

MotionExplorer: Exploratory Search in Human Motion Capture Data Based on Hierarchical Aggregation

Jürgen Bernard, Nils Wilhelm, Björn Krüger, Thorsten May, Tobias Schreck, and Jörn Kohlhammer

Abstract—We present MotionExplorer, an exploratory search and analysis system for sequences of human motion in large motion capture data collections. This special type of multivariate time series data is relevant in many research fields including medicine, sports and animation. Key tasks in working with motion data include analysis of motion states and transitions, and synthesis of motion vectors by interpolation and combination. In the practice of research and application of human motion data, challenges exist in providing visual summaries and drill-down functionality for handling large motion data collections. We find that this domain can benefit from appropriate visual retrieval and analysis support to handle these tasks in presence of large motion data. To address this need, we developed MotionExplorer together with domain experts as an exploratory search system based on interactive aggregation and visualization of motion states as a basis for data navigation, exploration, and search. Based on an overview-first type visualization, users are able to search for interesting sub-sequences of motion based on a query-by-example metaphor, and explore search results by details on demand. We developed MotionExplorer in close collaboration with the targeted users who are researchers working on human motion synthesis and analysis, including a summative field study. Additionally, we conducted a laboratory design study to substantially improve MotionExplorer towards an intuitive, usable and robust design. MotionExplorer enables the search in human motion capture data with only a few mouse clicks. The researchers unanimously confirm that the system can efficiently support their work.

Index Terms—Visual analytics, exploratory search, multivariate time series, motion capture data, data aggregation, cluster glyph

1 INTRODUCTION

Human motion capture data can be regarded an instance of multivariate time series data. It is applied in various research fields such as medicine, sports and animation. It may be acquired from human actors labeled with detectable markers using video tracking or obtained synthetically by simulation. Resulting human motion time series are stored in large data collections where they are potentially available for analysis and reusage. Motion capture data typically consists of frames that contain a high-dimensional vector representation of a human *pose*. For large data collections, manual analysis is very time-consuming and the retrieval of a motion sequence of interest is often done by manually screening the set of motion sequences. We target the user group of researchers in human motion synthesis, who are called *domain experts* in the following. They focus on an effective re-use of existing motion capture data to synthesize new human motions. Thus, an overview of existing data is important to understand the potential building blocks of existing motion. Our users distinguish motion on the macro and micro level. On the macro level, different types of motions ‘from A to B’ are denoted (e.g., activities such as cross-country-skiing in the classical or the skating technique). On the micro level, different style variations of the same motion type are distinguished (e.g., different speeds, dynamics, or skill levels of the same motion). Motion synthesis typically involves identifying first the motion type of interest and then selecting an appropriate style to use or adapt.

We identified the following problems we aim at supporting with MotionExplorer: 1.) As motion data collections are typically large, an overview is needed. To our knowledge, visual access to large collections of motion capture data has not become a common application field for visual analytics yet. In addition, the spatio-temporal variations in existing human motion make this complex data type inappropriate for many existing tools. Due to a lack of specialized tools, the domain experts report on working through their data collections manually in many cases. 2.) Visual-interactive functionality to efficiently obtain meaningful data subsets is scarce. While algorithmic search methods for indexing and matching of motion data exist [27], to date in practice, retrieval in motion data often still relies on textual metadata and not directly on the content. This also includes a lack of visual query formulation interfaces in practice. 3.) Identifying interesting sequences from the retrieved search results remains challenging. While in large motion databases, style variations typically exist, these are hard to distinguish automatically but often require the subjective interpretation by the expert. To that end, again, visual representations are expected to be useful. As yet however, they are not often leveraged in practice.

To help domain experts with these problems, we developed MotionExplorer as an ESS type system (Exploratory Search System, [43]) supporting large human motion capture data collections. The overall system is shown in Figure 1. The development was guided by a design study and the system provides overview first (via clustering) and detail on demand (via search and drill-down) functionality at its core:

- Jürgen Bernard is with Fraunhofer Institute for Computer Graphics Research Darmstadt. E-mail: juergen.bernard@igd.fraunhofer.de.
- Nils Wilhelm is with Fraunhofer Institute for Computer Graphics Research Darmstadt. E-mail: nils.wilhelm@igd.fraunhofer.de.
- Björn Krüger is with Institute of Computer Science, Universität Bonn. E-mail: kruegerb@cs.uni-bonn.de.
- Thorsten May is with Fraunhofer Institute for Computer Graphics Research Darmstadt. E-mail: thorsten.may@igd.fraunhofer.de.
- Tobias Schreck is with Data Analysis and Visualization Group, Universität Konstanz. E-mail: tobias.schreck@uni-konstanz.de.
- Jörn Kohlhammer is with Fraunhofer Institute for Computer Graphics Research Darmstadt. E-mail: joern.kohlhammer@igd.fraunhofer.de.

Manuscript received 31 March 2013; accepted 1 August 2013; posted online 13 October 2013; mailed on 4 October 2013.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

- An interactive dendrogram visualization supports exploration of a hierarchical clustering and the steering of the aggregation level of the data set. Human motion between pose aggregates can be explored in a node-link diagram where poses define the nodes and edges represent motion. The views are interactively linked, and enable drill-down exploration.
- Domain experts are able to visually search for motion sequences. The system provides a detail-on-demand visualization for human pose aggregates, which also serves as a query-by-example pool. Search results of human motion sequences can be interactively explored in detail for style variations.

We developed MotionExplorer in collaboration with domain experts, who provided the data set and defined the analytical challenges in their research workflow. With the domain experts, we acquired

