Stony Brook University

The official electronic file of this thesis or dissertation is maintained by the University Libraries on behalf of The Graduate School at Stony Brook University.

Interactive Visual Analytics In Medical Imaging

A Dissertation presented

by

Ievgeniia Gutenko

to

The Graduate School

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

in

Computer Science

Stony Brook University

May 2017

Stony Brook University

The Graduate School

Ievgeniia Gutenko

We, the dissertation committee for the above candidate for the
Doctor of Philosophy degree, hereby recommend
acceptance of this dissertation

Arie E. Kaufman - Advisor
Distinguished Professor, Dept. of Computer Science, Stony Brook University

Rong Zhao
Director, Software Division, CEWIT, Stony Brook University

Klaus Mueller
Professor, Dept. of Computer Science, Stony Brook University

Feng Qiu
Siemens, Princeton, NJ

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

Interactive Visual Analytics In Medical Imaging

by

Ievgeniia Gutenko

Doctor of Philosophy

in

Computer Science

Stony Brook University

2017

With the size of volumetric medical data constantly growing, demand is increasing for efficient compression, storage, and understanding of data. Along with an increase in the size of information come two new major forms of data consumption. First of all, doctors no longer need to be tied to the hospital room to read through a single patient data but can do it in comfort of their home or on their mobile device. Secondly, widespread of browser-based tools enables new collaborative ways of information analysis and research. In this work, we focus on these two novel ways of analysis of medical information: a mobile device for single-subject data delivery and reading and browser-based collaborative research tools for visual analysis of large groups of subjects.

Mobile devices have a number of unique characteristics that make them suitable platforms for mobile health applications. They are portable and can provide always-on connectivity with increasing throughput rates, effectively allowing a medical doctor to conduct the diagnostic process without being constrained to a static workstation. We construct and evaluate a dual pipeline for volumetric rendering on the mobile device and a remote rendering pipeline. Modern mobile devices also provide us new methods of input as an alternative to traditional mouse-keyboard interaction. We study two

new ways of selection within volumetric data via new methods of input and construct respective models of human interaction. Next, we take a step further and evaluate novel mixed reality platform for simple immersive analysis tasks and interactions with volumetric data.

With easier access to the data through protected hospital network, one can easily gather larger groups of subjects retrospectively for research in a specific problem domain. Often forgotten by laypeople, the spleen has risen as one of the examples of such problems. We develop a browser-based visual analytics application AnaFe to study changes in the spleen over time through multiple image-derived features. We take advantage of the amount of largely reproducible image features to enable construction of similarity-based queries through interaction with these features. Through rapid visual feedback, the user can compare several time-varying insights and their correlating features in a single overview.

Contents

List of Figures

List of Tables

List of Publications

JOURNAL AND REFEREED CONFERENCE PUBLICATIONS:

K. Mirhosseini, **I. Gutenko**, S. Ojal, J. Marino, A. Kaufman. Automatic Speed and Direction Control Along Constrained Navigation Paths. *IEEE Virtual Reality*, 2017

I. Gutenko, K. Dmitriev K, A. E. Kaufman, M.A. Barish. AnaFe: Visual Analytics of Image-derived Temporal Features - Focusing on the Spleen. *IEEE Transactions on Visualization and Computer Graphics*, 2016 [64]

C. Papadopoulos, **I. Gutenko**, A. E. Kaufman. VEEVVIE: Visual Explorer for Empirical Visualization, VR and Interaction Experiments. *IEEE Transactions on Visualization and Computer Graphics*, 2015 [111]

I. Gutenko, H. Peng, X. Gu, M. A. Barish, A. E. Kaufman. Maximal area and conformal welding heuristics for optimal slice selection in splenic volume estimation. *Proc. SPIE 9785, Medical Imaging: Computer-Aided Diagnosis, 97853V*, 2016 [65]

K. Dmitriev, **I. Gutenko**, S. Nadeem, A.E. Kaufman. Pancreas and Cyst Segmentation. *Proc. SPIE 9784, Medical Imaging: Image Processing, 97842C*, 2016 [44]

C. Papadopoulos, S. Mirhosseini, **I. Gutenko**, K. Petkov, A. E. Kaufman, B. Laha. Scalability Limits of Large Immersive High-Resolution Displays. *IEEE Virtual Reality*, 2015 [112]

I. Gutenko, K. Petkov, C. Papadopoulos, X. Zhao, J. H. Park, A. E. Kaufman, R. Cha. Remote Volume Rendering Pipeline for mHealth Applications. *SPIE Medical Imaging*, pp. 903904-903904, 2014 [67]

PUBLICATIONS IN SUBMISSION:

I. Gutenko, K. Mirhosseini, A. Kaufman. Touching the Volumetric Void: Interacting with Scientific Visualization in Mixed Reality. *IEEE Visualization*, 2017

K. Mirhosseini, **I. Gutenko**, S. Ojal, J. Marino, A. Kaufman. Immersive Virtual Colonoscopy. *IEEE Visualization*, 2017

M.A. Barish, **I. Gutenko**, K. S. Baker, A. E. Kaufman. Comparison of unidimensional and multidimensional splenic measurements as a predictor of change in volume over time. *In submission to Radiology*

J. H. Park, **I. Gutenko**, A. E. Kaufman. Saliency-aware compression of Volumetric Data. *In submission to IEEE Transactions on Multimedia*

POSTERS:

I. Gutenko, K. Mirhosseini, A. E. Kaufman. Angle and Pressure-based Volumetric Picking on Touchscreen Devices. *IEEE 3D User Interfaces Symposium*, 2017

Q. Sun, S. Mirhosseini, **I. Gutenko**, J. H. Park, C. Papadopoulos, B. Laha, and A. E. Kaufman. Buyer's Satisfaction in A Virtual Fitting Room Scenario Based on Realism of Avatar. *IEEE 3D User Interfaces Symposium*, pp 183-184, 2015.

I. Gutenko, M. A. Barish, A. E. Kaufman. Visual Exploration of Splenic Size andMorphology. *The 12th International Conference & Expo on Emerging Technologies for a Smarter World (CEWIT)*, 2015

I. Gutenko, A. E. Kaufman Semi-automatic Segmentation of Spleen from Abdominal Computed Tomography Images. *The 11th International Conference & Expo on Emerging Technologies for a Smarter World (CE-WIT)*, 2014

I. Gutenko, A. E. Kaufman. Volumetric Segmentation of Coronary Ar-

teries from Computed Tomography Angiography. *The 10th International Conference & Expo on Emerging Technologies for a Smarter World (CEWIT)*, 2013.

I. Gutenko, A. E. Kaufman. Volumetric Segmentation for Computed Tomography Angiography. *The 9th International Conference & Expo on Emerging Technologies for a Smarter World (CEWIT)*, 2012.

TALKS:

I. Gutenko, M. A. Barish, A.E. Kaufman. What does your spleen say? Visual Analytics for Medical Imaging. *Grace Hopper Celebration for Women in Computing*, 2015

I. Gutenko, K. Petkov, C. Papadopoulos, X. Zhao, JH. Park, A. E. Kaufman, R. Cha. Remote Volume Rendering Pipeline for mHealth Applications. *CDDA Workshop*, 2014.

I. Gutenko, K. Petkov, C. Papadopoulos, X. Zhao, JH. Park, A. E. Kaufman, R. Cha. Mobile-based Volume Rendering Pipeline for m-Health. *The 10th International Conference & Expo on Emerging Technologies for a Smarter World (CEWIT)*, 2013.

1 Introduction

1.1 Ubiquitous Visualization: In the Doctor's Pocket and on the Web

Modern visualization of the medical data is not limited to volumetric rendering of Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) data. While decades ago visualizing even a single stack of subject's scan images facilitated a challenge, nowadays such visualization is ubiquitous and accessible to almost anyone with an average computing device. Topics of the modern medical visualization research vary from analysis of novel interaction techniques with volumetric data on the new visualization platforms and devices to advanced computer-aided detection diagnoses. Researchers are concerned with questions on how to fit gigabytes of imaging data into doctor's pocket or how to squeeze only relevant multi-subject information on a single screen of the device.

This dissertation addresses two central topics in the modern medical visualization and analytics. In the first part of this dissertation, we talk about the question of "fitting" medical volumetric data on the mobile device of the doctor. This part of the work addresses visualization of a single-subject's imaging data. On the example of CT angiography data (CTA), we investigate two types of pipelines for mobile health (*mHealth*) applications. Moving beyond the desktop requires new types of interaction with the data. We look into interaction with volumetric data on a tablet device and in the mixed reality. In the second part of this thesis, we utilizes advances in visual analytics research for multi-subject data visualization with a focus on one medical domain application.

While some of the research presented here is made possible only due to advances in the hardware, some contributions to the field of radiology are independent of technology choices. One can observe rapid advances in visualization and computer-aided detection in certain areas of medicine, particularly those that facilitate life-threatening conditions and frequently require immediate attention and action. For example, cardiac imaging is a well researched topic with thousand of publications. At the same time, some other areas remain stagnant due to lack of focus of both medical and visualization research. One of the examples of such, is the spleen, organ often forgotten not only by the laypeople, but also by a number of doctors. It has only recently acquired attention of the radiologists. Monitoring of changes in spleen

is critical for observation of lymphoma and leukemia and variety of other diseases. We focus on the spleen as a domain application for visual analytics research on improving feature exploration in time-varying multi-subject data.

1.2 Visualization of Medical Data beyond the Desktop

Mobile devices (such as tablets and smartphones) have a number of unique characteristics that make them suitable platforms for mobile health applications. They are portable and can provide always-on connectivity with increasing throughput rates, effectively allowing a medical doctor to conduct the diagnostic process without being constrained to a static workstation. Current 4G mobile network technology can provide several megabits worth of downstream bandwidth, making it a viable channel for the transmission of medical data to a device. Additionally, the hardware capabilities of portable devices are increasing at a steady pace. Features like 3D-accelerated rendering and programmable shaders, that only were available to desktop computers a few years ago, are now common place in the mobile segment. Display quality and resolution has also been constantly improving. While it seems that we can do almost anything with modern devices, their viability as platforms for medical diagnosis has a number of key hurdles that must be overcome.

Let us consider CT Angiography (CTA) as a target medical application. CTA data currently can reach resolutions of approximately 512^3 that span over the time domain (while capturing different phases of a heart's beating). If one assumes 16 bits of precision and 10 snapshots, then a single data set can exceed 2.5 gigabytes in size. Transmitting this data naively over an LTE connection with at 10 megabits / second would take approximately 35 minutes. Additionally, even if the data eventually got on the mobile device, rendering it would pose several challenges due to memory constraints, lack of computational power and a shortage of support for key graphics processing features that are considered commonplace on the desktop computer. Thus we can identify two main bottlenecks in the utilization of a mobile devices for mobile health: transmission of the data to the device and rendering data on the device itself. Consequently, we address two different aspects of mHealth pipeline. In the first type of the pipeline, we look into data size reduction for efficient transmission of the data itself in order to perform rendering directly on the device. In the second type, we leave the data on the centralized server location, and construct a remote rendering pipeline.

1.2.1 Data Size Reduction for On-Device Mobile Rendering

The first bottleneck of the mHealth visualization applications the transmission of data from some centralized location to the mobile device. This transmission needs to be efficient and disregard non-essential data. For example, for CTA data, this implies that only the voxels associated with the heart chamber and surrounding vessels are transmitted, effectively necessitating their segmentation from the original volume returned by the scanner. Since the segmented data may demonstrate redundancy in the spatial and temporal domains, performing compression will also result in a reduction of the transmission time. This compression needs to be image-quality cognizant since artifacts in the volume may affect the medical doctor's ability to properly diagnose a potential condition.

1.2.2 The Need of Remote Rendering Visualization

The second bottleneck for mobile visualization is centered around the rendering capabilities. Since these devices are constrained in terms of raw graphics power, the algorithms that are utilized must be able to deliver high visual quality in an efficient manner that allows for fluid interaction. Additionally, they must be implementable on the feature set available to mobile devices, which is a subset of what can be found on the desktop. An alternative strategy would have the rendering actually take place on a computationally powerful workstation and the results transmitted wirelessly to the mobile device. This remote rendering approach would mitigate the computational constraints of a mobile device but impose an even greater burden on the compression/transmission aspect of the pipeline, since the rapid delivery of high quality rendering results would be imperative for an interactive visualization and a fluid user experience. In this dissertation, we show one of the approaches to the construction of the remote rendering pipeline that is essentially mobile-platform independent.

1.3 Immersive Analytics for Volumetric Data

The rapidly developing field of mixed reality environments is creating novel ecosystems for immersive analytics applications. High visual immersion has shown to be beneficial for data analysis and exploration with further benefits from immersion from interaction, such as direct touch. Hybrid systems utilize

tangible and tactile interaction to achieve high immersion in mixed reality. The use of tangible controls, for example, gamepads, is common to virtual reality environments.

However, not all application scenarios may accommodate this type of interaction, either by design or due to hygiene issues. Rising mixed reality applications present the use with a "virtual" hologram visualization. The user can interact with such hologram at a distance of up close in mid-air. However, mid-air interaction is often deemed insufficient for the lack of haptic feedback. In the absence of any such meaningful interactive experience, one must communicate the result of successful interaction to the user via other channels. One may use visual feedback to indicate and change visual properties of an interactive object or its parts. Another channel may be the use of audio feedback to signal the beginning or the end of a particular action.

Interaction with volumetric data consists of a set of goal-oriented and supporting tasks [92]. Most commonly, interactions with 3D volumetric data have been mapped to a 2D screen. Transition to immersive environments requires careful mapping of the input device properties to this set of required interactions.

1.4 Browser-based Visual Analytics

Improvements in the hardware of mobile devices go in parallel with an increase in the power of an average grade personal computers and software. In fact, modern mobile browsers that support WebGL specification allow to create powerful visualization tools without significant difficulties in implementation. Primarily, this is due to the simplicity of implementation in modern scripting languages (such as JavaScript), and ready availability of a set of visualization libraries. The existence of convenient libraries and tools fosters a rapid prototyping of a wide variety of visual analytics applications that target narrow domain specific problems with a focus on analytical approaches, approaches for data discovery and exploration.

1.4.1 Visual Analysis of Heterogeneous Medical Data

Heterogeneous medical data is frequently analyzed for retrospective studies. In such studies basic patient information from health records, and imaging data is obtained retrospectively for one or a sequence of patient visits. Such data presents a particular challenge for medical researchers. Frequently

the number of datasets available is limited in quantity, while the number of derived features to be analyzed is high. Due to the high dimensionality of features, experts might focus only on a subset of data, thus potentially limiting generation of a new hypothesis. Additionally, one might only focus on trends observed during data collection, thus making the hypothesis generation biased towards observations.

Interactive visual analytics introduces several new patterns into the workflow of researchers at various stages. First, one must select apart of population for an observational or exploration study. Cohort studies evaluate hypothesis in samples sharing common characteristics, such as age, gender, or specific risk factors. Visual analytics introduces an interactive cohort selection with visualization of resulting population. Using this new step in the workflow the researcher can select a sample of a population and get immediate overview of this sample and its characteristics.

Secondly, a concept of *hypothesis-free exploration* aims to help *hypothesis generation* through a system of complex and interactive data visualizations of the data itself and its features. Through interactive exploration of features via multiple linked data charts and plots, the researcher is able to find a suitable subset of features and get the feedback on whether the hypothesis requires further investigation. Hence, quality information visualization should help clinicians not only to discover previously unknown patterns, but also identify similarity in groups of subjects with certain symptoms or subjects that match specific criteria.

1.4.2   Focusing on the Spleen

Our work is concentrated around spleen, as one of the "forgotten" and under researched organs. Spleen is one of the organs that is often left behind by the medical imaging community due to the lack of immediate risk to patient's life imposed by splenic abnormalities. However, it is often looked at much more carefully when it comes to disease progression or treatment. In particular, increase in splenic size, also known as splenomegaly, is a common finding in a wide range of abnormal conditions, including immunologic, hematopoietic, infectious, and storage diseases. While splenic size and changes in organ's morphology are important, they have not been looked at in scale and comparison. We design a novel visual analytics system for exploration of multiple parameters involved into change of splenic shape and size. Our system is entirely web-based and designed with modern web technologies. Using our

system, in just few minutes, radiologist can gain insight into a large set of time-varying patient data that has been has been previously looked at only on the individual basis. In this work, we focus on hypothesis-free exploration of robust imaging-derived features. Thus, our work is not limited to one anatomic organ or one imaging modality. We show an extension of our work to prostate data derived from MRI scans.

2 Related Work

In this section we address related work in two domains of medical visualization. First, we talk about volumetric rendering of 3D data on mobile devices, challenges, and existing solutions. We also address specific aspects of interaction with volumetric data rendering, for which modern mobile devices provide unique solutions. Secondly, we talk about visualization of heterogenous medical data. Visualization of such data is frequently built to address a certain problem within a narrow domain application. In this section, however, we abstract from the problem domain and focus only the ways to visualize heterogenous data for groups of subjects for retrospective and cohort studies.

2.1 Mobile Visualization of Single-Subject Data

In this section we will overview some volume rendering techniques including seminal work and state of the art approaches to visualizing volumetric MRI and CT data, and, specifically, approaches to volumetric rendering on mobile devices.

2.1.1 Defining Volumetric Rendering

In 1988 Levoy has introduced the technique for visualizing surfaces from volume data with application to CT to first display a "transparent jello" [96]. The volume rendering scheme introduced in the paper is the first attempt to visualize surfaces without approximation by geometric primitives. This technique also allows to separate shading and classification operations. The main idea is to compute the light transport equation by sampling colors $c_\lambda(x_i)$ and opacities $\alpha(x_i)$ of the voxels with a constant step and performing trilinear interpolation with their neighbors, thus compositing the values to obtain a single color, or simply performing linear interpolation between two pixels along the ray:

$$C_{out,\lambda}(\mathbf{u_i}) = C_{in,\lambda}(\mathbf{u_i})(1 - \alpha(x_i)) + c_\lambda(x_i)\alpha(x_i), \tag{1}$$

where $C_{in,\lambda}$ and $C_{out,\lambda}\mathbf{u_i}$ are the color of the ray where it enters and leaves each sample location respectively, and $u_i = (u_i, v_j)$ is the pixel location. Thus

the color can be computed as:

$$C_\lambda(u_i) = C_\lambda(u_i, v_j) = \sum_{k=0}^{K} \Big(c_\lambda(x_i, y_j, z_k) \alpha(x_i, y_j, z_k) \prod_{m=k+1}^{K} (1 - \alpha(x_i, y_j, z_m)) \Big),$$
(2)

and color value at initial sampling location $c_\lambda(x_i, y_j, z_0)$ would be a fully opaque background color and opacity at the point is equal to 1. The discrete compositing scheme introduced in the paper is an approximation of a well-known rendering integral.

In our work we use only two of the volumetric rendering approaches: ray casting and texture slicing based on seminal works.

Ray Casting. The idea of ray casting is to calculate the volume rendering integral based on points sampled from the volume when rays are cast from the camera into the volume for every pixel on the image plane. This method has been introduced by Levoy in the simple form and was further extended to benefit from acceleration on the graphics hardware and its rapid development. Ray casting allows to render the data at interactive frame rates. It can also benefit from various techniques such as early ray termination or empty space skipping [91].

Texture Slicing. In texture slicing the volume is represented as a stack of adjacent polygons that are aligned with one axis of the volume. Image coordinates of a 2D slice are used to access points in the volume. However, this requires to store three stacks of texture images in the memory to allow for the change of the viewing direction for more than 90 degrees. When the hardware supports 3D textures, when the viewing direction is changed, slices can be recomputed in real time [25].

2.1.2 Native Volumetric Rendering on Mobile Devices

Can modern mobile devices actually achieve high enough frame rate and maintain quality of rendering?

In 2012, Noguera *et al.* [107] has considered a platform specific implementation of basic volume rendering pipeline by considering limitations of OpenGL ES 2.0. They have presented results of implementation of the ray casting and texture slicing volume rendering on iPad2, 4th generation iPod touch, and browser implementation with WebGL. Authors suggest packing 3D stack of images into a pattern of a 2D texture, thus addressing lack of 3D textures in OpenGL ES 2.0. Later, this technique has been improved by

utilizing a RGBA texture to store the volume [107]. Based on the experiments, it can be concluded that the resolution of the volumetric data does not affect rendering performance, as long as data fits in memory. Results obtained from the WebGL prototype on a desktop PC were significantly better since they do not have limitations similar to the mobile device, and are only constrained by the response of the browser.

2.1.3 Remote Volumetric Rendering

When the server is used for rendering the data, much higher interactive frame rates can be achieved. However, this model requires reliable connectivity which is not always possible in remote areas that rely on mobile devices as their primary visualization source. Hence, use of this model would significantly benefit mostly densely populated areas, where this reliable connection can be established.

Lamberti and Sanna [95] have proposed an approach to remote rendering of 3D graphics and have presented a deep analysis of the system based on two applications: surface and volume rendering. The method is a three-tier architecture with a Remote Visualization Server (RVS) cluster consisting of PCs driven by Chromium software. It showcases a **streaming-based approach** to remote rendering. This solution requires significantly smaller bandwidth, however the main bottleneck of it is a time needed for decompression on the client. The requirements of a streaming-based solution can be evaluated quantitatively in terms of total latency based on the amount of time required to perform rendering on the server, amount of time required to extract image data from the frame buffer, resampling and encoding, amount of time to stream new frame and updating resulting visualization on the client, amount of time required for a command from a client to reach the server side, influence of the frame rate on the server on the frame rate on the client, and even window size on both client and server side. Using these measurements authors evaluate the 3D volume rendering system with simple Mobile 3D Viewer on the client side [95]. For the volume size 512^3 the authors claim to achieve 30 fps when 8 rendering servers are involved and streaming video size is 240x240 pixels. By increasing streaming video size to 512x512 pixels, the same configuration loses more than twice in performance. The paper evaluates the system only in "perfect" laboratory conditions and claims packet loss of around 0.3 percent for wireless LAN and 0.5 percent in GPRS/UMTS environments. These conditions clearly do not correspond to the real world

scenario, where signal loss and deterioration, especially in remote areas, that could benefit from this kind of visualization, will play a significant role.

Campoalegre *et al.* [26] propose an image-based client-server model for visualization on a mobile device. The paper introduces a transfer function aware compression scheme that basically disregards fully opaque regions of the volume. Since the bricking scheme is used for volume representation, a convenient property of Haar wavelets can be utilized for transformation: applying n steps of the Haar transformation to entire volume gives the same result as locally applying these n Haar steps to each individual block. This method allows for multi-scale resolution representation and progressive rendering.

2.1.4 Interacting with Volumetric Rendering Data on Touchscreen Devices

As Butson *et al.* [23] noted, use of the devices that doctors are already familiar with in everyday life makes an educational process much easier due to ease of navigation compared to complicated workstation application. One of the contributing factors is a natural use of a finger as a pointing device. However, it is also expected that fast frame rate has to follow this fast finger movement. Assuming that we can achieve high enough frame rate on the modern mobile device, will the finger itself be sufficient for interaction with volumetric data?

In our work, we consider specifics of interaction with volumetric data. Particularly, we focus on the problem of picking in volumetric rendering. Volumetric picking is a method of a point acquisition within a volumetric rendering of the 3D data [96]. Such picking may be based on one or many rays cast through the volume from a 2D point on the screen. There are several types of known volumetric selection techniques: first hit-based, the user-defined threshold of opacity accumulation along the ray profile, intersection of rays, metadata-based, and based on a known structure of interest. First-hit ray picking results in "touching" a first non-transparent point in the volume. Thus, it constrains the user's ability to pick points inside the volume itself. Threshold-based selection [56] compares opacity accumulated along the ray to the user-defined threshold value.

Structure-aware or contextual picking is mostly used in a medical domain. For example, blobby structures within the volume can be selected using mean shift algorithm for the center of mass projection of the ray [115]. Peng *et al.*

use this method for direct pinpointing in gigabyte-sized microscopy images. Additionally, they introduce a method of picking based on the intersection of two rays cast from different angles.

Unknown structures within volumetric data can be revealed with opacity [125] or feature [103] peeling techniques. The first technique considers opacity accumulation along the ray. The second one considers extrema points of accumulation along several simultaneously traced rays. However, these methods do not allow picking of the target point within the volume. Wiebel *et al.* [147, 148] inspect opacity change along the ray profile to determine the selected object based on its visibility. They introduce a single slice rendering at the target point. Such rendering provides feedback to the user and allows to adjust the position of this slice.

