

GUCCI - Guided Cardiac Cohort Investigation of Blood Flow Data

Monique Meuschke, Uli Niemann, Benjamin Behrendt, Matthias Gutberlet, Bernhard Preim, Kai Lawonn

Abstract—We present the framework *GUCCI* (Guided Cardiac Cohort Investigation), which provides a guided visual analytics workflow to analyze cohort-based measured blood flow data in the aorta. In the past, many specialized techniques have been developed for the visual exploration of such data sets for a better understanding of the influence of morphological and hemodynamic conditions on cardiovascular diseases. However, there is a lack of dedicated techniques that allow visual comparison of multiple data sets and defined cohorts, which is essential to characterize pathologies.

GUCCI offers visual analytics techniques and novel visualization methods to guide the user through the comparison of predefined cohorts, such as healthy volunteers and patients with a pathologically altered aorta. The combination of overview and glyph-based depictions together with statistical cohort-specific information allows investigating differences and similarities of the time-dependent data. Our framework was evaluated in a qualitative user study with three radiologists specialized in cardiac imaging and two experts in medical blood flow visualization. They were able to discover cohort-specific characteristics, which supports the derivation of standard values as well as the assessment of pathology-related severity and the need for treatment.

Index Terms—Medical Visualization, Cohort Analysis, Measured Blood Flow Data, Cardiac Diseases.



1 INTRODUCTION

CARDIOVASCULAR diseases (CVDs) are correlated to abnormal morphological and hemodynamic features [1], [2]. Four-dimensional phase-contrast magnetic resonance imaging (4D PC-MRI) [3] enables non-invasive acquisition of time-dependent flow data in a volume. A more in-depth knowledge of these relationships based on 4D PC-MRI has great potential to improve the diagnosis of CVDs, to assess their severity more reliably and to support therapeutic decisions [4].

Several groups developed visual exploration techniques for single 4D PC-MRI data sets [5], [6], [7] involving different time-dependent flow features. Visual analytics support for analyzing cohorts instead of data from a single patient is the next step to support medical research in developing guidelines how to interpret 4D PC-MRI data. Comparative visual analysis of cohorts is needed to gain new insights into the relationships between changes in blood flow and vascular morphology. This would make it possible to identify correlations between the multitude of flow features and to reduce them to a meaningful subset. Furthermore, standard values for flow features could be determined.

Therefore, we introduce the framework *GUCCI* that enables a comparative visual analysis of multiple 4D PC-MRI data sets.

We will focus on aortic blood flow and related pathologies, which is one of the main applications for 4D PC-MRI measurements. Furthermore, we focus on cohorts that were predefined based on meta information such as gender, age, and pathology. We use machine learning techniques to automatically identify characteristic flow and wall features. The predefined cohorts and the selected features then serve as input for visual analysis. *GUCCI* provides novel visualization techniques, such as a 2D atlas of predefined cohorts based on the underlying flow data to assess their separability. Moreover, several diagrams and glyph-based techniques are provided to investigate differences and similarities of time-dependent flow features. Users are guided visually through the analysis of the data, with their attention drawn to the most important aspects through visual highlighting techniques and automatic determination of characteristics of interest. In summary, we make the following contributions:

- Multiple linked views to comparatively analyze predefined cohorts based on aortic 4D PC-MRI data, where the user's attention is drawn to interesting features.
- A novel 2D atlas of predefined cohorts, which serves as an overview depiction of cohort characteristics.
- Novel glyph-based techniques to compare multiple time-dependent features for cohorts on measuring planes.

2 RELATED WORK

Related work comprises the visualization of blood flow data, where we focus on quantitative analysis of 4D PC-MRI data (Sect. 2.1). Besides, there are approaches that used information visualization techniques to visually explore flow data, which are related to our work (Sect. 2.2). In addition, approaches dealing with the analysis of cohort data using visual analytics techniques are essential for our approach (Sect. 2.3).

- *M. Meuschke is with the Department of Simulation and Graphics, University of Magdeburg, and with the Department of Theoretical Computer Science, University of Jena, Germany. E-mail: meuschke@isg.cs.uni-magdeburg.de.*
- *Uli Niemann is with the University of Magdeburg, Germany. E-mail: uli.niemann@ovgu.de.*
- *B. Behrendt is with the University of Magdeburg, Germany. E-mail: behrendt@isg.cs.uni-magdeburg.de.*
- *M. Gutberlet is with the University of Leipzig – Heart Centre, Germany. E-mail: matthias.gutberlet@helios-gesundheit.de.*
- *B. Preim is with the Department of Simulation and Graphics, University of Magdeburg, Germany. E-mail: bernhard@isg.cs.uni-magdeburg.de.*
- *K. Lawonn is with the Department of Theoretical Computer Science, University of Jena, Germany. E-mail: kai.lawonn@uni-jena.de.*

Manuscript received September 16, 2020; revised September 16, 2020.

2.1 Quantitative Analysis of 4D PC-MRI Data

4D PC-MRI offers the reproducible quantification of clinically relevant flow features. Three methods are used: *cross-sectional*, *surface-based*, and *grid-based methods* [8]. For *cross-sectional methods*, measuring planes are placed in the vessel, aligned perpendicular to its centerline. Time-dependent flow curves describe the flow in the plane over the heartbeat. Features, such as heart volume and the percentaged back flow, which occurs when a valve does not close properly, depend on the size and inclination of the plane. *Surface-based methods* calculate measures on the vessel surface. A relevant factor is the wall shear stress, which can be invoked by near-wall vortex flow [9]. This category also includes derived values such as the vessel diameter along the centerline. *Grid-based methods* operate on the image data grid, e.g., to compare different PC-MRI sequences [10]. Relative pressure drops allow to assess the severity of vessel narrowings [11].

2.2 Flow Data Analysis using Information Visualization

Our work is related to previous attempts to the feature-based analysis of flow simulation results with information visualization techniques. An early example is the WEAVE system [12], where cardiac simulation data is mapped to the heart anatomy and to scatterplots and histograms. Thus, brushing allows to specify features in the feature space and linking highlights corresponding locations on the heart surface. A prominent example for this line of research is the SimVis system [13], [14] that provides advanced scatterplots, scatterplot matrices and parallel coordinate views to extract features of time-dependent vehicle simulations.

Another effective way to visualize several scalar values are glyphs. Ropinski et al. [15] introduced a taxonomy for glyph-based medical visualizations. They suggested to use simple and symmetrical glyph shapes, if possible 2D shapes, projected orthogonally to avoid size distortions. Therefore, we also use 2D symmetric glyphs, such as triangles, pentagons, heptagons, hexagons, and circles to provide an overview of cardiac features. Rocha et al. [16] applied sets of images, called *decal-maps*, to vascular surface models for visualizing different features. However, these glyphs are not appropriate to show temporal data. Van Pelt et al. [17] used glyphs to show the influence of various stent configurations for vessel dilatation on scalar features over time. The main drawbacks are the limited number of time steps displayed and the lack of presenting quantitative information. Meuschke et al. [18] used violin glyphs to show the temporal behavior of two time-dependent scalar features. Later, they extended their approach to visualize more than two time-dependent scalar features on surfaces simultaneously by using a combination of 2D and 3D bar charts, called *Skyline visualization* [19]. They showed that the violin glyphs are well suited to emphasize the temporal behavior of one scalar field over time. Therefore, we also use violin glyphs to show the temporal dynamics of one feature.

2.3 Visual Analysis of Cohort Data

An important application of visual analytics methods is the analysis of cohort data, where patients or healthy volunteers are investigated in several waves to assess their health status over a long period. Cohort studies allow to assess changes in the prevalence of diseases and health risks, such as obesity, as well as correlations between life style, laboratory values and diseases [20].

Cohort construction refers to analytic and interactive techniques for defining interesting subsets [21]. This process involves




Fig. 1. Structures and primary arteries of the human heart (left). Measuring planes placed semi-automatically within the aorta (right).

straightforward specifications, e.g., a cohort of males or females in a certain age range with a certain health disorder. Cohort construction is also relevant for patients suffering from chronic diseases or diseases that require long-term follow-up, such as cancer [22]. Steenwijk et al. [23] presented multiple linked views to support the user in extracting features from imaging and non-imaging data of patients with rheumatic disease. Visualizations, such as parallel coordinates and scatterplots, are provided for cohort analysis.

One essential development concerns better *user guidance* during the evaluation of complex medical data [24]. Visual analytics systems are often very complex because they enable a flexible combination and parametrization of visualization and analysis methods. However, this complexity often overwhelms the analyst, preventing the potential of the systems from being fully exploited. Guided visual analytics [25], [26] comprises strategies to guide users using predefined workflows. However, it is important to consider how much flexibility should be maintained and how strong the user guidance should be.

Contribution Summary. *GUCCI* is the first system to provide cohort-based analysis of blood flow data. Based on meta-information, cohorts are first predefined and machine learning methods are used to determine blood flow features that best distinguish these cohorts. We use combinations of overview visualizations and detailed views to analyze the complex time-dependent data, where the user is guided to strong differences between cohorts. The visualizations are enriched with statistical quantities to support both qualitative and quantitative data analysis.

3 MEDICAL BACKGROUND

The vascular system, with the heart as the central pump, ensures the blood supply of the body. The aorta is the largest artery in this system. Their caliber dimensions of about 2-3 cm allow the acquisition of a measurable signal for a meaningful blood flow analysis. During systole, oxygenated blood passes from the left ventricle (LV) via the aortic valve (AV) into the ascending aorta and is then supplied to the body, see Fig. 1(a). From the right ventricle, deoxygenated blood is pumped through the pulmonary valve into the pulmonary artery. During diastole, the AV is closed to prevent blood from flowing back into the LV.

3.1 Selected Pathologies and Related Flow Behavior

CVDs lead to vessel deformation, which increases the probability of quantitative and qualitative flow changes [27]. In the following, we explain selected pathologies and related flow changes.

Aneurysms. Aneurysms are local, permanent enlargements of the vessel diameter due to a weakened vessel wall. Hope et al. [28] found significantly increased volume of vortex flow in aneurysms.

Stenosis. A stenosis is a pathological narrowing of the vessel. As a result, the blood flow is blocked, which increases the flow velocity and vortex volume [29].

Bicuspid Aortic Valve. The bicuspid aortic valve (BAV) is the most common congenital heart defect, where two of the three valve leaflets are fused together. Ebel et al. [30] statistically analyzed flow features between volunteers and BAV patients. They found differences regarding flow velocities and vortical volume, but a detailed analysis of time-dependent flow features remains challenging. Moreover, the cohorts were quite small and there was no computer support for feature selection, which made the analysis process time-consuming and error-prone.

3.2 Data Acquisition and Preprocessing

A 4D PC-MRI data set contains each three time-resolved phase and magnitude images that describe the flow direction and strength, respectively. All temporal positions together represent one full heart beat. A 3 T Siemens Magnetom Verio MR scanner was used with a maximum expected velocity (V_{ENC}) of 1.5 m/s per dimension. The spatio-temporal resolution is $1.77 \times 1.77 \times 3.5 \text{ mm}^3 \setminus 50 \text{ ms}$ with a 132×192 grid for each of the 19 to 35 slices and 18 to 33 time steps. For our visualizations, all data sets must have the same number of time steps. For data sets with less than 33 time steps, missing time steps were interpolated linearly.

