Visual Analytics for Exploring Local Impact of Air Traffic

J. Buchmüller[1], H. Janetzko[1], G. Andrienko[2,3], N. Andrienko[2,3], G. Fuchs[2], D. A. Keim[1]

[1]University of Konstanz, Germany
[2]Fraunhofer IAIS, Germany
[3]City University London, UK

Abstract

The environmental and noise impact of airports often causes extensive political discussion which in some cases even lead to transnational tensions. Analyzing local approach and departure patterns around an airport is difficult since it depends on a variety of complex variables like weather, local and general regulations and many more. Yet, understanding these movements and the expected amount of flights during arrival and departure is of great interest to both casual and expert users, as planes have a higher impact on the areas beneath during these phases. We present a Visual Analytics framework that enables users to develop an understanding of local flight behavior through visual exploration of historical data and interactive manipulation of prediction models with direct feedback, as well as a classification quality visualization using a random noise metaphor. We showcase our approach using real world data from the Zurich International Airport region, where aircraft noise has led to an ongoing conflict between Germany and Switzerland. The use cases, findings and expert feedback demonstrate how our approach helps in understanding the situation and to substantiate the otherwise often subjective discourse on the topic.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications—General H.5 [Information Interfaces and Presentation]: User Interfaces—Graphical user interfaces (GUI) K.4.0 [Computers and Society]: General—

1. Introduction

Aircraft noise is hazardous to health and in affected areas, studies have proven negative effects on cardiovascular [Kni77] and mental health [AWAGH69] and even the reading comprehension and recognition memory of children [SBC*05]. Especially approaching aircraft have a high noise impact on the ground below, as they fly in slower and longer in lower altitudes than departing planes. Consequently, we decided to focus on approaching aircraft throughout this publication, as departures play no role in the discussion. Local conflicts around airports about the distribution of air traffic and noise are clearly comprehensible and a well-grounded discussion is crucial.

Since the 1990s, Germany and Switzerland have been discussing about a fair noise distribution. The international airport of the Swiss city of Zurich is located about 15 kilometers south to the Swiss-German border. Figure 1 shows the spatial distribution of on average over 700 arrivals and departures, a distribution that has not changed considerably since then. Remarkably, almost 80% of inbound air traffic comes in from

north, so that the predominant part of these approach routes leads over mostly German territory.

Figure 1: *Air traffic routes for arrivals (red) and departures (blue) to Zurich airport and average distributions of flights amongst them. Note the Swiss-German border to the north.*

Germans suffering from the aircraft noise intensely protest against these conditions. Despite attempts between the two nations to negotiate treaties regulating air traffic, no agreement

has been reached so far, which is why the German government in 2005 unilaterally imposed restrictions in the *220. Durchführungsverordnung zur Luftverkehrs-Ordnung* (220th executive order to air traffic regulations) that, in general, prohibits direct approaches over German territory between 9pm and 7am. One of the major problems in this extremely complex conflict with political, social, economic, and environmental impact is the lack of a neutral instance. Certainly, there is none helping in the discussion and providing facts rather than assumptions. Facts could enable different audiences to comprehend current and past air traffic, and to find reasons for apparent violations of effective rules. We present a working prototype specifically engineered to suit these information needs.

In comparison to a global view on a flight from one airport to another, on the local level a variety of factors influence the specific way an airplane approaches its destination, for example weather conditions, safety, and political regulations, air exclusion zones, and other traffic. Due to this complexity, we chose a Visual Analytics approach that enables the user to explore past and current flight movements in the region by landing direction, to find flights violating specified rules and to find explanations for apparent violations. Furthermore, we provide a visual interface for a prediction model allowing the forecast of the expected amount of flights with respect to weather conditions. Complemented with prediction comparison and quality measurement visualizations, users can learn to understand flight behavior in the region by testing assumptions against the prediction model.

The structure of the remaining work is as follows: After outlining related and existing approaches in Section 2, we present the basic structure and workflow of the proposed system in Section 3 with the three main tasks of data exploration in Section 3.1, prediction in Section 3.1 as well as comparison in Section 3.1. Afterwards, we show the usefulness of our approach based on real data with use cases and interesting insights in Section 4 and expert feedback in Section 5. We finally conclude with an outlook towards further research questions and possible fields of application in Section 6.

2. Related Work

We first discuss related work in general visual analysis of movement data in Section 2.1. Specific work in air traffic, prediction and uncertainty visualization are covered in Sections 2.2 and 2.3, while Section 2.4 concludes with the positioning of our work within the aforementioned.

2.1. Visual Analysis of Movement Data

Analyzing air traffic falls into the domain of spatio-temporal data analysis [AAB*13]. Common analysis techniques such as aggregation, segmentation, or clustering are applied in a variety of different application scenarios. Tominski et al. [TSAA12] visualize features of road traffic in 3D by stacking trajectories as color-coded attribute bands on the map and on top of each other. Amongst others, Andrienko

et al. [AAW07] also propose the use of Space-Time Cubes [KW05] to display the temporal evolution of a trajectory and clustering to identify common patterns. Flying entities have to be treated differently as they are less restricted by the topology than other moving objects such as, for example, trains or cars. Spretke et al. [SBJ*11] analyze bird movements using semantic segmentation to identify different states like day or night flights or stopovers, which enables to draw inferences about the animals' behavior. Both single trajectories [WAPW06] and sets of trajectories [GWY*11] can be visually explored.

Corresponding to the analysis of pure movement, it is also desired to analyze movement attributes along trajectories. Worner et al. [WE12] investigate the speed feature of public transportation vehicles. Bak et al. perform an automatic analysis of meeting points of objects [BMH*12] and investigate locations of stops [BPSD12]. Dykes et al. [DM03] explore the design space for detecting patterns and external influence factors to movement. The visual analysis of vessel movement is described by Willems et al. [WVDWVW09] with a focus on the visual object speed representation. Wood et al. analyze the beginning and the end of movement trajectories by spatially layouted small multiples of matrices [WDS10,WSD11]. The context of movement, in this case being weather the contextual dimension, is proposed in an interactive tool by Lundblad et al. [LEH09]. Similar to the applied spatio-temporal rule checking, Scheepens et al. [SWvdW*11] use so called blocks on a density map to emphasize regions that are of special interest depending on the underlying vessels' behavior.

