The latter is suitable when one of the attributes has a cyclic value range. Examples are spatial direction (Fig. 1), time of the day (Fig. 6), and relative time within a week. Points of trajectories are represented by points in the constructed space according to their attribute values, and these new points are connected in the chronological order. The result is a trajectory in the artificial space. The artificial space with the trajectories can be represented in a map display analogously to geographical space. There is no background map, but it is possible to create a map layer with labelled grid lines or axes. A benefit of constructing maps of artificial spaces rather than implementing other types of display to represent dynamic attributes is that all map-based visualisation and interaction techniques can be uniformly applied to any map, either with geographic or artificial space. Thus, in Fig. 5 (bottom), a qualitative attribute of trajectory segments (flight phase)

Fig. 5. Division of flights into phases: ascent (red), cruise (blue), and descent (yellow). Top: A 3D view with the horizontal plane representing the geographic space and vertical dimension representing the flight level. Bottom: An artificial space with the horizontal and vertical dimensions representing the relative time from the start to the end and the flight level, respectively. The cruise phase is represented by dashed lines for better visibility of the other two phases.

is represented on a map by colour coding, and Fig. 1 demonstrates the possibilities to build a density map and to set a “spatial” filtering window in an artificial space in the same way as it is done in a usual map. This uniformity is very convenient for analysis.

To obtain a meaningful representation of trajectories in an artificial space, it is reasonable to base one of its dimensions on an attribute that is monotonous along the trajectory, i.e., the value either increases or decreases. Examples include time (absolute or relative with respect to flight start and/or end) and relative position along the path. It can also make sense to take a non-monotonous attribute that changes gradually rather than abruptly, e.g., movement direction (Fig. 1) or spatial distance to the trip start or end. The attribute for the second dimension may be chosen more arbitrarily. It can characterise the movement, as the speed or flight level, or moving objects, as the weight, or local context, as the traffic intensity around the points (Fig. 6). It can also be one of the difference measures (section 3.2) derived on the basis of point matching results (section 4.1.1).

A set of trajectories can be mapped onto an artificial space two or more times using comparable attributes, i.e., having similar meanings, common units of measurement, and similar value ranges. Thus, the example in Fig. 6 involves two transformations of the same set of trajectories to an artificial space ‘time of the day - local traffic intensity’. One transformation is based on values of an attribute reflecting the expected traffic intensities and the other transformation is based on the actual traffic intensities. The traffic intensity is measured as the number of flights per hour in a grid cell of the size 10x10 km.

All existing maps based on geographic or artificial spaces are interlinked through interactive operations, such as mouse hovering, brushing, and filtering, which are performed in the same ways in any map and affect uniformly all other maps. As an example, Fig. 1 demonstrates how filtering in a map of an artificial space affects a geographic map.

4.3 Dynamic aggregation and filtering

Aggregation is a common approach when it is necessary to deal with large amounts of data. Dynamic aggregation means that data aggregates are automatically re-computed in response to data filtering to include only the data satisfying the current filters. This, in turn, triggers automatic updating of all displays showing aggregated information.

Since our analytical framework involves intensive use of maps based on real or artificial spaces, we are primarily interested in using spatial

Fig. 6. Left: Density distribution of the cruise parts of flight trajectories in polar coordinates consisting of the time of the day (angle) and expected local traffic intensity (radius). Right: Difference of the densities corresponding to the actual and to the expected traffic intensities.

aggregation methods in all such spaces. There are two types of spatial aggregation: continuous and discrete. Continuous aggregation [48] represents data as a surface in which the value of some attribute (particularly, a statistical summary, such as a density) gradually changes from position to position. A surface is visualised by encoding the values by colours. Continuous aggregation is demonstrated in Figs. 1 and 6, where sets of trajectories are represented by density surfaces built in artificial spaces. Examples based on the geographic space can be seen in Figs. 7 and 8.

Discrete spatial aggregation is based on dividing the space into regular or irregular compartments, which are usually larger than in a raster. For the compartments, various data summaries are derived, typically without involvement of spatial smoothing. Sufficiently large compartment sizes allow visual representation of the attached summaries not only by colour coding but also by symbols, diagrams, or glyphs, which can be applied to several summary attributes, e.g., the mean and the standard deviation. Discrete aggregation of movement data can produce not only place-based aggregates summarising visits of places (i.e., spatial compartments) by moving objects but also link-based aggregates summarising movements between the places [6, 13]. The latter are visualised on *flow maps* [36] using flow symbols (e.g., as in Fig. 8) with the widths proportional to aggregate values, such as move counts.

