

Visual Analysis of Multivariate Intensive Care Surveillance Data

N. Brich^{1,2} , C. Schulz³ , J. Peter^{2,4,5} , W. Klingert⁵ , M. Schenk⁵ , D. Weiskopf³ , and M. Krone^{1,2} 

¹Big Data Visual Analytics in Life Sciences, University of Tübingen, Germany

²Institute for Bioinformatics and Medical Informatics (IBMI), University of Tübingen, Germany

³Visualization Research Center, University of Stuttgart, Germany

⁴Translational Bioinformatics Research, University of Tübingen, Germany

⁵Department of General, Visceral and Transplant Surgery, University of Tübingen, Germany

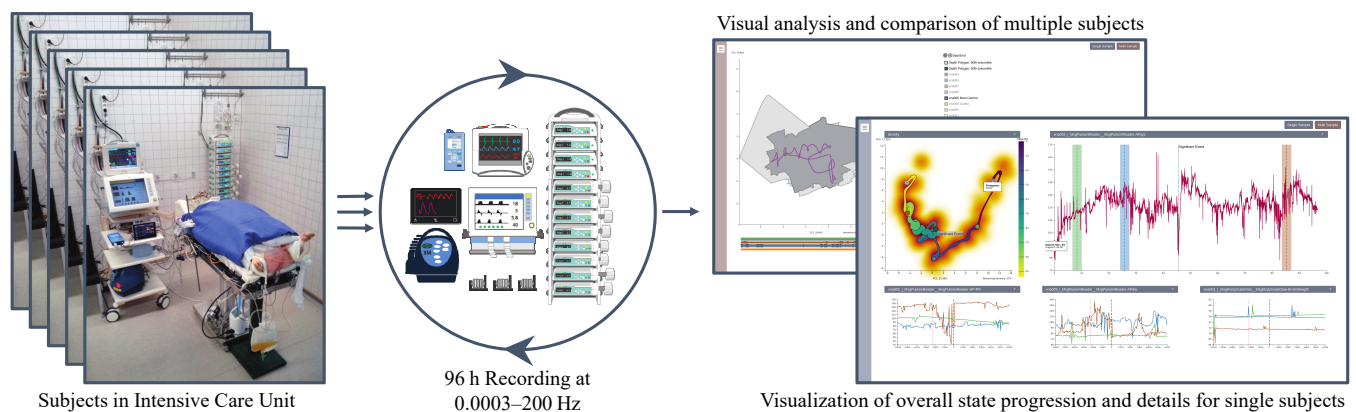


Figure 1: We analyze a multivariate data ensemble obtained during an experiment with pig subjects attached to multiple medical devices in an intensive care unit. The overall state progression over the set of time series per animal is shown as a time-curve-inspired plot. A line plot showing a selected measurement and small multiples composed of selection-dependent line plots allow for state comparison and in-depth analysis. For an ensemble of multiple subjects, a time curve boxplot allows to compare subjects. Photograph courtesy of Peter et al. [Pet18].

Abstract

We present an approach for visual analysis of high-dimensional measurement data with varying sampling rates in the context of an experimental post-surgery study performed on a porcine surrogate model. The study aimed at identifying parameters suitable for diagnosing and prognosticating the volume state—a crucial and difficult task in intensive care medicine. In intensive care, most assessments not only depend on a single measurement but a plethora of mixed measurements over time. Even for trained experts, efficient and accurate analysis of such multivariate time-dependent data remains a challenging task. We present a linked-view post hoc visual analysis application that reduces data complexity by combining projection-based time curves for overview with small multiples for details on demand. Our approach supports not only the analysis of individual patients but also the analysis of ensembles by adapting existing techniques using non-parametric statistics. We evaluated the effectiveness and acceptance of our application through expert feedback with domain scientists from the surgical department using real-world data: the results show that our approach allows for detailed analysis of changes in patient state while also summarizing the temporal development of the overall condition. Furthermore, the medical experts believe that our method can be transferred from medical research to the clinical context, for example, to identify the early onset of a sepsis.

CCS Concepts

• **Applied computing** → Health care information systems; • **Mathematics of computing** → Time series analysis; Dimensionality reduction; • **Human-centered computing** → Information visualization;

1. Introduction

In a modern clinical setting, patients in an intensive care unit (ICU) are closely monitored by multiple medical devices. In such a scenario, this multitude of measurements has to be communicated for interpretation, either to medical personnel for progressive analysis (monitoring), or—as in our context—to medical researchers for post-hoc analysis (experiments). However, humans have difficulties aggregating more than three quantities, especially under the presence of uncertainty [GASM17]. On the other hand, multiple parameters are often needed to provide an overall impression of the condition of a patient. Only considering a combination of specific changes in multiple parameters will hint toward important transitions of the whole homeostatic system. For instance, individual parameters might seem uncritical if considered independently, but in conjunction, can indicate an imminent circulatory shock.

Ideally, an illustration of patient state would be based on the entirety of time-varying interacting variables, without having to rely on a set of specific, pre-defined parameters. Moreover, such an approach would reveal high-dimensional processes clearly without being overly specific and still provide fine-grained details on demand. In practice, one of the primary problems with multivariate time-varying data is the conceptual and physical limitation of straightforward visualization techniques: most established approaches, such as scatterplot matrices, do not scale well enough to support the assessment of more than a few variables due to visual clutter and required mental effort to draw the right conclusions. In response to this observation, many expert users resort to univariate visualization techniques and (over-)simplifying derived characteristic quantities that take little advantage of the human visual system. An example of this approach is the commonly used circulatory shock index, which is derived from a set of relevant parameters [CBK*09; BGT*05; RSB*94] but does not provide further information about the underlying condition of the patient.

In this work, we propose a visualization approach that facilitates the exploration and analysis of multi dimensional, time-dependent ICU data based on time curves [BSH*16; vdEHBvW16], small multiples [Tuf92], and comparative views for an ensembles subjects. To this end, we studied various design choices in close collaboration over several months with three medical researchers from the surgical department, who are also co-authors of this work. We showcase our approach using real-world data obtained from an experimental surgery study performed on a porcine surrogate model (cf. Figure 1). To compute our time curves, we have analyzed various projection techniques regarding their suitability for our application. Using a generative data model, we were able to identify varying frequency and oscillation in general as a projection-quality-determining factor. In order to allow the user to explore different subsets of the complete data set, the user can select and deselect data sources and single variables stemming from these sources. To compare patient states and drill down into details, our application allows the user to select multiple data points in time that are then shown using small multiples. The small multiples view is composed of line plots that allow for the assessment of differences within a user-defined temporal context. The contributions within our web-based visualization application can be summarized as follows: the visual analysis approach to high-dimensional and time-varying ICU data, the adaptation of

various visualization techniques using z-score and data-depth-based statistics to ease outlier analysis, a domain-driven generative data model to evaluate the applicability of various dimensionality reduction algorithms to ICU measurements, and a discussion containing expert feedback as well as an outlook for clinical surveillance.

2. Related Work

Although the visualization of high-dimensional and time-varying data has been investigated for more than a decade and has been reviewed in many extensive surveys [AMST11; KH13; HW13; LMW*17], it is still an active area of research. Classical multivariate visualizations such as parallel coordinate plots and scatter plot matrices are either a poor fit for visualizing multivariate time series, or only suitable for specific tasks like finding correlations between single variables. Thus, tailored task-driven visualization methods are needed [Mun15] to facilitate an effective visual analysis of complex multivariate time-varying data. Since we are dealing with a number of time-varying input dimensions exceeding the capabilities of classical multivariate visualization, we subsequently briefly review suitable approaches and techniques.

Dimensionality Reduction A trend for the visualization of multivariate data is to use dimensionality reduction. The basic idea of dimension reduction with regards to multivariate data is the reduction of the total dimension amount while maintaining some or ideally most of the information content of the original data set. This is achieved by deriving new variables summarizing specific parameters of the data. In the most simple form, this derivation is performed by basic calculations generating context-dependent new variables like indices (e.g., the abovementioned shock index [CBK*09; BGT*05; RSB*94]). Since this way of deriving new variables depends on the context of the used variables, it is not particularly useful in a general context. Thus, more general ways of deriving new variables are needed. A typical example of such a generalized way is Principal Component Analysis (PCA) [Hot33; Pea01], which computes a linear projection based on variance. Another widespread technique is Multi-Dimensional Scaling (MDS) [Tor52; BG05], which optimizes for globally representing high-dimensional distances in a low-dimensional space. In all generality, there is neither a perfect solution nor a uniquely defined global optimum to this problem. Thus, a wide range of dimension reduction techniques has been developed. A popular choice is the computationally rather expensive t-Distributed Stochastic Neighbor Embedding (t-SNE) [vdMH08], which focuses on clustering similar data points by optimizing for adjacency and entropy instead of distance, generally leading to a misrepresentation of global distances in the embedding. A more recent approach is Uniform Manifold Approximation and Projection (UMAP) [MHM18], which aims at reducing this misrepresentation of distant data points. Since it is not clear in advance which of these approaches is suitable, our application offers the choice between PCA and MDS. After some initial experiments using t-SNE and UMAP, it quickly became apparent that their performance characteristics do not fit our requirements. This is primarily due to their reduced ability in preserving global distances, as mentioned by McInnes et al. [MHM18]. Moreover, our application guides the data-driven choice of dimension reduction technique, based on approaches from research on quantification and visualization of

projection error [SSK10; HAF13; WSA*16]. To gain a better understanding of the overall suitability of the dimension reduction techniques for our application, we use a domain-specific generative data model for evaluation purposes [SBK*16] (see Subsection 5.2).

