https://doi.org/10.2352/ISSN.2470-1173.2017.1.VDA-387

Visual-Interactive Semi-Supervised Labeling of Human Motion Capture Data

Jürgen Bernard[1], Eduard Dobermann[1], Anna Vögele[2], Björn Krüger[3], Jörn Kohlhammer[1,4], Dieter Fellner[1,4]
[1] ***Technische Universität Darmstadt, Germany***
[2] ***Institute of Computer Graphics, University of Bonn, Germany***
[3] ***Gokhale Method Institute, Stanford CA, United States***
[4] ***Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany***

Abstract

The characterization and abstraction of large multivariate time series data often poses challenges with respect to effectiveness or efficiency. Using the example of human motion capture data challenges exist in creating compact solutions that still reflect semantics and kinematics in a meaningful way. We present a visual-interactive approach for the semi-supervised labeling of human motion capture data. Users are enabled to assign labels to the data which can subsequently be used to represent the multivariate time series as sequences of motion classes. The approach combines multiple views supporting the user in the visual-interactive labeling process. Visual guidance concepts further ease the labeling process by propagating the results of supportive algorithmic models. The abstraction of motion capture data to sequences of event intervals allows overview and detail-on-demand visualizations even for large and heterogeneous data collections. The guided selection of candidate data for the extension and improvement of the labeling closes the feedback loop of the semi-supervised workflow. We demonstrate the effectiveness and the efficiency of the approach in two usage scenarios, taking visual-interactive learning and human motion synthesis as examples.

Introduction

Recording, storing, and analyzing human motion capture (MoCap) data has become common practice, e.g., to substantiate hypotheses about human body mechanics. One frequently applied approach is to track actors with multiple markers, yielding quantitative information depending on time, referring to MoCap data as a special type of multivariate time series data. Large volumes of temporal events, combined with the multivariate nature of recorded body configurations over time accounts for a complex data type that quickly exceeds human capabilities in manual processing and analysis.

In fact, the analysis of MoCap data is interesting for a variety of research and application areas such as medicine, sports, biomechanics, or animation. Typical analytical goals are to explore, compare, simplify, segment, classify, retrieve, or synthesize MoCap data. One important prerequisite for these goals is the characterization and abstraction of MoCap data, e.g., by an alphabet of labels representing different MoCap classes. Ideally, these classes match the semantics of an application domain and the information need of individual users. In this approach, our notion of labeling is to assign class information to vector data, i.e., to combine semantics with the kinematics of MoCap data.

Recently, *visual analytics* [1] was introduced for the visual-interactive analysis of MoCap data, combining the computational power of algorithmic models with human ability to identify and interact with patterns. Example approaches address exploratory search [2], the visual abstraction and aggregation [3], the visual comparison of gaits [4], or the visual-interactive segmentation [5] of MoCap data and patterns. These visual interfaces showed how MoCap patterns can be *explored visually* and how users can *interact* with the data.

However, labeling of MoCap data has not been subject to visual analytics. In fact, user-based labeling of MoCap data can serve as an effective way for the characterization and abstraction of large MoCap collections. Especially when the data collections are complex (such as the HDM05 [6] or CMU [7]) the abstraction of MoCap data can facilitate visual access, search, comparison, and exploration tasks in an effective and efficient way. In addition, downstream analysis goals such as the tedious process of human motion synthesis could be supported with powerful visual-interactive techniques.

The overall motivation of the approach is to enable broad ranges of user groups to assign meaningful labels to MoCap data. This research agenda comes with a variety of challenges. Ideally, user-defined labels preserve as much of the relevant information and likewise neglect irrelevant information of the data. A prerequisite is to provide visual-interactive access and control of the MoCap data to foster the labeling process. Further, related work shows that algorithmic models could be included supporting users in the labeling process. At a glance, these models are either supervised or unsupervised, which should be conflated in a unified approach. A related challenge is how syntactical and semantical information can be included. In addition, analytical questions arise whether the user-defined labels cover the data space in a meaningful way. Accordingly, the selection of meaningful MoCap candidates for labeling poses a challenge, depending on criteria such as avoiding overfitting, reducing the entropy produced by unlabeled data, and addressing the user task at hand.

We present an approach for the visual-interactive semi-supervised labeling of MoCap data. Overall, we propose three primary contributions. First, we present a conceptual workflow for the semi-supervised labeling process. In two core steps, users can label single MoCap sequences and subsequently explore the MoCap search space by taking the user-defined labels into account. Different mechanisms support users in the selection of alternative MoCap sequences (candidates) to further extend and improve the labeling which closes the feedback loop of the workflow. Second, we present a tool that assembles novel visual-interactive in-

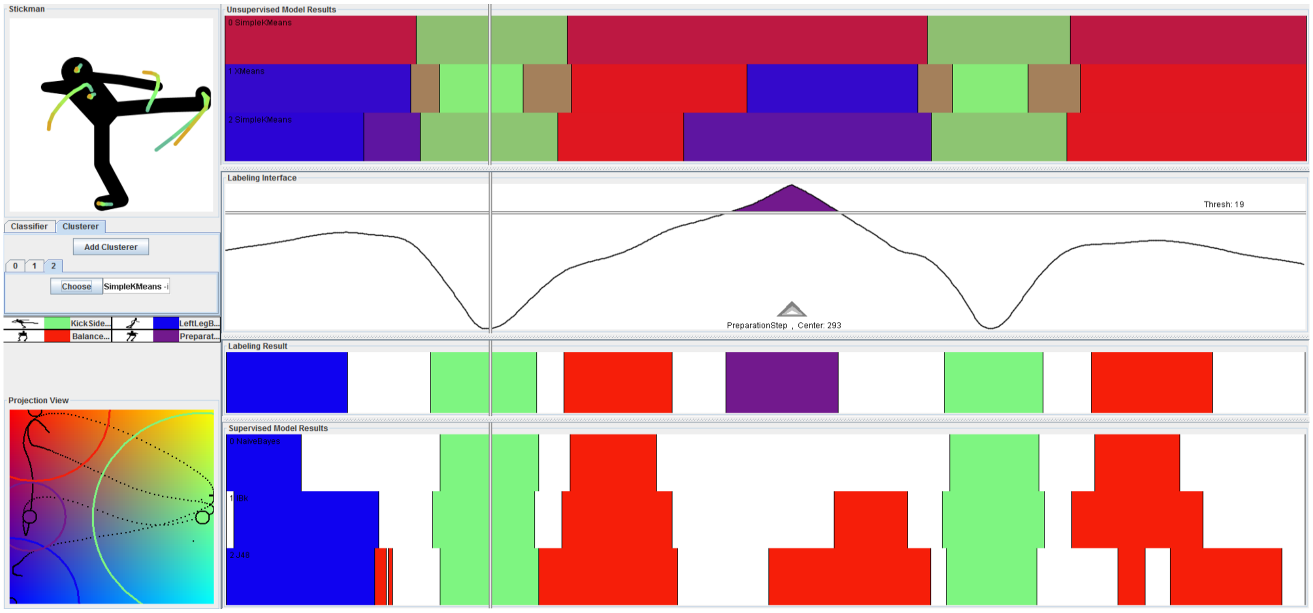

Figure 1. Overview of the visual-interactive semi-supervised labeling approach. A stick figure visualization (upper left) shows the current position within a MoCap sequence. Users can interactively create labels and define a threshold range for each label (center). Unsupervised (top) and supervised (bottom) algorithmic models support users in the labeling process. A projection-based view (lower left) relates labels to the multivariate input space. In the example a kick motion is labeled. The green label represents the kick execution while three labels (blue, red, purple) indicate preparation and balancing phases.

terfaces that enable users to assign labels to MoCap sequences. Unsupervised and supervised algorithmic models support users to identify meaningful *key poses* in the data which will be used for the labeling. Users can relate the set of defined labels with the multivariate MoCap data to assess data coverage in a projection-based interface. Our third contribution covers the exploration of labeling results in the entire MoCap collection. We provide two techniques that use labels to abstract MoCap data to sequences of event intervals allowing the exploration of event data in overview visualizations. Finally, we present two complementary techniques for the identification of MoCap candidates for the next labeling iteration. We demonstrate the effectiveness and the efficiency of the approach in two usage scenarios. The first scenario incorporates the principles of active learning with the goal to train a classifier for MoCap data on the basis of a user-defined set of labels. The second scenario follows the goal of human motion synthesis. A set of labels is created that allows the identification of MoCap sequences applicable for motion synthesis.

Related Work

This approach combines the techniques from the MoCap application domain with information visualization and visual analytics. We structure the related work in two main topics: the core analysis of MoCap data, and visual-interactive approaches for pattern analysis an labeling of multivariate time series.