The software *Bloodline* [6] is used to analyze individual 4D PC-MRI data sets. *Bloodline* works as follows. First, image artifacts are automatically corrected. Then, the thoracic aorta and pulmonary artery are segmented using the Graph Cut algorithm. From the segmentation masks, polygonal vessel surfaces are extracted. In addition, six measuring planes are placed semi-automatically at standardized positions [31], see Fig. 1(b). Based on the phase images, the flow within the vessels is automatically calculated. To evaluate the heart function, two types of global features commonly used in clinical studies are calculated automatically:

- 1) *Time-independent features* refer to the entire vessel and do not change during the heartbeat, e.g., the heart rate or maximum vessel diameter (197 features).
- 2) *Time-dependent features*, e.g., maximum velocities or vortex volumes. For each time step, the minimum, maximum, mean, and median value within the vessel segmentation is calculated (17 features).

For each measuring plane, two types of features are calculated:

- 1) *Time-independent features*, e.g., flow rate or cardiac output. For each plane, the minimum, maximum, mean, and median value is calculated (41 features).
- 2) *Time-dependent features*, e.g., flow jet angle. For each time step and plane, the minimum, maximum, mean, and median values are provided (7 features).

So far, 90 data sets of heart-healthy subjects and 22 BAV patients have been evaluated with *Bloodline*, which serve as input for *GUCCI*. Currently, *GUCCI* focuses on the analysis of aortic flow data, as we have only few patients with pathological changes in other vessels such as the pulmonary artery.

4 REQUIREMENT ANALYSIS

An important research focus is to establish medical guidelines based on the extracted flow and meta information. For this purpose, a comparison of data sets regarding the morphological and hemodynamic features is essential. However, *Bloodline* provides no possibility to compare data sets. Therefore, clinical experts currently perform statistical analysis with commercial software such as Excel or SPSS, to validate and adapt hypotheses in an iterative process [30], [32]. Due to the amount and diversity of the data, this requires a high cognitive effort to interpret their temporal behavior. The simultaneous visual analysis is limited to a few data sets, as there is a lack of efficient comparative visualizations.

To overcome these problems, we designed a workflow to comparatively analyze aortic flow data for cohorts in collaboration with a radiologist specialized in cardiac imaging with 30 years of work (co-author of the paper). The goal is to enable medical experts to check existing hypotheses and to form new hypotheses through interactive data analysis. The intended workflow is that the analyst selects predefined cohorts and compares them to detect significant differences and similarities between them. This information can then be used to form new hypothesis and plan more specialized clinical studies to verify them.

Together with our medical expert, we discussed how *GUCCI* could be used for the analysis of measured aortic flow data to support clinical research. Physicians are very interested in analyzing differences and similarities between cohorts, where the target variable is defined based on meta information such as age, gender, and pathology. Therefore, our targeted scenario involves a cross-comparison of predefined cohorts. This discussion results in three predefined cohort settings (CS):

- CS 1: Healthy volunteers (HV) versus BAV patients,
- CS 2: Older HV versus BAV patients,
- CS 3: Female HV versus male HV.

Based on this scenario and a literature search, we summarize the main requirements for *GUCCI*:

R. 1 Feature Selection. Due to the large number of features, methods are required that reduce the number of features automatically to a meaningful subset w.r.t. the target variable.

R. 2 Cohort Adaption. The analyst should be able to add features in which (s)he is interested to the automatically selected features of a predefined cohort.

R. 3 Guidance. Methods are needed that guide the user's attention to the most interesting aspects of the data, such as large differences between cohorts or outliers.

R. 4 Perceive Separability. The analyst should get an overview of whether the selected features allow a clear separation of the data sets w.r.t. the target variable.

R. 5 Select Subgroups. The analyst should be able to easily select subgroups with similar feature values within cohorts.

R. 6 Analyze Differences. The expert should be able to analyze differences and similarities of features within the predefined cohorts to further characterize subgroups, such as "healthy aging" and "normal aging" people in a cohort of healthy people.

R. 7 Handle Data Diversity. Quantitative and qualitative visualizations are needed to analyze the features, which also convey their temporal behavior. This includes statistical information and relationships about selected features.

R. 8 Export Findings. The framework should provide an option to export findings in a standardized format.

5 THE GUCCI SYSTEM

The design of *GUCCI* is guided by the defined requirements to enable a comparative analysis of predefined cohorts regarding aortic flow data. We reduce the large feature space by employing feature selection techniques w.r.t. predefined cohorts, see Sect. 5.1. Within *GUCCI*, the expert selects one of the three cohort settings for visual analysis of blood flow data. The user receives a histogram-based overview of the features that were automatically selected for the respective cohort, see Sect. 5.2. At this point, the user has the option to add features and thus expand the cohort definition. Then, the visual cohort analysis is carried out, described in Sect. 5.3 and Sect. 5.4. Starting from an overview depiction, the different data types can be explored in more detail. Four visualization techniques are provided to represent the four data types, see Fig. 2. These depictions are enriched with statistical information to get a better understanding of the data. During the visual analysis, the user's attention is drawn to large differences between cohorts, see Sect. 5.5. To support the exploration of subgroups, all views are connected by the "brushing & linking" concept (meet Req. 5). The user can brush data sets in one view, which are then highlighted in all other views. At the end, detected findings can be stored for further documentation or communication with other domain experts, see Sect. 5.6.

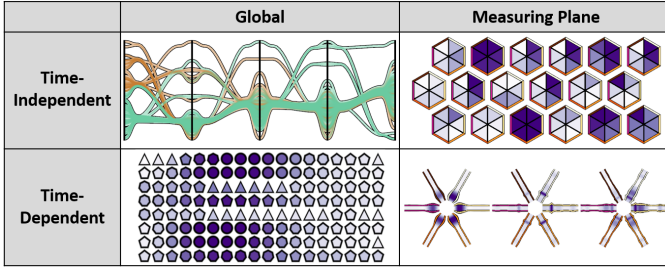


Fig. 2. Depiction of the four data types. Parallel coordinates and pixel plots show the global features, see Sect. 5.3. The plane-based data are depicted by hexagon and violin plots, see Sect. 5.4.

5.1 Feature Selection

First, we reduce the large feature space to focus on the most important ones. For the reduction, the 197 global time-independent features are considered. These represent the largest feature group in the data and contain many correlating features.

We use supervised classification of cohorts to find characteristic flow features. For this purpose, different feature selection methods and classifiers are available. To find a suitable set of features, we recently proposed a data analysis pipeline to classify the three predefined cohorts, encompassing feature selection, model training and hyperparameter tuning [33]. For this purpose, we combined several feature selection methods and classifiers to train separate models for the healthy subgroups and BAV patients. In the following, we briefly describe for all three cohort settings the automatically detected features that we discussed with our medical expert regarding their plausibility. More details regarding the methodical and evaluation setup can be found in our previous work [33].

CS 1. For the first classification, HV versus BAV patients, sequential forward search (SFS) combined with random forest (RF) classification [34] led to the highest accuracy. Three out of the 197 global, time-independent features were selected:

- **Time-to-Peak Vorticity:** represents the time during the cardiac cycle when the total volume of vortical blood flow in the aorta reaches its maximum.
- **Time-to-Peak In-Plane Velocity:** represents the time during the cardiac cycle when the maximum velocity of in-plane blood flow occurs in the aorta.
- **Peak Systolic In-Plane Mean Velocity:** quantifies the maximum of mean in-plane blood flow velocities across all levels during systole.

CS 2. For older HV versus BAV patients, SFS combined with Support Vector Machine classification yielded best performance. We determined the cutoff value as age with maximum information gain towards cardiac status [33]. Five features were selected:

- **Peak Systolic Mean Velocity:** quantifies the maximum of the mean blood flow velocities in the aorta during systole.
- **Time-to-Peak-Systolic-Through-Plane-Mean-Velocity:** indicates the systolic time at which the maximum of the mean velocity in the through-plane occurs over all planes.
- **Time-to-Peak-Diastolic-In-Plane-Mean-Velocity:** describes the diastolic time at which the maximum of the mean in-plane velocity occurs over all planes.
- **Diastolic-Median-Right-Rotation-Volume-Rel:** represents the median right-turn flow volume occurring over all diastolic time steps relative to total vessel volume.
- **Peak-Mean-Vorticity-Pressure:** is the maximum of the mean pressure within a vortex region over all time steps.

CS 3. SFS combined with RF also achieves best generalization performance to classify healthy men versus healthy women. Again, three features were selected:

- **Peak-Velocity:** represents the maximum flow velocity in the aorta during a cardiac cycle.
- **Peak-Systolic-VelocityQ99:** represents the maximum flow velocity (taken from the 99.5 % quantile range to discard outliers due to noise) occurring in the aorta during systole.
- **Time-to-Peak-Diastolic-Through-Plane-Velocity:** is the diastolic time at which maximum through-plane velocity occurs at one of the planes in the aorta.

Thus, we are able to automatically reduce the number of features to a meaningful subset w.r.t. the target variable (meet Req. 1).

5.2 Cohort Adaption

After the user loads a predefined cohort, we provide a histogram-based overview of the corresponding features. Fig. 3 illustrates the feature overview for the BAV cohort. In consultation with our clinical expert, the user should be able to add features of interest to the automatically selected feature set. Therefore, we allow the user to add individual features from the global time-independent features (meet Req. 2). For each feature, up to three ranges can be defined that a cohort must fulfill. Multiple feature specifications are evaluated as conjunction and the specification of different value ranges within a feature as disjunction.

Above the histogram, the number of records that meet the predefined criteria is displayed. The first number indicates how many data sets only fulfill the current value range and the second number in parentheses shows how many records match all given ranges. A summary shows the selected features and value ranges within a cohort. Here, the analyst can also remove value ranges or entire features. When the cohorts have been adapted, the analyst can switch to the comparative cohort visualization view. Moreover, the analyst can save and reload adapted cohorts.




Fig. 3. Histogram-based overview of pre-selected cohort features. The user can add new features, where up to three value ranges can be defined.

5.3 Cohort Visualization of Global Features

The simultaneous display of global and plane-based features would result in an overloaded visualization. Therefore, the visual exploration consists of two main steps: first, the visual analysis of the global features and second, the analysis of the plane-based data. The associated visualizations are arranged in two views.

The user interface to analyze the global time-independent and time-dependent features is divided into four subviews based on Shneiderman’s mantra [35] “Overview first, zoom and filter, then details-on-demand”. The first view gives a qualitative overview of the time-independent features w.r.t. the defined cohorts. Parallel coordinates together with statistical plots allow a more detailed exploration. In addition, a combination of an overview visualization and line charts is provided for a detailed analysis of the time-dependent features w.r.t. the predefined cohorts. In the following subsections, all subviews are explained in more detail.

5.3.1 Qualitative Cohort Overview

Due to the complexity of the data, first a qualitative overview is given to assess the separability of cohorts (meet Req. 4). Commonly used overview techniques for multivariate data are scatterplot matrices, table lens or glyph-based approaches such as star plots [36]. While these approaches do not scale well with the number of features, they are also less suited to assess the separability of cohorts.

Another frequent procedure for visualizing multivariate data is to apply dimension reduction techniques, where the data is projected to a lower dimensional feature space [36]. Standard techniques such as scatterplots can be used to depict the projected data sets, which are represented by circles colored to the target variable. Besides, the background of the scatterplot can be used to encode the distance between objects regarding the third main component [37]. Thus, the separability of cohorts can be assessed.