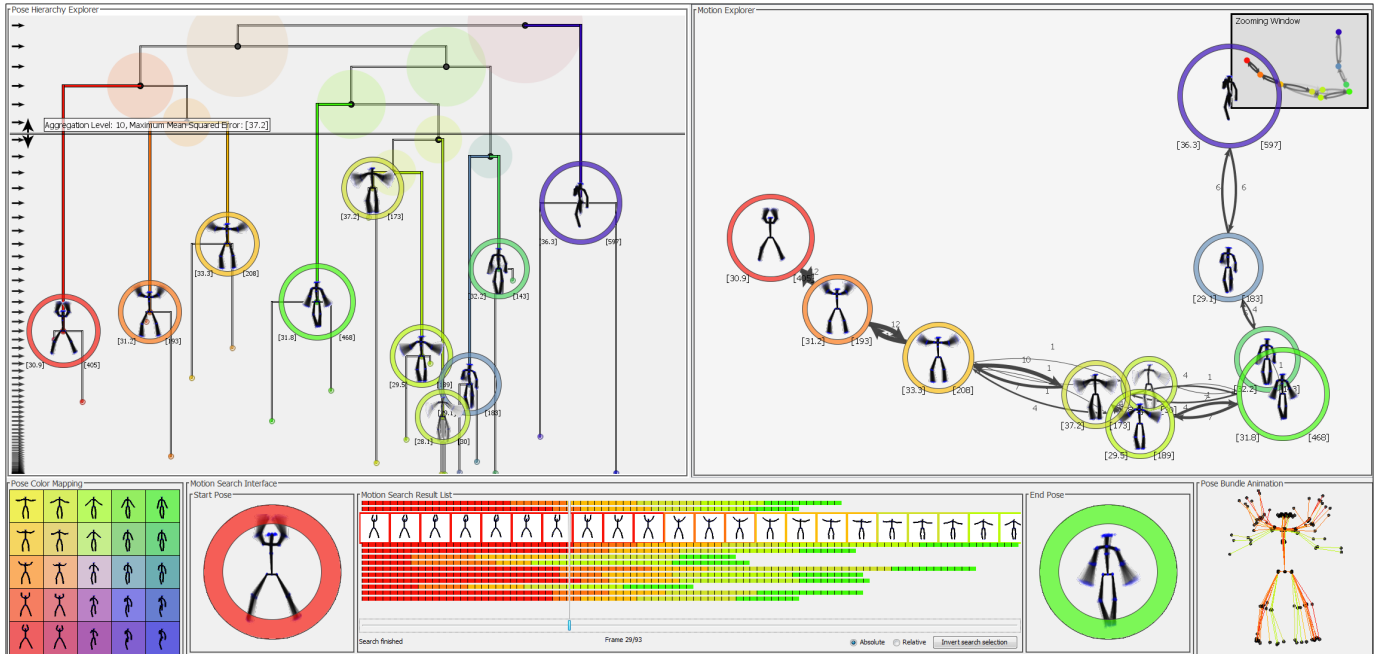


Fig. 1. MotionExplorer enables the exploratory search in human motion capture data. The pose hierarchy explorer (upper left) allows adjustment of the aggregation level of a hierarchical clustering and filtering of human poses. The motion explorer (upper right) shows human pose and movement aggregates. The search interface (below) allows for a visual query definition and an interactive search result exploration to identify style variations.

high-level requirements at an early stage of the design. The collaborations in this two-year design study are shown in Figure 2.

The remainder of this paper is structured as follows. Section 2 reports on the targeted user domain and the collaborative requirement definition process. We then present related work in Section 3. Data abstractions are discussed in Section 4. The implementation results of the visual encodings and interactive capabilities are described in Section 5. Additionally, we report on how the outcome of a formative laboratory study with non-experts helped us successively improve the system in the design process. Finally, we conducted a summative field study with the domain experts presented in Section 6. Section 7 discusses our results, and Section 8 concludes.

2 BACKGROUND

We conducted formative field observations and interviews with the domain experts to clarify the user needs on human motion synthesis, our targeted domain.

2.1 Analytical Tasks in Human Motion Synthesis

The motion synthesis products of our domain experts are often targeted towards computer animation. To this end, they need large numbers of camera takes to create human motion sequences. In the field of computer animation, the creation of realistic motion sequences by hand is a time-consuming task. For this reason, motion capturing of human actors has become a standard technique in computer animation. The motion sequences performed by the actor are recorded by a multi-camera system. The 3D positions of the markers can be reconstructed from the two-dimensional images via triangulation for each frame.

Reusing previously recorded data for motion synthesis has proven to be successful. These techniques are commonly based on the two following principles: *Parametric synthesis* generates motion spaces from a given set of example motion sequences. Allowing interpolation, the results can be adapted to specific spatial or temporal constraints [22]. *Concatenation of motion segments* links together short motion sequences at appropriate locations to create a longer sequence. Using the latter approach, motion graphs [23] have become a standard technique in data-driven computer animation. A motion graph is a model that has a prominent role in many applications as an automated method of generating character animation, especially since motion graphs have a range of benefits for interactive applications. Tech-

nically, all possible transitions between given motion sequences need to be computed in advance. The transitions and the corresponding motion sequences are subsequently stored in a graph structure. Consequently, to create a new motion, sequences are concatenated to a path through nodes of the graph. An overview of the most popular routines and approaches in the field of example-based motion synthesis is found in the state-of-the-art report of Pejša and Pandžić [30].

2.2 Identification of Current Problems / Limitations of Current Tools

The main drawback of data-driven motion synthesis approaches, such as motion graphs, is that the domain experts have no direct control over every single step that is performed during the synthesis.

Firstly, there is no overview of the underlying motion capture database. Thus, much time is spent on checking the database for the motions. Secondly, the domain experts want to select possible paths through the graph. Thirdly, more sophisticated techniques are based on local statistical models on motion sequences. Here the domain experts want to identify meaningful sequences in terms of number, length and selection. For research purposes, many test data sets have to be generated. Usually, data are recorded for a specific project only and only reused as long as persons that recorded the data are available.

To overcome these issues, a visual interface to the data is desirable. The domain experts should be enabled to get a quick overview of the data that can be refined for parts of special interest. In addition, the workflow could benefit from a simple searching interface to identify if motion sequences are available in the data set. Based on manually selected sequences, the domain experts would be able to apply local modifications to the motion graph.

2.3 Requirement Analysis

The requirements of the domain experts originally comprised four high-level questions that are addressed by our system: (1) Which variations of poses are available in the data set? (2) How can queries be specified intuitively? (3) Does a motion exist to re-use (or do I need to record new camera takes)? (4) Which style variations exist from start pose A to end pose B? In the course of the requirement analysis phase, we were able to report on the knowledge gain with respect to the research questions of the domain experts. In return, the domain experts helped us with the interpretation of the motion capturing data.



Fig. 2. Collaboration with users in the course of this design study.

In this phase, we discussed data sets, still images of rapid prototypes and continued with in-depth discussions. At the end of the requirement analysis phase after six months of collaboration, we were able to characterize the requirements on a finer level of detail.

- R₁ Overview:** The system should provide a ‘big picture’ of the data. A global ordering would be nice to assess similar data items.
- R₂ Aggregation:** To reduce the complexity, large data sets can be aggregated. The level of aggregation should be steerable.
- R₃ Cluster glyph:** Data aggregations should be displayed as visual structures. These should allow for an intuitive assessment.
- R₄ Motion graph:** This metaphor should be integrated. Poses should be displayed as nodes and motion sequences as edges.
- R₅ Level of detail:** Local behavior of the data should be explorable. This enables the domain experts to analyze details.
- R₆ Filtering:** It should be possible to exclude and re-include data. Filtering helps to concentrate on relevant parts of the data.
- R₇ Visual query formulation:** Formulating example queries for motion sequence search should be as intuitively as possible.
- R₈ Path search:** A state of the art retrieval technique should be integrated to search for motion sequences.
- R₉ Search result exploration:** Style variations of different retrieved sequences are interesting and should be recognizable.
- R₁₀ Multiple visual data representations:** Different aspects of the data should be shown side-by-side in the same view.

3 RELATED WORK

This work relates to the wider field of automated and visual-interactive analytics. We focus on the analysis of time-oriented data [2]. We start with methods and techniques for time-series aggregation and clustering. We then focus on visualization techniques for graphs, hierarchies and glyphs, which have been proposed specifically to enhance the creation and/or use of clustering techniques.

3.1 Aggregation of Time-Oriented Data

One can distinguish three different methods for time-series aggregation: descriptor-based techniques to obtain compact representations, projection-based techniques to reduce dimensionality and clustering techniques to group the data. Our approach relates to all of the three aggregation concepts.