Most of these techniques are concerned with the precision of picking of the exact location within the volume by resolving ambiguities associated with perceived rendering. It is important to notice that all of the above techniques perform picking by casting a straight ray through the volume. Our technique provides the user control over both angle and depth of the picking ray. We study whether both of these parameters improve volumetric target acquisition time.

Unlike our problem, a 2D target picking is well studied. Fitt's law [49] provides a quantitative model of a human motor system and speed-accuracy trade-offs in the tasks of pointing and selection. Paul Fitt notes that it is possible to determine experimentally the amount of noise that interferes with information transmission (selection). We believe that our study may serve as a base for deriving a quantitative model of human behavior for volumetric selection. Importantly, we point to perceptual "noise" (accumulation of volumetric intensities) along the picking ray.

In our work, we propose considering an angle of a digital stylus to alter picking ray direction. Kopper *et al.* [89] studied target acquisition at an angle for a distal pointing task. Trivariate target selection model includes depth of the target [59] but no angle. More importantly, the latter model is limited to a specific type of 3D environment, and might not generalize to volumetric rendering.

2.1.5 Interacting with Volumetric Data in Mixed and Virtual Environments

In immersive virtual reality environments and on large display surfaces, interaction with virtual objects has been widely studied [77] though various **types of inputs**: direct manipulation, tangible, tactile interaction, or the combination known as hybrid approaches. Overall, the choice of types of input and their mappings has always been evasive [51]. The combination of several interaction paradigms, frequently a touchscreen device, creates a hybrid type of interaction [12, 14, 99, 133] or can be used in a locally coupled interaction paradigm [80]. The recent study by Besancon *et al.* [14] discusses the hybrid interaction design through the combination of both tactile and tangible paradigms for 3D scientific visualization. Lopez *et al.* [99] propose to use a touch navigation on a mobile device (such as a tablet) to interact with rendering in a combined monoscopic and stereoscopic viewing environment. Such setup addresses the issue of combing high visual immersion with high immersion through interaction. What-you-see-is-what-you-feel (WYSIWYF) approach [133] integrates the natural user interface of a wall display and a handheld device with multi-touch and 3D-tilt sensing capabilities. Ballagas *et al.* [12] also propose to use a smartphone as a universal input device for a set of interaction to manipulate position, orientation, and selection techniques. Direct touch interactions provide an immersion from interaction in addition to immersion from the visualization [93]. Some of the carefully designed touch wall displays or touch surface interactions [37, 53, 87, 156] can be utilized in the hybrid interaction settings. Interaction with a medication data on a 3D touch table for for orthopedic surgery planning was introduced by Lundstrom *et al.* [101].

Common **interaction tasks** for exploration of 3D data, and specifically, volumetric data, require usage of a set of **interaction techniques**. Most of these techniques have been widely studied and provide excellent modeling of both experienced and inexperienced user behavior.

Object selection techniques, such as ray casting, have been evaluated both in virtual and stereoscopic environments [9]. Commonly, the user is provided with a cursor feedback during the selection task. For example, a ray cursor is often recognized to be superior to a 3D point cursor in a single target environment by Grossman *et al.* [60]. Hincapié-Ramos *et al.* introduce a ray casting technique for augmented reality HMDs [73] with the use of orientation sensor held in users' hand.

While a selection of 3D objects in a 3D environment is studied for subjects of accuracy, it may not always be straightforward in various data types, due to their complexities and ambiguities. In this sense, a selection is referring to **selection of an area/volume of interest** within the dataset itself. Such selection may be done by placing an encapsulating object such as bounding box or sphere within the dataset. More intricate selection techniques may be employed for selecting regions of particle data Yu *et al.* [154]. TeddySelection and CloudLasso techniques [154, 155] address the need to spatially select a subset of a 3D particle cloud by simply encircling the target particles on the screen using either the mouse or direct-touch input. SpaceCast, TraceCast, and PointCast [155] are able to infer a user's subtle selection intention from gestural input. These can deal with complex situations such as partially occluded point clusters or multiple cluster layers, and can all be fine-tuned after the selection interaction has been completed. Owada *et al.* [110] introduce a volume catcher technique for selection of the region of interest via user interaction with volumetric rendering. The user performs a free-form 2D stroke, and the system performs a sweeping plane selection and segmentation based on this selection.

Picking refers to selecting a single point within the dataset (check with reviewer feedback). Such interaction is essential for working with volumetric data. Previous section describes related works for volumetric picking. **Measurements** for volumetric data can be conducted by picking several points within the data or by utilizing a more complex set of tools [68, 123, 128].

View adjustment and manipulation can be as complex as the "world in miniature" (WIM) metaphor [36] for supporting interactive querying and data exploration. In such setting, large-scale detailed data visualization is accompanied with an interactive miniature which can be adjusted and manipulated to interrogate volume data. Alternatively, view adjustments can include removal of several "layers" of the volumetric data [103, 125], or a style setting. We refer to some of these works to show a possible way to adjust the view in the mixed reality setting.

Clipping / cutting plane positioning is a form of view adjustment that allows the user to clip part of the view. For example, in the hybrid setting, a cutting plane can be directly manipulated using a commonly available mobile device [133]. Issartel *et al.* [79] propose a method of clipping plane slicing for handheld augmented reality via a set of tangible tools. De Guzman *et al.* [41] develop two physical forks that allow school children to explore a 3D virtual model of the human body with a cutting plane.

Bare hand interactions are commonly generalized as a set of gestures or postures aimed at a set of data/widget manipulations. Among such bare hand interactions, Isenberg *et al.* [78] differentiate between gestures, postures, and quasi-postures for interaction in 3D context for a direct-touch environment. Camera- and sensor-based tracking (Microsoft Kinect[1] and Leap Motion[2]) with or without the help of fiducial markers is frequently used to recognize gestures and map them to manipulations required by the application [21, 85, 153]. Either one or both hands may perform the interaction. Based on such usage, manual activities can be classified into unimanual (performed with a single hand), bimanual symmetric, and bimanual asymmetric (where two hands perform same coordinated gesture or require a complex coordination, respectively) [63]. This theory of two-handed interaction design is followed by Cuttler *et al.* [40] through a synergistic combination of 3D tools and interactive techniques.

The "touching the void" metaphor applies to intangible displays [30] that essentially float in mid-air. Mid-air interactions with 3D data are most commonly found in stereoscopic environments or augmented reality settings [10, 16, 28]. As the dimensionality of the visualized data is the same as an interaction space, one can take advantage of an additional depth. In this scenario, the fundamental tasks of object manipulation no longer require mapping of 3DOF to 2D interactions [18]. When a 2D surface is used as an input, a common parallax problem occurs between the two images shown to both eyes [19, 20, 38, 77, 142, 143]. In some cases, the input is performed in the space of the data visualization itself. Such input may suffer from touch-through [30, 138] or invisible wall problems.

Lack of important tactile [127] and tangible feedback in mid-air interaction calls for a stronger emphasis on other types of feedback, for example, visual and audio. The importance of multimodal feedback in mobile augmented reality has been studied by Hurst *et al.* [75]. Alternatively, some studies aim to overcome the limitations of traditional planar displays and to provide missing haptic feedback through the use of commodity materials [122].

In the virtual reality environments, it might be important to visualize user's hands (i.e. in the form of an avatar) as means of providing feedback. In stereoscopic environment setting, visual feedback should work under all

---

[1]https://developer.microsoft.com/en-us/windows/kinect
[2]https://www.leapmotion.com

viewing conditions [8] and not cause a distraction to the user. In the mixed or augmented reality, the user can see their hands. However, during the direct touch interaction with objects, user's hand may become occluded, and the user may not have the knowledge of the precise hand position. In fact, even imprecision of tracking hardware may cause errors in a variety of tasks [102]. Thus, the system must provide visual feedback for the location of the user's hand.

2.2 Visual Analytics of Multi-Subject Data

The majority of related works have addressed the visualization and exploration of several types of medical information: heterogeneous data from cohort studies (including temporal information), data from electronic health records (EHR), and visual analysis of physical and feature spaces for medical data. Here we combine the analysis of image and non-image data, but focus on user interactions with imaging-derived features that change over time.

2.2.1 Interactive Visual Analysis of Heterogeneous Data

Recent visual analytics research focuses on the creation of highly interactive tools [88, 141] and data organization solutions [7, 135]. The power of such tools is in *hypothesis-free* [135, 141] exploration that allows domain experts to iterate over multiple data variables. Thus, researchers do not need to rely solely on intuition and observations from clinical practice, but can get a visual summary of multiple parameters at once.

Steenwijk *et al.* [135] have proposed a conceptual framework for heterogeneous temporal patient data organization analysis. This framework defines domains, features, mappings, and studies, and combines them into a relational database. The data-cube model [7] for cohort studies handles partially overlapping data subsets and provides higher computation efficiency in comparison to a relational database. While the proposed model accounts for the aggregation of multi-timepoint data (included in the study), there are no visualizations presented that allow direct comparison over the course of time. Their visual linking of spatial and non-spatial views is limited to viewing a single set in a 3D view at a time. We find the need to align temporal subject sequences based on imaging-derived features and allow simultaneous viewing of multiple 3D sets.

15

Klemm *et al.* [88] have introduced an interactive visual analysis (IVA) workflow for epidemiological cohort studies targeted at domain experts. Their work supports the definition of demographics subgroups, among subjects with lower back pain, driven by spine shape clustering features. The mean representative 3D spine shape is visualized over a familiar information visualization with statistical information. In the epidemiology domain, Chui *et al.* [35] study role of age, time in disease progression through a combination of outcome pyramids, time-series and image plots. Through such combination health professionals are able to gain insights into distribution of disease outcomes.

Keefe *et al.* [83] have highlighted the importance of combining a 3D data overview and 2D information visualization for muti-dimensional data analysis. They take advantage of an overview visualization of 3D small multiples of pig mandibles and related chewing traces. We find their overview display particularly helpful for rapid visual pinpointing of differences when compared to single [7] or mean shapes [88]. For our application, small multiples overview of temporal spleen data substitutes for traditional sequential analysis. Highly interactive visual analytics tools for cohort studies are particularly powerful with the addition of descriptive statistics of multiple data dimensions [141]. In this form they enable dual analysis for hypothesis generation. In our work we focus on a combination of imaging-derived features and physical objects to enable a rapid comparison of disease progression.

2.2.2 Visual Analytics of Physical and Feature Space

Researches have combined the visualization of physical and feature spaces for medical data. Feature space may be defined as measurements taken by clinicians or features derived from imaging data. For example, WEAVE [58] combines the visualization of measurements and an anatomical representation for cardiac simulation. Raicu [120] has summarized mining knowledge from medical imaging data based on features derived from CT. He pays particular attention to tissue classification, which includes classification of spleen tissues with lowest sensitivity and precision values due to its similar attenuation to liver.

The feature space can further be divided into higher level representations, for example models of shape, or lower level representations per voxel. Per-voxel features are explored in works by Fang *et al.* [47] and Blass *et al.* [17]. Per-voxel features enable the creation of time activity curves (TAC) which are

important in the areas of nuclear medicine [47]. Tightly coupled views of per-voxel feature projections enable pattern finding and interactive segmentation from multi-dimensional data [17]. In the same domain, Raidou *et al.* [121] proposes visual analysis of tumor characterization based on dimensionality reduction techniques.

Higher level representations allow for the exploration of larger collections of data. For example, shape variation can be explored through visualization of a given 3D shape within a projected shape space [22]. Busking *et al.* [22] introduce a framework comprised of three views. It supports the exploration of shapes over a population and individual shape progression. Statistical deformation models (SDM) (Hermann *et al.* [72]) are used to study anatomical shape covariation interactively. Caban *et al.* [24] give an overview on visualization of SDM. However, it is important to note that per-voxel feature visualization result in significant clutter and might no necessarily emphasize the differences among shapes.

2.2.3   Interactive Visual Analysis of Temporal Patient Data

Multi-subject visualization of temporal data necessitates the organization of a wide range of information (5 W's [157]), allows one to focus on patterns and specific scenarios of events in the emergency room during patient intake [45, 106], and finds differences and similarities in temporal patient data (CoCo [104]). Finding similarities in patient cohorts is addressed in: *LifeFlow* [150], *CareFlow* [117], *DecisionFlow* [57], and *OutFlow* [149]. Not only analysis, but also construction of the cohort is a challenging question when dealing with large amounts of information [90].

It is important to emphasize that while these works focus on visualization of similarities and differences in temporal patient histories, they do not include analysis of imaging-derived features.

Medical Visualization beyond the Desktop

The rapid advancement of mobile platforms and platforms for virtual and augmented reality brings unique challenges to the field of medical data analysis and visual analytics. In this chapter, we address two major bottlenecks associated with an expansion of the medical data analysis beyond the desktop or workstation: rendering and interaction of large medical data on mobile devices and in mixed reality. Briefly, contributions of this chapter can be summarized as follows:

- First, we consider possibilities of rendering volumetric medical data on the mobile device. We introduce transfer function guided saliency-aware lossy-to-lossless compression scheme.

- Secondly, we built a framework for remote rendering of volumetric data for thin clients. We leverage the rendering power of a server computer and portability of a mobile device for always-on connectivity.

- Next, we look into differences in the interaction paradigms between the 2D touchscreen devices and a traditional desktop workstation. We take advantage of the new pressure-sensitive touchscreen display to explore in depth interaction with volumetric data. Our work studies a problem in volumetric picking based on angle and pressure of the input stylus device.

- Finally, we explore mid-air interaction with scientific volumetric data in mixed reality. We map hand position and orientation to a set of natural mid-air gestures for volumetric data navigation and exploration.

3 Towards On-Device Mobile Volumetric Rendering

On-device volumetric rendering requires full transmission of the data to the device. Data compression is performed to reduce the size of information to be transmitted. Specifically for medical data loss of any information is not permissible. Hence, while lossy compression schemes provide much higher compression ratio, one cannot use them for data such as CTA. Progressive compression schemes enable transmission of the data first at low resolution and quality and up to fully lossless data reconstruction. Thus, the data can be initially presented to the user at low quality or resolution and further progressively refined. Additionally, often only a certain region/volume (ROI/VOI) of a large dataset is of interest to the user. Extraction of this region can help reducing compression, transmission, and visualization time, as well as storage space. For example, non-uniform encoding allocates fewer bits or completely disregards less significant regions, thus providing a higher compression ratio. Another option is the transmission of the VOI first at low quality or resolution and progressive refinement. Such techniques significantly reduce the amount of information that needs to be transmitted first and provide an initial subset of data for evaluation much faster.

3.1 Saliency-aware Compression of Volumetric Data

We consider the situation when the relevance of specific regions within the volume is unknown but can be derived from the underlying data values. In particular, we suggest reordering data based on a computational model of the transfer function-guided (TF) visual saliency. In the following work we first present a 3D visual saliency model for volumetric data, then integrate and evaluate it with a set of lossy-to-lossless compression schemes.

3.1.1 Role of the Transfer Function

Medical software used in clinical practice utilizes a set of pre-set TFs and facilitates a simplified interface for their adjustment. Analysis of CTA data follows a protocol, for example, CTA of thoracic aorta [50] includes pre-set TFs in the software, such as that of TeraRecon (San Mateo, CA). In the user interface of the software, the default value of the TF is specified in two ways. First, the default TF is grayscale, set as a linear ramp from black
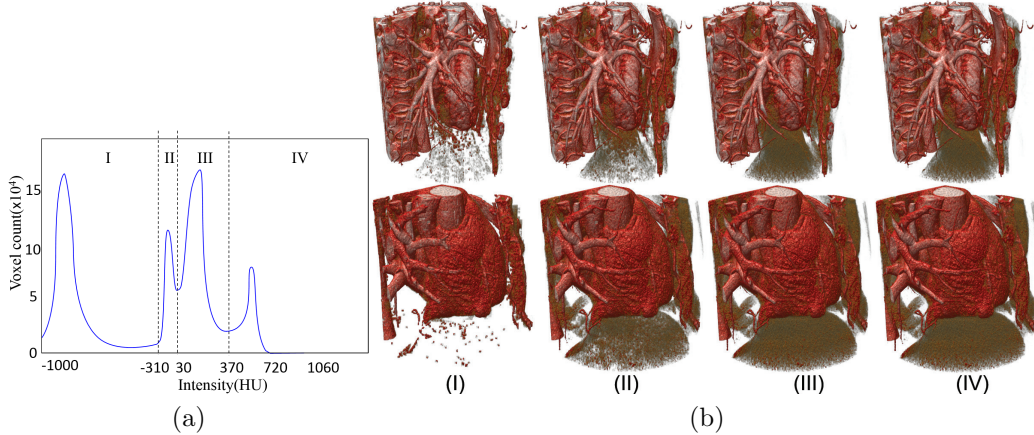
Figure 1: A histogram (a) for two 512^3 chest CTA datasets (b - top and bottom) with the same fixed TF used for both datasets. The histogram is separated into components based on the material properties: I - Pulmonary parenchyma; II - Pericardium, diaphragm; III - Myocardium, hepatica parenchyma, coronary arteries without contrast; IV - Ventricles, coronary arteries with contrast, vein, aorta, bone, pulmonary vessels. The results of progressive transmission and rendering at the first through fourth iteration (I-IV).

to white (window center 324, window length 186) at full opacity. Second, TF is set in a red to a yellow range on a triangle ramp (window center 200, window length 106) at opacity 0.4. Properties of CT allow for such fixed differentiation between tissues types (Figure 1 (a)) [61]. Figure 1 (b) shows a use case, where we select identical TFs for two datasets, similar to those used in the medical software, and produce progressive volume rendering, on the target device. Availability of pre-defined transfer functions, allow us to utilize this data to determine saliency of the data.

3.1.2   3D DCT-based Visual Saliency

We introduce a 3D block-based saliency and transfer function guided compression scheme of volumetric data that is content- and spatially scalable. Saliency map is constructed by weighted averaging the coefficients of the 3D Discrete Cosine Transform (DCT) by analyzing low-level features: intensity, color differences, texture energy, orientation, and gradient.

Saliency of each volumetric region is computed from the coefficients of the 3D discrete cosine transform on the 3D volume, that is fully in the transform domain. Our method of saliency detection is largely based on the work by Fang *et al.* [46]. We have extended the idea of the method to 3D and constructed the saliency map of the 3D volume based on the five features: intensity, color, opacity, orientation, and texture (Figure 2).
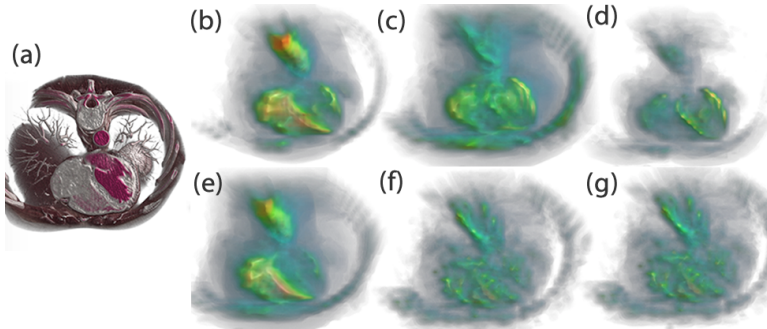
Figure 2: Saliency maps for given transfer function of a chest CT: (a) application of transfer function; saliency map features: (b) opacity, (c) Cr color component, (d) Cb color component, (e) intensity, (f) orientation, and (g) texture features.

3.1.3 Lossy-to-lossless Compression

Our method is integrated into a resolution scalable coding scheme with integer wavelet transform of the image, so it allows the rendering of each significant region at a different resolution and even lossless reconstruction can be achieved. At the target device, the received data is rendered progressively based on its saliency. The benchmarking is performed using three coding schemes with the proposed saliency detection method: SPIHT [84], EBCOT [137], and ESCOT [151] (Table 1).

Prioritized streaming of the data helps to achieve data reduction by rendering regions based on their saliency, and disregarding less essential components. Thus, only after one fifth of the data has been transmitted, the user is able to get overall representation of salient volume regions (Figure 3).

Table 1: Compression rate (bits per voxel) comparison. We compare our multi-iteration (It.) transmission based on our saliency-guided scheme, as well as it's total - combined sum of all iterations, with lossless compression ratio for the full dataset. Three coding algorithms (SPIHT, EBCOT, ES-COT) have been evaluated with Haar wavelet transform. All the datasets are in 8-bit precision.

Dataset	Size	It. 1	It. 2	It. 3	It. 4	1-4	1-3	Full
1 (SPIHT)	512×512×512	0.71	0.24	0.08	1.36	2.40	**1.03**	2.04
1 (EBCOT)	512×512×512	0.37	0.11	0.03	0.97	1.47	**0.51**	1.34
1 (ESCOT)	512×512×512	0.48	0.13	0.03	1.07	1.71	**0.64**	1.30
2 (SPIHT)	512×512×512	0.64	0.46	0.25	0.96	2.30	**1.35**	1.93
2 (EBCOT)	512×512×512	0.35	0.24	0.12	0.70	1.41	**0.71**	1.26
2 (ESCOT)	512×512×512	0.49	0.30	0.13	0.80	1.73	**0.93**	1.22
3 (SPIHT)	512×512×1024	0.78	0.14	0.03	1.43	2.27	**0.96**	2.03
3 (EBCOT)	512×512×1024	0.48	0.14	0.04	1.00	1.65	**0.65**	1.52
3 (ESCOT)	512×512×1024	0.60	0.13	0.04	1.01	1.78	**0.77**	1.34

3.1.4 Conclusions

In this work, we have presented a novel framework for progressive saliency-based visualization of the compressed volumetric data. Our method detects the 3D volumetric salient region by incorporating a block-based weighted averaging (3D DCT) with a predefined or user-defined TF. The size of the created saliency map is orders of magnitude smaller than the original volume, which results in a block structure of the map. To utilize these maps for progressive transmission and rendering, as a first step, we separate the volume into several parts. These salient volume parts are constructed based on the ranges of the values in the saliency map. The ranges can be set constant or varied dynamically. More importantly, we reduce additional redundancy in the transform coefficients before the coding step of the compression by combining blocks with their neighbors depending on the saliency similarity. Compression and transmission of the 3D volumetric data are scheduled based on the ranges of the saliency values starting from the highest. The receiver obtains only a set of data from one saliency range at a time without losing any details, decompresses and renders it after combining it with previously obtained data.
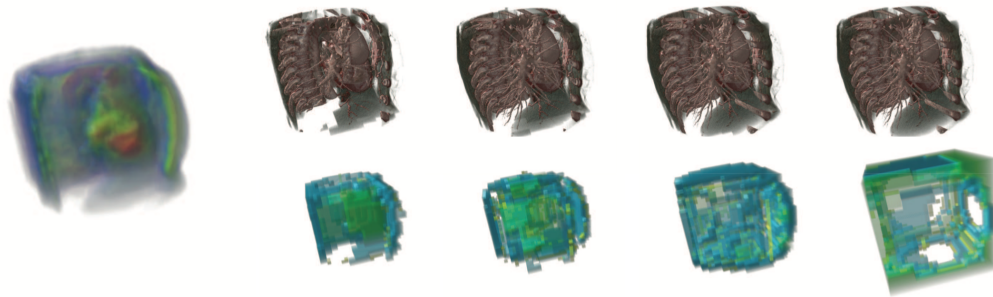
Figure 3: Progressive rendering of the chest CT: (Left) Saliency map, (Top) progressive rendering in four steps, (Bottom) corresponding energy maps.

4 Remote Volume Rendering Pipeline for mHealth Applications

4.1 Enabling Remote Volumetric Rendering

With widely available wireless network connectivity, the proliferation of mobile devices, and the constant increase in the graphics processing power, it is intuitive to remotely process large amounts of medical imaging data on the powerful server, while allowing fast and responsive streaming to mobile devices. The development of these components individually, enables the framework presented in this paper, which is significantly advanced when compared to previous attempts of remote rendering [27, 81]. First, we consider the server side for the processing of medical data that is essentially the driving force in our system. The commodization of GPUs has resulted in the development of clustered visualization applications. While many of those applications involve off-line rendering, for example when dealing with very large scientific datasets or complex movie scenes, remote visualization has been used in certain interactive applications as well. One such example is cloud-based gaming services, where high-quality images are generated in a data center, streamed over a network connection, and the user controls the game on the client side.

Secondly, due to the significant increase in the role of mobile computing over the past few years, we are able to consider advanced mobile devices as suitable clients for our system architecture. Current mobile devices, including various smartphones and tablet devices, have powerful CPU and GPU

characteristics that make them suitable platforms for various mHealth applications. Their configuration, computing capability, display quality and resolution are comparable to desktop counterparts available only a few years ago. However, due to the large amount of volumetric data, transferring the data and rendering it entirely on the device is not always possible [107, 108]. The portability and always-on connectivity of these devices allows our system to stream images rendered on the server to the mobile device, and permits a medical doctor or health care professional to conduct the diagnostic process and follow up without being constrained to a workstation computer in the hospital facility. Therefore, mobile devices are extremely promising as the client side for a remote volume rendering framework.