Analysis of MoCap Data

In the last decade, large data collections featuring a range of different activities are available [6, 7]. Methods addressing a number of goals related to the analysis of human motion have been advanced using such established data sets for comparison and reference. Two baseline analysis tasks for MoCap data are retrieval and labeling, supporting analysis goals such as segmentation, exploration, motion editing, or synthesis.

MoCap Retrieval

Retrieval systems allow for examination of motions present in MoCap data sets. They also enable locating sub-motions similar to a query example. Hence, retrieval techniques are especially related with the idea of interactive label refinement in the tool proposed. Contrary to our approach which enables experts to select and refine choices of classifiers, parameters, as well as suitable primer poses for training the labeling to fit their need, most established systems depend on predefined methods and parameter sets. To name one example, based on frame by frame comparison, Match Webs [8] are an index structure for MoCap retrieval. In contrast, Chai and Hodgins [9] use a neighbor graph as part of a pre-processing step on a motion data collection allowing fast nearest-neighbor search. Both approaches include pre-processing steps of quadratic complexity and thus can't be employed to search larger datasets at interactive speed. Retrieval of similar poses by classification is a strongly related idea to the proposed approach. Binary geometric features [10] can be used in order to achieve this. Also, efficient look-up methods reduce complexity and make for better speed (cf Müller et al. [10, 11]). On the downside, methods based on Boolean features cannot describe close numerical similarity which is one advantage of body position-based features used in the approach at hand.

Dimensionality reduction techniques such as principal component analysis (PCA) are well established [12] to compute compact representations of otherwise high-dimensional MoCap data. Another way to come up with low dimensional representations is computing feature sets in a normalized pose space [13]. Dimensionality reduction is achieved by using a subset of joint positions. As opposed to the feature space of all joint positions of the MoCap marker set as used in our approach, their a feature set consists of hand, feet and the head positions. Like in this paper, the geometric normalization in the pose space is achieved by adjusting the origin of each pose to an estimated position of the center of mass

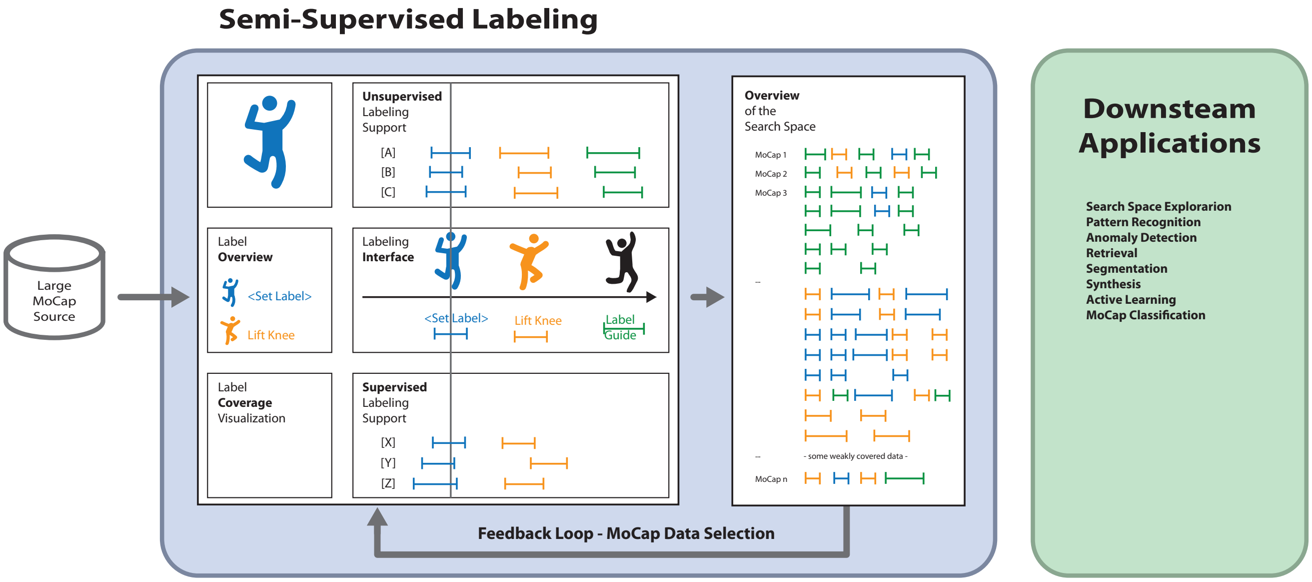

Figure 2. *Conceptual workflow of the approach. A large and possibly unknown MoCap data source is accessed. An assembly of visual-interactive interfaces supports users in the creation of label information including unsupervised and supervised algorithmic models. An overview visualization enables the exploration of the entire search space including label information. The selection of relevant MoCap candidates for labeling closes the feedback loop.*

given by the several hip markers. They also show, that it is possible to search these feature sets efficiently by employing a kd-tree as index structure. Extensions to online versions of this framework where proposed by Tautges et al. [14] and Riaz et al. [15], where the later allow for additional constraints in online search.

Annotation and Labeling of MoCap

Annotation and labeling of multivariate time series is needed for many applications in data mining or machine learning. A range of semi-supervised or unsupervised techniques have been developed in the last decade.

A synthesis application, similar to ours, was introduced by Arikan et al. [16] who describes how labels may be used in a synthesis framework, namely by training a support vector machine on set of motion data. Müller et al. [17] introduce the concept of motion templates representing semantic relations in the solution of the annotation problem as a classification task. The approach by Zhou et al. [18] solves a segmentation problem using hierarchical cluster analysis to find a partition of given multivariate time series into disjoint segments. Vögele et al. [19] describe an unsupervised method for finding activities and primitive motion units as well as learning pre-defined labels by clustering in multivariate time series. However, their method is based on neighborhood graphs solving the problem efficiently. Also, the primitive motion units found by their method correspond more exactly to repetitions and cyclic elements of the segmented motion. This concept was further extended to fine motor hand and finger movements by Stollenwerk et al. [20]. Bouchard et al. [21] present an automatic segmentation approach producing semantics-based motion units from general MoCap data by examining the qualitative properties using Laban Movement Analysis (LMA). This bridges the gap between high-level semantic features and low-level kinematic features. To our knowledge, a straightforward way for users to explore synthesis candidates interactively is missing to date.

Visual-Interactive Approaches

Various visual-interactive approaches involve the creation and analysis of clusters, patterns, or segments, all of which can support users in assigning label information. Moreover, some approaches focus on labeling techniques for segmented data. We focus on approaches for multivariate time series and MoCap data.

Visual-Interactive Temporal Pattern Analysis

Especially if the data collection is initially unknown, visual data analysis can help to gain an overview of the data. At a glance, users can be supported in the exploration of structural information, such as clusters, outliers, or other interesting patterns. Various approaches exist that reveal patterns in general multivariate time series. We refer to the surveys of Aigner et al. [22], Bernard [31], and Andrienko et al. [23] for systematic overviews of visual approachces for time-oriented and spatio-temporal data. In fact, the visualization of patterns can be seen as a labeling support, even if most approaches do not provide explicit labeling interaction. Examples for large and complex time series data are the LiveRAC tool designed for the visual exploration of system management data [24], the CloudLines approach for the detection of clusters [25], or the Gatherminer approach for visual discovery of patterns [26]. Another visualization technique for multivariate time series uses data projection to draw a path metaphor in 2D [27]. Data aggregation is applied in combination with color coding to highlight temporal patterns. Similarly, Hao et al. amplify temporal patterns with different color codings [28]. Our approach goes one step further using static 2D colormaps to encode labels in a similarity-preserving way [29]. Various visual-interactive tools directly address the exploration of MoCap patterns [30, 2, 4, 3], often facilitated with stick figure visualizations. The approaches build on data aggregation and dimension reduction [31]. MoCap patterns can be compared by position, shape, and color information. However, visual-interactive labeling is not provided. The retrieval-based approach by Krüger et al. [13] focuses on the

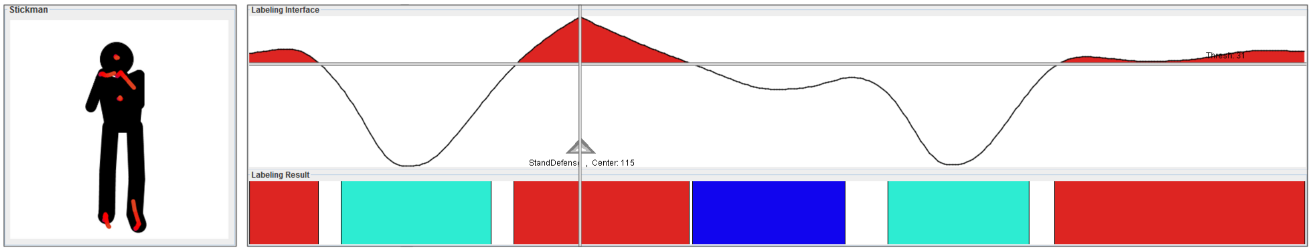

Figure 3. *Labeling interface. Users can localize key poses which can be added to the set of labels. Here, a red label was created representing the defense stand of a boxer. A horizontal line defines the threshold range of the label. The example demonstrates that labels can occur multiple times within a sequence.*

analysis of periodic MoCap patterns. Similar to one of our interfaces the MotionTrack tool uses data projection to visualize MoCap data in 2D [32]. However, the focus is on the visual comparison of motion patterns and sequences, while we visualize and compare user-defined labels.