2.2. Analysis of Single and Multiple Air Traffic Trajectories

Still, air traffic cannot fly freely and is subject to numerous regulations and flight rules. For the analysis of multiples of aircraft trajectories, Hurter et al. [HTC09] presented a comprehensive framework that allows the user to visualize a limited set of features interactively, supported by extensive Brushing- and Linking capabilities. In another publication, Hurter et al. [HCGT14] compare image-based information visualization techniques for their usefulness in aicraft trajectory data. Furthermore, Hurter et al. [HAG*14] propose a Visual Analytics System to derive wind information from flight trajectories. Klein et al. [KvdZT14] build upon the aforementioned solution and claim to furthermore be able to visually identify pattern changes in flight trajectories both locally and globally. Buschmann et al. [BTD14] present another 2D and 3D approach to the visualization of massive aircraft trajectories and provide responsive techniques for filtering, aggregation and mapping of attributes. Another approach by Albrecht et al. [ALP12] abstracts from single trajectories to density maps of air traffic, also making it possible to show path conflict probabilities in the density map.

2.3. Movement Prediction And Visual Analysis of Uncertainty

Part of this work deals with the visual representation of spatial predictions for flight density. Different approaches exist to predict individual movements, for example of cars in WhereNext by Monreale et al. [MPTG09]. An important finding is that the prediction of movement helps in understanding behavior, but that the type of objects to predict and the scope (e.g. global vs. local) change the requirements for predictions, as Nižetić et al. state in [NF10]. As predictions always incorporate uncertainty about the results, we apply a noise visualization to represent the prediction quality. While we apply it to a whole cell, Kinkeldey et al. [KMKS14] apply noise to the edges of a grid. Cedilnik et al. [CR00] also use the edges of a procedural grid to display uncertainty by applying distortion.

2.4. Positioning our Work

While general rules of spatio-temporal data analysis apply, visual approaches on movement data can be roughly distinguished by their scope either on large amounts or only few trajectories. As well, some approaches focus purely on the visualization of trajectories and/or attributes alone, while other approaches make use of techniques like aggregation, clustering and segmentation to add semantic meaning to movements. One work that is close to ours is [KvdZT14] providing possibilities for the exploration of recorded data and for the visual recognition and comparison of spatial and temporal changes in movement patterns. But in addition, our system helps by providing an overview visualization with stacked bar charts that point the user to times with unusual behavior as described in Section 3.1. Moreover, we provide means to find reasons for unexpected behavior. Also, we describe how our approach helps in understanding the general behavior of flights under adjustable weather conditions and detached from varying individual time intervals with visual interactive prediction feedback in Section 3.1. Drawing conclusions about behavior and testing them is also supported by the comparison visualization we introduce in Section 3.1.

3. System Design

The framework integrates a tool set combining the spatial and temporal exploration of movement data with a prediction view. Additionally our framework contains tools to evaluate and improve predictors on a general and local level. Furthermore, it supports the checking of rules that can implement any combination of spatial and temporal constraints for the recorded trajectories in order to find outliers or violations.

Our system is written in Java, and for the prediction part, we make use of the flexible WEKA machine learning suite provided by Hall et al. [HFH*09]. Please note that our focus was not to provide a classification as accurate as possible, but to determine methods for prediction and classifier error visualization, and to learn from movement predictions combined with interactive feature selection. Our approach partitions the area around Zurich airport in quadratic subcells. For each subcell, we train an individual, regression tree based WEKA-M5Rules classifier using the records of about 12 months of flight traffic. The input features are expressed in intervals of thirty minutes, and contain the time of day, flight density in absolute numbers and the weather features on the runway, which are wind direction, power, temperature, precipitation, and obscuration. We chose these features because they are not only easy to comprehend, but also because according to an expert, these factors are the most influential ones especially for approaching flights.

The air traffic data used throughout this publication is collected from the AirTraffic LIVE webservice by Kramarz and Loeber [KL07].

3.1. Exploration of Recorded Data

The first core feature is the exploration view in Figure 2. This view allows the analyst to review recorded data. Generally, the screen is divided into two parts: On the left (Figure 2 A) is the management area. Here, the user can choose the view to be presented, set visualization options, select data by date and time, and review statistics about the data and the weather situation. The right side (Figure 2 B) is a map providing the spatial reference for trajectories and is able to display more details and aggregated views on the data.

Regular Partitioning The grid to be seen in the background is a tesselation of squares with a length of ten kilometers, covering a square of 200 by 200 kilometers with Zurich airport at its center. The idea of tesselating the space in which the trajectories occur was inspired by Andrienko et al. in [AA08], where the authors used area tesselation to display cell-aggregated characteristics of car traffic in the city of Milan. In our case, the grid cells are used to aggregate values such as flight density. For this work, the side length of each cell has been set to ten kilometers, as this size roughly reflects the possible area of a single aircraft's surface noise footprint and provides a good tradeoff between accuracy and generalization. A grid resolution of 10 km is also used by the European Aviation Safety Agency (EASA) in a study on en-route aircraft noise [S. 09].