We leverage this technology to provide high-quality volume rendering targeted at the medical applications on low-power mobile devices. The main contributions of this work are:

- Fast remote volume rendering pipeline for mHealth applications capable of handling large volumetric datasets up to 2048^3;

- High quality responsive streaming of large amount of volumetric data to a mobile clients with a use of hardware accelerate encoding and decoding techniques on both server and client sides;

- Evaluation on mobile health applications such as Computed Tomography (CT) colonography, exploration of other large CT datasets.

4.2   Implementation Details

Our remote volume rendering system is an extension of our existing distributed visualization software and offers a variety of configurations. In the case of remote volume rendering, the two principal components of our system are the image-generating server and a thin client that accepts and transmits user input and also streams images from the server. In this section, we outline the implementation details of each component.

4.2.1   Server Implementation Details

The server component implements a volume rendering pipeline that supports a variety of modalities (raycasting, texture slicing, 1D and 2D transfer functions, different compositing methods, etc). The server can operate in local
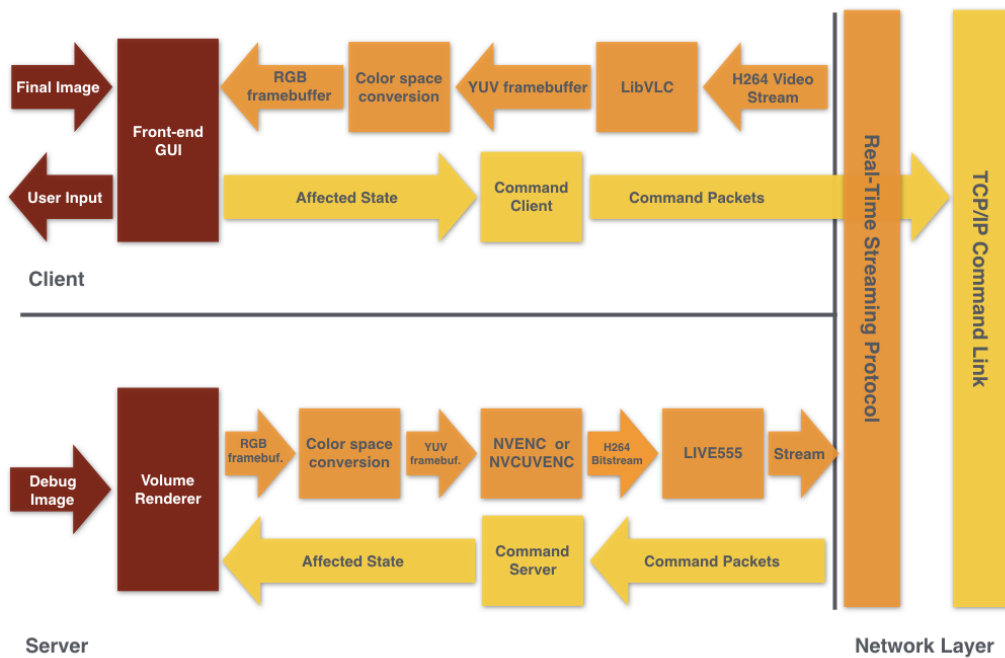
Figure 4: Block diagram outlining the connectivity and flow of data between components of our volume rendering pipeline. Blocks and arrows in orange relate to the streaming of volume renderings over an H.264 video stream, while the command infrastructure is represented in yellow. User-facing inputs and components are annotated in red.

mode, outputting the final image directly to the screen. For remote rendering, the server instead operates in off-screen mode. In this configuration, the server instead outputs the rendering result to a texture, via an OpenGL framebuffer object (FBO). Additionally, in this mode of operation, the server renders at a fixed internal resolution, rather than resizing the OpenGL viewport based on the size of the on-screen window. Also, the aspect ratio of the rendering is decoupled from the on-screen window and is instead transmitted to the server by the client application on startup and when changes in the client's window size occur. The remainder of the volume rendering pipeline is agnostic to the server's mode of operation.

In the final stage of the image synthesis kernel, the server performs a color-space conversion from the RGB output to YUV 4:2:0 on the GPU. YUV 4:2:0 is ubiquitously accepted as an input color space for a number of hardware-accelerated H.264 encoders. In our system's case, the NVENC (for modern Kepler GPUs) also supports the YUV 4:4:4 format that preservers the color at full resolution. Finally, the server reconverts the volume rendering result to RGB in a rapid, screen-space, pass and displays it locally for debug and visualization purposes.

The color-transformed framebuffer is passed on to the encoder subsystem for conversion into an H.264 video stream. This subsystem supports the NVCUVENC and NVENC [4] video encoders, which allow for performant coding of the images on the GPU. The video stream is exposed over the Real-Time Streaming Protocol (RTSP), implemented using the LIVE555 library [3]. A side-benefit of this approach is that the visualization session can be observed not only by the client application (outlined below), but also by almost any commodity device with media streaming support (such as an ARM-based smartphone). A disadvantage of RTSP-based streaming is the fact that it does not leverage the resiliency features of NVENC. We plan on developing a custom streaming protocol that takes advantage of these features, permitting clients to notify the server of transmission errors. In this situation, future frames would be encoded without references to past corrupted frames. In our current implementation, RTSP provides resiliency to such network transmission errors at the expense of input latency.

Finally, the server application instantiates a TCP/IP server, listening for incoming connections on a predetermined port. Successful connection establishes a command link between the server and a single client instance, over which command packets are exchanged. Incoming packets are processed based on a header byte, with different headers denoting packets of different

content. Our system decouples the transmission of different aspects of the render state in order to reduce network overhead. For example, manipulating the virtual camera results in only those parameters being transmitted to the server (rather than communicating the entire render state, including compositing settings, transfer functions, etc.). It is also worth noting that even a single camera manipulation can generate tens of command packets. Rather than repainting the scene once for every incoming command packet, the server maintains a steady refresh rate (30 or 60 frames per second in our implementation) and each frame is generated using the aggregate information from all command packets received up to that point in time.

4.2.2 Client Implementation Details

The client component of our system presents a graphical user interface to the user that is visually indistinguishable from a server instance running in local mode. All relevant volume rendering functionality is exposed via mouse, keyboard and multi-touch interaction modalities. The client maintains a local copy of the render state variables and transmits the appropriate command packets when parts of the state are affected over TCP/IP.

The key subsystem of the client is an H.264 network streaming client, implemented using LibVLC [2]. The utilization of this pervasive media playback framework allows for hardware acceleration of H.264 decoding on certain hardware platforms via the following APIs:

- Intel Media SDK [1] - interfaces with the Intel GPUs available in x86 tablets, such as the Microsoft Surface. The SDK provides hardware-accelerated H.264 decoding and integrates with OpenCL to provide additional processing of the video stream on the client. One major disadvantage is that this SDK is a hybrid system with major computation load on CPU.

- NVCUVID [4] - library for accessing the video decoding functionality of NVIDIA GPUs. Depending on the GPU capabilities, the decoding may be implemented in the CUDA language, or using power-efficient on-board hardware. The resulting images can be further processed in CUDA or combined with client-based rendering in OpenGL.

The decoded framebuffer is presented to the application in YUV format, which is color-converted to RGB on the GPU via an OpenGL 2.0 pixel shader
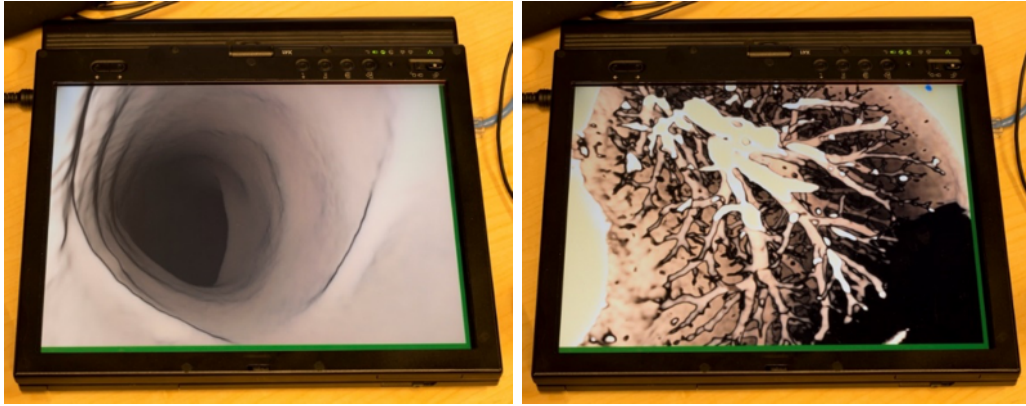
Figure 5: Representative frames of the movie: remote visualization of (left) CT colonography and (right) CT lung volume data

and then presented to the user. This color-space conversion is the most advanced GPU feature required by the client application and it is supported on all x86-compatible integrated GPUs that have been released in recent years. Consequently, the client application can be executed on a vast range of devices and poses minimal hardware requirements. A block diagram illustrating the architecture of our mobile volume rendering pipeline can be seen in Figure 4.

4.3 Results

Our volume rendering system supports a variety of visualization techniques, which are also available during video streaming. We focus our evaluation specifically on features that cannot be implemented on low-powered devices as follows:

- Smoother rendering during interaction - mobile devices often struggle with rendering volumes at high screen resolutions. The visuals of some immersive interactive applications, such as Virtual Colonoscopy, cover the entire viewport the majority of time, making local rendering option unattractive. Contrarily, our client-server approach is not constrained in this way.
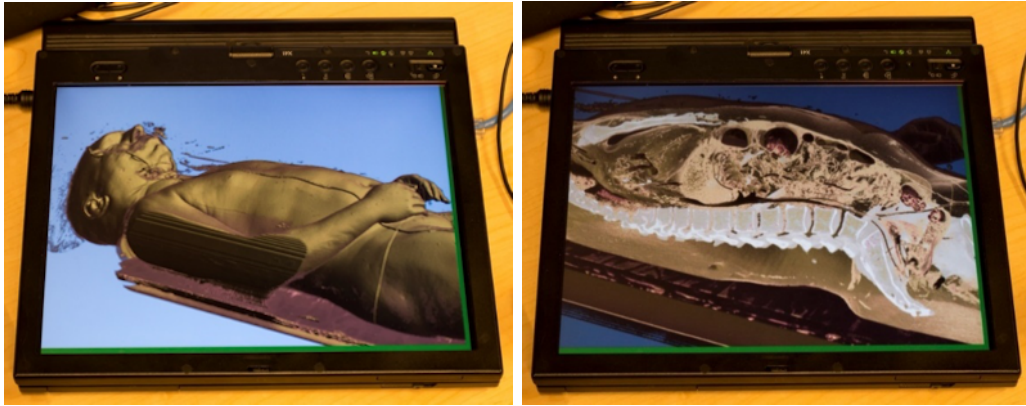
Figure 6: Remote visualization of the Visible Human dataset with (left) a full body view and (right) a sagittal view.

- Higher-order reconstruction filter - we use a cubic filter that is 3 times more expensive than the standard trilinear filter, but provides a higher image quality. The usage of such a filter would be prohibitive when rendering locally on a mobile device.

- Large datasets - our system handles volumes up to 2048^3 at interactive rendering speeds. Rendering such volumes locally on a mobile device would be very difficult, due to space and bandwidth constraints.

Our evaluation is performed on the following hardware: cluster of 6 Dell Precision T7600 workstations with dual 6-core CPUs, 64GB of memory and NVIDIA Quadro K5000 GPU as server; and IBM Thinkpad X41 tablet as client. The tablet is significantly slower than modern devices and does not offer any hardware video decoding support. However, it was chosen purposely to show that our framework is not constrained to the latest mobile devices as clients. We evaluate the results based on several potential applications. The prototype of a CT Colonography (Virtual Colonoscopy) applications (such as [74, 48]) is evaluated and shown in Figure 1(a). The volume rendering is performed on the server side, encoded in the H.264 format and streamed to the client side tablet device. The interactive exploration of the lung data is shown on the Figure 1(b). The data resolution is 512x512x431 which still allows us to provide a 60 fps video stream even with the high order reconstruction filter. The transfer function is optimized for visualization
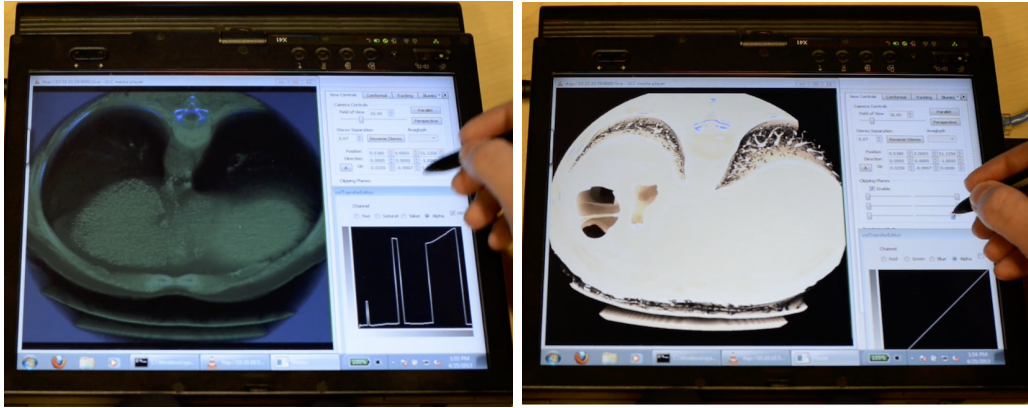
Figure 7: Remote visualization of the cardiac dataset. (left) 1D transfer function editing on the client. (right) Clipping plane positioning on the client.

of the colon wall, yet it also segments the lung structures. Figure 5 shows snapshots of the interactive exploration of the Visible Human [134], which is a 512x512x1877 volumetric dataset. The NVIDIA Quadro K5000 in the server side can provide smooth interaction with the visualization and we show two different rendering modes isosurface rendering of the skin and visualization of internal tissues. Figure 7 illustrates interaction of the user with a client device allowing her to change parameters of the visualization.

Additionally, we measured the network bandwidth required for transmitting the H.264 stream, as a factor of the internal rendering resolution of the server. For this benchmark, we utilized a higher-end tablet device (Microsoft Surface) that would benefit from the additional detail of increased rendering resolutions, due to its 1080p display. We ensured that the synthesized image covered the majority of the view frustum and applied a colorful transfer function to the volume (in order to introduce some entropy to the chroma channels of the framebuffer). We observed a mostly linear increase in bandwidth utilization. For an optimal rendering resolution of 1024^2, we recorded an average bandwidth utilization of approximately 12 megabits / second, while rendering at 30fps. It is worth noting that for higher rendering resolutions (such as 2048^2), network utilization actually drops, as the server is not able to generate frames at a high-enough framerate for particular rendering settings. Our measurements, along with an illustration of the benchmark
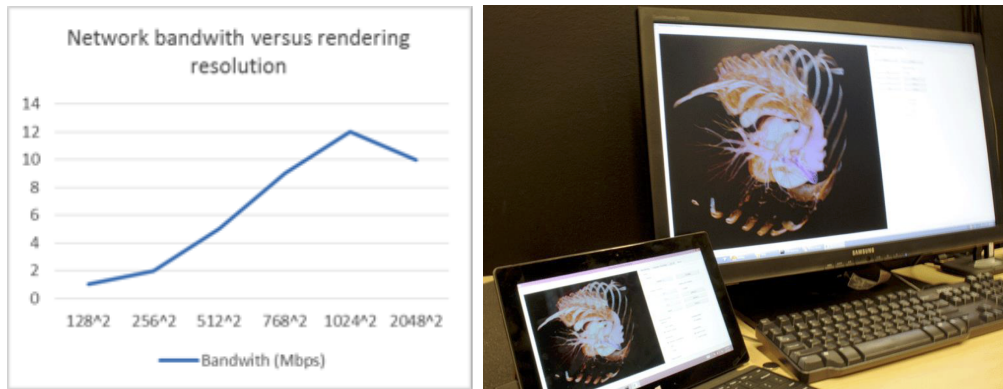
Figure 8: Left: Chart summarizing our network bandwidth utilization benchmark. Right: View of our test system, showing a Microsoft Surface tablet device and a high-end workstation, running the client and server components of our system with synchronized views.

setup are visible in Figure 8.

4.4 Conclusions

In this paper, we demonstrate a remote volume rendering system that leverages high-end workstations to generate the images based on user input and to also encode and stream the results in the H.264 video format to a mobile device. The volume rendering image synthesis happens on the server side, on a dedicated workstation or GPU cluster. Our system can generate multiple views of the volume data, and supports multiple video streams with different bitrates and resolutions for different client devices. Visualization clients can interface with a session either via a commodity media streamer (without support for user interaction) or through a thin-client application that exposes the full user interface and transmits inputs via TCP/IP. Our prototype system can handle various large and complex volume data and provide high quality volume rendering at real-time speeds.

Looking forward, this remote rendering framework can be expanded in two ways. First of all, one can explore novel techniques for natural interaction with the volumetric visualization (such as augmented reality and various forms of gestural interaction). Additionally, one can investigating various avenues for further increasing the rendering performance on the client side

and reducing the network overhead associated with H.264 streaming. For example, our current pipeline performs the volume rendering at a fixed internal resolution, which remains constant throughout the session. We are investigating adaptive rendering schemes that dynamically adjust the volume rendering resolution based on available bandwidth and the complexity of the image.

5 Interaction with Volumetric Data on Touchscreen Devices: Angle and Pressure-based Volumetric Picking on Touchscreen Devices

5.1 Introduction

A number of scientific applications uses the volumetric rendering of a 3D scalar data. Location picking inside this rendering on a 2D screen is an ambiguous task when given an input device such as a mouse. There is a high uncertainty of the target location due to its occlusion by less transparent structures. A lack of depth information from the user's input does not allow to derive intended location to be picked. Additionally, the user is unable to select objects outside the picking ray that is perpendicular to the screen. The need of precise point selection limits visualization of such data to a simple 2D slice rendering or multi-planar reprojection.

Modern input devices, such as a digital stylus pen, allow us to capture both angle and pressure of the user's input. These parameters can be mapped to the angle and depth of the picking ray. In this work, we study the benefits of such mapping to improve volumetric picking. We perform a study of usability of three methods of picking on the touchscreen devices. In the first one, we allow picking based solely on the pressure of the input. In the second method, in addition to pressure mapping, we alter the direction of the picking ray based on azimuth and altitude of the stylus. Finally, we compare these to a single finger picking method with a radius of touch mapped to depth. To the best of our knowledge, this is the first attempt to evaluate picking task for volumetric data formally including depth of the target. Hence, our work provides insights to specifics of interaction with volumetric data. We suggest classification of user strategies for pressure and angle-based picking. We also study the impact of opacity accumulation along the ray profile that contributes to ambiguities of volumetric picking. Our contributions can be summarized as follows:

Our contributions can be summarized as follows:

- a technique that allows users to select 3D targets in the volumetric rendering at any depth and angle on touchscreen devices using both finger and a digital stylus pen input;

- an analysis of usability of three volumetric picking methods;

33
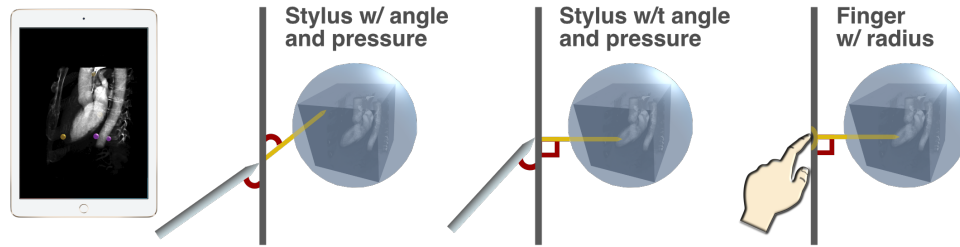
Figure 9: Volumetric rendering on an iPad Pro with picking targets ("target" - yellow, "purple" - selected) and three introduced techniques.

- analysis of user strategies based on 432 observations collected from 12 participants.

5.2 Volumetric Picking

Below we describe our method of volumetric picking on a touchscreen device. The motivation for this work is an overall lack of precision and quality of volumetric target selection, and most importantly, lack of formal evaluation of selection, target pointing or acquisition. We describe a common application setup required to facilitate selection. Here we also define a mapping of the input device parameters to the picking ray.

5.2.1 Technique

Volumetric picking is based on the ray cast from a 2D point on a screen to the point of intersection with a 3D target. Determining the location of the intended 3D target in the volume is an ambiguous task due to the lack of well-defined object boundaries. Volumetric data rendering presents primarily as a bounded cube object [96]. There are no boundaries of objects within the rendering itself. Preliminary data segmentation or iso-surface extraction may provide such boundaries. However, the segmentation of the objects requires data domain knowledge and a series of sophisticated algorithms. Thus, to simply pick a volumetric target along the picking ray, one must be able to specify the length of the ray at the time of the selection.

    Ray picking starts from an invisible surface of origin and perturbs into the volume to a certain depth. The manipulation of the ray direction should

remain consistent regardless of the rotation of the volume cube. We ensure such consistency though the selection of the ray starting point on the minimal bounding sphere to the volume cube. Using an arbitrary surface in the virtual world, such as camera frustum would complicate the task of picking to the user because this invisible surface would alter the direction of the ray arbitrarily, thus making the selection not intuitive.

Such setup is common to all three studied methods:

- stylus picking with an angle and with pressure parameters,

- stylus picking with pressure only (no angle),

- single finger picking.

These methods differ in the input parameters used and how they are mapped to obtain the end point of the ray or the intended location of the target.

5.2.2   Mapping Angle and Pressure

A digital stylus pen provides us three parameters: azimuth and altitude angles of the tilt, as well as the pressure used. Altitude of $\pi/2$ corresponds to a perpendicular angle to the screen surface. Azimuth reaches 0 when the stylus is pointing to the right. The orientation of the stylus is agnostic from the point of ray origin. Hence, it is not affected by our choice of the bounding sphere object. For the ease of implementation, we perform the calculation of the required ray transformation in the screen space. Altitude and azimuth magnitudes correspond to the amount of the right vector rotation around Y and Z axis with a resulting vector in the screen space. This vector is projected to the virtual space based on the screen to world projection matrix.

Pressure of the stylus input determines the magnitude of the picking ray. Thus, the end point of the picking ray specifies the intended location of the selection.

Perspective camera correction must be performed in order to present the user with a picking ray visually extending the stylus. Unlike orthographic cameras, perspective cameras use non-parallel rays for rendering the scene, thus will visually distort the ray. To solve this issue, we perform correction on the rays beside the center of the camera before applying rotation based on tilt.
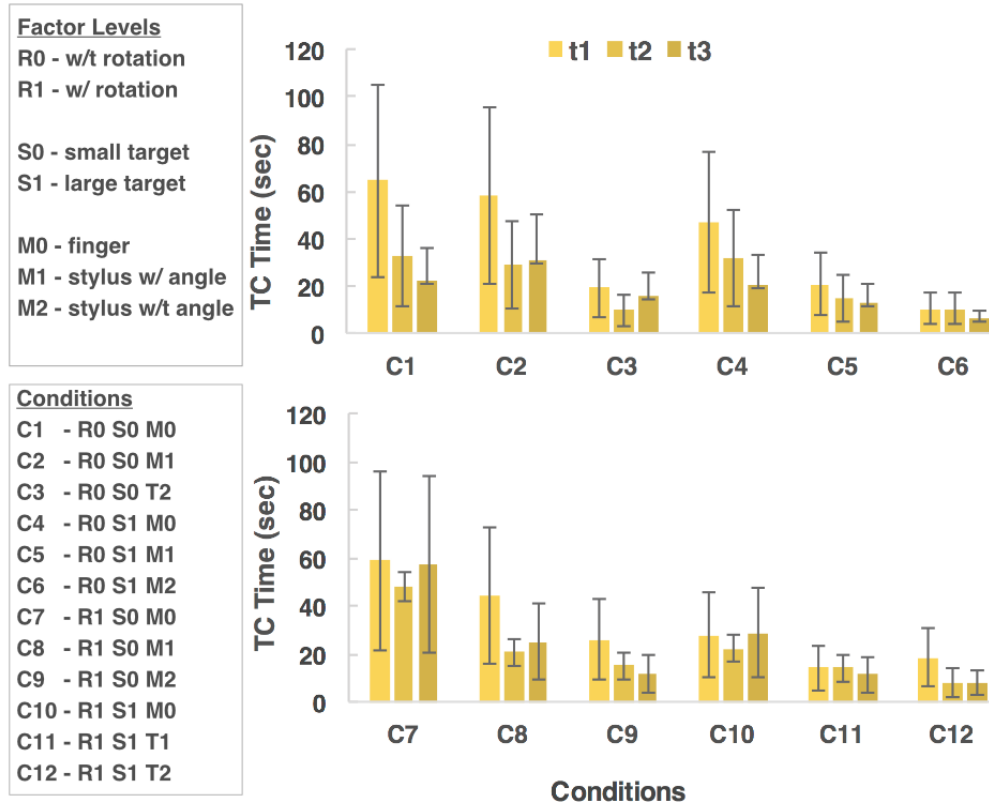
Figure 10: Changes in the mean task completion times over three trials per condition. The error bars indicate the 95% confidence intervals.