Popular aggregation techniques based on time-series descriptors have been surveyed by Ding et al. [12]. However, most time-series descriptors have been defined for univariate time-series. In this work, we use a descriptor for multivariate time-series based on the recommendation of the domain experts [24]. Projection-based techniques reduce the dimensionality of the input data by the least significant components [16]. In contrast to descriptor-based techniques, projection-based techniques are well suited for multivariate time-series [42]. In this approach, we project high-dimensional vector representations of human poses in 2D via the PCA projection technique just like it was presented for high-dimensional earth observation data [8]. In addition, we use the Self-organizing Maps method as a projection technique to create a similarity-preserving color legend for human poses [7]. Specifically visualization-driven approaches based on clustering multivariate segments are relatively scarce. Berkhin [6] provides an overview of clustering on multidimensional data in general, while Liao [25] analyzes how existing clustering techniques have been modified to support time-series data. For time-series clustering we can further distinguish techniques which cluster (a) complete series, (b) subsequences, or (c) single timestamps of input time series. Andrienko

and Andrienko [3] present a raw-data-based approach for movement trajectory data. Scheepens et al. [32] present an aggregation method for multivariate trajectories based on continuous density surfaces. A technique for visual-interactive clustering of 2D time-dependent subsequences has been presented by Schreck et al. [34]. Turkey et al. [39] provide yet another perspective by examining the clusters as time-dependent entities, based on a frame-by-frame clustering of all trajectories. For univariate time series, popular ‘single-frame’ aggregations like the SAX approach [26] have been proposed by Lin et al. They show that a time series can also be represented by a discrete symbolic alphabet, which can transform the analysis task to a state-transition analysis. Scherer et al. [33] present an exploratory search system where single frames of bivariate time series are clustered. In our approach, we apply hierarchical divisive clustering on single frames of multi-variate human motion capture data (human poses). In combination with the PCA-projection technique, we enhance multivariate time series (motion capture data) analysis by node-link-oriented state transition analysis. Using the vocabulary of the domain experts, we present an aggregation of both, the *spatial* (descriptor, clustering) and the *temporal* (projection, state-transition) *domain* of the data.

3.2 Visual Analysis of High-Dimensional Aggregated Data

Visualization techniques have been used in combination with time-series clustering mostly for three purposes, all of which are relevant for our approach. The first purpose is the interactive steering to *create or modify* the clustering. The second purpose is to *view* its results. The third purpose is to *use* the clustering results to refine and filter views. Heinrich et al. [20] present a technique for the visual-interactive aggregation of genome sequence data, allowing the user to control the aggregation strategies. Guo et al. [19] combine multivariate and spatio-temporal information in a system for the semi-automatic generation of aggregations. Hierarchies created from clusters can be used to show the aggregated data on multiple levels-of-details, as it has been presented for vector field data [38]. Following this idea, our visualization of the state-transitions is also related to graph visualization [40]. In our node-link-diagram, we adopt Balzer and Deussen’s idea of surface-metaphors to represent cluster hierarchy levels as graphs [5]. In the works of Archambault et al. [4] and Abello et al. [1], steerable graph hierarchy construction and drawing techniques are presented. Similar to our work, a top-down approach is used for the hierarchical aggregation. However, while we apply a global aggregation level, the authors chose unbalanced drill-down techniques. Most tree-visualization techniques [35] can naturally be applied to show cluster hierarchies. They have e.g., been adopted by biologists to show genome sequence clusters [13]. Bisson and Blanch [9] propose *stacked trees* for showing cluster hierarchy and content in a single visualization for cluster comparison. Seo and Shneiderman [36] present an example of a dendrogram-based view coupled with a heatmap showing the details of the cluster. We adopt this approach with respect to how dendrograms can be used for cluster inspection, while differing by the underlying data type.

Ward [41] presents a taxonomy of glyph-based visualizations for high-dimensional data. He describes the strength, but also ascertains limitations of glyph placement strategies. The DICON framework [11] is a promising approach for cluster glyph visualization. Similar to our approach, the glyphs are used to provide meaningful high-level statistical information to represent multidimensional clusters. There are projection-based visualizations based on path metaphors for sequential n-grams [42], text and music collections [29, 37], multivariate time series [8], metabolic pathways [10] and motion capture data [21, 31]. While these approaches convince in the differentiation of multivariate data items, challenges in preventing occlusion for large data sources may remain. In our approach, we determine the layout of the state-transition graph by a projection, providing a map-metaphor for an ‘information landscape’ for high-dimensional data. We discarded the alternative use of graph layout methods (see, for example [18] for a recent survey), because they do not convey any additional information on the position of the nodes.




Fig. 3. Data types provided: (a) single human pose/frame, (b) pose aggregate/cluster, (c) hierarchical clustering tree (d) cluster motion graph.

4 DATA ABSTRACTION AND FUNCTIONAL SUPPORT

In this section, we describe the data processing capabilities of the system. We present a brief overview of the choices made for the feature extraction, the data structures, the aggregation techniques, and the search functionality. For research purposes, multiple motion capture databases are publicly available. In our experiments we use data obtained from the HDM05 motion capture database [28]. This database is a systematic recording of a wide variety of motions performed by various actors, in multiple repetitions. Many instances of similar human movement are also available as sets of extremely varied behaviors. Thus, the database provides a good basis for investigating the effectiveness of our approach.

4.1 Feature Vectors of Human Poses

In agreement with the domain experts, we define a *single human pose* as one ‘atomic’ data object, see Figure 3a. After the import of the data source (binary C3D-files), motions are represented as sequences of single poses. In the following, we present preprocessing routines and feature extraction steps as suggested by the domain experts. In a normalization step, the pelvis of each human pose is shifted to the origin of the coordinate system. To obtain poses that are always viewed ‘head-on’, each pose is normalized by a rotation around the z-axis. For overview and search, the domain experts recommend the extraction of a compact feature vector (FV) representation which keeps most of the original information. For this reason, the set of 3D-markers of each human pose is reduced to a relevant subset [24]. We obtain a FV with 48 dimensions containing 16 3D marker coordinates in absolute scale (cm unit). FVs will be the basis for subsequent aggregation techniques and the search functionality, respectively. In agreement with the domain experts, we use the Euclidean distance measure as the default similarity function.

4.2 Hierarchical Aggregation of Human Poses

A *pose cluster* (Figure 3b) contains large numbers of similar human poses. In the vocabulary of our targeted domain, a cluster represents a spatial aggregation of the data. It is the most important data structure in the analytical and exploratory components of the system. The domain experts can directly interact with clusters using the visualizations of MotionExplorer. MotionExplorer also provides statistical information to assess cluster quality.

MotionExplorer uses a hierarchical clustering algorithm to create aggregations, which can be inspected on multiple levels of detail (\mathbf{R}_2). The *pose cluster hierarchy* (see Figure 3c) is calculated in a separate thread using a divisive clustering algorithm to scale for large data sets [14]. That way, the coarsest clusters are promptly available for inspection. The domain experts may control two clustering parameters [6]:

- (a) which cluster is to be split next
- (b) the splitting strategy for a particular cluster

Since we want to obtain compact clusters, we define (a) splitting the cluster with the maximum standard deviation as a default. Alternatively, we suggest splitting the cluster with the highest number of elements. Concerning (b) we follow a representative approach and apply a k-means clustering with $k = 2$ as a default. Based on an inquiry with the domain experts, we provide both parameters as steerable. Hence, the data aggregation process becomes interactively adaptable.

In an early prototype of MotionExplorer, the algorithm performance of the divisive hierarchical clustering was compared with a k-means




Fig. 4. Cluster Glyph for Human Poses. The style variations of human poses around the cluster centroid are shown by opacity bands.

(partitioning) and DBScan (density-based) implementation. Two advantages of the divisive clustering became apparent: (a) the level-of-detail concept can be achieved with a single calculation and (b) the hierarchical structure allows for multiple levels of detail with the same data elements in respective sub-trees. We did not detect differences in the results, whether we tested with data with similar human poses or sets of extremely varied behavior.

4.3 Domain-specific Data Structures and Functionality

In addition to the described spatial aggregation, we provide an aggregation of the temporal data component. The clusters are used to create the adapted motion graph metaphor as a *directed cluster graph* (\mathbf{R}_4). With the clusters defining the nodes, two nodes are connected if there is at least one motion sequence in the database, which connects the corresponding poses (see Figure 3d). For each level of detail in the cluster hierarchy, an individual graph is calculated.