Data Selection Workflow At the beginning of the exploration workflow, the user chooses the trajectories he wants to analyze either by time or by the distribution of certain attributes from a dropdown menu, shown in detail on the right-hand side of Figure 3. We provide the data on a daily basis to the user. As flight traffic usually follows daytime patterns, the user can look for times with specific traffic behavior: In the date dropdown menu, the tool supports the display of horizontally stacked bar charts displaying the proportions between either attribute values of the dataset or other proportions related to the moving objects' behavior. By stacking these bar charts vertically, the specific bars can be compared easily. Thus, the analyst is able to quickly determine changes in behavior between the time steps and to spot times with divergent
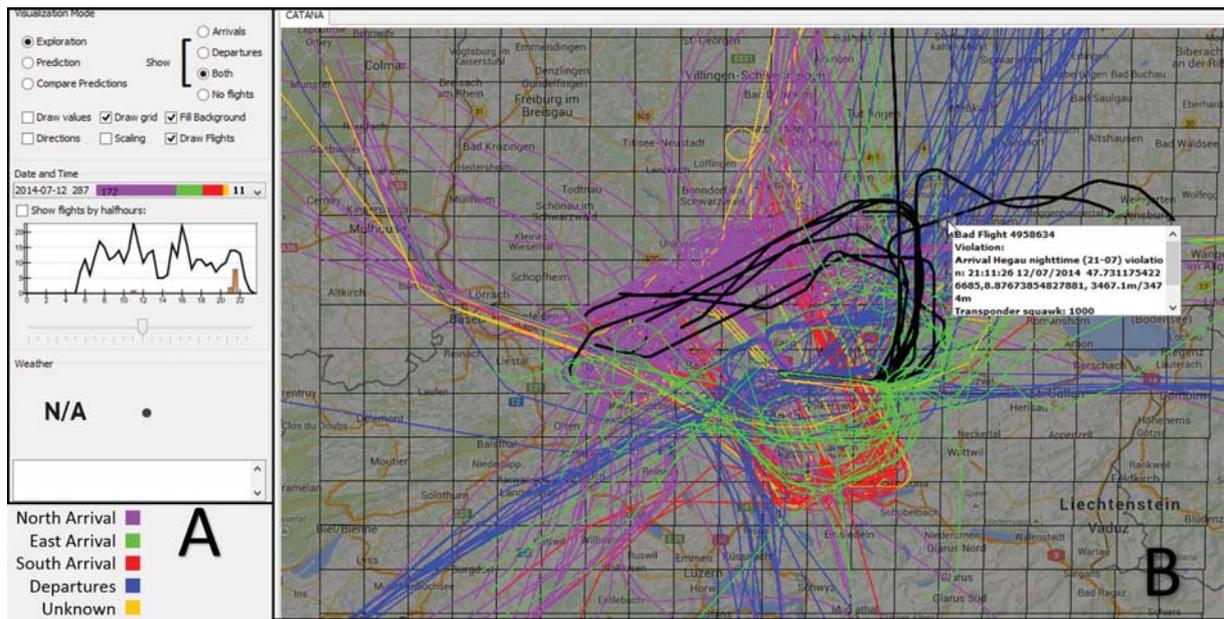
Figure 2: *The Exploration View. Part A denotes the management area, whereas part B contains a map view depicting all selected data and showing details on a flight violating a rule in a textbox-overlay. Black tracks denote flights that violate rules.*
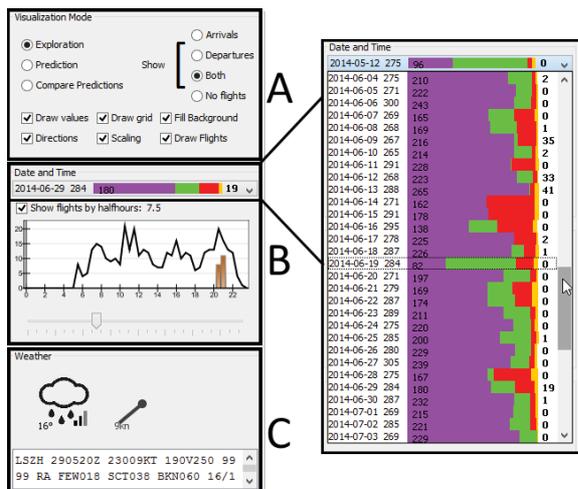

Figure 3: *Management interface with visualization options and behavior filter in part A, date and time selection in part B, weather display in Part C and the overview visualization in the right part.*

or otherwise special behavior. The decision, which properties to display in these bar charts is highly task-dependent and cannot be generalized for arbitrary kinds of data. In this work's example, the proportions reflect the amount of arriving planes either approaching from the north (purple, runways 14 and 16), the east (green, runway 28) or the south (red, runway

34). Due to terrain restrictions, no flights can approach from the west. Orange color indicates a flight were the landing direction could not be determined automatically. As shown in Figure 1, the distribution patterns of arriving flights are usually the same. Thus, deviations from the usual distribution are indicative of extraordinary behavior in the data that the analyst can then examine in detail.

To further aid the analyst, the bar chart can be annotated with more information. Here, the number next to the date reflects the total number of arrivals recorded for the chosen time frame. The second number shows the amount of flights approaching from the north, thus also giving an idea on how many flights the other proportions roughly reflect. The last number shows the amount of flights violating the predefined rules on that day, which can be another indicator for unexpected behavior. We chose to visualize the landing directions as the key point of the bilateral discussion.

In-Depth Analysis In the next step, the analyst reviews the selected data: A combined line and bar chart, as displayed in Figure 3 B, provides information on the temporal distribution of the amount of objects on screen over the chosen time frame. The bar chart indicates the time and amount of flights that violate the predefined rules. By comparing bar and line chart, the analyst can assess the proportion between special, in our case rule-violating, objects and the total amount of present objects.

The analyst has now two choices to investigate the selected trajectories on the map. First, he can choose a cell-wise aggregated view like in Figure 4 displaying the averaged principal

direction. Arrivals and departures are averaged separately, as naturally, arrivals and departures usually take different directions even in the same area. In cells where no flights have taken place, no arrows are displayed. Furthermore, the user can scale the arrow sizes by the amount of flights that occurred in the respective cells compared to all others.

