Trend-Centric Motion Visualization: Designing and Applying a new Strategy for Analyzing Scientific Motion Collections

David Schroeder *Student Member, IEEE*, Fedor Korsakov *Student Member, IEEE*, Carissa Mai-Ping Knipe, Lauren Thorson, Arin M. Ellingson, David Nuckley, John Carlis, and Daniel F Keefe, *Senior Member, IEEE*

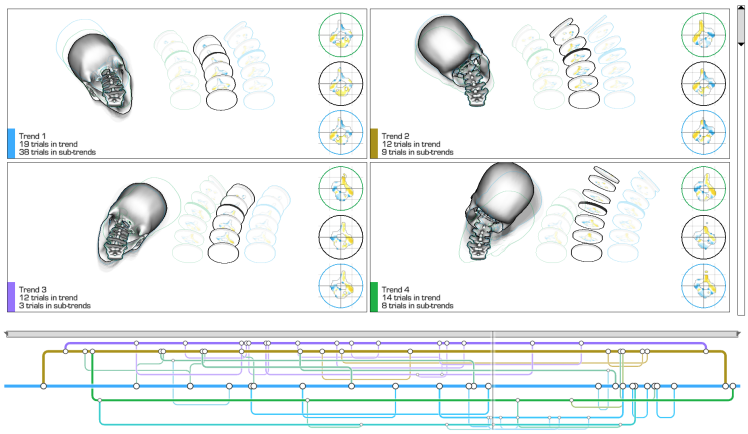

Fig. 1. This trend-centric motion visualization helps scientists analyze human neck kinematic data in a new way. Rather than analyzing a single specific trial in detail, what scientists would most like to do with these data is to identify similarities and differences across the entire collection of trials in order to classify motions as healthy vs. non-healthy or suggest appropriate courses of treatment. Trend-centric motion visualization helps scientists accomplish this analysis by: 1. identifying the trends in a motion collection; 2. displaying the trends in the form of the 2D timeline shown at the bottom of the figure; 3. displaying the 4D anatomical context needed to interpret the trends, as shown at the top of the figure; and 4. combining all these elements into an exploratory visualization system that supports interactive selection and querying.

Abstract—In biomechanics studies, researchers collect, via experiments or simulations, datasets with hundreds or thousands of trials, each describing the same type of motion (e.g., a neck flexion-extension exercise) but under different conditions (e.g., different patients, different disease states, pre- and post-treatment). Analyzing similarities and differences across all of the trials in these collections is a major challenge. Visualizing a single trial at a time does not work, and the typical alternative of juxtaposing multiple trials in a single visual display leads to complex, difficult-to-interpret visualizations. We address this problem via a new strategy that organizes the analysis around motion trends rather than trials. This new strategy matches the cognitive approach that scientists would like to take when analyzing motion collections. We introduce several technical innovations making trend-centric motion visualization possible. First, an algorithm detects a motion collection's trends via time-dependent clustering. Second, a 2D graphical technique visualizes how trials leave and join trends. Third, a 3D graphical technique, using a median 3D motion plus a visual variance indicator, visualizes the biomechanics of the set of trials within each trend. These innovations are combined to create an interactive exploratory visualization tool, which we designed through an iterative process in collaboration with both domain scientists and a traditionally-trained graphic designer. We report on insights generated during this design process and demonstrate the tool's effectiveness via a validation study with synthetic data and feedback from expert musculoskeletal biomechanics researchers who used the tool to analyze the effects of disc degeneration on human spinal kinematics.

Index Terms—Design studies, focus + context techniques, integrating spatial and non-spatial data visualization, visual design, biomedical and medical visualization

1 Introduction

New abilities to analyze multidimensional time-varying data are critical to solving many of today's important problems in science, engineering, and medicine. Imagine if doctors could record highly detailed spinal movements for a patient suffering from intervertebral disc degeneration. New insights from these data could positively impact health, informing not only patient-specific surgical interventions but also non-invasive therapies. Today, emerging technologies make it possible for doctors to collect and simulate such data. Unfortunately, analyzing these new data remains a major challenge, impeding possible clinical and research applications.

Consider analyzing spinal motions, where 3D vertebral shapes and orientations, and inter-vertebral distances are all complex and important. These constantly change as a patient moves, as do velocities, smoothness, coordination, and other data values needed by scientists.

- *David Schroeder, Fedor Korsakov, Carissa Mai-Ping Knipe, Arin Ellingson, John Carlis, and Daniel F Keefe are at the University of Minnesota. E-mails: {daschroe,korsakov,keefe}@cs.umn.edu, knipec@carleton.edu, {ellin224,carlis}@umn.edu.*
- *Lauren Thorson is with the Minneapolis College of Art and Design. E-mail: lthorson@mcad.edu.*
- *Arin M. Ellingson is with Department of Orthopedic Surgery at the Mayo Clinic. Email: Ellingson.Arin@mayo.edu*
- *David Nuckley is with Zimmer Spine. E-mail: dnuckley@umn.edu.*

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014. Date of publication 11 Aug. 2014; date of current version 9 Nov. 2014.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.
Digital Object Identifier 10.1109/TVCG.2014.2346451

This makes analyzing even a single trial of the data challenging. Furthermore, current studies of human spinal mechanics involve hundreds of patients, each performing multiple repeated trials (e.g., pre- and post-treatment) of the same motion exercise (e.g., a single flexion-extension of the neck), for a total of hundreds to thousands of trials. The result is a difficult-to-analyze multidimensional time-varying dataset, which we call a *motion collection*.

Our work is motivated by the needs of scientists studying motion collections. Specifically, these scientists need to analyze how biokinematics differ in healthy people vs. people with injuries or people in various stages of disease, and they need to study how biokinematics differ pre- and post-treatment. Since a motion collection's trials all are based on the same type of exercise, it is common practice to use time-warping or subsampling to align all the trials to a common time axis [12]. Unfortunately, despite warping, the spatial complexity of the motions together with the large size of motion collections make it impossible to do a thorough comparative analysis of these data using current statistical or visual strategies.

This paper introduces a new strategy for analyzing motion collections called trend-centric motion visualization. This strategy differs from previous motion visualization strategies (e.g., [20, 34, 18, 5]) in that it focuses on *trends rather than trials*. Through feedback from our collaborators, musculoskeletal biomechanics researchers studying the human spine, we know that a trend-centric strategy makes sense; for example, one scientist reports that basing his analysis on trends rather than trials is "absolutely correct" in terms of matching the mindset with which he would like to conduct his analyses. The technical challenge we address in this paper is redesigning data processing and visualization algorithms to support trend-centric visualization.

This research makes several contributions. (1) We introduce an algorithm for detecting a motion collection's trends via time-dependent clustering. This algorithm does not simply cluster together whole trials that are similar; rather it identifies *trends* – sequences of consecutive frames that are similar across multiple trials. For example, the algorithm might detect that a single trial follows a trend that is consistent with healthy neck motion during the first half of the trial, when the neck flexes, but follows a trend indicative of severe disc degeneration in the second half, when the neck extends. (2) We introduce a novel graphical technique inspired by transportation network maps for visualizing a motion collection's trends, including how trends come and go over time (Figure 1, bottom). (3) We contribute an illustrative graphical technique for visualizing a group of trials as a median with variance (Figure 1, top). (4) We demonstrate how widgets based on these two techniques can be combined into a single interactive tool, and apply this tool to analyze two different motion collections. The first is a set of 200 simulated motions of the human spine – this simulated data enables us to verify that the tool works as expected on known data. The second is from a recent cadaveric study of human spinal kinematics at different degrees of disc degeneration – we report on how our domain science collaborators used the tool to analyze these data. (5) Finally, we report on insights from our iterative design process, which included collaboration with both a traditionally-trained graphic designer and the domain scientists.

