

Dynamic Visual Abstraction of Soccer Movement

D. Sacha¹, F. Al-Masoudi¹, M. Stein¹, T. Schreck², D. A. Keim¹, G. Andrienko³, and H. Janetzko⁴

¹University of Konstanz, Germany ²Graz University of Technology, Austria

³Fraunhofer Institute IAIS, Germany & City University London, UK ⁴University of Zurich, Switzerland




Figure 1: Abstracted movements are shown together with overviews on varying levels of abstraction. On the left, all attack turns of both teams with a shot on goal event of an entire soccer match are shown revealing the different attacking styles of the teams. On the right, several player trajectories are shown for a few seconds of movement revealing similarly moving groups on the soccer pitch.

Abstract

Trajectory-based visualization of coordinated movement data within a bounded area, such as player and ball movement within a soccer pitch, can easily result in visual crossings, overplotting, and clutter. Trajectory abstraction can help to cope with these issues, but it is a challenging problem to select the right level of abstraction (LoA) for a given data set and analysis task. We present a novel dynamic approach that combines trajectory simplification and clustering techniques with the goal to support interpretation and understanding of movement patterns. Our technique provides smooth transitions between different abstraction types that can be computed dynamically and on-the-fly. This enables the analyst to effectively navigate and explore the space of possible abstractions in large trajectory data sets. Additionally, we provide a proof of concept for supporting the analyst in determining the LoA semi-automatically with a recommender system. Our approach is illustrated and evaluated by case studies, quantitative measures, and expert feedback. We further demonstrate that it allows analysts to solve a variety of analysis tasks in the domain of soccer.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques H.5.2 [Information Interfaces and Presentation]: User Interfaces—User-centered design

1. Introduction

The recent advance of modern motion tracking technologies makes movement analysis practically applicable in many different domains. An example is professional competitive team sports (such as soccer) where movement is bound by a specific area, restricted to specific rules, and coordinated by tactical and reactive behavior. Visual Analytics (VA) techniques for movement analysis can support users from media, sports analysis, coaching, and management to extract valuable insights from this data. However, raw tracking data is hard to process and interpret due to severe overplotting and clutter when visualized. This is true even for a few minutes of movement to

be visualized (see Figure 2 for an illustrative example). A generic solution to handle the trajectory overplotting problem is to simplify, abstract, or aggregate the trajectories, a concept that is known as *Generalization* in cartography [MRS07]. A plethora of techniques for trajectory generalization have been proposed and applied to a variety of data types and application domains [DBC*15]. However, the techniques are often specific and hard to generalize. For example, individual techniques may be useful for a specific amount of tracked data points (e.g., hundreds vs. thousands of tracked trajectories), or designed for a specific analysis task and application domain, or restricted by computational complexity. Further, analysis goals are often ill-defined requiring adaptive exploration of abstractions.




Figure 2: Raw trajectory plots of 22 player and the ball movements for increasing time intervals. The visualizations suffer from overplotting. This paper tackles the problem of applying an adequate generalization technique that dynamically adapts to the data that needs to be visualized.

To address this problem, we propose a novel VA approach that combines several state of the art generalization methods, and allows an analyst to interact and explore the space of trajectory abstractions, to find the most appropriate for the analysis task and data set at hand. Specifically, we implement trajectory simplification and aggregation techniques as layers of a generalization hierarchy. As a running use case, we apply and evaluate our approach to soccer trajectories. As the basis for steering the generalization techniques, we introduce a *level of abstraction* (LoA) parameter that is mapped to specific abstraction layers and configurations. This enables us to provide the analyst with smooth transitions between the used techniques and to support analysts in building a mental model of the underlying computations. Our application to soccer movement was inspired by tactical board drawings (sketches) of soccer experts, where abstraction is naturally done. We implement interactive visualizations supporting a variety of analysis task and further support the analyst in finding an adequate configuration of the technique by adding a recommender system that learns and predicts the desired LoA based on captured measures from explicit user feedback. The applicability and usefulness of our technique is illustrated and evaluated for the domain of soccer movement analysis with two qualitative and one quantitative user studies. Especially the interactive and dynamic computations in real-time helped the analysts to understand and assess the generated visualizations.

2. Related Work

We review general related work in visual analysis of trajectory data and applications including team sport. Afterwards, we detail methods for trajectory reduction and aggregation, followed by a discussion of existing works in user guidance and recommending for visual exploration. Finally, we emphasize the novel aspects of our approach as compared to existing works.

2.1. Visual Analysis of Trajectory Data and Applications

As a basis, object movement can be described by trajectories, i.e., positions in a movement space as a function of time. Due to advances in sensing technologies, trajectories can be recorded at high accuracy and spatio-temporal resolution, e.g., for traffic vehicles, animals, or the members of a sports team during a match. The recent textbook on visual analysis of movement [AAB*13a] proposes a comprehensive approach covering multiple analysis perspectives: focusing on moving objects, locations, and time intervals and proposing corresponding methods suitable to various application domains.

Analysis of trajectory data is important in many application contexts. For example, the study of *movement ecology* in biology concerns understanding movement patterns of organisms in nature. An

encompassing overview of movement ecology analysis is given in [DBC*15]. The authors distinguish tasks in analyzing spatio-temporal patterns, classification and identification of behavior, and the relation between movement and the environmental surrounding. Also, trajectory-based analysis of team sport data has recently gained interest, with a number of works addressing the soccer case. For example, Perin et al. [PVF13] introduce visual designs to present and summarize game situations in soccer matches. Soccer pass analysis and methods to cluster player trajectories are addressed by Gudmundsson and Wolle [GW14]. In own previous work, we describe feature-based techniques to visualize, segment, and classify soccer data [JSS*14]. We also describe work for the visual exploration of soccer player interactions and free space situations in an interactive system [SJB*16] and, furthermore, introduce an approach for sketch-based search and comparison of soccer movements [SSN*16]. There are manifold other applications for visual trajectory analysis, including analysis of moving persons in buildings [IWSK07] or vehicles in a street network [ADKZ*16].