5.2.3 Mapping Finger Touch

Picking with a finger is naturally less precise than that with a stylus. The most important implication for our application is a lack of pressure sensitivity to the finger touch on the touchscreen device used in the experiment. To emulate a similar experience to that of a pressure-sensitive digital stylus, we use a range normalized radius of the finger touch point to approximate the pressure applied by the user's finger. In the naïve implementation, due to the coarseness of the finger touch, this might result in an inability to reach targets at arbitrary depth in the volume. However, with an application running at 60 frames per second, we can interpolate between the frame depth values,
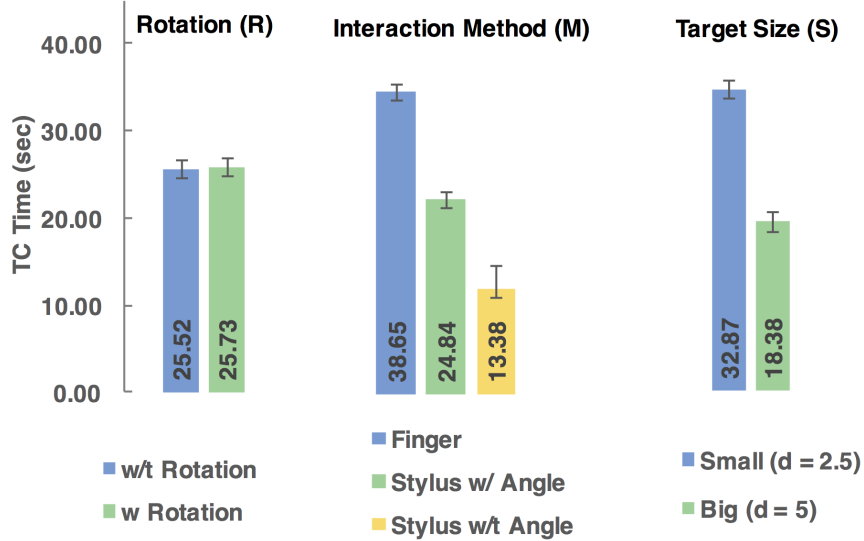
Figure 11: Mean task completion times for each factor. The error bars indicate the 95% confidence intervals.

thus picking along the part of the ray path. With this implementation, we are able to approach selection depth technique of a digital stylus.

5.3 Implementation

Volumetric rendering implementation is preferred with a support of 3D textures. Modern mobile devices deliver OpenGL ES 3.0 and Metal APIs, which enable an implementation of a volume renderer. Our system is made using Unity rendering engine in C# and a Volumetric Raytracing sample [6] based on Nvidia example [5].

To evaluate picking methods, we add a set of **opaque mesh targets** to the volume data. These targets get rendered before the volume data. The resulting rendering shows an overlay effect of the volumetric data. Thus, it provides correct depth information to the user. Selection of the mesh and not volumetric targets enables the use of the default ray collision method of the rendering engine. It also satisfies the requirement of dynamically changing target color upon successful picking. Embedding volumetric targets would not be feasible at run-time due to bandwidth and memory size limitations of

a mobile device.

5.4   Experiment

5.4.1   Apparatus

The experiment was performed on an iPad Pro tablet with a 9.7-inch display. Apple Pencil was used as a digital stylus input device. All subjects sat at a table with a tablet located horizontally on the table.

5.4.2   Task

Participants were asked to perform a target selection task. They were presented with a volumetric rendering of a contrast-enhanced cardiac computed tomography angiography (CTA). 5 yellow target spheres were placed within the dataset at different depth resulting in different amount of occlusion. Upon successful selection, targets changed color to purple to provide user sense of accomplishment.

5.5   Subjects

Our user study had 12 participants, consisting of 9 males and 3 females between the ages of 22 and 35. Participation in the study was voluntarily, and the sample consisted primarily of the university students. Participants filled out a pre-study questionnaire for the knowledge of visualization, use of touchscreen devices, and use of the stylus. All of the users were daily users of touchscreen devices. All but one subjects were right-handed. One participant has admitted color blindness. However, she has indicated that she could perform the test as there was sufficient differentiation among colors of "target" and "selected" states of spheres.

5.5.1   Procedure and Design

A consistent procedure was used for all trials and participants. A repeated measures within-subject factorial design was used. The independent variables were R - rotation of a dataset (full 3D rotation or no rotation), S - target size (diameter of a sphere: 2.5 and 5 unit length targets), M - interaction modality (stylus with or without angle, or finger). A fully crossed design

resulted in 12 combinations of R, S, and M. Each combination was used to create a separate iOS application.

The experiment was performed in one session of 25-50 minutes duration. First, each user was given instructions using a sample application. This application featured a rotation of a dataset and a target picking with a pressure and angle of a stylus. The user was also provided with instructions on the use of the finger for target selection. Following this, each user performed three randomly generated sequences of 12 trials. Participants could take breaks between trial sequences, but not within each sequence. After the first and the last trial sequence, the user filled out general feedback form about the preferred method of navigation and experienced difficulties. Additionally, users had to answer a question about their depth perception of 3D targets and volumetric data.

5.5.2 Measures

For each trial, an iOS application logged sequences of events performed by the user:

- timestamped target picking event and target ID;

- timestamped finger touch event and coordinates; beginning and end of a touch phase;

- timestamped stylus touch event, coordinates, azimuth and altitude of the tilt, pressure; beginning and end of a touch phase;

- rotation quaternion of the volume data.

The dependent variable task completion time (TC) was defined as a difference in time between starting a trial and a last selected target.

5.6   Results

5.6.1   Task Completion Time Analysis

We performed a three-way analysis of variance with repeated measures (RA-NOVA) for task completion time. We average observations obtained from
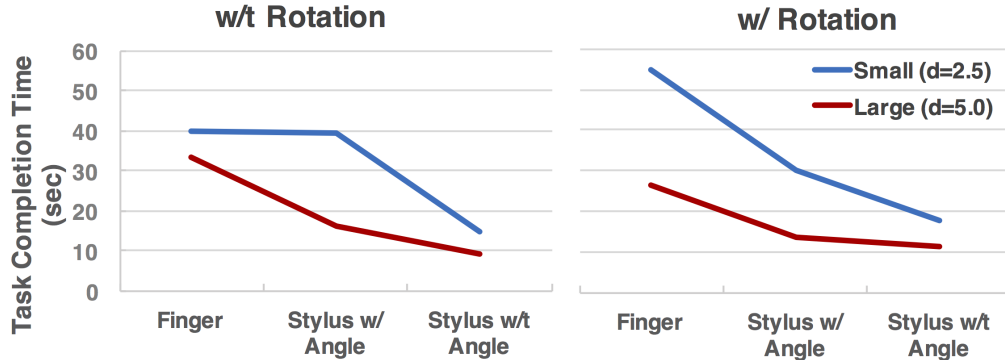
Figure 12: Analysis of interaction between factor levels for independent variables: method of interaction, target size, and rotation mode.

three repeated trials per condition for each subject to reduce variation associated with learning a new interaction technique. Figure 10 shows that there were changes in mean task completion time per condition. The magnitude of these changes was unequal among conditions, which shows different learning rate of techniques. The assumption of normality of the data for task completion time was not met. Thus, the data was transformed using $sqrt(x)$. The transformed data has met the assumption of homogeneity of variances.

5.6.2   Main effects

We found significant main effects for independent variables S ($F_{1,11} = 78.6, p < 0.0001$) and M ($F_{2,22} = 30.4, p < 0.0001$). There was no significant main effect for R ($F_{1,11} = 0.01, p = 0.94$). Figure 11 shows mean task completion times for all the variables.

We have performed multiple contrasts RANOVA test of three interaction modes M. The test has revealed a significant difference between finger-based and both types of stylus interactions ($F_{1,22} = 48.4, p < 0.0001$). Additionally, there was a significant difference between two types of stylus interaction ($F_{1,22} = 12.5, p < 0.002$).
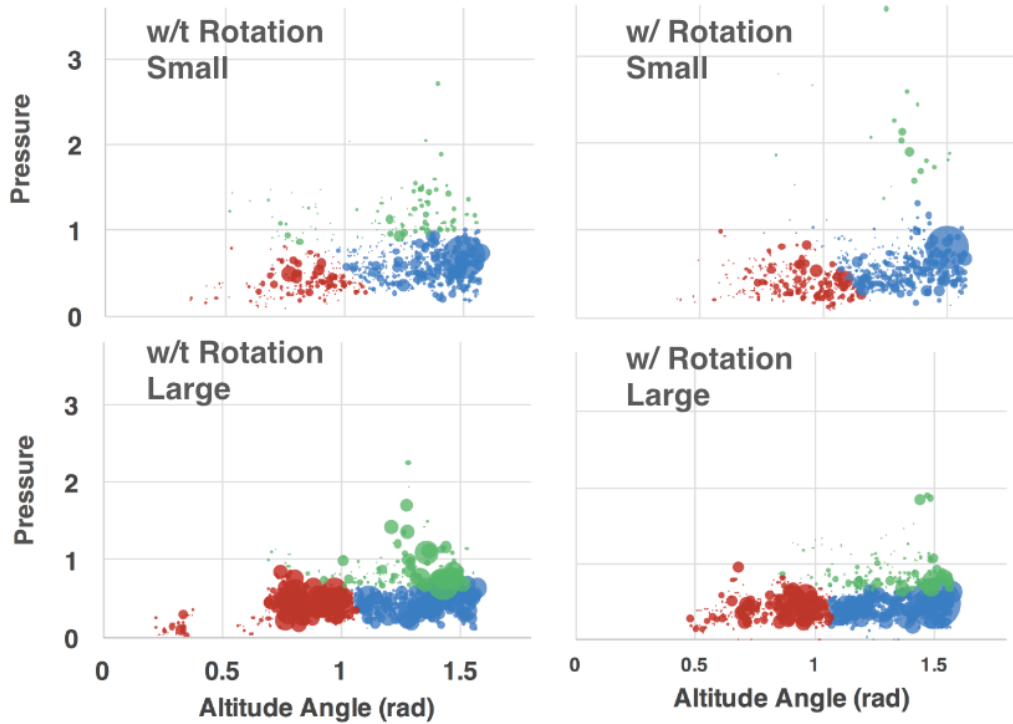
Figure 13: Clustering pressure and angle of the input device during interaction.

5.6.3 Interactions

There was a significant interaction effect between S and M ($F_{2,22} = 5.69, p = 0.01$). Based on further analysis, one can observe that users have experienced challenges when picking smaller targets using stylus with an angle technique 12.

There was a three-way interaction between all three independent variables S, M, and R ($F_{2,22} = 5.68, p = 0.01$). It was observed that when rotation was introduced, users took longer time to pick smaller targets. It can be explained by the fact that most of the users were daily users of touchscreen devices. Hence, they were used to the task of a small 2D target selection, which is essentially similar to our picking task without rotation.
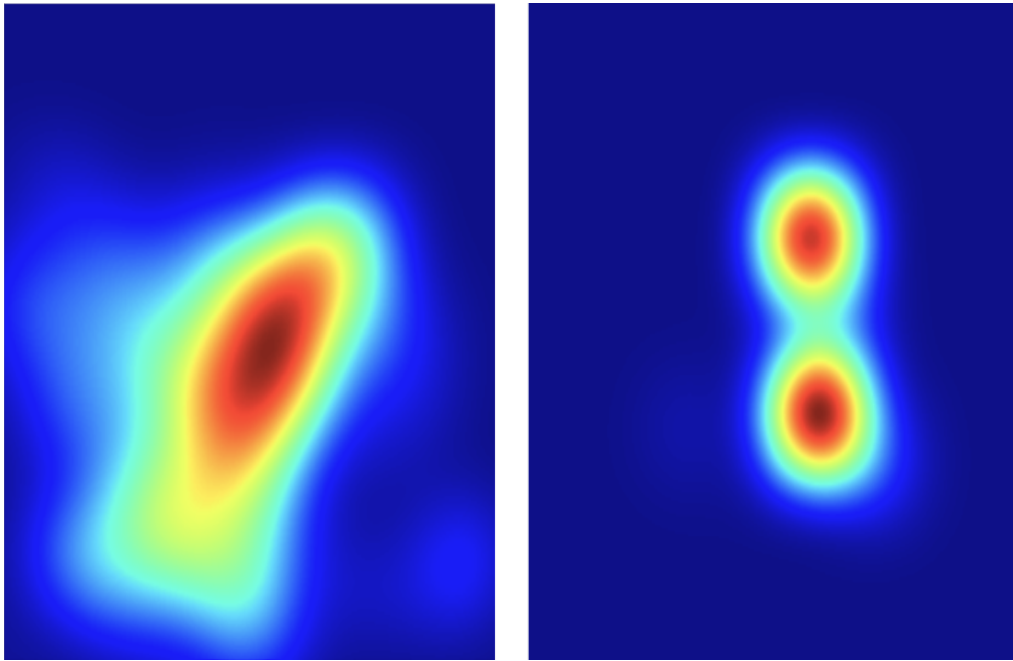
Figure 14: Varying interaction patterns for distinct picking tasks: (left) interaction with rotation, large targets, stylus with an angle, (right) interaction without rotation, with a finger, and small targets. The pattern on the right illustrates difficulties of reaching reaching shallow and deep targets.

5.6.4 Analysis of User Strategies

During experiment sessions, we had observed various user strategies when a digital stylus was used with an angle. While some users have actively used an angle-based picking, some have experienced noticeable difficulties. Determining preferred stylus interaction modes and, in particular, altitude angle of the device is necessary for pursuing research on models of human motor behavior for volumetric data.

For the four conditions of interaction with an angle of the stylus, we have analyzed altitudes of the device. We clustered this data (using *k-means*) into three groups based on altitude angle and pressure (Figure 13). One can observe two dense clusters (red and blue), indicating different user interaction strategies. The third cluster (green) is different for two interaction methods: while it resembles noise for small targets, it lies in proximity to the blue

cluster for the larger targets. The absence of the data rotation mode has also caused difficulties for the user to reach targets that are very far away or very close. Figure 14 illustrates such differences through a heat map of user's finger touches for two distinct conditions.

5.6.5 Analysis of Opacity Accumulation Along the Picking Ray

In our work, we assume that user will aim to pick only visible points. Thus, accumulation of volumetric opacities along the ray should always be less than 1. Intuitively, if the user picks shallow point that is not occluded, the opacity accumulation will be low. If the point is deep, the opacity accumulation will be higher, but still less than one. We have analyzed the data to compare opacities accumulation across the techniques. For the finger picking technique, the accumulation of opacities along the picking ray was the lowest and increased with the smaller target size (mean value of 0.11 for large targets, and 0.13 for small). The reverse trend was observed for stylus-based techniques (for the stylus with no angle technique, for example, mean of 0.38 for the large target, 0.27 for small). The accumulation along the ray was the highest for the stylus with an angle techniques with no data rotation (0.48 for small targets, and 0.41 for large).

5.7 Conclusion

In this work, we present a study on the usability of interaction methods for volumetric picking on the touchscreen devices. We propose and evaluate three techniques for target acquisition: single finger, pressure and angle stylus picking, the pressure only stylus selection. Our study shows significant differences among all three proposed methods of interaction. Importantly, we observe a significant effect when comparing user task completion times for selection techniques with a digital stylus pen with and without an angle. This study provides directions for development of models of human motor interaction for volumetric picking and target acquisition.

6 Interaction with Volumetric Data in Immersive Environment: Touching the Volumetric Void

6.1 Introduction

This part of the dissertation explores interaction with volumetric data via direct touch in mid-air for mixed reality applications. It focuses on an intuitive mapping of input hardware properties to mixed reality interaction techniques. We use Microsoft HoloLens mixed reality head mounted display (HMD) that superimposes "virtual" holograms onto the real world. Out of the box, it provides several ways for users to interact with these holograms. Head-tracking approximates the direction of the user's gaze, while user's hand position is also being tracked. To enrich the set of possible interactions and to create more natural ways to interact with the data, we would like to be able to track orientation of the user's palm. A variety of off-the-shelf camera units and wristbands can provide such information, either based on hand reconstruction from the depth image of the camera, or via a gyroscope sensor. In our work, we take a wristband approach for precision of the values and construct custom wearable piece with a gyroscope sensor for one or both hands.

Mid-air interaction requires consistent visual and audio feedback. Both of these feedback channels have been shown to play an important role [30] in interaction techniques. At the same time, direct mid-air touch provides a clean and hygienic way to interact with the data, which is essential for certain application scenarios, such as surgery planning [101] and assistance [42]. Interaction with CT and MRI data in the current clinical settings is limited to desktop style interaction, thus essential 3 DOF interactions are mapped to 2D input tasks. We aim to find a balance of intuitive and efficient ways to interact with volumetric data through direct mid-air touch. It is important to note that an excess of visual feedback may clutter an already visually complex volumetric rendering. Therefore, a moderate amount of audio feedback may be used to complement it.

Visual analytics has been always of particular interest to data exploration in domain specific applications, especially when it comes to creation of tools and techniques. For example, scientific applications may source data from medicine, material sciences, astronomy, physics and other related fields.

Nearly in isolation, user interface researchers focus on development and evaluation of interaction techniques and their precise evaluation [11]. However, immersive analytics applications require seamless collaborative work from both research fields. In our opinion, to ensure such seamless work one must evaluate modular components of interaction that comprise complex data manipulations. Such evaluation will ensure usability of provided interaction designs. In this work we focus on modular components of the proposed complex gestures. We evaluate types of visual feedback for a volumetric picking task. Additionally, we evaluate the use of gyroscopes for simple rotation tasks to ensure that such rotation can be extended to more complex data manipulations.

Our contributions can be summarized as follows:

- First, we present a mapping of a traditional set of interactions with scientific volumetric data to a mixed reality application using Microsoft HoloLens[3].

- Secondly, we focus on rotation techniques and incorporate that into gestures of rotation using custom hardware.

- Then, we address the importance of visual feedback for mid-air interactions and provide a set of visual and sound feedback metaphors.

- And finally we perform a detailed quantitative study of the selected techniques which comprise complex interactions.

6.2 Interaction Tasks in Spatial 3D Data Visualization

Laha *et al.* [92] survey and classify common tasks for volumetric data exploration across multiple domains. We summarize interaction types for a typical scenario of volumetric data exploration:

- *Navigating with 7 DOF*: Ideally, navigation of volumetric data in mixed-reality settings should be consistent with other 3D data spatial interaction. Hence, we need to provide 3 DOF for translation, 3 DOF for rotation, and 1 DOF for uniform zooming/scaling. Unlike in the projection-based visualization, we should be able to utilize more natural navigation and gestures. To comply with this principle, we propose to use user's hand orientation for data rotation.

---

[3]https://www.microsoft.com/microsoft-hololens/
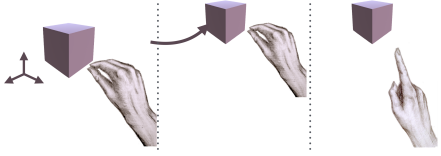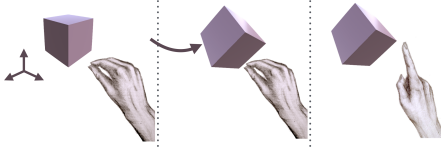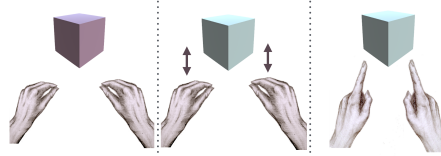
- Data manipulation via *cutting plane* positioning with 3 DOF. Cutting plane positioning is essentially similar to navigation, as discussed above. It is constrained to 1 DOF for control of the depth of the plane, and 2 DOF for its orientation control (azimuth and elevation angles). The plane positioning in mid-air can be performed unimanually (position of a hand mapped to the depth of the plane, and orientation of the hand respectively to the orientation of the plane) or bimanually (by dividing the task between two hands).

- *3D point picking* in 2D projection rendering is difficult due to inability of determining the depth of a "pick" operation. As such, to differentiate simple navigation with point selection and to convey the explicit intent to select a given point with a deterministic depth, the user will need to perform an additional gesture. We have chosen to represent this with a "tapping" gesture to finalize the selection after orientation within the hologram.

- *Selection of a subvolume in volumetric or particle cloud data* also suffers from the limitations of 2D projection based visualization. Special techniques are required to resolve ambiguities of depth specification. For example, when a 3D widget is used to specify a subvolume, one must adjust its boundaries in the helper views. Such interaction can also be simplified to a more natural type of 3D selection widget placement in mixed reality. For instance, the user can approach the volume and specify the selection directly within the volumetric data rendering.

- *Custom viewing modes (with varying DOF)* may require interaction during the visualization. For example, opacity [125] and feature [103] peeling techniques require specification of maximum threshold for the data and the layer. The mapping of such modes should be customized to unimanual or bimanual input.

- *Visualization mode switching* also requires to provide a complete set of gestures while preserving the context of the current interaction. In many applications, the user can switch between visualization modes with a system of widgets, buttons, menus, or sliders. *Voice control* enables easy switching between the gesture recognition modes. Thus, the user does not have to be taken out of current interaction context to select a widget.

- *Interaction mode switching* is another necessity of such systems, and they need a similar level of seamlessness as visualization mode switching, as mentioned above. Switching between various gestures should not take the user out of interaction context. Voice commands make it easy to perform such switching.

It is important to note that we provide sound feedback on the successful start of each gesture.

Table 2: Mapping of the input properties of HoloLens HMD and custom wearable unit to volumetric data interactions in mixed reality.

Interaction	Mapping	Gesture
Voice control - visualization and interaction mode switching		
Gaze / head tracking - selection of an object to interact		
Hand position - 3D volumetric picking		
3D point picking. Hand position in the world space is directly transformed into the volume data space.	*Start*: Position a hand inside the volume in "ready" gesture. *End*: Perform picking with "air tap".	
Translation. User's hand position is directly mapped to the new position of an object.	*Start*: Perform "tap" to select a volume. *In*: Move hand in "hold" gesture to a new location. *End*: Release "hold" gesture.	
Rotation by dragging. The difference in user's hand position is mapped to an arcball rotation identical to that of a 2D mouse input.	*Start*: Perform "tap" to select a volume. *In*: Move hand in "hold" gesture to a rotate. *End*: Release "hold" gesture.	
Custom view modification. Opacity peeling example. Maximum threshold of the data is defined by the location of the user's left hand (optional), while the number of the peeling layer is controlled by the right hand.	*Start*: Perform "tap" with right hand to select a volume. *In*: Move hand in "hold" gesture to peel the layers. *End*: Release "hold" gesture.	

+ Hand orientation - manipulate rotation of the data or a widget

Rotation by hand. Orientation of the user's palm, as measured by the gyroscope sensor, fully defines rotation of an object.

Start: Perform "tap" to select a volume. *In*: Rotate hand in "hold" gesture to a rotate. *End*: Release "hold" gesture.
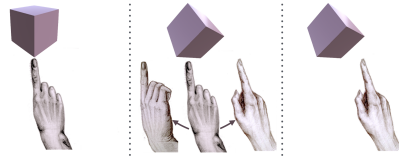
Cutting plane positioning. Orientation of the user's palm, as measured by the gyroscope sensor, is directly mapped to orientation of the cutting plane, while hand position is mapped to the depth of the plane.

Start: Perform "tap" and "hold" gesture. *In*: Rotate the hand with a gyroscope to select orientation of the plane. Move hand to and from the camera to select depth. *End*: Release "hold".
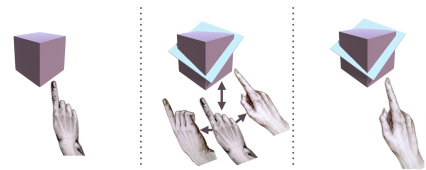
VOI. 3D box widget positioning with two hands in the "ready" gesture. Distance between two hands determines size of the bounding box.

