

Visual-Interactive Segmentation of Multivariate Time Series

Jürgen Bernard^{1,2}, Eduard Dobermann², Markus Bögl³, Martin Röhlig⁴, Anna Vögele⁵, and Jörn Kohlhammer^{1,2}

¹TU Darmstadt, Germany; ²Fraunhofer IGD, Germany

³Vienna University of Technology, Austria; ⁴University of Rostock, Germany; ⁵University of Bonn, Germany

Abstract

Choosing appropriate time series segmentation algorithms and relevant parameter values is a challenging problem. In order to choose meaningful candidates it is important that different segmentation results are comparable. We propose a Visual Analytics (VA) approach to address these challenges in the scope of human motion capture data, a special type of multivariate time series data. In our prototype, users can interactively select from a rich set of segmentation algorithm candidates. In an overview visualization, the results of these segmentations can be compared and adjusted with regard to visualizations of raw data. A similarity-preserving colormap further facilitates visual comparison and labeling of segments. We present our prototype and demonstrate how it can ease the choice of winning candidates from a set of results for the segmentation of human motion capture data.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

1. Introduction

Segmenting time-oriented data is a highly relevant analytical task in many application domains. Prominent examples are interpreting EEG data in medicine, analyzing movement exercises in sports, and investigating causes and effects of weather phenomena in earth observation. The segmentation of time series follows the general idea of dividing complex information into smaller and more manageable units. These ‘atomic’ temporal units are crucial for a number of applications such as detection of specific sub-sequences and analysis of repetition and variation. Downstream tasks include identifying recurring patterns, detecting abnormal behavior, or simply making more sense of long time series. Visual analytics (VA) has become an indispensable means for supporting users in situations where multivariate time series data occur along with the above tasks, but has only been sparsely applied to the segmentation of time series.

The segmentation of time series is non-trivial. The suitability of automated segmentation algorithms heavily depends on the data and the application task, and thus cannot be determined a priori. Similarly, most algorithms provide at least one parameter, further complicating the search for adequate segmentations. Finally, time series are a complex data type, allowing various characterizations of both the temporal domain and the value domain [AMST11]. Here, we focus on multivariate time series, a data type that can describe phenomena of multiple variables developing over time by taking the example of human motion capture (MoCap) data. Recent advances in VA indicate a great potential for support-

ing users in segmenting motion data. Examples include the classification of different gaits of horses [WVZ*15], comparing retrieved motion subsequences [BWK*13], or interactively organizing unstructured motion data [JER16]. However, the segmentation of human MoCap data has received only marginal attention in VA research to date.

The following challenges have to be addressed for a successful approach. First, the aforementioned VA techniques indicate that different classes of segmentation algorithms have to be integrated in a unified approach **Ch₁**. More precisely, we select classification, clustering, and binning as a valuable starting point. Second, to stand out from the great variety of blackbox approaches, these algorithms need to be selectable by visual-interactive means **Ch₂**. Third, it is important to consider multiple candidates **Ch₃**. Visually comparing segmentation results stemming from different algorithms and parametrizations eases the selection of winning candidates in the first place. Fourth, the ‘raw’ multivariate time series data has to be visualized in order to allow analytical reasoning for local differences in segmentation results **Ch₄**. Finally, any label of any produced time series segment needs to be clearly visually distinguishable **Ch₅**. By default, if users do not explicitly assign labels to segments, this process needs to be supported in an automated way. This is a prerequisite for segmentation results being comparable.

To address these requirements, we introduce a novel VA approach for supporting users in the segmentation of multivariate time series, combining the following characteristics:

Analytical: Segmenting and labeling the raw time series




Figure 1: Overview of the visual-interactive prototype for the segmentation of human motion capture data. Users can select segmentation algorithms and parameters at the left. A similarity-preserving color mapping (top right) enables the visual comparison of segment labels. The view at the center enables the visual comparison of multiple segmentation results at a glance.

data based on a rich set of segmentation algorithms, including classification, clustering, and binning; facilitating the configuration of algorithm-specific parameters to obtain meaningful results.

Visual: Presenting a novel visual design that integrates (i) an overview of all segmented time series, (ii) detail views reflecting certain aspects of segments back onto the raw data, and (iii) similarity-preserving color coding for automated labeling and visual comparison of segments.

Interactive: Providing interaction techniques facilitating the analysis of individual labeled segments in detail, the exploration of temporal patterns, and the comparison of labeled segments corresponding to different algorithm configurations.