Visual-Interactive Segmentation and Labeling

Recently, different approaches examined the segmentation and labeling of multivariate time series. Similar to our classification and clustering algorithms the approach by Alsallakh et al. [33] supports users with automated segmentation results with the ability to improve labels manually. Röhling et al. [34] compare the results of various segmentations using similar visual variables for the labels. In contrast to our work, the focus is on the analysis of parameter spaces for segmentation algorithms. An inspiring visualization approach for MoCap data is presented by Alemi et al. [35] using color codings to represent speed and acceleration information. Our approach is also related to a tool that allows users to run multiple clustering and classification algorithms [5]. More precisely, our approach adopts and extends two features. First, the applicability to calculate supervised (classification) and unsupervised (clustering) algorithms, which we use as labeling support. Second, the ability to visually compare multiple results of such supportive algorithms, similar to the approach of Röhling et al. [34]. Our approach differs in the analytical focus. We enable users to interactively define and analyze an alphabet of labels on the basis of key poses, while the core focus in [5] is the segmentation and segment comparison.

Concept

Conceptual Workflow

This work conflates principles from information retrieval, machine learning, data mining, information visualization, and visual analytics. Before we present our technical approach, we first outline a concept that assembles different principles to a general workflow, depicted in Figure 2. A large and unexplored MoCap data source represents the situation at start. A visual-interactive interface enables users to analyze a single MoCap sequence in detail and to assign label information. The results of supervised and unsupervised algorithms can be used to create meaningful labels. Sets of user-defined labels can be observed in overview visualizations showing the distribution of existing labels across the entire data set. Thus, the visualization supports the exploration of the search space including label information. A feedback loop closes the semi-supervised process, the next iteration can, e.g., be triggered with the guided selection of meaningful MoCap candidates.

Requirements to the Approach

We formalize the problem statement of this work provided in the introduction section to derive a set of functional requirements. The following ten requirements reflect the design goals of our implementation and thus, helped us to build an appropriate approach. The requirements can directly be used to rebuild a semi-supervised labeling approach. As an alternative they build a baseline for further extension. Finally, we make the set of requirements explicit for future approaches working on methodology.

- **Req$_1$** Preprocessing – meaningful feature extraction and selection, meaningful similarity measures
- **Req$_2$** Algorithms and Parameters – allow interactive selection and control of models and model parameters
- **Req$_3$** Visual Representation – provide visualizations that help users to access and understand the data at hand
- **Req$_4$** Interaction – enable users to interact with the data and define labels in an intuitive way
- **Req$_5$** Label Adoption – provide a means to apply the label information on the entire data set
- **Req$_6$** Overview – provide an overview of the entire search space including label information
- **Req$_7$** Candidate Selection – visual-interactive support for users selecting new MoCap candidates for labeling
- **Req$_8$** Refinement – enable users to modify the labeling to improve the (subjective) labeling quality
- **Req$_9$** Few Labels – automated transfer of labels should be possible on the basis of very few manual labels
- **Req$_{10}$** Labeling Progress – provide a means that reflects the coverage of the labeling with respect to the data set

Approach

We present an approach that enables users to label MoCap data in a visual-interactive semi-supervised way. Our approach addresses the principle steps of our conceptual workflow (cf. Figure 2) as well as the requirements proposed in the concept section. Building on this, we present a set of visual-interactive interfaces providing different perspectives on the data and supporting different subtasks within the labeling process.

In a section on data characterization, we introduce the data collection and the operations applied to preprocess the data. Next, we present an overview of our novel visual-interactive interfaces assembled to a tool. The following section introduces a solution to integrate unsupervised and supervised labeling support into the approach, before we present the visual-interactive interface for labeling MoCap data. A core feature is supporting the definition of *key poses* and *thresholds* which characterize the labels used in this approach. The section on labeling optimization presents the

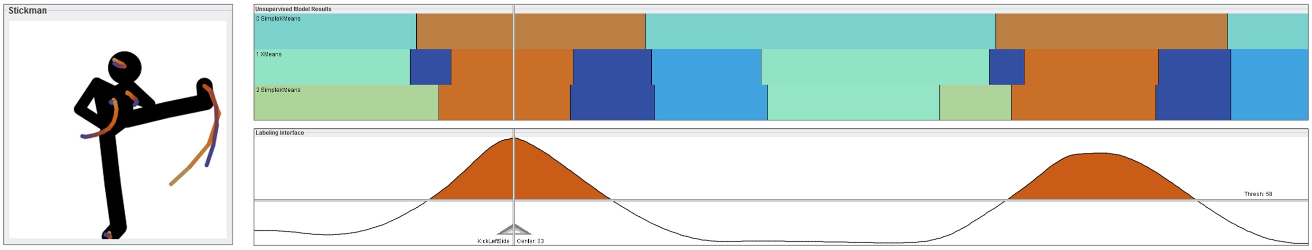

Figure 4. *Labeling support by unsupervised algorithmic models. We provide an interface where users can execute clustering algorithms and compare multiple results visually. In the example all three clustering algorithms (SimpleKMeans, XMeans, SimpleXMeans) match at a brown region indicating a kicking pose. We decide to add a label ("KickLeftSide"). Interestingly the label recurs later in the sequence, indicated by the clustering results and the labeling algorithm.*

visual-interactive interface where users can relate the set of labels with the multivariate input space of the MoCap data. Next, we present a solution providing an overview of the labeling results applied for the entire MoCap data collection. Finally, the feedback loop of the proposed workflow is closed with different techniques enabling users to identify meaningful candidates of MoCap data for the next labeling iteration.

Data Characterization and Feature Selection

The data at hand are collections of MoCap data such as [6, 7]. In particular, the input to our system consists of binary C3D data of human actors performing motion sequences from different categories. Motion sequences consist of a pose per actor per frame which is given by the marker positions present in the trial. Often, the frequency of frames per second is 100 Hz or above. The data is parsed and loaded into an internal file format including the data content, as well as metadata about marker positions and joints. Each pose is then characterized by the set of 44 3D positions of the individual markers. Each pose and every motion passes several preprocessing steps such as outlier removal and missing value handling, to guarantee data quality (**Req**$_1$). One important step is normalization to make individual poses comparable form a kinematic and also a visual perspective. Each pose is translated to its estimated center of mass, i.e., to the center of mass of all markers describing the actor's hip. The Euclidean distance serves as a measure of similarity when comparing different poses.

Overview of the Visual-Interactive Tool

We present a tool that conflates several novel visual-interactive interfaces, Figure 1 provides an overview. At the upper left the visual representation of single human poses is shown providing an intuitive visual access to MoCap data (**Req**$_3$). We created a visualization on the basis of a stick figure, which is a best-practice metaphor to represent MoCap poses visually [2, 5, 10, 13, 36]. Our stick figure is special in the way that it reflects all 44 marker positions available in the data set, connected with joints according to the metadata. In addition, we represent

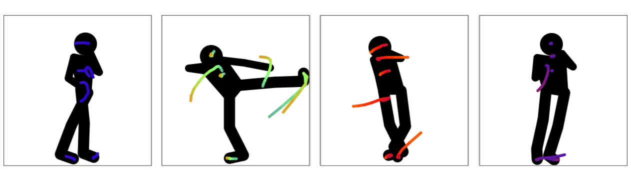

Figure 5. *Stick figure visualizations showing four user-defined labels. Colored bands represent the temporal domain of the MoCap sequence and thus, motions of individual body parts.*

the temporal domain indicated with colored bands at any marker (angle) of the pose. The length of the bars correlates with the speed of the respective body part (see Figure 5). At the center of the visual-interactive labeling interface (cf. Figure 1) a slider control is provided. The slider allows to browse through the MoCap sequence supporting the fine-grained analysis of input data. At the top of the tool the results of unsupervised models support users in labeling unknown data. Moreover, the visualization can be used to stick the labeling towards the kinematic information, which complements the semantical perspective to the data (expressed by the notion of the user). At the bottom, an interface for the results of supervised MoCap labeling support is provided. Users can add a variety of classifiers learning the user-defined labels. Individual characteristics of the model results can be assessed by visual comparison. Finally, at the lower left, we present a visual interface that combines the labeling information with the multivariate input data. Data projection is applied to relate and compare a 2D representation of the MoCap data with the label information.