2.2. Visual Aggregation, Reduction and Simplification of Trajectory Data

Next, we detail several existing approaches to reduce the amount of rendered information for trajectory data. In *visual abstraction*, the goal is to summarize large trajectory data using appropriate visual representations. Willems et al. [WvdWvW11] compare the effect of density-based aggregation, animated dots, and the space-time-cube technique regarding movement pattern understanding. *Data reduction* techniques can be applied to reduce data size prior to the analysis and visualization. Data clustering is a well-known data reduction tool that forms groups of similar data items. Based on appropriate trajectory similarity functions or feature-based representations [PKM*07], clustering algorithms, e.g., k-means, hierarchical or density-based clustering can be applied [HKP11]. For example, Andrienko et al. [AAR*09] use density-based clustering to analyze a large set of traffic trajectories via a smaller number of clusters. Schreck et al. [SBTK09] employ the Self-Organizing Map algorithm for trajectory clustering to generate overview visualizations and support pattern comparison. Besides clustering, also filtering approaches can reduce trajectory data. We note that clustering is usually applied to complete trajectories, while filtering finds sub-trajectories. Von Landesberger et al. [vLBSF14] apply a moving average analysis to certain time-dependent trajectory features, with the goal to filter for potentially interesting sub-trajectories. Trajectory filtering based on predefined trajectory features or user-sketched trajectory outlines is further proposed by Hurter et al. [HTC09]. Additionally, Janetzko et al. [JSS*14] propose trajectory feature

visualizations for interactive filtering and identification of relevant sub-trajectories.

While clustering and filtering reduce the number or size of trajectories, *simplification* techniques [MRS07] are available to reduce the level of detail of trajectories. Andrienko and Andrienko [AA11] propose a tessellation of the space by Voronoi polygons that reflect the density of key points of trajectories, and represent trajectories as flows between the polygons. Several works consider trajectory simplification to analyze trajectories of groups of objects or subjects. The work of Laube et al. [LIW05] represents trajectories as sequences of turns, supporting the analysis of coordination of moving objects. Andrienko et al. [AAB*13b] propose to re-map trajectory coordinates to a so-called group-space, thus allowing to find different roles of group members. Furthermore, a similar transformation is used for characterizing football situations in time intervals and clustering time intervals by so defined feature vectors [AAB*16].

2.3. User Guidance and Recommending

Our approach includes learning of trajectory abstraction levels from user feedback, hence we relate to recommending and relevance feedback. The main goal of recommender systems [JZFF10] is to support user search by suggesting previously unseen yet potentially relevant information. This can be seen as a problem of classification of relevance. Recommender systems often rely on information from a given user basis, such as user profile (e.g., classes of users), explicit or implicit user feedback information (e.g., provided ratings or reviews), and log data. In Information Retrieval, relevance feedback techniques [BYRN11] distinguish relevant and irrelevant information items based on explicit user feedback. Recent works have applied approaches from recommending and relevance feedback to support visual-interactive exploration. Relevance feedback can be applied to predict relevance of previously unseen scatter plot views, effectively narrowing the search space [BKSS14]. Healey and Dennis [HD12] train a classifier from user-selected views in a larger geospatial data set, supporting navigation to unexplored data areas. Finally, in our previous work we train classifiers to find potentially interesting scenes in soccer matches based on trajectory features [JSS*14] and explicit user feedback.

2.4. Summary and Novel Aspects of our Work

Related work on the visual analysis of trajectory data provides useful building blocks for an interactive exploration system, but not yet allows for real-time interactive data reduction, simplification, and aggregation. Our work addresses this challenge by combining such existing trajectory abstraction techniques with the aim to enable real-time interactive search of appropriate data reduction and simplification levels for visual trajectory exploration. With such a formulation we bridge the gap between the prior work and our novel developments. We also support this exploration by applying a simple recommender module that can suggest potentially useful levels of aggregation and abstraction, based on an existing set of training observations. Our approach hence can improve the efficiency of trajectory analysis and dynamically adapt to user preference and analysis context. By application, our work contributes to data analysis for soccer data, however, is not limited to this.

3. Dynamic Visual Abstraction of Soccer Movement

Our approach combines three fundamental principles of data or movement generalization, which arose from our discussions and have been identified in related works:

- *Focusing on Relevant Information* Detail
- *Trajectory Simplification* Simplification
- *Summarizing Similar Movements* Aggregation

Following our “focus - simplify - aggregate” mantra, we perform visual abstraction and describe in the remainder of the paper how we realized and implemented them in the domain of soccer movement.

3.1. Design Study Methodology

We conducted a user-centered and iterative design study methodology process [SMM12] to build our visual abstraction system. We used real data from professional soccer matches and work in close collaboration with a domain expert (E1) that already worked with us in previous studies. He works for the German soccer club FC Bayern München as a certified coach in the youth sector. Furthermore, he has been an active soccer player for 24 years. During our previous projects, we came across the problem of abstracting the movement depending on the selected amount of data.

To identify the challenges and requirements of a meaningful visual representation as described in the next section, we invited the expert for a first round of interviews followed by a research group discussion about framing and defining the project. We used the resulting abstraction requirements as an input for our initial design phase choosing proper (from data perspective) simplification and aggregation algorithms. With this bundle of methods and a first prototypical implementation, we were able to discuss with E1 in a more concrete manner. On the one hand, E1 was able to provide feedback about the implemented abstraction layers, used techniques, and parameterizations. On the other hand, we were able to elaborate on a variety of concrete analysis tasks (see Section 4.2). We iteratively implemented and revised the order of simplification and aggregation methods until the sequence of abstraction steps felt natural to the domain expert. In a subsequent and longer implementation phase, we built the major parts of our final system.

Hence, we conducted a qualitative user study with the aim to validate our design. For this study, we were able to recruit a second (unbiased) soccer expert (E2). He was invited from the group of active amateur players and is playing soccer actively for 16 years and has been part of several teams and leagues (e.g., the third-highest league of Switzerland). We decided for this second expert to have a “broad” representation from people of different levels of professionalism and education that could make use of such a system. Additionally, we decided to fine-tune the detailed parameterizations with another qualitative user study with E1 (see Section 5). Based on the feedback and our observations of both studies, we were able to apply final adaptations and to define the “default” parameterizations.

We also observed that a recommender system could be a useful addition to support exploration. After the implementation of a recommender system, we conducted a quantitative user study with both experts to measure its accuracy and were able to identify next steps (see Section 6). Finally, we reflected our design process (see Section 7) during the writing of this paper.