2. Related Work

Recently, a conceptual framework and a conceptual VA approach for segmenting and labeling of multivariate time series have been proposed [ABG*14, GSB*15]. Röhlig et al. [RLK*15] present a first VA approach for supporting activity recognition based on both of these conceptual designs. They focus on supporting users in understanding the complex relationships among parameters, recognized activities, and associated accuracies. In contrast, we additionally aim at the interactive selection and steering of segmentation algorithms along with complementary visual analysis of raw measurements. In the following, we give a brief overview of the related work in time series segmentation algorithms, as well as the visual analysis and comparison of segmented time series and of temporal patterns.

2.1. Segmentation of Time Series

Our approach is based on a great variety of existing algorithms for segmenting time series. Temporal segmentation has a background in a number of scientific areas ranging from data mining [Fea06] to audio and speech processing [PM14], and behavioral pattern recognition [XM07].

One popular strategy for motion segmentation is based on clustering the poses in multivariate series. This can be done by either extracting motion motifs [BCvdPP08] or clustering pose information temporally [LMGCG12] or performing exploratory searches [BWK*13]. Alternative approaches are based on training annotated segments and classification by pre-computed templates, e.g., Lv et al. [LN06] and Müller et al. [MBS09]. There is a number of methods designed for fully automatic motion segmentation. Barbič et al. [BPF*04] propose a PCA-based method focused on detecting activities. The works of Zhou et al. [ZITH13] use (hierarchically) aligned cluster analysis (H)ACA for temporal clustering of poses. Vögele et al. [VKK14] present a graph-based approach for segmentation of activities and detection of primitive motion units by evaluation of self-similarity within MoCap data. Automatic segmentation is also achieved as a by-product of synthesis tasks, e.g. motion concatenation [KGP02] and motion parameterization [KG04] or Min and Chai's Motion Graphs++ [MC12]. Finally, functional constraints for segmenting time series can be pre-defined, e.g., by rule-based or binning-based approaches. Examples of this type of approach are found in the works of Keogh et al. [KCHP04].

2.2. Visual Analysis of Segmented Time Series

The visual-interactive analysis of multivariate time series data allows users to identify temporal patterns and thus, supports subsequent segmentation of the data. Several techniques for visually analyzing segmented time series have been proposed, e.g., TimeSeriesPaths [BWS*12] and MotionExplorer [BWK*13]. Moreover, different approaches based on projections and path metaphors have been applied [WG11, SKK04, BWS*12]. In LiveRAC [MMKN08], one-dimensional categorical color coding is used for interactive visual exploration of system management time series data.

Generally, segmented and labeled time series can be referred to as categorical data over time, cf. [WPS*09]. Vi-

sually comparing such data requires suitable visual encoding to overview multiple time series at once. Typically, the different categories are mapped to colors or symbols. The challenge of visualizing this kind of data was tackled in different areas. Examples include, assessing genome sequences [HVB*12], analyzing streams, clusters, and patterns of events [SWM*11, FMK12], exploring large spatio-temporal datasets [dNO15], and searching and comparing color-coded subsequence retrieval results [BWK*13, BSM*15].

In summary, several individual techniques exist for either segmenting, labeling, or visually analyzing multivariate time series data. Yet, none of these techniques supports an integrated solution for visualizing and comparing hundreds of time series segmentation results stemming from different algorithms as well as varying parameter settings. We aim at such an integrated solution covering all aspects: (i) selecting and parameterizing segmentation algorithms, (ii) visually comparing multiple segmentation results, and (iii) interactively analyzing labeled segments, raw measurements, and associated uncertainties over time.

3. The Visual-Interactive System

We present a VA approach for the segmentation of time series, using the example of human MoCap data. Users can select from a variety of segmentation algorithms and steer respective parameter values. An overview visualization allows the comparison of multiple segmentation results on the basis of color-coded labels. Two visualizations for raw multivariate time series data support users with the interpretation of segments. We show the usefulness of the prototype in a usage scenario enclosed in the system description, motivated by an expert interested in the visual comparison of the Decorate classifier, the RandomForest classifier, and the SimpleKMeans clusterer for a *rotating-arms motion*.