As second option, the user can display all individual trajectories in the selected time. In this view, the coloring of the trajectories matches colors of the proportions from the drop-down menu. Additional properties that have not been included into the bar chart can also be shown on the map. For example, departing flights are represented in blue and flights that violate aforementioned rules in black (compare trajectories in Figure 2). This color coding enables the analyst to quickly distinguish, which areas of the map are distinctive for certain behaviors and in which areas space is shared between the different groups of trajectories. For example, areas traversed by arrivals landing from the north with purple trajectories can be clearly distinguished from areas reserved for eastern or southern approaches in green and red. The plain view on the unfiltered data might still involve too many trajectories to analyze individual behavior. The analyst can narrow down the displayed trajectories by selecting only single time steps. Our tool provides the ability to skip through the data in steps of thirty minutes by adjusting the slider beneath the line chart (compare Figure 3 B). Steps of half an hour seem appropriate, as this relates to the resolution of the available weather information that can be crucial to the aircraft's behavior. As a side effect, the amount of displayed flights also gets reduced to approximately 30 at a time. Another option to reduce the amount of displayed trajectories is the possibility to chose between arrivals and departures (Figure 3(A)), as arriving flights behave differently.

Highlighting of single or multiple trajectories for better comparison by click-selection is supported. Around the mouse pointer, a selection area is provided, and selected trajectories appear in yellow or dark gray (for illicit flights), as shown in Figure 4. As the used map underlay provides information in different colors (e.g., amber for main streets, green for woods, and blue for waters), the contrast between trajectories and map can sometimes be low, making it difficult to distinguish single paths. To avoid this, we chose the color of the trajectories specifically to improve discriminability. In addition, the user can turn on a transparent, dark background overlay between the map and the trajectories to improve contrast and to reduce visual clutter introduced by the map, as well to be seen in Figure 4.

Prediction Model The prediction view includes another feature helping advanced users to actively influence the classifiers' training data by either in- or excluding flights. By default, all regular flights are included, while all flights violating a rule are excluded from the training data. Still, with his knowledge, an expert could be aware of movements that follow other requirements than regular traffic and thus are not helpful or even harmful to the classifiers prediction model.
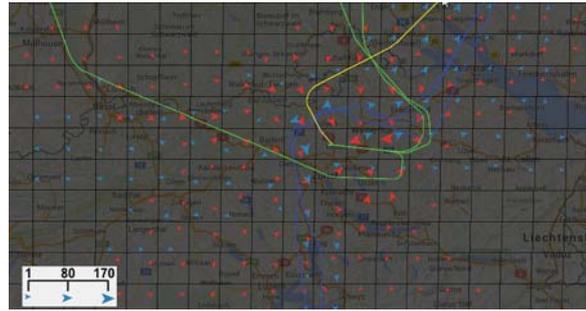

Figure 4: *Per-cell averaged flight direction visualization with the average direction of arrivals in red and of departures in blue. Selected trajectories appear in yellow.*

He can now choose to exclude these flights by selecting them and choosing the respective option. On the other hand, the predefined rules might not regard special permits or exceptions and falsely mark an object as violator. For example, there are no temporal restrictions on arrivals over German territory on the Swiss national holiday - knowledge that should be incorporated into the prediction model.

Finally, the analyst can start to draw conclusions by reviewing the trajectories and accessing additional details about them. Flights are heavily influenced by local weather, so in Figure 3 C, icons visually represent the most important features of temperature, wind strength and direction, and possible kinds and intensity of precipitation and obscuration. The analyst can also access more details about flights by hovering over them. As shown in the upper half of Figure 2, an overlay window will provide further information of the underlying trajectories. The example shows additional information on a flight that violated a predefined rule, giving the user hints for the reason for the violation. The exploration view provides a first overview and to build hypotheses on the behavior of the displayed flights based upon their trajectories and associated attributes. From here, the analyst can now generalize his hypotheses by comparing reality to a prediction model in the prediction view. Alternatively, the user can try to test his hypothesis in the comparison view, where he is able to compare predictions for different weather situations at the same time.
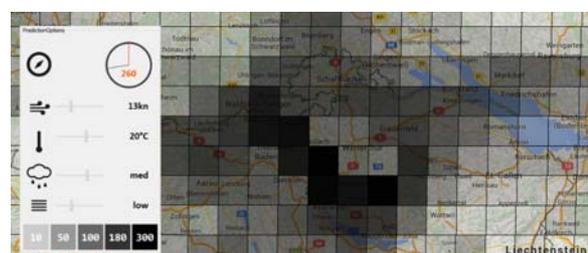

Figure 5: *Prediction view, with input feature selection controls and density representation.*

Prediction of Flight Density And General Behavior In the exploration view, the user is able to review only historical data. In the prediction view, shown in Figure 5, the analyst can find out about the general behavior interactively by visualizing the prediction of a classifier under different conditions. The first step in the prediction view is to initialize the prediction with the desired input features. The classifier was trained with the amount of flights in a given cell's area and the prevalent weather conditions. This allows to predict the flight density in each cell for different weather situations, which can be selected by the user, as he can change the input features of wind direction by turning the compass dial, and, using the sliders, of wind speed, temperature, intensity of precipitation and/or obscuration. Using these controls, shown in Figure 5 on the left, the user can create a custom situation of his wishes or he can recreate a situation he reviewed in the exploration view for comparison. The user gets immediate feedback on his actions, as the result is directly displayed in the map view. The predicted values are visually represented by a transparency-modified black overlay for each cell (Figure 5). The higher the value, the higher the opacity, making cells with high prediction values appear darker and less transparent than cells with only few predicted flights. The result is a density map of the tesselated area in the resolution of the grid cells' side length. If needed, the user can also display the values as text labels directly in each cell. The user experiences an animation-like effect by the immediate feedback and can understand transitions from one state to another. For example, the user can follow the in- or decrease of flights in certain areas directly while sliding the wind strength control back and forth.

At this point it is noteworthy that at first sight, it could appear useful to smooth the predicted values into a soft-edge density map with smooth transitions between the specific density regions. The problem is that such a smoothing would create a false impression of accuracy. The amount of flights going through a single cell is directly related to the impact these flights have noise- or otherwise on the ground below. Smoothing between cells would mislead the viewer to think that the density areas are precisely defined and the impact is modeled exactly, while the usage of square cells still communicates the notion that the impact is averaged for the whole cell area and thus contains a certain displacement error.

