

Combining Automated and Interactive Visual Analysis of Biomechanical Motion Data

Scott Spurlock¹, Remco Chang², Xiaoyu Wang¹, George Arceneaux IV¹,
Daniel F. Keefe³, and Richard Souvenir¹

¹ Department of Computer Science, The University of North Carolina at Charlotte

² Department of Computer Science, Tufts University

³ Department of Computer Science and Engineering, University of Minnesota

Abstract. We present a framework for combining automated and interactive visual analysis techniques for use on high-resolution biomechanical data. Analyzing the complex 3D motion of, e.g., pigs chewing or bats flying, can be enhanced by providing investigators with a multi-view interface that allows interaction across multiple modalities and representations. In this paper, we employ nonlinear dimensionality reduction to automatically learn a low-dimensional representation of the data and hierarchical clustering to learn patterns inherent within the motion segments. Our multi-view framework allows investigators to simultaneously view a low-dimensional embedding, motion segment clustering, and 3D visual representation of the data side-by-side. We describe an application to a dataset containing thousands of frames of high-speed, 3D motion data collected over multiple experimental trials.

1 Introduction

As CT technology becomes more mature, scientists are now able to capture high speed motion of bones and joints at the rates of 250 to 500 frames per second with sub-millimeter accuracy [1]. This new imaging modality has allowed the scientists to examine animal kinematics movement in ways that have not been possible before. However, with such rich information, new methods for analyzing biomechanical motion data become increasingly important.

Some work has already considered exploring this data both visually and interactively [2, 3]. These systems allowed the scientists to examine raw biomechanical data, but with limited automated analyses. Visualizations have been developed to reanimate the motion data by displaying bones moving in space and overlaying additional data, such as instantaneous helical axes computed from the motion [2]. Recently, Keefe et al. developed an interactive system that combined these 3D motion visualizations with complementary 2D visualizations to better capture the higher dimensionality of the data [3]. These two systems have aided scientists in characterizing the data and finding patterns within the animals' behaviors. There are limitations to these systems, however, based on their reliance on user interactions and visual inspection. The sheer volume and complexity of the data may obscure patterns or relationships from manual discovery.

In this paper, we extend the previous work on biomechanical motion analysis to incorporate automated methods for analysis. Specifically, we incorporate dimensionality reduction and clustering techniques to reduce complexity and highlight patterns within the data. We then integrate the result of these unsupervised learning approaches into an interactive tool to enable visual analysis.

While these automated analysis techniques for biomechanical motion are common in the machine learning and computer vision fields, they have rarely been integrated with highly interactive visual analytical systems. By leveraging a blended model of learning approaches with interactive visual analysis, we aim to enable a new style of hybrid investigation [4, 5]. Using this system, an investigator can utilize automated computational methods to identify hidden temporal patterns embedded within the data while retaining full interactive exploration capabilities to analyze the data either in raw or post-processed form. In our system, a low-dimensional graphical representation of the data can be viewed concurrently with clustering output, and side-by-side with relevant video clips. We believe that this synthesis is an advancement over previous systems in which the burden of investigation is solely dependent upon the user, and that by using this hybrid approach the user can identify new patterns more quickly and in a repeatable fashion.

2 Related Work

Much research has focused on analyzing biomechanical motion through 3D visualizations. Most of these 3D visualization systems focus on rendering the spatial changes of different parts of motion data, for example, the trajectory of jaws or the rotation of joints [3]. Using direct manipulations, most of these systems allow the user to control the viewpoint and to focus on movements of specific parts. In recent work, Keefe et al. [2] presented an effective 3D visualization framework for biomechanical motion analysis with interactive visualizations with detailed anatomical features.