Based on this concept, we created a 2D cohort atlas, where the preselected global time-independent features serve as input, see Fig. 4. In the first step, a dimension reduction technique is applied to these features. For this purpose, we first tried a principal component analysis (PCA). However, this did not lead to a visual separation of the cohorts. Therefore, we also applied two frequently used nonlinear dimension reduction techniques, *UMAP* [38] and *t-SNE* [39]. We investigated the separability of cohorts in the resulting atlas visualization, where the most important hyperparameters of each method were tuned. These experiments showed that *UMAP* is better suited to visually separate cohorts, where we set the number of nearest neighbors to five. For all

other parameters, we kept the default values. For more details, we attached the resulting scatter plots of the *UMAP* and *t-SNE* experiments as supplemental material.

Subsequently, all data sets were projected into the space spanned by the vectors of the *UMAP* projection. Each data set is represented by a colored circle, where color reflects the cohort. Here, orange indicates BAV patients, while green represents healthy volunteers. The background color represents the density distribution of the data sets. In dark gray areas, many data sets lie close together, while a light background depicts sparser regions.




Fig. 4. Qualitative cohort overview. The data sets are projected to the 2D space, where each record is represented by a colored circle. The background color shows their density. The darker, the more similar the data sets are. Moreover, the user can brush specific data sets (red regions), which are highlighted in the other views.

5.3.2 Visual Analysis of Time-Independent Features

Starting from the 2D atlas overview, the user should be able to study qualitative and quantitative differences and similarities of time-independent features between the cohorts to further characterize subgroups (meet Req. 6). Thus, a quantitative and qualitative comparison between cohorts should be possible, where the visualization makes efficient use of space. There are various methods for the comparative visualization of multivariate data [36], e.g., a scatterplot matrix as well as multiple or stacked bar charts. However, these techniques need a lot of space, where correlations between features cannot be perceived at a glance. In addition, multiple bar charts can handle different ranges of values poorly, which would require logarithmic scaling, whereas components of stacked bar charts are difficult to compare.

Parallel Coordinates. In contrast, parallel coordinates (PC) make efficient use of space and enable both qualitative and quantitative data analysis [40], which are therefore used for the visual

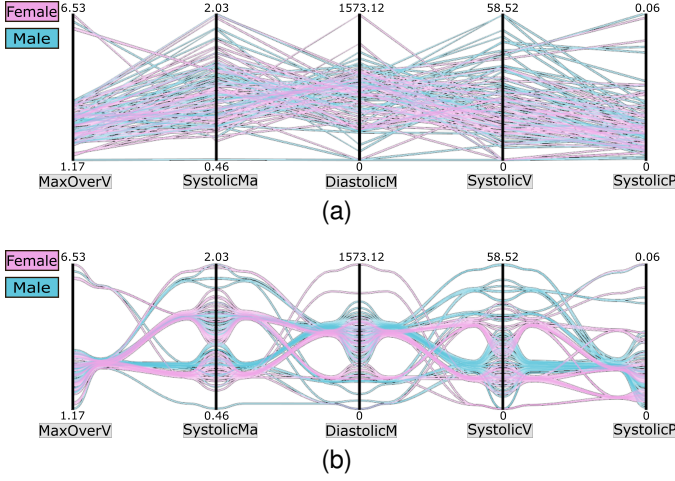


Fig. 5. PC without (a) and with line bundling (b). The grouping supports the perception of similarities and differences between cohorts. The rendering order can be changed using the colored panels on the left.

exploration of the time-independent features. Each feature from the preselected feature subset represents a vertical axis in the PC. Since the value ranges of the features differ, they must be scaled to the interval $[0,1]$ for display.

Since the order of axes affects the revealed patterns, we applied the technique by Hurley et al. [41] to find an appropriate ordering. The axes and their pair-wise neighborhood are treated as a graph, where nodes represent the axes and the edges represent two possible neighboring axes. This results in a complete graph. The edges are weighted using the Pearson correlation between associated features. To position strongly correlated features adjacently to each other, we chose the path with maximum accumulated weight as default axes order. If the user wants to investigate other feature pairs, the arrangement of the axes can be swapped.

In the PC plot, the lines are colored according to their cohort membership, which can lead to visual clutter, see Fig. 5(a). Therefore, we apply the line bundling technique by Palmas et al. [42], which leads to a grouped representation based on a density-based clustering of the polylines, see Fig. 5(b). We use the default parameter settings recommended by the authors [42], which lead to convincing groupings for our data. Moreover, a black contour is added to each bundle to increase visual separation.

Since the line bundling cannot completely bypass occlusions, we applied transparency to the bundles. This communicates the frequency of the data sets with certain feature values by overlaying several transparent bundles. However, in case of multiple cohorts, the interpretation of the blended colors can be difficult. Therefore, we provide an interactive changing of the order of the rendered polylines depending on the cohort membership. The analyst can change the order of the colored panels on the left margin to decide which cohort is visible on the top layer and therefore not hidden.

Statistical Diagrams. Since the PC give less insights about the distribution of two features, we provide statistical diagrams to further support the quantitative analysis (meet Req. 7), see Fig. 6. Here, the analyst can select two features from the preselected global time-independent features using the drop-down menus. Both features are visualized using a scatterplot to obtain more detailed information about their dependency structure. For this purpose, besides *linear* relationship between features, more




Fig. 6. Quantitative cohort visualization. Two user-selected time-independent features are visualized using a scatterplot combined with a marginal histogram. Above, statistical values are listed for each cohort.

complex bivariate relationships (quadratic and logarithmic) are considered. We determine which of the three relationships fits the underlying data best:

- simple linear regression: $Y = \beta_0 + \beta_1 \cdot X$
- quadratic regression: $Y = \beta_0 + \beta_1 \cdot X + \beta_2 \cdot X^2$
- logarithmic regression: $Y = \beta_0 + \beta_1 \cdot X + \beta_2 \cdot \log X$

The parameters β_i are estimated using the least-squares optimization. The scatterplot is enhanced by a colored curve visualizing the best fitting type of regression to describe the bivariate relationship between the X and Y features. To evaluate the quality of a fitted regression line, we use the *Bayesian information criterion* (BIC) value, a statistical measure, which describes the estimation of a function of the posterior probability that a model is true under a given Bayesian setup. The lower the BIC value the better the fit. Hence, the analyst can explore U-shaped, J-shaped [43] or other relations between features to analyze previously unknown associations and to generate new hypotheses.

Furthermore, marginal histograms display the distributions of the features as univariate histograms on the horizontal and vertical axes of the scatterplot, respectively. They support the comparison of relative frequencies between the cohorts. Thus, the height of each bar represents the relative amount of data sets w.r.t. the size of the corresponding cohort. When the user moves the mouse over a bar, the percentage of the total size of the associated cohort is displayed. In addition, the median, mean and standard deviation (SD) for each cohort is listed in relation to the features.

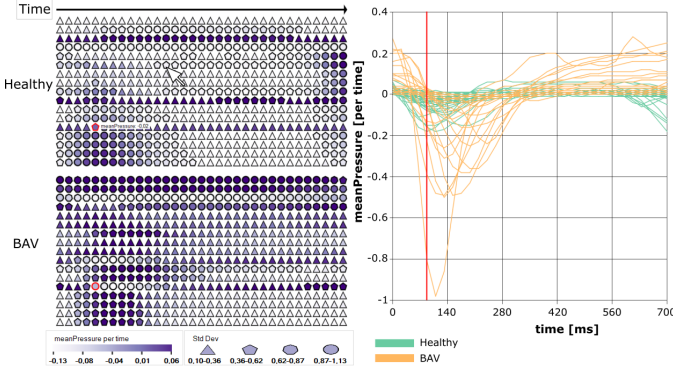


Fig. 7. Visualization of global, time-dependent features using pixel plots combined with a line plot. Each row of a pixel plot represents a feature, where the columns show the temporal behavior of the mean and SD within a cohort. The line plot shows the temporal course of a selected feature of all data sets considered in the cohorts.

5.3.3 Visual Analysis of Time-Dependent Features

To explore the 17 time-dependent features, we again apply Shneiderman’s mantra [35]. First, we give an overview of all features and their temporal progression. Then, the user can filter interesting features to analyze them in detail.

There are several methods to create a visual overview about multiple time-dependent features [44], e.g., conventional line plots. However, displaying all features simultaneously in line plots would be too space-intensive. Other alternatives would be *Horizon Graphs* [45] or *ThemeRiver* [46]. Horizon Graphs are based on the mean of the mean values. This is not useful in our case because flow features in systole usually have larger values than in diastole. Hence, interesting information would be averaged out. Similarly, ThemeRiver diagrams are inconvenient for comparing individual series [47] and thus of little use for our purposes.

Similar to Boyandin et. al [48], we adapt a heat map to give an overview about the global time-dependent features, called *pixel plot*, as shown in Fig. 7. Each row represents a time-dependent feature, where the columns represent the measured time steps. A pixel plot is drawn for each cohort. Within the cohort, the mean value of the respective feature is determined for each time step. These mean values are normalized across all cohorts based on the maximum of the respective feature and are mapped to a light to dark purple scale. In addition, the shape of the primitives (triangle, pentagon, heptagon, and circle) indicates discretized ranges of the SD to provide further statistical insights (meet Req. 7). Triangles represent a small SD, which increases with the number of primitive edges, where a circle shows a high SD.

To filter an interesting feature the user moves the mouse over the pixel plot, where the corresponding feature is highlighted in all pixel plots. The remaining features are grayed out. For a selected feature, a line plot is adjusted, which shows the temporal course of the data sets. The lines are colored in relation to their cohort, the red line indicates the currently selected time step. Here, green represents a healthy and orange a BAV cohort. Moreover, the feature name and its mean value for the cohort are shown. We also support the comparison of feature values at specific time steps. If the user selects a single primitive within a pixel plot, the corresponding primitives in the other cohorts get a red outline w.r.t. the corresponding time step. Thus, feature values can be easily compared between cohorts at certain time points (meet Req. 6).




Fig. 8. Visualization of time-independent features on the six measuring planes using hexagon plots combined with PC. The hexagons show the mean values of the features. The user can select individual hexagons for which the neighboring PC adapt themselves.

5.4 Cohort Visualization of Plane-Based Features

The analysis of plane-based data is also divided into the visualization of time-independent and time-dependent features. For both data types a combination of overview and detailed view is used.

5.4.1 Visual Analysis of Time-Independent Features

Within the aorta, six measuring planes are defined, with each plane being represented by an individual color. For each plane, 41 time-independent flow features are defined. Our goal was to find a space-saving overview of these features at each plane.

Typical methods for displaying multivariate data are multiple bar charts, stacked bar charts or pie charts [49]. However, these diagrams cannot be arranged in a particularly space-saving manner for the 41 features to be displayed [50]. Moreover, it is difficult and error-prone to compare the different sections of pie charts [49].

Another option would be to use glyphs. A commonly used glyph for multivariate data visualization are radar charts. However, they scale not well for more than five data entries and become quite space-intensive [49]. A more intuitive way would be to use aortic shaped glyphs where the planes are represented by lines, similar to Fig. 1(b), and color-code these lines according to the corresponding feature values. However, the perception of feature values would strongly depend on the line thickness and the major part of the aorta would remain unused. Another possibility would be to color the area between two planes. This would suggest that the feature value is defined for a larger aortic section, which is not true. In addition, such glyphs would be very wasteful of space.

Inspired by cartographic map visualization [50], we use hexagon plots, where each feature is represented by a hexagon, see Fig. 8. The hexagons allow a very space-saving arrangement, where the space for a cohort can be tessellated without any white area. Thus, 41 hexagons are drawn per cohort. Each triangle within a hexagon represents one of the six planes, with its colored border matching the plane colors. The plane name is displayed when the user moves the mouse over the border of a triangle. Within the triangles, the cohort mean value of each feature is color-coded. The hexagon plot is used to quickly locate features that are very similar or different between cohorts in terms of the measurement planes. To further assist in finding similarities and differences, individual hexagons can be selected, with corresponding hexagons in the other cohorts also highlighted (meet Req. 6).