Although it is possible technically to generate link-based aggregates in an artificial space, this may not be a valid idea. Links in real space are associated with meaningful spatial directions, while directions in an artificial space may be meaningless; hence, a flow map in an artificial space can hardly be interpretable and useful for analysis.

Both continuous and discrete aggregates can be dynamic, i.e., sensitive to data filtering. However, this responsiveness is not always desirable. Thus, it may be necessary to compare aggregates obtained from a subset of data to overall aggregates or to aggregates of another data subset. Hence, there should be a mechanism to suspend dynamic re-aggregation for some of the derived datasets (map layers) while the others preserve their dynamic behaviour. This enables three modes of comparison [26]: by juxtaposition, as in Fig. 8, superposition, as in Fig. 13, or explicit encoding of differences or ratios, as in Fig. 6. The latter example, in particular, demonstrates that actual flights tend to deviate from their planned routes to spaces with low traffic intensity. The raster of density differences shows high increase of the computed trajectory density for values of traffic intensity that are close to zero and decrease of the density for higher traffic intensities.

Different types of interactive filtering applicable to spatio-temporal data have been described elsewhere [6, 15]. Among them, the filter type called ‘time mask’ [15] is suitable for supporting analysis with regard to the global context (section 3.3). The filter selects data from time intervals satisfying specified query conditions. The latter can be based on attributes characterising the global context.

5 CASE STUDY

The approach was validated in a case study performed by a team of data analysts and aviation domain experts collaborating remotely with the use of synchronous and asynchronous communication techniques.

Fig. 7. Left: The lines drawn with 2% opacity represent actual trajectories in blue and planned trajectories in red. Right: The density map summarizes the actual flight segments that deviated from the plans by 2.5 km or more.

The experts stated the overall analysis problem and set various specific questions throughout the study. The analysts generated visualizations containing answers to the experts’ questions and provided explanations. The experts interpreted the patterns observed, and the team made inferences and drew conclusions. The analysis was done on a dataset describing 32,736 flights over the territory of Spain in the period from April 7 till April 14, 2016 (8 days). For each flight, its planned and actual trajectories are available. The map in Fig. 7, left, represents the planned and actual trajectories by red and blue lines, respectively. It is easy to note that the routes of the actual flights often deviate from the planned routes. The analysts applied Algorithm 1 to match points of the actual and planned trajectories and computed various difference measures, including the differences in the lengths of the path segments between matched points that spatially coincide with 1 km tolerance.

In the following, all illustrations representing subsets of data in real and artificial spaces (section 4.2) were obtained by means of interactive filtering and dynamic aggregation (section 4.3). In particular, segments of trajectories were filtered based on the aforementioned difference measure ‘difference of path lengths between coinciding points’. Apart from the techniques presented in section 4, the study involved other visual analytics techniques, in particular, time series clustering [9].

5.1 Problem statement

The domain experts wish to investigate the deviations of actual trajectories from the planned routes such that the lengths of the corresponding path segments significantly differed. Fig. 7, right, shows the densities of the actual flight segments that were at least 2.5 km longer or shorter than planned. The highest densities occurred around the major airports, but the domain experts are not interested in these deviations. They know that path changes in airport vicinities are typically caused by changes of the take-off and landing directions, which, in turn, correspond to changes of the wind. What is really interesting to the experts is the path length changes during the cruise phase of the flights. Therefore, the analysts divided both the flight plans and actual flights into phases (section 4.1.2) and selected only the cruise phase by means of filtering. There were 29,343 flights whose cruise phase took place at least partly over the territory of Spain. Path reductions by 2.5 km or more occurred in 10,695 of them (36.4%), and path increases by at least 2.5 km occurred in 3,139 flights (10.7%). Hence, path reductions, i.e., straighter path segments than it was planned, occurred quite frequently. Occurrences of path reductions and extensions in a trajectory are not mutually exclusive; there were 725 flights (2.5%) involving both.

For the domain experts, more interesting are the cases when aircraft fly more directly than it was planned. They explain that, when negotiating their plans with the flight management services, airlines usually ask for straighter paths but in most cases do not get such a permission and have to plan longer paths. However, during the actual flights, pilots are often allowed and even commanded by flight controllers to fly more directly. On the one hand, this increases the flight efficiency; on the other hand, it also entails serious drawbacks. When longer paths are planned, more fuel has to be taken on board, which increases the aircraft weight. Flying with extra fuel is not just useless but expensive and environmentally negative. Besides, unplanned path shortening can decrease the flight duration making the aircraft come too early to its destination airport, where it may have to wait in the air for a permission

Fig. 8. Left: The cruise parts of the flight plans have been aggregated in a density map (top) and in a flow map (bottom). Right: The density and flow maps represent the planned path segments such that the corresponding actual paths were shorter by 2.5 km or more.

to land. Therefore, the main improvement desired by airlines is the flight plan compliance with straighter path segments put in the plans.