Start: Position both hands inside the volume data to see the bounding box. *In*: Adjust hands to select size of the box. Perform "tap" to clip the volume, "hold" to rotate. *End*: Release "hold".
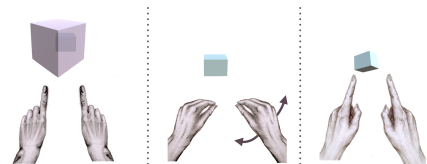
6.3   Design Space of Property Mapping

The hardware used in this prototype implementation provides us with several inputs that can be mapped to the common volumetric data interactions. Here we define these inputs and outline their usage in our application:
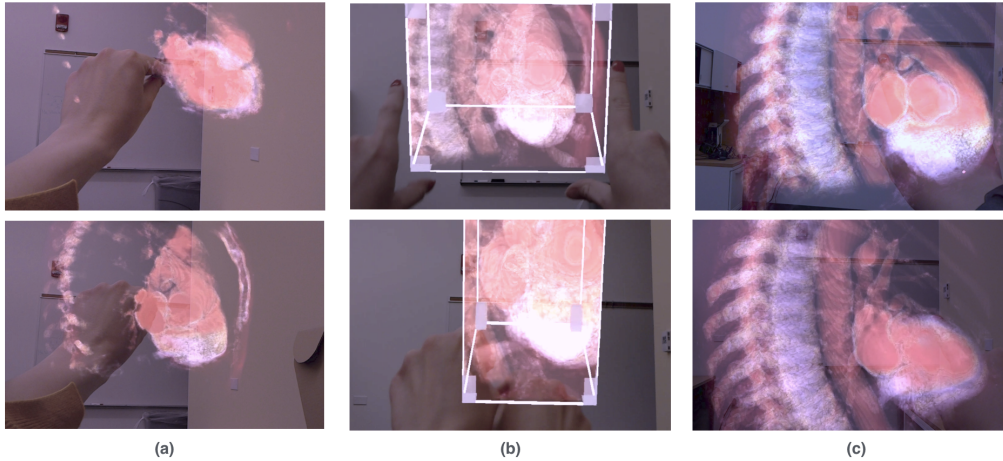
Figure 15: Examples of user interactions with volumetric rendering in mid-air: (a) view modification with opacity peeling, (b) volume of interest selection, (c) clipping plane positioning.

- *Voice controls* - recognition of simple commands triggered by the user. Such commands are an intuitive way to switch between interaction widgets and interaction types. In our work, voice control allows us to switch between a variety of gesture recognizers to allow seamless transition between interactions.

- *Gaze*, defined as user's head orientation, selects the object in focus. Once a certain object is in focus, the user can interact with it until the focus changes.

- *Hand position* - an absolute position of user's hand in the world space as tracked by HMD. We use the hand position to enable visual feedback for interaction at the point of contact with the volume.

- *Hand (palm) orientation* - an orientation of user's palm that is tracked by a custom wearable unit. Hand orientation is used to control the orientation of the entire dataset or the orientation of the virtual widgets used to interact with the volume. For example, in our work, we use the orientation of the user's hand to orient the clipping plane, rotate the box when selecting VOI, and rotate the entire dataset.

Having discussed the requirements of a set of intuitive interactions, as

well as the available input sensors, we can refer to Table 6.2 to see how these input devices can be put to use to fulfill the gestures as discussed above. For instance, positioning a cutting plane (as a gesture to perform data manipulation in a spatial manner) requires a sort of "dragging" gesture, which in and of itself is a multi-part gesture: (1) picking a starting point, (2) rotating the plane and selecting the depth within field of view, and (3) ending the specification. We have opted to keep the gestures consistent throughout our methods, and for this particular example — as seen in the table — we start the selection process by the tap gesture (which corresponds with data point selection), and we use the intuitive gesture of holding onto the point to signify the breadth of the operation. Some of the interactions are captured in Fig 15 for better demonstration.

6.4 Visual Feedback for Mid-air Volumetric Interaction

In this work, we address the importance of visual feedback on volumetric data itself. Through the definition of the transfer function, some of the data may become transparent.

In the mixed reality "virtual" holograms, users' hand may get occluded with a semi-transparent hologram. With an accumulation of opacity from the volumetric data, a user's hand may become fully invisible. Thus, it is important to provide visual feedback on the location of the area of contact. Immersive applications use 3D cursors or rays to indicate the area of interaction. A direct-touch interaction may be accompanied with a touch shadow to show the area of contact or its intensity. Such feedback might not be feasible due to inherent presence of shadows in volumetric data. In volumetric data representation, we aim to provide visual feedback in the location of contact with the user's hand without occluding the data. In the mixed reality scenario, such feedback has to be visible and intuitive for multiple participating users. For example, the feedback upon interaction with mesh surface data may be indicated by the use of opacity or color correction.

We evaluate three types of 3D direct touch feedback for interaction with volumetric data (Fig 16):

- *3D cursor.* We visualize a 3D cursor in the form of a sphere located approximately in the middle of the user's finger. As the interaction is
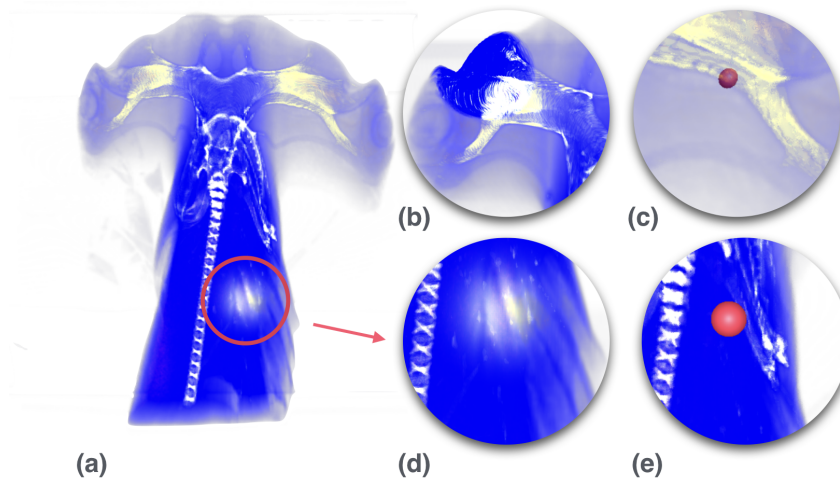
Figure 16: Visual feedback metaphors: (a, d) volumetric "candle"; (b) iso-surface based visualization; (c, e) 3D cursor in front of the volume (e) and inside the volume with data occlusion (c).

happening as a direct touch, there is no need to provide feedback in the form of the ray.

- *Volumetric color "candle"*. We visualize the feedback of the user's finger touch with a highlight as a candle halo centered at the location of contact.

- *Iso-surface "candle"*. We visualize the feedback by amplifying data opacities based on the iso-surface of the underlying data.

Not all interactions require visual feedback. For example, basic navigation provides an immediate feedback by showing a new view of the data. An additional visual feedback is required for the selection of the *volume of interest (VOI)*. The user can select a VOI by extending two hands into a "ready" gesture. The distance between two hands will indicate the initial size of the box in either vertical or horizontal dimensions (depending on the alignment of the user's hands). The orientation of the box can be changed based on the orientation of the user's right hand. The selected VOI box will be shown as a semi-transparent object that does not occlude the volume itself. On "select" gesture — indicated by tapping and holding with both hands — the volume outside of the box will be clipped and only the VOI will be shown.
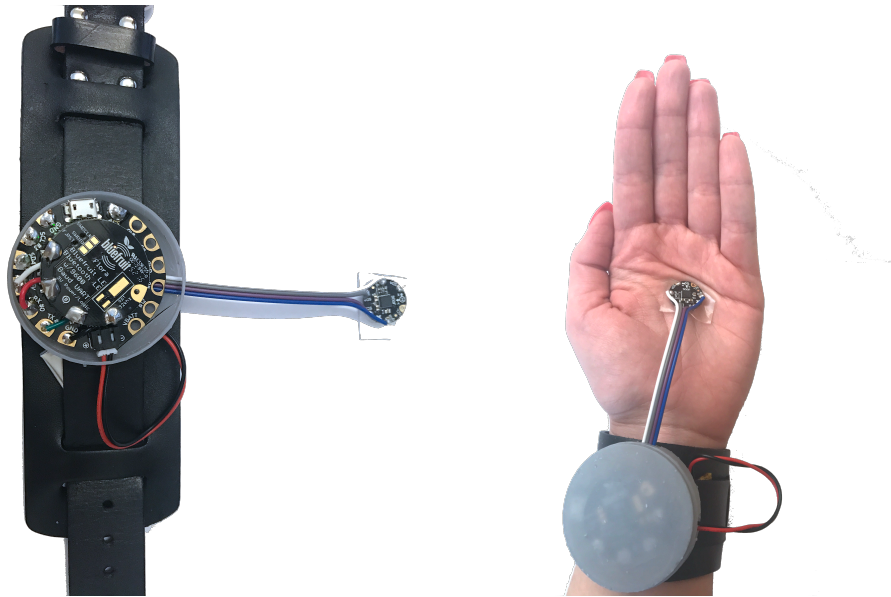
Figure 17: Design of a custom wearable for capturing palm orientation.

6.5 Hand orientation capture

We have built our custom wearable unit using Adafruit's Flora electronic wearable platform[4] (Figure 17). We are using wearable gyroscope sensors (FLORA 9-DOF Accelerometer/Gyroscope/Magnetometer - LSM9DS0 - v1.0[5]) and low energy Bluetooth (Flora Wearable Bluefruit LE Module[6]) for communicating the data back to the HMD.

The communication of the data back to the HMD can be performed in two ways: (1) directly via Bluetooth, or (2) via MQTT, which is a machine-to-machine (M2M)/"Internet of Things" connectivity protocol[7]. Based on our experimentation, we have followed the second setup with M2M connectivity protocol. Such setup resulted in virtually no latency and came with the benefit of preserving the full precision of the data.

We have assembled the custom wearable unit into a 3D printed enclosure and attached it to a leather wristband, which conveniently packs the battery required to power up the device. The 9DOF sensor is attached to the user's

---

[4]https://www.adafruit.com/product/659
[5]https://www.adafruit.com/product/2020
[6]https://www.adafruit.com/product/2487
[7]http://mqtt.org
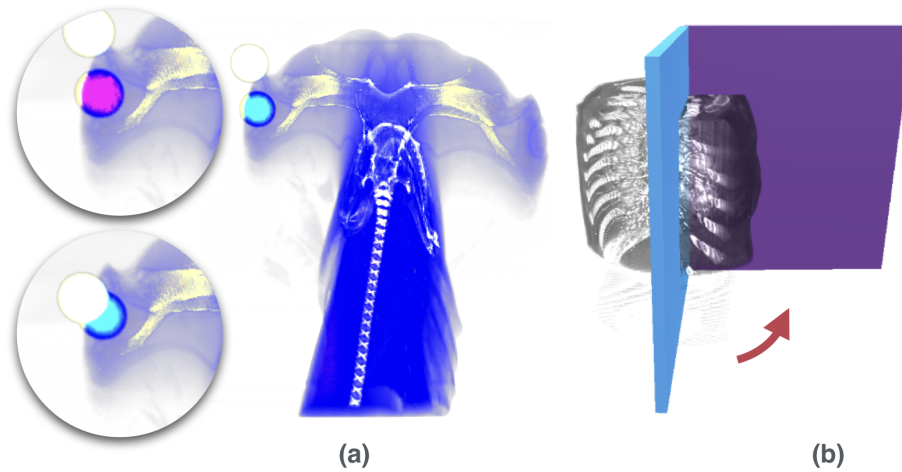
Figure 18: User study tasks. (a) 3D volumetric picking task: user selection highlighted with white volumetric sphere, successful selection of target - cyan, failed attempt - magenta. (b) 1 DOF rotation task of 3D box alignment.

palm, thus providing a precise orientation of the palm and not just the wrist (as one may capture with off-the-shelf wearable units).

6.6 User Study

In this study we evaluate the building blocks of the introduced natural gestures. First, we evaluate 3D volumetric picking with three types of visual feedback. Secondly, we assess feasibility of the two types of rotation. Our study had 12 participants between ages 23 and 39. Participation in the study was voluntary and consisted primarily of graduate and undergraduate students at the university. Despite the fact that one of the participants was left handed, they were able to perform tasks with their right hand. Each user was asked regarding their scientific visualization experience, ranking it on a scale from 1 to 5, five being the most satisfactory. Participants were also asked about their experience with VR/AR headsets, and particularly their use of Microsoft HoloLens. Half of the participants expressed visualization experience greater than rank 4. Only 7.7% did not have any previous VR/AR experience, and 69.2% have tried HoloLens HMD before at least once.

6.6.1 Tasks

Picking a point/location in 3D volume. The first task is a simple picking task, where the user is asked to pick a location on the "virtual" hologram of a volume. During the training session, the user was shown the target location: left eye of a smalleye hammerhead shark[8]. Unlike the usual picking task, the location was purposely selected to be fairly large (Figure 18 (a)), which might seem questionable to the reader. In our opinion, selection with a full hand gesture in mid-air may encounter a problem similar to that of a fat finger on a touchscreen device. To compensate for that, the target was chosen to be a large subvolume.

At the start of the task, the user was seated in the chair. Once the task started, the volumetric rendering appeared in front of the user about 4 feet away. To perform picking the user had to stand up, walk up to the volumetric rendering, and explore the dataset with one of the available techniques (ISO highlight, candle, 3D cursor, or no highlight) with their hand in the "ready" position. Once the location was found, the user could perform a selection gesture to pick the target. After the selection gesture was performed, the location of the pick would be highlighted in white, while the target location would light up with cyan or magenta highlight based on the success or fail of the task completion (Figure 18 (a)).

Rotation of the dataset. In this second task, the user was asked to perform a rotation task around the vertical (Y) axis (1 DOF). The rotation could be carried out with either free hand connected to the gyroscope wristband or by "dragging" an arc ball. To indicate the start and end point of the rotation, two rectangles were placed orthogonally: blue indicating the start of the rotation and purple - final position after rotation (Figure 18 (b)).

6.6.2 Procedure and Measures

Each user underwent a brief training session on an example of one picking task and was trained to perform a rotation task. A consistent experimental procedure was used for all trials and all participants. In total, users completed six randomly ordered tasks, repeating each three times, resulting in a total of 18 trials.

In both rotation and picking tasks, independent variables were the types of interaction technique. Time and accuracy (success/failure and distance to

---

[8]http://morph3d.eu

target; difference angle at the final alignment of rectangles) were recorded for both tasks.

NASA-TLX questionnaire [70] was used to evaluate the subjective task load. In addition to an average questionnaire score, we were interested in the subjective evaluation of performance and users' efforts. The experiment was performed in one session of approximately 30 minutes.

6.6.3   Analysis

We performed a repeated measures analysis of variance (ANOVA) *to determine the effect of each test condition on task completion time and performance accuracy*. In such a setting, we have two within-subject factors: trial number/time (first, second, or third) and test condition (task). We choose a mixed design and do not average results from three trials per task as we want to observe the effect of trial number on user performance.

The transformation $log(x)$ or $1/x$ was applied to the data and normality has been verified with Shapiro-Wilk and Anderson-Darling tests at 0.05 significance. We performed a repeated measures ANOVA *to determine the effect of each test condition on users' subjective task load score, as well as performance and effort subscores*.

6.7   Results

6.7.1   Task Completion Time

Picking task. Mean time completion for each picking task was found to be approximately similar (Figure 19). ANOVA revealed no significant main effect of the test condition on time to complete picking task ($F_{(3,128)} = 1.081, p = 0.360$). However, there was a significant effect of trial number ($F_{(1,128)} = 6.067, p = 0.015$) and a trend for both test and trial number ($F_{(3,128)} = 2.267, p = 0.083$) on time to complete picking task. The mean time for task completion per trial for all test conditions is shown in Figure 21.

Rotation task. The mean time completion for each rotation task is shown in Figure 19. ANOVA has revealed no significant main effect of test condition on the time required to complete the rotation task ($F_{(1,64)} = 0.920, p = 0.341$) or on either time, or trial number ($F_{(1,64)} = 0.038, p = 0.846$). There was no significant main effect of trial number on the time of
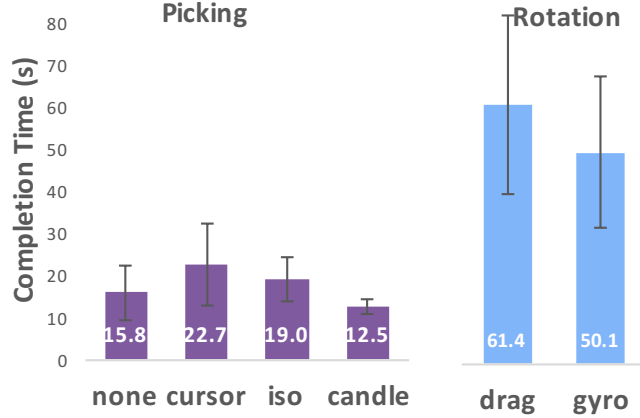
Figure 19: Means task completion time (seconds) for (left) picking task with four highlighting techniques and (right) rotation. Error bars indicate 95% confidence intervals.

completion of the rotation task ($F_{(1,64)} = 1.535, p = 0.220$).

6.7.2 Accuracy

Picking task. For the picking task, we evaluate both the success and failure of the picking, as well as the distance to the target location. Test condition did not have any statistically significant effect on participants' ability to *successfully select the target* (test: $F_{(3,128)} = 0.630, p = 0.596$; trial number: $F_{(1,128)} = 0.071, p = 0.790$; test and trial: $F_{(3,128)} = 0.412, p = 0.745$). Also there was no statistically significant effect on *distance to the target location* (test: $F_{(3,128)} = 0.126, p = 0.944$; trial number: $F_{(1,128)} = 0.018, p = 0.892$; test and trial: $F_{(3,128)} = 0.288, p = 0.834$).

Rotation task. Our analysis has revealed no significant main effect of test conditions on the accuracy of the rotation task (test: $F_{(1,64)} = 0.826, p = 0.367$; trial number: $F_{(1,64)} = 0.275, p = 0.602$; test and trial number: $F_{(1,64)} = 0.002, p = 0.968$).

6.7.3 Subjective Task Load

We have found no statistically significant effect of either rotation or picking task test conditions on subjective task load. Additionally, we have evalu-
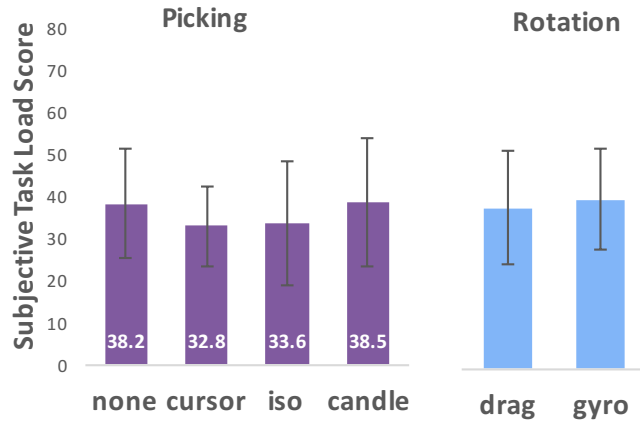
Figure 20: Subjective task load scores (NASA-TLX) for (left) picking and (right) rotation tasks. Error bars indicate 95% confidence intervals.

ated the effort and performance subscale scores, which have not revealed significant main effect ($p > 0.05$). The scores are shown in Figure 20.

6.8 Discussion

In our user study, we were attempting to verify two basic skills that build up to more complex gestures to mid-air volume interaction. For example, cutting plane positioning within the volumetric data requires adjusting the depth of the plane and its orientation. Thus, both hands of the users are involved in the task: one hand (i.e. left) picking the depth, while the other — equipped with a gyroscope sensor — is adjusting the orientation.

Based on the analysis of task completion time, we have observed that given time, subjects were adjusting to picking technique. As a relatively significant period of the trial, time was given to the subject standing up and walking to the spatial representation of the data, it might be unclear whether there is a significant difference in task completion time when using different visual feedback for the volumetric picking.

Our analyses have shown that there has been no significant difference in completion time for the rotation task. We find this result interesting, as a user had to perform significantly more "drag" operations to complete the first type of rotation task, while required only a single hand movement to complete the free hand rotation. During the study, we have observed that
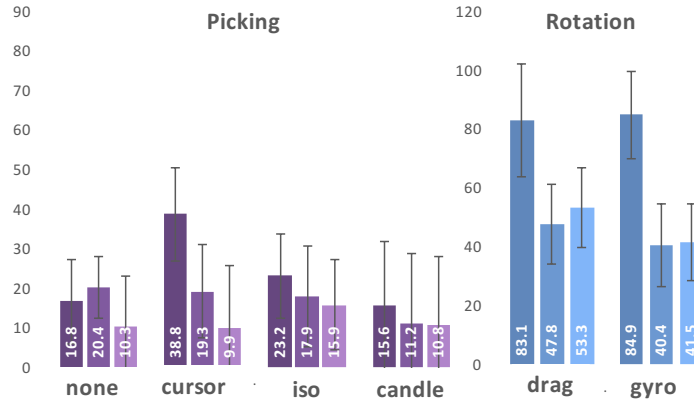
Figure 21: Mean task completion time for three repeated trials for (left) picking (right) rotation tasks.

the user took longer to adjust to the free hand rotation.

While we have found no statistically significant effect of the test conditions on the accuracy of the picking task, we have observed that among participants who have been more successful in the picking attempts, there was a trend for differences among the techniques. The participants who have shown to be more successful in the picking also had greater experience in visualization and had discussed their experience in gaming. In our opinion, a larger study is needed to confirm if such a correlation can be deduced. We believe that differences in techniques might have remained obscured to inexperienced participants.

6.9 Conclusions

In this work, we present a set of mid-air interactions with volumetric "virtual" holograms in mixed reality. We capture user input using Microsoft HoloLens and a custom built wearable wristband. We also design a detailed visual feedback that helps the user to locate their hand within the volumetric rendering even when it gets occluded by the data. Our study considers basic blocks of user interactions comprised of hand position and orientation. The experiments with volumetric picking have shown that careful use of visual feedback metaphors may result in overall better user experience without increasing subjective task load and without a negative effect on performance.

Visual Analytics for Medical Imaging: Focusing on the Spleen

The spleen, the largest organ in the lymphatic system, is often forgotten by laypeople but is of significant importance to clinicians. An increase in splenic size, splenomegaly, accompanies immune response in a wide range of abnormal conditions, including immunologic, hematopoietic, infectious, and storage diseases [13, 126, 109]. The evaluation, staging and response assessment of Hodgkin and Non-Hodgkin Lymphoma incorporate the spleen in the diagnostic process and it underlines a lack of consensus on standard splenic metrics [32, 132]. Additionally, spleens vary in shape and size across the patients, further complicating the task of finding a suitable evaluation criteria.

Diseases that affect the spleen pose particular challenges since they lack the clear-cut characterization between healthy and unhealthy subjects, which is often found in other organs (e.g. kidney). Specifically:

- Splenic maladies can manifest as variations in both shape and size. Thus, domain experts who are characterizing such abnormalities often find it difficult to do so based on a single metric.

- Traditionally, splenic disease is determined based on simple linear measurements (width/length) or volumetric estimates. These metrics often fail to characterize disease, since unhealthy spleens can fall within normal ranges. Consequently, it is important to consider other parameters, such as shape.

- Groups of patients can exhibit patterns of disease progression (e.g. changes in organ volume, shape, or other features) over time, which are of value in identifying the efficacy of treatment regimens and the accuracy of certain measurements for characterization.

With widespread adoption of multiple scanning modalities, such as computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound (US), splenic size can be estimated qualitative and quantitatively [39, 55, 118, 124, 152]. While the adoption of computer-aided techniques is progressing in various areas of radiology, majority of existing studies on splenic volumetry still resort to primitive single [15, 94] and multi-dimensional measurements

[86]. Among these measurements the most commonly referenced are craniocaudal splenic length, width, and thickness measured in the plane with maximal organs width [39, 86, 118, 152] or in axial organ's cross-section with maximal product of orthogonal length and width [124].

Multiple reports show correlation of single- and multi-dimensional indices to volume at a single point in time across a population [39, 55, 118, 124, 152]. However, to the best of our knowledge, none of these works have established how accurately these one- and multi-dimensional indices reflect change in splenic volume over time. Complex variations in splenic morphology [31, 76, 105] might not be portrayed accurately with ellipsoid-based formulation, and cause the error in estimate to propagate across the time domain.