The other major aspect of biomechanical motion analysis is the temporal patterns. Visualizing trends in time-varying and multi-variate data has been considered in depth within the information visualization community [6]. However, for visualization of biomechanical motion over time, 3D views of the model are often animated, and additional data attributes are often visualized using color, texture, streamlines, and 3D data glyphs [2, 7, 8]. While these annotated 3D views can be quite powerful, it has been suggested that understanding trends over time through animation may not be the most effective strategy [6]. Our system follows the framework developed by Keefe et al. [3] in that we also utilize a multiple coordinated visualization approach and support analysis using both the 3D model view and 2D information visualizations.




Fig. 1. Example images from the pig chewing motion data set

3 Biomechanical Motion Data

The techniques and example application in this paper are presented with the goal of applying broadly across different types of motion analysis of interest in the biological sciences. The specific datasets driving the framework presented here come from a study at Brown University that made use of marker-based X-ray Reconstruction of Moving Morphology (XROMM)⁴ to capture measurements of the lower jaw movements of miniature swine (Sinclair strain) during mastication.

Although these data describe the motion of just two bones, these bones connect to each other at two joints and also whenever the teeth come into contact with each other; thus, the data serve as an ideal springboard for complex high-dimensional, multi-joint analyses of other biomechanical structures. The chewing motion itself is quite unique among animals, hence the interest from an evolutionary biology perspective in studying the coordinated motion of these bones under different experimental conditions.

From the raw data collected from multiple high-speed fluoroscopic videos captured experimentally, 4×4 transformations can be derived to describe the rigid body transformations (translation and rotation) of the pig's mandible in relation to its skull (see Figure 1). Previous research has identified groupings of particular sequences of frames into segments, which identify related, temporally proximate frames.

4 Automated Analysis of Biomechanical Motion Data

As described in Section 3, the biomechanical motion data is represented as 4×4 transformation matrices describing the positions of the individual bones. This high-dimensional data describes the configuration of the animal at a particular timestep and the time-series describes the biomechanical motion. In order to better understand and visualize the motion and discover any underlying patterns, we employ *dimensionality reduction* and *data clustering*. In this section, we describe the steps for the automated analysis of this data.

⁴ <http://xromm.org>

4.1 Dimensionality Reduction

Most data analysis techniques on high-dimensional points and point sets do not work well. One strategy to overcome this problem is to find an equivalent lower (typically 2 or 3) dimensional representation of the data. Dimensionality reduction is the technique of automatically learning a low-dimensional representation for data. The most well-known techniques are Principal Component Analysis (PCA) [9] and Independent Component Analysis (ICA) [10]. These methods seek to represent data as linear combinations of a small number of basis vectors. However, many data sets, including the transformation matrices of the biomechanical motion data considered in this work, tend to vary in ways which are very poorly approximated by changes in linear basis functions.

Techniques in the field of manifold learning embed high-dimensional data points which lie on a *nonlinear* manifold onto a corresponding lower-dimensional space. There exists a number of automated techniques for learning these low-dimensional embeddings, such as Isomap [11] and LLE [12]. These methods have been used in computer vision and graphics for many applications, including medical image segmentation [13] and light parameter estimation from single images [14]. In this paper, we use the Isomap algorithm, but the general approach could be applied with any of the other nonlinear dimensionality algorithms.