For a more detailed investigation of the diversity of selected features, PC are shown to the right of the hexagon plots. The vertical axes correspond to the six planes. Again, each data set is represented by a polyline in the PC, colored according to the cohort membership. When the user moves the mouse over a hexagon, the PC for the selected feature adjust.

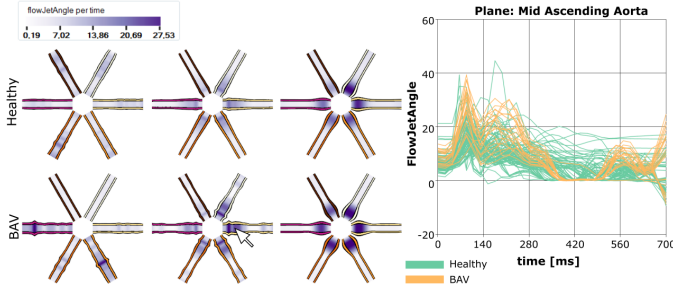


Fig. 9. Visualization of time-dependent plane features using violin plots. Each violin shows a feature, where the six arms represent the planes. The temporal behavior is encoded along the arms, where color and width indicate the mean value of the cohort. A line plot shows the temporal behavior of a selected feature at the corresponding plane.

5.4.2 Visualization of Time-Dependent Features

Seven time-dependent features are defined at each plane. Similar to the hexagon plots, different representations are possible to give an overview of their temporal behavior. A typical approach is to apply animation to a basic geometrical representation to show the temporal behavior of features [51]. However, the user would have to watch the animation several times to get a temporal overview, whereby there is an increased risk to overlook interesting aspects. Besides, more abstract visualizations such as *Silhouette Graphs* [52] or circle plot-based depictions [53], [54] could be used. However, similar to ThemeRiver, with Silhouette Graphs low feature values can be hard to perceive, whereby in plot-based depictions seven features already lead to very small, hardly perceptible subdivisions of the segments [18].

Similar to the hexagons, we wanted to use a representation in which each plane is represented by a partial geometry. The temporal progression should be represented statically to avoid perception problems with animations. Inspired by Meuschke et al. [18], we use *violin plots* to give an overview of the time-dependent features, see Fig. 9. Violin plots are a glyph-based technique, where each plot represents a cohort and each violin represents a feature. However, in contrast to Meuschke et al. [18], we use only 2D violins instead of a 3D representation because we do not have a 3D surface as an underlying domain.

Although the violin plots are not as space-saving as the hexagons, their design is based on the hexagons, which allows for an easier mental transition from the hexagons to the violin plots to grasp the added time component. Similar to the hexagons, they consist of six parts (arms), which together represent the planes. Thus, each arm corresponds to one plane, with the color of the arm's border corresponding to the plane colors. The measured time steps are plotted along an arm from inside to outside. Dark purple wide areas represent high mean values and thin light areas show time steps with low mean values. To compare features between cohorts, individual violins can be selected, while corresponding violins in other cohorts are highlighted (meet Req. 6).

To examine the temporal course of a feature at the planes more detailed, a line diagram adapts itself besides the plots. This shows the temporal behavior of a selected feature at a specific plane. When the user moves the mouse over the arms of the violin plots, the line diagram adapts depending on which feature and which plane corresponds to the current mouse position.

5.5 Guided Exploration of Cardiac Data

Purely manual data exploration can be very tedious. Our medical expert explained that large differences regarding features are very interesting as well as data sets that can be classified as outliers. Therefore, we automatically detect and highlight dissimilarities between cohorts as well as outliers. This draws the user's attention to interesting characteristics (meet Req. 3).

5.5.1 Automatic Determination of Cohort Dissimilarity

Due to the complexity of the pixel plot as well as the representations for the plane-based features, the user is directed to strong dissimilarities between cohorts. Dissimilarities between cohorts can be viewed at two levels. First, on the basis of the glyphs (mean values and SD of the features) and second, on the basis of the line plots for a selected feature.

Dissimilarity of Glyph Representations. We compute a similarity measure to sort the features in the pixel plot, starting with the feature that differs most between the cohorts. For this purpose, we consider the color-coded mean and SD values of a feature over time as curves. There are numerous algorithms to determine similarities between curves [55]. A challenge in any case are the strongly varying value ranges and units of the time-dependent features. The goal was to choose a measure that medical experts could easily understand and that represents the perceptual differences in the pixel plot. Strong deviations in shape and color of the primitives suggest large differences. This should be reflected in the similarity measure.

Therefore, we use the average differences of the mean and SD values of a feature over time between two cohorts as measure. For a given (averaged) feature $f_i^c(t)$ of the cohort c at a point in time t , we determine the interpolation factor $\lambda_i^c(t)$ by $\lambda_i^c(t) = \frac{f_i(t) - \min_t f_i(t)}{\max_t f_i(t) - \min_t f_i(t)} \in [0, 1]$. Afterwards, we calculate the average of the absolute difference of both cohorts: $\Lambda_i = \frac{1}{|T|} \sum_{t \in T} |\lambda_i^1 - \lambda_i^2|$. Note that T is the set of discretized points in time, where $|T|$ denotes the cardinality of T , i.e., the number of elements in T . In case of more than two cohorts, we change the equation to: $\Lambda_i = \frac{1}{|T| \cdot \binom{|C|}{2}} \sum_{t \in T, j, k \in C, j > k} |\lambda_i^j - \lambda_i^k|$, where $\binom{|C|}{2}$ is the binomial coefficient indicating the number of pairwise combinations of cohorts. We apply this to the SD values, too and obtain Γ_i . Finally, for every feature i , we average both values $\frac{1}{2}(\Lambda_i + \Gamma_i)$ and obtain a similarity value, which can be used to sort the features.

In the same way, the similarities for the hexagon and the violin plots are calculated to sort the features. For each plane, the similarity is computed. Finally, the resulting six similarity values are averaged to get an overall similarity value per feature.

Dissimilarity of Line Plots. While the sorted glyph-based plots (pixel, hexagon and violin plots) provide an overview of large differences between cohorts, the corresponding line plots highlight temporal regions where strong differences w.r.t. a selected feature occur between cohorts. This gives the user a detailed insight in which sections of the cardiac cycle cohorts differ strongly.

To highlight temporal intervals with strong differences in feature values, we again consider the mean values over time as a curve for each cohort, see Fig. 10. Next, an absolute difference curve of the mean curves is determined (blue curve). In the case that two cohorts are compared, we calculate the absolute difference of the associated mean values for each time step, resulting in a difference curve. In the case where more than two cohorts are considered, we average the pairwise absolute difference of all




Fig. 10. Computation of temporal intervals, where large differences between the presented cohorts arise for the selected feature.

cohorts. Afterwards, we divide the difference curve into regions by determining its local maxima. The maxima are used to define the boundaries of the regions. Between two local maxima there has to be a local minimum. Accordingly, a region is generated for each local minimum, which is bounded by local maxima. An average difference value is then calculated for each region. The region-based difference values are mapped to a magenta color scale. Thus, the user can immediately see in which temporal areas strong differences (low similarity) between the cohorts occur.

5.5.2 Automatic Determination of Outliers

We use scatterplots and PC to visualize the time-independent features as well as line plots to depict the time-dependent features. Within these depictions, we automatically detect outliers that can be highlighted on demand. By detecting outliers, medical experts can explore whether, e.g., HV have similar feature values to BAV patients. In the following paragraphs, we describe the outlier detection within the corresponding visualization.

Scatterplot. To automatically determine outliers in scatterplots, an outlier definition must first be established. Either univariate outliers (data points that take on extreme values in one variable) or bivariate outliers (data points that are extreme in both represented variables) are determined. In consultation with our clinical expert, we detect univariate outliers using the boxplot outlier definition. Accordingly, values above $Q3 + 1.5 \times IQR$ or below $Q1 - 1.5 \times IQR$ are considered as outliers with $Q3$ and $Q1$ being the third and first quartile that span the interquartile range IQR .

Parallel Coordinates. To detect outliers in PC, we use the method by Novotny and Hauser [56]. First, the data points between two adjacent axes in the PC are transformed to a scatterplot, where the axes in the PC represent the x and y axis in the scatterplot. Then, a regular grid is superimposed on the scatterplot and the number of data points (population) per grid cell is determined. This yields a binned representation of data points, where we use the suggested resolution of 256×256 bins. On each computed bin-map a median filter is applied to detect outliers. For this purpose,

a population threshold has to be set between 1 and 10 % of the maximum population. We follow the suggestions of the authors and use a population threshold of 7 % for central bins as well as 2 % for the corners and 4 % for the borders. Finally, all detected outliers in all bin-maps are united into the final outlier set.

Line Plots. Similar to the computation of the similarity measure for ordering the features in, e.g., the pixel plot, multiple metrics could be used to detect outliers in line plots. Again, we use a simple measure that is easy to understand for clinical experts. Accordingly, for each curve in the line plot, we determine the absolute mean distance to the already determined average curve. Then, we again apply the boxplot-based metric to the calculated distance values. According to this, outliers show strong deviations from the mean value over time. Of course, more advanced methods such as mean-of-closest-point distance [57] could be applied at this point. However, a detailed comparison of different measures and their results would exceed the scope of this paper.

5.6 Documentation of Findings

Saving results of the visual analysis is important for knowledge discovery and collaboration between experts. To export findings, GUCCI takes screenshots of all subviews, where another wizard opens to guide the user through the process of exporting the results (meet Req. 8). The user can add a specific title and a textual description to each image. In addition, user-specific markers can be added to the images using simple drawing interactions to highlight individual structures. Finally, the results can be saved. For this purpose, a `LaTeXscript` is called which generates a PDF document from the images together with the text input. In addition, the PDF document contains meta information, such as the number of records and features of the cohorts. Thus, the user gets a compact summary of the exploration results.

5.7 Implementation

GUCCI is a *Windows Forms application*, which is a toolkit of the *Microsoft .NET Framework* to design *GUIs* written in `C#`. For rendering, the *open toolkit library* (OpenTK) is used that provides access to graphics tools contained in the *open graphics library* (OpenGL). For feature selection and computation of the correlation coefficients, we use *R-Scripts* integrated into the `C#` setup with the help of *R.Net assemblies*. GUCCI requires a graphics card that supports OpenGL 4.3.

2D Atlas. The 2D atlas is generated in a two-step rendering process. First, the projected data sets resulting from the UMAP are drawn into a texture in the form of colored circles using the OpenGL extension `SHADER_IMAGE_LOAD_STORE`. The size of the image comprises 512×512 pixels. In the second run, the density values for determining the background color are computed. For this purpose, a screen-filling quad is drawn using the texture generated in the first pass. Moreover, a Gaussian distribution is created for each data set, where the SD is set to the width of the used texture. Then, we determine for each data set the records in its vicinity and the corresponding Gaussian distributions are accumulated, where the size of the neighborhood is set to 100×100 pixels. This results in a scalar field which is mapped to the light to dark gray color scale.

PC and Statistical Diagrams. The applied line bundling technique [42] on the PC results in a set of Bézier splines. These are rendered on the GPU, where each spline is assigned a transparency value of 0.5. For blending, the function `glBlendFunc(SRC`

ALPHA, *ONE MINUS SRC ALPHA*) is used to reduce occlusion problems between bundles of different cohorts. For rendering the scatterplot, a Windows Forms diagram is used, where each record is rendered as a colored circle. Since the diagram belongs to the graphic objects of C#, its paint event is used to render the marginal histogram. Here, colored rectangles are drawn by using built-in functionalities to draw standard geometries.