Unsupervised and Supervised Labeling Support

We make use of unsupervised and supervised algorithmic models from data mining and machine learning to support visual-interactive labeling. For this purpose, a concept is applied presented in the context of segmenting multivariate time series data [5], which we adopt for semi-supervised labeling of MoCap data.

Especially if the given MoCap sequence is previously unknown unsupervised labeling support is crucial. For this purpose, we provide an interface where users can add multiple clustering routines within the labeling process in an interactive way (**Req**$_2$), see Figure 4. The results of these routines are visualized line-by-line, with different color codings showing partitions of the data set. Visual comparison enables the identification of subsequences where the results of different clusterings fairly match each other. In this way, the results of multiple clusterings can be a means to identify key poses for robust and generalizeable MoCap labels. A right click within the interface automatically adds the respective key pose to the label set (**Req**$_4$). With the unsupervised labeling support, we also provide a solution to solve cold start problems. Moreover, it represents the algorithm perspective applied on the high-dimensional feature space (kinematics). This information can be combined with the information need of the user (reflecting semantical information), building a valuable basis for meaningful MoCap labels.

The interface at the bottom of the system uses similar visual encodings, an example is depicted in Figure 6. Users can add multiple classifiers as a means to conduct supervised labeling support within the labeling process (**Req**$_2$). The classifiers automatically

Figure 6. Labeling support by supervised algorithmic models. An interface allows adding classifiers learning the use-based labels. The results of the classifiers can be visually compared. In the example a boxing pose is shown, the yellow label represents an upright stand. Three classifiers are tested with the current MoCap sequence. It can be seen that BayesNet and IBk fairly match the labeling, while the DecisionTable produces a bad classification result.

learn the user-based labeling information. The results of the multiple classifiers can be compared visually, allowing the identification of matching and but also of conflicting classification results. This enables users to assess the quality of supervised models with respect to the given MoCap sequence and the set of labels. Good candidates can subsequently be tested with the remaining MoCap collection, e.g., to create a compact and representative abstraction of the MoCap collection (**Req**$_5$).

Visual-Interactive Definition of Key Poses (Labels)

A core functionality of the approach is the definition of labels. Figure 3 presents the visual-interactive interface enabling users to define meaningful key poses within sequences of MoCap data (**Req**$_4$). Meaningful key poses can be expressed by user needs, data properties, or be inspired by the information of the unsupervised labeling support. A simple right click within the interface adds (removes) a key pose. Every user-defined key pose builds the basis for one label. In the example in Figure 3 a user defined a key pose represented with a red label. The key pose shows a boxer in a defense position looking towards the camera. A line chart visualization within the labeling interface shows the proximity of the MoCap data with respect to the user-defined key pose over time, extending the technique of Müller et al. [17]. The metaphor of filled areas (here above the threshold) is inspired by horizon graphs [37]. A black horizontal threshold line indicates which parts of the MoCap sequence will be assigned to the red event interval when the abstraction algorithm will be executed. This threshold line can be dragged in vertical axis allowing users to adjust the threshold value to a meaningful level. In this way, users can individualize the label definition for every given label. It can be seen that the red label in Figure 3 recurs at different times of the MoCap sequence. In other words, the red label covers large parts of the sequence and can, e.g., be used to detect recurring patterns. However, it can also be seen that some parts of the MoCap sequence are not in the proximity of the red key pose requiring additional labels for a meaningful data abstraction.

Label Analysis and Optimization

The interface at the lower right of the tool explicitly relates key poses with the multivariate input of the MoCap data. A dimension-reduction algorithm represents the structure of the multivariate MoCap data in 2D. Every pose of the MoCap data is reflected with a small dot in the plot, referred to as a scatterplot metaphor. The black curved path line represents the MoCap sequence as a trajectory in 2D. Every key pose is represented as a black small circle. For every key pose a colored outline reflects the extent of a label in the vector space, which we call *orbits*. The example in Figure 7 shows four labels (red, purple, blue, and green), the green label has the largest orbit. The extent of an orbit complies with the user-defined threshold, visually approximated with a circular shape. The interface, combining multivariate MoCap poses, key poses, and orbits enables users to assess the coverage of the input space, as well as intersections and gaps between orbits. Users can use this information to modify and refine the labeling result (**Req**$_8$). To further support the exploration of the input space, the slider control shown in Figure 3 can be used to drag through the given MoCap sequence. The current position of the slider is highlighted with a plus shape, in Figure 7 the plus is at the right of the interface. We call the dragging interaction concept the ‘rollercoaster animation’ [27]. Another interaction design provided in the interface allows the direct definition of labels on the basis of the input space (**Req**$_4$). Users can right-click any location in the 2D space and assign the nearest neighboring pose to the set of key poses used for labeling. The color information of the interface is used to color the new label and thus, to visually discriminate the label from the set of existing labels.

We consider dimension reduction approaches (e.g., based on data projection) a non-trivial analytical task. To avoid issues of presentation quality and other side effects, we build on best-practice approaches for projecting multivariate time series into 2D [27] and related works [38, 39, 40]. One subject to future work would be the definition of synthetic key poses, e.g., by inverse projection of the selected 2D coordinate (would limit the set of available projections to linear candidates). We neglect this strategy since a) the use synthetic body positions is a research field on its own and b) our approach explicitly works on real data with labels based on really existing key poses.

Exploration of the Large MoCap Collections

The workflow proposed in the concept section emphasizes the iterative character of the labeling process. If users have finished labeling a single MoCap sequence, it is crucial to apply the labeling on the entire data set. Two key requirements to a visual interface supporting this task are to provide an overview of the entire search space including label information (**Req**$_6$) and visual-interactive support for the identification of candidates for the next labeling iteration (**Req**$_7$). One additional challenge is to condense a large and complex MoCap data collection and represent it in a visualization. In this way, we provide a means to apply the label information on the entire data set, and present it visually (**Req**$_5$).

In our approach, we use the label information to abstract MoCap data to a compact and yet faithful representation. For

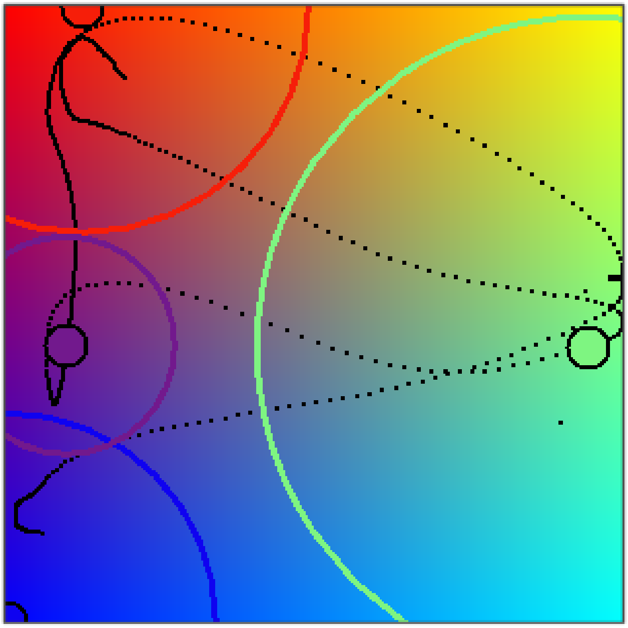

Figure 7. Interface for the analysis and optimization of labels. A dimension-reduction algorithm projects the poses of the multivariate MoCap data into 2D, yielding a path metaphor (black dotted line). Key poses are represented as black circles. Colored circular outlines ('orbits') indicate the threshold for every label. The interface can be used to assess the coverage of the data set by the labels, as well as for the identification of overlaps and gaps. In the example the green label ("KickSide", cf Figure 1) has the largest extent.

this purpose, we borrow the notion of *event intervals* [41] for the visualization of large collections of labeled MoCap sequences. Typically, an event interval is defined by the start and the end (time/index) of an event label, according to a given object depending on time or other sequential information. In addition, color and textual information can be provided, e.g., to communicate semantical context. We provide two different mechanisms for the transformation of MoCap data into sequences of event intervals (**Req**$_5$). The first variant is based on the user-defined labels and the orbits (thresholds). A blackbox algorithm based on a sliding window approach calculates when a sequence enters and leaves the orbits of the labels in the vector space. Accordingly, new event intervals are added to the sequence of event intervals. The second variant builds on user-defined classifiers. Within the labeling process these classifiers have learned the label information on the basis of the candidate MoCap data. As a result the classifiers can be tested with the remaining MoCap collection. An optional postprocessing step is to increase the accuracy of the result by neglecting MoCap data weakly covered by the selected classifier. We recommend the Simpson's diversity index as a criterion of class diversity which can be used as a threshold.