Comparison of Predictions And Recorded Data As predictions inherently involve uncertainty about the predicted values, our system enables the analyst to check whether the prediction model actually behaves as expected and in which regions the prediction quality is low. The density map view can be switched from displaying predicted values to the display of real values (for example in Figure 7A and B), but the immediate transition makes it hard for the user to compare the values. That is why, to be able to compare the values in situ, the available space in each cell has to represent both values. As a result, we introduce a picture-in-picture visualization for direct value comparison (Figure 6). An obvious segmenta-
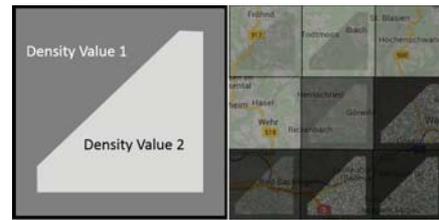

Figure 6: *Example of a cell containing two different density values for comparison (left) and examples for cells with all possible difference combinations and noise-based uncertainty visualization in detail (right).*
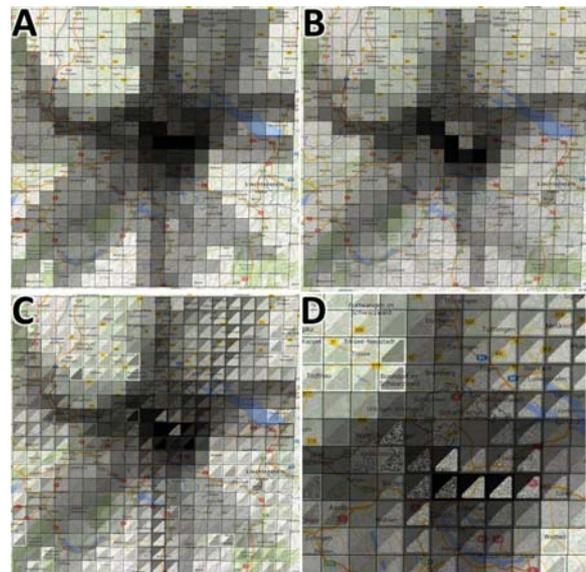

Figure 7: *View on aggregated real-world (A) and predicted density values (B) of the same day. Comparison visualization thereof in (C) and (D) with noise visualization for predictor error.*

tion could be a complete diagonal cut, or a sub-region either centered in the cell or aligned with one corner. Yet, most shapes share a problem: the regularity of the shapes has a perceptional side-effect, creating the strong impression of cohesiveness, distracting from the individual cells. To balance this effect, the chosen shape is clipped from a diagonal cut. One value's surface visually engulfs the other, creating the impression of a fore- and background. This design alleviates the effect of the perception of cohesive areas between neighboring cells and at the same time still allows the user to quickly spot the tendency and coarse amount of the difference between the two values. The design strongly boosts the visual distinction between areas where predicted values match the real ones and areas where this is not the case. Another application for the comparison view is to compare two different predictions for different weather conditions. Learning from different input feature sets and being able to compare them helps the user to get an understanding of the way air traffic
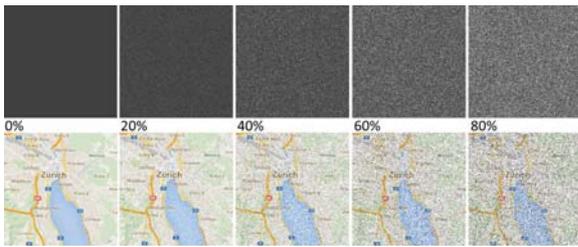
Figure 8: *Noise applied to a solid, dark gray image and a map excerpt showing Zurich. With increasing noise, the incremental loss of information can be observed, while the main features of the map excerpt stay visible.*

works in that region, as he is able to relate density development to changes in features. Comparing two situations helps him to find general hypotheses instead of slowly having to build one by iterating through many days manually.

Finally, the advanced user might also be interested in the prediction confidence of the applied classifier for each cell. In the field of machine learning, the quality of classifiers usually is measured by statistical methods like relative or mean squared error. Basically, measures like these allow assertions over bias and variance of the prediction results, allowing to judge whether the predictor produces consistently realistic values or not. But these quality measures usually do not regard the spatial dimension of spatio-temporal data. Yet, for the analyst it could be very interesting to see exactly where a classifier is confident to predict values close to reality, but still displays deviant values for the currently selected conditions. This information can help to find areas where the objects behave differently than expected/modeled, as areas with high classifier uncertainty indicate that the prediction input features do not help in explaining the moving objects' behavior in these regions.

Displaying this information in the prediction view is challenging, because most visual variables are already mapped or cannot be used since position, size, shape and orientation are preordained by the grid cell and difference visualization shapes, and color cannot be applied because this might interfere with colors of the map underlay. However, one visual variable is left: Texture. Textures cover an area with more or less recurring patterns, for example grids of lines or dots. We apply random noise to the image, much as one would expect for example from a mistuned analog TV. The amount of noise applied to an image can be regulated, and increasing noise gradually reduces the information in the image the noise is applied to. An example for different amounts of random noise applied to an image can be seen in Figure 8. The noise gradually degenerates the underlying image, thus reducing details step by step. This seems to be an ideal metaphor for uncertainty, as with increasing noise, the visual quality of the image decreases together with the classifier confidence. In other words, the less the viewer can perceive of an area, the less sure is the classifier about the moving objects' behavior in that area. While the specific amount of noise cannot be

easily estimated by a viewer, the user can tell the difference between small and large amounts and estimate the dimension of noise easily and can also spot areas with several cells where the error in- or decreases. Areas affected by larger uncertainty than others indicate that the used classifier or its parameterization, respectively, cannot sufficiently explain the behavior of moving objects in this area, which itself can be important to know for the analyst. In Figure 7 D, different levels of noise with different prediction values can be observed. We use Gaussian Noise to calculate a randomized gaussian distribution of gray values to be applied to the image. The amount of resulting noise can be parameterized because the function is dependent on the standard deviation σ that can be modified as needed, as for example explained in [Bar13].