Pixel Plot. Initially, for each feature and time step within a cohort, a 2D position is uploaded to the GPU using *Vertex Buffer Objects*. These positions are extended to squares by using *geometry shaders*, where each of the four corner points receives texture coordinates. This information is used in the *fragment shader* to transform the squares into one of the primitives used to represent the SD. For the line plot, again a Windows Forms diagram is used, where each line is defined by a set of connecting points and the line color represents the corresponding cohort.

Hexagons. The hexagon plots are also generated on the GPU by using *geometry shader*. First, the hexagon mid points are uploaded to the GPU using *Vertex Buffer Objects*. In addition, the color-coded mean values are transferred to the GPU using *Shader Storage Buffer Objects* that allow reading operations at arbitrary buffer positions. This is necessary to read the six mean values of a hexagon. For each uploaded mid point, six triangles are generated in the *geometry shader*, where the final color-coding happens in the *fragment shader*. Here, a discretized white-to-purple color-scale is defined, which maps each incoming mean value to color.

Violins. For rendering the violins, we use the OpenGL *Instanced Rendering (IR)*, because this effectively generates the glyphs by reducing the amount of OpenGL draw calls. Using *IR* allows to allocate memory for one violin geometry that is reused to generate the following glyphs. The geometry of one violin is generated using Blender, where each arm is represented by a triangle strip. This geometry is loaded in GUCCI and reused for *IR*. The color-coded mean values over time are transferred to the GPU by using *Shader Storage Buffer Objects* to read the mean value for each vertex position of each violin arm.

6 EVALUATION

In the evaluation sessions that lasted between 60 and 90 minutes five domain experts used *GUCCI* with remote assistance to analyze cardiac blood flow data. We video-recorded the sessions and analyzed them afterwards. At the beginning of each session, we gathered some background information on the experts who formed two expert groups (EG). EG₁ comprises three radiologists P₁ to P₃ with an expertise in heart imaging (3, 5, 23 years of experience in radiology). P₃ also supported us with defining the requirements, see Sect. 4. They have a background in statistics and are familiar with statistical plots such as histograms and scatterplots. Their experience with interactive visualization tools is limited. EG₂ consists of two experts for medical visualizations P₄, and P₅ (2 and 8 years of experience), which were quite familiar with interactive and advanced blood flow visualizations.

First, we introduced the basic workflow of *GUCCI* and explained the interaction possibilities. Then, a prevalence study was carried out with EG₁ for the predefined cohorts, see Section 6.1. After a 30-minute introduction the physicians were able to interact independently with *GUCCI*. Then, we asked all EGs to fill out a questionnaire to assess the suitability of our visualizations, see Sect. 6.2. For this purpose, the predefined cohorts are presented to EG₂, since they did not participate in the prevalence study.

6.1 Prevalence Study

Our main goal is to support a comparative analysis of measured blood flow data of predefined cohorts to investigate the range of normal values and to assess the severity of specific pathologies. To check whether *GUCCI* is suitable for carrying out this scenario, we conducted a prevalence study with EG₁, examining different cohorts w.r.t. striking differences and similarities. The physicians were asked to investigate the three predefined cohort settings (CS):

- CS 1: Healthy volunteers (HV) versus BAV patients,
- CS 2: Older HV (> 47 yr.) versus BAV patients,
- CS 3: Female HV versus male HV.

Next, we discuss the most interesting findings for each setting.

CS 1. The first comparison should show whether HV and BAV patients differ regarding the extracted features. In addition to the automatically selected features, the experts were interested in the maximum vortical flow volume and pressure values in both cohorts. The physicians detected clear differences regarding the global and plane-based features between both cohorts, see Figure 11. According to their expectations, the maximum flow velocity, vortical flow volume and pressure values in BAV patients are about two to three times higher than in HV (Fig. 11(a) (first, fourth and fifth axis)). Regarding the global time-dependent features, BAV patients have a lot of increased values, which shows the increased pumping capacity of the heart due to the valve defect (Fig. 11(b)). In addition, Fig. 11(c) shows that the volume of the backflow (red bordered hexagons) in BAV patients is up to six times larger at the first three planes (planes close to the AV) compared to the HV, which reduces the cardiac output (cyan bordered hexagons) in the BAV patients up to 30 %. For the time-dependent plane-based data (Fig. 11(d)), the BAV patients also show significantly higher flow rates during systole (cyan bordered violins). Moreover, the angle of the main flow (*flow jet*), which usually follows the vessel centerline, is significantly higher for BAV patients (red bordered violins), which increases the risk of additional pathologies such as aneurysms.

However, against their expectations, the HV show large variances in the time points, where the maximum inplane velocity and vortex volume occurs (Fig. 11(a) (second and third axis)). To further explore these variances within the HV cohort, the experts brushed HV data sets in the PC that show very large and low values regarding the maximum inplane velocity and vortex volume. The selected data sets are highlighted in dark green, which the experts then explore in the remaining views. Due to the similarity-based ordering of the features, the experts immediately find out that some of these HV data sets show very high values for the global time-dependent features, e.g., increased volume values of the right rotating flow (Fig. 11(b)). In contrast, however, these selected HV data sets do not show particularly elevated values in the plane-based features, which was very surprising for the experts (Fig. 11(c, d)). The selection of conspicuous data sets helps the physicians to better understand the variety of flow features of HV.

CS 2. Since HV have partly similar values as BAV patients, the physicians were interested whether these HV are of an older age. The goal was to automatically select predictive features that could discriminate between elderly HV and BAV patients. This question is motivated by the challenge in clinical practice to distinguish between a non-pathological decline in cardiac function with age (healthy aging) and the onset of early stages of BAV. Therefore, the experts brushed older HV (> 67 yr.) in the PC plot (Fig. 12(a)).

It can be seen that these older HV have a similar vortex volume as BAV patients (third axis). Nevertheless, BAV patients show significantly higher and lower values, e.g., for diastolic flow velocity (fourth axis). Similarly, BAV patients have higher values in global time-dependent data. Region-based similarity analysis was activated in the pixel plot. It can be seen that older HV show increased values, which partly correspond to those of BAV patients. However, especially during systole, many BAV data sets have higher values, e.g., for mean pressure in regions with vortical flow (Fig. 12(b)). Thus, BAV patients have more vortical flow due to the valve defect. This helps the physicians to define hypotheses regarding the severity of BAV patients. An advanced stage BAV patient often shows an aortic valve stenosis, which seems to lead to significant differences in flow values. BAV patients with a lower severity seem to show similar values as elderly HV.

CS 3. As expected by the physicians, the male and female HV show in general only minor differences regarding the global time-independent features (Fig. 13(a)). However, based on outlier detection in the pixel plot, the physicians were surprised that the proportion of left- and right-handed vortical flow is about 25 % higher in men at all time steps (Fig. 13(b)). Men also show a higher SD for these features. By highlighting the affected male HV, it can be seen that they also have higher values for the plane-based features, e.g., twice as high velocities of the flow jet (Fig. 13(c)). The physicians suspect that the increased proportion of vortical flow and higher velocities result from the larger heart and blood volume of men. However, it remains to be explored whether this leads to a higher risk of secondary diseases. Thus, a detailed analysis of HV is important to derive normal values and reliably detect abnormal characteristics.

6.2 Expert Feedback based on Questionnaire

After the prevalence study, the experts filled out a questionnaire about the visualizations and their medical relevance using a five-point Likert scale (—, —, —, +, ++). We also analyzed their oral feedback. In the following, we summarize the main results. The questionnaire including the answers is shown in Fig. 14.

First, we evaluated the *comparative cohort analysis* in general. All experts confirmed the importance to analyze large cohorts for collecting standard flow values and to better understand relations between morphological and hemodynamic changes. Then, we inquired the *adaption of cohorts*. Basically, the list-based presentation of features was rated as suitable for adapting cohorts. P3 and P4 would have liked a more compact overview of available features. Features could be grouped according to categories such as pressure-related, speed-related and flowjet-related to present the features more clearly and to facilitate their selection. The experts found the histogram-based representation of the features as well as their combination and linking to be well suited for adapting cohorts (fulfill Req. 2). However, they had difficulty selecting meaningful feature combinations or value ranges for flow and wall features, resulting in the rather mixed ratings. The reason for this is a lack of knowledge about which features and value ranges might be characteristic for a cohort. Here, it would help to first determine normal values for a healthy cohort, which are displayed in *GUCCI*, in order to define cohorts with abnormal feature expressions.

Afterwards, the *qualitative cohort summary* was evaluated. We demonstrated the concept of the 2D atlas to get an overview of similarities and differences of the cohorts, which was described as really useful. P3 and P4 noted that comments or visual hints

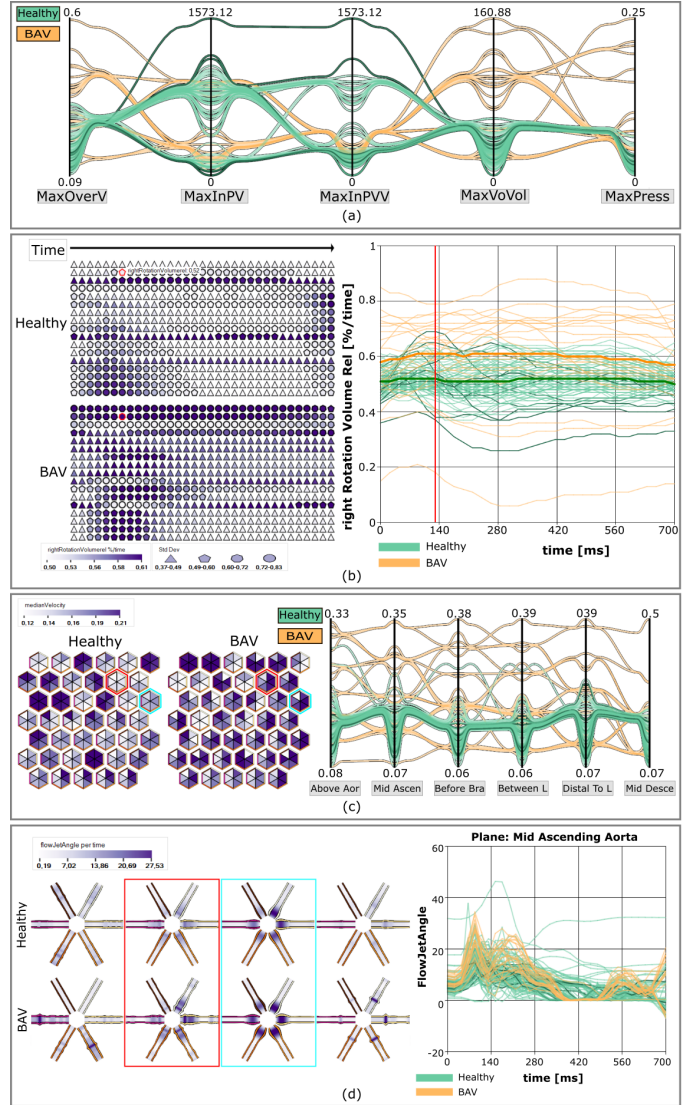


Fig. 11. Comparison of HV and BAV patients. HV and BAV patients show clear differences in the extracted features. However, HV show partly large variances in the features, which are highlighted in the views.

on how the similarities were calculated or what is coded on the axes would be helpful to better interpret the plot. In addition, the coloring of the atlas background supports the recognition of subgroups (fulfill Req. 4). However, P2-P4 would have liked a legend explaining the background color. Also the brushing of data sets was evaluated as an important feature to select subgroups (fulfill Req. 5). P3 would have liked different brushing colors to select different subgroups.