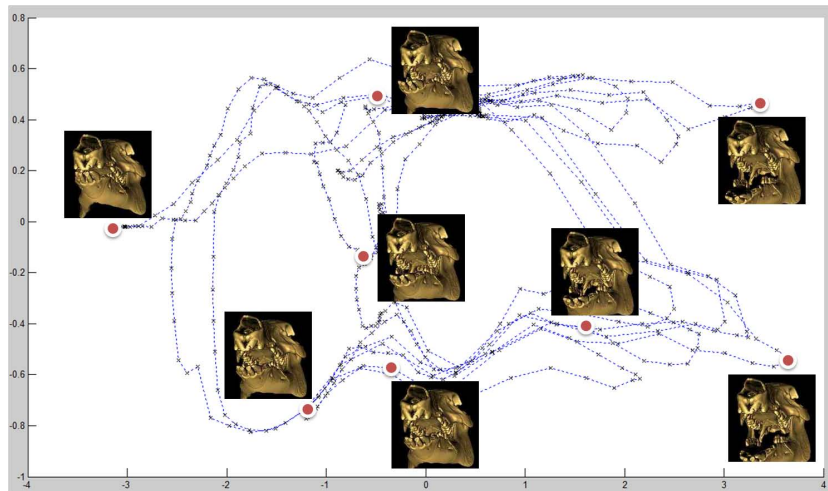


Fig. 2. 2D Isomap embedding of the pig chewing data set. Each point corresponds to an original frame of data. For the indicated points, the corresponding image is shown.

Isomap embeds points in a low-dimensional Euclidean space by preserving the geodesic pair-wise distances of the points in the original space. To estimate the (unknown) geodesic distances, distances are calculated between points in a trusted neighborhood and generalized into geodesic distances using an all-pairs

shortest-path algorithm. With most manifold learning algorithms, discovering which points belong in the trusted neighborhood is a fundamental operation. Typically, the Euclidean distance is used, but in certain cases other distance measures have been shown to lead to a more accurate embedding of the original data [15]. Due to the structure of the transformation matrices (most notably the rotational component) used in our data, we use a distance metric based on exponential matrix mapping described in [16].

Figure 2 shows the Isomap embedding of one of the data sets. Each 2D point represents one of the original frames from the video data and the corresponding images of selected points are shown. One can observe a *perceptual organization* of the original data in the embedded space. The x - and y -axes correlate with the two major components of the motion: opening and closing of the mouth and lateral motion of the mandible, respectively.

This embedding step describes the relationship between each of the frames in the original data set, but doesn't yet describe the known structure within the data. Each of the biomechanical motions considered in the work are comprised of multiple, short, temporal segments that correspond to distinct phases of motion. The shape of these segments in the embedded space can be used compare multiple motion patterns. In order to discover the similarity among segments within a data set, we apply hierarchical clustering.

4.2 Clustering

The process of comparing and grouping temporal segments from the embedding into clusters can be automated using well-known unsupervised data clustering techniques. We use agglomerative clustering to iteratively and hierarchically merge similar segments into clusters. Depending on the goal of the analysis, single-link, complete-link, or average-link can be used. In all three methods, the distance measure [17] applied between segments uses a combination of curve fitting and string matching techniques to calculate the similarity of two trajectories. The metric is scale-, rotation-, translation-, and duration-invariant. Figure 3 shows four segments from the data shown in Figure 2. Our clustering process iteratively groups the segments shown from left to right as being the most similar in shape. From this process we generate a dendrogram, a hierarchical tree structure, to allow the end-user to interactively choose the level of grouping that most meaningful to the investigation.

5 Interactive Visual Analysis of Patterns in Motion Data

In this section, we describe an interactive visual interface for analyzing 3D motion data. This interface displays the raw data using an animated 3D model view, as well as the automated analysis output described in section 4. By integrating these views using a multiple-coordinated-views approach [18], our system provides methods to perform interactive analysis across modalities. Shown together in Figure 4, these views allow the user to explore the 3D motion sequence in space and time, in its embedding space, as well as at a clustering level.




Fig. 3. Separate phases of motion represented as segments from the Isomap embedding. The two on the left were measured to be the closest in shape.

5.1 Exploring the Motion Space

Given our initial dataset of 3D transformations, we utilize an interactive 3D visualization and timeline view to help the user analyze the motion data from both spatial and temporal perspectives.

3D Visualization for Motion Exploration Following related work on biomechanical motion visualization, our system supports motion playback, viewpoint adjustment, and motion sequence comparison. Figure 4 (B) shows the 3D visualization in our system. The user can interactively change the viewpoint of the model and focus on the movement at each time step of the data. Coordinated with the embedding space, the 3D visualization enables the users to explore the precise spatial location of each motion sequence. If any specific motion sequences have been selected, our system will also show a side-by-side playback panel to help the user compare the motion difference between these sequences.