Figure 3: The image illustrates our basic idea and implementation of chaining several different trajectory abstraction techniques. The different layers and used techniques are shown on top. In the center of the figure, we can see how the techniques are parameterized and mapped to a global level of abstraction parameter. The bottom shows the visual outputs for some of the steps within this abstraction space.

3.2. Data and Domain Requirements

We derive our work motivated by soccer analysis. The typical data set of interest contains 22 captured trajectories (one trajectory for each player on the soccer pitch) and manually annotated event streams. Events are captured for passes, shots, throw-ins, fouls, cards etc., each event containing a timestamp, position (x,y), actor-id, and an event-type. For some matches, the ball trajectory is available but it is possible to derive and interpolate the ball movements from respective events (e.g., passes and shots). The trajectories are captured with a resolution of 100ms resulting in approximately 1,242,000 data points (10 points per second × 60s × 90 (minutes per match) × 23 (trajectories)) per match. Our system is able to handle soccer specific changes to the dataset as for example player substitutions or team changeovers after the half-time break. We implemented an optional pre-analysis routine that rotates all movements of the second half by 180 degrees to enable an half-time independent analysis.

Our visualization methods support the understanding of complex and dense movement behavior by visual abstraction. While short time windows in principle allow very detailed analysis views, we also need means of abstraction for single and multiple time spans. Of course, abstraction techniques suitable for movement patterns are highly dependent on the length of the selected time span. Dealing with intervals ranging from seconds over minutes to hours requires adaptive and dynamic techniques. However, the design space of visual abstraction is large and needs to be narrowed down. Consequently, we involved a domain expert and discussed the needs for abstraction methods. We derived the following four requirements:

- R1** The abstracted movement needs to express that we do not visualize the raw movement data but rather an approximation.
- R2** All data points, including outliers as they represent important tactical features, need to be reflected in the abstraction.
- R3** The overall abstraction algorithm needs to be parametrized and allow for varying levels of abstraction.
- R4** Computing the abstracted visualization must be possible in interactive time.

Our initial idea was to visualize the ball movement enriched with important events and players. We therefore focus on these developments in the following paragraphs and discuss in detail our chosen abstraction methods depicted in Figure 3. Note, that we selected the methods based on the application needs and expert feedback. Nevertheless, the proposed techniques are exchangeable depending on the respective analysis and visualization requirements.

3.3. Detail Layer - Focusing on Important Information

We define a “team-turn” as the sequence of actions (or events) of a specific team in ball possession. Obviously, the ball trajectory is the most important information in this case, as it determines all the other movements. Consequently, we visualize the movement of the ball and numbered glyphs for the involved players (Detail in Figures 3 and 4). Movement types (e.g., passes, dribbling, ball reception) are displayed by common used strokes in the soccer domain (straight lines for passes and curly lines for dribbling). By focusing on the ball movement, we are able to reduce the amount of data visualized. Details on demand (e.g., movement of the involved players) are enabled and triggered by mouse hovering. The movement of involved players is visualized as small triangles pointing in direction of the movement. It is also possible to show (and abstract) specific player trajectories of interest permanently.

3.4. Simplification Layer - Trajectory Simplification

We employ several simplification techniques ranging from low to high simplification illustrated by Simplification in Figures 3 and 4. We specifically selected methods that use data points as control points for the simplification enabling outlier-aware abstraction as stated in our second requirement R2.

We use Catmull-Rom-Splines as an abstraction method resulting in trajectories very close to the observed ball movement. As an intermediate abstraction step we apply Smoothing via Iterative Averaging (SIA) [MH12]. SIA offers two parameters SI (Smoothing Iterations) and SS (Smoothing Sensitivity) that allow to control




Figure 4: Hovering interactions are provided in each layer: In the detail layer, hovering specific events will reveal tool-tips and player trajectories. In the simplification layer, the trajectory is highlighted and the events are shown. In the aggregation layer, the cluster is highlighted and different tool-tips reveal the involved players and event counts. Further, the sample points of the cluster representation are shown and the cluster member trajectories are visualized in the background including their convex hull.

the line smoothing. Increasing these parameters will increase the degree of simplification. The strongest simplification is performed by Bézier-curves. Rational Bézier-curves allow us to add adjustable weights (w) to control points determining the force of attraction for the shape simplification. Applied to our movements, we can decrease w to increase the degree of abstraction. In our soccer case, we linearly select ten control points. The first and last points are the most important information for analysts: The begin tells where and how the team captured the ball, while the end represents the outcome of a team-turn. Consequently, to retain these important features of the trajectory, we do not smooth the begin and end of the trajectory and keep their weightings fixed ($w = 200$) while the other control points will be decreased. Note, that other application domains might require different control point weightings.

We combine all the mentioned simplification techniques and parameterizations to a chain of increasing simplification levels exemplified in Figure 3. The specific ordering and parameterizations of the techniques have been developed and evaluated in close collaboration with domain experts. It is worth mentioning that we incorporated the SIA technique in a later stage based on the feedback that the transition between Catmull-Rom-Splines and Bézier-curves was not smooth enough. In our development, we first focused on simplifying the ball movements of team-turns. Although ball and players have different physics of movement with very different movement properties, it was possible to apply our implementations also to the player movements in a later stage.

3.5. Aggregation Layer - Summarizing Similar Movements

We aggregate trajectory segments according to similarity employing clustering (**Aggregation** in Figures 3 and 4). We chose algorithms that produce k clusters. This parameter perfectly fits to our intended LoA as stronger abstraction can be obtained by decreasing k . For our application domain, two different well-established algorithms are proposed: k-Means and k-Medoids with standard Euclidean distance measures. For k-Medoids we additionally implemented further state of the art distance measures in the area of trajectory clustering (Frechét, Hausdorff, and dynamic time warping - DTW). These clustering methods are easily explainable and per se define cluster prototypes. K-Means creates an artificial cluster prototype while k-Medoids determines a cluster member to this end. From

the trajectories we sample the same amount of points resulting in same-length feature vectors. For our application, ten sampling points are sufficiently describing usual observed team-turn segments. A cluster is rendered by a simplified cluster centroid based on the ten sampling points. Both methods are easily controllable by the number of desired clusters (k) and efficient enough to support real-time interactions required by **R4**. The granularity of decreasing k is dynamically obtained: according to our soccer experts the abstraction levels from $k = 1$ to $k = 5$ have to be always available. Of course, if there are for example only two trajectories to be aggregated, the maximum supported k will be 2. Additionally, we provide five additional clustering levels $calc_k(X)$ depending on the number of trajectory segments as defined in Equation 1.