A last step concerns the search functionality (\mathbf{R}_5). A crucial step for data-driven methods are fast searching techniques which allow the domain experts to efficiently and effectively retrieve motion capture data of interest. The domain experts mostly apply a searching technique with a complexity of $k \log(n)$, where n is the size of the database, and k is the number of nearest neighbors. We adapt a technique by Krüger et al. [24] and exploit the generated cluster data structure as an index. First, all sequences are identified that contain the start cluster and the end cluster of the query in the right order. Then, the sub-sequences surrounded by the two cluster centroids are returned as the search result.

The result of this phase is a data abstraction and a functional adaptation to the domain. Each decision in the development phase was made in collaboration with the domain experts. In the following section, we report on the visual encodings of the targeted system.

5 VISUAL MAPPINGS AND INTERACTIVE FUNCTIONALITY

Based on the data abstraction described in Section 4 we now report on the different visualizations (Sections 5.1 - 5.4) and interaction methods (Section 5.5) of MotionExplorer. We conclude this section with a summarization of the design process during the implementation phase.

Our approach to the visual mapping process is in line with Fuchs and Hauser’s recommendation to create multiple visualizations to support multiple operations and show all important aspects of the data [17]. Our data models and visual representations are also inspired by Elmqvist and Fekete [14], where a model for building, visualizing and interacting with multi-scale representations of high-dimensional data is presented. The ability of meaningful data aggregate visualizations is regarded crucial to deal with highly aggregated data. Moreover, providing multiple levels of detail is seen important in current visual aggregation approaches.

MotionExplorer is developed for our domain experts, but there is potential for use beyond this domain. However, we continue to denote the *users* of the visual-interactive capabilities as *domain experts*.

5.1 A Cluster Glyph for Human Poses

Inspired by the human anatomical drawing by Leonardo da Vinci (Vitruvian Man), we decided to use a circular design for the cluster glyph (see Figure 4). We encode the information for cluster interpretation and visual data summary based on the preference of the domain experts (\mathbf{R}_3). We show the cluster centroid as a human stick-figure pose,




Fig. 5. Pose Hierarchy Explorer. The result of a hierarchical clustering is visualized as a dendrogram. Domain experts can adapt the number of shown clusters (currently 26) by dragging the aggregation level slider.

and the set of poses in the cluster as deviating, transparent figures, which is important for the domain experts. We limit the number of shown poses for each cluster to a maximum of 500. Based on an inquiry with the domain experts, the sampling method is in favor of the nearest neighbors to the centroid. For cluster quality assessment, we use the quantization error of a cluster as a measure of compactness, represented by a label at the bottom left. The relative cluster size is displayed by the size of the surrounding circle, the absolute size is labeled at the bottom right. Finally, we color the cycle of each cluster glyph to illustrate similarity among the clusters. The color value depends on the cluster centroid vector, and follows the system-global, similarity-preserving color legend at the bottom left corner of the system for a straightforward look-up (see Figure 1). The grid of the color legend is the result of a SOM that is trained using all feature vectors in the manner of a vector quantization scheme [7]. Thus, the most prominent human poses are arranged on the grid structure for a quick overview. We use the 2D colormap proposed by Ziegler et al. [44] to make effective use of the RGB color space.

5.2 Exploration of Hierarchically Clustered Human Poses

The *pose hierarchy explorer* meets the following requirements:

- it serves as a global overview (\mathbf{R}_1) by representing the tree structure of the cluster hierarchy
- it enables filtering to narrow down the number of elements (\mathbf{R}_6)
- it provides interactive level of detail functionality (\mathbf{R}_2 , \mathbf{R}_5)

Starting with the root cluster at the very top, sub-tree structures are displayed top-to-bottom. Clusters are represented with the pose aggregate glyph that helps to provide the global overview of the data set (see Figure 5).

We provide an aggregation slider to define the number of shown clusters (see Figure 7). Depending on the slider position, the highest cluster of each subtree below the current slider position is rendered. Elmqvist and Fekete describe this approach as ‘level traversal’ [14]. The level of each cluster (the depth in the tree) depends on the rank of their standard deviations and thus on the order of the split criterion in the clustering. As a consequence, the vertical axis of the hierarchy explorer encodes the standard deviation ranking of each cluster. This ‘position-by-ranking’ method was preferred by the domain experts in favor of an allocation by the binary depth. It provides the additional analytical benefit of identifying the cluster compactness ranking along the vertical axis. This is also the main advantage of the dendrogram over a Windows Explorer-like visual encoding of the clustering tree.




Fig. 6. The motion explorer shows human poses as nodes and motion sequences as edges. In this example skiing poses are represented from a lateral and a frontal view. The projection-based layout clearly distinguishes between the two postures. The prevent-overlap function is active in the sub-graph at the bottom right, based on the mouse focus.

An interaction with the vertical aggregation slider drills-down or rolls-up the number of visible clusters, respectively. Thus, the domain expert can obtain different levels of granularity. The effects of this aggregation level hold for the complete system. While approaches with local aggregation capability in single sub-trees exist [1] [4], Motion-Explorer provides a global aggregation level. With this specification, domain experts can always analyze clusters of equivalent standard deviations, which serves as an orientation in the assessment of style variations. If a local drill-down is needed, two approaches are possible: 1) present a detailed overview first and then exclude irrelevant data by filtering, or 2) start with a coarse aggregate and then explicitly open subtrees for more detail. This interesting design question also depends on the application goal. We plan to further examine the applicability of local and global aggregation level adaption in the future work.

The domain experts may exclude or include parts of the hierarchy from the view to focus on relevant parts of the dataset. Filter operations always apply to a cluster and its subtree. Filtered clusters are indicated as gray pose aggregates.

5.3 Analysis of Human Motion Sequences

The *motion explorer* (see Figure 6) serves as a complementary view to the hierarchy explorer. It shows human motion sequences between the aggregated poses that exist in the data set. The visual encodings fulfill the following requirements:

- overview of all poses (\mathbf{R}_1) at the current aggregation level (\mathbf{R}_5) following the motion graph metaphor (\mathbf{R}_4)
- global order of all poses to intuitively assess the similarity
- interactive means to avoid local overplotting
- support of visual query formulation (\mathbf{R}_7)

A node-link diagram provides a global overview of the aggregation level (\mathbf{R}_1), where nodes represent the human pose aggregates and edges show existing motion sequences. Domain experts are able to explore the diversity of poses and the available motion transition between them, thus meeting requirement \mathbf{R}_4 .

Concerning the node layout, the domain experts consider the global ordering of human poses as most important. Although close arrangement of two poses produces visual overplotting, the domain experts confirm the analytical importance of identifying dense regions on the map. Thus, we allocate the node positions via the linear and deterministic PCA projection in favor of force-directed layout algorithms. High-dimensional cluster information is projected in 2D by means of



Fig. 7. Pose Hierarchy Explorer. The result of a hierarchical clustering is visualized as a dendrogram visualization. Domain experts can change the level of detail by dragging the aggregation level slider. Here, poses of skiing movement were drilled-down to the aggregation levels of 5, 15, and 30.