To find out if there is potential for this improvement, the aviation experts wish to uncover patterns (regularities or dependencies) in the path reductions, which would show where and/or when and/or under what conditions shorter paths could be allowed at the planning stage.

5.2 Seeking spatial patterns

The density and flow maps in Fig. 8, left, summarize all segments of the flight plans belonging to the cruise phase. In the flow map, the widths of the flow symbols are proportional to the flight counts. The maps on the right represent only the parts of the plans that were shortened in the real flights by at least 2.5 km. The overall spatial patterns differ quite much. It is seen that path reductions occurred mostly on certain routes, moreover, in particular parts of these routes. Furthermore, it appears that there may be directional patterns of path reductions.

To investigate this, the analysts transform the planned trajectories to an artificial space with a polar coordinate system in which the angle represents the movement direction and the radius corresponds to the spatial distance from the beginning of the cruise phase or from the boundary of the studied territory if the cruise phase began outside of it. In Fig. 1, the trajectories are represented by lines drawn with 1% opacity (A) and by a density map (B). Both maps clearly show the prevalence of the north-eastern and south-western movement directions. After selecting only those segments that were substituted by shorter paths in the real flights, the density map changes its appearance as shown in Fig. 1C. There is a hot spot of very high density at the angle 225 and distance 0-15 km. The density range mapped onto the colour shades in Fig. 1C has been interactively reduced to make high densities beyond the hot spot more prominent. As a result, the density map has a hole in the position of the hot spot, but now it is easy to see that the densities of the path-reduced segments are generally higher in the directions north - north-east - east-north-east than in the opposite directions, except for the hot spot and two wings originating from it.

To see what trajectories correspond to the hot spot, the analysts create a filtering window enclosing it (shown in an inset in the lower right corner of Fig. 1C). The filter selects all trajectories that cross the window in the artificial space. A geographic map display shows the selected trajectories in the geographic space (Fig. 1D). They are drawn with 5% opacity; hence, brighter lines correspond to multiple overlapping trajectories. The map shows that many trajectories come from the south-west of France, but there are also bunches of trajectories whose cruise phase started elsewhere, e.g., south from Madrid.

The spatial pattern exploration has revealed the existence of particular traffic lines and directions where more path length reductions

Fig. 9. The links of the generalized traffic network (Fig. 8) have been clustered by the similarity of the time series of path reductions. The 2D histograms below the map (rows: days, columns: hourly intervals with step 30 minutes) represent the average per hour counts of all flights (left) and reduced segments (right) going through the links of the clusters 3-7.

occurred than elsewhere. The next task is to find out when and/or under what conditions this happened.

5.3 Seeking spatio-temporal patterns

To investigate the temporal aspect of the path reductions, the analysts use the discrete aggregation of the planned trajectories into flows in a generalized traffic network, as shown in Fig. 8. For the links of this network, they compute time series of flow magnitudes, i.e., flight counts per time unit. Specifically, they take time intervals of 1 hour length with 30 minutes shift, i.e., the intervals overlap by 30 minutes for temporal smoothing and diminishing boundary effects. The time series are computed for all flight segments and separately for those which were reduced in the actual flights. Partition-based clustering (k-means) is applied to the latter time series. The analysts perform the interactive clustering procedure [9] in which a projection display of the cluster centres is used to choose a suitable value for the parameter k .

The clustering results are demonstrated in Fig. 9. Colours have been assigned to the clusters based on the positions of their centres in the projection space [9]. The links in the flow map are painted in these colours, and the same colours are used in the 2D time histogram displays below the map. Each histogram represents in an aggregated way the time series for the links belonging to one cluster. The rows correspond to the days and the columns to the hourly intervals within a day. The sizes of the painted rectangles in the cells represent statistical aggregates, such as the means, as in Fig. 9, of the respective values from the time series. The histograms on the left correspond to the whole trajectories, and those on the right correspond to the reduced segments. The histograms are shown for five out of the seven clusters produced. The remaining two clusters contain time series with very low values, which have no practical interest for the domain experts.