The visualization of large volumes and varieties of event intervals is a research field on its own. Various methods and techniques have been presented for the visualization, exploration, as well as for the simplification of event intervals (see, e.g., [42] for a survey). A review on existing approaches reveals that our approach does not require a novel visual interface to address requirements **Req**$_6$ and **Req**$_7$. In turn, we utilize the EventFlow [41] tool, a best-practice approach for the visual-interactive exploration of large collections of event intervals. After the transformation of

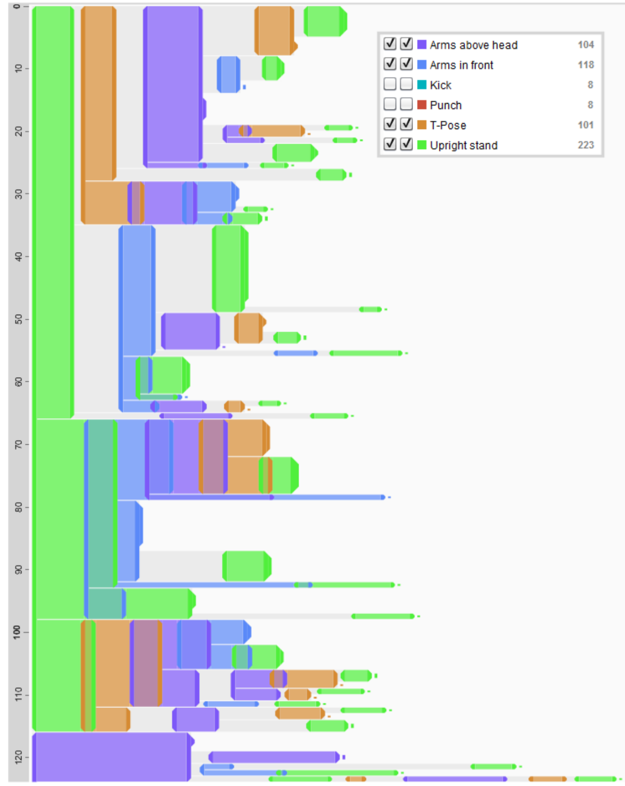

Figure 8. Exploration of MoCap labels represented as event intervals. The EventFlow [41] tool provides an overview of all labeled sequences of the search space, using color coding for the labels. Grouping and ordering similar event sequences helps to structure the MoCap collection.

MoCap data into sequences of event intervals, the interface is used to show the output result. In Figure 8, an overview of 140 MoCap sequences is represented, all sequences contain at least one of the user-defined labels. Users are able to group similar event sequences which helps to structure the MoCap sequences and thus, to ease the visual access. In turn, our abstraction of MoCap data to event sequences, combined with EventFlow [41], can be used to provide an overview of large collections of labeled MoCap collections (**Req**$_6$). In addition, the interface reveals both frequently occurring patterns as well as unique or outlier patterns. Both types of analytical tasks are relevant for the identification of candidates for the next labeling iteration (**Req**$_7$).

Identification of MoCap Candidates for Labeling

The identification of a labeling candidate closes the feedback loop proposed in the workflow. Users can load the particular MoCap sequence in the labeling interface, and refine, improve, or extend the labeling in another iteration (**Req**$_8$). We provide different mechanisms for the selection of meaningful candidates for the next labeling iteration (**Req**$_7$).

Visual-Interactive Identification

One class of mechanisms for the identification of interesting candidates follows a visual-interactive strategy. The interface that implements the overview task described in the workflow can be used to provide drill-down and detail on demand capability. Based on the identification of MoCap patterns and detailed in-

Labels given (12):	'PunchLeftSide', 'PunchLeftFront', 'PunchRightSide', 'PunchRightFront', 'StandDefends', 'StandWide', 'KickRightFront', 'KickLeftSide', 'KickRightSide', 'KickLeftFront', 'StepRight', 'StepLeft'
Motion Class Used for Labeling	**# Related Labels**
Kicks forward (right foot)	'KickRightFront', 'StandDefends', 'StepLeft'
Kicks to side (right foot)	'KickRightSide', 'StandDefends', 'StepLeft'
Kicks forward (left foot)	'KickLeftFront', 'StandDefends', 'StepRight'
Kicks to side (left foot)	'KickLeftSide', 'StandDefends', 'StepRight'
Punches forward (right arm)	'PunchRightFront', 'StandDefends'
Punches to side (right arm)	'PunchRightSide', 'StandDefends'
Punches forward (left arm)	'PunchLeftFront', 'StandDefends','StandWide'
Punches to side (left arm)	'PunchLeftSide', 'PunchRightFront', 'StandDefends'

Coverage by 12 labels trained on motion classes related to Martial Arts (cf [6], Section 3-2)

formation, it is possible to identify meaningful candidates for the next labeling iteration (**Req**$_7$). Two driving aspects for the identification of candidates for labeling are a) overlapping labels and b) gaps between labels. Overlapping labels may call for action of label improvement. Potentially large gaps between labels makes MoCap sequences good candidates for the definition of additional labels. Both examples (overlaps and gaps) can be seen in Figure 8.

Entropy-based Candidate Identification

One class of candidates exists that cannot be qualified with visual interfaces for the exploration of existing event intervals; the set of unlabeled MoCap sequences. Requirement **Req**$_7$ suggests algorithmic support for the propagation of MoCap candidates to be used for labeling. We refer to such models as *coverage models* since our aim is that user-defined labels cover as large volumes of the search space as possible. Requirement **Req**$_7$ is borrowed from research in active learning where the problem is, e.g., addressed by entropy-based sampling methods, aiming to reduce as much entropy as possible with as few iterations as possible (see, e.g., [43] for an overview). Entropy is an information-theoretic measure describing the amount of information needed to encode a distribution. The data space spanned by the chosen feature vector approach has $44 * 3 = 132$ dimensions. Thus, labeling such a high-dimensional vector space without algorithmic support is a tedious process. To improve scalability, our coverage model only considers volumes within the space actually populated by existing MoCap poses. For every given pose in the feature space the coverage model calculates the arithmetical distance to the set of user-defined labels. In this way, the coverage model seeks poses (i.e., volumes) within the vector space with high remaining entropy. The output of the coverage model is a list of unlabeled poses, ordered with respect the calculated distances.

As a positive side effect of the coverage model, it is possible to assess the progress of the labeling process (**Req**$_{10}$). The remaining distances propagated by the coverage model can be used as a an algorithmic stopping criterion. Thus, the approach can be facilitated with a means to inform users about the labeling progress. The progress is assessed in one of the usage scenarios in the next section, Figure 10 depicts the decrease of the average distance calculated by the coverage model.

Usage Scenarios

We demonstrate the usefulness of our techniques by example of two usage scenarios. The examples show how users are enabled to define labels for large collections of MoCap data for different analysis goals in an effective and efficient way.

As a source of data, the HDM05 MoCap data collection [6] is employed. This version is publicly available and contains C3D data of 5 actors performing roughly 1,500 motion trials from 5 different categories (roughly 70 motion classes with 10 to 50 realizations per class). More precisely, the 'cut' version of the C3D motion takes is used which contains the full set of trials that have been preprocessed by cutting away the T-poses that are usually the beginning and end of each take. An overview of the data used in the scenarios is provided in Table 1. As one additional challenge, we assume that the contents of the data collection are to some degree unknown at start, i.e., none of the available documentation on the trials or actors is taken into account in the experiments. As a result the provided unsupervised algorithms will guide users visual-interactively towards initial labels. As a second additional challenge, we are interested in a labeling result that only contains very few labels. Likewise, it is our aim to choose the labels such that the data collection is well reflected. Please note that more labels are likely to produce even better results, which is beyond the scope of the scenario.

Visual-Interactive Learning of a MoCap Classifier

We demonstrate how the the visual-interactive semi-supervised labeling interface can be used to learn classifiers for MoCap data. A set of user-defined labels is created in combination with various supportive classifiers. In the course of the labeling iterations the classifiers are trained with the labels, referring to the scenario as an active learning approach. At the end of every iteration the output of the classifiers is used for the exploration of the MoCap data collection. With the identification of MoCap sequences which are still weakly covered with the current state of the label set the feedback loop is closed, building the basis for a the next labeling iteration.