$$calc_k(X) = \frac{\max(|S|, kMax)}{5} \cdot X \quad (1)$$

In Equation 1, X is the abstraction level ranging from 1 to 5, S is the set of trajectory segments to be clustered, and $kMax$ is the maximal number of clusters. In several expert discussions, we found a value of $kMax = 30$ suitable to the analysts needs, being a good compromise between details and overplotting. The decision whether to apply k-Means or k-Medoids (as well as the different distance functions) is dependent on the respective analysis task as shown later in the evaluation (Section 5.2). Please note that the trajectory segments are obtained based on a ball possession for team-turns. Further segments (for each moving object) can be obtained from manual time interval selections or in combination with feature-based segmentation techniques (provided by Janetzko et. al [JSS*14]).

3.6. Method Chain - Combining Generalization Techniques

We combine and chain the methods according to their ability to handle growing amounts of data and to their degree of abstraction as postulated in the third requirement **R3**. We combine the methods described above into a linear abstraction process as shown in Figure 3 controllable by mouse wheel. The crucial point of our sequence of different simplification and aggregation techniques is the proper setting of parameters. We iterated several times over the proper order and parameters with subject matter experts in soccer analysis. Consequently, we are able to visualize smooth model transitions and let the analyst interactively explore the *Abstraction Space*. It is further possible to generate an **Overview** of all LoA parameterizations

(e.g., shown in Figure 1) that lets the analyst compare and choose an adequate visualization. At this point we want to refer the reader to the video in the supplemental materials illustrating our approach.

We visualize the different abstraction layers in different ways conveying that we do not show the raw movement as required by R1. In the **Detail** layer, the raw team-turn data is visualized using the ball trajectory with specific events and movement types (pass, dribbling) where we use the color of the team in ball possession. Raw player trajectories are visualized as little triangles pointing into the movement direction. In the **Simplification** layer, the trajectory thickness is mapped to the LoA and we use lighter colors and textures for simplified player trajectories. In the **Aggregation** layer, the trajectory/cluster thickness is mapped to the number of cluster members, which is also shown by a label. We again apply different textures to distinguish players from team-turns (players with lighter color and texture, team-turns with darker color and texture). We treat the team-turns as well as the players of the particular teams as separate aggregations (e.g., it is not intended to aggregate player trajectories with team-turns or team-turns of opposing teams).

4. Analyzing Soccer Movement

This section describes our visual abstraction implementation integrated into an existing soccer analysis system described by Janetzko et al. [JSS*14] and showcases how it supports soccer experts in their analysis. The default user interface comprises a soccer pitch with the trajectory (layer)-rendering, an option panel to define configurations, and a timeline view that shows selected time intervals. For our visual abstraction work, we added the abstraction layers with its configurations in the option panels and a snapshot view that allows the analyst to define bookmarks. Additionally, we enable the analyst to create specific filters using the Move-Filter (see Section 4.1). Finally, it is possible to annotate and add visualizations to a note-taking interface developed by Sacha et al. [SBFK16], which is also able to capture user interactions.

4.1. Interactions

We provide the analyst with several interactions and configurations to enable interactive exploration of soccer movements.

Object & Time Selections: The analyst is able to select multiple time intervals in the timeline. The option panel lets the analyst select the moving objects to be visualized (turns of team A or B, and players). We also developed a filter component enabling the analyst to define conditions to select specific team-turns of interest. It is, for example, possible to add players that have to be involved in a team-turn or specific events that have to occur (e.g., crosses). Team-turn clusters can be selected (or removed) or removed directly on the soccer-pitch if the **Aggregation** layer is shown.

Abstraction Layer Interactions: It is possible to navigate within the abstraction space (adapting the LoA-parameter) by simple mouse scrolling. It is further possible to bring up an **Abstraction Overview** - overlay that computes and visualizes all steps of the global LoA-parameter. This enables the analyst to compare and select different levels of abstractions (see Figure 1). Detail on Demand interactions are provided in each layer: In the **Detail** layer, the analyst is able

to hover on events and players to reveal a tool-tip and the player trajectory (Figure 4-left). In the **Simplification** layer, hovering a trajectory will reveal the important events (stations) of a team-turn (Figure 4-center). In the **Aggregation** layer, hovering a cluster will reveal all the cluster members and the convex cluster hull in the background (Figure 4-right). This supports the analyst in evaluating the cluster quality and to fine-tune different clustering configurations. Additional tool-tips show the frequency of events and involved players. The analyst can choose the clustering algorithm (k-Means or k-Medoids) and the used distance metrics for k-Medoids (Fréchet, Hausdorff, Euclidean, DTW) in the option panel.

Meta Interactions: Analytic provenance interactions let the analyst save and add visualizations to the snapshots bar and to the note-taking interface. The analyst can re-arrange all the views and open other analysis capabilities of the soccer analysis system.

4.2. Soccer Analysis Tasks

The presented system supports several different kinds of analysis. On the one hand, it is possible to focus on different moving objects (team-turns, players). On the other hand, the analysis can focus on continuous (temporal sequence is preserved and important) or discrete movements (spatial similarity is more important). We observed for each abstraction layer advantages and drawbacks. The **Detail** layer is good for fine-grained and exact analysis as no distortion is applied but limited to the amount of data. The **Simplification** layer works well when the analyst wants to follow a continuous movement of the ball or players and simplifies complex “branched” movements. However, the events and movement points get distorted and do not refer to the exact positions anymore. Finally, the **Aggregation** layer works well in identifying similar movements and to generate overviews of discrete movements, but obviously, a lot of details are lost. Putting all these methods together enables us to overcome a lot of these drawbacks and to satisfy many analysis tasks. In this study, we identified four different kinds of investigations that can be performed with our implemented system.