Next, the visual analysis of the global, time-independent features using *parallel coordinates* together with *statistical diagrams* was evaluated. Basically, the combination of scatterplot and histograms gives an overview of the distribution of features between the cohorts. Especially the coloring of the histogram based on cohort membership was emphasized as helpful, as it allows a quick interpretation of the distributions. However, P4 found the rotated histogram on the y-axis difficult to read. P3 and P5 would also have preferred a display of statistical graphs of multiple features simultaneously to better compare features. In addition, the experts assessed the specification of statistical values (median, mean

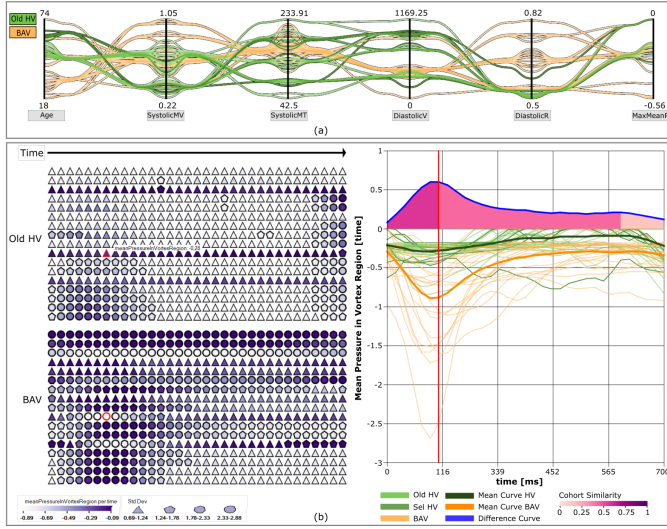


Fig. 12. Comparison of older HV and BAV patients. The experts brushed older HV which show partly similar values compared to BAV patients.

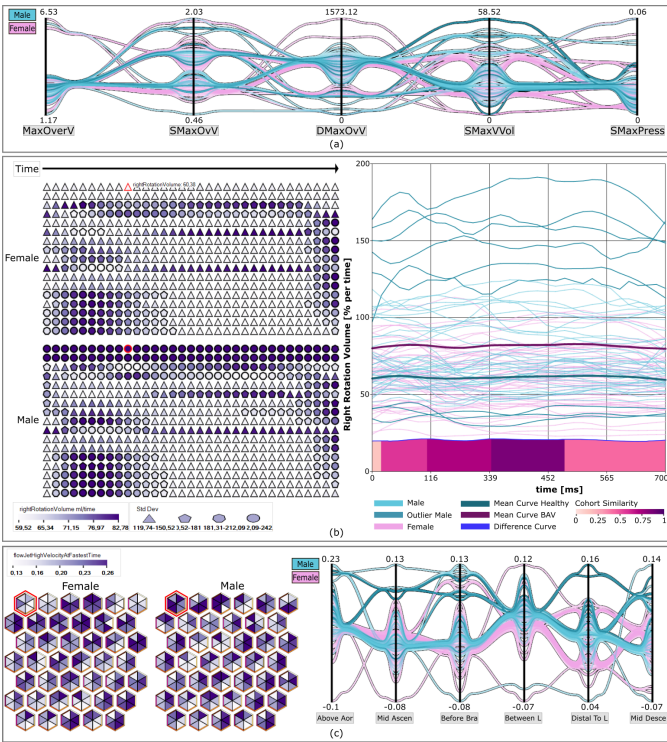


Fig. 13. Comparison of female and male HV. While both cohorts have in general similar values, detected outliers in the male cohort show increased values in the global time-dependent and plane-based features.

and SD) as helpful for finding characteristic differences between cohorts. P5 mentioned that automatic sorting of features based on statistical differences would further support finding anomalies.

Moreover, the experts confirmed that the regression curve is important to analyze more complex relations between two features (fulfill Req. 7). However, they would have liked more explanatory information on the concept of regression, e.g., in the form of info boxes. Users who are not familiar with more advanced statistical aspects would then have an opportunity to first learn more about this in order to interpret the representations.

PC make it possible to identify features in which the cohorts are similar or different, also for the plane-based features (fulfill Req. 6). However, similar to the statistical concepts, two medical experts here would have liked more information about the basic concept of PC. The usefulness of the polyline bundling was rated differently by the experts. While P1 and P2 confirmed that it supports the perception of subsets with similar feature values, P5 prefers the simple line-based ending since in his opinion too many details about individual data sets get lost with the bundling. Similarly, P3 and P4 explain that the bundling gives a good overview of possible subgroups, but for the detailed investigation of individual records the line-based rendering was preferred. Exchanging the rendering order reduces occlusion problems and supports the exploration of cohort-specific characteristics.

Next, we evaluated the suitability of the *pixel plot*. The experts confirmed that it gives a fast overview about the global time-dependent data. Some experts would have liked an explanatory example of strong and weak differences based on the representations to reduce the learning curve. Due to the guidance, features that require detailed analysis can be quickly identified within the pixel plot (fulfill Req. 3). However, the mapping of SD to shape was discussed controversial. Two physicians and one visualization expert had no problems quickly identifying features with high and low diversity. In contrast, P3 and P5 noted that the combination of color and shape to simultaneously encode mean and SD required a lot of concentration to identify areas of interest. In addition, P5 felt that coding the mean would be sufficient for data analysis. Besides, the interactive adaptation of the line plot for a detailed analysis of cohort-specific differences was positively highlighted. This also applies to the adjustment of line plots to display time-dependent plane data. Only P4 suggested to adjust the line plot only when the user clicks on a specific feature in the plot.

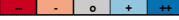
Then, we evaluated the techniques to *visually analyze the plane-based features*. Both the hexagon plot and the violin plot provide a quick overview of the features, where the interactive highlighting allows a quick comparison of cohorts (fulfill Req. 6). Moreover, the combination of color and width of the violin arms conveys the temporal behavior of a feature. Similar to the previous presentations, there were deductions in the rating due to missing examples or info boxes to better introduce the presentations.

Finally, we inquired the *export of findings*. Generating screenshots with individual descriptions and drawings assists documentation and collaborative data analysis (fulfill Req. 8). The experts would also like to see the integration of voice input to comment on the findings. Related to this, the saving and loading of adapted cohorts was rated as important for collaborative exploration.

7 DISCUSSION

The evaluation indicates that *GUCCI* supports a cohort-based exploration of aortic flow data. The physicians were able to perform a comparative analysis, which led to a better understanding of normal value ranges and revealed considerable differences in BAV patients, which supports the assessment of severity.

Manual Cohort Definition. Currently, we focus on the analysis of predefined cohorts for which a target variable can be formulated. Our medical cooperation partners are interested in whether other interesting cohorts are coded in the data whose target variables are defined based on flow parameters. In a preliminary study, to find such cohorts, we provided a correlation matrix of the global time-independent features color-coded by the Pearson correlation



	Clinicians			Vis Experts	
	P1	P2	P3	P4	P5
Comparative Cohort Analysis					
The analysis of large cohorts of patients and healthy persons is important for the collection of standard values regarding hemodynamic features.	++	++	++	++	++
Comparative analyses between different cohorts allow a better understanding of the relationships between morphological and hemodynamic changes.	++	++	++	++	++
Adaptation of Cohorts					
The list-based feature presentation is suitable to adapt cohorts.	++	++	o	+	++
The histograms help to narrow down the feature ranges.	++	+	o	+	++
The combination of features is suitable for adjusting cohorts.	++	+	o	o	+
The linking possibilities are sufficient to adjust cohorts.	++	+	o	+	++
Qualitative Cohort Overview					
The atlas shows cohort similarities and differences.	++	++	+	+	++
The atlas background color helps to recognize subgroups.	++	+	+	+	++
Brushing within the views and highlighting the selected objects supports data exploration.	++	++	+	++	++
Visual Analysis of Global Time-Independent Features					
The combination of scatterplot and histogram gives an overview of the distribution of the selected features between the cohorts.	++	++	+	o	+
The display of statistical parameters (median, mean and SD) is helpful to find characteristic differences between cohorts.	++	++	++	++	+
The additional display of the quadratic and logarithmic regression curve allows analyses of complex relationships between two features.	++	+	o	+	+
PC allow to identify similar or different features between cohorts.	+	++	+	++	++
The line bundling supports the perception of similar subsets.	++	++	+	+	o
Switching the order supports the exploration of cohort-specific characteristics.	++	++	+	++	+
Visual Analysis of Global Time-Dependent Features					
The pixel plot gives a quick overview of the time-dependent global features.	++	++	+	+	+
Features that require detailed analysis can be quickly identified with the pixel plot.	++	+	++	++	++
The mapping of the SD via the shape is helpful to recognize the diversity in the data.	+	+	o	+	o
The interactive adjustment of the line plot allows a detailed analysis of differences between cohorts w.r.t. the feature.	++	++	++	+	++
Visual Analysis of Plane-based Time-Independent Features					
The hexagon plot gives a quick overview of the time-independent, local features on the planes.	++	++	++	++	+
The interactive highlighting of hexagons between different cohorts allows a quick comparison of them.	++	+	+	+	+
PC allow to identify planes where features between cohorts are similar or different.	+	++	+	++	++
Visual Analysis of Plane-based Time-Dependent Features					
The violin plot gives a quick overview of the time-dependent, local features on the planes.	++	+	++	+	++
The arms of the violins convey the temporal feature behavior.	++	+	+	+	+
The interactive highlighting of individual violin plots between different cohorts allows for a quick comparison.	++	+	+	+	+
Interactive adjustment of the line plots allows detailed analysis of differences between cohorts w.r.t. the selected feature and the associated plane.	++	+	++	++	++
Export of Findings					
The generation of screenshots with individual descriptions is important for the documentation of the exploration results.	++	++	+	++	++
The generation of screenshots with individual descriptions is important for the collaborative analysis of the data.	++	++	+	++	++
Drawing additional information into the result images supports the documentation of the results.	++	++	+	++	++
Saving and loading adapted cohorts is important for collaborative exploration.	++	++	++	++	++

Fig. 14. The results of the informal expert feedback with color-coded Likert scores. Each box represents the answer of one expert.

coefficient. Users had the option to remove correlated features based on a threshold. The remaining features served as input to the cohort definition. New cohorts and associated features could be created using the histogram-based feature visualization. However, the experts had difficulty defining cohorts based on flow features. In the future, more advanced techniques, such as subspace clustering [58], will be needed to manually define interesting cohorts.

However, if the experts detect other interesting features and value ranges during the visual exploration, *GUCCI* allows an iterative cohort refinement. The experts can go back to the cohort adaption at any time and adjust information. In addition, for the relations between different features, more complex links would also be conceivable, which would allow disjunctions or negations. However, this would result in difficult to understand links, which, in consultation with our physicians, would not be appropriate.

Limited Number of Cohorts. Currently we only have access to a few data sets with pathologies other than BAV. Therefore, we discussed the basic limitation to two cohorts to be analyzed with our experts. This would allow the use of visualization techniques, such as differential color coding, which specifically highlights

differences between two cohorts. However, our medical experts were not convinced to limit *GUCCI* to the analysis of two cohorts. They expect that in the near future other data sets with other pathologies of cardiac vessels will be available. Our experts are interested in comparing other aortic pathologies with a healthy cohort and the BAV cohort to find out if there are pathology-specific features that characterize the corresponding pathology. Therefore, we have not limited the tool to the simultaneous analysis of two cohorts. For the exploration of more than two cohorts, the experts have to interact with the camera settings of some views (such as the pixel, hexagon and violin plot).