2 Related Work

2.1 Visualizing Motion

Visualizations of motion have been useful in medicine [26], sports [28], animation [5], migration studies [9], evolutionary biology [18], traffic analysis [3], and fluid dynamics [35]. Furthermore, the visualization research community has emphasized understanding both human perception of motion [8, 36] and how motion can be used as a visual cue to encode data [4, 33].

The motion visualization research that most closely relates to ours supports analysis of motions relevant to scientists. Visualization research in this area has focused on creating accurate animated (and often interactively controlled) 3D computer graphics representations for a single trial of motion data, for example, understanding the intricate kinematics, 3D geometries and spatial relationships, of the bones in the human wrist [26]. In this example, the researchers used a number of intelligent visual design decisions: visualizing distance data by displaying color-maps and 3D contour lines directly on the bones, enabling users to interactively show or hide individual bones to better view the small volumes between the bones, and, later [17], displaying the axes of rotation derived from experimental data in the same spatial context as the bones.

Other visualization tools follow a similar approach, introducing new 3D computer graphics techniques to optimize scientists' ability to analyze one trial's complex motions [30, 18, 35]. However, our scientist collaborators want to study multiple trials, and, while the tools for single-trial visualization can be pushed somewhat via overplotting (i.e., displaying data for multiple trials in the same 3D window) or side-by-side views, beyond a few trials such visualizations become too complex. Thus, these tools do not scale to visualizing motion collections.

2.2 Visualizing Motion Collections

Only a few recent visualization tools specifically address the problem of analyzing motion collections. Motion Explorer [5] combines a number of information visualization widgets (e.g., interactive dendrograms, node-link diagrams of motion graphs, a visual querying interface) with a pose-based clustering technique to help users analyze human motion capture databases as employed in computer game and movie animation.

We also use multiple interactive linked views, but our user requirements differ significantly. Motion Explorer helps users identify variations in stick figure poses based on motion capture data for about a dozen marker points distributed across the human body. Thus, the motion variations that it helps to detect are relative large (e.g., defensive vs. offensive boxing moves). In contrast, our scientists study detailed changes in the 3D spatial relationships, coordination of movement, etc. for multiple bones each of which have complex 3D shapes. This led us to analyzing trends rather than poses and utilizing detailed animated 3D views rather than thumbnail images.

Two recent visualization tools specifically help scientists understand motion collections. The first is applied to analyzing surgical training data [34], and the second to analyzing chewing motions in pigs, a topic studied by evolutionary biologists [18]. Like our work, both utilize multiple coordinated windows to convey multidimensional motion data. However, rather than manually selecting patterns in the data (e.g., via interactive brushing in a parallel coordinates plot), our tool first automatically classifies trends within the motion collection; then, these trends serve as the building block for interactive data analysis. Since individual trials can belong to different trends over time, this trend-centric strategy is a major departure from current clinical and research practice in biomechanics.

2.3 Other Related Visualization Techniques

Illustrative rendering and visual abstraction have been used to convey motion, for example, adding cartoon-like motion lines to show the trajectory of moving objects [15], see also [13, 16, 23, 27]. These and other illustrative visualization techniques (e.g., [2]) effectively use visual bandwidth by simplifying unneeded detail while emphasizing important detail. Our work adds two components missing from current illustrative renderings: depicting trends and variance across motion collections instead of a single motion, and driving these depictions with experimentally-collected motion data. Conceptually, visualizing differences between multiple motions is similar to uncertainty visualization [22, 24], which attempts to visualize not only the primary variables in a dataset but also the "error bars" around them. Our visualization techniques provide something similar to "box and whisker plots" for 3D motion data.

Our work also has some similarity to the broad topic of ensemble visualization. Ensemble visualization has typically been applied to sets of simulations run on high-performance computing architectures, as in a parameter study for engineering design [29] or an exploration of what if scenarios for disaster management [38]. We share a common goal of visualizing a collection rather than a single instance.

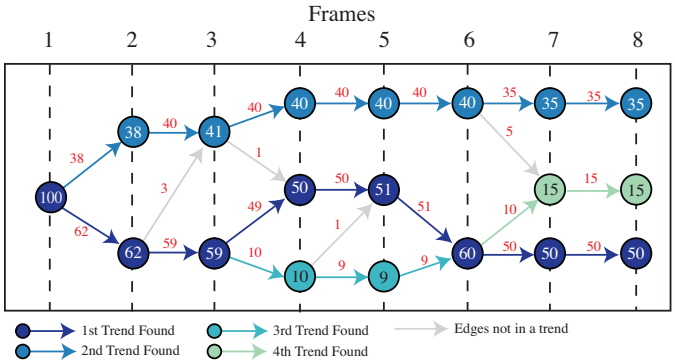

Fig. 2. Trends are identified using a directed graph data structure and a longest path algorithm that gives preference to the strongest trends present in the data.

3 Trend-Centric Motion Visualization

Our trend-centric motion visualization strategy begins by clustering the multidimensional data (e.g., distances, velocities, and bone orientations) in each trial. Clustering is first computed separately for each frame. Then, the clusters are connected to form coherent "trends" over time. This section begins by describing the trend-detecting algorithm in detail. Then, we present the series of 2D and 3D interactive visualization widgets developed to enable scientists to analyze trend-based motion collections.

3.1 Detecting Trends in Motion Collections

Although clustering is well studied within the visualization community, the problem of identifying trends that come and go over time within a set of many related trials is different than typical applications of clustering. Our goal is to identify important sequences of similar consecutive frames that exist across multiple trials in the collection. An example scenario that we want to detect is one where many trials all follow a similar pattern up to a certain point in time (e.g., until the neck reaches close to a maximum angle of flexion) and only then do significant differences begin to emerge in the data (e.g., healthy necks follow one trend, moderately diseased follow another, and severely diseased follow yet another). To accomplish this, we require a localized approach to clustering that operates at the frame-to-frame level.

Our algorithm assumes that each trial in the motion collection is comparable along a common time axis. Note that since the example applications discussed later in the paper deal with either synthetic data or *in vitro* data collected with the aid of a robotic device, each trial in these motion collections is already of the same duration.

For *in vivo* studies or other situations where data are less easily aligned to a common time axis, our trend-centric visualization strategy would require some form of temporal alignment as a pre-processing step. While time alignment is a difficult problem to solve in the general case, there are many common biomechanics scenarios where the time-alignment problem can be effectively solved via time-warping, time-normalization, segmenting based on key poses, or other strategies [12].

TODO: Feature vectors are defined at each frame of each trial, but the definition of the components of the vector depend on the dataset used. In our *in vitro* application, the feature vector includes the components of the frame-to-frame helical axis that describes the movement of L4 relative to L5. The biomechanics literature has shown that helical axes' utility for describing spinal motions [11]. In our synthetic dataset, the feature vector includes physiologically meaningful quantities, such as amount of flexion/extension and lateral bending. Once feature vectors are computed for every frame of every trial, the values are normalized per component to lie in the range $[0,1]$.

The algorithm's first step using these data is frame-wise trial clustering. Each frame is considered in isolation, and trials are clustered based on the similarity of their feature vectors. Our implementation uses a threshold-limited agglomerative clustering, and we found that

Fig. 3. The trend timeline illustrates how trend split and converge over time by using a style visual style similar to transportation network maps. Time progresses horizontally, and trends are separated vertically.

it is a useful practice when working with, often noisy, experimental data to also apply some subtle data filtering in this step. For example, in our implementation, the instantaneous velocity at a given frame is calculated using a sliding window to compute a weighted average of the velocities at neighboring frames.