Exploratory Team-Turn Analysis: The system visualizes all the team-turns of an entire match allowing the analyst to iteratively select and explore clusters of interest in order to drill down into specific turns. Figure 5 shows all the turns of a soccer match for the red team A (A1) and the blue team B (B1). We can clearly spot differences. While the red team has more turns through the middle, the blue team tries to attack via the wings (a typical V-shape is visible in Figure 5-B1). The analyst is now interested in the upper red cluster, because it ends at the center/midfield line of the soccer pitch representing “unsuccessful” turns where the red team loses the ball too early. By selecting this cluster the analyst is provided with another abstraction of the selected turns (Figure 5-A2) and in a subsequent selection of the upper red cluster with a more detailed representation (Figure 5-A3). Hovering the trajectories revealed that player 18 of the red team was involved in most of these turns losing the ball too early and player 3 of the blue team was able to capture the ball. A similar approach of interactively selecting clusters allowed us to identify and investigate all right wing attacks that are initiated by a long ball of the blue goal keeper (Figure 5-B1, B2, and B3).

Analyzing Specific Team-Turns: The move filter lets analysts




Figure 5: Exploratory Team-Turn Analysis: All turns of the teams are shown in A1 and B1 revealing different tactical attacking styles. The system lets the analyst explore and select clusters of interest interactively drilling down into particular team-turns of interest. In A1-A2-A3, the analyst investigated “unsuccessful” turns. In B1-B2-B3 the analyst investigated turns initiated by a long ball of the goal keeper.




Figure 6: Analyzing Player Segments: Player trajectory segments are created based on movement speed, acceleration, and straightness (A). All the sprint segments are filtered (B). The resulting sprint segments of this player can be abstracted using our approach (C). The aggregated sprint view allows us to discover the main sprint areas and a few outliers (D).

filter the dataset for specific team-turns. Figure 1-left shows all attacks with a “shot on target” event. We can identify two different tactics. The blue team shows long turns with a zic-zac pattern (a tactic with short passes and long ball possessions), whereas the red team shows relatively short turns with the ball capturing near the opponents goal and early finish (also known as “pressing”).

Analyzing Player Segments: Another analysis task is to analyze the movement segments of a particular player. Player trajectory segments can be obtained e.g., in combination with existing feature-based approaches implemented by Janetzko et al. [JSS* 14]. We used their implementation to segment the movement of a midfielder player based on speed, straightness, and acceleration (Figure 6-A) and filtered the segments for high speed and acceleration (Figure 6-B) in order to analyze sprints. We can apply our techniques to these trajectories (Figure 6-C) and investigate an aggregated representation of the players sprints (Figure 6-D) which are mainly in the opposing teams half of the soccer pitch on the right wing with a few outliers. It further reveals that this player always attempts to reach the left full-back when sprinting. This player might be advised to occasionally take the ball outside the full-back as predictable play is easy to defend against. Note, that the segmentation and filtering have to be applied according to the analysis task at hand. We could also apply our abstractions to all the trajectory segments or filter for different movements (e.g., stop moments with a high negative acceleration). However, such segmentations are essential preprocessings for our abstractions to support meaningful analysis tasks.

Analyzing Collective Movements: Finally, we can apply the same technique to identify and investigate collective movements of several players. To do so, we selected a short time span but all player trajectories and apply our trajectory abstraction technique (Figure 1-

right). The visualization allows analysts to identify different groups of players. Obviously, the two goal keepers are shown as separate clusters as their movement is different from the other players.

5. Evaluations

We conducted two qualitative user studies to validate our design and collect feedback about our implementations.

Apparatus: The studies were conducted in a lab setting using a 24 inch screen to show video sequences or the soccer analysis system. For the first study, the participants were provided with a sheet of paper for each task, showing an empty soccer pitch or a printout with trajectories that should be abstracted. The participants were provided with a pen to draw on the sheets. For the second study, the only input device was a common computer mouse. The participants were seated approximately 50 cm away from the screen. The experimenter was present during the study for answering questions, introducing the study procedure, and collecting feedback.

5.1. Visual Design

We conducted a qualitative user study with both domain experts to validate our abstraction layers and if the visual results meet their expectations. Expert E1 was involved in the design process of our visual abstraction technique, expert E2 was just involved during the evaluation phase. The general idea of this experiment was to let the experts draw the abstractions on a sheet of paper and to compare the results with our computed visualizations afterwards.

Tasks and Procedure: The participants had to produce drawings that are comparable to each of our abstraction layers. In the first




Figure 7: The participants had to perform the visual abstraction tasks manually on a sheet of paper. A few examples of the participants drawings are shown on the left and the computed visualizations are shown on the right (one example for each layer).

task (T1), both participants were shown six specific team-turns from a video recording. We showed them a video sequence and paused the video when the participants requested it. For each sequence the experts had to draw the team-turn in an empty soccer pitch on a piece of paper. The experts were asked to focus on the sequential events to draw an abstract representation of the team-turn. We told them that they would not need to provide an exact drawing. The rationale behind this task was to let the experts produce team turn drawings that can be used to evaluate our **Detail** and **Simplification** layer. In the second task (T2), both experts were provided with several printouts of a soccer pitch containing many trajectories representing all turns of a single team. The participants had to draw the general flow of movement on the sheet of paper on top of the trajectories. The experts got no further restriction or hints on performing the task. The aim of this task was to let the experts produce trajectory summarizations that can be used to evaluate our **Aggregation** layer.

Results: The generated drawings are provided in the supplemental material and three representative results for each abstraction layer are shown in Figure 7. In T1, expert E2 (who was not involved in the design process) produced drawings that are very similar to our **Detail** layer basically connecting annotated circles of important events. Note, that this layer has been designed together with expert E1 who produced simplified representations of the team-turns that come very close to our **Simplification** layer. In T2, both experts produced results similar to the **Aggregation** layer by creating thick arrows above the main movements that can be seen on the soccer pitch. Both experts used *thickness* of the arrows to communicate aggregation of the underlying trajectories. This study helped us to demonstrate our abstraction and visualization concepts that are in line with the mental models of the soccer experts.

5.2. Comparing Clustering Results

We implemented two different clustering methods (k-Means and k-Medoids) and specifically for k-Medoids, we also implemented




Figure 8: Two clustering representations for the same trajectories. The k-Means method was considered useful to analyze an aggregated movement direction, whereas k-Medoids was considered more useful to identify soccer-specific patterns.

four different distance functions. We conducted another deeper qualitative study with one of the experts (E1) to evaluate the distance functions and to collect feedback about the clustering techniques.