The histogram displays allow the team to make the following observations. The path reductions on the links of cluster 5 (dark purple)

Fig. 10. The space-time cube shows time series of path reductions computed for a fine regular division of the territory. The hourly counts 3 or higher are represented by proportional circle sizes.

occurred quite regularly; the time series of the reduced paths look similar to the complete time series. It appears that some proportion of the flights using these links can be safely allowed to plan more direct paths. Cluster 6 (light blue) has the highest regularity of reductions among the remaining clusters, which may also entail some potential for optimizing flight planning, although the flight counts here are much lower than in cluster 5. Cluster 4 (orange) could be of high interest from the perspective of planning optimization, as it involves many flights, but its temporal pattern of path reductions has low regularity. The same applies to the other clusters. There were some days and/or hours when more path reductions occurred than in the other days and/or hours. The experts suppose that some air traffic controllers may be more inclined than others to straighten paths in the flights under their control.

The experts want to verify their observations using a less aggregated representation than in Fig. 9. The analysts compute time series of path reductions for cells of a regular grid with square cells of the size 10x10 km and visualize the result in a space-time cube, as in Fig. 10, where the temporal dimension is oriented upwards. In this view, the values 3 or higher are represented by proportional sizes of circle symbols. Concentrations of circles signify spatio-temporal clusters of path reductions. The experts rotate the cube and observe where reductions were practised regularly and where they occurred more occasionally. The space-time cube offers them a more refined view of the spatio-temporal patterns of path reductions. They conclude that the patterns can mostly be considered individual-independent with a high level of confidence, making optimisation in strategic planning feasible.

The 2D time histograms also allow the experts to check if the path reduction patterns could be related to the traffic intensity, as it can happen that some flights are requested to deviate from their planned routes for decreasing the traffic intensity in the parts of the air space where it is expected to be high. The experts could not find such a relationship in their earlier studies with other datasets and other methods, but they wanted to check this result using visual analytics techniques. This can be done by comparing the histograms on the left and on the right of Fig. 9. The histograms on the left show the traffic intensities represented by the flight counts per time unit. The existence of a relationship would manifest through high path reduction counts in the times of high traffic and very low values in the remaining times, which is not the case. For clusters 5 and 6, the reduction counts look proportional to the flight counts; for the other clusters, the patterns of path reductions are less regular and thus dissimilar to those of flight counts.

5.4 Seeking relationships to temporal deviations

The domain experts want to see how the deviations from the planned routes are related to flight delays computed as the differences between the time till the actual flight end and the time till the planned flight end. Positive differences mean delays of the actual flights and negative differences mean the opposite. It could be hypothesized that air traffic controllers may tend to allow straighter paths to retarding flights to help them reduce the delays. For revealing the relationships and checking this hypothesis, the analysts build an artificial space where the X-axis

Fig. 11. A-C: An artificial space with the X-axis representing time and Y-axis representing the delays of the actual flights compared to the plans. A: density of all segments of the actual trajectories; B: density of their shortened segments; C: the average differences of the lengths of the actual and planned segments. D: Analogous to C, but the Y-axis represents the differences in the time till the flight end in the actual and planned trajectories.

Fig. 12. Flow maps obtained by aggregation of the flight plan segments such that the corresponding actual segments take longer (left) and shorter (right) time by at least 5 minutes.

represents time and the Y-axis represents the delays. After converting the actual trajectories to this coordinate system, the analysts create a density map (Fig. 11A). Generally, the densities are symmetric with respect to the line $Y=0$, which means that delays and advances are mostly balanced. The density map in Fig. 11B summarises the length-reduced parts of the actual trajectories. The density distribution pattern also looks quite symmetric with respect to $Y=0$, but there are no regions of very high densities around this line, as in the overall density distribution. This means that the flights involving path reductions are less frequently on time but more often deviate from the planned times. The symmetry of the density distribution signifies that retarding flights do not use shorter paths more often than others.

To verify this observation, the analysts produce a raster with average differences of the path lengths in the actual and planned trajectories. In Fig. 11C, the differences are represented using a diverging colour scale with shades of blue for negative differences (shorter segments in actual flights) and shades of brown for positive differences (longer segments in actual flights). The visible prevalence of negative values above the line $Y=0$ may mean that paths in retarding flights tend to be straightened to a higher degree than in other flights.

In a similar manner, the team investigates the existence of a relationship between the path straightening and the tendency to longer or shorter flight durations than planned. They consider the differences between the times till the respective flight ends in the actual and planned trajectories. Positive differences for actual flight segments mean that

Fig. 13. Comparison of path reductions for different airlines. Purple: aggregate flows from the complete trajectories (cruise phase) of each airline. Yellow: aggregates from the path-reduced parts of the trajectories.

they take longer time than planned, and vice versa. The analysts build an artificial space similar to that in Fig. 11A-C, but its vertical dimension represents the differences in the times till the flight end in the actual and planned trajectories. The distribution of the average differences between the actual and planned path lengths in this space is shown in Fig. 11D. It is seen that positive values prevail above the line $Y=0$ and negative values below this line, i.e., path extensions are more associated with the flights taking longer time and path reductions with the flights taking less time than planned. A natural explanation is that path extensions increase and path reductions decrease flight durations.