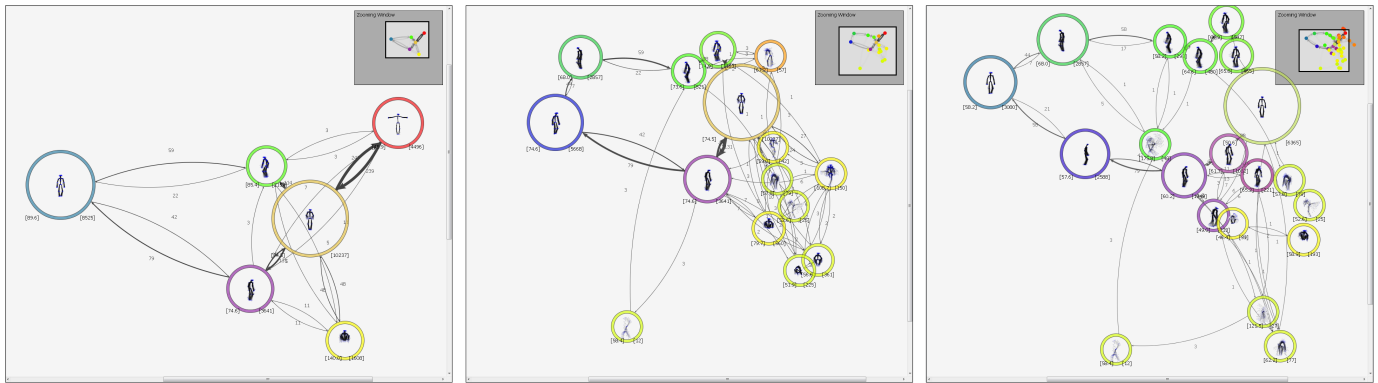


Fig. 8. Exploration of style variations. Based on an overview of the large data collection at aggregation level 6, the domain expert drilled down the level to 17, and 39, respectively. Zooming and filtering was applied to focus on relevant human motion. Various cycles of human movement exist. For example, at the aggregation levels of 17 and 39, a ‘cartwheel’ motion beginning with an upright stand was identified at the bottom left.

topology-preservation. A benefit of the deterministic projection is the reduced transformation-cost when the set of displayed clusters is modified by the domain experts. To avoid overplotting, a zooming technique and a local prevent-overlap functionality were selected. The zooming technique is combined with an overview at the upper right. This enables the domain experts to navigate through the human poses and to explore relevant poses on various levels of detail (\mathbf{R}_5). The local prevent-overlap functionality (also called ‘flower-metaphor’) rearranges colliding human poses as an overlap-free local layout around the pose in focus. An example is shown by the green and blue poses in Figure 6.

We provide a detail-on-demand view on a single pose, triggered by a mouse click. The enlarged visual representation of a cluster glyph allows the domain experts to explore details of a single pose cluster.

5.4 Exploratory Search

The *motion search interface* enables the domain experts to:

- run the human motion retrieval algorithm of the domain experts (\mathbf{R}_8) based on visually formulated queries (\mathbf{R}_7)
- analyze style variations in the search result (\mathbf{R}_9) based on exploratory analysis capability
- drill down the number of retrieved sequences to the number of relevant results (\mathbf{R}_6)

The search interface at the bottom of the system (see Figure 1) is divided into (a) a field that represents the start pose, (b) the motion search result list, (c) the end pose field and (d) the pose bundle animation. To enable visual query definition for the search interface, we provide a drag-and-drop technique by which example poses can be dragged into the start and the end pose of the search panel to start the search. This allows for visual search operations on large data sets with the effort of only two mouse interactions. Furthermore, the detail-on-demand view also provides an interface to define the pose either as

start or as end pose for the query. As soon as the domain experts have defined a query by example (\mathbf{R}_7), the retrieval algorithm of the domain experts is executed (\mathbf{R}_8). Resulting motion subsequences are arranged in the motion search list interface (see Figure 9). While the visual query formulation for the search is based on aggregated poses (clusters), the search result provides motion sequences on the granularity of single frames of the data set. The frames are separated by black vertical lines. Since the query formulation only assigns the start and the end pose, there are usual variations in the course of the retrieved motion sequences.

In consultation with the domain experts, resulting motion sequences are warped to fit with the horizontal display space per default. This enables an easy comparison of retrieved sequences along the relative time axis. The display of sequences in an absolute arrangement of time is also possible. In this case, each frame in the result visualization has an identical length but only the longest retrieved sequence spans the complete display space along the time axis, see Figure 9. Inspired by the idea of coloring queried sequences by an intrinsic property [20] [15], we again apply the similarity-preserving color legend and color the resulting subsequences according to the cluster affiliations. For an interactive exploration of various resulting sequences, a time slider is provided. By dragging the slider, an animation of all corresponding result poses is shown in the pose bundle animation view (see Figure 1 bottom right) and the domain experts can observe style variations between different motions (\mathbf{R}_9). A filter function enables the domain experts to reduce the number of style variations (\mathbf{R}_6). Inverting the filter status is also possible. For additional information on single subsequences, a focus-and-context interaction is implemented. By focusing the result list, the currently accessed subsequence is enriched by a key frame representation of the particular human motion.

This search interface allows our domain experts to find interesting human motion style variations across different results (\mathbf{R}_9). The application examples and Figure 9 further illustrate this aspect.




Fig. 9. The search interface with search results in a relative (upper) and an absolute (lower) representation. Two style variations can be identified. One short sequence from light green to dark green and a long sequence from light green to yellow to dark green. Focus and context is applied to show the frames of a single search result. The animation slider enables the domain experts to analyze the temporal variations of the result.

5.5 Linking Views

The system links together the visual encodings of clusters and poses in all views (\mathbf{R}_{10}). If domain experts change the aggregation level in the pose hierarchy explorer, the motion explorer is automatically adjusted to the new level of detail. Excluded clusters are grayed out in the hierarchy explorer and removed from the motion explorer. With a click on one of the clusters in the color legend, the system automatically adapts the aggregation level to the granularity that fits with the clicked pose of interest. Thus, the domain experts are able to start their analysis with a click on an interesting pose at the color legend. If a cluster in one of the views is focused, the cluster pose is also highlighted in every view where the respective cluster is visualized. This is especially helpful if the number of clusters is large and the color information is no longer sufficient to identify identical clusters in multiple views. We designed the search algorithm to ignore the current filter status, and always search on the entire data set. Finally, we attach tool tips to all active elements of the system to make the system self-explanatory.

5.6 Design Process

The design process was highly iterative. We conducted a formative laboratory design study during the implementation phase of Motion-Explorer in order to optimize the design choices. As a result, a variety of visual mappings and interactive capabilities were refined in the course of the process. 14 non-experts were asked to complete a questionnaire with 57 questions about visual encodings and interactivity preferences. Based on still images, we conducted informal interviews with the participants, which partially led to creative discussions beyond the concerns of the questionnaire. We observed the participants, using the prototypical system to solve analytical tasks or to communicate perceptual and cognitive preferences. We now describe improvements based on the results of the qualitative study.

Pose hierarchy explorer One of the heavily discussed design aspects was the visual representation of the pose hierarchy explorer. During the development process, Windows Explorer-like and treemap-like representations were presented to the participants. In the end, the latter representation was rejected because of the aggregation slider concept. Different visual representations of a cluster with respect to the changing aggregation level were another matter of concern. For parent and child clusters in the hierarchy, a trade-off between showing important information and an over-representation was achieved.

Motion explorer Since the domain experts preferred the similarity-preserving projection in favor of a force-directed layout, we discussed local overplotting avoidance strategies in the interviews. We list six considered design choices in order of preference from worst to best: (6) a global slider to adjust the glyph sizes, (5) focus and hide neighbors, (4) focus and highlight, (3) semantic zooming by locally adapting the aggregation level, (2) focus for an overlap-free local layout - the 'flower-metaphor', (1) geometric zooming + small zooming window. In the end, we decided for the last two choices.