4. Case Studies

For our prototype, we identified two different groups of users: First, there are professional and semiprofessional users that already possess a certain degree of domain knowledge. Such users could for example comprise journalists or PR officials. The primary tasks of these users are the analysis of long-term data and identification and explanation of outstanding flight patterns.

The second user group is formed by the interested general public, e.g., local residents or real estate customers. These users are more interested in the current situation, as they want to check instantly, why their area experiences so much traffic at this very moment. As well, they are interested in using the prediction to model typical weather situations to check for the expected traffic in their vicinity. In the following, we will demonstrate how the presented prototype can fulfill these information needs based on flight data from June to September 2014 for the two groups of users we identified. The derived insights are real examples for interesting findings.

For both groups, no services exist that can provide functionality to display recorded flights over longer periods of time, or any possibility to check for rule violations, to view statistics about them or to support the creation of a mental flight behavior model.

4.1. Professional Users

As the regional conflict has been ongoing for years now, many municipalities in the area maintain a close watch on the situation, and react to citizen's complaints about possible violations. Thus, commissaries on the issue need to be able to quickly review the time period since their last analysis, finding out whether the air traffic complied to existing rules or violations did indeed occur. This functionality is provided in the exploration view in Figure 9.

The analyst uses the date selection menu to check whether on mentioned days any violations, represented by the rightmost number, occurred (Figure 9 (1)), and thus can already dismiss unjustified reports about days with no violations. On days where violations happened, our prototype supports the analyst by showing when the violations happened at that day
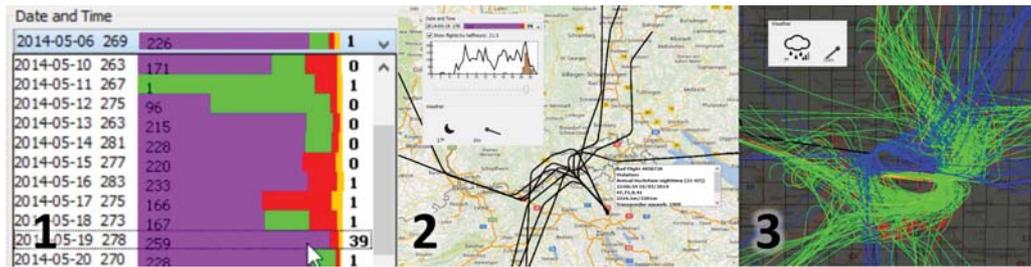
Figure 9: *Exploration with aggregated overview (1), flights violating regulations (2) and a day with no northern approaches (3).*

in the bar chart to the left (Figure 9 (2)) and by providing the weather report for the time the violations happened. Furthermore, the analyst can view details of each suspected illicit flight that might be of interest, showing the precise rule that was violated, for instance the current altitude and the type of aircraft. An example can be seen in Figure 9 (2), where landings during night time have to take the northern route, as a six knot south-eastern wind results in unsafe landing conditions for aircraft approaching from the south or east.

Finally, the analyst has found sound explanations for most violations, but also a few flights are left for which he cannot find any explanation, leaving room for further research. In the same manner, the analyst can as well validate official explanations for certain exceptions or verify the statements made in reports on rule violations published by airport authorities.

While the analyst is still browsing the dates, a certain specific date catches the attention of the analyst in Figure 9 (1), May 11th: On this date, no flights have approached the airport from the usual northern direction, as approaches via the north would appear in purple. In fact, it seems only one flight landed from the north and all other flights landed from either the east (green) or the south (red). Simultaneously, the number of recorded flights is not lower than on other days. As a next step, the analyst can select the day to check out more details (Figure 9 (3)). The overview visualization shows where flights occurred and where not. The analyst notices that all incoming flights are taking other routes than the northern one. The airplanes land from the east and depart towards the north. Checking the weather, the analyst realizes that strong winds were blowing most of the day effectively preventing approaches from the north. While there is a sound explanation for the unusual behavior, this is as well a very surprising and outstanding finding, as it proves that the default operating concept is not the only possible one. Other operational concepts can support the same amount of daily flights and are not as far out of the question for regular operations as sometimes claimed.

Summarizing, the analyst is able to conveniently check several dates for reported incidents, filtering out dates not worth looking at without the need to take a direct look at that data. Furthermore, it is possible to compare the provided environmental conditions for any given time, thus allowing the user to draw conclusions or to make decisions, for example whether an exception was valid or not. Finally, already while

browsing the available dates, the stacked bar chart visualization helps the analyst to find interesting behavior on the fly, and he can directly take a look and find out more.
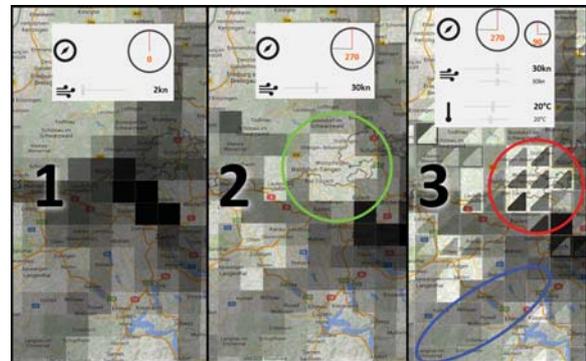
4.2. General Public

Figure 10: *Two density predictions, showing differences between (1) and (2). (3) shows the comparison visualization with two opposing weather situations.*

Especially for an audience with less experience on the subject, it is very hard to understand all the factors that have an influence on flight density in their direct vicinity. Just from taking a look at the recorded data day by day, people can neither learn about the general aircraft behavior in their area, nor can they learn about the conditions under which pattern changes occur.

Our prototype enables users like future home owners and interested citizens to interactively model desired conditions and to check the expected flight density. The derived knowledge could influence someone in buying a house or, on the other hand, could help someone to apprehend and adapt to situations he previously did not understand.