The team also looks at the spatial patterns of the increased and decreased segment durations (Fig. 12) and makes an interesting observation: increases and decreases of the durations with respect to the plans are related to the flying directions. The durations tend to increase in the directions toward the west, south-west, and south and to decrease in the opposite directions. In fact, the prevailing directions of the segment duration decrease correspond quite well to the prevailing directions of flight path shortening that were seen in Fig. 1, bottom left; however, the spatial patterns differ (cf. Figs. 12, right and 8, bottom right). This means that the decreases of the durations cannot be fully attributed to the path shortenings but there should be a different reason, such as an impact of the jet stream.

5.5 Comparing path reductions between airlines

The domain experts would also like to investigate whether different airlines are equally treated by air traffic controllers with regard to path straightening, or some of them have higher chances to get straighter paths than the others. The proportions of the number of path-reduced segments to the number of all segments differ significantly among the airlines; thus, for 8 major airlines, the proportions range from 10.6% to 33.5%. However, these differences are not indicative of unequal treatment, because different airlines have different routes, and it was observed earlier that path reductions occur on some links of the traffic network more often than on others. Conclusions concerning equal or unequal treatment can only be made based on comparing relative frequencies of path reductions for different airlines on the same links. Such comparisons can be done using flow maps, as in Fig. 13, where the maps represent aggregated data from the flight plans of four airlines. The dark purple flow symbols represent the aggregates obtained from the whole trajectories (cruise phase), and the yellow flow symbols drawn on top represent the aggregates from the shortened segments. The maximal absolute values of the flow magnitudes, which are represented by the maximal widths of the flow symbols, differ among the maps. This supports comparison of the relative frequencies.

Generally, the flow maps in Fig. 13 show that on the same links

different airlines have approximately equal frequencies of straighter paths relative to the total frequencies of flying through these links. The large differences in the proportions of path reductions are due to differences in the routes served by the airlines. Hence, there are no significant evidences of possible unequal treatment of different airlines by the air traffic controllers.

5.6 Overall conclusion from the study

The study allowed the domain experts to gain valuable insights into the functioning of the air traffic management (ATM) system. Their questions regarding flight plan compliance were answered. They saw possibilities for optimising planning on certain important routes and links of the traffic network, found that route changes do not depend on personal preferences of traffic controllers, and observed that different airlines receive equal treatment. They confirmed their previous finding concerning the independence of the path reductions of the traffic intensity and learned that path reductions are also not used massively as an instrument to decrease flight delays. These findings basically show that path shortening does not happen due to special reasons, but its purpose may be convenience of traffic controlling and better distribution of the staff workload. Hence, shorter paths can be put in flight plans more often than now.

6 DISCUSSION AND CONCLUSION

What concerns the case study, the domain experts were very much satisfied by their collaboration with the analysts, who showed them the power of visual analytics techniques. The experts believe that the capability of visual analytics to reveal spatio-temporal patterns has a great potential for reducing inefficiencies in the ATM system by learning the actual behaviour of the system and being able to understand it better. Results as shown in this paper are thus of great interest not just scientifically but also from the operational perspective. The experts emphasise the importance of the problem addressed in the case study and the need for better flight planning and plan compliance, especially for reducing negative impacts on the environment. The study yielded relevant findings, which, of course, require further testing on other datasets. Nevertheless, the experts acknowledged that it showed

- the capability of visual analytics to detect patterns that can be applied for global benefit of the system,
- the ability to provide a better insight into the ATM system and thus support better decision making in strategic/pre-tactical phase,
- the usability of the results achieved, providing direct benefit in the operational environment.

Although our research on comparative analysis of trajectories was motivated by practical needs existing in the aviation domain, we strove to obtain results that could have a broader area of applicability. Based on our previous experiences in analysing diverse kinds of movement data from different domains, we considered comparative analysis as a general problem and developed a generic conceptual framework, which centres around the concept of pairwise difference measures. Accordingly, we developed a method to derive these measures through pairwise point matching between trajectories. The framework defines synoptic comparison tasks as seeking patterns in distributions of the differences with respect to space, time, trajectory structure, movement characteristics, and spatio-temporal context. Accordingly, we composed a suite of visual analytics techniques that can support exploration of these distributions and discovery of existing patterns. The description of our case study provides an example that facilitates understanding of the general procedure and shows how to use the proposed techniques.