Motion search interface The capability of the visual-interactive definition of queries was an important concern, the choices were: (4) left click and right click to assign start / end pose, (3) an interactive context menu, (2) detail-on-demand + assign start / end pose, (1) drag-and-drop. Again, the last two choices were implemented.

Linking views We recognized the need to choose visual encodings for the clusters that are as similar in all views as possible. In the end, we used the cluster glyph wherever possible. Additionally, brushing and linking is now used to indicate clusters in multiple views. Based on an inquiry with many participants of the design interviews, the color legend was made 'clickable'. As a result, the system now adjusts the aggregation level to fit best to the clicked cluster. This functionality can serve as a starting point for data exploration.

6 CASE STUDY

The motivation for the case studies was the question of how Motion-Explorer helps the domain experts with their research. We were interested in finding out if the domain experts would be able to identify their data collections and if the visual encodings and the interactive functionality is intuitive and usable. We were also interested in whether the domain experts would accept the system as a tool for decision-making in their scientific work.

6.1 Field Study

We conducted an insight-based summative field study with five domain experts of human motion synthesis at their lab. Due to the small number of participants, we will not discuss our findings in terms of quantitative measures. The procedure was as follows: After a familiarization, we introduced the domain experts to the purpose of the field study. We presented an introduction to the tool and to the supported analysis tasks. We divided the study into two parts. Each participant was asked to (a) perform four specific tasks and after that (b) was invited to run the system in an exploratory manner on her own data sets without supervision.

The data collection for the task-based tests is presented in Table 1. The four tasks were to (1) become confident with the system, to (2) identify a 'T-pose', to (3) formulate a visual query of a 'jumping-jack' motion, and to (4) identify style variations in the search result. Concerning (1) we collected informal feedback as to how the domain experts familiarized themselves with the system. The identification of the 'T-pose' (2) was successfully accomplished by all domain experts, the slowest participant needed two minutes. Recalling the number of human poses in the collection, this means a 'browsing speed' of at least 150 human poses per second when MotionExplorer is used for the identification of poses. A 'T-pose' is shown at the upper right of Figure 8 (red pose). Most domain experts clicked on the color legend as the entry point to focus on a particular pose. The result of a 'jumping-jack' search can be seen in the title Figure 1. The task of identifying style variations in the search result proved that both the absolute and the relative search result representation was used. Subsequent to the

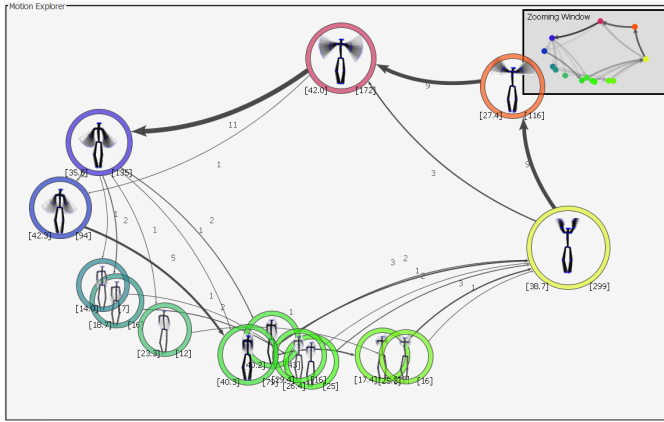


Fig. 10. Exploration of micro style variations (rotating arms). Many micro-variations within the green poses can be explored. However, the yellow and the magenta poses are traversed by each motion sequence.

task-based tests, the participants were requested to use the system with their own data.

6.2 Results of the Data Exploration

The exploratory part of the field study was conducted on data collections recommended by the domain experts. Thus, we were able to observe the domain experts using the system in a real-world setup. To enhance the outcome of this insight-based method, we collected still images, photos and videos. Most interesting for us was determining how fast the domain experts would familiarize themselves with the system, and what data findings or insights would occur. We present five notable observations made when the domain experts performed complex analysis tasks in Table 2.

7 SUMMARY

We present a qualitative summary of our findings in the field study and subsequently discuss the approach of this design study in general.

7.1 Generalization of Field Study Results

In our summative field study, we applied MotionExplorer in a variety of search tasks. The exploratory search results differed due to variations in the applied data sets, candidates and tested task completion tests. We believe that the decision to allow users the exploration of datasets of their own choice paid off. We obtained valuable qualitative feedback. The basic idea of the case study was to observe whether the system meets the domain experts' requirements. In addition, we hoped that the domain experts would acknowledge benefits of the system.

New Data Findings *'I have found way more style variations in my data than expected. For me, locally adapting the motion graph model has become even more important'*

The domain experts report on a surprisingly large number of style variations within their search results. Filtering unwanted retrieved motion sequences in the search result helped the domain experts to obtain motion sequences of interest. They liked the exploratory nature of the

Table 1. HDM05 motion classes used for task completion tests in the field study. The applied collection contains 19,598 human poses in total.

Motion Class	#	Motion Class	#
clapAboveHead	17	kickLFront	29
elbowToKneeLelbowStart	27	kickRFront	30
elbowToKneeRelbowStart	27	punchLFront	30
grabFloorR	16	punchRFront	30
grabHighR	29	skierLstart	30
grabMiddleR	28	squat	52
jumpingJack	52	walk2StepsLstart	31
		walk2StepsRstart	31




Fig. 11. Observation of a domain expert exploring offense and defense behavior of boxing movement. The search result presents two sequences that contain style variations with similar behavior. Here, filtering was heavily applied to better resolve boxing motions in detail. The red pose (focused in the search result) is highlighted in the upper views.

motion explorer. Based on the provided overview of their data, the domain experts identified the capability to decide which motion sequences are still missing in their data collection and should be recorded in the near future.

Communication *'Can I take screen shots with this tool? I would like to show my analysis results to my colleagues'*

The domain experts identified the ability of the system to help communicating search and analysis results. Domain experts reported on the need to perform plausibility checks and identified MotionExplorer as a suitable communication medium. Hypothesis generation on pose and motion variations is a common task for the domain experts. The possibility to visually verify hypotheses and to communicate the outcome was appreciated.

Usability *'It simply fits very well. Usually we start with reading tutorials'*

The domain experts were pleasantly surprised that they learned to use the system very quickly. They welcomed the intuitive description of the cluster glyph of human poses and the interplay of the different views. The domain experts considered most of the visual encodings appropriate and identified most of the interactive capabilities without assistance. The attachment of tool tips turned out to be beneficial to lookup the functionality of previously unknown active elements. After the task completion tests, all domain experts familiarized with the system and were thus able to work autonomously.

Effectiveness *'I've been working with motion capture data for eight years now, but this perspective on the data was really enriching. I had a lot of fun with the tool'*

We recognized a curiosity about the system among the domain experts, who recognized the components of the system and their interplay as helpful. The pose hierarchy explorer and the motion graph in combination form a helpful solution to obtain a global overview. Filter and zooming functionality were readily used to drill-down to local features in the data. The domain experts liked visualizing the data in detail to analyze problems with their currently applied models. As demonstrated in the case study, domain experts successfully performed a re-identification of known data aspects. Before, they had to choose statistical means to approach similar problems.

Efficiency *'No need to analyze poses and sequences by hand. Producing an effective result in only a few minutes. I like that'*

The increased efficiency is maybe the most substantial benefit of the system. A reduced effort in performing analytical tasks was recognized by all domain experts. Gaining fast insight into the complex data type of human motion sequences by the overview visualizations was widely accepted, especially if the data collection is large. The identification of human poses, motion sequences and style variations thereof

Table 2. Complex analysis tasks performed by domain experts in the observational phase of the case study. Mapping to the system requirements.