Advanced Visualizations. The complexity of the underlying data requires advanced visual exploration techniques to provide a detailed insight into their behavior. Since physicians are generally unfamiliar with such visualizations, *GUCCI* required a certain amount of learning to use them. Initially, they found concepts like the pixel plot very abstract. After an introductory explanation of these concepts, the physicians were able to explore cohorts within minutes. The concept of overview and detail display (like pixel plot together with line plot) was then described as very convincing to find conspicuous features which are then explored in more detail. The two-part splitting into the visualization of the global features and the display of the plane-based data was also emphasized as well suited. The integrated display of global and plane-based features would lead to severely overloaded views.

Performance. *GUCCI* was tested on a standard desktop computer with an Intel Core i5 with 2.8 GHz, 16 GB RAM, and an Nvidia GeForce GTX 1080 Ti. Based on the 122 data sets, we achieve a real-time performance of approximately 40 FPS, except for computing the 2D atlas, the line bundling, and the outlier detection. The performance test was conducted with a resolution of 2560×1440 pixels so that all views exhibit an appropriate size. More complex calculations like the line bundling were outsourced to external threads to keep the user interface interactive. We also artificially increased the number of data sets to test the stability of *GUCCI*. In consultation with our medical expert, the number of data sets is expected to increase to 1000 in the next five years. Even based on this number, no performance problems occurred.

Limitations. At the moment, the number of patient records is rather low. This negatively affects the quality of the supervised feature selection and excludes the use of machine learning for small cohorts. Moreover, PC have known disadvantages. The number of displayable features is limited. We recommend not to select more than ten features. Another problem are the varying value ranges of features. The minimum and maximum are therefore displayed on each axis. Furthermore, certain visualizations are easier to extend than others. While new features could be easily integrated in the pixel plot, the addition of new measurement planes in the hexagon plot is not easy without losing the space-saving representation. In principle, other n-gons can be used for more planes, but they cannot be arranged as efficiently as hexagons.

Generalizability. To use *GUCCI* for other applications, the corresponding data should comprise a set of numerical features, which can be classified into global and local features. Global features are defined in a larger area of the underlying domain and do not have a fixed spatial reference point. Local features are defined at a fixed spatial location within the domain. Additionally, it is advantageous if both classes can be divided into time-independent and time-dependent features.

Our study has shown that depending on the number of features, feature reduction methods should be provided to minimize the

effort to define and analyze cohorts. In the case of sufficiently large data sets and a clearly formulated target variable, automatic feature reduction methods should be used. Otherwise, manual cohort definition should be supported by, e.g., correlation-based approaches combined with clustering techniques to remove correlated features and thus enable analysis of small cohorts.

The comparative visual analysis of cohorts w.r.t. several global and local features, which can change over time, is a challenging problem. Multiple views are needed to show qualitative and quantitative data aspects, where it is important that all views are linked to support the detection of outliers and interesting subgroups. Our case study shows that the integration of high-level abstractions for viewing both time-independent and time-dependent features combined with visualizations that allow a more detailed analysis of individual features is a powerful way to investigate cohorts. However, the visualizations alone do not ensure an efficient exploration process. The visual analysis should be guided by providing automatic methods that highlight outliers as well as features, which are very similar and dissimilar between cohorts. This tracks the user's attention to the most important data aspects and avoids an open-end exploration. In summary, *GUCCI* provides useful techniques to the visualization community that can also be used in other domains.

GUCCI could be used, e.g., to investigate other vessel diseases regarding their blood flow features. A growing field is the analysis of cerebral aneurysms, which are pathological bulges of the cerebral arteries. The analysis is based on simulated blood flow data with the aim of assessing the risk of a rupture of the vessel wall. The goal is to better understand the interplay of morphology and blood flow in order to find characteristic differences between ruptured and non-ruptured aneurysms. *GUCCI* could be used for the analysis of these cohorts, since time-dependent and time-independent features are also analyzed, whereby the number of features should be reduced first. The representations for the plane features would have to be adapted for this purpose, so that, for example, features at different locations on the aneurysm could be analyzed comparatively to find endangered wall regions.

In addition to blood flow data, *GUCCI* could also be easily adapted for the analysis of other cohort data, such as the *SHIP* data [59]. Important aspects would also be the reduction of features and the definition of interesting cohorts. Furthermore, time-independent and time-dependent features are given, which could be analyzed with *GUCCI*. The specific analysis of local features in a certain anatomical region would be less focused here. Similarly, *GUCCI* could be used in domains outside of medicine, if the underlying features can be grouped into local and global features with a time-independent or time-dependent behavior. Possible applications would be, e.g., the analysis of weather or social media data.

8 CONCLUSION AND FUTURE WORK

Medical experts are strongly interested in a cohort-based analysis of cardiac blood flow data to derive normal value ranges and establish medical guidelines. For this purpose, we presented the *GUCCI* framework, which allows a comparative analysis of measured blood flow data in the aorta for predefined cohorts. The user is guided through visual analysis of the complex data by highlighting strong differences and similarities between cohorts. In this way, *GUCCI* enables a detailed study of CVDs as well as the diversity of HV w.r.t. aortic flow information.

In the future, we want to integrate and visualize blood flow data in other heart vessels. Diseases, such as stenosis, also occur, e.g., in the pulmonary artery. Thus, a comparative data analysis in these vessels is important to better understand and treat these diseases. In addition to quantitative flow properties, qualitative features seem to influence the development and progress of CVDs. Therefore, we plan to integrate the 3D vessel geometry and patterns of internal blood flow into the data analysis to extend the qualitative cohort analysis. Furthermore, we plan to improve the discovery of new, interesting cohorts based on the hemodynamic and morphological features. For this purpose, suitable clustering methods and interactive visualizations of the cluster results are required. Moreover, we want to perform a larger user study with more data sets and experts using *GUCCI* over a longer period of time. Finally, we are planning a merger of *Bloodline* and *GUCCI*, based on a database that manages all data sets. The combination of both tools would result in a very powerful software. For newly added data sets, the pre-processing and segmentation process from *Bloodline* would be performed, as well as the calculation of flow features. Subsequently, the exploration of predefined cohorts presented in *GUCCI* would follow with the proposed visualizations, whereby newly added data sets would be inserted automatically.

ACKNOWLEDGMENTS

This work was partially funded by the Carl Zeiss Foundation. The authors would like to thank Sebastian Ebel and Boris Riekens for the fruitful discussions on these topics as well as Benjamin Köhler for providing the *Bloodline* system, the basis for this research.

REFERENCES

- [1] C. J. François, S. Srinivasan, M. L. Schiebler, S. B. Reeder, E. Niespodzany, B. R. Landgraf, O. Wieben, and A. Frydrychowicz, "4D cardiovascular magnetic resonance velocity mapping of alterations of right heart flow patterns and main pulmonary artery hemodynamics in tetralogy of Fallot," *J Cardiovasc Magn R*, vol. 14, no. 1, p. 16, 2012.
- [2] R. Lorenz, J. Bock, A. Barker, F. von Knobelsdorff-Brenkenhoff, W. Wallis, J. Korvink, M. Bissell, J. Schulz-Menger, and M. Markl, "4D flow magnetic resonance imaging in bicuspid aortic valve disease demonstrates altered distribution of aortic blood flow helicity," *Magn Reson Med*, vol. 71, no. 4, pp. 1542–53, 2014.
- [3] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben, "4D Flow MRI," *J Magn Reson Imaging*, vol. 36, no. 5, pp. 1015–36, 2012.
- [4] K. Suwa, O. A. Rahman, E. Bollache, M. J. Rose, A. A. Rahsepar, J. C. Carr, J. D. Collins, A. J. Barker, and M. Markl, "Effect of aortic valve disease on 3D hemodynamics in patients with aortic dilation and trileaflet aortic valve morphology," *J Magn Reson Imaging*, vol. 51, no. 2, pp. 481–91, 2020.
- [5] E. Heiberg, C. Green, J. Toger, A. M. Andersson, M. Carlsson, and H. Arheden, "FourFlow-open source code software for quantification and visualization of time-resolved three-directional phase contrast magnetic resonance velocity mapping," *J Cardiovasc Magn R*, vol. 14, no. S1, p. W14, 2012.
- [6] B. Köhler, M. Grothoff, M. Gutberlet, and B. Preim, "Bloodline: A system for the guided analysis of cardiac 4D PC-MRI data," *Computers & Graphics*, vol. 82, pp. 32–43, 2019.
- [7] E. M. Semaan, M. Carr, M. Gulsun, P. van Ooij, A. Stalder, J. Carr, J. Collins, M. Markl, and B. SPOT-TISWOODE, "Evaluation of an optimized post-processing tool for 4D flow MRI data analysis in healthy volunteers and patients with aortic stenosis, aortic insufficiency, and aortic aneurysm," in *Proc. of ISMRM*, 2014, p. 3948.
- [8] B. Köhler, S. Born, R. F. van Pelt, A. Hennemuth, U. Preim, and B. Preim, "A survey of cardiac 4D PC-MRI data processing," vol. 36, no. 6, pp. 5–35, 2017.
- [9] W. V. Potters, H. A. Marquering, E. VanBavel, and A. J. Nederveen, "Measuring wall shear stress using velocity-encoded mri," *Curr Cardiovasc Imaging Rep*, vol. 7, no. 4, p. 9257, 2014.