We also tested our approach using other data examples. Thus, a supplement to this paper [16] demonstrates an application to analysing vehicle traffic along a motorway. Except for dividing flight trajectories into phases, all other techniques and the whole analytical procedure are domain-independent and can be used for comparative analysis problems in any domain where differences between trajectories following similar routes need to be quantified and analysed.

This work was supported by Fraunhofer Cluster of Excellence ‘Cognitive Internet Technologies’ and by EU in projects datAcron (grant 687591), VaVeL (grant 688380), and Track&Know (grant 780754).

REFERENCES

- [1] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of time-oriented data*. Springer Science & Business Media, 2011.
- [2] G. H. Albrecht, H. T. Lee, and A. Pang. Visual analysis of air traffic data using aircraft density and conflict probability. In *Infotech@ Aerospace 2012*. 2012.
- [3] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. 5:75–91, 03 1995.
- [4] G. Andrienko and N. Andrienko. Spatio-temporal aggregation for visual analysis of movements. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (IEEE VAST) 2008*, pp. 51–58, 2008. doi: 10.1109/VAST.2008.4677356
- [5] G. Andrienko and N. Andrienko. Poster: Dynamic time transformation for interpreting clusters of trajectories with space-time cube. In *2010 IEEE Symposium on Visual Analytics Science and Technology*, pp. 213–214, Oct 2010. doi: 10.1109/VAST.2010.5653580
- [6] G. Andrienko, N. Andrienko, P. Bak, D. Keim, and S. Wrobel. *Visual Analytics of Movement*. Springer, 2013. doi: 10.1007/978-3-642-37583-5
- [7] G. Andrienko, N. Andrienko, W. Chen, R. Maciejewski, and Y. Zhao. Visual analytics of mobility and transportation: State of the art and further research directions. *IEEE Transactions on Intelligent Transportation Systems*, 18(8):2232–2249, 2017. doi: 10.1109/TITS.2017.2683539
- [8] G. Andrienko, N. Andrienko, G. Fuchs, and J. M. Cordero-Garcia. Clustering trajectories by relevant parts for air traffic analysis. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):34–44, Jan 2018. doi: 10.1109/TVCG.2017.2744322
- [9] G. Andrienko, N. Andrienko, G. Fuchs, and J. Wood. Revealing patterns and trends of mass mobility through spatial and temporal abstraction of origin-destination movement data. *IEEE Transactions on Visualization and Computer Graphics*, 23(9):2120–2136, Sept 2017. doi: 10.1109/TVCG.2016.2616404
- [10] G. Andrienko, N. Andrienko, and M. Heurich. An event-based conceptual model for context-aware movement analysis. *International Journal of Geographical Information Science*, 25:1347–1370, 2011. doi: 10.1080/13658816.2011.556120
- [11] G. Andrienko, N. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel. Scalable analysis of movement data for extracting and exploring significant places. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):1078–1094, July 2013. doi: 10.1109/TVCG.2012.311
- [12] N. Andrienko and G. Andrienko. *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005. doi: 10.1007/3-540-31190-4
- [13] N. Andrienko and G. Andrienko. Spatial generalization and aggregation of massive movement data. *IEEE Trans. Vis. Comput. Graph.*, 17(2):205–219, 2011. doi: 10.1109/TVCG.2010.44
- [14] N. Andrienko, G. Andrienko, L. Barrett, M. Dostie, and P. Henzi. Space transformation for understanding group movement. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2169–2178, Dec 2013. doi: 10.1109/TVCG.2013.193
- [15] N. Andrienko, G. Andrienko, E. Camossi, C. Claramunt, J. M. Cordero-Garcia, G. Fuchs, M. Hadzagic, A.-L. Joussetme, C. Ray, D. Scarlatti, and G. Vouras. Visual exploration of movement and event data with interactive time masks. *Visual Informatics*, 1(1):25 – 39, 2017. doi: 10.1016/j.visinf.2017.01.004
- [16] N. Andrienko, G. Andrienko, J. M. Cordero Garcia, and D. Scarlatti. Analysis of flight variability: a systematic approach. Appendix. <http://www.geoanalytics.net/and/papers/vast18appendix.pdf>. Accessed: 20.06.2018.
- [17] N. Andrienko, G. Andrienko, G. Fuchs, and P. Jankowski. Scalable and privacy-respectful interactive discovery of place semantics from human mobility traces. *Information Visualization*, 15(2):117–153, 2016. doi: 10.1177/1473871615581216
- [18] B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic. Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):559–568, Jan 2016. doi: 10.1109/TVCG.2015.2467851
- [19] J. Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983.