In the execution of the experiment, visual-interactive definition of key poses in the first MoCap sequence and clustering the results guides the user towards the first set of meaningful labels. Figure 4 shows the results of three clustering routines calculated for the first MoCap sequence (a kick sequence). We set our first label to reflect the brown clusters depicting the kick poses ("KickLeftSide") of the sequence. In the following, we add two additional labels describing the defensive stand within the kick sequence ("StandDefends", "StepRight"). We load another kick sequence into the interface and refine the "StandDefends" label, Figure 3 details. With the labeling in hand, the supervised classification is executed and the results can be evaluated on test sequences in an explorative way. In Figure 6 the results of three classifiers are shown leart with the three labels and the training data. It can be seen that the classifiers fairly match for "StepRight" (here: yellow), while the other two labels are classified more heterogeneously. In addition, the DecisionTable seems to produce weak results compared to BayesNet and IBk.

For the visual evaluation of first results, the EventFlow tool can be applied. Thus, major gaps across the investigated sequences that will need to be eliminated in the next iteration may be identified. In this way, a suitable training set can be identified and iteratively enhanced while using a small number of labels only. At the same stage, our *coverage model* is used suggesting yet unlabeled MoCap sequences that could contribute to optimal coverage when labeled. This feedback loop supports our active learning strategy. We proceed with the labeling process with other included in the data set. The process enables the user to come up with a set of 12 labels which represent large parts of the data collection. In the course of the iterative process, the cov-

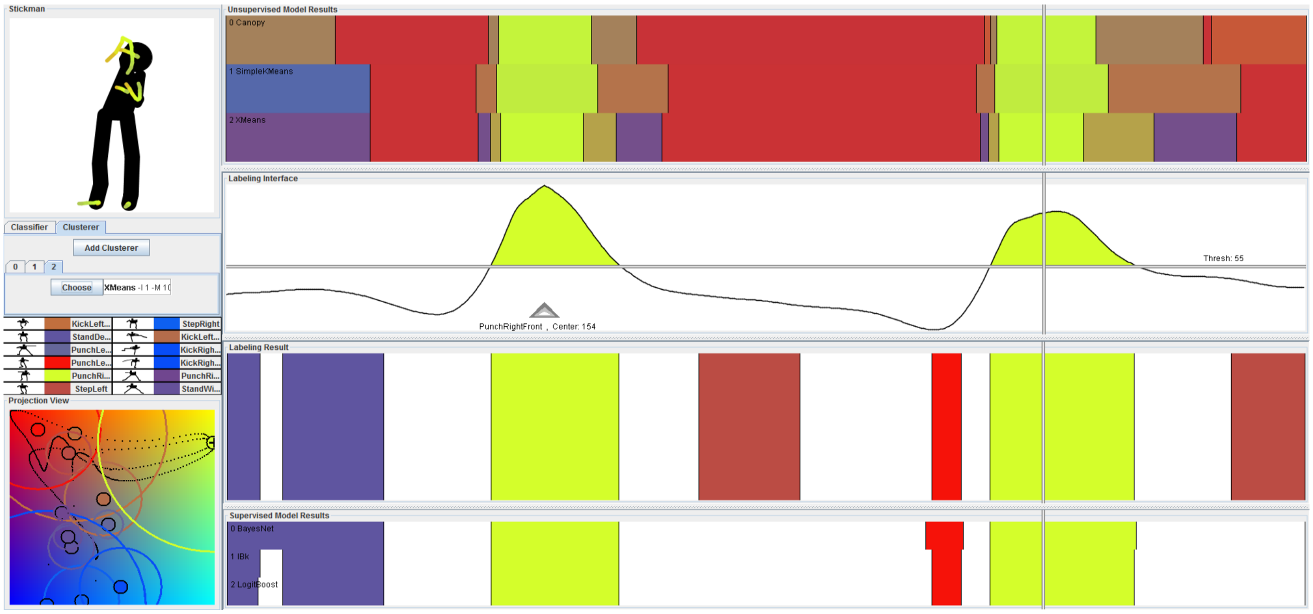

Figure 9. *The visual-interactive labeling system showing a process with 12 labels. In the current iteration a "PunchRightFront" label is added (yellow), strongly supported by three clustering results. The current MoCap sequence is covered by the classifiers with up to five labels (bottom). The projection-based label overview (lower left) shows the global label distribution in the vector space. It can be seen that the yellow label covers a new area of the MoCap space.*

erage model indicates the progress of the labeling process. Figure 10 shows a barchart representing the progress assessed by the coverage model. It can be seen that the remaining average distance of all poses to their closest labels decreases in the course of the process. With an increasing number of labels the numeric output of the coverage seems to converge. This demonstrates that the coverage model can be used to assess the process of the labeling. However, we ascertain that more labels may would further improve the coverage. The labeling interface with 12 labels is presented in Figure 9. In the current iteration a yellow label was added ("PunchRightFront"). The result of the labeling process is now explored with EventFlow to identify different MoCap patterns associated with motion classes (such as instances of locomotion, exercise). Figure 11 shows an overview of the labeled and abstracted MoCap collection as a typical result of a labeling scenario. The given example contains 12 different labels assigned to trials of 8 different motion classes. The result shows that the active learning process has provided initial labeling for more than 40 motion classes that are related to the ones used in the labeling phase. In particular, it can be seen that many of the labels occur in various MoCap sequences. As an example many sequences share the green "StandDefends" label. Similar, the "StepLeft" and "StepRight" poses seem to connect various classes of recorded motions. We select a set of MoCap sequences sharing "StandDefends" and "StepLeft". In Figure 12 these sequences are depicted in detail showing micro variations of the different motion styles. Users can clearly distinguish overlaps and gaps for in-depth motion analysis or further label improvement.

We conclude the first usage scenario. Note that the choice of a relatively small number of labels (see Table above) does not reflect the total number of motion classes in the HDM05. However, the usage scenario demonstrates that even a small number of labels can achieve high coverage on the full data set, and could be easily extended by using the *coverage model*.

Visual-Interactive Labeling for MoCap Synthesis

The second usage scenario demonstrates the usefulness of the semi-supervised labeling workflow and the implemented visual-interactive interface for the analysis goal of MoCap synthesis. Data-driven MoCap synthesis applications are commonly based on a set of pre-annotated motion data as well as a larger, possibly unstructured, heterogenous data pool. From there, subsequences of data may be retrieved, modified, and rearranged in order to come up with new motion sequences, optionally satisfying a set of additional constraints. In the execution of the synthesis, at least a single known MoCap sequence needs to be previously annotated or discretized by labels. This sequence serves as a primary piece in the synthesis workflow. Its labels allow for extraction of segments or subsequences desired for the synthesis result. Iteratively generating more annotation on the data pool such as in the first usage scenario is an optional but not a necessary step. However, the optimization of the synthesis primer sequence is an important matter. The expert user pursuing the synthesis task needs to select the most representative motion for the envisioned synthesis (e.g., the most representative 'jumping jack' motion). The user-defined labels need to be refined in this way, particularly, because the choice of the primer motion has a number of effects on the different stages of the outcome. First of all, picking a motion that is not representative may result in poor outcome of the classification of other motions, e.g., there might be larger gaps between two chosen pose labels which in the primer are directly subsequent (e.g., lower part of match window in Figure 13). Eventually, this may lead to less coverage of the used data set than expected. Second, poor choice of primer pose labels (or specification of too many different pose labels) will result in overlap of classes for relevant motions (e.g., top of Figure 13). Such overlaps are generally not a problem, since there are in fact transitions that may correctly fall into more than one category (e.g., natural transitions between two activities which may also occur as isolated poses). However, selection of well-suited centroids as representative poses in the labeling phase helps avoid too much overlap. Consequently, the EventFlow tool may be applied to other motions in the data pool for the identification of synthesis candidates that might fit the primer. Exploring these candidates further with the search option helps rule out unsuitable subsequences such that results suffice given constraints.

Figure 13 illustrates how EventFlow supports the synthesis idea by its search option of subsequences. The displayed data are based on labeling and learning of punches and kicks (as described in the Usage Scenario 'Visual-Interactive Learning of a MoCap Classifier'). As can be seen, the queery for the synthesis application is simply a kick pose ('KickRightFront', seen in the upper right part of the image). The matching results show that, besides motions from several kick trials, also other subsequences such as stepping forward or kicking the other leg might fit in a synthesized sequence. However, the search option allows for further refining the queery by constraining the remaining options.