The result of Step 1 is a set of one or more clusters for each frame of the motion data, as illustrated in Figure 2. Here, each cluster is represented as a circle. Notice that just one cluster is detected at Frame 1, and this single cluster contains all 100 trials in the motion collection. At Frame 2, two clusters are detected; one contains 38 trials, and the other contains 62. At frame three, a few trials switch from the lower to the upper cluster. Then, at Frame 4, three clusters are detected, one containing 40 trials, one containing 50 trials, and one containing 10 trials. Note that, the same clustering threshold is used for all the frames, but depending upon the data, this may result in a different number of clusters identified at each frame.

The second step is to connect these clusters to form trends that exist across multiple frames by identifying correspondences between the clusters at Frame i with those at Frame $i+1$. A graph data structure is used to accomplish this. We treat the clusters identified in step 1 as the vertices of the graph and connect the vertices with weighted, directed edges, as shown by the arrows in Figure 2. Edges are added to connect every cluster detected at frame i to every cluster detected at frame $i+1$. The edge weights are set to the number of trials that are common between the two clusters, and zero weight edges are omitted. For example, in Figure 2, look at the edges connecting the clusters at Frame 3 to the clusters at Frame 4. The top cluster in Frame 4 contains 40 trials, and we determine (by simply comparing the id's of the trials contained in each cluster) that all 40 of these came from the top cluster in Frame 3. The middle cluster in Frame 4 contains 50 trials, and we determine that 49 of these came from the bottom cluster of size 60 in Frame 3. The bottom cluster in Frame 4 contains 10 trials, and these all came from the bottom cluster in Frame 3.

With this graph complete, the trends can be found in a way that gives preference to the strongest trends in the data by iteratively calculating the longest path on the graph. The overall longest path found is identified as the first trend. Then, the vertices and edges that belong to this trend are removed from the graph and a new longest path is found. This second longest path is the second trend, and so on. Since this graph is trivially topologically ordered by frame index, each longest path calculation can be completed in linear time.

Each trend is defined by three data attributes: (1) the index of the frame at which the trend starts; (2) the index of the frame at which the trend ends; and (3) for each frame that the trend is active, the set of trials that belong to the trend. Our algorithms sometimes also consider the "strength" of each trend. We define the *frame-wise strength* of a trend as simply the number of trials contained in the trend at that frame. By extension, the *overall strength* of a trend is the sum of all its frame-wise strengths.

The entire clustering and trend-detection process is fast enough to allow users to adjust the clustering threshold interactively. This takes under one second on the *in vitro* dataset and under five seconds on the synthetic dataset. By adjusting the clustering threshold, users can examine the resulting trends at either higher or lower levels of detail.

3.2 Visualizing Trends on a Timeline

Figure 3 shows trends visualized with the horizontal axis representing time and different trends separated vertically. A trend can branch into several weaker trends or several trends can join together to form a new, stronger trend. When a new trend is formed, we define its main parent

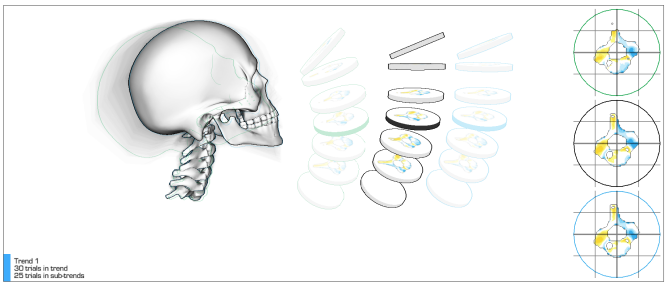

Fig. 4. Trends are displayed using a 3D visualization technique that functions like a box-and-whiskers plot, enabling scientists to understand both the key characteristics of the trend's median motion as well as a sense of the variation that exists across all of the trials in the trend.

as the trend contributing the most trials to the new trend; this is easily determined from the graph's edge weights. If the trend ends before the end of the motion sequence, this implies that the trials contained within it have all joined other trend(s); in this case, the trend to which the majority of the trials move is considered a second parent.

Each trend is rendered as a line with a unique vertical position. The first, strongest trend is assigned to the middle of the viewport. Then, a screen-space quality metric is minimized to find the best positioning of each successive trend. The quality metric $Q(y)$ is given below, where y is the proposed position, F is the set of frames for which the trend is active, T_f is the set of already laid-out trends that exist at frame f, y_t is the position of trend t, s_t is the strength of trend t, and ε is a small positive constant to avoid division by zero. We use $y \in [0,1]$ to correspond to the height of the viewport, and set ε to 0.05.

$$Q(y) = \sum_{f \in F} \sum_{t \in T_f} \frac{\varepsilon\, s_t}{\varepsilon + |y - y_t|} \tag{1}$$

In practice, we found that the overall trend strengths for the biokinematic datasets that interest us tend to follow a bimodal distribution, with a relatively small number ($\tilde{5}$) of strong *main trends* plus a much larger set of weaker *sub-trends* with not much in-between these two groups. We use this insight to our advantage in assigning colors and line weight to the trends.

We test each new motion collection algorithmically to see if the trend strengths do indeed follow a bimodal distribution. If they do, we divide the trends into the two classes: *main trends* and *sub-trends*.

Trend colors are then assigned based on strength and on the colors of a trend's parents (if they exist). All main trends are assigned a color in the CIELAB color space on the plane $L^* = 60$, and on a circle centered on $(a^*, b^*) = (0,0)$ with radius 110. These coordinates provide distinct, vividly saturated colors for the main trends. Then, all subtrends with at least one parent are assigned a desaturated, lightened version of the average color of the parent(s).

The line weight (width) of each trend line is similarly based on the trend's strength and computed using the equation

$$w_i = (0.7\sqrt{s_i/s_1} + 0.3) w_{max}, \tag{2}$$

where i is the index of the trend, w_{max} a the maximum line width based on the size of the viewport, and s_1 is the strength of the strongest trend. The square root allows for differences between trend strength to be seen across the full range of strengths.

Trends are drawn at their assigned vertical position using the computed color and line weight as a horizontal line that spans all of the frames at which the trend is active. Finally, curved lines are added to connect each trend to its parent(s) when appropriate.

3.3 Visualizing Trends with Appropriate Anatomical Context

The trend timeline overviews a motion collection's trends, but interpreting the motions also requires seeing the anatomical context. Building upon existing motion visualization techniques, our 3D visualization includes techniques useful for analyzing individual trials (e.g., contour lines, and axes of rotation), but also must depict a set of related motions.

To address this problem, we introduce a median plus variance visualization technique that, as shown in Figure 4, juxtaposes three trials in the same 3D view. First, the most representative trial of the trend, its median trial, is rendered. Then, two additional trials are rendered, one plus and one minus a standard deviation away from the median.

The median trial is computed taking into account all the available data recorded for each frame (positions, orientations, distances, speeds, etc.). First, the feature vector defined in Section 3.1 is computed for each frame of each trial in the trend. Then, the mean feature vector is computed for each frame. The median is computed by comparing each trial's sequence of feature vectors to the sequence of mean feature vectors. The trial with the minimum summed squared difference is considered the median. Note that although we compute an "average motion" during this calculation, at the request of our collaborators, we deliberately display the median motion rather than the mean, since the mean is not guaranteed to be physically possible, while the median comes from an actual observation.

The trials that are +/- one standard deviation away from the median are computed using the same feature vectors. A standard deviation is calculated from the distribution of each trial's total distance from the mean trial; this distance is based on the feature vectors summed over all frames of the trial. With a standard deviation calculated, it is then possible to identify the subset of trials that fall within one standard deviation of the median. From this subset of trials, the two trials with the largest mutual dissimilarity are selected as the trials that are 'plus' and 'minus' one standard deviation away from the median.