Variation in organ shapes in measurements used is common among many medical sub-domains. Most medical research studies which focus on the analysis of time-varying imaging data are faced with this obstacle. The goal of these studies is to determine universally applicable metrics for use in clinical practice. However, the process of finding even an initial set of candidate metrics requires a comparison of volumetric imaging data across patients and studies over the time of disease progression.

Contributions of our work in this chapter can be summarized as follows:

- As a first step towards precise measurements for observation of changes in splenic volume and morphology, we evaluate a number of traditionally used splenic measurements. Unlike other methods that only use a single snapshot of data per patient, we evaluate uni- and multi-dimensional measurements based on several timepoints of data per subject.

- As a second step, we evaluate a novel heuristic used to determine the best possible measurement that can be performed given only standard formulations of "splenic index" can be used. Conformal welding heuristic allows us to select best possible plane of measurement and correlate all other estimations to this selection. Importantly, the conformal welding shape signature can be utilized as a biomarker of organ shape change.

- Finally, we explore additional imaging features that could provide a better characterization of the spleens. In the last section of this work we build a visual analytics tool for exploration of time-varying imaging-derived features. Our tool provides an overview of the study popula-

tion, allows to view registered organs, and preform similarity search based on several groups of imaging features. We propose several use cases for our tool and showcase an additional application for time-varying prostate data.

The following three sections of this dissertation, will address the above topics that resulted in the contributions to the fields of radiology, computer-aided detection in medical imaging, and visual analytics.

7 Comparison of Unidimensional and Multidimensional Splenic Measurements as a Predictor of Change in Volume Over Time

7.1 Introduction

Multiple studies on splenic volume evaluation use craniocaudal length, width, and thickness in correlation to splenic volume. While there is lack of consensus where the measurement should be taken, some studies have shown better correlation of length [15, 118] or width [39, 124] within their samples. In our study among one-dimensional measurements both splenic width and thickness show higher correlation to volume compared to organ's length. In our study we compare four different axial planes of measurement determined automatically, thus establishing a standard for uniform comparison and in attempt to eliminate errors arising from subjective evaluation. Estimation of splenic volume has been limited to regression-based methods combining products of one-dimensional diameters [15, 39, 118, 124, 152].

In this work, we analyze how change in single and combined splenic measurements predicts change in splenic volume: previously established single measurements, as well as product of two and three measurements. We evaluate differences in the magnitude of change of these indices and splenic volume and compare which plane of axial measurement is best for each of the index measurements, as well as what type of index measurement should be used.

7.2 Materials

7.2.1 Study Sample

The retrospective study was approved by Institutional Review Board (IRB), with a waiver of informed consent. The PACS system in the hospital was used to identify subjects with at least two supine abdominal CT scans performed at least 4 weeks apart. 34 adults (81 scans) without known splenic pathology and 38 adults (120 scans) with known pathology were examined and included in the study. Known pathology included: chronic lymphoid leukemia, low-grade B-cell lymphoma, marginal zone lymphoma, mantle cell lymphoma, or hairy cell leukemia. The study sample included 25 (35%) female subjects and 47 (65 %) male. Median age of the subject was 58. Total mean volume of the

organ at the first subjects visit was $436.68pm396.89$ mL (female - 309.65 mL, [73.46, 1100.87]; male - 504.25 mL, [87.23, 2254.81]). Mean time difference between the scans - $418.73pm451.49$ days, [2, 2430].

Slice thickness on the images varied (1.25 (1 scan), 2 mm (94 scans), 2.5 mm (17 scans), 5 mm (89 scans)) and was acquired with or without IV contrast. Data regarding phase of contrast and scanner type was not collected.

7.2.2   Spleen Segmentation

All images were loaded and manually segmented by drawing outlines along the organ boundary in a 3D software package (Alice software, Parexel Informatics; Waltham, MA). The region of interest (ROI) outlines were drawn by a trained assistant under supervision of radiologist with 20 years of experience. The assistant was trained by drawing ROIs on a set of test subjects from a prior separate dataset which was then subsequently verified by an expert. Training continued until error in volumes remained under 10% with maximum error of pm 30 ml. Once trained, verification of accuracy was assessed by comparing the volumes obtained by the assistant to the volumes obtained by MAB in the first 10 patients. If volumes error exceeded 10% or pm 30 ml, the assistant would be retrained on that case. This occurred in none of the 10 cases. Ground truth volume of the organ was determined by the software by multiplying the area of the ROI by the slice interval.

7.3   Methods

7.3.1   Measurements

The following splenic measurements were obtained: craniocaudal splenic length (L), width (W) and thickness (T). Width and thickness are orthogonal diameters measured on the same axial plane of the spleen. Width is defined as the longest diameter completely within the organ boundary, and thickness is defined as the longest orthogonal diameter completely contained within the organ boundary. Three axial planes of measurement were selected: axial plane of maximal width (MW), axial plane of maximal cross-sectional area (MA), and the axial plane through the midpoint in craniocaudal dimension (MP). For the purpose of this study in order to ensure consistent comparison of measurements and eliminate human error, all diameters and
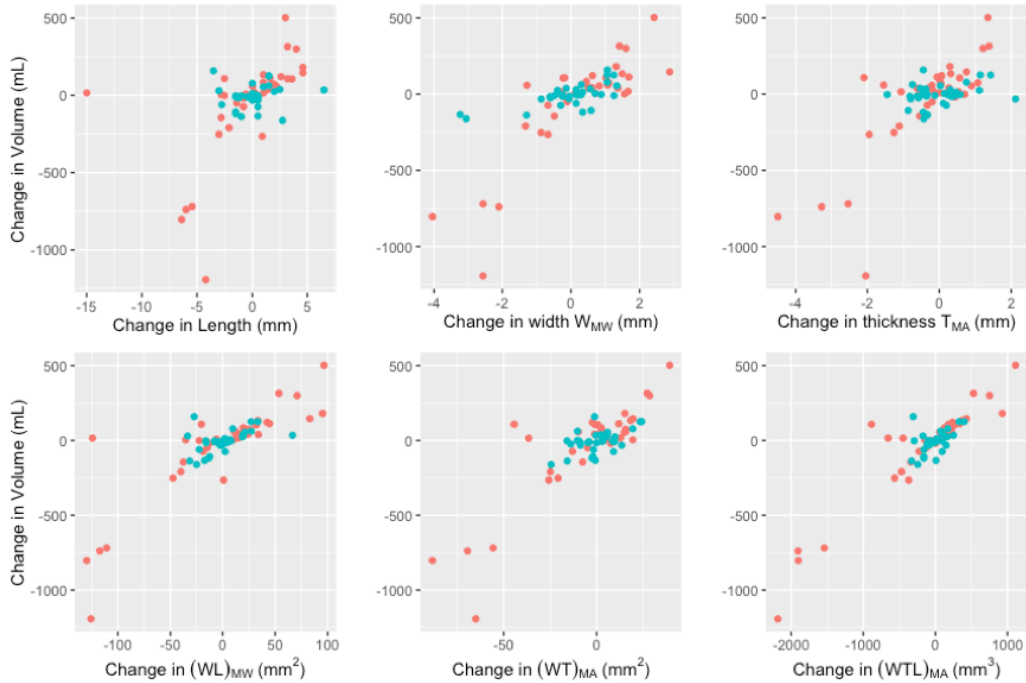
Figure 22: Correlation of the statistically significant change in 1D, 2D, and 3D metrics to change in splenic volume. Green color indicates healthy subjects, red - with a known pathology.

planes of measurement were obtained automatically by Alice software and a computer program written in Python (Python, version 2.7, Python Software Foundation). Slice numbers of organ's cross-sections were stored along with respective measurements.

7.3.2 Statistical Analysis

Statistical analysis of the data was performed in R language and environment for statistical computing and graphics [119] using standard statistical methods and the cocor package [43]. Basic statistic summary was obtained for all single splenic diameters (1D index), product of two ($W \cdot T$, 2D index), as well as product of all three diameters ($W \cdot T \cdot L$, 3D index). Correlation of all diameters and indices to the ground truth volume and its change was established by Pearson correlation. Our main goal is to show which of the metrics (change in 1D index, change in 2D index, or change in 3D index)

shows the highest correlation to change in volume. We formulate this as a follows:

H_0: There is no difference between correlation of change in metric A to change in volume and change in metric B to change in volume.

H_1: Correlations of change in metrics A and B to the change in volume are not equal.

Thus, we compare two overlapping correlations based on dependent groups (the same group of subjects) by employing two statistical tests. First, we show significant difference in magnitude of correlations using Pearson's correlation [114]. Secondly, we employ Zou's confidence intervals [158], known as superior to significance testing [43]. Unlike other methods that evaluate correlation of the volume of the organ to the metric, we evaluate correlations of the change in splenic volume and *change* in respective metrics. All results below are shown for the change (difference) between last and first visits for each subject.

7.4 Results

Figures 22 shows largest correlations of 1D, 2D, and 3D index metrics to the true volume of the organ. Mean, SD, and 95% CI are reported in Tables 7.4.1 and 7.4.1. Significant differences between two correlations and respective confidence intervals are indicated in bold in the Tables 7.4.2, 7.4.3, 7.4.4, 7.4.5.

7.4.1 Correlation of Splenic Volume and 1D, 2D and 3D Measures

In our study among one-dimensional measurements both splenic width (W_{MW}: $r = 0.842$; $R^2 = 0.709$) and thickness (T_{MA}: $r = 0.807$; $R^2 = 0.651$) show higher correlation to static volume compared to organ's craniocaudal length (L: $r = 0.771$; $R^2 = 0.594$). Among 2D measures, products of bi-orthogonal measurements that include mid-point cross-section WT_{MP} ($r = 0.912$; $R^2 = 0.832$) and $(W \cdot L)_{MP}$ ($r = 0.937$; $R^2 = 0.877$) show higher correlation to static volume, with maximal area cross section being second largest (Table 7.4.1). However, change in measures on MA and MW show higher correlation to change in volume ($(W \cdot T)_{MA}$: $r = 0.833$; $R^2 = 0.694$; $(W \cdot L)_{MW}$: $r = 0.814$; $R^2 = 0.662$) (Table 7.4.1). All of the 3D measures showed excellent correlation to the static volume of the organ, with mid-point cross-sectional plane ones being the highest ($(W \cdot T \cdot L)_{MP}$: $r = 0.971$; $R^2 = 0.942$) (Table 7.4.1).

Table 3: Summary statistics of splenic measurements for the first subject's visit. Range, mean, standard deviation, and 95% CI, correlation coefficient and coefficient of determination for the measurement/index to the volume of the organ.

Metric	Range	Mean	SD	95% CI	r^2	R^2
Volume	[73.46, 2254.809]	436.676	396.889	[345.001, 528.351]		
Length (L)	[4.5, 24.2]	11.024	3.958	[10.11, 11.939]	0.771	0.594
Width (W)						
W_{MA}	[6.17, 20.373]	11.226	2.789	[10.582, 11.87]	0.825	0.68
W_{MW}	[7.23, 20.782]	11.624	2.768	[10.985, 12.263]	0.842	0.709
W_{MP}	[4.808, 20.14]	11.099	2.774	[10.458, 11.74]	0.833	0.694
Thickness (T)						
T_{MA}	[2.966, 14.141]	7.121	2.217	[6.609, 7.633]	0.807	0.651
T_{MW}	[2.791, 14.254]	6.419	2.165	[5.919, 6.919]	0.778	0.606
T_{MP}	[2.825, 13.718]	6.698	2.151	[6.201, 7.195]	0.792	0.627
$W \cdot T$						
$(W \cdot T)_{MA}$	[21.245, 288.095]	84.636	46.068	[73.995, 95.277]	0.911	0.829
$(W \cdot T)_{MW}$	[20.179, 296.227]	78.956	45.109	[68.536, 89.375]	0.894	0.799
$(W \cdot T)_{MP}$	[21.117, 276.281]	78.719	44.456	[68.45, 88.988]	0.912	0.832
$L \cdot W$						
$(L \cdot W)_{MA}$	[41.648, 391.162]	130.783	74.005	[67.542, 101.73]	0.932	0.87

Table 3: Summary statistics of splenic measurements for the first subject's visit. Range, mean, standard deviation, and 95% CI, correlation coefficient and coefficient of determination for the measurement/index to the volume of the organ.

Metric	Range	Mean	SD	95% CI	r^2	R^2
$(L \cdot W)_{MW}$	[42.932, 399.014]	135.059	75.189	[117.691, 152.426]	0.94	0.884
$(L \cdot W)_{MP}$	[32.454, 386.688]	129.393	73.494	[112.417, 146.369]	0.937	0.877
$W \cdot T \cdot L$						
$(WTL)_{MA}$	[148.718, 5531.416]	1060.73	941.078	[843.356, 1278.104]	0.967	0.935
$(WTL)_{MW}$	[141.253, 5687.551]	989.206	919.085	[776.912, 1201.5]	0.953	0.909
$(WTL)_{MP}$	[137.26, 5304.586]	987.182	907.444	[777.577, 1196.787]	0.971	0.942
Estimates						
$Rezai(0.36 WTL + 28)$	[88.407, 2075.518]	421.302	345.712	[341.448, 501.156]	0.974	0.95
$Ell(0.524 \cdot WTL)$	[74.016, 2980.277]	518.344	481.6	[407.102, 629.586]	0.953	0.909

Table 4: Summary statistic of splenic measurements for the difference of last and first subject's visit. Range, mean, standard deviation, and 95% CI, correlation coefficient and coefficient of determination for the measurement/index to volume of the organ.

Metric	Range	Mean	SD	95% CI	r^2	R^2
Volume	[-1194.362, 502.279]	-27.291	239.727	[-82.664, 28.082]		
Length (L)	[-15, 6.5]	-0.086	3.002	[-0.779, 0.607]	0.53	0.281
Width (W)						
W_{MA}	[-3.771, 3.738]	0.155	1.341	[-0.154, 0.465]	0.602	0.362
W_{MW}	[-4.046, 2.882]	0.101	1.257	[-0.189, 0.391]	0.753	0.567
W_{MP}	[-4.14, 2.867]	-0.007	1.335	[-0.315, 0.301]	0.597	0.357
Thickness (T)						
T_{MA}	[-4.491, 2.118]	-0.195	1.078	[-0.444, 0.054]	0.718	0.516
T_{MW}	[-4.748, 3.52]	-0.256	1.388	[-0.576, 0.065]	0.442	0.196
T_{MP}	[-3.949, 4.117]	-0.175	1.271	[-0.469, 0.118]	0.456	0.208
$W \cdot L$						
$(W \cdot L)_{MA}$	[-125.499, 91.41]	-0.402	44.09	[-10.586, 9.783]	0.764	0.583
$(W \cdot L)_{MW}$	[-129.924, 96.663]	-0.471	44.619	[-10.777, 9.835]	0.814	0.662
$(W \cdot L)_{MP}$	[-129.579, 94.319]	-2.434	43.429	[-12.466, 7.597]	0.775	0.6
$W \cdot T$						
$(W \cdot T)_{MA}$	[-88.445, 39.214]	-2.544	22.062	[-7.64, 2.552]	0.833	0.694

Table 4: Summary statistic of splenic measurements for the difference of last and first subject's visit. Range, mean, standard deviation, and 95% CI, correlation coefficient and coefficient of determination for the measurement/index to volume of the organ.

Metric	Range	Mean	SD	95% CI	r^2	R^2
$(W \cdot T)_{MW}$	[-93.455, 42.917]	-3.59	23.544	[-9.028, 1.849]	0.7	0.49
$(W \cdot T)_{MP}$	[-84.106, 36.185]	-2.853	19.438	[-7.343, 1.637]	0.789	0.622
$W \cdot T \cdot L$						
$(WTL)_{MA}$	[-2185.922, 1111.098]	-77.897	544.692	[-203.712, 47.919]	0.911	0.831
$(WTL)_{MW}$	[-2369.325, 1131.317]	-80.464	519.617	[-200.487, 39.559]	0.874	0.764
$(WTL)_{MP}$	[-1972.807, 1052.037]	-72.02	487.353	[-184.591, 40.551]	0.901	0.812
Estimates						
$Rezai \cdot (0.36 WTL \cdot 28)$	[-852.957, 407.274]	-28.967	187.062	[-72.175, 14.241]	0.874	0.764
$Ell(0.524 \cdot WTL)$	[-1033.751, 551.267]	-37.738	255.373	[-96.725, 21.249]	0.901	0.812

7.4.2 Change in Splenic Volume and 1D Predictors

We have found a significant difference between correlations of change in width of the organ measured on MW versus MA planes ($p = 0.007$, 95% CI (W_{MA}, W_{MW}) [-0.282, -0.049]; $p = 0.014$, 95% CI (W_{MP}, W_{MW}) [-0.299, -0.037]) and width on MW plane versus thickness on MP plane ($p = 0.003$, 95% CI (W_{MW}, T_{MP}) [-0.506, -0.108]). This indicates that change in width of the spleen measured on the axial plane with maximal width (WMW) of the organ is a better predictor of change in volume.

Similarly, there are significant differences in correlation of change in volume and change in thickness of the organ on different planes of measurement ($p = 0.003$, 95% CI (T_{MA}, T_{MP}) [0.094, 0.452]; $p = 0.002$, 95% CI (T_{MA}, T_{MW}) [0.108, 0.466]). Thus, there is an indication that change in thickness of the organ on the axial plane with maximal cross-sectional area (T_{MA}) is more predictive of change in volume than thicknesses on other planes of measurement.

Finally, there is a significant difference between correlation of change in volume and splenic craniocaudal length versus width on MW plane ($p = 0.013$, 95% CI (L, W_{MW}) [-0.415, -0.052]), as well as thickness on MA plane ($p = 0.040$, 95% CI (L, T_{MA}) [-0.382, -0.01]). Thus, there is an indication that change in width (W_{MW}) or thickness (T_{MA}) are better predictors of change in volume than length (Table 7.4.2).

Table 5: Significance of the difference between correlations of 1D measurements (change in length (L), thickness (T), and width (W) measured on three axial planes (MA, MW, MP) to volume of the organ.

		W			T		
		MA	MP	MW	MA	MP	MW
L	MA	0.463	0.507	0.013	0.040	0.526	0.463
		[-0.125, 0.274]	[-0.137, 0.276]	[0.052, 0.415]	[0.01, 0.382]	[-0.308, 0.159]	[0.052, 0.415]
W	MA	-	0.941	0.007	-	0.149	0.113
		-	[-0.125, 0.135]	[-0.282, -0.049]	-	[-0.054, 0.354]	[-0.04, 0.367]
	MP	-	-	0.014	-	-	0.124
		-	-	[-0.299, -0.037]	-	-	[-0.045, 0.363]
T	MA	-	0.116	0.586	-	0.003	0.002
		-	[-0.03, 0.286]	[-0.171, 0.096]	-	[0.094, 0.452]	[0.108, 0.466]
	MP	-	-	0.003	-	-	0.891
		-	-	[-0.506, -0.108]	-	-	[-0.189, 0.216]

7.4.3 Change in Splenic Volume and 2D Predictors

There was a significant difference in correlations between change in product of width and thickness on MA versus MW planes of measurement ($p = 0.010$, 95% CI $((W \cdot T)_{MA}, (W \cdot T)_{MP})$ [0.04, 0.254]), that indicated change in $(W \cdot T)_{MA}$ to be a better predictor of change in volume.

There was a significant difference in correlations between change in volume and change in product of width and length on all three planes $p = 0.003$, 95% CI $((W \cdot L)_{MA}, (W \cdot L)_{MP})$ [-0.109, -0.019]; $p = 0.030$, 95% CI $((W \cdot L)_{MP}, (W \cdot L)_{MW})$ [-0.096, -0.002]). Confidence intervals indicate that change in $(W \cdot L)_{MW}$ to be a better predictor than respective measurement products on MA and MP planes of measurement (Table 7.4.3).

Table 6: Significance of the difference between correlations of 2D index metrics (product of change in length and width $(W \cdot L)$, product of change in thickness and length $(W \cdot T)$ measured on three axial planes (MA, MW, MP) to the change in the true volume of the organ.

	S		WL	
	MP	MW	MP	MW
MA	0.157	0.010	0.534	0.003
	[-0.018, 0.12]	[0.04, 0.254]	[-0.064, 0.036]	[-0.109, -0.019]
MP	-	0.068	-	0.030
	-	[-0.004, 0.202]	-	[-0.096, -0.002]

74

7.4.4 Change in Splenic Volume and 3D Predictors

There was no significant difference in correlation between change in volume and change in the product of any of the three sets of 3 orthogonal measurements. However, there was a trend between change in volume and change in the product measured from the MA and MW planes (Table 7.4.4).

Table 7: Significance of the difference between correlations of 3D index metrics (product of change in width, thickness, and length ($W \cdot T \cdot L$) measured on three axial planes (MA, MW, MP) to the change in the true volume of the organ

-	Volume	
-	MP	MW
MA	0.358	0.074
	[-0.016, 0.042]	[-0.002, 0.091]
MP	-	0.162
	-	[-0.012, 0.077]

7.4.5 2D and 3D Predictors of Change in Segmented Splenic Volume

There was a significant difference in correlation between change in volume and all of the 3D indices versus any of the evaluated 2D indices (Table 7.4.5) which indicates that all of the 3D indices are a significantly better predictor of change in volume.

Table 8: Significance of the difference between correlations of 3D index metrics (product of width, thickness, and length ($W \cdot T \cdot L$)) and 2D indices ($W \cdot T$, $W \cdot L$) measured on three axial planes (MA, MW, MP) to volume of the organ.

		S	WL
V	MA	0.003	0.000
		[0.035, 0.148]	[0.097, 0.241]
	MP	0.001	0.000
		[0.058, 0.199]	[0.077, 0.214]
	MW	0.000	0.057
		[0.105, 0.286]	[0.001 0.137]

7.5 Conclusions

In this study we evaluate only basic indices based on simple measurements of width, length, and thickness of the organ. Additionally, other measurements (for example, using different landmarks) can be used to evaluate volumetric measurement. For the purpose of this study, we focus only on index-based measurements.

We perform our evaluation on a limited sample. Thus, the results can benefit from further evaluation on larger samples. Our study evaluates two methods based on 3D splenic index, it can be further extended to other methods that rely on single dimensional measurement to predict change in volume.

8 Maximal Area and Conformal Welding Heuristics for Optimal Slice Selection in Splenic Volume Estimation

In this work, we propose two heuristics for the determination of the plane of axial measurement, which drives volumetric estimation calculations. The first is termed the "maximal area" heuristic, which corresponds to the plane, which defines the maximal area among all slices enclosed in the splenic volume. The second is termed the "conformal welding" heuristic, which is a novel heuristic that utilizes the axial plane that bifurcates the splenic mesh in a way that maximizes the symmetry between the two halves.

We perform evaluation of these heuristics on a set of meshes extracted from manually segmented volumes of the organ, and provide guidance with respect to selection of the axial plane of measurement that can be used in Rezai and prolate ellipsoid methods without the need of full segmentation.

Maximal Area Heuristic. Previous splenic volume estimation methods find the maximal axial width of the organ, and further find the longest perpendicular thickness within the same axial plane. As shown in Figure 23, some errors of the observers result in actually choosing slices with greater areas. In the work on visual perception, Wagner [144] reports that on average length of the object is overestimated, whereas area is underestimated. However, we were not able to find sources stating that this is the case for objects that are "wedge-shaped" or "ellipsoid-like", and how task of comparing such shapes affects observer's performance. We evaluate maximal area-based heuristic volumetric estimation and report that it actually results in better volumetric estimation.