- [20] P. C. Besse, B. Guillolet, J. M. Loubes, and F. Royer. Review and perspective for distance-based clustering of vehicle trajectories. *IEEE Transactions on Intelligent Transportation Systems*, 17(11):3306–3317, Nov 2016. doi: 10.1109/TITS.2016.2547641
- [21] W. Bohte and K. Maat. Deriving and validating trip purposes and travel modes for multi-day gps-based travel surveys: A large-scale application in the Netherlands. *Transportation Research Part C: Emerging Technologies*, 17(3):285 – 297, 2009. doi: 10.1016/j.trc.2008.11.004
- [22] J. Buchmüller, H. Janetzko, G. Andrienko, N. Andrienko, G. Fuchs, and D. A. Keim. Visual analytics for exploring local impact of air traffic. *Computer Graphics Forum*, 34(3):181–190, 2015. doi: 10.1111/cgf.12630
- [23] S. Buschmann, M. Trapp, and J. Döllner. Animated visualization of spatial-temporal trajectory data for air-traffic analysis. *The Visual Computer*, 32(3):371–381, 2016. doi: 10.1007/s00371-015-1185-9
- [24] W. Chen, F. Guo, and F. Y. Wang. A survey of traffic data visualization. *IEEE Transactions on Intelligent Transportation Systems*, 16(6):2970–2984, Dec 2015. doi: 10.1109/TITS.2015.2436897
- [25] O. Ersoy, C. Hurter, F. Paulovich, G. Cantareiro, and A. Telea. Skeleton-based edge bundling for graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2364–2373, Dec 2011. doi: 10.1109/TVCG.2011.233
- [26] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, Oct. 2011. doi: 10.1177/1473871611416549
- [27] M. Harrower, A. L. Griffin, and A. MacEachren. Temporal focusing and temporal brushing: assessing their impact in geographic visualization. In *Proceedings of 19th International Cartographic Conference, Ottawa, Canada*, vol. 1, pp. 729–738, 1999.
- [28] M. Harrower, A. MacEachren, and A. L. Griffin. Developing a geographic visualization tool to support earth science learning. *Cartography and Geographic Information Science*, 27(4):279–293, 2000. doi: 10.1559/152304000783547759
- [29] C. Hurter, R. Alligier, D. Gianazza, S. Puechmorel, G. Andrienko, and N. Andrienko. Wind parameters extraction from aircraft trajectories. *Computers, Environment and Urban Systems*, 47:28 – 43, 2014. doi: 10.1016/j.compenvurbsys.2014.01.005
- [30] C. Hurter, Y. Brenier, J. Ducas, and E. L. Guilcher. Cap: Collaborative advanced planning, trade-off between airspace management and optimized flight performance: Demonstration of en-route reduced airspace congestion through collaborative flight planning. In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, pp. 1–9, Sept 2016. doi: 10.1109/DASC.2016.7777947
- [31] C. Hurter, M. Serrurier, R. Alonso, G. Tabart, and J.-L. Vinot. An automatic generation of schematic maps to display flight routes for air traffic controllers: Structure and color optimization. In *Proceedings of the International Conference on Advanced Visual Interfaces, AVI '10*, pp. 233–240. ACM, New York, NY, USA, 2010. doi: 10.1145/1842993.1843034
- [32] C. Hurter, B. Tissoires, and S. Conversy. Fromdady: Spreading aircraft trajectories across views to support iterative queries. 15(6):1017–1024, 2009. doi: 10.1109/TVCG.2009.145
- [33] T. Klein, M. van der Zwan, and A. Telea. Dynamic multiscale visualization of flight data. In *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, vol. 1, pp. 104–114, Jan 2014.
- [34] M.-J. Kraak. The space-time cube revisited from a geovisualization perspective. In *Proceedings of the 21st International Cartographic Conference (ICC)*, pp. 1988–1996, Aug. 2003.
- [35] M.-J. Kraak and O. Huisman. Beyond exploratory visualization of space time paths. In H. Miller and J. Han, eds., *Geographic data mining and knowledge discovery*, pp. 431–443. Taylor & Francis, United Kingdom, 2 ed., 2009.
- [36] M.-J. Kraak and F. Ormeling. *Cartography: Visualization of Spatial Data*. Guilford Press, 2011.
- [37] M.-P. Kwan. Gender, the home-work link, and space-time patterns of nonemployment activities. *Economic Geography*, 75(4):370–394, 1999. doi: 10.1111/j.1944-8287.1999.tb00126.x
- [38] O. D. Lampe and H. Hauser. Interactive visualization of streaming data with kernel density estimation. In *2011 IEEE Pacific Visualization Symposium*, pp. 171–178, March 2011. doi: 10.1109/PACIFICVIS.2011.5742387
- [39] P. Laube, S. Imfeld, and R. Weibel. Discovering relative motion patterns in groups of moving point objects. *International Journal of Geographical Information Science*, 19(6):639–668, 2005. doi: 10.1080/13658810500105572
- [40] H. Liu and M. Schneider. Similarity measurement of moving object trajectories. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on GeoStreaming, IWGS '12*, pp. 19–22. ACM, New York, NY, USA, 2012. doi: 10.1145/2442968.2442971
- [41] Y. Lou, C. Zhang, Y. Zheng, X. Xie, W. Wang, and Y. Huang. Map-