Finally, an underpainting is added to the visualization to convey the full range of a trend's trials. All of the trend's trials are drawn in a flat, semi-transparent color. By analogy to a box-and-whisker diagram, the three 3D renderings of the bones are similar to the box in that they prominently display values for something like the first, second, and third quartiles in the data. The underpainting is similar to whiskers in that it displays in a less prominent manner the full range of the data, including any outliers. The final result is an animated rendering, as shown in Figure 4 and the accompanying video, that conveys not just the characteristic motion that defines the trend but also a sense of the variance within the trend.

3.4 Interactively Exploring a Complete Motion Collection

Figures 1 and 5 illustrate how the trend timeline and the 3D trend visualizations are combined into an interactive tool. The trend timeline at the bottom of the screen not only provides an overview of how all the trends in the motion collection come and go over time, but also provides a way to drill down through the data to examine specific trend(s) in detail. Clicking on a trend selects the trend and its subtrends for further analysis. Each of the selected trends is then displayed in its own 3D visualization viewport on the top portion of the screen. A range bar on the right side of the visualization enables users to narrow the subset of trends to display in 3D windows, which is useful during the most detailed stages of analysis.

To complement these interactive zooming and filtering techniques, scientists use mouse and keyboard interactions inside both the 3D and 2D windows to explore the data more closely. Scientists can rotate, pan, and zoom within the 3D views. By default, the median trial is fully shaded, but this can also be adjusted interactively by simply hovering the mouse over either the plus or minus one standard deviation renderings. Mouse hovering, in both the 2D and 3D views, also causes the timeline visualization to indicate the provenance of all of the trials in that trend, as in Figure 6.

4 Applications and Evaluations

Our research has been conducted in collaboration with a group of musculoskeletal biomechanics researchers. Both of the applications described in this section are derived from real data analysis problems faced in their research. In the next section, we describe a verification

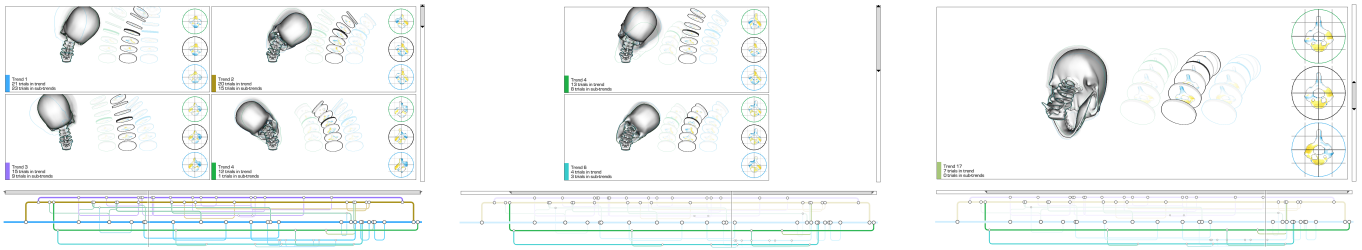

Fig. 5. A typical analysis sequence begins with an overview of all trends in the motion collection (left). Then, scientists select a subset of interest (middle). Then, the range bar is adjusted to further focus the visualization by displaying just a subset of the selected trends in the 3D view; here, analysis is focused on a sub-trend of interest (right).

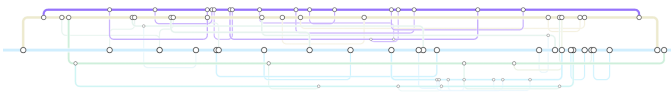

Fig. 6. When a single trend is selected by clicking, the trend timeline visualization updates to convey the provenance of each of the trials in that trend. Here, a trend near the top of the motion sequence is selected – notice how line width is used to convey where the trials in this trend came from and are going "upstream and downstream" from the selected trend.

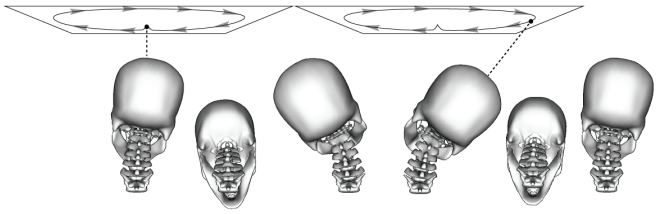

Fig. 7. The simulated skull motion. The skull is constrained to point toward a point on a spline on a plane above the skull. The spline moves the skull forward, where it then rolls counterclockwise and returns to the resting position.

study, which uses simulated data that we created in order to verify that the visualizations work as expected under known conditions. Although we generated this data synthetically, the type of motion is motivated specifically by a recent study conducted by our collaborators. The second application is based on actual experimental data collected by our collaborators during a cadaveric study of 18 human lumbar spine specimens.

4.1 Verification Study

When patients with neck and/or back pain visit a doctor they are often asked to perform physical exercises from which measurements are made. The simplest is a flexion/extension exercise – the patient bends his or her neck as far forward and as far backward as possible, and the maximum angles reached without pain are recorded. In the biomechanics research community, there is now a push to incorporate data from more complex exercises into clinical decision making. One such exercise is neck circumduction (Figure 7). As an example, a recent study measured head-to-torso kinematics using a lightweight linkage device. The superior end was worn on the head and inferior end was positioned over the T1 vertebra and snuggly attached to the torso. The study contained over 100 patients, and each patient performed the neck circumduction exercise three times at two different sessions, for a total of 600 trials [11]. Researchers are excited about these datasets and their potential impact on healthcare because they capture much more information than current 2D descriptions of motion commonly used in current clinical practice. Even more exciting than the external tracking data described above is the ability to collect and simulate the internal

Table 1. Simulated motion datasets.

Dataset	Num. Motions	Direction and Speed	SD of Noise	Speed Variation
Dataset 1	40	clockwise, quickly at first	0.1 m	2%
	20	clockwise, slowly at first	0.1 m	2%
	27	counter-clockwise, quickly at first	0.1 m	2%
	13	counter-clockwise, slowly at first	0.1 m	2%
Dataset 2	40	clockwise, quickly at first	0.4 m	8%
	20	clockwise, slowly at first	0.4 m	8%
	27	counter-clockwise, quickly at first	0.4 m	8%
	13	counter-clockwise, slowly at first	0.4 m	8%

motions of the vertebrae. These data can be collected experimentally using biplane fluoroscopy (e.g., [7]) or simulated using a neck kinematics model (e.g., [37]); however, since this type of data has only recently become available, there is not yet consensus on what patterns will be most useful for clinical decision making.

4.1.1 Synthetic Data

With this understanding of the type of motion datasets that investigators are now collecting and struggling to analyze, we developed a verification study based on our own simulated dataset of neck motions. Although our simulation method is inspired by models published in the biomechanics community [37], we must begin our discussion by emphasizing that our goal in this application was *not* to develop a new physically-accurate neck kinematics model. Rather, we set out to generate a synthetic dataset of physically-plausible motions of the vertebrae with programmatically controlled variance. This enables us to verify that the visualization techniques work as expected on synthetic data.

The synthetic data simulate the motion of the spinal column from the skull to the T1 vertebra. We use the Bullet physics engine, and each bone is modeled as a rigid body connected to neighboring bones via springs. To calculate vertebra positions and orientations, we fix the position and orientation of T1 and apply a constraint to the skull to force its "up" vector to point at a moving target located on a plane 1.3 meters above the skull. The spring constraints then adjust the position and orientation of the skull and the vertebrae between the skull and T1, producing a physically plausible pose for each bone. The characteristics of each simulation can be adjusted by adjusting the moving target's path. To create an initial path, we captured an example motion experimentally using an optical motion tracker, and then simplified the data to convert it into a spline with 11 control points. By adjusting the spline path programmatically, we generate motion collections where each motion is slightly different.