Visualization of Multivariate Time Series Our medical visual analysis application was inspired by several techniques to visualize multivariate time series that leverage dimension reduction: Jaeckle et al. [JFSK16] use a technique called temporal-MDS to project non-temporal dimensions into a single dimension while retaining one display dimension for the mapping of the time component. This approach results in an optimal depiction of the progressing time but impedes the distinction of different states since they are distinguished though position on a single axis. However, we want to provide an intuitive overview of differences in patient states. Thus, the utilization of the two major axes dimension for non-temporal components seems like a better choice. Bach et al. [BSH*16] concurred with this assessment and developed the time curve concept. They used MDS to project the high-dimensional data in two dimensions while keeping the temporal succession obvious to the spectator by connecting consecutive points using a Bézier curve. A similar approach was previously presented by Bernard et al. [BWS*12], who used PCA instead of MDS. Concurrently to Bach et al., van den Elzen et al. [vdEHBvW16] developed a similar technique to visualize temporal evolution of graphs. We stick to the term time curve due to its intuitiveness. These two techniques can be classified as connected scatterplots, which follow the same idea of connecting temporally consecutive data points, but do not use dimensionality reduction and date back to 1790 [HKF16]. An example is the work of Grottel et al. [GHWG14], who map time to line density in traditional scatter plots and parallel coordinates. Recently the concept of time curves has been used by Hinterreiter et al. [HSS*20] to analyze solutions in problem solving tasks. We adopt the concept of time curves and integrate it into our exploratory data analysis application. Due to the oscillating nature of the ICU data, we extend the original concept of time curves with filtering. We use this filtering step to declutter the data—and, consequently, the resulting visualization—and to remove outliers. Our application can handle not only single data sets but also multiple ones in order to compare the ensemble members (i.e., pigs). For the analysis of multiple data sets, we extend the time curve visualization with the notion of data depth and curve boxplots [MWK14]. Data depth acts as a measurement to describe how central a single data point is the entirety of a larger data set: the higher a data depth, the more central a given data point is [CN08]. While data depth was already applied to time series in the past [LR09; LSLG14], we explore the application of this concept to high-dimensional time series projected to two dimensions.

Medical Visualization Ordonez et al. [OdL*10] explored the usage of star charts for visualizing ICU data. While this works reasonably well for a small number of variables, visual scalability is limited. In the course of the increasing digitalization, the visualization of electronic health records (EHR) has gained much attention. Zhang et al. [ZMG*10] proposed a visual analytics framework that unifies EHR information for use in emergency rooms. While it also shows time-varying multidimensional data, their approach is not scalable and assumes discrete patient states and transitions, not numerical measurements. Similarly, Rind et al. [RAM*11] presented VisuEx-

plore, a tool for the visual analysis of long-term medical records of patients with chronic diseases. Their application case is different from ours since their data is more sparsely sampled over much longer periods of time. Ten Caat et al. [CMR05] introduced the concept of tiled parallel coordinates for multichannel EEG data. While their data is more similar to ours, they just have multiple sensors of the same type that all measure at the same frequency. They also use aggregation over the whole measurement time, which is not suitable in our case.

3. Medical Background

In this section, we provide a summary of the experiment and required clinical background knowledge—details about the framework and experimental setup can be found in works by Peter et al. [PKK*14; Pet18; KPT*18]. Preliminary research showed insufficient knowledge regarding the assessment of the volume state (i.e., the state of an organism with regards to the amount of fluids inside the body) in clinical practice for hospitalized patients. Our research was guided by observations in volume state changes of individual pigs during the previously conducted *auto pilot study* (AP). The subsequent *volume-need-analysis study* (VNA) was conducted to answer research questions regarding the reliable assessment of volume states in a clinical context. For example, in clinical practice, one indicator for a general lack of fluid is poor elasticity of the skin. While testing through pinching is easy and cost-efficient, it requires the patient to be in a hazardous state for the test to work. Thus, one interesting questions is whether it is possible to predict states and potential problems by observing measurable quantities of the organism so that doctors can proactively intervene, instead of being forced to react. To this end, a set of ICU devices commonly found in a clinical context was used to obtain the measurements. However, the setup was far more comprehensive than a regular setup in a hospital.

The VNA study was conducted using 10 fully anesthetized pig surrogates in accordance with ethical guidelines. In total, the experimental setup measured roughly 250 parameters (omitting non-relevant parameters for error logging) over 96 hours, with two experiments being ended prematurely. Recorded parameters include a variety of cardiovascular parameters such as blood pressure and heart rate, blood gas analysis results such as ion and carbohydrate concentrations, as well as general parameters like weight and diuresis volume. All measurements were collected from a set of nine devices via manufacturer interfaces and wiretapping, depending on the device. Each sample was stored as a tuple consisting of the measurand, time, and value. While some devices sample adaptively, the average sampling rate varies between $1/3600$ Hz (blood gas analysis) and 500 Hz (electrocardiogram). All devices are synchronized to an accuracy of approximately one second. The subset of the data that we used consisted of 1.6 GB of CSV files per subject, with each CSV file containing one time series with a size 1 kB to 12 MB.

Our domain scientists want to identify yet unknown relations between the volume state and other physiological parameters to guide further clinical research and development. The experimental setup continuously monitors infusion rates (in-flow), diuresis (out-flow), and the bodyweight of the pigs to obtain almost complete knowledge about the volume state of subjects. Moreover, the setup automatically manages homeostasis (internal steady-state) to en-

sure reproducible conditions. During the study, usage of forced-air warming blankets was implemented and suspect to data collection to measure and control the influence of temperature from the fifth subject onward.

Initially, the animals were given defined amounts of fluid using intravenous infusions to reach an increased fluid state. Then, fluid supply was stopped and the body fluid was reduced through diuretics. The fluid state was assessed at defined states by performing a variety of tests, including a Trendelenburg maneuver (movement compensation test), a respiratory hold test, laboratory analysis, and ultrasound measurements of the diameter of the vena cava (large veins that return deoxygenated blood from the body into the heart).

4. Requirements

The measurement data already defined the quantity structure (see Section 3); therefore, we can focus the requirement analysis on the functionality. The high-level goal of our collaboration is to provide an *orthogonal approach to visual analysis of individual measurements*. In line with our experts, we argue that a loss of information to gain insights into high-dimensional processes is justifiable: as our primary use case is the post-hoc analysis of ICU data, our system cannot cause harm to the animal subjects. Although this is out of scope for the post-hoc analysis use case, our application would be backed up by validated alarm systems in a real-time hospital setting, minimizing potential hazards caused by misinterpretation and loss of information. Thus, the following goals and requirements were defined iteratively during multiple design sessions and on-site meetings together with our medical domain experts:

R1 The main objective is to be able to tell whether the patient is in a stable state or not. Since a patient does not necessarily have to be within the ideal parameter ranges, communication of change rates and state comparison is more important than absolute values.

R2 Both the aggregate of change as well as individual changes can be of interest. Therefore, the visualization application should allow for both a top-down as well as a bottom-up exploration of the data.

R3 Robust statistical methods should be used to guide the search for interesting patterns and relevant behavior—some parameters are heavy-tailed distributed and oscillating.

R4 Medical researchers have high expectations regarding sanity-checking and explanation of phenomena due to the nature of the medical domain—easy interpretation without a strong mathematical background is of major interest.

R5 Lastly, it should be possible to put individual patients in the context of the entire ensemble to assess whether a patient is representative or exceptional.

5. Visualization of ICU Data

Our goal was to combine several visualization approaches in one application tailored to medical experts. In this section, we describe the techniques we use as components and their adaption to our specific application domain. In accordance with the visualization pipeline [Mor13], this includes data cleansing, an evaluation of dimensionality reduction algorithms, and a description of the visual mappings, which consist of filtered time curves and small multiples composed of line plots for in-depth analysis and comparison of

patient states, as well as time curve boxplots for ensemble analysis (Figure 1).

5.1. Data Cleansing and Preprocessing

We cleanse our data beforehand since our users are not interested in raw data and unambiguous measurement errors. Given a set of raw time series, we drop categorical data and time series that have a constant value, since these are not of interest. Moreover, we remove erroneous, as defined by a list of user defined error codes, and filter inaccurate samples. To identify and remove such outliers, we have explored various techniques, such as the Savitzky-Golay filter [SG64] that approximates a variable using a polynomial. While this filter removes severe outliers, it also introduces artificial values and thus conflicts with R4. Ultimately, we opted for a z-score-based filter that classifies values as outliers that are not contained in the hull spanned by the moving average and standard deviation. In particular, we use z-score filtered data to make non-robust visualization techniques more resilient regarding outliers and oscillating data values.