Analysis Goal	Data Set	Analysis Process and Analysis Result	Requirement
Re-identification of poses and motion in a large data collection	Long uncut camera take of an actor; previously known motion	One expert explored a previously known uncut camera take. With the motion explorer, she identified style variations on a macro level. Different paths were identified to get 'from A to B', Figure 8 details. The domain expert was affirmed when she identified the expected 'cartwheel' motion cycle at an aggregation level of 17.	R₁, R₂, R₃, R₄, R₅, R₆, R₁₀
Exploration of 'micro-style variations'	12 cut-sequences of 'rotating-arms' motion	One analyst was interested in style variations of 'rotating-arms' motion at the micro-level. As can be seen in Figure 10, there are variations in the green phase of the motion sequence. Apparently the actor of the cut-sequences performed various <i>body configurations</i> while her arms pointed at the ground.	R₁, R₂, R₃, R₄, R₅
Exploration of boxing poses	Collection of cut-sequences, including boxing	Another domain expert was interested in the exploration of cut sequences containing boxing motions, the observation is shown in Figure 11. After an information drill-down, a query with two different defensive stance poses was executed (orange and purple). The first and the third retrieved sequence show similar behavior.	R₁, R₂, R₃, R₄, R₅, R₆, R₇, R₈, R₉, R₁₀
Exploration of skiing movement; finding missing camera takes	Cut-sequences; different variations of skiing motions	We observed a domain expert looking for missing cross-country skiing motion data. She set the aggregation level to 26 (see Figure 5), two sub-graphs are visible in the motion explorer (see Figure 6). Apparently, no human motion data is available that connects the two sub-graphs, i.e. body rotation. It might be recorded in future.	R₁, R₂, R₃, R₄, R₅, R₁₀
Searching for a pose to connect two different motion sequences	Cut-sequences of two different motion sequences	In this use case, a domain expert explored two different human motions for a common pose. By identifying this pose, a synthesis of both motions could be created. The result is shown in Figure 1, the green upstand pose at the bottom right fulfills this criterion. Based on this finding, the domain expert may now connect the red pose and the purple pose to a new synthesized human motion in a subsequent work step.	R₁, R₂, R₃, R₄, R₇, R₈, R₉, R₁₀

was considered simple, which was quoted as very time-consuming in the past. We acknowledge that not all participants were able to perform a search with the minimal costs of five clicks. Nevertheless, we report on a significant improvement of efficiency, especially for large collections. E.g. the identification of human 'T-poses' in an unknown data set (and a new system) took between seconds and minutes. The domain experts stated that going through 20,000 human poses would usually take the domain experts several hours. To conclude, the aggregation of multiple occurring poses and motions considerably reduces the work load for going through such large data collections.

7.2 Discussion and Extension Possibilities

Our developed system supports overviewing and searching in a complex type of data; human motion capture data. To the best of our knowledge, our system defines for the first time exploratory search and analysis for large motion data based on a customized and interactive data clustering approach. Our defined visual display integrates cluster visualization by a custom glyph visually aggregating sets of poses. It uses a graph representation to show similarities of pose clusters together with state transitions. We believe visual-interactive analysis in motion data can help to better understand and re-use large repositories of motion data. Considering new technologies for motion tracking, we expect much more motion data, inducing the need for scalable retrieval and analysis techniques to become available in the future.

As shown qualitatively, our system could be operated by expert and non-expert users in an intuitive and effective way. The received comments indicate that the approach can benefit tasks in motion analysis. The motion data representation in our case is an instance of multivariate, time-dependent data. This data is typically hard to visualize. In our case, we could draw on representing the motion poses by a stick-figure-oriented glyph, which is easy to interpret by users, as opposed to more abstract high-dimensional visualizations, which may be harder to interpret by users, especially in case of large data sets.

We also identified a number of limitations of our current study that should be investigated as a next step. First, pose clustering in our approach is done by interactive clustering. This relies on the user being able to find useful cluster parameters and interpret the obtained cluster. Also, automatic classification of poses should be integrated. Based on training data, one could train classifiers to automatically, and possibly more robustly, segment different poses. Also, the identification of style variations at the micro level may be difficult to do only visually, especially if variations are all mapped to the same glyph. In densely populated glyphs it may be difficult to visually discriminate the styles. Again, automatic classification could be useful to this end.

Our state transition view currently uses the PCA to project the pose clusters to 2D space. This may induce overlap and make it difficult to get an overview of complex state transition graphs. Constrained layout or space filling approaches should be considered to scale for larger

transition graphs. In effect, loosing some positional similarity may not be expensive, considering the gain in display space for larger data sets. Also, the motion glyph or transitions could be enriched with meta data (if available) of the motion recordings. If certain poses have been pre-classified, this should be included in the visualization. Interactive capability for efficiently labeling glyphs and transitions is also possible. We can also imagine several promising extensions to increase the analytical potential of the system. For example, if we consider motion data captured in a soccer game, additional data could be captured on the *result* of a certain motion. E.g. a players motion led to a goal, the player missed the goal, or a player was injured. We imagine that through appropriate visual overlays, correlation analysis between motion and respective results can be useful for training purposes.

Finally, we performed a qualitative study involving mainly university students on the non-expert side. This may be seen as a bias in the considered user population. More formal studies, including also quantitative measurements on more precisely defined tasks, are desirable to further assess the system capabilities. For example, we may more closely incorporate our analysis and search system into workflows for interactive motion synthesis and observe the improvement potential.

8 CONCLUSION AND FUTURE WORK

We presented MotionExplorer, an exploratory search system for large data collections of motion capture data. The system provides an overview of human poses in a dendrogram visualization that represents the result of a hierarchical clustering. A node-link diagram enables the user to analyze human poses as nodes, where each node shows a collection of similar pose instances by a stick-figure-oriented glyph. Links between the pose glyphs indicate motion sequences and allow exploration of the poses in context. Visual interfaces for query formulation and the exploratory search result analysis support the user in retrieving human motion sequences and style variations thereof. We evaluated the system by laboratory design interviews with non-experts, and in a field study with domain experts. Based on the field study, the experts confirm the usability and efficiency of the system to support their day-to-day work. Analysis results within the field study also presented new aspects of the data that had not been taken into consideration or had not been visually communicated before. A number of limitations and extension possibilities of our work that should be considered for future work have been discussed in the preceding section.

ACKNOWLEDGMENTS

The authors thank the 'Multimedia, Simulation and Virtual Reality Group' at the University of Bonn, Germany. We appreciated the collaborative work and the bilateral exchange in knowledge and expertise during the complete duration of this design study. We also thank the 14 non-experts that participated in the laboratory design interviews for their time and feedback.