Tasks and Procedure: We chose eight different time and team selections of a soccer match M1: Whole match, first/second half, and team-turns with crosses (all for team A and B). We first showed the **Simplification** layer representation to illustrate all the trajectories that have to be clustered. Subsequently, the participant had to compare and provide feedback about the different distance functions (T1) and the preferred clustering technique (T2) within the **Aggregation** layer. The first task (T1) was to compare distance functions with the k-Medoids method. Expert E1 had to configure the clustering and compare the results of all the distance functions. He was further asked to provide specific feedback and decide for a preferred distance function. In the second task (T2), the expert was asked to compare the k-Medoids with the k-Means clusterings. The aim of these tasks was to collect feedback about the used methods (metric, clustering) for different match situations.

Results: In T1, the expert preferred for six situations DTW and for two situations the Euclidean distance. The Hausdorff distance did not provide good results because it was not possible to spot the desired tactical patterns in the visualizations (*V-*, or *cradle*-shapes). The DTW turned out to be the expert's favorite as it was possible to reveal the desired patterns for most of the situations. However, the expert also emphasized that he likes the ability to compare and explore the different distance functions. Therefore, we provide the DTW distance metric as a the default configuration but retain the ability to change it. In T2, the expert reported advantages and disadvantages for both clustering methods depending on the analysis task and selected situation. The k-Means method was useful to estimate an aggregated direction of the movements, whereas k-Medoids was considered more useful to detect and analyze soccer-specific patterns. For example, Figure 8 visualizes the cluster representation (foreground) and the aggregated trajectories (background) for k-Means (a) and k-Medoids (b). The general movement direction is shown with k-Means while the *cradle*-shape pattern is preserved by k-Medoids. As a result we provide the analyst the ability to change the clustering method on demand.

6. Recommending the Level of Abstraction

We observed that setting the LoA parameter (by scrolling the mouse wheel) is useful for understanding the abstractions, however, it also

turned out to be time consuming in exploratory analysis tasks. Additionally, further parameters (e.g., clustering, distance function) have to be chosen by the analyst. To support such parameterization tasks we started to build a recommender system that provides the analyst with automatic pre-configurations. We further observed that such a goal is user dependent (they may prefer more/less abstraction).

Measures and Recommender Training: A domain expert provides explicit feedback about specific useful visualizations by pressing a “learn from current view” button. This feedback will create a record for the current visualization in an (individualized) recommender database. We defined several measures to capture data and visual characteristics in combination with configuration information. Firstly, we measure the sum of selected time interval lengths (data). Secondly, we measure the amount of trajectory crossings, the crossed surface area, and the sum of trajectory surfaces (visual). Finally, we capture the used abstraction layer with its local and global abstraction parameter (configuration). The system offers a 50 step training-process covering randomized and specific situations (randomized time intervals as well as match specific team turns). An example csv file for such captured data is provided in the supplemental materials. During the actual analysis process, further explicit feedback will individualize and adapt the recommender system. Instead of performing the personalized initial training, it is also possible to choose a “default” (pre-trained) recommendation database, which was created during our evaluation sessions.

Providing Recommendations: We analyzed a captured training dataset and found a significant positive Pearson correlation between the LoA and the selected duration ($r(48) = .68, p = .00$). This simply means the more time and data is selected, the more abstraction is needed. As a result of this statistical test we select the time intervals as a primary feature to predict the global LoA (which is composed of the abstraction layer and the local abstraction parameter). We decided for a simple and efficient *lazy learning* classifier [AMS97] as they are able to handle and solve multiple and changing problems. Furthermore, additional training data shall be considered during runtime. The *k-Nearest-Neighbor* (kNN) algorithm satisfies these requirements and has been implemented to determine the abstraction layer and its local abstraction parameter. Note that we cannot simply predict the global LoA because its range is dynamic (dependent on the number of trajectories, see Section 3). The process is realized as follows: First, all the learned records in the training database are sorted according to the time selection length (Euclidean distance) to identify the k closest neighbors. The number of considered neighbors is determined by Equation 2, where T is the total amount of training data, 5 an upper and 1 a lower bound.

$$f(T) = \max \left\{ \min \left\{ \left\lfloor \frac{|T|}{10}, 5 \right\rfloor, 1 \right\}, 1 \right\} \quad (2)$$

The kNN-records are used to determine the most frequent abstraction layer and uses the remaining records to calculate the mean value of the local abstraction parameters. The results are then used to pre-configure the LoA.

Accuracy: We conducted a quantitative user study with both soccer experts to evaluate the usefulness and precision of our recommender system. The basic idea of this experiment was to compare the recommended LoA with a user defined LoA and to measure the difference as an error e . The participants were presented with a




Figure 9: The error rates between the manually chosen and recommended LoA parameters for both participants and two evaluation cases (personalized vs. “default” training).

specific time selection and had to set the desired LoA parameter by scrolling the mouse-wheel and to save their decision using the “next” button. In the first “personalized” evaluation (T1), our goal was to measure the accuracy of a previously trained classifier for two different matches. The participants had to train the recommender using the 50 test cases for a specific match M1. In a subsequent iteration, the participants had to perform another 50 steps for a different match M2. For every step, we calculate the difference as error e and added a penalty of 10 ($e = e + 10$) when two different layers were used as they support different analysis goals (e.g., **Aggregation** instead of **Simplification** layer). In the second “default” evaluation (T2), we tested the accuracy of our “default” classifier that is not trained by the participant. Therefore, we created a “default” database trained by several users for matches M1 and M2 with a size of 400 records. We then evaluated match M3 with 50 test steps for each participant and calculated e similar as in the previous case.

Results: The results for both participants and study tasks are shown in Figure 9. For the “personalized” cases (T1), we found an average error of $e = 1.43$ for expert E1 and $e = 1.45$ for expert E2. Hence, we are able to state that even our simple kNN approach is able to provide a “good” pre-configuration and is in average not worse than 1.5 steps (within the 21-step abstraction space). However, in Figure 9 we can also spot three outliers with very high values for the cases when the the wrong abstraction layer was recommended. For the “default” cases (T2), we found an average error of $e = 1.92$ for expert E1 and $e = 2.06$ for expert E2. In this case, the recommender system performed not as good as in the personalized case. However, the error is in average still as close as 2 steps to the user preference for both participants. Similar to the previous case, we also spot three outliers greater 10 where the wrong abstraction layer was chosen (out of 100). This study showed that good recommendations can be provided even without a personalized training before usage. However, the personalized still outperforms the default case.