- [10] S. Ebel, J. Dufke, B. Köhler, B. Preim, S. Rosemeier, B. Jung, I. Dähnert, P. Lurz, M. Borger, M. Grothoff *et al.*, “Comparison of two accelerated 4d-flow sequences for aortic flow quantification,” *Scientific reports*, vol. 9, no. 1, pp. 1–10, 2019.
- [11] J. M. Tyszka, D. H. Laidlaw, J. W. Asa, and J. M. Silverman, “Three-dimensional, time-resolved (4D) relative pressure mapping using magnetic resonance imaging,” *J Magn Reson Imaging*, vol. 12, no. 2, pp. 321–29, 2000.
- [12] D. L. Gresh, B. E. Rogowitz, R. L. Winslow, D. F. Scollan, and C. K. Yung, “WEAVE: A system for visually linking 3D and statistical visualizations applied to cardiac simulation and measurement data,” in *Proc. of IEEE Visualization*, 2000, pp. 489–92.
- [13] H. Doleisch, M. Gasser, and H. Hauser, “Interactive Feature Specification for Focus+Context Visualization of Complex Simulation Data,” in *Proc. of IEEE/EG VisSym*, 2003, pp. 239–48.
- [14] H. Doleisch, M. Mayer, M. Gasser, R. Wanker, and H. Hauser, “Case Study: Visual Analysis of Complex, Time-Dependent Simulation Results of a Diesel Exhaust System,” in *Proc. of IEEE/EG VisSym*, 2004, pp. 91–6.
- [15] T. Ropinski, S. Oeltze, and B. Preim, “Survey of glyph-based visualization techniques for spatial multivariate medical data,” in *Computers & Graphics*, 2011, pp. 392–401.
- [16] A. Rocha, U. Alim, J. Silva, and M. Sousa, “Decal-maps: Real-time layering of decals on surfaces for multivariate visualization,” *IEEE Trans Vis Comput Graph*, vol. 23, no. 1, pp. 821–30, 2017.
- [17] R. van Pelt, R. Gasteiger, K. Lawonn, M. Meuschke, and B. Preim, “Comparative blood flow visualization for cerebral aneurysm treatment assessment,” *Comput Graph Forum*, vol. 33, no. 3, pp. 131–40, 2014.
- [18] M. Meuschke, T. Günther, R. Wickenhöfer, B. Preim, and K. Lawonn, “Visual Analysis of Aneurysm Data using Statistical Graphics,” *IEEE Trans Vis Comput Graph*, vol. 25, no. 1, pp. 997–1007, 2018.
- [19] M. Meuschke, S. Voß, F. Gaidzik, B. Preim, and K. Lawonn, “Skyscraper visualization of multiple time-dependent scalar fields on surfaces,” *Computers & Graphics*, 2021.
- [20] B. Preim and K. Lawonn, “A survey of visual analytics for public health,” *Comput Graph Forum*, vol. 39, no. 1, pp. 543–80, 2020.
- [21] J. Krause, A. Perer, and H. Stavropoulos, “Supporting Iterative Cohort Construction with Visual Temporal Queries,” *IEEE Trans Vis Comput Graph*, vol. 22, no. 1, pp. 91–100, 2016.
- [22] C. Turkey, A. Lex, M. Streit, H. Pfister, and H. Hauser, “Characterizing cancer subtypes using dual analysis in Caleydo StratomeX,” *IEEE Comput Graph*, vol. 34, no. 2, pp. 38–47, 2014.
- [23] M. D. Steenwijk, J. Milles, M. Buchem, J. Reiber, and C. P. Botha, “Integrated visual analysis for heterogeneous datasets in cohort studies,” in *IEEE VisWeek Workshop on Visual Analytics in Health Care*, 2010, p. 3.
- [24] D. Ceneda, T. Gschwandtner, and S. Miksch, “A review of guidance approaches in visual data analysis: A multifocal perspective,” *Comput Graph Forum*, vol. 38, no. 3, pp. 861–79, 2019.
- [25] E. Packer, P. Bak, M. Nikkilä, V. Polishchuk, and H. J. Ship, “Visual analytics for spatial clustering: Using a heuristic approach for guided exploration,” *IEEE Trans Vis Comput Graph*, vol. 19, no. 12, pp. 2179–88, 2013.
- [26] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski, “Characterizing guidance in visual analytics,” *IEEE Trans Vis Comput Graph*, vol. 23, no. 1, pp. 111–20, 2016.
- [27] P. Kilner, G. Yang, R. Mohiaddin, D. N. Firmin, and D. B. Longmore, “Helical and Retrograde Secondary Flow Patterns in the Aortic Arch Studied by Three Directional Magnetic Resonance Velocity Mapping,” *Circulation*, vol. 88, pp. 2235–47, 1993.
- [28] T. A. Hope, M. Markl, L. Wigström, M. T. Alley, D. Miller, and R. Herfkens, “Comparison of Flow Patterns in Ascending Aortic Aneurysms and Volunteers using Four-Dimensional Magnetic Resonance Velocity Mapping,” *J Magn Reson Imaging*, vol. 26, no. 6, pp. 1471–9, 2007.
- [29] A. Frydrychowicz, M. Markl, D. Hirtler, A. Harloff, C. Schlensak, J. Geiger, B. Stiller, and R. Arnold, “Aortic Hemodynamics in Patients With and Without Repair of Aortic Coarctation: In Vivo Analysis by 4D Flow-Sensitive Magnetic Resonance Imaging,” *Invest Radiol*, vol. 46, no. 5, pp. 317–25, 2011.
- [30] S. Ebel, J. Dufke, B. Köhler, B. Preim, B. Behrendt, B. Riekens, B. Jung, C. Stehning, S. Kropf, M. Grothoff *et al.*, “Automated Quantitative Extraction and Analysis of 4D flow Patterns in the Ascending Aorta: An intraindividual comparison at 1.5 T and 3 T,” *Scientific reports*, vol. 10, no. 1, pp. 1–13, 2020.
- [31] J. Schulz-Menger, D. A. Bluemke, J. Bremerich, S. D. Flamm, M. A. Fogel, M. G. Friedrich, R. J. Kim, F. von Knobelsdorff-Brenkenhoff, C. M. Kramer, D. J. Pennell *et al.*, “Standardized image interpretation and post processing in cardiovascular magnetic resonance: Society for Cardiovascular Magnetic Resonance (SCMR) board of trustees task force on standardized post processing,” *J Cardiovasc Magn Reson*, vol. 15, no. 1, p. 35, 2013.
- [32] Z. Zhang, D. Gotz, and A. Perer, “Iterative cohort analysis and exploration,” *Information Visualization*, vol. 14, no. 4, pp. 289–307, 2015.
- [33] U. Niemann, A. Neog, B. Behrendt, K. Lawonn, M. Gutberlet, M. Spiliopoulou, B. Preim, and M. Meuschke, “Cardiac cohort classification based on morphologic and hemodynamic parameters extracted from 4d pc-mri data,” *arXiv preprint arXiv:2010.05612*, 2020.
- [34] L. Breiman, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [35] B. Shneiderman, “The eyes have it: A task by data type taxonomy for information visualizations,” in *Proc. of IEEE Symposium on Visual Languages*, 1996, pp. 336–43.
- [36] J. Kehler and H. Hauser, “Visualization and visual analysis of multifaceted scientific data: A survey,” *IEEE Trans Vis Comput Graph*, vol. 19, no. 3, pp. 495–513, 2012.
- [37] C. Seifert, V. Sabol, and W. Kienreich, “Stress Maps: Analysing Local Phenomena in Dimensionality Reduction Based Visualisations,” in *EuroVAST@ EuroVis*, 2010.
- [38] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv preprint arXiv:1802.03426*, 2018.
- [39] L. van der Maaten and G. Hinton, “Visualizing data using t-sne,” *J Mach Learn Res*, vol. 9, no. 86, pp. 2579–2605, 2008.
- [40] J. Heinrich and D. Weiskopf, “State of the art of parallel coordinates,” in *Eurographics (STARs)*, 2013, pp. 95–116.
- [41] C. B. Hurley and R. Oldford, “Pairwise display of high-dimensional information via eulerian tours and hamiltonian decompositions,” *J Comput Graph Stat*, vol. 19, no. 4, pp. 861–86, 2010.
- [42] G. Palmas, M. Bachynskyi, A. Oulasvirta, H. P. Seidel, and T. Weinkauff, “An edge-bundling layout for interactive parallel coordinates,” in *Proc. of IEEE Pacific Vis*, 2014, pp. 57–64.
- [43] A. M. Mood, “Introduction to the theory of statistics.” McGraw-hill, 1950.
- [44] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of time-oriented data*. Springer Science & Business Media, 2011.
- [45] H. Reijner *et al.*, “The development of the horizon graph,” *Electronic Proc. of the VisWeek Workshop From Theory to Practice: Design, Vision and Visualization*, 2008.
- [46] S. Havre, E. Hetzler, P. Whitney, and L. Nowell, “Themeriver: Visualizing thematic changes in large document collections,” *IEEE Trans Vis Comput Graph*, vol. 8, no. 1, pp. 9–20, 2002.
- [47] J. Heer, N. Kong, and M. Agrawala, “Sizing the horizon: the effects of chart size and layering on the graphical perception of time series visualizations,” in *Proc. of the SIGCHI conference on human factors in computing systems*, 2009, pp. 1303–12.
- [48] I. Boyandin, E. Bertini, P. Bak, and D. Lalanne, “Flowstrates: An approach for visual exploration of temporal origin-destination data,” in *Comput Graph Forum*, vol. 30, no. 3, 2011, pp. 971–80.
- [49] C. Ware, *Information visualization: perception for design*. Morgan Kaufmann, 2004.
- [50] D. B. Carr, A. R. Olsen, and D. White, “Hexagon mosaic maps for display of univariate and bivariate geographical data,” *Cartography and Geographic Information Systems*, vol. 19, no. 4, pp. 228–36, 1992.
- [51] B. Preim and M. Meuschke, “A survey of medical animations,” *Computer and Graphics*, vol. 90, pp. 145–68, 2020.
- [52] R. L. Harris, *Information graphics: A comprehensive illustrated reference*. Oxford University Press, 2000.
- [53] B. Köhler, M. Meuschke, U. Preim, K. Fischbach, M. Gutberlet, and B. Preim, “Two-Dimensional Plot Visualization of Aortic Vortex Flow in Cardiac 4D PC-MRI Data,” in *Proc. of Bildverarb Med*, 2015, pp. 257–61.
- [54] M. Meuschke, B. Köhler, U. Preim, B. Preim, and K. Lawonn, “Semi-automatic Vortex Flow Classification in 4D PC-MRI Data of the Aorta,” *Comput Graph Forum*, vol. 35, no. 3, pp. 351–60, 2016.
- [55] A. S. Cheam and M. Fredette, “On the importance of similarity characteristics of curve clustering and its applications,” *Pattern Recognition Letters*, vol. 135, pp. 360–67, 2020.
- [56] M. Novotny and H. Hauser, “Outlier-preserving focus+ context visualization in parallel coordinates,” *IEEE Trans Vis Comput Graph*, vol. 12, no. 5, pp. 893–900, 2006.
- [57] I. Corouge, S. Gouttard, and G. Gerig, “Towards a shape model of white matter fiber bundles using diffusion tensor MRI,” in *IEEE International Symposium on Biomedical Imaging*, 2004, pp. 344–47.

- [58] S. Alemzadeh, T. Hielscher, U. Niemann, L. Cibulski, T. Ittermann, H. Völzke, M. Spiliopoulou, and B. Preim, "Subpopulation Discovery and Validation in Epidemiological Data," in *Proc. of EuroVA Workshop at EuroVis*, 2017, pp. 43–7.
- [59] H. Völzke, D. Alte, C. O. Schmidt, D. Radke *et al.*, "Cohort profile: The Study of Health in Pomerania," *Int J Epidemiol*, vol. 40, pp. 294–307, 2011.



Monique Meuschke is a Post Doctoral Research Fellow at the University of Magdeburg (UoM), and the University of Jena, Germany. In 2015, she received a master in Computational Visualistics from the UoM. In 2019, she received a Ph.D. in Computer Science from the UoM. Her research interests are the visual analysis and exploration of medical blood flow data as well as information visualization and visual analytics for medical applications.



Benjamin Behrendt is a research associate at the Department of Simulation and Graphics of the UoM. In 2015, he received a Master's Degree in Computer Science from the UoM. Since then, he works on explorative visualizations of cardiac and cerebral blood flow data for medical applications.



Uli Niemann is a research associate at the UoM. He studied Business Information Systems and is currently pursuing his Ph.D. in Computer Science. He develops machine learning methods and systems for clinical and epidemiological applications.



Matthias Gutberlet is the chief physician of the department of Radiology of the Heart Centre Leipzig and he holds a professorship for cardiological imaging at the University of Leipzig.



Bernhard Preim heads the Visualization Group at the UoM since 2003. Before, he was a staff member at MEVIS Bremen, Germany. He received a Ph.D. at the UoM in 1998 and a habilitation (venia legendi) in 2002 from the University of Bremen. His research interests are in medical visualization for diagnosis and treatment as well as virtual reality for medical training and visual analytics of medical data.



Kai Lawonn is an assistant professor for medical visualization at the University of Jena. Formerly, he was an assistant professor at the University of Koblenz-Landau. He studied Mathematics at the Freie University of Berlin and in 2014 he received his Ph.D. in Computer Science from the UoM. In 2017, he received a habilitation (venia legendi) at the UoM. His research interests are illustrative visualizations for medical and biological applications.