Apart from the abovementioned robustness, all considered dimensionality reduction techniques require a uniform sampling, i.e., all components of a high-dimensional point have to be defined. Due to the high variance in sampling rates (ranging from 1/3600 Hz for the blood gas analysis to 3 Hz for the heart rate in our data subset), a simple downsampling approach would lead to massive loss of data. Note that we do not use linear interpolation to fill in missing values since this would introduce an invalid assumption about the measured quantities. Therefore, we decided on a more data-driven, assumption-free strategy to unify sampling rates across all time series by forward-filling missing values (i.e. the last valid value will be used until the next valid value is reached). Our medical experts also approved of this strategy. Thus, we effectively down-sample time series with a higher sampling rate than a chosen reference sampling rate and up-sample time series with a lower sampling rate, respectively (i.e., database left join). As a result, our entire data processing pipeline is also in line with R4, as it does not modify historical samples.




Figure 2: Synthetic sample data set, consisting of six curves which each corresponding to a single synthetic subject, of our generative data model projected using PCA (left) and MDS (right). For two subjects a constant value was added to the function, while for another two subjects, a constant value was subtracted. Different colors are used for easier visual separation. Note the visually prominent loops as well as the increased distortion at beginning and end of the MDS time curves. The filled circles (●) correspond to the time curves starting point and the dotted circles (○) to their endpoint.




Figure 3: Extract from the measurements of our dimension reduction evaluation. (a) The plot shows the influence of variation of model parameters with increasing amounts of standard deviation (SD) on the x-axis for MDS and PCA. Visually speaking, SD is the degree of deviation around an ensemble of randomly chosen centroids. The y-axis shows the percentage of maximum remaining variance not embodied by PC1 and PC2 for PCA, and the percentage of maximum stress induced by MDS. As observable, varying frequency is the decisive factor for projection quality with both methods. The frequency histograms at SD 0.1 (b) and SD 0.9 (c) over $t/2$ are shown for reference.

5.2. Dimensionality Reduction

Time curves project a time-varying high-dimensional space into a two-dimensional space, where time is not part of the dimensionality reduction, i.e., $(t, \mathbb{R}^n) \mapsto (t, \mathbb{R}^2)$. Time is added again after the dimensionality reduction by connecting the points in chronological order using a line. Therefore, we had to choose a suitable dimension reduction technique. However, the VNA data set is an experimental data set, containing unknown events, trends, and patterns. Thus, we use a generative data model [SBK*16] to evaluate the chosen dimensionality reduction. Our data generator emulates the properties of the expected real-world multivariate time series using a combination of sinusoidal functions and piecewise linear ones. To justify the validity of our generator, we refer to the parameter range of the devices used during the experiment, as well as expert feedback. Thus, the resulting time series are representative regarding experimental measurements, as they contain variables with repeating patterns of varying frequency and amplitude (e.g., pulse, respiratory monitoring) in conjunction with aperiodic, linearly changing quantities (e.g., temperature, weight). This approach allows us to establish a link between the parameters of the generative model and the output of the dimensionality reduction algorithm. This link is established through an evaluation metric: for PCA, it is the length of the first two components; for MDS, it is stress. This form of sensitivity analysis would not be possible using only a few data sets with unknown properties. Moreover, the generative data model helps us to do sanity checking and allows users to familiarize themselves with the application using easy-to-understand data (R4).

The example shown in Figure 2 uses six synthetic time series to illustrate differences between PCA and MDS. In order to model physiological processes, which often exhibit oscillating behavior, the synthetic data sets are sinusoidal curve ensembles with varying frequency, amplitude and phase. In general, MDS represents long distances better than short ones, causing the beginning and end to be slightly deformed. Moreover, MDS tends to introduce loops that are not present in the data, which might be caused by non-optimal flipping and rotation while iterating (unfavorable local optima). PCA, on the other hand, shows the oscillating behavior of the data while separating dissimilar clusters of time series. Figure 3 shows the relative robustness of PCA and MDS regarding variation of the model parameters: frequency, amplitude, and phase. While

amplitude and phase have little impact on projection quality, varying frequencies quickly degrades projection quality. Considering both aspects, we rate PCA as the suitable default and more defensive choice for dimensionality reduction in our application (R3).

Our recommendation to use PCA is motivated by caution and not in line with the time curve implementations of Bach et al. [BSH*16] and Van Elzen et al. [vdEHBvW16], who use MDS for their time curves. MDS minimizes stress $\sigma = \sum_{i < j} w_{ij} (\hat{d}_{ij} - d_{ij})^2$, where w_{ij} is a weight (usually set to 1.0), d_{ij} is the high-dimensional Euclidean distance, and \hat{d}_{ij} is the two-dimensional Euclidean distance. Minimizing σ has several drawbacks: First, small distances are less well represented than large distances. Second, distance is optimized globally, which is undesirable in higher-dimensional spaces since all distances tend to become equally near and far apart—this is also known as the curse of dimensionality [HAK00]. Third, most algorithms that minimize stress do so in a point-centric manner, e.g., SMACOF [LM09]. The last issue often leads to suboptimal local optima without proper and expensive initialization of the optimization process. In contrast, PCA avoids the first and last issue entirely while being linear, which is favorable regarding R4. We assume that issues regarding MDS are less prominent in the use cases of Bach et al. and Van Elzen et al., since there is less oscillation in their data than in our use case.

The ensemble analysis (R5) requires a collective dimensionality reduction of multiple subjects. Here, we apply Z-transform individually to each subject prior to the PCA. The PCA is then performed for all subjects at once, ensuring identical loadings and, in turn, principal components, allowing for comparisons between the subjects (see Figure 4). Consequently, and for a meaningful comparison, the same set of dimensions must be used for each subject.

5.3. Time Curve Boxplots and Scarf Boxplots

When plotting multiple time curves in a single visualization, the result is often a spaghetti plot dominated by visual clutter even for relatively low numbers of curves. Especially the overlapping curves are challenging to assess, rendering it difficult to see which one of the curves is representative. Therefore, we extract a representative hull that encloses a centrally-outward ordered share of curves. To establish this order, we resort to a class of order-statistical metrics called

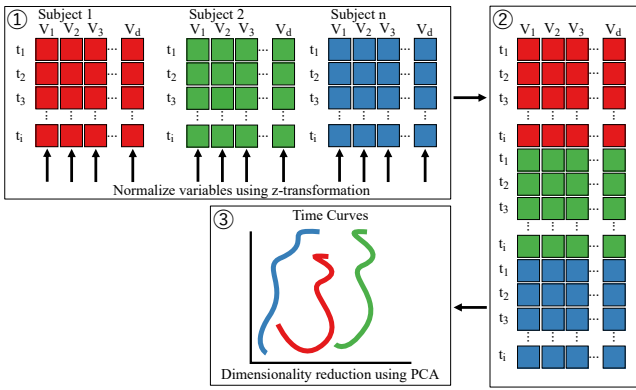


Figure 4: Schematic depiction of the normalization and PCA projection for multiple subjects. Subjects are normalized individually using a z-transformation. The resulting data frames are then concatenated and PCA is subsequently performed on this concatenated data frame ensuring identical loadings allowing comparison of individual subjects.




Figure 5: Schematic depiction of functional band depth. The band (gray), indicating the most representative functions (bounding functions, black). The blue function is included, i.e., belonging to the group of most representative function, whereas the red one is not included, and thus is less representative for the ensemble of functions.


data depth [CN08] that describe how central an element is within an ensemble. In this context, centrality can be understood as representativeness. We use a metric called *functional band depth* [LR09], which builds on functional bands. A functional band B_k consist of $k \geq 2$ functions $f_{ir}(t \in I)$ with restriction interval I :

$$B_k \left(\begin{matrix} f_{i1} \\ \vdots \\ f_{ik} \end{matrix} \right) = \{ (t, y) : t \in I, \min_{r \in [1, k]} f_{ir}(t) \leq y \leq \max_{r \in [1, k]} f_{ir}(t) \} \quad (1)$$

A function lies between other functional bands if it is enclosed by the respective minimum and maximum values, as shown in Figure 5. Formally, we can count this condition using an indicator function χ with *true* $\mapsto 1$ and *false* $\mapsto 0$ to define band depth $BD_{n,k}$.

$$BD_{n,k}(f_i) = \sum_{j=2}^k \binom{n}{j}^{-1} \sum_{l=1}^j \chi \left(\{ (t, f_i(t)) \mid t \in I \} \subset B_l \left(\begin{matrix} f_{i1} \\ \vdots \\ f_{il} \end{matrix} \right) \right) \quad (2)$$

This formula first counts how often a function f_i is contained in bands of length l . The number of actual inclusions is then put in relation to the possible ones. The process is repeated from bands of length 2 to up to bands of length k . A smaller k is more sensitive to variation, whereas a larger k is the more resilient to small fluctuation.



(a) time curve ensemble (twelve curves) without time curve boxplot (b) time curve ensemble (twelve curves) with time curve boxplot

Figure 6: Illustration of decluttering and outlier analysis using summarization via functional data depth boxplots. Left: twelve cluttered time curves. Right: the same curves with the upper 50% quantile merged into a boxplot area. The box indicates the set of most representative curves, while the outliers remain individual curves.