Discussion and Future Work

In the course of this work, we identified several issues that allow design alternatives and possible extensions in future works. Currently, the labeling interface accepts a single MoCap sequence to assign labels. A possible extension would be to accept a set of MoCap sequences. As an example it would be interesting to *visually compare* various repetitions of a particular motion - or instances of the same scripted motion by different actors - and the micro variations of respective user-defined labels. When optimizing key poses in our 2D projection, we assume that all frames belonging to one class are arranged within a circular orbit around the corresponding key pose. Although this assumption worked well for the examples presented in this work, a circular arrangement of poses within one class is not the general case. One strand of future research will be to adapt and refine the 2D mapping of the high dimensional data on the basis of the given labels. Another matter that could be looked into as an objective for future research is the creation of a hierarchy of labels in order to identify latent relations in the data. As of now, multiple labels for the same frame are not explicitly considered even though incidents of overlapping labels have been observed. A hierarchical structure of labels (L1: Exercises, L2: squat, push up, ...) would reflect the relation between representative poses and motion classes that have been identified before. Moreover, labels associated with metadata or descriptions of the motion performed, with the actor (gender, age), or with medical information could enable pattern detection. As indicated in the second usage scenario the synthesis of MoCap data is an interesting but challenging field of research. One additional benefit could be to directly combine our approach with visual-interactive synthesis techniques. We assume that the effectiveness of such a complex analytical and computational workflow heavily depends on the requirements posed by the involved user group.

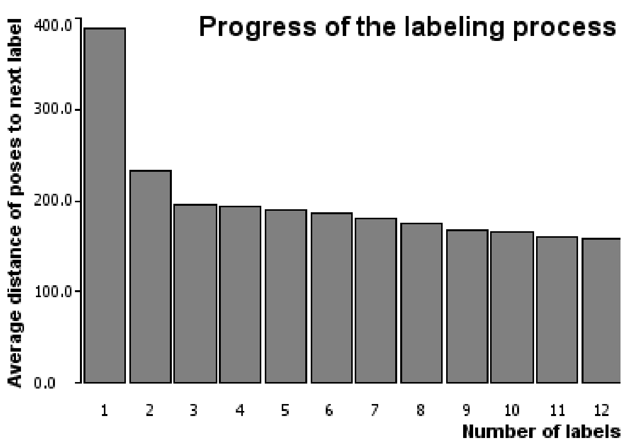

Figure 10. *Progress of the labeling process calculated by the coverage model. The remaining average distance of poses to the label set decreases.*

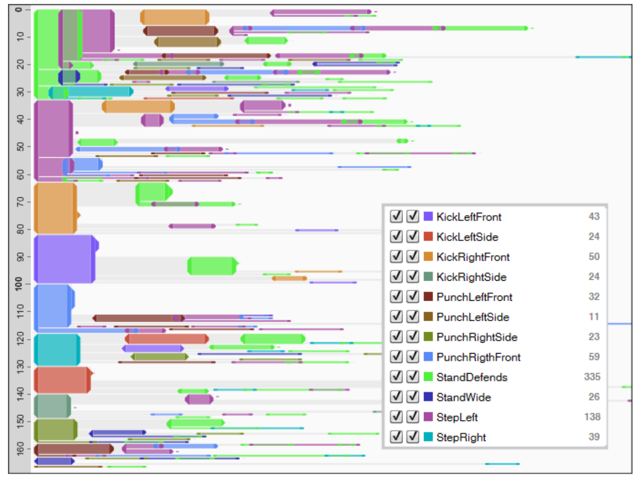

Figure 11. *Overview of the MoCap collection after 12 user-defined labels. The abstraction to sequences of event intervals provides a compact visualization of a large number of sequences using the EventFlow [41] tool. Sorting of sequences with respect to the labels additionally structures the data.*

Conclusion

We presented a visual-interactive two-step workflow for the semi-supervised labeling of large collections of human motion capture (MoCap) data. We described the core steps in a conceptual section and suggested requirements to approaches implementing such a workflow. In the first step of our implementation of the workflow, users can execute and compare multiple results of unsupervised and supervised algorithms for the segmentation of the data. A visual-interactive interface allows the definition of key poses in MoCap data which are subsequently used as labels. In a projection-based interface users can compare different labels with the multivariate input for the identification of overlaps and gaps to improve the labeling result. In the second step of the approach, users can explore the possibly large search space on the basis of user-defined set of labels. We present different mechanisms for the selection of new MoCap candidates or labeling, thus closing the feedback loop of the semi-supervised labeling approach. Two usage scenarios demonstrated the effectiveness and efficiency of our approach. We conducted an active learning strategy to identify a small set of labels which were used to characterize and abstract a large MoCap data collection. As a result, users can easily identify various classes of subsequences of existing MoCap data, e.g., as a prerequisite for effective exploratory search in MoCap data. In a second usage scenario addresses the goal of MoCap synthesis. Based on labeling an existing MoCap

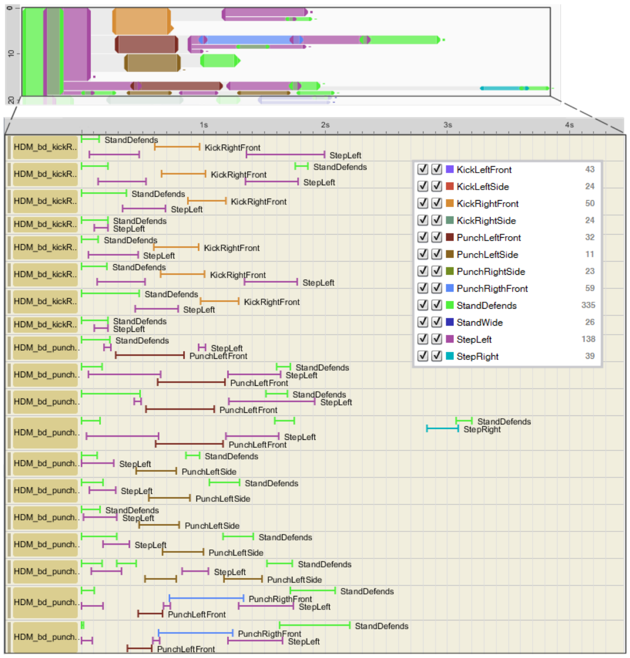

Figure 12. *Selection of sequences all starting with a defense pose (green) for a detailed analysis with EventFlow [41]. In many cases the defense pose is followed by a step left, as well as kicking or punching movements.*

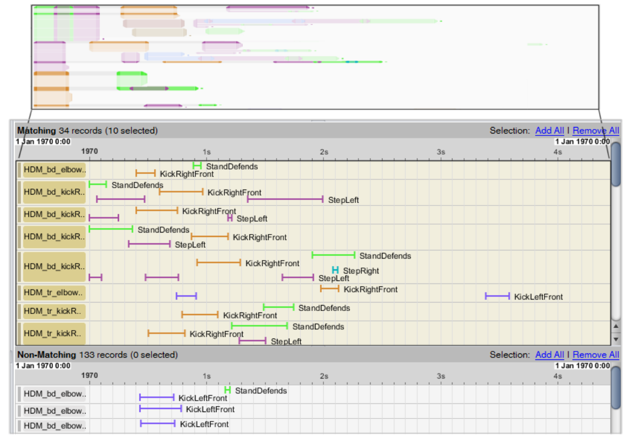

Figure 13. *Selection of sequences matching the information need for MoCap synthesis (associated with label KickRightFront). At the bottom of the EventFlow [41] tool non-matching sequences are filtered out (KickLeftFront).*

sequence, meaningful subsequences in a large MoCap collection are identified to carry out synthesis tasks.

References

[1] D. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann, Eds., *Mastering the Information Age: Solving Problems with Visual Analytics*. VisMaster, http://www.vismaster.eu/book/, 2010.

[2] J. Bernard, N. Wilhelm, B. Krüger, T. May, T. Schreck, and J. Kohlhammer, "MotionExplorer: Exploratory search in human motion capture data based on hierarchical aggregation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2257–2266, 2013.

[3] S. Jang, N. Elmqvist, and K. Ramani, "Motionflow: Visual abstraction and aggregation of sequential patterns in human motion tracking data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 21–30, 2016.

[4] N. Wilhelm, A. Vögele, R. Zsoldos, T. Licka, B. Krüger, and J. Bernard, "Furyexplorer: visual-interactive exploration of horse motion capture data," in *SPIE, Visualization and Data Analysis (VDA)*, 2015, pp. 93 970F–93 970F–15.

[5] J. Bernard, E. Dobermann, M. Bögl, M. Röhlig, A. Vögele, and J. Kohlhammer, "Visual-Interactive Segmentation of Multivariate Time Series," in *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2016.

[6] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber, "Documentation mocap database hdm05," Universität Bonn, Tech. Rep. CG-2007-2, 2007.

[7] CMU, "Carnegie-mellon mocap database," http://mocap.cs.cmu.edu, 2016.

[8] L. Kovar and M. Gleicher, "Automated extraction and parameterization of motions in large data sets," *ACM Transactions on Graphics (TOG)*, vol. 23, no. 3, pp. 559–568, 2004. [Online]. Available: http://doi.acm.org/10.1145/1015706.1015760

[9] J. Chai and J. K. Hodgins, "Performance animation from low-dimensional control signals," in *ACM SIGGRAPH SIGGRAPH*. New York, NY, USA: ACM, 2005, pp. 686–696. [Online]. Available: http://doi.acm.org/10.1145/1186822.1073248

[10] M. Müller, T. Röder, and M. Clausen, "Efficient content-based retrieval of motion capture data," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 677–685, 2005.