Figure 10 shows the results of a users interactions with our prototype. The user wants to investigate the influence of the weather conditions on flight patterns, so he uses the prediction view. Our prototype presents a density visualization for an average day with low winds and about 20 degrees Celsius (Figure 10 (1)). Dense and sparse areas can be identified immediately by comparing the cell opacity. Changes to the selected weather conditions using the respective sliders will invoke real-time transitions in saturation, when prediction

values are changing as in Figure 10 (1) to (2). The transitions directly help the user to understand in which direction the expected density values change when the weather changes. In this example, the user can learn how strong westerly winds force flights to land from the east, effectively reducing the number of flights approaching from the areas north to the airport (area marked in the green circles).

To enable the user to compare two predictions for different situations, for example two common, but different weather situations, he can use the prediction comparison view (Figure 10 (3)), in this example westerly wind vs. easterly wind. The outer part of each cell represents the predicted amount of air traffic for the first selected weather situation, the inner for the second (compare Figure 6). The user can now immediately spot regions with large differences in predicted flight numbers, marked in red, thus being able to identify regions heavily influenced by the modified input features. The region marked in blue shows an area where the changed input features do not influence the density much or at all, depicted by grid cells with uniform colors.

Summarized, the analyst is able to easily identify areas that experience less flights than others. Our system provides options to check the predictions for different weather situations and allows to directly compare them. In addition, for both single and compared predictions, the user is able to consult the uncertainty visualization as depicted in Figure 6 with different amounts of uncertainty, thus being able to judge how accurate the predictions are. Showing the uncertainty to the user helps avoiding hasty conclusions and cognitive biases by the user, as the user now knows where he can trust the data and where not.

5. Expert Feedback

We are in contact with a political journalist, who works for the local newspaper Südkurier and has covered the conflict on a local and national level for now more than 15 years. He is certain, that our approach would be very beneficial for his work: Violation detection, traffic pattern overview, the exploration and comparison of data of longer time spans up to several years, and the respective analysis will help journalists to verify claims made by involved parties. Up to now, there is no easy way to detect unusual situations, which is an important part of his work. We presented and explained our prototype and use-cases to the journalist and asked for feedback and suggestions for improvement. After a short introduction, our expert was able to use the prototype by himself. He explored different days and time intervals and started speculating about different days with clearly differing flight behavior. While Browsing flights violating regulations, he quickly started to try to find reasons, assisted by the additional weather and flight information displayed. In particular, our expert found the stacked bar chart overview to be very intuitive and helpful to decide which days to investigate further, e.g. in order to search for signs of new operational flight approach concepts.

While the overall feedback was very good, our expert had

improvement suggestions: Foremost, a live overview would be desirable for him as well as the possibility to visualize more features like aircraft type and altitude directly on the trajectories. The system should also hint probable causes for violations and present the results to the user. Incorporation of noise level information and prediction would be desirable. According to the expert, there will be a future challenge with new and more complex flight regulations that will be more flexible, hindering a simple automatic rule violation reasoning. Our system will have to be adjusted to still be able to determine reasons for violations.

6. Conclusion

We presented a Visual Analytics approach to explore and investigate local air traffic behavior. We enable users to develop an understanding of how flight traffic evolves in his specific area of interest. The visual exploration together with our weather-dependent prediction enables users to get further insights. As prediction always comprises uncertainty, we can call the users' attention to regions with lesser prediction quality visually using Gaussian noise image modification. Additionally, we can both detect flights violating regulations and it is often possible to find explanations.

We were able to identify some flights that violated rules for no apparent reason, mostly in the evenings. But we were also able to give a sound explanation for other violations. This could prove very beneficial for the political discussion, because allegations can be checked neutral point of view. Also contributing to the conflict discussion is the possibility to quickly find days with divergent flight behavior showing different modi operandi. This is clearly shown in the surprising finding described in section 4.1 that had a notable impact on the local debate. Backed by our expert feedback, we are convinced that our approach is a valuable contribution not only to the Zurich Airport conflict, but for other regions around airports worldwide. The local community is extremely interested in our work and we have received several requests to publish our tool. For this reason, we plan to provide access to the tool to more people, for example with a web service that can specifically adress different user groups, also possibly expanding to other areas as well. Until then, we want to work on the integration of further data sources like flight schedules and aircraft noise models in order to improve reasoning for rule violations and traffic impact, enabling a fact-driven discussion between citizens and air traffic planners.

7. Acknowledgement

References

[AA08] ANDRIENKO G., ANDRIENKO N.: Spatiotemporal aggregation for visual analysis of movements. In *In Proceedings of IEEE Symposium on Visual Analytics Science and Technology (VAST 2008)* (2008), Society Press. 3

[AAB*13] ANDRIENKO G., ANDRIENKO N., BAK P., KEIM D. A., WROBEL S.: *Visual Analytics of Movement*. Springer, 2013. 2

[AAW07] ANDRIENKO G., ANDRIENKO N., WROBEL S.: Visual analytics tools for analysis of movement data. *SIGKDD Explor. Newsl. 9*, 2 (Dec. 2007). 2

[ALP12] ALBRECHT G. H., LEE H. T., PANG A.: Visual analysis of air traffic data using aircraft density and conflict probability. In *Infotech@ Aerospace 2012*. 2012. 2

[AWAGH69] ABEY-WICKRAMA I., A'BROOK M., GATTONI F., HERRIDGE C.: Mental-hospital admissions and aircraft noise. *The Lancet 294*, 7633 (1969). 1

[Bar13] BARBU T.: Variational image denoising approach with diffusion porous media flow. In *Abstract and Applied Analysis* (2013), Hindawi Publishing Corporation. 7

[BMH*12] BAK P., MARDER M., HARARY S., YAELI A., SHIP H. J.: Scalable detection of spatiotemporal encounters in historical movement data. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 915–924. 2

[BPSD12] BAK P., PACKER E., SHIP H. J., DOTAN D.: Algorithmic and visual analysis of spatiotemporal stops in movement data. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems* (2012), ACM, pp. 462–465. 2