Note that Equation 1 can also be expressed as a convex combination, which emphasizes the enclosing nature of the isosurfaces that can be defined on this metric, i.e., our enclosing hull. While data depth has been used to reduce visual clutter and design boxplots for quite some time [SG11; MWK14], it has not been applied to dimension-reduced curves yet to our knowledge. In this application the calculation of the functional band depth is performed on the high dimensional, z-transformed, data. The enclosing hulls are then generated, according to the order established by data depth, from unsmoothed time curves after dimensionality reduction. Thus, included curves do not have to lie entirely in the convex shape of the enclosing hull and the convex hull does not perfectly describe the smoothed curves. Figure 6 shows our time curve boxplot that covers the 50% quantile of most representative time curves, allowing the omission of all time curves included in this quantile. Consequently, recognizing which curves are outliers and which ones are representative for the whole ensemble is easily possible. We additionally can use a second boxplot area covering a wider quantile range, thus decluttering the visualization even further. Figure 7 shows this concept applied to an ensemble of eight curves. By showing the two quantiles (50% and 80%) and the most representative curve (which corresponds to the median), we follow the well-known design of a standard boxplot.

As the containment of a curve in the box area is not only dependent on being at the correct position, but on being at the correct position at a specific point in time, it becomes difficult to judge when a given curve is not included in the curve box area in comparison to another curve at that time. Hence, we provide an accompanying scarf boxplot (see Figure 8) that depicts box inclusion over time. Since the scarfs are temporally aligned, it is easy to judge when a curve is representative and or not in comparison to other curves.

5.4. Filtered Time Curves

Since all dimensionality reduction algorithms emit rather cluttered time curves for real-world as well as generated data, it became clear that we had to reduce the complexity of the curves to cope with visual clutter.

Therefore, we smoothed the emitted data using a Bartlett-window moving average. The width of the moving window determines the degree of smoothing, which can be adjusted by the user. A wider window results in stronger smoothing, with higher deviations from




Figure 7: Visual metaphor of the boxplot as applied to time curves. The cluttered view of eight time curves is simplified by representing curves as areas corresponding to different quantiles: The 50% quantile (dark gray) corresponds to the box of the boxplot, while the 80% quantile (light gray) corresponds to the whiskers. The red time curve is the most representative one (by functional data depth), and thus corresponds to the median line in the boxplot.




Figure 8: Scarf boxplot illustrating representativeness over time from left to right for each subject that is not included in the lesser quantile box plot area. Colors and labels correspond to the ones in the main time curve boxplot, while the dark gray and gray markers indicate inclusion in the respective boxplot area.

the original data, whereas a smaller window in less smoothing and stronger adherence to the data emitted from the dimension reduction algorithm. Smoothing results in a cleaner curve and that is overall visually more appealing. Obviously, filtering introduces a loss of information (R4). Hence, we devised different ways to compensate for this loss of information.

One option is a heat map that shows the density of data for a given region before smoothing (Figure 9 left). We use a Gaussian kernel in conjunction with kernel density estimation to generate a smooth heat map. This map allows assessing how many time points lie in the proximity of a time curve segment and, thus, enables the viewer to draw conclusions about the velocity of changes in the system.

The other option is a mapping of variables onto the time curve itself (Figure 9 right). Since line color was already used to display the elapsed time since the start of the experiment, we chose line width as a secondary channel for information. To ensure comparability with the heat map approach, the density of measurements in the vicinity of a time curve segment can be used as line width. That is, a thicker time curve indicates less change in the patient state (i.e., many projected points that are close together) and vice versa.

5.5. Details Using Small Multiples

The purpose of the time curve visualization described above is to give an overview of the progression of patient state (R1). While

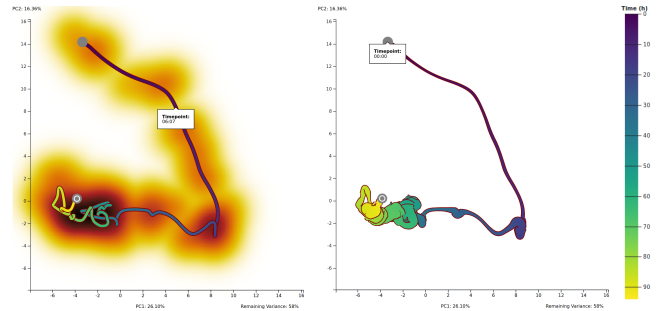


Figure 9: Time curve with constant width and heat map (left) and time curve with data point density in proximity encoded as width (right) of subject VNA004. The curves are color-coded according to the time in hours since the start of the measurement using the viridis color map shown to the right. Both curves are filtered using the average of a moving triangular window. The heat map uses the black body color map (■) [Mor16]. To match the white background, we apply to color mapping inversely, i.e., high values correspond to dark colors, low ones to light colors. When using width encoding a higher width corresponds to a higher density of data points in the proximity

it allows for identifying potentially interesting points in time, it is not intended for a detailed analysis of individual parameters (R2). Given the unavoidable projection error, loss of information, and the sanity checking requirements (R4), we have added a details-on-demand approach that shows small multiples of unprojected data dimensions. For this purpose, there is an auxiliary line plot that depicts values over time and several small line plots that allow for easy superimposition of small timespans. A user may select between 1 to N points in the time curve or auxiliary plot. For each small line plot, the user can select different data dimensions. Each line plot has one polyline per selected point, with the respective time series being centered to the origin (highlighted by a dashed vertical line in the auxiliary plot) based on the selected point (cf. Figure 10). The user can also select the time delta, as multiples of the sampling interval, that is used for the abscissa, which shows the relative time. That is, each polyline shows the temporal development of one data dimension forward and backward in time, starting from the selected point. To allow an easier perception of this temporal interval, it is also drawn as a semi-transparent window extending from the vertical indicators in the auxiliary line plot (cf. Figure 10). The line plots allow the user to get an impression of projection error and temporal change in individual dimensions by comparing the values at different points in time. It also allows comparing the progression of different data dimensions and, consequently, to infer interdependency and correlation. The line plots relate to requirement R2, since they facilitate the individual inspection of each line plot while the small multiples view also allows for investigating localized trends and patterns.

6. Description of the Visualization System

We implemented our visualizations described in Section 5 as a web-based application. The resulting client-server-based approach offers

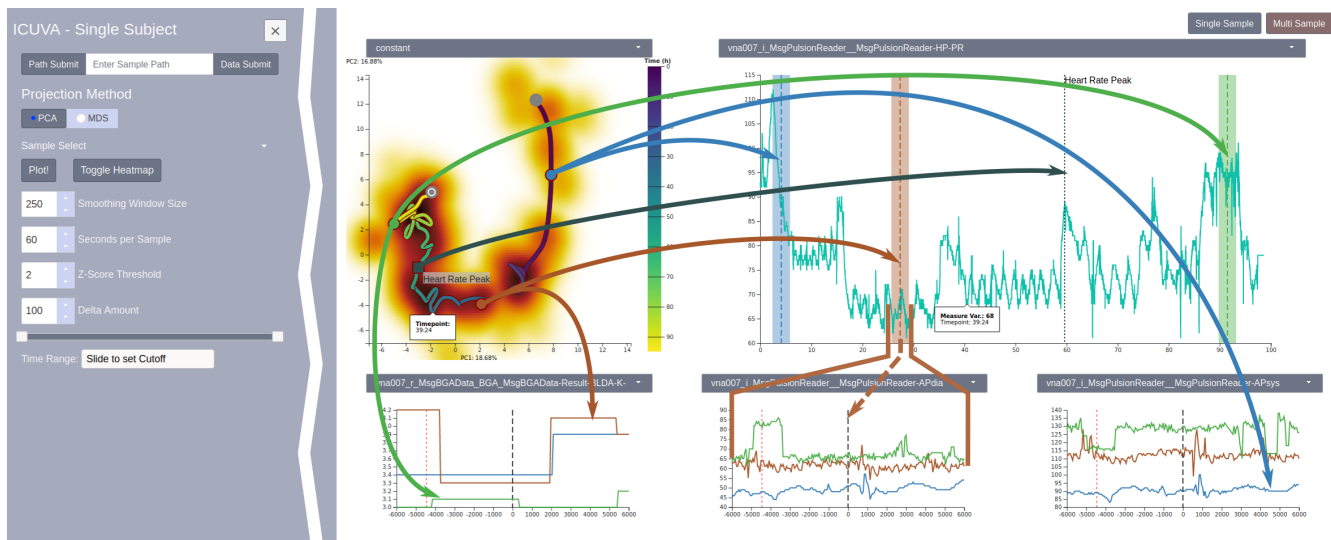


Figure 10: Overview of the single-subject mode (right) and expandable menu (left). The time curve shows patient state progression (top left). An auxiliary line plot shows a single time series (top right). Three samples are selected (brown, blue, and green) to showcase linking. With a dark-grey line and marker indicating linking with custom labels (in this case a local maximum in heart rate was marked and labeled), useful for event highlighting (e.g. peaks in specific measurement variables) Small line plots (bottom) show superimposed time intervals (in seconds) based on selected points; the width of the bar in the auxiliary line plot corresponds to the, identically colored, time segment in the small line plots. The shown variables are, blood K- concentration (as given by the blood gas analysis), arterial diastolic blood pressure and arterial systolic blood pressure (as recorded by the "PulsionReader" device) shown as examples.

several benefits for our use case: First, costly calculation can be performed on the server, removing the need for potent client hardware and enabling the use of relatively thin clients like tablet PCs for viewing and interaction with the visualization system. Second, this approach allows multiple users to access the visualization from different devices simultaneously, which is for example convenient for collaborating medical researchers that are working off-site. Third, data from multiple experiments—recorded at different sites—can be easily observed remotely, potentially saving time and, thus, increasing efficiency. Furthermore, the web-based approach also allows analyzing data from any computer in the hospital, regardless of user permissions, which is convenient in everyday use.