We generated two different synthetic motion collections using this strategy. Table 4.1.1 summarizes the characteristics of each. Each contains 100 trials, which we break into four subgroups based on relatively significant changes in the spline path, such as changing the direction of the movement (clockwise vs. counter-clockwise) and the speed profile (quickly at first vs. slowly at first). Additional variation is then added by applying a Gaussian noise function to the positions of the spline's control points and applying random variations to the speed

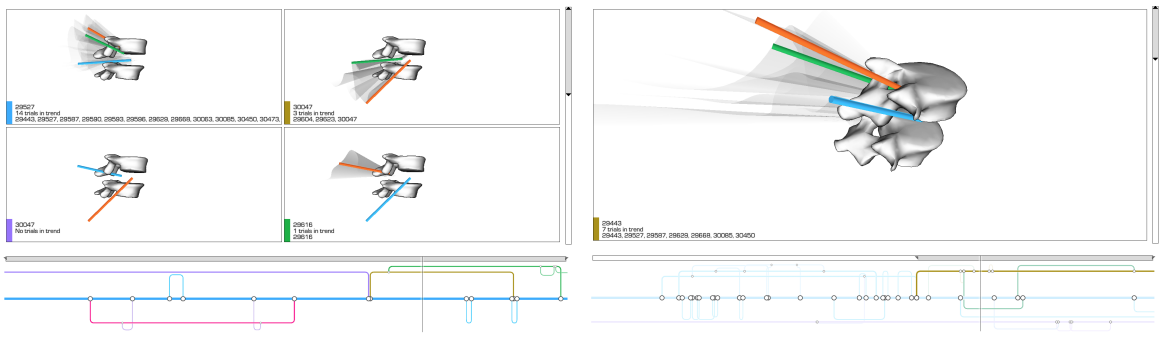

Fig. 8. Two screen shots of the trend-centric motion visualizations used to analyze data from a recent *in vitro* study of 18 lumbar spine specimens. The visualizations highlight the helical axes computed for the motions and the 3D paths that these axes sweep out over time. The motion of the vertebra is more subtle in this example, as compared to the example pictured in Figure 1. Here, the variation in the helical axis orientation is the most important factor to analyze. So, in this example, we apply the median plus variance visualization technique described earlier directly to the helical axis data.

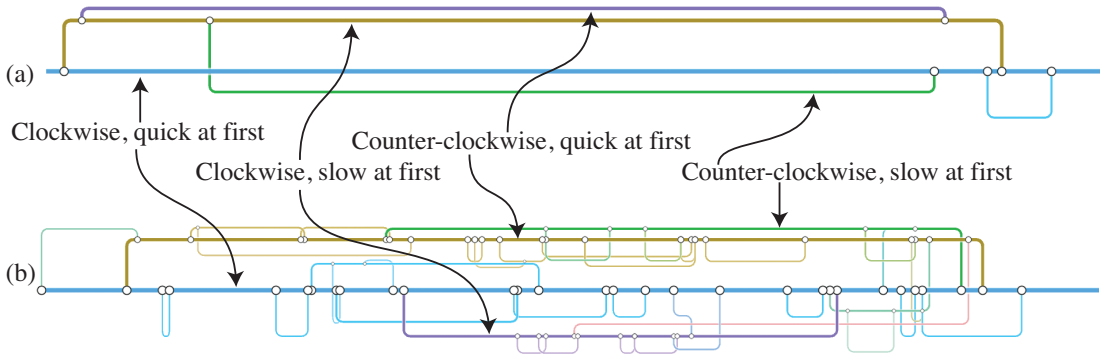

Fig. 9. Clusters identified in two motion collections differing only by the amount of noise present in the data. The trials in (b) have 4 times as much noise applied. This has a clear impact on the trends detected, in particular a number of new "sub-trends" are detected; however, the major expected patten in the synthetic data is still visible: all trials start following a single trend branch into 4 main patterns with some noise added, then come back together to follow a single trend.

at which the moving target traverses the spline. Since the moving target is positioned on a plane 1.3 meters above the skull, it sweeps out a path that covers a relatively large area (about 3 x 4 meters). In Dataset 1, the Gaussian noise added to the control points has a standard deviation of 0.1 meters, so significant variation can be observed. In Dataset 2, the standard deviation is 0.4 meters. The variations in Dataset 2 are significant enough that some trials include major changes from the expected pattern, such as a "hitch" in the motion or a short sequence where there is a reversal in the direction of rotation.

4.1.2 Visualization Results and Analysis

Figure 1 shows an example trend-centric visualization of these data, and Figure 9 shows a detailed comparison of the trendlines that are generated for each of the two synthetic datasets. The feature vector used for clustering in these examples was composed of four physiologically meaningful quantities for the motion of the skull relative to T1: (1) the amount of forward/backward bending, (2) the amount of left/right bending, (3) the positional speed, and (4) the rotational speed. All of these were normalized to lie on the range (0,1) to ensure equal weighting.

As shown in Figure 9 (top), the trend line produced for Dataset 1 closely matches the trends we expect to discover given the properties of the data described above. For the majority of the frames, four distinct trends are clearly identifiable. If we look at the composition of these trends, we find that each trend corresponds to one of the subgroups in Table 4.1.1. All of the trials also cluster into a single trend at both the beginning and end of the entire sequence. Each simulated exercise begins from the same neutral upright pose, then the head lowers so that the eyes would be looking straight down. Then, the rotational motions begin. The timing of the branching point observed in the trend line matches the transition to rotational movement. Likewise, at the end of each exercise, the head returns to the same neutral upright pose; thus, each trial is very similar at this point in the sequence.

Figure 9 (bottom) shows a trendline produced with the same clustering threshold, but this time for Dataset 2, which includes far more substantial variation between trials. This provides an interesting point of comparison. The major trends in the data are still identified, but additional "sub-trends" are also identified as the variation causes individual trials to sometimes join or detach from a trend over time. These sub-trends do not cross over major boundaries. For example, notice that there is a major branch early on in the sequence that divides the trials that run clockwise from the trials that run counter-clockwise. New trends and sub-trends come and go after this point, but clockwise and counter-clockwise trials are never characterized as belonging to the same trend until the end of the motion sequence at which time all of the trials follow a similar pattern of returning to the neutral upright pose.

We interpret these results, augmented with our observations from interactive use of the tool, as a verification that the visualization system functions as expected with well-understood data, displaying major trends clearly and merging into a single trend when appropriate. When the variation across the trials increases, the complexity of the trends identified also increases, but major patterns remain visible and smaller patterns provide useful pointers to interesting nuances of the data.

4.2 Experimental Study and Expert User Feedback

The second application is to an *in vitro* biomechanics study conducted by our collaborators. The study compares motion of the lumbar spine for 18 human spine specimens donated through an anatomy bequest program. All musculature was removed and the spine was mechanically tested using a Spine Kinetic Simulator, which is a six-axis servo hydraulic testing apparatus capable of reproducing spinal motion for *in vitro* specimens. The superior (L3) and inferior (Sacrum) vertebral bodies bodies were fixed to the device and pure moments were applied in all bending directions: flexion/extension, left/right lateral bending, and left/right axial rotation. The motion of the two vertebrae of interest for this study (L4 and L5) were then tracked optically using a marker-based Vicon motion capture system, calibrated to record motion data at 100 Hz with an accuracy of 0.02 mm.[10]

4.2.1 Results and Analysis

Figure 8 shows the visualization results achieved for the left/right lateral bending trials recorded for the 18 specimens, each selected because it is representative of a different degree of disc health. The helical axes between L4 and L5 were of primary interest to our collaborators. The previously presented 3D visualization strategy is maintained: the helical axes of the median trial and the two "+/- one standard deviation" trials are shown in distinct colors, and helical axis sheets for all trials are shown in the background.