[BTD14] BUSCHMANN S., TRAPP M., DÖLLNER J.: Real-time animated visualization of massive air-traffic trajectories. In *Proceedings of CyberWorlds* (2014), IEEE Computer Society. 2

[CR00] CEDILNIK A., RHEINGANS P.: Procedural annotation of uncertain information. In *Visualization 2000. Proceedings* (2000), IEEE, pp. 77–84. 3

[DM03] DYKES J. A., MOUNTAIN D. M.: Seeking structure in records of spatio-temporal behaviour: visualization issues, efforts and applications. *Computational Statistics & Data Analysis 43*, 4 (2003), 581–603. 2

[GWY*11] GUO H., WANG Z., YU B., ZHAO H., YUAN X.: Tripvista: Triple perspective visual trajectory analytics and its application on microscopic traffic data at a road intersection. In *Pacific Visualization Symposium (PacificVis), 2011 IEEE* (2011), IEEE, pp. 163–170. 2

[HAG*14] HURTER C., ALLIGIER R., GIANAZZA D., PUECHMOREL S., ANDRIENKO G., ANDRIENKO N.: Wind parameters extraction from aircraft trajectories. *Computers, Environment and Urban Systems 47* (2014), 28–43. 2

[HCGT14] HURTER C., CONVERSY S., GIANAZZA D., TELEA A.: Interactive image-based information visualization for aircraft trajectory analysis. *Transportation Research Part C: Emerging Technologies 47* (2014), 207–227. 2

[HFH*09] HALL M., FRANK E., HOLMES G., PFAHRINGER B., REUTEMANN P., WITTEN I. H.: The weka data mining software: an update. *ACM SIGKDD explorations newsletter 11*, 1 (2009), 10–18. 3

[HTC09] HURTER C., TISSOIRES B., CONVERSY S.: Fromdady: Spreading aircraft trajectories across views to support iterative queries. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (Nov. 2009). 2

[KL07] KRAMARZ D., LOEBER A.: Visualisierung von Transponder-Daten mittels Mashup. *Diplomarbeit, Zürcher Hochschule für Angewandte Wissenschaften* (2007). 3

[KMKS14] KINKELDEY C., MASON J., KLIPPEL A., SCHIEWE J.: Evaluation of noise annotation lines: using noise to represent thematic uncertainty in maps. *Cartography and Geographic Information Science 41*, 5 (2014), 430–439. 3

[Kni77] KNIPSCHILD P.: V. medical effects of aircraft noise: Community cardiovascular survey. *International Archives of Occupational and Environmental Health 40*, 3 (1977). 1

[KvdZT14] KLEIN T. R., VAN DER ZWAN M., TELEA A.: Dynamic multiscale visualization of flight data. In *VISAPP (1)* (2014), pp. 104–114. 2, 3

[KW05] KAPLER T., WRIGHT W.: Geotime information visualization. *Information Visualization 4*, 2 (2005), 136–146. 2

[LEH09] LUNDBLAD P., EURENIUS O., HELDRING T.: Interactive visualization of weather and ship data. In *Information Visualisation, 2009 13th International Conference* (2009), IEEE, pp. 379–386. 2

[MPTG09] MONREALE A., PINELLI F., TRASARTI R., GIANNOTTI F.: Wherenext: a location predictor on trajectory pattern mining. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* (2009), ACM, pp. 637–646. 3

[NF10] NIŽETIĆ I., FERTALJ K.: Automation of the moving objects movement prediction process independent of the application area. *Computer Science and Information Systems/ComSIS 7*, 4 (2010), 931–945. 3

[S. 09] S. L., ANOTEC CONSULTING: Background noise level and noise levels from en route aircraft (banoerac). Final report to the European Aviation Safety Agency, 2009. 3

[SBC*05] STANSFELD S. A., BERGLUND B., CLARK C., LOPEZ-BARRIO I., FISCHER P., ÖHRSTRÖM E., HAINES M. M., HEAD J., HYGGE S., VAN KAMP I., ET AL.: Aircraft and road traffic noise and children's cognition and health: a cross-national study. *The Lancet 365*, 9475 (2005). 1

[SBJ*11] SPRETKE D., BAK P., JANETZKO H., KRANSTAUBER B., MANSMANN F., DAVIDSON S.: Exploration through enrichment: A visual analytics approach for animal movement. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (2011), GIS '11, ACM. 2

[SWvdW*11] SCHEEPENS R., WILLEMS N., VAN DE WETERING H., ANDRIENKO G., ANDRIENKO N., VAN WIJK J. J.: Composite density maps for multivariate trajectories. *Visualization and Computer Graphics, IEEE Transactions on 17*, 12 (2011), 2518–2527. 2

[TSAA12] TOMINSKI C., SCHUMANN H., ANDRIENKO G., ANDRIENKO N.: Stacking-based visualization of trajectory attribute data. *Visualization and Computer Graphics, IEEE Transactions on 18*, 12 (2012). 2

[WAPW06] WARE C., ARSENAULT R., PLUMLEE M., WILEY D.: Visualizing the underwater behavior of humpback whales. vol. 26, IEEE Computer Graphics and Applications, pp. 14–18. 2

[WDS10] WOOD J., DYKES J., SLINGSBY A.: Visualisation of origins, destinations and flows with od maps. *The Cartographic Journal 47*, 2 (2010), 117–129. 2

[WE12] WÖRNER M., ERTL T.: Visual analysis of public transport vehicle movement. In *EuroVA 2012: International Workshop on Visual Analytics* (2012), The Eurographics Association. 2

[WSD11] WOOD J., SLINGSBY A., DYKES J.: Visualizing the dynamics of london's bicycle-hire scheme. *Cartographica: The International Journal for Geographic Information and Geovisualization 46*, 4 (2011), 239–251. 2

[WVDWVW09] WILLEMS N., VAN DE WETERING H., VAN WIJK J. J.: Visualization of vessel movements. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library. 2