The decision to implement our system as a web application influenced the selection of the programming languages and associated packages. The server was implemented using Python and the Flask framework [MRLU] along with several other packages for data wrangling. Flask was used as a means to serve and communicate with a thin HTML5 single-page client, effectively hiding data management and computationally expensive calculations on the server-side. The data was stored using Pandas [McK*10] data frames for filtering and sampling. The client-side visualization and UI was done using HTML5, JavaScript and the D3 package [BOH11].

The system features two distinct, but similarly structured views: the first view is meant for analyzing a single subject and offers features for an in-depth post-hoc analysis. The second view is tailored to the analysis of an ensemble and, thus, provides features for the comparison of multiple subjects.

Figure 10 shows the single-subject mode. The collapsible menu

on the left is used to select a data set and configure most of the visualization pipeline. This includes settings that influence data processing and selection, such as the data source, restriction to specific variables, and the filtering. Moreover, the user can choose between PCA and MDS for dimensionality reduction; however, in line with our evaluation presented in Subsection 5.2, PCA is the default setting. Furthermore, the interval of the time curve used for dimensionality reduction can be shortened using a slider element.

The single-subject mode features a patient state progression visualized as a filtered time curve, a large auxiliary line plot, and several small selection-relative line plots. A drop-down menu above the time curve visualization allows for the selection of a variable contained in the data set that is encoded onto the width of the time curve. Alternatively, constant width and spatial point density can be chosen. To scrutinize the progression of individual variables, a large auxiliary line plot is shown next to the filtered time curve. The visualized variable can be chosen in the drop-down menu above the plot. Below, small selection-centric line plots are shown. The time curve and the auxiliary line plot are linked views that support hover and click events to maximize the ability to pinpoint specific events in the patient state progression. To support this drill-down process, we provide interactions to annotate time points with highlights and labels. Selecting points on either the time curve or in the auxiliary line plot will determine the point in time that is shown in the small selection-relative line plots. These plots show the temporal surroundings of the selected time points over the chosen variables.

The multi-subject mode has an analogous menu for data loading and selection and also shows a time curve plot as the main view.

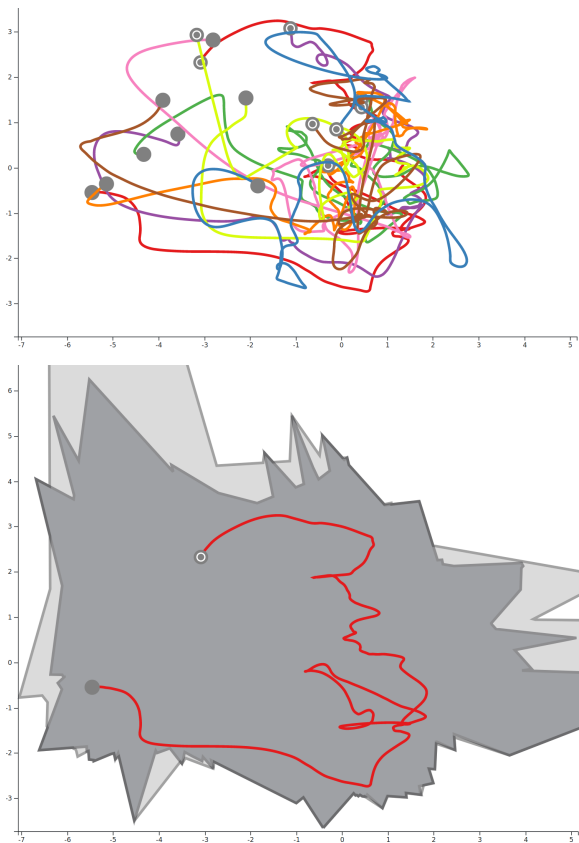


Figure 11: An ensemble of time curves. To prevent biasing the projection, all parameters were normalized per subject using the z -transform. Top: eight superimposed animal subjects shown as filtered time curves. Bottom: a de-cluttered representation of the same data as time curve boxplot. The most representative subject is shown as red filtered time curve. To provide additional context, the unfiltered 50% and 80% quantiles are shown as dark-gray and light-gray areas. Additionally, a scarf plot as shown in Figure 8 is generated for the timecurve boxplot

This plot can visualize an ensemble of (filtered) time curves as well as the time curve boxplot. It also supports toggling the visibility of individual ensemble members and the enclosing curve boxplots (cf. Subsection 5.3). Since the superimposed time curves depict the similarity of the samples instead of a temporal alignment, we show an accompanying scarf plot to illustrate box inclusion over time. The scarf plot located below the time curve plot only visualizes curves that are not included in the inner quantile range (cf. Figure 11). Thus, representativeness can easily be analyzed both in terms of values and time. Additional controls in the menu allow the user to adjust the quantiles for the time curve boxplot computation.

7. Results and Discussion

In this section, we detail several examples and investigations that showcase that our presented visual analysis approach is a promising tool for medical research concerned with ICU data. We also

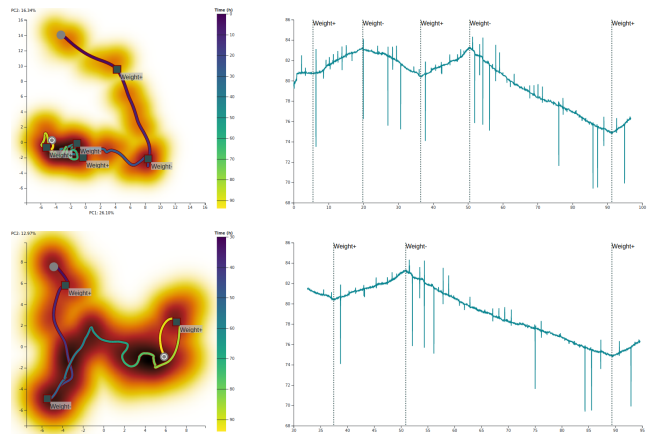


Figure 12: Juxtaposition of the time curve and the net-weight of the subject over the whole 96-hours duration (top) and a shortened time frame of 65 hours (bottom).

highlight how these use cases relate to the user requirements. We evaluated our approach using the data described in Section 3. For the single-subject mode, a set of 82 parameters was used. For the multi-subject mode, we present the results for nine animal subjects using a set of nine relevant parameters each (i.e., 81 parameters in total). This subset of nine out of ten subjects from the VNA study was chosen since one subject had a shorter study duration and was not part of the temperature control experiment (cf. Section 3). Furthermore, we also wanted to investigate the potential of our visual analysis approach together with our domain experts, who are senior researchers in the field of experimental surgery. Feedback for our application was gathered during in-person demonstration meetings with these domain experts, where the computer scientists acted as the *drivers* and the domain experts as the *navigators* [AKGF11].

Single-Subject Mode Several observations were made while investigating the VNA data using our application. With respect to the single subject visualizations, some key aspects of the data can be seen: Figure 12 shows the filtered time curve of subject VNA004. With respect to insights gained about the data, the most prominent feature is the significant and rapid change in status during the first 20 hours, indicated by the comparatively small temporal change as visible by the small change in color and the large difference in state, as visible by the large distance between the points (Patient Stability: R1). While it is difficult to pinpoint the exact reason for this activity, our medical experts have identified this segment of the curve as the onset of anesthetics and the implantation of measuring devices that put a strain on the homeostatic system of the animal. The following seven parameters that might explain this rapid change in the time curve have been identified using auxiliary plot and the small multiples in our application (R2): a peak in heart rate and a relaxation to lower levels during this period; a drop in blood glucose levels accompanied by the start of intravenous glucose administration; an increase in intravenous volume, measured by weight and liquid infusions; low blood pressure, as well as a low blood pH level.