Conformal Welding Heuristic. In our method, we propose to compute the global correlations between two surface mesh contours of the top and bottom part, of the spleen separated by the axial plane of measurement. The shape signature demonstrates the global geometric features encoded into these two parts, and can be utilized as a biomarker for the organ shape change. Such method has been previously proposed for analysis of CHD progression and pathology in the brain [116]. The shape signature is computed based on the conformal structure of the organ surface [62, 146] and can be accurately computed using the surface Ricci flow method [82, 145]. Sharon and Mumford [131] have used conformal mappings for modeling the 2D shape
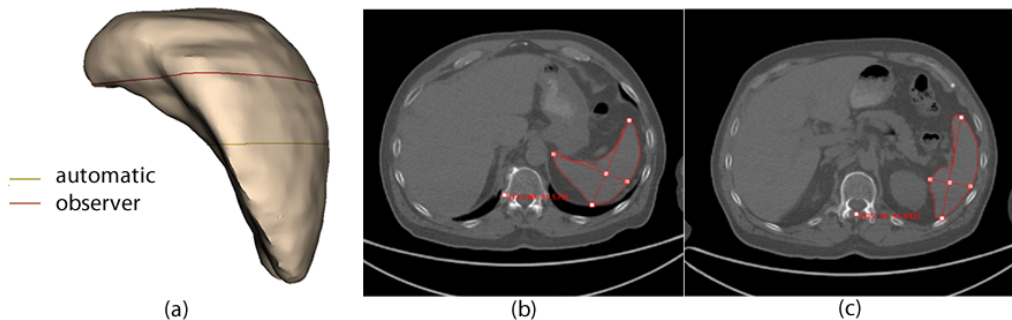
Figure 23: Comparison of observer (red) and automatic (yellow) measurement using maximal axial length as heuristic for plane selection. (a) Rendering of the manually segmented mesh. (b) Measurement determined by the observer, and (c) measured automatically. Volume calculated is 632.6 and 462.6 mL for the observer measurement, 368.6 and 281.2 mL for automatic, prolate ellipsoid and Rezai methods, respectively. True volume of the organ 381.3 mL.

space. We apply the 3D generalization of Sharon and Mumford's 2D shape space to 3D surface of the spleen. This method builds a Teichmüller space for 3D shapes by using conformal mappings. In this Teichmüller space, every 3D contour (a simple closed curve) is represented by a point in the space; each point denotes a unique equivalence class of diffeomorphisms up to Möbius transformation. For a 3D surface, the diffeomorphisms of all the contours form a global shape representation of the surface. By using this signature, the similarities of 3D shapes can be quantitatively analyzed, therefore, the classification and recognition of 3D objects can be performed from their observed contours. While conformal welding heuristic requires an input of a mesh of a segmented organ, it allows us to select best possible plane of measurement and correlate all other estimations to this selection.

We evaluate these heuristics and compare them against the well-established maximal length approach with two estimation methods, Rezai and prolate ellipsoid. Our evaluation is based on two sample groups: subjects with no disease that affects splenic size, as well as the group with leukemia and lymphoma that were undergoing treatment expected to result in decrease of the splenic volume. We expose and discuss the results of our statistical analysis, which demonstrates that our heuristics are superior to prior methods in certain situations. Finally, we conclude with guidance on how these heuristic

determinations can be reached without requiring ground-truth organ segmentation.

8.1 Theoretic Background

In this section, we briefly introduce the theoretical foundations necessary for the computation of the conformal welding heuristic. For more details, we refer readers to the classical books [54, 71].

8.1.1 Riemann Mapping

Conformal mapping between two surfaces preserves angles. Suppose $(S_1, \mathbf{g}_1)$ and $(S_2, \mathbf{g}_2)$ are two surfaces embedded in $\mathbb{R}^3$, $\mathbf{g}_1$ and $\mathbf{g}_2$ are the induced Euclidean metrics. A mapping $\phi : S_1 \to S_2$ is called *conformal*, if the pull back metric of $\mathbf{g}_2$ induced by ϕ on S_1 differs from $\mathbf{g}_1$ by a positive scalar function: $\phi^* \mathbf{g}_2 = e^{2\lambda} \mathbf{g}_1$, where $\lambda : S_1 \to \mathbb{R}$ is a scalar function, called the *conformal factor*. For example, all the conformal automorphisms of the unit disk form the *Möbius transformation group* of the disk, each mapping is given by $z \to e^{i\theta} \frac{z-z_0}{1-\bar{z}_0 z}$. A genus zero surface with a single boundary is called a *topological disk*, which can be conformally mapped onto the planar unit disk by a *Riemann mapping*, all such mappings differ by a Möbius transformation.

The Riemann mapping can be computed by *Ricci flow*, which is the process of deforming Riemannian metric $\mathbf{g}$ proportional to the curvature, such that the curvature K evolves according to a heat diffusion process, eventually the curvature becomes constant everywhere. Suppose the metric $\mathbf{g} = (g_{ij})$ in local coordinate. Hamilton [69] introduced the Ricci flow as $\frac{dg_{ij}}{dt} = -K g_{ij}$. Surface Ricci flow conformally deforms the Riemannian metric, and converges to constant curvature metric [33].

8.1.2 Conformal Welding Shape Descriptor

Suppose γ is a closed curve on a genus zero closed surface S, segments S into two components $\{\Omega_0, \Omega_1\}$. The Riemann mapping $\phi_k : \Omega_k \to \mathbb{D}$ maps each segment Ω_k to the disk. Let $f := \phi_1 \circ \phi_0^{-1}|_{\mathbb{S}^1} : \mathbb{S}^1 \to \mathbb{S}^1$ be the diffeomorphism from the circle to itself. We called the term diffeomorphism f the *signature of γ*. The closed curve γ on a genus zero closed Riemannian surface S is determined by its signature, unique up to a conformal automorphism of the surface.

8.2 Algorithm

In this section, we describe the process of calculating the conformal welding shape signature used in our heuristic. This process is illustrated in Figure 24 in detail.

8.2.1 Riemann Mapping

The surface is represented as a triangular mesh $\Sigma(V, E, F)$, on which we apply discrete Ricci flow method [82].

Definition 1 (Discrete Conformal Factor). The discrete conformal factor function is defined on the vertex set $u : V \to \mathbb{R}$, such that for each edge $[v_i, v_j]$, the edge length $l_{ij} = e^{u_i} \beta_{ij} e^{u_j}$, where β_{ij} is the initial edge length.

The discrete Gaussian curvature on each vertex v_i is defined as angle deficit: $K_i = \begin{cases} 2\pi - \sum_{ij} \theta_i^{jk} & v_i \notin \partial\Sigma \\ \pi - \sum_{ij} \theta_i^{jk} & v_i \in \partial\Sigma \end{cases}$, where θ_i^{ij} is the corner angle at v_i in the face $[v_i, v_j, v_k]$. Then, the discrete Gauss-Bonnet theorem holds: $\sum_{v_i \in \Sigma} K_i = 2\pi\chi(\Sigma)$.

Definition 2 (Discrete Ricci Energy). The discrete Ricci energy is given by $E(\mathbf{u}) = \int^{\mathbf{u}} \sum_i K_i du_i$, where $\mathbf{u}$ is the vector of conformal factors $(u_1, u_2, \cdots, u_n)$.

The discrete Ricci energy is convex on the space $\sum_i u_i = 0$. The Hessian matrix of the Ricci energy is given by $\frac{\partial^2 E}{\partial u_i \partial u_j} = \frac{\partial K_i}{\partial u_j}$ which has explicit geometric interpretation. Suppose $[v_i, v_j]$ is an interior edge on Σ, which is adjacent to two faces $[v_i, v_j, v_k]$ and $[v_j, v_i, v_l]$, $\frac{\partial K_i}{\partial u_j} = -(\cot\theta_k^{ij} + \cot\theta_l^{ji})$, if $[v_i, v_j]$ is a boundary edge adjacent to $[v_i, v_j, v_k]$, then $\frac{\partial K_i}{\partial u_j} = -\cot\theta_k^{ij}$. Furthermore, $\frac{\partial K_i}{\partial u_i} = -\sum_{[v_i, v_j] \in \Sigma} \frac{\partial K_i}{\partial u_j}$.

Definition 3 (Delaunay Triangulation). A closed triangle mesh is Delaunay, if for each edge $[v_i, v_j]$ adjacent to faces $[v_i, v_j, v_k]$ and $[v_j, v_i, v_l]$, $\theta_k^{ij} + \theta_l^{ji} \leq \pi$.

Given target curvature $\bar{K} : V \to \mathbb{R}$, satisfying the Gauss-Bonnet theorem, the discrete Ricci flow is given by $\frac{du_i}{dt} = \bar{K}_i - K_i$, which is the gradient flow of the following energy $F(\mathbf{u}) = \int^{\mathbf{u}} \sum_i (\bar{K}_i - K_i) du_i$. This energy is strictly concave in the space $\sum_i u_i = 0$, and can be optimized directly using Newton's method, $\nabla F = (\bar{K}_1 - K_1, \bar{K}_2 - K_2, \cdots, \bar{K}_n - K_n)^T$. The Hessian matrix is the

negative of that of $E(\mathbf{u})$. Furthermore, during the flow, we preserve the mesh to be Delaunay all the time. This guarantees the existence of the solution.

Discrete Riemann Mapping. Given a topological disk, by puncturing a small hole in the center, we convert it to a topological annulus. Then, we set the target curvatures to be zeros everywhere, and run Yamabe flow to get a flat metric. By applying the complex exponential map, we can map the annulus with the flat metric to a planar annulus, and then fill the center hole. This gives the Riemann mapping from the original topological disk to the unit disk.

8.2.2 Computing Shape Descriptor

After the computation of the conformal mapping, each connected components is mapped to the disk. We use a Möbius transformation to map the mass center of each disk to the origin. Then we compute the signature directly. Suppose $\gamma \subset S_1 \cap S_2$, where S_1 and S_2 are two segments. $\varphi_1 : S_1 \to \mathbb{D}$ and $\varphi_2 : S_2 \to \mathbb{D}$ are two conformal mappings, then, the signature of γ is given by $f_\gamma := \varphi_2|_\gamma \circ \varphi_1^{-1}|_\gamma$.

8.3 Methods

The data for this study was collected retrospectively from patient charts and PACS. The study includes the data of the adult patients (over 18 years of age), with 2 to 4 additional abdominal CT scans over a three year period. For this work we have selected the datasets with the highest number of patient follow-ups available, which resulted in 12 datasets of normal subjects (31 scans total), and 11 of abnormal (31 scans). Manual volumetric segmentation of the splenic contour on each axial slice was performed using Alice (Paraxel Informatics, Waltham, MA) in a blinded fashion under supervision of a fellowship-trained attending radiologist. Ground truth computation of the segmented volume of interest was performed by the software. Overall the splenic volume of subjects in healthy group ranged from 73 to 705 mL (mean 306, $\pm$ 173 mL). Subjects in sick group demonstrated ground truth splenic volume ranging from 108 to 2255 mL (mean 578, $\pm$ 462 mL). CT slice thickness varied from 2.5 to 5 mm in the healthy group, and from 1.25 to 5 mm in a sick group with majority being 2 mm. From the ground truth segmentation, we also generated 3D meshes which were used to drive the conformal welding heuristic. These mesh had an average granularity of approximately
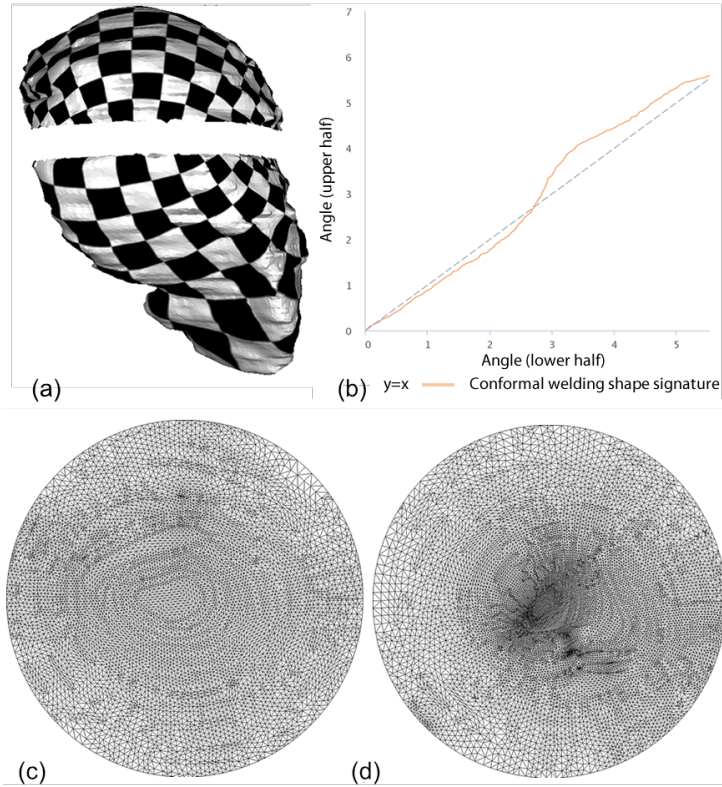
Figure 24: Conformal welding shape signature pipeline: (a) Surface cut by the plane with selected conformal welding-based heuristic, (c-d) top and bottom parts of the mesh conformally mapped to circle, (b) resulting signature of the shape.

10000 - 16000 vertices. Overall, our evaluation was structured as follows. For each splenic dataset, we first determined the maximal axial length of the organ. Following that, we determined the axial plane that resulted in a maximization of organ cross-section. Finally, we determined the axial plane through the conformal welding heuristic. For the purpose of determining the latter, we proceed with the following steps:

- We cut the surface mesh at every slice starting from the maximum area slice up to ± 5 slices.

- For each cut, we find the signature of global correspondence between upper and lower surfaces of the mesh.

82
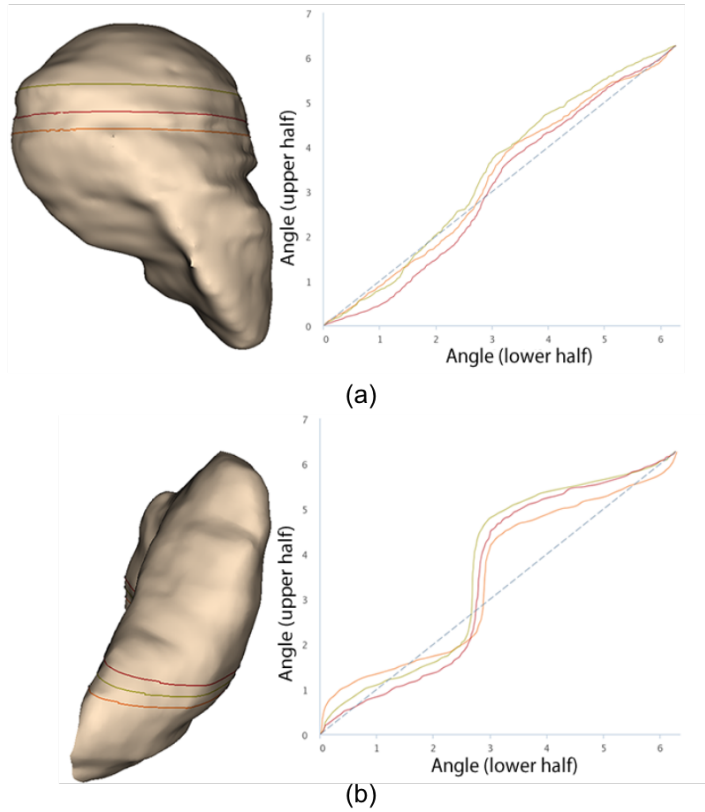
Figure 25: Shape signatures resulting from different cut plane selections: maximal length of the organ's cross section (yellow), maximal area (red), and optimal plane based on conformal welding heuristic (orange). Two conditions are illustrated: (a) Surface mesh of a spleen with known disease; (b) surface mesh of a healthy spleen and respective shape descriptors for mesh cut with different heuristics.

- Further, we find l^2-norm of each signature to the perfect ellipsoid signature (Figure 25).

- For our conformal welding shape signature heuristic, we select the axial slice that with the minimum l^2-norm that is only up to 3 cross-sections (6 - 15 mm) away from the maximal area cross-section. We do not examine the signatures of the cuts that are close to top or bottom slices since at those points the cross-section of the organ becomes naturally

Table 9: Lin's concordance correlation for compared heuristics.

Method	Normal	Abnormal
Maximal diameter product heuristic		
Rezai	ccc = 0.953 [0.872,0.983]	ccc = 0.982 [0.951,0.993]
	mean = -17.96 ± 39.6	mean = -19.44 ± 52.03
Pr. Ellipsoid	ccc = 0.896 [0.742,0.96]	ccc = 0.884 [0.705,0.957]
	mean = 72.38 ± 64.14	mean = 194.36 ±191.88
Maximal width-based heuristic		
Rezai	ccc = 0.931 [0.821, 0.974]	ccc = 0.977 [0.940, 0.991]
	mean = -35.48 ± 42.2	mean = -54.75 ± 40.6
Pr. Ellipsoid	ccc = 0.913 [0.778, 0.967]	ccc = 0.907 [0.748, 0.968]
	mean = 46.88 ± 64.0	mean = 142.96 ± 193.4
Maximal area-based heuristic		
Rezai	**ccc = 0.953 [0.877,0.982]**	**ccc = 0.980 [0.949,0.993]**
	mean = -25.1 ± 31.36	**mean = -29.12 ± 34.99**
Pr. Ellipsoid	ccc = 0.910 [0.778,0.965]	ccc = 0.894 [0.725,0.961]
	mean = 61.99 ± 52.6	mean = 180.26 ± 183.67
Conformal welding-based heuristic		
Rezai	ccc = 0.854 [0.665,0.939]	ccc = 0.974 [0.933,0.99]
	mean = -65.79 ± 45.54	mean = -48.92 ± 58.31
Pr. Ellipsoid	**ccc = 0.933 [0.833,0.973]**	**ccc = 0.917 [0.778,0.97]**
	mean = 2.76 ± 37.84	**mean = 151.45 ± 160.24**

close to the ellipse resulting in the signatures with the smallest l^2-norms, yet not representing global shape of the organ.

We compare four heuristics for selecting the plane of measurement: maximal width of the organ, maximal product of orthogonal diameters, maximal area, and our novel conformal welding based heuristic. For each one of these heuristics, we estimated the splenic volume utilizing the Rezai and prolate ellipsoid formulations. Overall, we end up with 4×2 (heuristic × estimation method) volumetric measurements per subject group.

8.4 Results

As is standard in volumetric estimation against the ground truth volume in biomedical studies, we should compare concordance correlation coefficient (CCC) of the established metric with a ground truth volume, which is traditionally done with Lin's CCC [98]. However, since our data consists of repeated measurements for every subject over time, we utilize a modified version of concordance estimation by Carrasco *et al.* [29]. In our evaluation we utilize their R package implementation, in particular the estimation of CCC from a variance components and fixed effects that are estimated using the restricted maximum likelihood (REML) approach.

The results of this analysis, broken down between the normal and abnormal (or sick) dataset samples, can be seen in Table 9 for the maximal orthogonal product of diameters, maximal width, maximal area and conformal welding heuristics respectively. Our analysis has shown that in all of the cases, prolate ellipsoid method tends to overestimate organ volume, while Rezai underestimate. Substantial correlation is observed for Rezai method for both heuristic based on the maximal product of orthogonal diameters, and for maximal area-based heuristic. With respect to prolate ellipsoid method, our novel conformal welding-based heuristic shows substantial correlation and establishes the best possible estimation.

Our evaluations show that existing formulations for the volume estimation do not necessarily show previously reported quality of fit to the new population. Our observations suggest that it is the high shape variability of the organ and deviation from assumed ellipsoid-like shape that results in high deviation from the ground truth volume, in particular for prolate ellipsoid method. Overall, estimations were more concordant with true volumes using the Rezai method. While correlation between both methods and true volume was not weak, occasionally changes in morphology over time significantly altered the true orthogonal width, resulting in falsely large calculated change in volume.

First, it is important to note that while maximal diameter product-based heuristic shows substantial correlation, it is highly unlikely for the observer to select the cross section of the organ with greatest product of two orthogonal diameters. Our suggested maximal area and conformal welding-based heuristics, report moderate results for Rezai and prolate ellipsoid volume estimations. The maximal area-based heuristic outperforms traditional width-based heuristic, also known as "calipers" length, for Rezai method for both sick and healthy organs. This can be explained with more frequently occurring wedge shapes of organ cross-sections, such that maximal axial width

heuristic does not capture the thickest part of the organ, and might erroneously select concave wedge shape.

In addition, our conformal welding-based heuristic achieves moderate concordance for the ellipsoid method. It outperforms the established maximal width heuristic for the prolate ellipsoid method. In particular, this is due to selection of the organ cross-section whose signature is closest to the perfect ellipsoid assumed by the formula, and which represents the global correspondence between the upper and lower parts. Both traditional maximal width and maximal area heuristics can be used for abnormal cases due to "swollen" and ellipsoid-like shapes. While conformal-welding based heuristic, similar to maximal diameter product heuristic, is not detectable by the observer without a fully segmented organ, it does provide us the best possible plane for closest estimation by prolate ellipsoid method due to the assumed ellipsoid shape.

8.5   Conclusions

We have performed analysis of splenic volume estimation techniques based on two groups of subjects: subjects with no known disease to affect the splenic size and subjects with leukemia and lymphoma, diseases that frequently result in splenomegaly. Our results suggest that traditional methods tend to over- or underestimate the volume, which results in inaccurate analysis of volumetric changes over time. We have analyzed the shape of the spleen and proposed two heuristics that outperform traditional width-based measurements. In case one-dimensional measurements remain a standard, our suggestion is to further re-evaluated existing methods and categorize metrics and heuristics based on the shape of the organ.

9 AnaFe: Visual Analytics of Image-derived Temporal Features

9.1 Introduction

Diseases that affect spleen pose particular challenges based on the interplay of multi-dimensional, time-varying data, and data across populations. Visual analytics approach has the potential to support medical analysis of the problem domain of splenic disease. Hence, In this work, we propose a visual analytics tool, *AnaFe*, for the analysis of changes in the organ based on imaging-derived features. The current quality of image acquisition and analysis allows for the extraction of reproducible quantitative imaging features. We take advantage of such features derived from the field of radiomics [113] to enable an open-ended similarity search and trend exploration in splenic imaging data. Such features have several benefits: simplicity in derivation, mapping to visualization, and most importantly reproducibility and versatility across imaging types and domains.

AnaFe supports a similarity search and comparison workflow based on a set of robust radiomics features. Our proposed workflow enables a user to construct custom similarity-based queries through interaction with these features. Through a number of visual queries and a rapid visual feedback, the user is able to compare several time-variant imaging sets and their correlating features in a single overview. Driven by demands of medical research analysis, *AnaFe* combines a set of linked visualization views. Thus, researchers can concentrate on exploring and characterizing changes in data and corresponding features.

The utility of our tool is demonstrated through two case studies conducted by our collaborating radiologist on a set of 189 datasets. Based on the target application domain we outline a set of visualization requirements for our visual analytics system and describe the resulting design decisions. Then, we illustrate the implementation of our tool and focus on the integration of various feature types and their mapping to visualization views. Next, we describe two application use-case scenarios. We conclude by offering directions for future work.
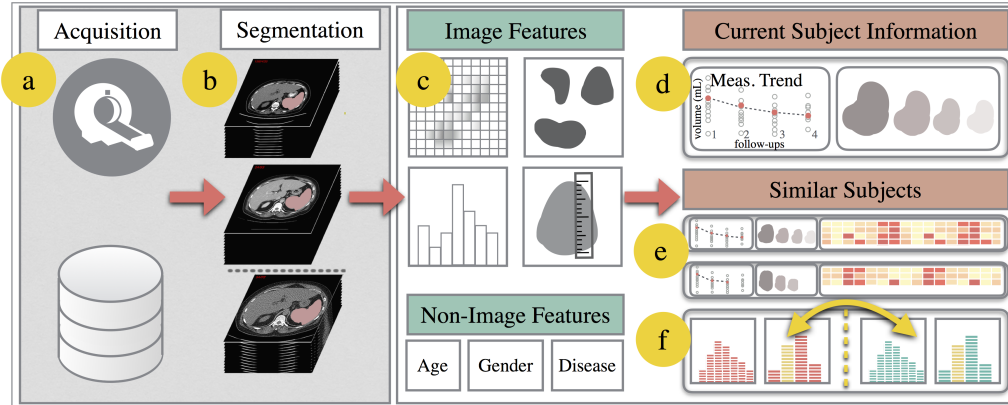
Figure 26: An overview of the steps required for temporal feature analysis by *AnaFe*. The data is sourced from follow-up CT scans (a) that have been segmented (b). In this work, we first derive robust image features from the field of radiomics (c): texture, shape, intensity, and measurements. These features, along with other study information (non-image features), are used by our web-based application *AnaFe*. After selection of the current subject (d), the application performs a search for subjects with similar progression. Search results filtering is performed by brushing a selection over the distribution charts (f) of average features and average change in progression (interchangeable views shown with a double yellow arrow.)

9.2   System Design and Implementation

AnaFe is a framework for analytical exploration of 3D organ data and associated medical-imaging derived features. We focus on parts (c)-(f) in Figure 26: deriving the imaging features and their visual analysis in combination with non-image features from the study data.

9.2.1   Input Data

Our application consumes several types of data that stems from the medical research study of spleen variation. Each study is comprised of multiple subjects observed over several visits.

Subject Information. Each study contains information with respect to gender, age, and disease status of the subject (sick or healthy).

Imaging-derived Features. Feature vectors for each study are derived

from the original DICOM data after applying segmentation masks. We describe these features in detail in Section 9.2.2. The effect of segmentation on quantifiable radiomics features has been previously evaluated [113]. For the purpose of unbiased evaluation we used features derived only from manually segmented data.