The results both confirm some expected patterns in the data and suggest new directions for further analysis. In Figure 8 (left) notice that there are two regions where almost all of the trials are grouped together into a single trend. These groupings make sense given the investigators' current hypotheses and the analyses they have conducted to date using other methods. Moving from left to right across the trendline, the spines begin in a neutral position, then bend all the way to the right, then all the way to the left, then return to neutral. When the trials come together into a single trend, this corresponds to the sections of motion when the spines are transitioning from being unloaded to loaded in the other direction. The likely explanation of these groupings is that during these sequences, the spine is going through the "neutral zone" – it is not fully loaded. The investigators hypothesize that the trials will be most similar during these transitions, and this hypothesis is supported by the trend-centric analysis. In contrast, as seen in Figure 8 (left), at other frames, many distinct trends are detected. What we believe the trend-centric analysis is demonstrating in this case is that the investigators did an excellent job of picking a wide range of specimens for their study. In this study, the specimens were specifically chosen to cover a wide range of disc health, age, and other factors that may impact the motion. The analysis suggests that these factors do, indeed, yield different motion patterns but only in a loaded state at the end ranges of motion.

Together with our collaborators, we then explored the logical follow-on question, "how do the trends change if we adjust our notion of how similar the trials must be to be included in the same trend?" By decreasing the clustering threshold so that clusters form more readily, the three most prominent trends at one point in the data match the scientists hypotheses, for example, the two specimens that are considered "most healthy" are clustered together in one of these trends. At this point, it is less clear what to make of the two other trends, but the scientists are excited to follow up on this analysis. For example, perhaps the other trials that were clustered together share a similar pattern of disc degeneration that could be detected with MRI.

4.2.2 User Feedback

High-level feedback from the domain scientists was positive. Typically, kinematic data is collected using motion capture software and analyzed either within this software or via custom MATLAB scripts. Traditional measures of spinal biomechanics primarily include scalar metrics that characterize the endpoints of the motion. These data are analyzed for the purpose of statistical comparisons between different test groups. In contrast, our trend-centric visualization strategy enables scientists to study not just summary statistics for the endpoints of the motion but also detailed patterns that are visible at any time during the motion. Our collaborators report, "since there are infinite pathways of motion to reach [the] endpoints, it is of critical importance to quantify the quality of motion throughout the entire trial."

Since few visual data analysis tools in this style exist, our first concern in evaluating the tool was to understand whether the metaphors adopted in the visualization, especially the notion of motion trends that may change over time, actually fit the way these scientists think about their analyses. The response from scientists was that the approach is "Absolutely correct on separating into groups."

Another high-level indication of the success of the tool is that our collaborators have repeatedly told us that they are going to change the way that they collect their data in the future in order to better facilitate similar visualization efforts. Since this will bring visualization into the earliest stages of their investigative process, we take this as an important indication of the potential value of this tool.

Several lower-level points of feedback also emerged during discussion and tool use together with our collaborators. The visual aesthetic and overall display of the data was appreciated. We believe this increases engagement with the tool; moreover, we believe the presented design has a significant usability advantage over traditional multi-view visualization systems that require manual window positioning. The ability to adjust the 3D viewpoint was evaluated as absolutely critical, and different views (e.g., top down) were preferred depending on the particular question being discussed. There was much discussion about the trend visualization and how well it did or did not capture characteristics of the data. During this discussion, we realized how critical it is to support the ability to interactively adjust the threshold used in clustering, which is possible for this dataset. (The synthetic datasets, which are larger, require a few seconds to cluster; however, exploiting per-frame clustering's inherent parallelism could speed this up dramatically and allow real-time interaction.)

5 Iterative Design and Insights

Our trend-centric visualization techniques were designed through an iterative approach involving domain scientist collaborators plus a traditionally trained and experienced graphic designer. This work builds upon recent research that seeks to understand how best to integrate visual artists into designing and evaluating sophisticated data visualizations [19, 1, 21, 14]; researchers and practitioners working in any area of data visualization can adapt our collaborative design process.

This section describes our team's iterative design process and lessons learned, and is oriented around three specific aspects of trend-centric visualization: (1) design of the trend lines; (2) design of the discs; (3) design of the median + variance visualizations. For each of these aspects, the graphic designer developed, literally, hundreds of "design sketches", often using traditional graphic design tools, such as Adobe Illustrator. This sketching was the key activity in our design process's ideation phases. Several sketches are pictured in Figures 10 and 11; however, please refer to the Supplemental Material accompanying the paper to better understand the depth of this process, including the large number of visual possibilities considered. Each of the sketches, along with additional inspirations (e.g., Figure 12), were evaluated by the team using critique. The graphic designer and visualization researchers participated in these critique sessions once a week, for nearly a full year, with the domain scientists joining the group once every 4-8 weeks. Based on the insights generated during critique, the graphic designer refined the design and/or handed the design off as a specification to the team's programmers. As full-featured data visualizations were developed through the programming process and programmers added their own creative insights, they were likewise critiqued. The three specific examples described next illustrate the types of visual design insights that resulted from the process and justify some key design decisions we made.

5.1 Design of the Timeline

Our trend timeline designs started with visuals conveying the frame-wise similarity between multiple trials.

Figure 10a displays four trials with colored dots arranged below the timeline at the display's bottom to denote specific frames of each trial

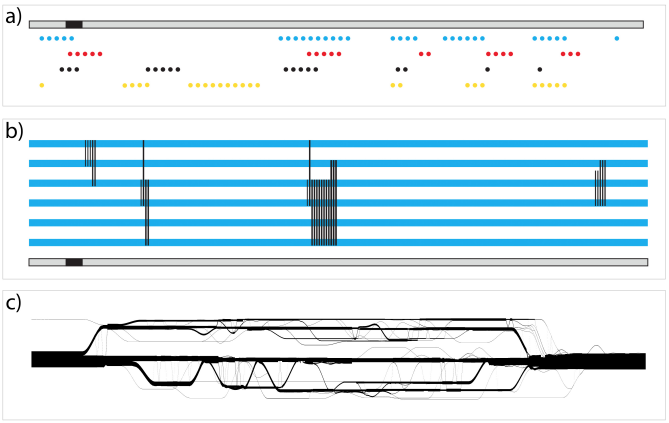

Fig. 10. Three design sketches used to prototype and evaluate alternative representations for a the trend timeline.

that have similar characteristics. The visual idea of representing each frame as a colored dot was evaluated as an interesting way to support detailed, frame-level analysis on a timeline, but the idea was evaluated as limited in its applicability to large datasets since it would necessitate many small dots (or a scrollable timeline) for trials with many frames.

Figure 10b illustrates an example follow-on design. Here, the timelines have a continuous visual representation (a solid blue horizontal line), which supports large datasets well. To convey similarity between the trials a black vertical line connects them at the frames where they are most similar. The main strengths of this design are the continuous representation for the trials and that the viewer's eye is immediately drawn to the strong visual links between the trials. The critical weakness is that this vertical line strategy does not work unless the trials are sorted in some way so that the links only need to reach adjacent trials and do not need to skip over lines; thus, this design was judged as impractical for use with any realistic data but was useful for pushing the team toward visuals that provide strong visual links between the trials.