[11] M. Müller, A. Baak, and H.-P. Seidel, "Efficient and robust annotation of motion capture data," in *ACM SIGGRAPH/EG Symposium on Computer Animation (SCA)*. New York, NY, USA: ACM, 2009, pp. 17–26. [Online]. Available: http://doi.acm.org/10.1145/1599470.1599473

[12] K. Forbes and E. Fiume, "An efficient search algorithm for motion data using weighted pca," in *ACM SIGGRAPH/EG Symposium on Computer Animation (SCA)*. New York, NY, USA: ACM, 2005, pp. 67–76. [Online]. Available: http://doi.acm.org/10.1145/1073368.1073377

[13] B. Krüger, J. Tautges, A. Weber, and A. Zinke, "Fast local and global similarity searches in large motion capture databases," in *ACM SIGGRAPH/EG Symp. on Comp. Animation (SCA)*. Eurographics Association, 2010, pp. 1–10.

[14] J. Tautges, A. Zinke, B. Krüger, J. Baumann, A. Weber, T. Helten, M. Müller, H.-P. Seidel, and B. Eberhardt, "Motion reconstruction using sparse accelerometer data," *ACM Trans. Graph.*, vol. 30, no. 3, pp. 18:1–18:12, 2011. [Online]. Available: http://doi.acm.org/10.1145/1966394.1966397

[15] Q. Riaz, G. Tao, B. Krüger, and A. Weber, "Motion reconstruction using very few accelerometers and ground contacts," *Graph. Models*, vol. 79, no. C, pp. 23–38, 2015. [Online]. Available: http://dx.doi.org/10.1016/j.gmod.2015.04.001

[16] O. Arikan, D. A. Forsyth, and J. F. O'Brien, "Motion synthesis from annotations," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 402–408, 2003. [Online]. Available: http://doi.acm.org/10.1145/882262.882284

[17] M. Müller and T. Röder, "Motion templates for automatic

classification and retrieval of motion capture data," in *SIGGRAPH/EG Symposium on Computer Animation (SCA)*. Aire-la-Ville, Switzerland, Switzerland: The Eurographics Association, 2006, pp. 137–146.

[18] F. Zhou, F. D. la Torre, and J. K. Hodgins, "Hierarchical aligned cluster analysis for temporal clustering of human motion," *IEEE Trans. on PAMI*, 2013.

[19] A. Vögele, B. Krüger, and R. Klein, "Efficient unsupervised temporal segmentation of human motion," in *ACM SIGGRAPH/EG Symposium on Computer Animation (SCA)*. Aire-la-Ville, Switzerland, Switzerland: The Eurographics Association, 2014, pp. 167–176.

[20] K. Stollenwerk, A. Vögele, B. Krüger, A. Hinkenjann, and R. Klein, "Automatic temporal segmentation of articulated hand motion," *Lecture Notes in Computer Science, Comp. Science and Its Applications (ICCSA)*, 2016, accepted.

[21] D. Bouchard and N. I. Badler, "Segmenting motion capture data using a qualitative analysis," in *ACM SIGGRAPH Conference on Motion in Games (MIG)*. New York, NY, USA: ACM, 2015, pp. 23–30. [Online]. Available: http://doi.acm.org/10.1145/2822013.2822039

[22] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of Time-Oriented Data*, 1st ed., ser. Human-Computer Interaction. Springer Verlag, 2011.

[23] N. Andrienko and G. Andrienko, *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[24] P. McLachlan, T. Munzner, E. Koutsofios, and S. North, "Liverac: Interactive visual exploration of system management time-series data," in *SIGCHI Conference on Human Factors in Computing Systems (CHI)*. New York, NY, USA: ACM, 2008, pp. 1483–1492.

[25] M. Krstajic, E. Bertini, and D. A. Keim, "Cloudlines: Compact display of event episodes in multiple time-series." *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2432–2439, 2011.

[26] A. Sarkar, M. Spott, A. F. Blackwell, and M. Jamnik, "Visual discovery and model-driven explanation of time series patterns," in *IEEE Symp. on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 2016, pp. 78–86.

[27] J. Bernard, N. Wilhelm, M. Scherer, T. May, and T. Schreck, "TimeSeriesPaths: Projection-Based Explorative Analysis of Multivariate Time Series Data," *Journal of WSCG*, vol. 20, no. 2, pp. 97–106, 2012.

[28] M. C. Hao, M. Marwah, H. Janetzko, U. Dayal, D. A. Keim, D. Patnaik, N. Ramakrishnan, and R. K. Sharma, "Visual exploration of frequent patterns in multivariate time series," *Information Visualization*, vol. 11, no. 1, pp. 71–83, 2012. [Online]. Available: http://dx.doi.org/10.1177/1473871611430769

[29] J. Bernard, M. Steiger, S. Mittelstädt, S. Thum, D. Keim, and J. Kohlhammer, "A survey and task-based quality assessment of static 2d colormaps," vol. 9397. SPIE Press, 2015, pp. 93 970M–93 970M–16.

[30] S. Spurlock, R. Chang, X. Wang, G. Arceneaux, D. F. Keefe, and R. Souvenir, "Combining automated and interactive visual analysis of biomechanical motion data," in *Advances in Visual Computing - 6th International Symposium, ISVC 2010, Las Vegas, NV, USA, November 29 - December 1, 2010, Proceedings, Part II*, 2010, pp. 564–573.

[31] J. Bernard, "Exploratory search in time-oriented primary data," Ph.D. dissertation, Technische Universität, Darmstadt, December 2015. [Online]. Available: http://tuprints.ulb.tu-darmstadt.de/5173/

[32] Y. Hu, S. Wu, S. Xia, J. Fu, and W. Chen, "Motion track: Visualizing variations of human motion data," in *Pacific Visualization Symp. (PacificVis)*. IEEE, 2010, pp. 153–160.

[33] B. Alsallakh, M. Bögl, T. Gschwandtner, S. Miksch, B. Esmael, A. Arnaout, G. Thonhauser, and P. Zöllner, "A visual analytics approach to segmenting and labeling multivariate time series data," in *EuroVis Workshop on Visual Analytics (EuroVA)*, The Eurographics Association. The Eurographics Association, 2014, pp. 31–35.

[34] M. Röhlig, M. Luboschik, H. Schuman, M. Bögl, B. Alsallakh, and S. Miksch, "Analyzing parameter influence on time-series segmentation and labeling," in *IEEE Vis. Conference (VIS) (Poster Paper)*, 2014.

[35] O. Alemi, P. Pasquier, and C. Shaw, "Mova: Interactive movement analytics platform," in *Int. Workshop on Movement and Computing (MOCO)*. New York, NY, USA: ACM, 2014, pp. 37:37–37:42. [Online]. Available: http://doi.acm.org/10.1145/2617995.2618002

[36] M. Kapadia, I.-k. Chiang, T. Thomas, N. I. Badler, and J. T. Kider, Jr., "Efficient motion retrieval in large motion databases," in *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D)*. New York, NY, USA: ACM, 2013, pp. 19–28. [Online]. Available: http://doi.acm.org/10.1145/2448196.2448199

[37] T. Saito, H. N. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, and T. Kaseda, "Two-tone pseudo coloring: Compact visualization for one-dimensional data," in *IEEE Information Visualization (INFOVIS)*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 23–. [Online]. Available: http://dx.doi.org/10.1109/INFOVIS.2005.35

[38] Y. Mao, J. Dillon, and G. Lebanon, "Sequential document visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1208–1215, 2007. [Online]. Available: http://dx.doi.org/10.1109/TVCG.2007.70592

[39] Y. Hu, S. Wu, S. Xia, J. Fu, and W. Chen, "Motion track: Visualizing variations of human motion data," in *IEEE Pacific Visualization Symposium (PacificVis)*, 2010, pp. 153–160.

[40] M. O. Ward and Z. Guo, "Visual exploration of time-series data with shape space projections," in *EG/IEEE VGTC Conference on Visualization (EuroVis)*. Chichester, UK: The Eurographs Association, 2011, pp. 701–710.

[41] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman, "Temporal event sequence simplification," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2227–2236, 2013.

[42] F. Du, B. Shneiderman, C. Plaisant, S. Malik, and A. Perer, "Coping with volume and variety in temporal event sequences: Strategies for sharpening analytic focus," *IEEE Transactions on Visualization and Computer Graphics*, vol. PP, no. 99, pp. 1–1, 2016.

[43] B. Settles, "Active learning literature survey," University of Wisconsin–Madison, Computer Sciences Technical Report 1648, 2010.