3.1. Data Abstraction

We apply human MoCap data from a publicly available database [MRC*07] containing over 70 motion classes, each with up to 50 repetitions performed by different actors. In our scenario, we choose a *rotating-arms motion*. The domain expert provided a labeling of the data dividing the circular motion into six segments ('UpToPeak', 'UpToMid', 'MidToDown', 'DownToMid', 'MidToUp', and 'UpToPeak'). Every human pose in the data (a time frame) is characterized by series of 3D markers in absolute scale. According to Krüger et al. [KTWZ10], we reduce the number of 3D markers to 16, leading to a static *pose vector* with 48 dimensions. In order to cover the temporal domain, we create pose vector triples by combining $[t_{i-k}, t_i, t_{i+k}]$, where i is the current frame and k a steerable number of frames. This results in a 144-dimensional feature vector for each time stamp.

3.2. Visualization of Raw Multivariate Time Series

Part of the challenge in this work is the visual representation of raw multivariate time series data allowing analytical reasoning about segmentation results and labels **Ch4**. Following a generalizable approach, we use a bundled line chart

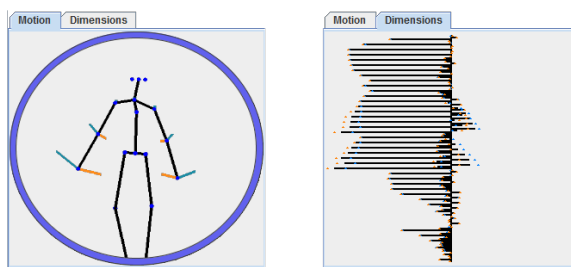


Figure 2: Visualizations of multivariate input data for single time stamps. A domain-specific and a general technique. Blue and orange colors represent the temporal domain.

visualization including one chart for every dimension (see Figure 1, top). To avoid overplotting, we suggest to filter dimensions with low variances, given the assumption that the degree of variation has high implications on the time series segmentation process. In addition, the label information of previously annotated time series can be visualized in the notion of colored event sequences (cf. Section 3.4). A slider control enables users to interact with the raw time series data. To further enhance this functionality, we provide two visual representations of single time stamps (see Figure 2). Users can choose between a domain-specific (left) and general visualization (right). The domain-specific visualization depicts the semantics of human poses with a stick figure icon inspired by the MotionExplorer system [BWK*13], but extends the design with colored lines indicating the temporal domain (blue = frames at t_{i-k}, \dots, t_{i-1} ; orange = frames at t_{i+1}, \dots, t_{i+k}). A second visualization can be used in a more general application setting. Horizontal bars show the value domains of all dimensions of the multivariate states. Per default, zero values are aligned at the 50% position of the horizontal axis.

3.3. Algorithm and Parameter Choice

To resolve **Ch2**, users can select algorithms and algorithm parameters in the left view (see Figure 1). We distinguish between classification, clustering, and binning approaches (cf. Section 2.1). At the moment, we support the selection of algorithms from the WEKA machine learning package [HFH*09], as well as our own algorithms **Ch1**. In general the interface for the algorithm choice can also be implemented for accessing alternative libraries. The context menu for the algorithm and parameter choice is adopted from the WEKA library. Taking the example of Figure 1, a SimpleKMeans clustering algorithm was chosen. After its execution, the result of the algorithm was automatically added to the result set (center of the system) and other coordinated views.

3.4. Visualization of Segmentation Results

Figure 3 shows three segmentation results in detail. Similar to the visualization of annotated labels in raw data (cf. Section 3.2), we choose a colored sequence bar for the label visualization. In addition, users can switch to linechart visualizations to analyze the probability distribution of segments over time. In this way, we propagate the uncertainty information of the algorithms to the visual interface. We take a closer