During the VNA study, all subjects were infused with high

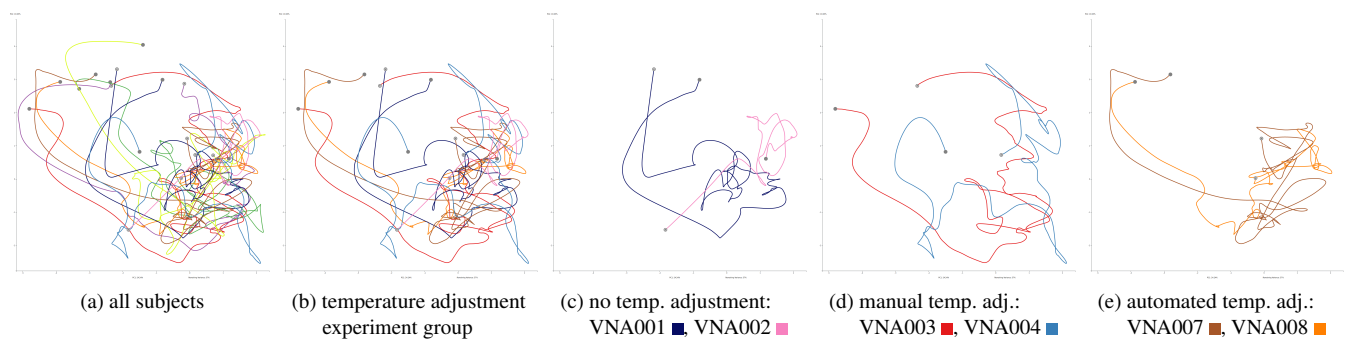


Figure 13: (a) An ensemble of nine animal subjects shown as filtered time curves projected to one space. (b) The six subjects that were part of the temperature control experiment. (c) Subjects VNA001 ■ and VNA002 ■, which did not receive temperature control. (d) Subjects VNA003 ■ and VNA004 ■, which received manual temperature control. (e) Subjects VNA007 ■ and VNA008 ■, which received automated temperature control. Each time curve results from projecting the same nine parameters of each subject.

amounts of liquids over an interval to be then deprived of these liquids. As mentioned above, this change in liquids can be seen in the continuous weight monitoring of the subject. The points at which the administration of liquids is either started or stopped leads to noticeable bends in the time curve. As shown in Figure 12, three of the five administration changes are in a hairball-like region, roughly between hour 32 and 96. In these 64 hours, the observed subject reaches a stable state, indicating that the variables do not change considerably. This stabilization can also be seen when individually observing the heart rate steadied at a lower level, a return to higher levels of glucose in the blood, as well as more physiological blood pressure and pH levels. However, the state of the subject during this period differs significantly from the initial state—illustrated by the large distance between the start of the time curve and this region—which was to be expected.

To disambiguate the hairball, the time curve was restricted to this temporal window of 64 hours (see Figure 12, bottom). The time range can be set by the user via a slider in the menu (see Figure 10). Note that this results in a reapplication of the dimensionality reduction to the original data. This is a deliberate choice as the omission of data will affect the result of the dimensionality reduction. The aforementioned bends in the time curve can be narrowed down to the fact that the amount of liquids excreted from or added to the organism are substantial. The smaller difference amounts to 2.7 kilograms, while the large drop amounts to 8.3 kilograms. Such a drastic change in liquid volume will place a strain on the organism. Thus, it is to be expected that the state as a whole—represented by the time curve—will show a noticeable reaction (cf. R2).

Multi-Subject Mode As shown in Figure 11, the time curve box plot allows the user to assess the representativeness of the subjects (R5). While the red curve shows the central—i.e., the most representative—subject, the scarf plot shows the centrality of the subjects that are not included in the user-adjustable inner quantile range (dark gray area). For an comparison of the patient state progression and similarities between the patients, the user can also choose to display individual curves. Figure 13(a) shows the multi-subject view, where each time curve represents one of the nine animal subjects. Six of these nine subjects were part of a temper-

ature control experiment [Pet18]: VNA001 ■ and VNA002 ■ had no additional temperature control management, VNA003 ■ and VNA004 ■ had their temperature adjusted manually by ICU staff, and VNA007 ■ and VNA008 ■ were subject to automated temperature control using a hot air blanket. As subject VNA002 ■ had a reduced experiment duration of 57 hours [Pet18]—compared to 69 hours for the remaining subjects—the analysis is also restricted to this 57-hours window.

When hiding the four subjects that were not part of the temperature control experiments, the resulting picture is already less cluttered (Figure 13b). The starting point of three of the six subjects (VNA003 ■, VNA007 ■, VNA008 ■) are grouped notably in the left of the figure, whereas two others are close (VNA001 ■, VNA004 ■), indicating the similarity of the initial conditions for these subjects. The last one (VNA002 ■) starts more to the right of the figure, closer to the end points of the other curves.

When turning the focus to the experiment groups as mentioned above, an even clearer trend is visible. The group for which the temperature was not adjusted (VNA001 ■ and VNA002 ■) exhibit rather different progressions of their respective time curves (Figure 13c). The time curves of the manually adjusted group (VNA003 ■ and VNA004 ■) show a more similar trend (Figure 13d). Finally, the progression of the subjects with automatically adjusted temperature (VNA007 ■ and VNA008 ■) is the most similar (Figure 13e). It is worth mentioning that the time curves of the manually and automatically temperature-adjusted subjects resemble each other, while the unadjusted subjects exhibit a vastly different progression.

Although temperature is only one of the nine included parameters, it seems that it is possible to group the subjects by the applied temperature adjustment. This might be due to the fact that body temperature plays a pivotal role in the homeostasis of an organism and, thus, has an impact on a wide variety of other parameters (e.g., blood pressure or heart rate). Failing to adequately adjust the temperature might lead to an imbalance in the homeostasis, which is in turn visible in the time curve representing the organism state.

Expert Feedback Overall, the time curves were intuitively interpreted as a state in a two-dimensional plane showing the similarity

and progression of the patient. This interpretation relates to requirement R1, due to the nature of the medical domain. Thinking in terms of relative change and progression is a common and natural concept for medical experts (e.g., patient health is declining or improving). Therefore, a suggestion for a future improvement of our application was that it would be useful to automatically annotate the time curve with landmarks that provide a reference for interpretation (e.g., the automatic administration of drugs). The choice of the underlying mathematical structure of the projection was not considered crucial for the interpretation. The possibility to investigate and compare the values of individual parameters at salient points in the time curve using the small multiples view was also rated as beneficial. This is in line with the sanity checking requirement (R4) and the need for aggregated as well as detailed information (R2). An earlier suggestion by the medical experts has already been implemented: the smoothing of time curve segments via averaging (see [Subsection 5.4](#)).

Although our current application is tailored to the post-hoc analysis of ICU data, the experts rated the potential of our application for patient surveillance using live data streaming as very high. In particular, they proposed to build a *sepsis radar* based on our application that allows inferring a life-threatening medical condition from the time curves. Sepsis is one of the leading causes of death in hospitals. Therefore, the early detection of the onset of sepsis is of prime interest.

Time Curve Perception While there is a clear benefit of utilizing two dimensions for the differentiation of non-temporal variables, the recognition of time intervals becomes more difficult compared to visualizations that map the time to one of the major axes. Therefore, we color-coded the curve by the elapsed time. Additionally, a tooltip with the exact time is shown when hovering the curve.

We used an averaged progression in lieu of the Bézier curves utilized by Bach et al. to connect the data points and indicate a flow of time. For our application, this approach is viable since we have a high number of sample points, requiring us to rather filter points instead of adding new ones via interpolation. Furthermore, the averaged progression is not required to pass through each data point, which increases projection quality for highly varying frequencies. As discussed in [Subsection 5.2](#), we compared our dimensionality reduction results using PCA to the ones of Bach et al. and Van Elzen et al., who used MDS. This comparison to previous approaches led to the conclusion that the characteristics of the input data need to be considered when choosing a suitable projection methods.

Applicability and Transfer to Other Domains Although our presented visualization application is tailored to the medical use case of ICU data analysis, we believe that some of the concepts we developed are also applicable in other domains. Applying data depth to an ensemble of time curves for decluttering and judging representativeness can be useful, e.g., in the visual analysis of performance workload data from an ensemble of high-performance computing (HPC) cluster nodes. This could show which of the nodes exhibits abnormal behavior. Furthermore, for data that contains high frequencies that are less important, where the user needs to see the overall trend, our proposed simplification and smoothing of time curves can be valuable. An example would be an ensemble of companies set into a high-dimensional context (e.g., stock price, sales per revenue,

net promoter score, etc.), where seasonal effects could be misleading if the trend over several years is of interest.

8. Conclusion and Outlook

We presented a web-based visual analysis application for multivariate, time-varying intensive care surveillance data. Due to the large amount of data measured in such an intensive care scenario, a comprehensive view on the patient and their progression becomes very difficult and traditional visualization methods like line charts fail to adequately present the data. Our application uses dimensionality reduction to visualize this high-dimensional measurement data as time curves [[BSH*16](#)] to provide an overview of this progression of the state of a patient over time. Due to the nature of the data, which has missing or erroneous values and oscillating parameters with varying sampling rates, we extended the time curves using filtering and smoothing. To prevent misinterpretation when the time curve is highly smoothed and to show the density of the input data, the underlying data can be shown as a heat map in the background. A small multiples view shows details on demand about the individual dimensions as line plots and allows users to select and compare different points in time. Our visual analysis application also provides methods to facilitate the comparison of multiple patients by projecting their individual filtered time curves into one two-dimensional space. Furthermore, the system can depict an ensemble as time curve boxplots that summarize typical time curves while highlighting outliers. The time curve boxplots are complemented by a scarf boxplot that shows at which points in time the outlier curves deviate from the box.

The design of our application was developed in close collaboration with medical researchers. Requirements were defined to ensure the effectiveness of the employed visualizations and concepts. We evaluated our application with real-world data and iteratively improved it based on feedback from the medical experts. As discussed in [Section 7](#), the domain experts rated our application as useful and see the definite potential for further developments.

We have shown the utility of our application for medical research using ICU data from a controlled animal experiment. As a next step, we want to apply our application to clinical data collected during the 2020 COVID-19 pandemic in the local university clinic. We hope to identify groups of patients with varying severity of the progression of the disease. Linking these groups to measured quantities might lead to insights into how to improve the treatment and, consequently, the chances of survival for future patients.