Next Steps: We are well aware that more sophisticated recommender systems and approaches do exist and that we just provide a simple and early proof of concept to illustrate its usefulness in our setting. Our results reveal a multitude of interesting research areas. Firstly, we will leverage more measures that can be obtained

from a visualization. Examples include data-related metrics, such as the number of moving objects, but also visual measures. Note, that we can pre-compute visual measures for different parameterizations (e.g., the used clustering technique) in order to optimize the visual configuration beyond the LoA parameter. Secondly, it will be very interesting to train separate recommender systems for different analysis tasks. We envision that we could even leverage such task-databases to detect user intent (determining the “closest” task) and adapt the visualizations accordingly. Finally, our approach learns from explicit user feedback. Another interesting research direction is to obtain the training data from implicit user feedback (e.g., once a visualization is bookmarked or annotated).

7. Discussion

We developed the system together with domain experts that iteratively helped us to adapt our implementations and refine our ideas. Our design process and evaluations focused on the visual abstraction of team-turns and a domain requirement was to strictly separate between the opposing teams and players. It will be interesting to evaluate and fine-tune our system with focus on player movements and apply our computations across teams.

By reflecting our design process we further noticed that it was “easier” for users to map the abstraction parameters to a LoA for the **Simplification** layer techniques than for the **Aggregation** (the analyst has to set more parameters). That is also why we had to conduct evaluations for the clustering. However, we also did not implement and evaluate all the visual design alternatives. We mapped, for example, the LoA intuitively to trajectory thickness (based on **R1**), however, thick arrows introduce additional overplotting. It would be interesting to parameterize thickness and investigate how much is needed and appropriate. Similarly, one could focus on the used textures and colors. Another configuration for the cluster rendering could be to show the means and medoids concurrently (if k is small).

Furthermore, we designed the system with two domain experts and derived some methodological choices that worked well in our setting. However, further requirements might become relevant for a broader set of users and more specific analysis tasks. For example, the number of sampled control points and their weighting (for both the **Simplification** and **Aggregation** layer) could be parameterized and further investigated. We note our approach supports an adaptive level of simplification per given situation, and that all trajectories within that situation are simplified using the same parameters. In future work, we may also consider adaptive simplification on a *per-trajectory* level. An idea to do so is to classify the type of trajectory, e.g., pertaining to different player moves or roles, and then to adapt the number of sampled control points and weightings to this class. Another way to improve the abstraction techniques is to take semantics or additional features into account. For example, in the **Simplification** layer, we can select and weight control points based on the events that are contained in the trajectory. Similarly, for the **Aggregation** layer, we can include further meta-data (e.g., events) and features (e.g., speed) into the similarity calculation.

Finally, we want to emphasize that we selected methods that naturally matched our domain requirements. In other application domains it could be interesting to investigate other abstraction techniques (e.g., grid-based or edge-bundling). We tried to complement

our visualizations with heatmaps or concav/convex-hulls but we had to abandon these investigations to keep our work focused. It will be interesting future work to generalize our approach and to enrich it with further abstraction techniques.

8. Conclusion

This paper presents a novel approach to define and combine visual abstraction techniques interactively to overcome over-plotting and clutter. We specifically focused on the domain of soccer movement data and designed a system that supports soccer analysts in solving a variety of analysis tasks. The interactive navigation and smooth model transitions allow the analyst to track and understand the underlying abstraction computations. A recommender system has been added as a proof of concept to automatically pre-configure the LoA based on the amount of data to be shown. In the future, it will be interesting to improve the obtained recommendations for a variety of tasks. Furthermore, we will apply this approach to other movement data types (e.g., animal movements) and adapt (or replace) the abstraction techniques with the ultimate goal to come up with a reusable and extensible abstraction framework.

Acknowledgments

We wish to thank Andreas Riederer for his valuable feedback and discussions. The soccer data used in this publication was generously provided by Prozone/Stats. This research was supported by the German Research Foundation DFG within Priority Research Program 1894 “Volunteered Geographic Information: Interpretation, Visualisation and Social Computing” (DFG SPP VGI).

References

- [AA11] ANDRIENKO N., ANDRIENKO G.: Spatial generalization and aggregation of massive movement data. *IEEE Trans. on Visualization and Computer Graphics* 17, 2 (2011), 205–219. doi:10.1109/TVCG.2010.44. 3
- [AAB*13a] ANDRIENKO G. L., ANDRIENKO N. V., BAK P., KEIM D. A., WROBEL S.: *Visual Analytics of Movement*. Springer, 2013. doi:10.1007/978-3-642-37583-5. 2
- [AAB*13b] ANDRIENKO N., ANDRIENKO G., BARRETT L., DOSTIE M., HENZI P.: Space transformation for understanding group movement. *IEEE Trans. on Visualization and Computer Graphics* 19, 12 (2013), 2169–2178. 3
- [AAB*16] ANDRIENKO G. L., ANDRIENKO N. V., BUDZIAK G., VON LANDESBERGER T., WEBER H.: Coordinate transformations for characterization and cluster analysis of spatial configurations in football. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD* (2016), pp. 27–31. doi:10.1007/978-3-319-46131-1_6. 3
- [AAR*09] ANDRIENKO G. L., ANDRIENKO N. V., RINZIVILLO S., NANNI M., PEDRESCHI D., GIANNOTTI F.: Interactive visual clustering of large collections of trajectories. In *IEEE Conf. on Visual Analytics in Science and Technology (VAST)* (2009), pp. 3–10. doi:10.1109/VAST.2009.5332584. 2
- [ADKZ*16] AL-DOHUKI S., KAMW F., ZHAO Y., MA C., WU Y., YANG J., YE X., WANG F., LI X., CHEN W.: Semantictraj: A new approach to interacting with massive taxi trajectories. *IEEE Trans. on Visualization and Computer Graphics* PP, 99 (2016), 1–1. doi:10.1109/TVCG.2016.2598416. 2