Timeline Visualization for Temporal Patterns To aid in finding temporal patterns, a timeline visualization shows time in relation to segments of motion frames, as well as hierarchical clusters of segments. Figure 4 (E) shows grouping of the motion data based on the predefined segments. These segments are dynamically color-coded to correspond to the cluster into which they have been grouped. As the user traverses the clustering hierarchy, the number of clusters will change accordingly. Higher thresholds results in fewer clusters, but allows the user to identify high-level behavioral patterns in the timeline; whereas a lower threshold produces more clusters, but allows the user to examine the segments in greater detail.

5.2 Analyzing the Embedding Space

The low-dimensional embedding space and the hierarchical clustering provide a foundation for further interactive analysis. When a dataset is first loaded, an overview of the data is displayed using the three coordinated views seen in Figure 4: a small-multiples view, an embedding view, and a 2D plot of each segment whose positions are based on their similarities. These three views have been carefully chosen for their analytical capabilities in analyzing different aspects of a 3D motion sequence.

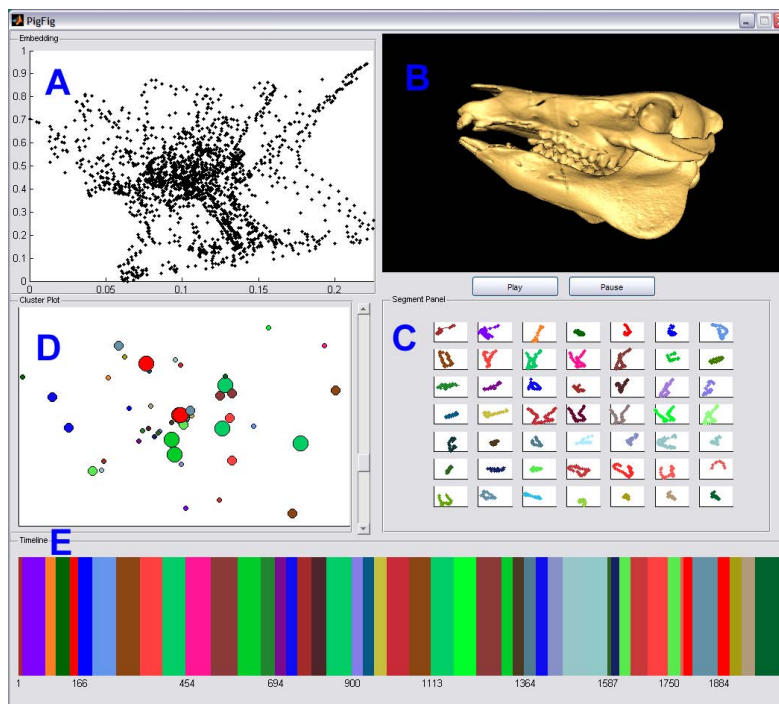


Fig. 4. The overview of our interactive visual analytics system. (A): The embedding view. (B): The 3D visualization view. (C): The segment view. (D): The cluster view. (E): The timeline view

Small Multiple Views We designed the small-multiples view to represent individual motion sequences in the embedding space. Because biomechanical motion data typically contains cycles, a key feature of our system is to help group and analyze the cyclic motions. Figure 4 (C) shows how these motion segments are assigned to small-multiple images corresponding to a trajectory in the embedding space. Each segment represents one cycle of frames, where each frame is a point in the embedding. Clicking on a segment will draw the trajectory on the relevant points in the embedding and show an animation of the corresponding frames in the 3D visualization view (see figure 4 (B)). The background color of each multiple encodes the data segment and corresponds to those in the cluster and timeline views (see Figure 4 (D)).