In the future, we also want to extend the scope of our application to live data, which would allow for live monitoring of patients in the hospital. To this end, the application has to be improved and retargeted in several ways: First, it has to be ensured that the time curve does not change too dramatically when adding new samples and recalculating the dimensionality reduction. For this, one possible avenue is the adaption of projection techniques to allow for the progressive addition of data points. Several interesting approaches to avoid frequent, global recalculations have been proposed for streaming data [[FCS*19](#); [CMX16](#); [GHN13](#)].

Acknowledgments

NB and MK are funded by Carl-Zeiss-Stiftung. CS and DW acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 251654672 – TRR 161.

References

- [AKGF11] ARIAS-HERNANDEZ, RICHARD, KAASTRA, LINDA T., GREEN, TERA M., and FISHER, BRIAN. “Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics”. *44th Hawaii International Conference on System Sciences*. 2011, 1–10. DOI: [10.1109/HICSS.2011.3399](https://doi.org/10.1109/HICSS.2011.3399).
- [AMST11] AIGNER, WOLFGANG, MIKSCH, SILVIA, SCHUMANN, HEIDRUD, and TOMINSKI, CHRISTIAN. *Visualization of Time-Oriented Data*. Human-Computer Interaction Series. London: Springer, 2011. ISBN: 978-0-85729-078-6. DOI: [10.1007/978-0-85729-079-3_2](https://doi.org/10.1007/978-0-85729-079-3_2).
- [BG05] BORG, I. and GROENEN, P. J. F. *Modern Multidimensional Scaling: Theory and Applications*. 2nd ed. Springer Series in Statistics. New York: Springer, 2005. ISBN: 978-0-387-25150-9. DOI: [10.1007/0-387-28981-X_2](https://doi.org/10.1007/0-387-28981-X_2).
- [BGT*05] BIRKHAHN, ROBERT H, GAETA, THEODORE J, TERRY, DOUG, et al. “Shock index in diagnosing early acute hypovolemia”. *The American Journal of Emergency Medicine* 23.3 (2005), 323–326. DOI: [10.1016/j.ajem.2005.02.029](https://doi.org/10.1016/j.ajem.2005.02.029).
- [BOH11] BOSTOCK, MICHAEL, OGIEVETSKY, VADIM, and HEER, JEFFREY. “D³ data-driven documents”. *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), 2301–2309. DOI: [10.1109/TVCG.2011.1858](https://doi.org/10.1109/TVCG.2011.1858).
- [BSH*16] BACH, BENJAMIN, SHI, CONGLEI, HEULOT, NICOLAS, et al. “Time Curves: Folding Time to Visualize Patterns of Temporal Evolution in Data”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), 559–568. DOI: [10.1109/TVCG.2015.2467851](https://doi.org/10.1109/TVCG.2015.2467851) 2, 3, 5, 11.
- [BWS*12] BERNARD, JÜRGEN, WILHELM, NILS, SCHERER, MAXIMILIAN, et al. “TimeSeriesPaths: Projection-Based Explorative Analysis of Multivariate Time Series Data”. *Journal of WSCG*. 2012, 97–106 3.
- [CBK*09] CANNON, CHAD M, BRAXTON, CARLA C, KLING-SMITH, MENDY, et al. “Utility of the shock index in predicting mortality in traumatically injured patients”. *Journal of Trauma and Acute Care Surgery* 67.6 (2009), 1426–1430. DOI: [10.1097/TA.0b013e3181bbf728](https://doi.org/10.1097/TA.0b013e3181bbf728) 2.
- [CMR05] CAAT, MICHAEL TEN, MAURITS, N. M., and ROERDINK, J. B. T. M. “Tiled Parallel Coordinates for the Visualization of Time-Varying Multichannel EEG Data”. *EUROVIS 2005: Eurographics / IEEE VGTC Symposium on Visualization*. Ed. by BRODLIE, KEN, DUKE, DAVID, and JOY, KEN. The Eurographics Association, 2005. DOI: [10.2312/VisSym/EuroVis05/061-068](https://doi.org/10.2312/VisSym/EuroVis05/061-068) 3.
- [CMX16] CHENG, SHENGHUI, MUELLER, KLAUS, and XU, WEI. “A framework to visualize temporal behavioral relationships in streaming multivariate data”. *2016 New York Scientific Data Summit (NYSDS)*. New York, NY, USA: IEEE, 2016, 1–10. DOI: [10.1109/NYSDS.2016.7747808](https://doi.org/10.1109/NYSDS.2016.7747808) 11.
- [CN08] CUESTA-ALBERTOS, J.A. and NIETO-REYES, A. “The random Tukey depth”. *Computational Statistics and Data Analysis* 52.11 (2008), 4979–4988. DOI: [10.1016/j.csda.2008.04.021](https://doi.org/10.1016/j.csda.2008.04.021) 3, 6.
- [FCS*19] FUJIWARA, TAKANORI, CHOU, JIA-KAI, SHILPIKA, et al. “An Incremental Dimensionality Reduction Method for Visualizing Streaming Multidimensional Data”. *IEEE Transactions on Visualization and Computer Graphics* (2019). DOI: [10.1109/TVCG.2019.2934433](https://doi.org/10.1109/TVCG.2019.2934433) 11.
- [GASM17] GREIS, MIRIAM, AVCI, EMRE, SCHMIDT, ALBRECHT, and MACHULLA, TONJA. “Increasing Users’ Confidence in Uncertain Data by Aggregating Data from Multiple Sources”. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. 2017, 828–840. DOI: [10.1145/3025453.3025998](https://doi.org/10.1145/3025453.3025998) 2.
- [GHN13] GANSNER, EMDEN R., HU, YIFAN, and NORTH, STEPHEN C. “Interactive Visualization of Streaming Text Data with Dynamic Maps”. *Journal of Graph Algorithms and Applications* 17.4 (2013), 515–540. DOI: [10.7155/jgaa.0030211](https://doi.org/10.7155/jgaa.0030211).
- [GHWG14] GROTTTEL, SEBASTIAN, HEINRICH, JULIAN, WEISKOPF, DANIEL, and GUMHOLD, STEFAN. “Visual Analysis of Trajectories in Multi-Dimensional State Spaces”. *Computer Graphics Forum* 33.6 (2014), 310–321. DOI: [10.1111/cgf.12352](https://doi.org/10.1111/cgf.12352) 3.
- [HAF13] HEULOT, N, AUPETIT, M, and FEKETE, J-D. “ProxiLens: Interactive Exploration of High-Dimensional Data using Projections”. *EuroVis Workshop on Visual Analytics using Multidimensional Projections*. Ed. by AUPETIT, M and van der MAATEN, L. The Eurographics Association, 2013. DOI: [10.2312/PE.VAMP.VAMP2013.011-015](https://doi.org/10.2312/PE.VAMP.VAMP2013.011-015) 3.
- [HAK00] HINNEBURG, ALEXANDER, AGGARWAL, CHARU C., and KEIM, DANIEL A. “What is the nearest neighbor in high dimensional spaces?”. *Proc. of the 26th Internat. Conference on Very Large Databases, Cairo, Egypt*. 2000, 506–515. DOI: [10.5555/645926.671675](https://doi.org/10.5555/645926.671675) 5.
- [HKF16] HAROZ, STEVE, KOSARA, ROBERT, and FRANCONERI, STEVEN L. “The Connected Scatterplot for Presenting Paired Time Series”. *IEEE Transactions on Visualization and Computer Graphics* (2016). DOI: [10.1109/TVCG.2015.2502587](https://doi.org/10.1109/TVCG.2015.2502587) 3.
- [Hot33] HOTELLING, HAROLD. “Analysis of a complex of statistical variables into principal components.” *Journal of educational psychology* 24.6 (1933), 417. DOI: [10.1037/h0071325](https://doi.org/10.1037/h0071325) 2.
- [HSS*20] HINTERREITER, ANDREAS, STEINPARZ, CHRISTIAN, SCHÖFL, MORITZ, et al. “Exploring Visual Patterns in Projected Human and Machine Decision-Making Paths”. *arXiv:2001.08372 [cs]* (Jan. 20, 2020). arXiv: [2001.08372](https://arxiv.org/abs/2001.08372). (Visited on 08/17/2020) 3.
- [HW13] HEINRICH, JULIAN and WEISKOPF, DANIEL. “State of the Art of Parallel Coordinates”. *Eurographics 2013 - State of the Art Reports*. Ed. by SBERT, M. and SZIRMAY-KALOS, L. The Eurographics Association, 2013. DOI: [10.2312/conf/EG2013/stars/095-116](https://doi.org/10.2312/conf/EG2013/stars/095-116) 2.
- [JFSK16] JACKLE, DOMINIK, FISCHER, FABIAN, SCHRECK, TOBIAS, and KEIM, DANIEL A. “Temporal MDS Plots for Analysis of Multivariate Data”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), 141–150. DOI: [10.1109/TVCG.2015.2467553](https://doi.org/10.1109/TVCG.2015.2467553) 3.
- [KH13] KEHRER, JOHANNES and HAUSER, HELWIG. “Visualization and Visual Analysis of Multifaceted Scientific Data: A Survey”. *IEEE Transactions on Visualization and Computer Graphics* 19.3 (2013), 495–513. DOI: [10.1109/TVCG.2012.110](https://doi.org/10.1109/TVCG.2012.110) 2.
- [KPT*18] KLINGERT, WILFRIED, PETER, JÖRG, THIEL, CHRISTIAN, et al. “Fully automated life support: an implementation and feasibility pilot study in healthy pigs”. *Intensive Care Medicine Experimental* 6.1 (2018), 2. DOI: [10.1186/s40635-018-0168-3](https://doi.org/10.1186/s40635-018-0168-3) 3.
- [LM09] LEEUW, JAN DE and MAIR, PATRICK. “Multidimensional Scaling Using Majorization: SMACOF in R”. *Journal of Statistical Software* 31.1 (2009), 1–30. DOI: [10.18637/jss.v031.i03](https://doi.org/10.18637/jss.v031.i03) 5.
- [LMW*17] LIU, SHUSEN, MALJOVEC, DAN, WANG, BEI, et al. “Visualizing High-Dimensional Data: Advances in the Past Decade”. *IEEE Transactions on Visualization and Computer Graphics* 23.3 (2017), 1249–1268. DOI: [10.1109/TVCG.2016.2640960](https://doi.org/10.1109/TVCG.2016.2640960) 2.
- [LR09] LÓPEZ-PINTADO, SARA and ROMO, JUAN. “On the Concept of Depth for Functional Data”. *Journal of the American Statistical Association* 104.486 (2009), 718–734. DOI: [10.1198/jasa.2009.01083](https://doi.org/10.1198/jasa.2009.01083) 3, 6.
- [LSLG14] LÓPEZ-PINTADO, SARA, SUN, YING, LIN, JUAN K., and GENTON, MARC G. “Simplicial band depth for multivariate functional data”. *Advances in Data Analysis and Classification* 8.3 (2014), 321–338. DOI: [10.1007/s11634-014-0166-6](https://doi.org/10.1007/s11634-014-0166-6) 3.
- [McK*10] MCKINNEY, WES et al. “Data structures for statistical computing in python”. *Proceedings of the 9th Python in Science Conference*. Vol. 445. Austin, TX, 2010, 51–56. DOI: [10.25080/Majora-92bf1922-012](https://doi.org/10.25080/Majora-92bf1922-012) 8.