Figure 10c shows the result of a critical innovation in the design. The idea demonstrated here switches from representing individual timelines for each trial to representing multiple trials together in a single "trend". This was evaluated as such an exciting concept that we advanced well beyond design sketches with this idea and actually produced a full implementation – Figure 10 is a result for the same dataset as used in the validation study described earlier. The main strength of this design is that bundling the trials into trends provides the clearest depiction yet of the similarity of multiple trials over time. We evaluated using line width to encode the strength of each trend, and deemed it too difficult to implement the gradual transitions that can be seen as a trend branches off. We implemented several alternatives but could not develop a satisfactory answer to the question of how these trend lines should be drawn when several sub-trends branch in and out from each other during transition periods. Another weakness is that it is unclear exactly where one trend starts and another ends. This makes it difficult to design a good visual strategy for linking the trend line display with the 3D visualizations at the top of the screen.

The final trend line design pictured throughout the paper solves both of these problems. We were inspired in this final design by Roberts' studies of transportation network maps [31], borrowing the visual concepts of: explicitly representing junctions; using clearly distinguishable colors for main lines and similar, but less saturated, colors for branch lines; and constraining angles when lines branch and join, thereby increasing readability and making implementation easier. We reinterpreted each of these visual concepts within our own driving applications in order to advance from the designs pictured in Figure 10 to those seen in Figures1 and 3.

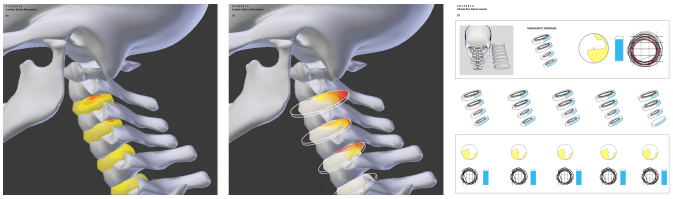

Fig. 11. Three design sketches used to prototype and evaluate visual techniques for conveying the motion of individual vertebrae and the distances between the vertebrae; this type of ideation led to the offset-simplified-discs visualization technique pictured in Figures 4 and 1.

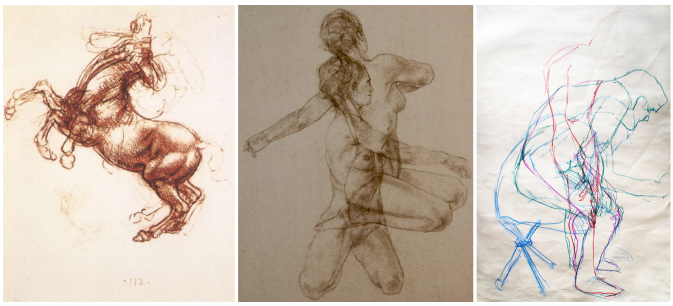

Fig. 12. These three drawings demonstrate the effective illustrative technique of juxtaposing multiple poses in artistic studies of motion. These motivate the median + variance rendering strategy developed. Center image copyright Gary Kaemmer, used with permission. Right image copyright Aleksandra Kulecka, used with permission.

5.2 Design of the Discs

The design decision to use a set of simplified cylindrical discs to better convey the orientation and distances between the vertebrae is one that we think we would not have reached without an interdisciplinary collaborative design process. Figure 11 shows three design sketches produced during refinement of the idea. As computer graphics researchers, we found that our tendency was to create visualizations based upon the most detailed and realistic 3D models of the vertebrae that we could acquire. The sketch in Figure 11a, builds on this approach, exploring a visual technique for conveying the distance between the vertebrae though color mapping directly on the vertebrae. However, this region is so tight that it is difficult to read the data. Figure 11b is the first of many sketches introduced by the graphic designer based on her idea to simplify the vertebrae by representing them as cylinders, making it easier to read the data displayed on them. As shown here, the first of these ideas were based on a two-pass rendering approach: draw 3D vertebrae with a traditional 3D rendering and then draw the disc geometries on top. The design was refined through many iterations (e.g., Figure 11c) to include offsetting the discs to the right of the accurate 3D rendering and several additional visual strategies for facilitating reading data off the discs. Our critique-based evaluations of this idea identified several strengths of the approach. First, the disc motion is easily seen because the geometry is simpler than the spine's geometry and occlusion is minimized. Second, it is much easier to read visual widgets drawn on discs (flat tops and bottoms and uniformly curving sides) than ones drawn on bone surfaces.

5.3 Design of the Median + Variance Visualizations

The visual technique used in the median + variance 3D visualizations is inspired directly by a common illustration technique used in traditional drawing: artists, to study and convey motion, juxtapose two to three different subject poses in one drawing. Figure 12 shows three examples. On the left, Leonardo daVinci emphasizes one pose while also suggesting several others that are drawn with less detail. Like daVincis horse, our design aims to highlight one primary pose while simultaneously conveying a more complete sense of related poses us-

ing a less detailed rendering style. We accomplish this through the underpainting layer. In the center, contemporary artist Gary Kaemmer juxtaposes two detailed poses, demonstrating that even detailed renderings of human anatomy in action can be effectively understood when a small number of poses are juxtaposed together. Like Kaemmer's example, we include a small number of detailed renderings of the anatomy when these poses are critical to understanding the motion. We use three rather than two poses, one for the median trial and one each for the trials that are plus and minus one standard deviation away from the median. On the right, another contemporary artist, Aleksandra Kulecka, employs color to distinguish the multiple poses depicted in this quick gesture sketch. Like Kuleckas work, we utilize color as a cue to distinguish the different poses.

6 Conclusions and Future Work

Trend-centric motion visualization is a new strategy for improving the way that scientists analyze motion collections. The key conceptual advance is to base analysis on motion trends rather than trials. To implement this strategy, we developed a series a visualization techniques. The technical innovations include: a time-dependent motion clustering technique, a 2D trend lines visualization inspired by transportation network maps, a 3D illustrative visualization technique for depicting a cluster of trials as a median with variance, and interactive techniques for linking these visualizations together to create an exploratory visualization tool. The work is evaluated through several means. First, the data processing and visual techniques are validated using a synthetic motion collection with known properties. Second, the new visualization tool is applied to a cutting-edge cadaveric study of the impact of disc degeneration on the biokinematics of the human lumbar spine; collaborators in this study confirm the benefits of the trend-centric strategy and the specific tool implemented. Finally, our interdisciplinary team presents insights from our own critique-based evaluations of the specific visualization techniques developed during the iterative design process.

In the future, we plan to apply trend-centric visualization to additional synthetic and *in vivo* datasets. This will enable us to study the scalability of trend-centric visualization, specifically with respect to analyzing more subtle motion patterns and larger data collections. New spinal kinematics studies are planned for the coming years; however, we are also keen to explore other application areas, such as gait recognition, Fitts' law tasks, and sports medicine; time-alignment strategies have been succesfully employed in each of these areas previously (e.g., [6, 25, 32]).

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1054783. This work was also partially supported by NIH/NIAMS grants T32 AR050938 and T32 AR056950.

References

[1] D. Acevedo, C. Jackson, F. Drury, and D. Laidlaw. Using visual design experts in critique-based evaluation of 2D vector visualization methods. *Visualization and Computer Graphics, IEEE Transactions on*, 14(4):877–884, 2008.

[2] M. Agrawala and C. Stolte. Rendering effective route maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01*, pages 241–249, New York, New York, USA, 2001. ACM Press.

[3] G. Andrienko, N. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel. From movement tracks through events to places: Extracting and characterizing significant places from mobility data. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 161–170. IEEE, 2011.

[4] L. Bartram, C. Ware, and T. Calvert. Moticons:: detection, distraction and task. *International Journal of Human-Computer Studies*, 58(5):515–545, 2003.

[5] J. Bernard, N. Wilhelm, B. Kruger, T. May, T. Schreck, and J. Kohlhammer. Motionexplorer: Exploratory search in human motion capture data based on hierarchical aggregation. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2257–2266, 2013.

[6] N. Boulgouris, K. Plataniotis, and D. Hatzinakos. Gait recognition using dynamic time warping. In *Multimedia Signal Processing, 2004 IEEE 6th Workshop on*, pages 263–266, Sept 2004.