Embedding Space View Within the embedding (unlike the original high-dimensional data space), the Euclidean distance represents the dissimilarity between data points. To highlight the data distribution and correlation between different motions, we naturally display this as a 2D scatterplot where each frame of the motion data is encoded as a point. Previously defined segments encapsulate particular sequences of motions, which correspond to a sequence of points

in the embedding. Using standard mouse interactions, the user can analyze the low-dimensional representation to explore the perceptual organization of a particular data set. The user can examine each frame by mousing over the points, which will automatically update the 3D visualization.

Zoomable Cluster View Our system provides an interactive, zoomable cluster view. This view is a graph visualization that shows a point for each cluster. The points are positioned using multi-dimensional scaling (MDS) with the distance metric described in Section 4. Each bubble in this view indicates one cluster and is color-coded in accordance with the small-multiple and timeline views. Each cluster varies in size with the number of frames it encapsulates. As shown in Figure 4 (D), the cluster view enables user to interactively zoom in and out on different levels of grouping. The user can use the scroll bar to choose different levels of the grouping results, which automatically updates the size of individual clusters and corresponding coloring scheme.

5.3 Connecting the Embedding Space and the Data Space

The interaction and linked views are the keys for the user to simultaneously explore and analyze both the high-dimensional motion space and the lower-dimensional embedding. Since all views are coordinated visually and interactively, they collectively provide a cohesive exploration environment and support analysis of both spatial and temporal perspectives. For example, the embedding space view may be animated either through interaction in the embedding view or by selecting individual segments from the small multiple view. Also, the timeline can depict the temporal relationships at multiple clustering levels.

6 Discussion

In this paper, we introduce an integrated system that combines automated analytical methods with interactive visual analysis. Compared to existing work, our approach is innovative in that automated analysis can reduce the amount of ambiguity introduced through a user’s interactions. For instance, one of the most important features in the work by Keefe et al. [3] is the small-multiples view showing a trace of a point on the pig’s teeth plotted over time. This view is similar to the small-multiples view shown in Figure 4 (C). However, the key difference is that the tracer view requires the user to manually interact with the 3D view such that the front of the pig’s model is facing the user’s viewpoint. Only using this particular perspective can the tracer view show a pig’s bilateral chewing behavior, which is a sideways grinding motion during a pig’s chewing cycle. However, examined from the side-view, this subtle motion in the chewing would have been unnoticed. In contrast, with our method, we learn an embedding of the pig’s motion, which is analytically justified and is without the ambiguity of a user’s interactions. Since every user and analysis session presents the same embedding visualization, multiple analyses will be more consistent, and

the users are more likely to detect the same behaviors, thus providing a more defensible analysis result each time [19].

The agglomerative clustering we use provides the user a hierarchical structure to explore possible repeating phases within the motion. The user can interactively choose an appropriate threshold given a specific analytical goal. For instance, the user can choose a high threshold in the clustering view, which will produce fewer clusters of similar chewing segments to discover high-level behaviors. On the other hand, to identify low level differences between the segments, a lower threshold can be used to examine which of the segments are the most similar. When used in conjunction with the 3D comparison view, the user can then discover minute difference between the segments.

Without the combined use of automated methods and interactive visual representations, such analyses and discoveries may not be possible. Even with systems that allow for highly interactive visual analysis, the types of analyses are usually limited to visual comparisons by the user, which can be ambiguous depending on the user’s selections of viewpoints or segments of interest. Furthermore, the user cannot perform analyses across multiple levels of abstractions such as the features that our system can provide through clustering and interactive selection of thresholds.

7 Conclusions and Future Work

With new advances in scientific imaging, an increasing amount of high-resolution, high-quality biomechanical movement data is becoming available. With this opportunity comes the challenge of enabling scientists to make sense of information that is complex, temporal, multiphase, and cyclic. We presented a framework targeted at helping researchers meet this challenge. By combining machine learning methods with interactive visualization techniques, we provided users with a multi-pronged, hybrid approach to investigation. We demonstrated the combination of multiple, simultaneous views of the data where each view supports independent interaction, but work in concert to support more complex analysis.