- [MHM18] MCINNES, LELAND, HEALY, JOHN, and MELVILLE, JAMES. “UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction”. *arXiv:1802.03426 [cs, stat]* (2018) 2.
- [Mor13] MORELAND, KENNETH. “A Survey of Visualization Pipelines”. *IEEE Transactions on Visualization and Computer Graphics* 19.3 (2013), 367–378. DOI: [10.1109/TVCG.2012.1334](https://doi.org/10.1109/TVCG.2012.1334).
- [Mor16] MORELAND, KENNETH. “Why We Use Bad Color Maps and What You Can Do About It”. *Electronic Imaging* 2016.16 (2016), 1–6. DOI: [10.2352/ISSN.2470-1173.2016.16.HVEI-1337](https://doi.org/10.2352/ISSN.2470-1173.2016.16.HVEI-1337).
- [MLRU] MÖNNICH, ADRIAN, RONACHER, ARMIN, LORD, DAVID, and UNTERWADITZER, MARKUS. *Flask*. <https://palletsprojects.com/p/flask/> (last accessed 29.04.2020) 8.
- [Mun15] MUNZNER, TAMARA. *Visualization Analysis and Design*. AK Peters Visualization Series. CRC Press, 2015. ISBN: 978-1-4987-5971-7. DOI: [10.1201/b175112](https://doi.org/10.1201/b175112).
- [MWK14] MIRZARGAR, MAHSA, WHITAKER, ROSS T., and KIRBY, ROBERT M. “Curve Boxplot: Generalization of Boxplot for Ensembles of Curves”. *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), 2654–2663. DOI: [10.1109/TVCG.2014.234645536](https://doi.org/10.1109/TVCG.2014.234645536).
- [Od*10] ORDÓÑEZ, PATRICIA, DESJARDINS, MARIE, LOMBARDI, MICHAEL, et al. “An Animated Multivariate Visualization for Physiological and Clinical Data in the ICU”. *Proceedings of the 1st ACM International Health Informatics Symposium*. IHI '10. Arlington, Virginia, USA: Association for Computing Machinery, 2010, 771–779. DOI: [10.1145/1882992.18831093](https://doi.org/10.1145/1882992.18831093).
- [Pea01] PEARSON, KARL. “LIII. On lines and planes of closest fit to systems of points in space”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), 559–572. DOI: [10.1080/147864401094627202](https://doi.org/10.1080/147864401094627202).
- [Pet18] PETER, JÖRG. “Assessment of goal-directed closed-loop management in intensive care medicine”. Dissertation. Tübingen, 2018. URL: <https://publikationen.uni-tuebingen.de/xmlui/handle/10900/84091> (visited on 11/22/2019) 1, 3, 10.
- [PKK*14] PETER, JÖRG, KLINGERT, WILFRIED, KÖNIGSRÄINER, ALFRED, et al. “TICoMS – A Modular and Message-Based Framework for Monitoring and Control of Medical Devices”. *2014 IEEE 27th International Symposium on Computer-Based Medical Systems*. 2014, 473–474. DOI: [10.1109/CBMS.2014.963](https://doi.org/10.1109/CBMS.2014.963).
- [RAM*11] RIND, ALEXANDER, AIGNER, WOLFGANG, MIKSCH, SILVIA, et al. “Visual Exploration of Time-Oriented Patient Data for Chronic Diseases: Design Study and Evaluation”. *Information Quality in e-Health*. Ed. by HOLZINGER, ANDREAS and SIMONIC, KLAUS-MARTIN. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2011, 301–320. DOI: [10.1007/978-3-642-25364-5_223](https://doi.org/10.1007/978-3-642-25364-5_223).
- [RSB*94] RADY, MOHAMED Y, SMITHLINE, HOWARD A, BLAKE, HEIDI, et al. “A comparison of the shock index and conventional vital signs to identify acute, critical illness in the emergency department”. *Annals of Emergency Medicine* 24.4 (1994), 685–690. DOI: [10.1016/S0196-0644\(94\)70279-92](https://doi.org/10.1016/S0196-0644(94)70279-92).
- [SBK*16] SCHULZ, CHRISTOPH, BUTT, MIRIAM, KEIM, DANIEL A, et al. “Generative Data Models for Validation and Evaluation of Visualization Techniques”. *Proceedings of the Beyond Time and Errors on Novel Evaluation Methods for Visualization - BELIV*. ACM Press, 2016, 112–124. DOI: [10.1145/2993901.299390735](https://doi.org/10.1145/2993901.299390735).
- [SG11] SUN, YING and GENTON, MARC G. “Functional Boxplots”. *Journal of Computational and Graphical Statistics* 20.2 (Jan. 2011), 316–334. DOI: [10.1198/jcgs.2011.092246](https://doi.org/10.1198/jcgs.2011.092246).
- [SG64] SAVITZKY, ABRAHAM. and GOLAY, M. J. E. “Smoothing and Differentiation of Data by Simplified Least Squares Procedures”. *Analytical Chemistry* 36.8 (1964), 1627–1639. DOI: [10.1021/ac60214a0474](https://doi.org/10.1021/ac60214a0474).
- [SSK10] SEIFERT, CHRISTIN, SABOL, VEDRAN, and KIENREICH, WOLFGANG. “Stress Maps: Analysing Local Phenomena in Dimensionality Reduction Based Visualizations”. *European Symposium Visual Analytics Science and Technology (EuroVAST)*. 2010. DOI: [10.2312/PE/EuroVAST/EuroVAST10/013-0183](https://doi.org/10.2312/PE/EuroVAST/EuroVAST10/013-0183).
- [Tor52] TORGERSON, WARREN S. “Multidimensional scaling: I. Theory and method”. *Psychometrika* 17.4 (1952), 401–419. DOI: [10.1007/BF022889162](https://doi.org/10.1007/BF022889162).
- [Tuf92] TUFTE, EDWARD ROLF. *The visual display of quantitative information*. Graphics Press, 1992. ISBN: 978-0-9613921-0-9 2.
- [vdEBvW16] Van den ELZEN, STEF, HOLTEN, DANNY, BLAAS, JORIK, and van WIJK, JARKE J. “Reducing Snapshots to Points: A Visual Analytics Approach to Dynamic Network Exploration”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), 1–10. DOI: [10.1109/TVCG.2015.2468078235](https://doi.org/10.1109/TVCG.2015.2468078235).
- [vdMH08] Van der MAATEN, LAURENS and HINTON, GEOFFREY. “Visualizing data using t-SNE”. *Journal of Machine Learning Research* 9.Nov (2008), 2579–2605 2.
- [WSA*16] WANG, YONG, SHEN, QIAOMU, ARCHAMBAULT, DANIEL, et al. “AmbiguityVis: Visualization of ambiguity in graph layouts”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), 359–368. DOI: [10.1109/TVCG.2015.24676913](https://doi.org/10.1109/TVCG.2015.24676913).
- [ZMG*10] ZHANG, ZHIYUAN, MITTAL, ARUNESH, GARG, SUPRIYA, et al. “A Visual Analytics Framework for Emergency Room Clinical Encounters”. *IEEE Workshop on Visual Analytics in Health Care* (2010) 3.