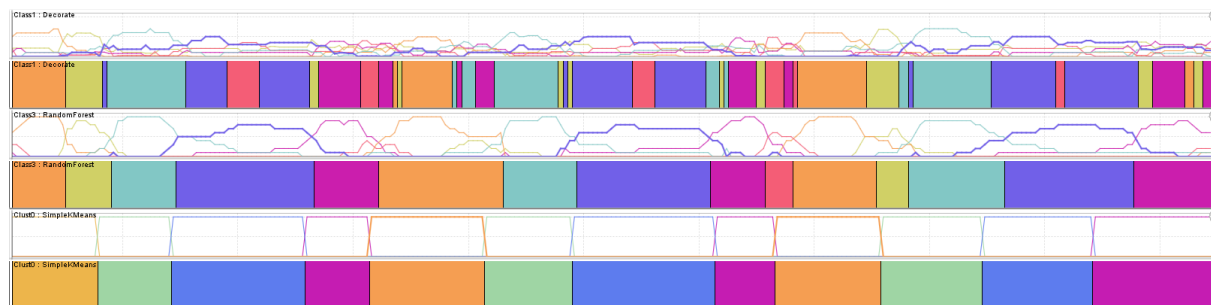


Figure 3: Visual comparison of three segmentation results including detail information about the distribution of segment probabilities. At the top the result of the Decorate classifier is shown (artificial examples: 1, ensemble size: 10, iterations: 10). It differs strongly from the result of the RandomForest classifier (number of trees: 10) in the middle row and the SimpleKMeans clustering algorithm ($k:5$) at the bottom, which produce similar outputs. Line charts indicate the probabilities of segments.

look at the three segmentation results. The result at the top has a rather inconsistent label distribution. Users can interpret this analyzing the probability distribution which is noisy in this case. The two results below are quite similar. Interestingly, one is the result of a classifier and the other one of a cluster algorithm. The three detailed visualizations of segmentation results are located at the center of the system (cf. Figure 1). In addition, we provide a pixel-based visualization at the bottom for the comparison of multiple result instances **Ch₃**. Using the example of Figure 1, users can visually compare the output of ten or even more segmentation results. Some of them are similar, such as the NaiveBayes classifier, the RandomForest classifier, and the LogiBoost classifier. However, many results pose different characteristics, such as different levels of abstractions. The selection of a result triggers all coordinated views of the prototype (cf. Section 3.5).

3.5. Making Segmentation Results Comparable

The usefulness of the approach comes with the ability to compare multiple segmentations intuitively **Ch₅**. Hence, we provide an automated method for coloring segments of any result in a meaningful way. In this regard, we require similar segments to have similar colors while dissimilar segments should contain different colors. By that, the visual representation of labels can be used for analytical reasoning without user annotation. We use dimension reduction to project human poses into the 2D space and apply a 2D colormap [BSM*15] to encode high-dimensional data with color in a similarity-preserving way (see the upper right view in Figure 1). Moreover, in two views at the right of the prototype, we provide information of every segmentation result and in a compact way. The selection of a single result shows details of individual segments at the bottom right (see Figure 1). Using the example of MoCap data, users can analyze the information of human poses for every segment with a visual pose bundle metaphor. The selection of single segment labels automatically highlights label occurrences in coordinated views. In Figure 1 an orange segment is selected. Black bars highlight the selection at the center of the prototype, while a cluster visualization highlights all human poses in the projection-based visualization at the top right.

4. Limitations and Future Work

While we addressed all proposed challenges, there are some limitations in the current approach. The chosen feature spaces in the given scenario are designed to allow for suitable representation of pose information given by MoCap marker position. However, other features such as angles or different ways of stacking information temporally are available and may also be suited as representatives. In addition, feature design and parameter contribution will vary with the choice of different capturing modalities, such as accelerometer recordings or depth data. While we support the choice of algorithm parameters, users are currently left alone with the implications of any given parameter value. Possible future extensions address the access to the parameter space, and to the expected output of a particular segmentation algorithm. According to [GSB*15] there are four types of uncertainties involved in segmenting and labeling of time series data. In our approach, we currently communicate the probabilities of the decision on a specific label, if the chosen algorithm provides this number. In future work we need to consider the other types of uncertainty, like the composition of alternative results, temporal uncertainty, and the uncertainty caused by different algorithms and parametrization.

5. Conclusions

In this work, we presented a VA approach for the segmentation of multivariate time series by example of human motion capture data. A similarity-preserving color coding enables the visual comparison of segment labels within and between segmentation results. Different visualizations for the raw time series data support the association of segmentation results with the semantics of the data and vice versa.

References

- [ABG*14] ALSALLAKH B., BÖGL M., GSCHWANDTNER T., MIKSCH S., ESMAEL B., ARNAOUT A., THONHAUSER G., ZÖLLNER P.: A visual analytics approach to segmenting and labeling multivariate time series data. In *EuroVis Workshop on Visual Analytics (EuroVA)* (2014), Eurographics, pp. 31–35. doi:10.2312/eurova.20141142. 2
- [AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: *Visualization of Time-Oriented Data*. Human-