[7] E. L. Brainerd, D. B. Baier, S. M. Gatesy, T. L. Hedrick, K. A. Metzger, S. L. Gilbert, and J. J. Crisco. X-ray reconstruction of moving morphology (XROMM): precision, accuracy and applications in comparative biomechanics research. *Journal of Experimental Zoology Part A: Ecological Genetics and Physiology*, 313A(5):262–279, 2010.

[8] D. Coffey, F. Korsakov, M. Ewert, H. Hagh-Shenas, L. Thorson, A. Ellingson, D. Nuckley, and D. Keefe. Visualizing motion data in virtual reality: Understanding the roles of animation, interaction, and static presentation. *Computer Graphics Forum*, 31(3pt3):1215–1224, 2012.

[9] T. Crnovrsanin, C. Muelder, C. Correa, and K. Ma. Proximity-based visualization of movement trace data. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 11–18. IEEE, 2009.

[10] A. M. Ellingson, H. Mehta, D. W. Polly, J. Ellermann, and D. J. Nuckley. Disc degeneration assessed by quantitative T2*(T2 star) correlated with functional lumbar mechanics. *Spine*, 38(24):E1533–E1540, 2013.

[11] A. M. Ellingson, V. Yelisetti, C. A. Schulz, G. Bronfort, J. Downing, D. F. Keefe, and D. J. Nuckley. Instantaneous helical axis methodology to identify aberrant neck motion. *Clinical Biomechanics*, 28(7):731–735, 2013.

[12] N. E. Helwig, S. Hong, E. T. Hsiao-Wecksler, and J. D. Polk. Methods to temporally align gait cycle data. *Journal of biomechanics*, 44(3):561–566, 2011.

[13] W.-H. Hsu, J. Mei, C. D. Correa, and K.-L. Ma. Depicting time evolving flow with illustrative visualization techniques, 2009.

[14] B. Jackson, D. Coffey, L. Thorson, D. Schroeder, A. M. Ellingson, D. J. Nuckley, and D. F. Keefe. Toward mixed method evaluations of scientific visualizations and design process as an evaluation tool. In *Proceedings of the 2012 BELIV Workshop: Beyond Time and Errors - Novel Evaluation Methods for Visualization*, BELIV '12, pages 4:1–4:6, New York, NY, USA, 2012. ACM.

[15] A. Joshi and P. Rheingans. Illustration-inspired techniques for visualizing time-varying data. In *Visualization, 2005. VIS 05. IEEE*, pages 679–686. IEEE, 2005.

[16] A. Joshi and P. Rheingans. Evaluation of illustration-inspired techniques for time-varying data visualization. *Computer Graphics Forum*, 27(3):999–1006, 2008.

[17] R. N. Kamal, M. J. Rainbow, E. Akelman, and J. J. Crisco. In vivo triquetrum-hamate kinematics through a simulated hammering task wrist motion. *The Journal of Bone & Joint Surgery*, 94(12):e85–1, 2012.

[18] D. Keefe, M. Ewert, W. Ribarsky, and R. Chang. Interactive coordinated multiple-view visualization of biomechanical motion data. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1383–1390, 2009.

[19] D. F. Keefe, D. Acevedo, J. Miles, F. Drury, S. Swartz, and D. Laidlaw. Scientific sketching for collaborative VR visualization design. *Visualization and Computer Graphics, IEEE Transactions on*, 14(4):835–847, 2008.

[20] D. F. Keefe, T. M. O'Brien, D. B. Baier, S. M. Gatesy, E. L. Brainerd, and D. H. Laidlaw. Exploratory visualization of animal kinematics using instantaneous helical axes. 27(3):863–870, 2008.

[21] R. Kosara. Visualization criticism - the missing link between information visualization and art. In *Information Visualization, 2007. IV '07. 11th International Conference*, pages 631–636, 2007.

[22] H. Li, C.-W. Fu, Y. Li, and A. Hanson. Visualizing large-scale uncertainty in astrophysical data. *IEEE transactions on visualization and computer graphics*, 13(6):1640–7, 2007.

[23] A. Lu and H.-W. Shen. Interactive storyboard for overall time-varying data visualization. *Pacific Visualization Symposium, 2008 IEEE*, 2008.

[24] C. Lundstrom, P. Ljung, A. Persson, and A. Ynnerman. Uncertainty visualization in medical volume rendering using probabilistic animation. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1648–1655, 2007.

[25] C. L. MacKenzie, R. G. Marteniuk, C. Dugas, D. Liske, and B. Eickmeier. Three-dimensional movement trajectories in Fitts' task: Implications for control. *The Quarterly Journal of Experimental Psychology Section A*, 39(4):629–647, 1987.

[26] G. E. Marai, C. Demiralp, S. Andrews, and D. H. Laidlaw. Jointviewer-an interactive system for exploring orthopedic data. In *Proceedings of the*

conference on Visualization'04, pages 598–35. IEEE Computer Society, 2004.

[27] N. J. Mitra, Y.-L. Yang, D.-M. Yan, W. Li, and M. Agrawala. Illustrating how mechanical assemblies work. *ACM Transactions on Graphics*, 29(4):1, July 2010.

[28] G. Pingali, A. Opalach, Y. Jean, and I. Carlbom. Visualization of sports using motion trajectories: Providing insights into performance, style, and strategy. In *Proceedings of the Conference on Visualization '01*, VIS '01, pages 75–82, Washington, DC, USA, 2001. IEEE Computer Society.

[29] H. Piringer, S. Pajer, W. Berger, and H. Teichmann. Comparative visual analysis of 2D function ensembles. *Comp. Graph. Forum*, 31(3pt3):1195–1204, 2012.

[30] D. Riskin, D. Willis, J. Iriarte-Díaz, T. Hedrick, M. Kostandov, J. Chen, D. Laidlaw, K. Breuer, and S. Swartz. Quantifying the complexity of bat wing kinematics. *Journal of theoretical biology*, 254(3):604–615, 2008.

[31] M. J. Roberts. *Underground maps unravelled: Explorations in information design*. 2012.

[32] M. B. Sabick, M. R. Torry, Y.-K. Kim, and R. J. Hawkins. Humeral torque in professional baseball pitchers. *The American Journal of Sports Medicine*, 32(4):892–898, 2004.

[33] A. P. Sawant and C. G. Healey. Visualizing multidimensional query results using animation. In *Electronic Imaging 2008*, pages 680904–680904. International Society for Optics and Photonics, 2008.

[34] D. Schroeder, T. Kowalewski, L. White, J. Carlis, E. Santos, R. Sweet, T. S. Lendvay, T. Reihsen, and D. Keefe. Exploratory visualization of surgical training databases for improving skill acquisition. *IEEE Computer Graphics and Applications*, 32(6):71–81, Nov. 2012.

[35] J. S. Sobel, A. Forsberg, D. H. Laidlaw, R. C. Zeleznik, D. F. Keefe, I. Pivkin, G. E. Karniadakis, P. Richardson, and S. Swartz. Particle flurries. *Computer Graphics and Applications, IEEE*, 24(2):76–85, 2004.

[36] B. Tversky, J. Morrison, and M. Betrancourt. Animation: can it facilitate? *International Journal of Human-Computer Studies*, 57(4):247–262, 2002.

[37] A. N. Vasavada, S. Li, and S. L. Delp. Influence of muscle morphometry and moment arms on the moment-generating capacity of human neck muscles. *Spine*, 23(4):412–422, 1998.

[38] J. Waser, R. Fuchs, H. Ribicic, B. Schindler, G. Bloschl, and E. Groller. World lines. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1458–1467, 2010.