- [AMS97] ATKESON C. G., MOORE A. W., SCHAAL S.: Locally weighted learning for control. *Artif. Intell. Rev.* 11, 1-5 (1997), 75–113. doi:10.1023/A:1006511328852. 9
- [BKSS14] BEHRISCH M., KORKMAZ F., SHAO L., SCHRECK T.: Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. *IEEE Conf. on Visual Analytics in Science and Technology (VAST)* (2014), 43–52. doi:10.1109/VAST.2014.7042480. 3
- [BYRN11] BAEZA-YATES R., RIBEIRO-NETO B.: *Modern Information Retrieval*. ACM Press. ACM Press, 2011. Second Edition. 3
- [DBC*15] DEMŠAR U., BUCHIN K., CAGNACCI F., SAFI K., SPECKMANN B., VAN DE WEGHE N., WEISKOPF D., WEIBEL R.: Analysis and visualisation of movement: an interdisciplinary review. *Movement Ecology* 3, 1 (2015), 5. doi:10.1186/s40462-015-0032-y. 1, 2
- [GW14] GUDMUNDSSON J., WOLLE T.: Football analysis using spatio-temporal tools. *Computers, Environment and Urban Systems* 47 (2014), 16–27. doi:10.1016/j.compenurbsys.2013.09.004. 2
- [HD12] HEALEY C. G., DENNIS B. M.: Interest driven navigation in visualization. *IEEE Trans. Vis. Comput. Graph.* 18, 10 (2012), 1744–1756. 3
- [HKP11] HAN J., KAMBER M., PEI J.: *Data Mining: Concepts and Techniques*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2011. 2
- [HTC09] HURTER C., TISSOIRES B., CONVERSY S.: Fromdady: Spreading aircraft trajectories across views to support iterative queries. *IEEE Trans. on Visualization and Computer Graphics* 15, 6 (2009), 1017–1024. doi:10.1109/TVCG.2009.145. 2
- [IWSK07] IVANOV Y., WREN C., SOROKIN A., KAUR I.: Visualizing the history of living spaces. *IEEE Trans. on Visualization and Computer Graphics* 13, 6 (2007), 1153–1160. doi:10.1109/TVCG.2007.70621. 2
- [JSS*14] JANETZKO H., SACHA D., STEIN M., SCHRECK T., KEIM D. A., DEUSSEN O.: Feature-driven visual analytics of soccer data. In *IEEE Conf. on Visual Analytics in Science and Technology (VAST)* (2014), pp. 13–22. doi:10.1109/VAST.2014.7042477. 2, 3, 5, 6, 7
- [JZFF10] JANNACH D., ZANKER M., FELFERNIG A., FRIEDRICH G.: *Recommender Systems: An Introduction*, 1st ed. Cambridge University Press, New York, NY, USA, 2010. 3
- [LIW05] LAUBE P., IMFELD S., WEIBEL R.: Discovering relative motion patterns in groups of moving point objects. *International Journal of Geographical Information Science* 19, 6 (2005), 639–668. doi:10.1080/13658810500105572. 3
- [MH12] MANSOURYAR M., HEDAYATI A.: Smoothing via iterative averaging (sia) a basic technique for line smoothing. *International Journal of Computer and Electrical Engineering* 4, 3 (2012), 307. 4
- [MRS07] MACKANESS W. A., RUAS A., SARJAKOSKI T. (Eds.): *Generalisation of Geographic Information: Cartographic Modelling and Applications*. Elsevier, 2007. 1, 3
- [PKM*07] PELEKIS N., KOPANAKIS I., MARKETOS G., NTOUTSI I., ANDRIENKO G., THEODORIDIS Y.: Similarity search in trajectory databases. In *Proc. of the 14th International Symposium on Temporal Representation and Reasoning* (2007), IEEE Computer Society, pp. 129–140. doi:10.1109/TIME.2007.59. 2
- [PVF13] PERIN C., VUILLEMOT R., FEKETE J.: Soccerstories: A kick-off for visual soccer analysis. *IEEE Trans. on Visualization and Computer Graphics* 19, 12 (2013), 2506–2515. doi:10.1109/TVCG.2013.192. 2
- [SBFK16] SACHA D., BOESECKE I., FUCHS J., KEIM D. A.: Analytic Behavior and Trust Building in Visual Analytics. In *EuroVis 2016 - Short Papers* (2016), Bertini E., Elmqvist N., Wischgoll T. (Eds.), The Eurographics Association. doi:10.2312/eurovisshort.20161176. 6
- [SBTK09] SCHRECK T., BERNARD J., TEKUŠOVÁ T., KOHLHAMMER J.: Visual cluster analysis of trajectory data with interactive Kohonen maps. *Palgrave MacMillan Information Visualization* 8 (2009), 14–29. doi:10.1057/ivs.2008.29. 2
- [SJB*16] STEIN M., JANETZKO H., BREITKREUTZ T., SEEBACHER D., SCHRECK T., GROSSNIKLAUS M., COUZIN I. D., KEIM D. A.: Director's cut: Analysis and annotation of soccer matches. *IEEE Computer Graphics and Applications* 36, 5 (2016), 50–60. doi:10.1109/MCG.2016.102. 2
- [SMM12] SEDLMAIR M., MEYER M. D., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. on Visualization and Computer Graphics* 18, 12 (2012), 2431–2440. doi:10.1109/TVCG.2012.213. 3
- [SSN*16] SHAO L., SACHA D., NELDNER B., STEIN M., SCHRECK T.: Visual-interactive search for soccer trajectories to identify interesting game situations. In *Visualization and Data Analysis VDA* (2016), pp. 1–10. 2
- [VLBSF14] VON LANDESBERGER T., BREMM S., SCHRECK T., FELLNER D.: Feature-based automatic identification of interesting data segments in group movement data. *Sage Information Visualization* 13, 3 (2014), 190–212. doi:10.1177/1473871613477851. 2
- [VvdWvW11] WILLEMS N., VAN DE WETERING H., VAN WIJK J. J.: Evaluation of the visibility of vessel movement features in trajectory visualizations. In *Computer Graphics Forum* (2011), vol. 30, pp. 801–810. doi:10.1111/j.1467-8659.2011.01929.x. 2