- Computer Interaction. Springer Verlag, 2011. doi:10.1007/978-0-85729-079-3. 1
- [BCvdPP08] BEAUDOIN P., COROS S., VAN DE PANNE M., POULIN P.: Motion-motif graphs. In *Proc. of ACM SIGGRAPH/SCA* (2008), pp. 117–126. 2
- [BPF*04] BARBIČ J., PAN J.-Y., FALOUTSOS C., HODGINS J., POLLARD N.: Segmenting motion capture data into distinct behaviors. In *Proc. of Graphics Interface* (2004), pp. 185–194. 2
- [BSM*15] BERNARD J., STEIGER M., MITTELSTÄDT S., THUM S., KEIM D., KOHLHAMMER J.: A survey and task-based quality assessment of static 2D colormaps. In *SPIE, Visualization and Data Analysis (VDA)* (2015), vol. 9397. doi:10.1117/12.2079841. 3, 4
- [BWK*13] BERNARD J., WILHELM N., KRUGER B., MAY T., SCHRECK T., KOHLHAMMER J.: MotionExplorer: Exploratory search in human motion capture data based on hierarchical aggregation. *IEEE TVCG* 19, 12 (2013), 2257–2266. doi:10.1109/TVCG.2013.178. 1, 2, 3
- [BWS*12] BERNARD J., WILHELM N., SCHERER M., MAY T., SCHRECK T.: Timeseriespaths: Projection-based explorative analysis of multivariate time series data. *Journal of WSCG* 20, 2 (2012), 97–106. 2
- [dNO15] DO NASCIMENTO OLIVEIRA G.: *Ordered stacks of time series for exploratory analysis of large spatio-temporal datasets*. PhD thesis, Universidade Federal do Rio Grande do Sul, 2015. 3
- [Fea06] FEARNHEAD P.: Exact and efficient bayesian inference for multiple changepoint problems. *Statistics and Comp.* 16, 2 (2006), 203–213. doi:10.1007/s11222-006-8450-8. 2
- [FMK12] FISCHER F., MANSMANN F., KEIM D. A.: Real-time visual analytics for event data streams. In *Symposium on Applied Computing* (2012), ACM, pp. 801–806. doi:10.1145/2245276.2245432. 3
- [GSB*15] GSCHWANDTNER T., SCHUMANN H., BERNARD J., MAY T., BÖGL M., MIKSCH S., KOHLHAMMER J., RÖHLIG M., ALSALLAKH B.: Enhancing time series segmentation and labeling through the knowledge generation model. In *Poster Proceedings of EuroVis* (2015), Eurographics, p. 3. 2, 4
- [HFH*09] HALL M., FRANK E., HOLMES G., PFAHRINGER B., REUTEMANN P., WITTEN I. H.: The weka data mining software: An update. *SIGKDD* 11, 1 (2009), 10–18. 3
- [HVB*12] HEINRICH J., VEHLow C., BATTKE F., JÄGER G., WEISKOPF D., NIESELT K.: iHAT: interactive hierarchical aggregation table for genetic association data. *BMC Bioinform.* 13, 8 (2012), 2. doi:10.1109/BioVis.2011.6094049. 3
- [JER16] JANG S., ELMQVIST N., RAMANI K.: Motionflow: Visual abstraction and aggregation of sequential patterns in human motion tracking data. *IEEE TVCG* 22, 1 (2016), 21–30. doi:10.1109/TVCG.2015.2468292. 1
- [KCHP04] KEOGH E., CHU S., HART D., PAZZANI M.: Segmenting time series: A survey and novel approach. In *Data Mining In Time Series Databases*, vol. 57 of *Series in Machine Perception and Artificial Intelligence*. World Scientific Publishing Company, 2004, ch. 1, pp. 1–22. doi:10.1142/9789812565402_0001. 2
- [KG04] KOVAR L., GLEICHER M.: Automated extraction and parameterization of motions in large data sets. *ACM Trans. on Graph.* 23, 3 (2004), 559–568. doi:10.1145/1015706.1015760. 2
- [KGP02] KOVAR L., GLEICHER M., PIGHIN F.: Motion graphs. *ACM Trans. on Graph.* 21, 3 (2002), 473–482. doi:10.1145/566654.566605. 2
- [KTWZ10] KRÜGER B., TAUTGES J., WEBER A., ZINKE A.: Fast local and global similarity searches in large motion capture databases. In *SIGGRAPH/EG Symp. on Computer Animation* (2010), Eurographics, pp. 1–10. doi:10.2312/SCA/SCA10/001-010. 3