This system is a positive first effort towards tighter integration of the user experience with the underlying analysis methods. In the future, we plan to conduct case studies with domain experts and apply our system to broader studies across data sets from more diverse domains. In addition, we plan to investigate additional analysis methods for the automated analysis of this type of temporal data. For example, Hidden Markov Models (HMM) could be used for data where the segments are not explicitly defined but can be learned based on the original data sequence. These potential additions combined with the general approach of blending automated and multi-view, interactive visual analysis open the door to new insights from scientific data analysis and exploration.

Acknowledgements We wish to thank Elizabeth Brainerd and the XROMM group at Brown University for the insight, infrastructure, and data that enabled us to explore this research within the context of the pig mastication application.

Thanks also to David Laidlaw for growing the initial interdisciplinary collaborations that made this research possible.

References

1. You, B., Siy, P., Anderst, W., Tashman, S.: In vivo measurement of 3-d skeletal kinematics from sequences of biplane radiographs: Application to knee kinematics. *MedImg* **20** (2001) 514–525
2. Keefe, D.F., M. O’Brien, T., B. Baier, D., M. Gatesy, S., Brainerd, E.L., Laidlaw, D.H.: Exploratory visualization of animal kinematics using instantaneous helical axes. *Computer Graphics Forum (EuroVis Special Issue)* **27** (2008) 863–870
3. Keefe, D.F., Ewert, M., Ribarsky, W., Chang, R.: Interactive coordinated multiple-view visualization of biomechanical motion data. *IEEE Transactions on Visualization and Computer Graphics* **15** (2009) 1383–1390
4. Thomas, J., Kielman, J.: Challenges for visual analytics. *Information Visualization* **8** (2009) 309–314
5. Ribarsky, W., Fisher, B., Pottenger, W.: Science of analytical reasoning. *Information Visualization* **8** (2009) 254–262
6. Robertson, G., Fernandez, R., Fisher, D., Lee, B., Stasko, J.: Effectiveness of animation in trend visualization. *IEEE Transactions on Visualization and Computer Graphics* **6** (2008) 1325–1332
7. Marai, G.E., Laidlaw, D.H., Andrews, S., Grimm, C.M., Crisco, J.J.: Estimating joint contact areas and ligament lengths from bone kinematics and surfaces. *IEEE Transactions on Biomedical Engineering* **51** (2004) 790–799
8. Jan, S.L.V.S., Clapworthy, G.J., Rooze, M.: Visualization of combined motions in human joints. *IEEE Computer Graphics and Applications* **18** (1998) 10–14
9. Jolliffe, I.T.: *Principal Component Analysis*. Springer-Verlag (1986)
10. Hyvärinen, A., Karhunen, J., Oja, E.: *Independent Component Analysis*. John Wiley and Sons (2001)
11. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290** (2000) 2319–2323
12. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290** (2000) 2323–2326
13. Zhang, Q., Souvenir, R., Pless, R.: On manifold structure of cardiac MRI data: Application to segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Volume 1., IEEE Computer Society (2006) 1092–1098
14. Winnemöeller, H., Mohan, A., Tumblin, J., Gooch, B.: Light waving: Estimating light positions from photographs alone. *Comp. Graphics Forum* **24** (2005) 433–438
15. Souvenir, R., Pless, R.: Image distance functions for manifold learning. *Image Vision Comput.* **25** (2007) 365–373
16. Alexa, M.: Linear combination of transformations. *ACM Trans. Graph.* **21** (2002) 380–387
17. Hsieh, J., Yu, S., Chen, Y.: Motion-based video retrieval by trajectory matching. *IEEE Trans. Circuits and Systems for Video Technology* **16** (2006) 396–409
18. Roberts, J.C.: State of the art: Coordinated & multiple views in exploratory visualization. In: *CMV ’07: Proceedings of the Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization*, Washington, DC, USA, IEEE Computer Society (2007) 61–71
19. Thomas, J.J., Cook, K.A.: *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Ctr (2005)