- matching for low-sampling-rate gps trajectories. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '09, pp. 352–361. ACM, New York, NY, USA, 2009. doi: 10.1145/1653771.1653820
- [42] J. D. Mazimpaka and S. Timpf. A visual and computational analysis approach for exploring significant locations and time periods along a bus route. In *Proceedings of the 9th ACM SIGSPATIAL International Workshop on Computational Transportation Science*, IWCTS '16, pp. 43–48. ACM, New York, NY, USA, 2016. doi: 10.1145/3003965.3003970
- [43] N. Pelekis, G. Andrienko, N. Andrienko, I. Kopanakis, G. Marketos, and Y. Theodoridis. Visually exploring movement data via similarity-based analysis. *Journal of Intelligent Information Systems*, 38(2):343–391, Apr 2012. doi: 10.1007/s10844-011-0159-2
- [44] L. Rabiner and B.-H. Juang. *Fundamentals of Speech Recognition*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
- [45] M. Sakr, G. Andrienko, T. Behr, N. Andrienko, R. H. Güting, and C. Hurter. Exploring spatiotemporal patterns by integrating visual analytics with a moving objects database system. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '11, pp. 505–508. ACM, New York, NY, USA, 2011. doi: 10.1145/2093973.2094060
- [46] S. Sankararaman, P. K. Agarwal, T. Mølhave, J. Pan, and A. P. Boedihardjo. Model-driven matching and segmentation of trajectories. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL'13, pp. 234–243. ACM, New York, NY, USA, 2013. doi: 10.1145/2525314.2525360
- [47] T. Schreck, T. Tekušová, J. Kohlhammer, and D. Fellner. Trajectory-based visual analysis of large financial time series data. *SIGKDD Explor. Newsl.*, 9(2):30–37, Dec. 2007. doi: 10.1145/1345448.1345454
- [48] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, London New York, 1986.
- [49] C. Tominski, H. Schumann, G. Andrienko, and N. Andrienko. Stacking-based visualization of trajectory attribute data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2565–2574, Dec 2012. doi: 10.1109/TVCG.2012.265
- [50] K. Toohey and M. Duckham. Trajectory similarity measures. *SIGSPATIAL Special*, 7(1):43–50, May 2015. doi: 10.1145/2782759.2782767
- [51] S. van den Elzen, D. Holten, J. Blaas, and J. J. van Wijk. Reducing snapshots to points: A visual analytics approach to dynamic network exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):1–10, Jan 2016. doi: 10.1109/TVCG.2015.2468078
- [52] M. Vlachos, D. Gunopulos, and G. Kollios. Discovering similar multidimensional trajectories. In *Proceedings 18th International Conference on Data Engineering*, pp. 673–684, 2002. doi: 10.1109/ICDE.2002.994784
- [53] C. Ware, R. Arsenault, M. Plumlee, and D. Wiley. Visualizing the underwater behavior of humpback whales. *IEEE Computer Graphics and Applications*, 26(4):14–18, July 2006. doi: 10.1109/MCG.2006.93
- [54] N. Willems, H. Van De Wetering, and J. J. Van Wijk. Visualization of vessel movements. *Computer Graphics Forum*, 28(3):959–966, 2009. doi: 10.1111/j.1467-8659.2009.01440.x
- [55] M. Wörner and T. Ertl. Visual analysis of public transport vehicle movement. In K. Matkovic and G. Santucci, eds., *EuroVA 2012: International Workshop on Visual Analytics*. The Eurographics Association, 2012. doi: 10.2312/PE/EuroVAST/EuroVA12/079-083