- [LMGCG12] LÓPEZ-MENDEZ A., GALL J., CASAS J. R., GOOL L. J. V.: Metric learning from poses for temporal clustering of human motion. In *BMVC* (2012), pp. 1–12. 2
- [LN06] LV F., NEVATIA R.: Recognition and segmentation of 3-D human action using hmm and multi-class adaboost. In *Computer Vision (ECCV)*, vol. 3954 of *LNCS*. Springer, 2006, pp. 359–372. doi:10.1007/11744085_28. 2
- [MBS09] MÜLLER M., BAAK A., SEIDEL H.-P.: Efficient and robust annotation of motion capture data. In *SIGGRAPH/EG Symp. on Computer Animation* (2009), ACM, pp. 17–26. doi:10.1145/1599470.1599473. 2
- [MC12] MIN J., CHAI J.: Motion graphs++: A compact generative model for semantic motion analysis and synthesis. *ACM Trans. Graph.* 31, 6 (2012), 153:1–153:12. doi:10.1145/2366145.2366172. 2
- [MMKN08] MCLACHLAN P., MUNZNER T., KOUTSOFIOS E., NORTH S.: Liverac: interactive visual exploration of system management time-series data. In *SIGCHI Conf. on Human Factors in Computing Systems* (New York, NY, USA, 2008), ACM, pp. 1483–1492. doi:10.1145/1357054.1357286. 2
- [MRC*07] MÜLLER M., RÖDER T., CLAUSEN M., EBERHARDT B., KRÜGER B., WEBER A.: *Documentation: mocap database HDM05*. Computer Graphics Technical Report CG-2007-2, Universität Bonn, 2007. Data available at www.mpi-inf.mpg.de/resources/HDM05. 3
- [PM14] PRÄTZLICH T., MÜLLER M.: Frame-level audio segmentation for abridged musical works. In *Proc. of Int. Conf. on Music Information Retrieval (ISMIR)* (Taipei, Taiwan, 2014). 2
- [RLK*15] RÖHLIG M., LUBOSCHIK M., KRÜGER F., KIRSTE T., SCHUMANN H., BÖGL M., ALSALLAKH B., MIKSCH S.: Supporting activity recognition by visual analytics. In *IEEE VAST* (2015), pp. 41–48. doi:10.1109/VAST.2015.7347629. 2
- [SKK04] SAKAMOTO Y., KURIYAMA S., KANEKO T.: Motion map: image-based retrieval and segmentation of motion data. In *SIGGRAPH/EG Symp. on Computer Anim.* (2004), Eurographics, pp. 259–266. doi:10.1145/1028523.1028557. 2
- [SWM*11] SCHAEFER M., WANNER F., MANSMANN F., SCHEIBLE C., STENNETT V., HASSELROT A. T., KEIM D. A.: Visual Pattern Discovery in Timed Event Data. In *Proc. of Conf. on Visualization and Data Analysis* (2011), SPIE. doi:10.1117/12.871870. 3
- [VKK14] VÖGELE A., KRÜGER B., KLEIN R.: Efficient unsupervised temporal segmentation of human motion. In *SIGGRAPH/EG Symp. on Computer Animation* (2014). 2
- [WG11] WARD M. O., GUO Z.: Visual exploration of time-series data with shape space projections. *CGF* 30, 3 (2011), 701–710. doi:10.1111/j.1467-8659.2011.01919.x. 2
- [WPS*09] WANG T. D., PLAISANT C., SHNEIDERMAN B., SPRING N., ROSEMAN D., MARCHAND G., MUKHERJEE V., SMITH M.: Temporal summaries: Supporting temporal categorical searching, aggregation and comparison. *IEEE TVCG* 15, 6 (2009), 1049–1056. doi:10.1109/TVCG.2009.187. 2
- [WVZ*15] WILHELM N., VÖGELE A., ZSOLDOS R., LICKA T., KRÜGER B., BERNARD J.: Furyexplorer: visual-interactive exploration of horse motion capture data. In *SPIE VDA* (2015), pp. 93970F–93970F–15. doi:10.1117/12.2080001. 1
- [XM07] XUAN X., MURPHY K.: Modeling changing dependency structure in multivariate time series. In *Proc. of Int. Conf. on Machine Learning* (2007), ACM, pp. 1055–1062. doi:10.1145/1273496.1273629. 2
- [ZITH13] ZHOU F., LA TORRE F. D., HODGINS J. K.: Hierarchical aligned cluster analysis for temporal clustering of human motion. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 35, 3 (2013), 582–596. doi:10.1109/TPAMI.2012.137. 2