REFERENCES

- [1] J. Abello, F. van Ham, and N. Krishnan. ASK-GraphView: a large scale graph visualization system. *IEEE Transactions on Visualization and Computer Graphics*, page 2006, 2006.
- [2] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Springer, 2011.
- [3] N. V. Andrienko and G. L. Andrienko. Spatial generalization and aggregation of massive movement data. *IEEE Trans. Vis. Comput. Graph.*, 17(2):205–219, 2011.
- [4] D. Archambault, T. Munzner, and D. Auber. Tuggraph: Path-preserving hierarchies for browsing proximity and paths in graphs. In *Pacific Vis*, pages 113–120, Washington, DC, USA, 2009. IEEE Computer Society.
- [5] M. Balzer and O. Deussen. Level-of-detail visualization of clustered graph layouts. In S.-H. Hong and K.-L. Ma, editors, *Int. Asia-Pacific Symposium on Visualization 2007, Sydney*, pages 133–140. IEEE, 2007.
- [6] P. Berkhin. A survey of clustering data mining techniques. *Grouping Multidimensional Data*, pages 25–71, 2006.
- [7] J. Bernard, T. Ruppert, M. Scherer, J. Kohlhammer, and T. Schreck. Content-based layouts for exploratory metadata search in scientific research data. In *JCDL*, pages 139–148, New York, NY, USA, 2012. ACM.
- [8] J. Bernard, J. Wilhelm, M. Scherer, T. May, and T. Schreck. TimeSeries-Paths: projection-based explorative analysis of multivariate time series data. In *Journal of WSCG*. University of West Bohemia, Plzen, 2012.
- [9] G. Bisson and R. Blanch. Improving visualization of large hierarchical clustering. In *Proceedings of the 2012 16th International Conference on Information Visualisation, IV '12*, pages 220–228, Washington, DC, USA, 2012. IEEE Computer Society.
- [10] U. Brandes, T. Dwyer, and F. Schreiber. Visualizing related metabolic pathways in two and a half dimensions. In G. Liotta, editor, *Graph Drawing*, volume 2912, pages 111–122. Springer, 2003.
- [11] N. Cao, D. Gotz, J. Sun, and H. Qu. DICON: interactive visual analysis of multidimensional clusters. *Visualization and Computer Graphics*, 17(12):2581–2590, 2011.
- [12] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh. Querying and mining of time series data: experimental comparison of representations and distance measures. *VLDB Endow.*, 1(2):1542–1552, 2008.
- [13] M. B. Eisen, P. T. Spellman, P. O. Brown, and D. Botstein. Cluster analysis and display of genome-wide expression patterns. *Nat. Acad. of Sc.*, 95(25):14863–14868, 1998.
- [14] N. Elmqvist and J.-D. Fekete. Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Transactions on Visualization and Computer Graphics*, 16(3):439–454, 2010.
- [15] F. Fischer, F. Mansmann, and D. A. Keim. Real-time visual analytics for event data streams. In *ACM Symposium on Applied Computing, SAC '12*, pages 801–806, New York, NY, USA, 2012. ACM.
- [16] T. Fu. A review on time series data mining. *Eng. Appl. Artif. Intell.*, 24(1):164–181, 2011.
- [17] R. Fuchs and H. Hauser. Visualization of multi-variate scientific data. In *Computer Graphics Forum*, volume 28, pages 1670–1690, 2009.
- [18] H. Gibson and P. Vickers. A survey of two-dimensional graph layout techniques for information visualisation. *Information Visualization*, 2012.
- [19] D. Guo, J. Chen, A. M. MacEachren, and K. Liao. A visualization system for space-time and multivariate patterns (vis-stamp). *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1461–1474, 2006.
- [20] J. Heinrich, C. Vehlou, F. Battke, G. Jger, D. Weiskopf, and K. Nieselt. iHAT: interactive hierarchical aggregation table for genetic association data. *BMC Bioinformatics*, 13(S-8):S2, 2012.
- [21] Y. Hu, S. Wu, S. Xia, J. Fu, and W. Chen. Motion track: Visualizing variations of human motion data. In *PacificVis*, pages 153–160, 2010.
- [22] L. Kovar and M. Gleicher. Automated extraction and parameterization of motions in large data sets. *ACM Transactions on Graphics*, 23(3):559–568, 2004. SIGGRAPH 2004.
- [23] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. *ACM Trans. Graph.*, 21(3):473–482, 2002.
- [24] B. Krüger, J. Tautges, A. Weber, and A. Zinke. Fast local and global similarity searches in large motion capture databases. In *SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 1–10. Eurograph. Ass., 2010.
- [25] W. T. Liao. Clustering of time series data—a survey. *Pattern Recogn.*, 38(11):1857–1874, 2005.
- [26] J. Lin, E. Keogh, S. Lonardi, and B. Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *SIGMOD workshop on DMKD*, pages 2–11, New York, NY, USA, 2003. ACM.
- [27] M. Müller. *Information Retrieval for Music and Motion*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.
- [28] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation: mocap database HDM05. Computer Graphics Technical Report CG-2007-2, Universität Bonn, 2007. Data available at www.mpi-inf.mpg.de/resources/HDM05.
- [29] E. Pampalk, A. Rauber, and D. Merkl. Content-based organization and visualization of music archives. In *Tenth ACM International Conference on Multimedia*, pages 570–579, New York, NY, USA, 2002. ACM.
- [30] T. Pejsa and I. Pandzic. State of the art in example-based motion synthesis for virtual characters in interactive applications. *Computer Graphics Forum*, 29(1):202–226, 2010.
- [31] Y. Sakamoto, S. Kuriyama, and T. Kaneko. Motion map: image-based retrieval and segmentation of motion data. In *ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 259–266, Aire-la-Ville, Switzerland, 2004. Eurographics Association.
- [32] R. Scheepens, N. Willems, H. van de Wetering, G. Andrienko, N. Andrienko, and J. J. van Wijk. Composite density maps for multivariate trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2518–2527, 2011.
- [33] M. Scherer, J. Bernard, and T. Schreck. Retrieval and exploratory search in multivariate research data repositories using regressional features. In *Proceedings of the 11th annual international ACM/IEEE joint conference on Digital libraries, JCDL '11*, pages 363–372, New York, NY, USA, 2011. ACM.
- [34] T. Schreck, J. Bernard, T. Tekušová, and J. Kohlhammer. Visual cluster analysis of trajectory data with interactive Kohonen maps. *Information Visualization*, 8:14–29, 2009.
- [35] H.-J. Schulz. Treevis.net: A tree visualization reference. *IEEE Comput. Graph. Appl.*, 31(6):11–15, 2011.
- [36] J. Seo and B. Shneiderman. Interactively exploring hierarchical clustering results. *Computer*, 35(7):80–86, 2002.
- [37] S. Stober and A. Nürnberger. Musicgalaxy: A multi-focus zoomable interface for multi-facet exploration of music collections. In *Exploring Music Contents*, volume 6684 of *LNCS*, pages 273–302, Berlin / Heidelberg, 2011. Springer Verlag.
- [38] A. Telea and J. J. van Wijk. Simplified representation of vector fields. In *Conference on Visualization*, pages 35–42, Los Alamitos, CA, USA, 1999. IEEE Computer Society Press.
- [39] C. Turkay, J. Parulek, N. Reuter, and H. Hauser. Interactive visual analysis of temporal cluster structures. *Computer Graphics Forum*, 30(3):711–720, 2011.
- [40] T. von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. J. van Wijk, J.-D. Fekete, and D. W. Fellner. Visual analysis of large graphs: State-of-the-art and future research challenges. *Computer Graphics Forum*, 30(6):1719–1749, 2011.
- [41] M. O. Ward. A taxonomy of glyph placement strategies for multidimensional data visualization. *Information Visualization*, 1:194–210, 2002.
- [42] M. O. Ward and Z. Guo. Visual exploration of time-series data with shape space projections. *Computer Graphics Forum*, 30(3):701–710, 2011.
- [43] R. W. White and R. A. Roth. Exploratory search: Beyond the query-response paradigm. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 1(1):1–98, 2009.
- [44] H. Ziegler, T. Nietzsche, and D. A. Keim. Visual exploration and discovery of atypical behavior in financial time series data using two-dimensional colormaps. In *Proceedings of the 11th International Conference Information Visualization*, pages 308–315, Washington, DC, USA, 2007. IEEE Computer Society.