3D Surface Mesh Pre-processing. *AnaFe* requires a 3D mesh for each organ for rendering, which is obtained in two steps.

Mesh generation. From the manually segmented data we derive a 3D surface mesh of the organ via marching cubes [100]. The surface geometry is then post-processed with Taubin smoothing [136].

Mesh registration. For comparison of several organs' surfaces, we align the meshes using the Iterative Closest Point (ICP) algorithm[97]. Next, meshes are registered using non-rigid ICP to obtain correspondences between the vertices of consecutive timepoint meshes. This type of registration has not been previously used for the spleen. We have selected these methods based on the preliminary evaluation of their performance and ready availability.

Thus, any study of a group of subjects with a series of repeated measurements (dependent variables) over time (independent variable), that contains similar information can be analyzed by *AnaFe*.

9.2.2 Feature Types

In clinical practice determining the spleen size and response to treatment is often a qualitative rather than quantitative process. For the purpose of quantitative comparison of inter-subject similarity of organ change along with basic measurements, we integrate a set of imaging-derived features. With our application and other potential domains in mind, we group these features into the following four categories:

- *Measurements.* In practice, volume, craniocaudal length, width, and thickness serve as primary comparisons for spleens. Therefore, we define them as a separate category. For many other applications, measurements can be defined as descriptors of shape.

- *Shape.* Shape descriptors characterize spleen shapes, which include elliptical, triangular, wedge, and tetrahedron [126].

- *Intensity.* In CT scans, the spleen is often described as an organ of homogeneous density [126]. Thus, intensity is not extremely important

for this application. We include this feature group for the purpose of extensibility to other domains.

- *Texture.* Texture features are based primarily on a gray-level co-occurrence matrix (GLCM). Such features describe cluster prominence and can be used for tumor characterization. Similar to intensity, this feature is included for extensibility.

The detailed definitions of these features can be found in the related literature [113]. In addition, we provide information on the features used in our case studies in Table 9.2.2. These imaging-derived features, along with study information (gender, age, and disease status), comprise the full feature vector for each dataset.

Table 10: Summary of image-derived features and measurements used by *AnaFe*.

Feature	Definition
Measurements	
Volume	Total volume (V_{total}) of the organ determined by voxel count (I_{voxel}) and voxel volume (V_{voxel}): $$V_{total} = I_{voxel} \cdot V_{voxel} \qquad (3)$$
Craniocaudal length	Length of the organ from top first axial slice to the last multiplied by slice thickness.
Width	Maximal diameter of the organ on the axial slice (with maximal area)
Thickness	Diameter of the organ perpendicular to width
Shape	
Compactness	Reflects how close the shape of a 2D object to a circle: $$Compactness = \frac{2\sqrt{Area \cdot \pi}}{Perimeter} \qquad (4)$$

90

Convexity	Characterizes the convexity of a 2D object: $$Convexity = \frac{Perimeter_{convex\ hull}}{Perimeter} \quad (5)$$
Rectangu-larity	Reflects how close the shape of a 2D object to a rectangle : $$Rectangularity = \frac{Area}{Area_{mimnum\ bouding\ box}} \quad (6)$$
Ellipticity	Shows how much the shape of a 2D object differs from an ellipse: $$Ellipticity = \frac{Area}{Area_{fitted\ ellipse}} \quad (7)$$
Triangu-larity	Reflects how close the shape of a 2D object to a triangle: $$Triangularity = \frac{Area}{Area_{minimum\ bounding\ triangle}} \quad (8)$$
Surface-area-to-volume ratio	Characterize the complexity of a 3D object: $$A - V Ratio = \frac{\sum_i Area_{voxel_i}}{V_{total}} \quad (9)$$

Texture

Gray Level Co-occurrence Matrix (GLCM)P indicates frequencies of combinations of gray level co-occurrences

	Measures the local variations in the GLCM:
Contrast	$$Contrast = \sum_{i,j=0} P_{i,j}(i-j)^2 \qquad (10)$$

	Measures the closeness of the distribution of elements in the GLCM:
Homo-geneity	$$Homogeneity = \sum_{i,j=0} \frac{P_{i,j}}{1 + (i-j)^2} \qquad (11)$$

	Measures the variation of gray level pairs in the image:
Dissimilar-ity	$$Dissimilarity = \sum_{i,j=0} P_{i,j}|i-j| \qquad (12)$$

	Measures local homogeneity:
Energy	$$Energy = \sum_{i,j=0} P_{i,j}^2 \qquad (13)$$

	Measures the gray level linear dependencies in the image:
Correla-tion	$$Correlation = \sum_{i,j=0} P_{i,j}\frac{(i-\mu_i)(j-\mu_j)}{\sqrt{\sigma_i^2 \cdot \sigma_j^2}}, \qquad (14)$$ where $\mu_{i,j}$ and $\sigma_{i,j}$ are the mean and std of P_x and P_y

Intensity	
Maximum	Maximum intensity value
Minimum	Minimum intensity value

Mean & Standard Deviation	Mean and standard deviation of intensity values
Range	The range of intensity values
Kurtosis	Measures the sharpness of the histogram

9.2.3 Similarity Comparison

Comparing changes across subjects (**RQ2**) requires analyzing data from multiple timepoints. The number of visits varies from subject to subject based on the condition and progression of the disease. It was noted by our medical collaborators that the time difference between the two consecutive patient visits is not of essence for this comparison. Varying duration of treatment and the subject's initial conditions can affect the quantitative difference in the measurement/features. These compounding variables about the subject's original state were not available during the current study. Regardless of the velocity of change in the organ's condition, it must be detected. The focus of our study is visual identification and correlation of predictors of change among the large number of features (measurements and imaging-based) for future analysis and statistical testing.

We obtain a vector of imaging-derived features for each subject's visit. For similarity computation purposes, these features are normalized to the [0, 1] range based on maximum and minimum values. The original value of each feature is stored for display purposes (in labels and tooltips). Thus, each subject is described by several feature vectors over multiple visits. Similarity comparison between two studies is performed by computation of two measures: *cosine similarity* of feature vectors and *dynamic time warping* (DTW) of time-series.

Cosine similarity is defined as the similarity between two feature vectors for a given timepoint. Specifically, it is an angle between two feature vectors. In our application, we use a special case of weighted cosine similarity. Initial feature weights are assigned to be equal and can be changed by the user in the process of similarity-query construction as described in Section 9.3.7. Based on the weight, some features can be fully excluded from the comparison. For example, texture and intensity of the spleen can be

described as homogeneous, and thus the user can fully exclude such features if not found to be meaningful.

Time-series similarity is defined by the DTW distance measure [129, 130] algorithm with the above described cosine similarity as a distance cost at each point. DTW distance can be determined for sequences of unequal lengths and does not account for the difference in time between two points, which satisfies our requirements. As we have mentioned earlier, in the series of subject's scans the difference between two timepoints can be ignored. For the datasets where this assumption does not hold, two possible variations can be considered. In a simple case of evenly spaced events, Euclidean distance between time series with modification for missing data can be used. For irregular time series, variations of DTW or model-based search methods can be used. Additionally, DTW has been previously applied to the analysis of medical data with incomplete series [139].

AnaFe performs a similarity comparison of the selected multi-timepoint study to all other studies loaded in the system. The results of the similarity search update two synchronized views.

9.3 Visualization Design

AnaFe implements our collaborators' requirements through a number of highly interactive linked views described in this section:

- *Demographics Overview (DO)* - an overview of basic demographic information (age, gender, and disease types).

- *3D Small Multiples Objects over Time (SMO)* - an overview visualization of the organ progression via 3D small multiples that simplifies direct comparison of the organs as per **RQ2** and **RQ3**.

- *Feature Distribution Overview (FDO)* - a visualization of several groups of quantitative imaging features via familiar information visualization plots with interactive capabilities as per **RQ4**.

- *3D Object Detail (OD)* - a detailed progression view of 3D organ mesh shapes highlighting the changes between timepoints as an expansion to **RQ1**.

- *Measurement Progression over Time (MPT)* - an overall progression *trend* for the patient based on measurement metrics (volume, length, width, and thickness) (**RQ1**).

- *Feature Progression over Time (FPT)* - a heatmap-style visualization of feature progression over time (**RQ4**).

The layout of the application (Figure 27) is broken down into four main views and two control panels (main and feature selection). From the main control panel, the user can select a multi-subject study to be loaded from the back-end. The *DO*, *SMO*, *FDO*, and *MPT* views will be populated based on the number of subjects and the computed feature distributions as described in Section 9.3.5. The *OD* and *FPT* views alternate based on the user's selection in the control panel. *MPT*, *SMO*, and *FPT* (if selected) comprise a single table/list style scrollable view. The implemented capability of "virtual" scrolling allows for rendering of only currently visible items. This implementation alleviates the bottleneck of simultaneously handling multiple WebGL canvases and contexts for 3D small multiples rendering.

The selected study is displayed at the top of the table list and highlighted accordingly. The similarity comparison is performed based on the user defined query through filtering options in the FDO. We describe each of the views below.

9.3.1 Demographics Overview (DO)

The demographics overview consists of a bar chart visualization showing the distribution of gender, age, and disease status (sick or healthy) in the population. Spleen volumetry studies analyze the distribution of the organ's volume and other measurements (length, width, thickness) in the population for different demographic groups. Through interactions with this view, the user can filter which sets are shown and analyzed. By clicking on the column in the bar chart (or by brushing a selection over a range of columns), the user can select a category or a range of values. For example, only sick patients in the age group of 40-80 can be selected for analysis.

9.3.2 Measurement Progression over Time (MPT)

In clinical practice, the response to treatment or disease progression for the spleen is evaluated quantitatively based on length or volume of the organ
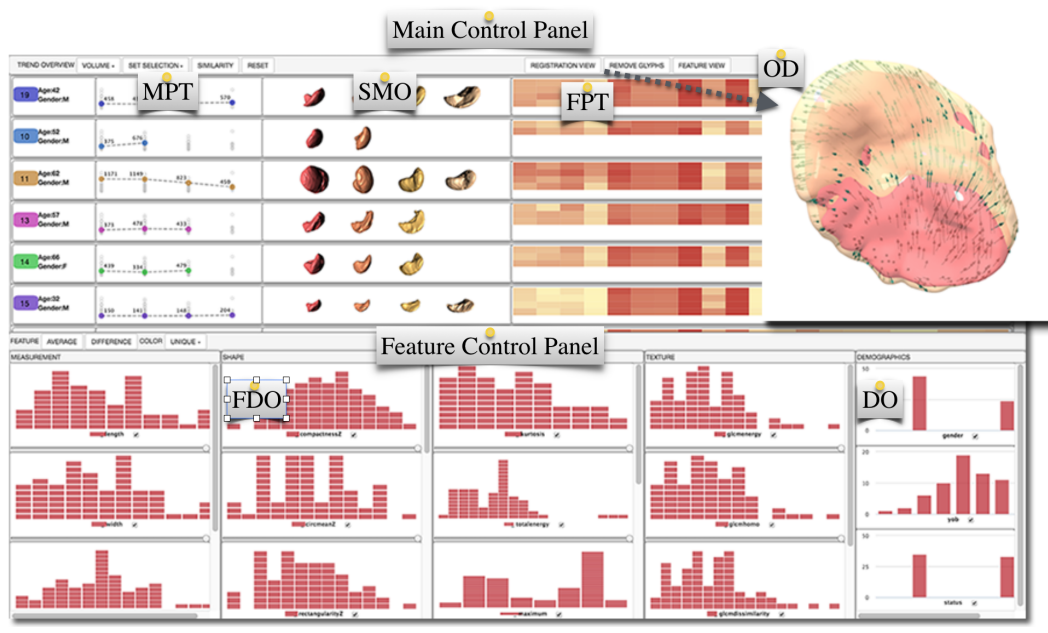
Figure 27: An application layout showing locations of all views, as described in Section 9.3. The arrow indicates interchangeable views.

[32, 124]. Similarly, in many applications, tumor growth is measured by a single diameter or volume. The progression trend (MPT view) should show an overall increase or decrease in this metric over a series of the subject's follow-up visits.

This view is implemented as a scatterplot (Figure 28) that shows the trajectory of the measurement for a given subject at each period of time (color filled markers). The dashed line connects two consecutive measurements into the progression trend per subject. In order to enable fast comparison of the trend of one subject with respect to others, this chart shows distribution of measurements across the population (grey empty markers). All measurements of one subject are highlighted on mouse interaction, thus providing trajectories for comparison.

As previously discussed [66], there is a lack of consensus on which measurements are used to compare disease progression. This panel allows the user to define a trend metric, e.g. splenic volume, and visually compare the change to the SMO and the OD views with the 3D rendering of the organ. Evaluating and understanding the manner in which spleens change in size,

96
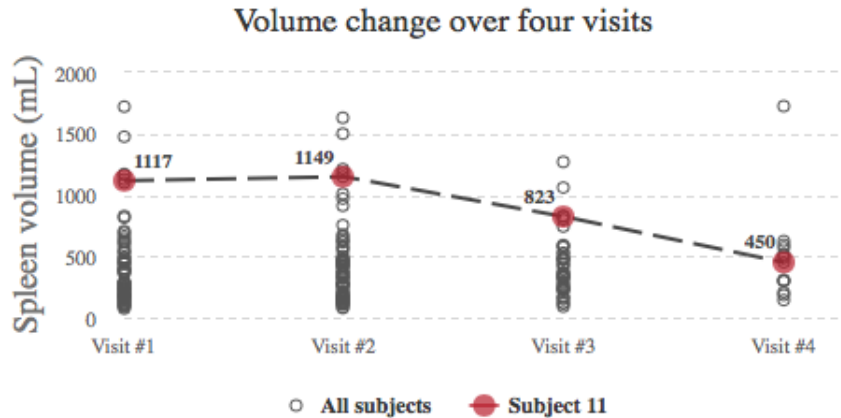
Figure 28: An example scatterplot of Measurement Progression over Time (MPT). Red markers show measurements for a given subject (#11) at each follow-up visit. The dashed line connects the measurements showing the progression trend for this given subject. Grey markers show measurements within the population at each visit.

for example shrink faster in length or in width, can characterize a disease progression or treatment. In this way, an abnormal spleen can be characterized even when it falls within the normal range for splenic size.

9.3.3 3D Small Multiples Objects over Time (SMO)

In a traditional clinical setting, only a few imaging results can be rendered simultaneously side-by-side. In such a scenario, the user can compare the change only for a single patient. To compare organ change among several subjects, we employ Tufte's principle of small multiples [140]. Keefe *et al.* [83] have used a similar approach for the visualization of 3D pig mandibles.

In our 3D small multiple views, the organ model is shown sequentially for each timepoint. It allows the user to quickly compare multiple organ models between several subjects (Figure 32 (1)). The time-series 3D object overview visualization allows individual rotation of each model for comparison from multiple angles. Furthermore, each of the organs can be selected with a double click for detailed investigation in the OD.

The SMO and MPT views are positioned side-by-side in the form of the table/list view. In this way the measurement scatterplot provides the re-
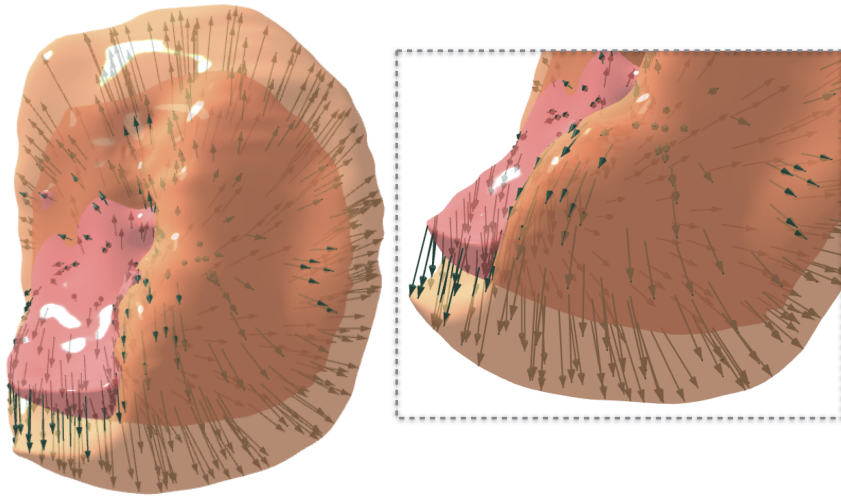
Figure 29: An example of the OD view showing two spleens rigidly aligned. Small black arrows indicate the direction of change and are obtained from the result of non-rigid registration.

quired trend visualization when change is not immediately visible on the 3D object. The 3D SMO view serves as a 3D thumbnail preview of the organ shapes. Correlating a trend metric from the MPT view (e.g. volume, length, or width) to the actual shapes can be done without context switching to multiple images. Thus, we combine both scientific and information visualization for analytics purposes.

9.3.4   3D Object Detail (OD)

Traditionally, organ progression observation and analysis is limited to statistical 1D measurements. Further investigation of 1D measurements and correlated spatial changes requires manual browsing through 2D slices of DICOM data for each dataset. Our collaborators have noted that this comparison demands substantial context-switching between several applications (statistical packages, DICOM viewers, etc.), incurring a significant time and effort overhead.

AnaFe renders the 3D surface mesh of the organ, thus enabling close investigation of local changes in the organ. Based on rigid alignment (ICP), organs from consecutive timepoints are aligned. Non-rigid registration (non-rigid ICP) provides point-by-point correspondences to render arrow glyphs. This type of representation can show the direction of the greatest change in
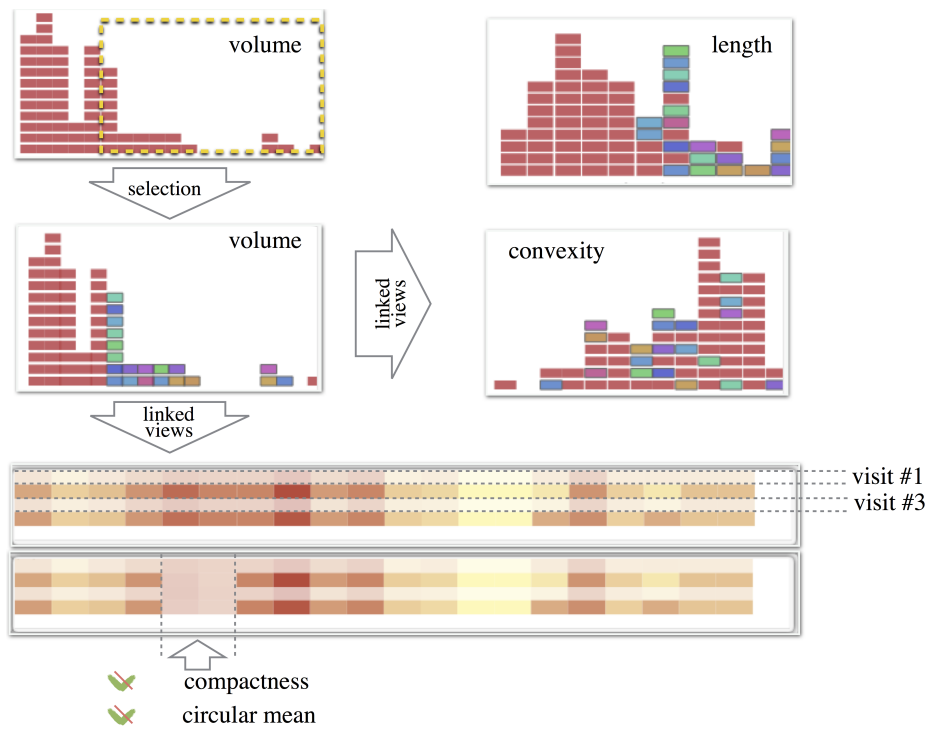
Figure 30: An example of linking via brushing support between "volume," "length," and "convexity" in the FDO view. By brushing in the "volume" chart visualization, the user can compare how selected subjects are distributed in other charts. Subjects and timepoints outside the selected range are removed from the similarity comparison (row). If the feature is fully deselected via the checkbox, the respective column in the FPT heatmap is dimmed out.

the spleen (Figure 29). From the semi-transparent rendering view the user can see whether the decrease was uniform in all dimensions or whether one dimension increased/decreased more rapidly. In many organs, a disease is defined in terms of distortion of their shape. Diseases that uniformly enlarge the organ without shape changes tend to fall into one group, while diseases that fundamentally change the shape tend to belong to another category. For the spleen, this view helps the user to understand whether the change is uniform or irregular.
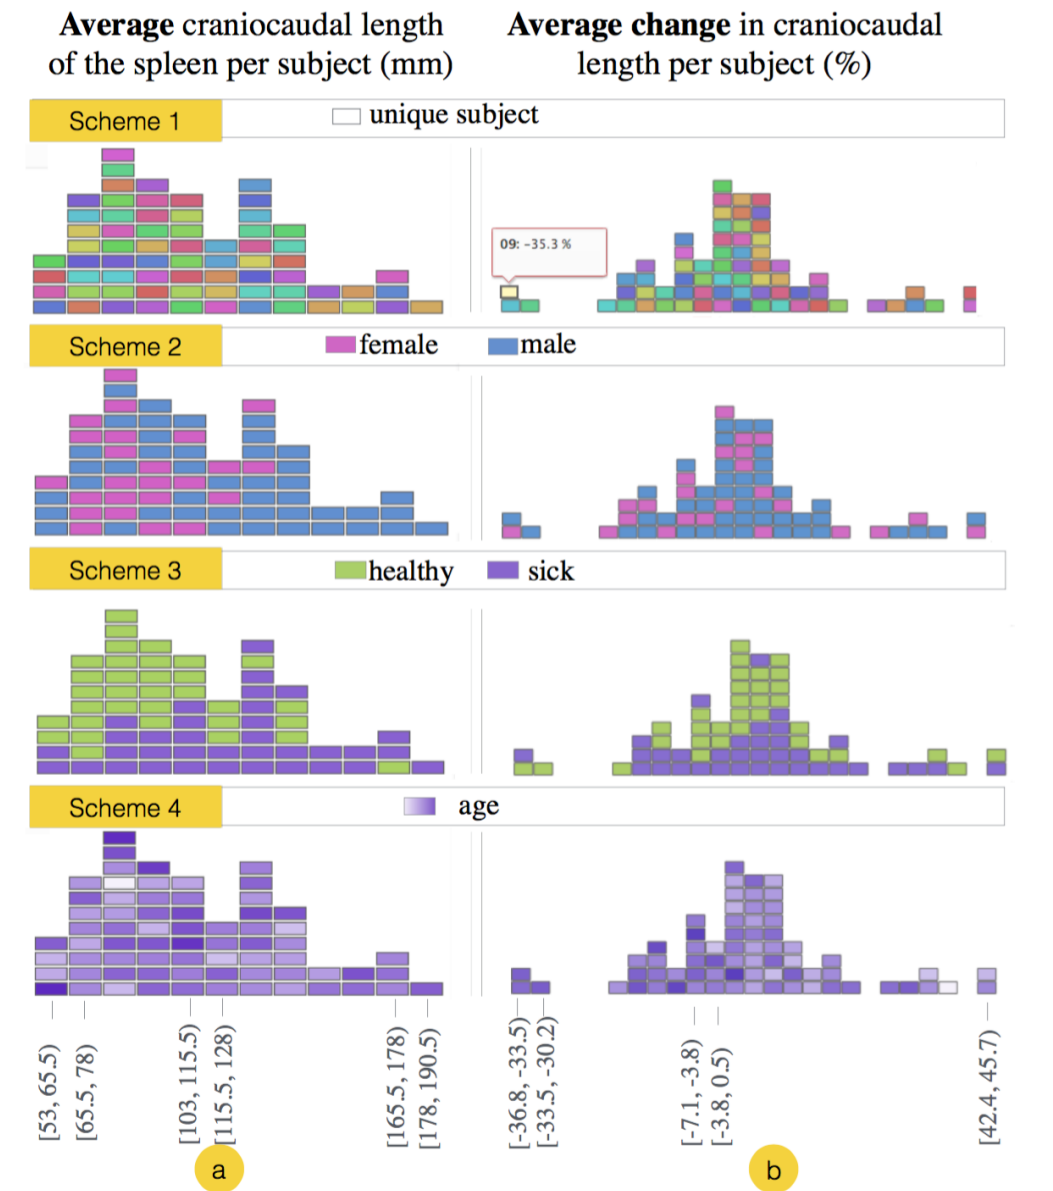
Figure 31: An example of feature histogram charts for an average craniocaudal length of the spleen per subject, and average relative change over time. The exact value of the feature is available in the tooltip when hovering over a subject's rectangle. The color scheme of the chart is selected by the user and is applied to each feature histogram. The figure presents 4 available color schemes: unique by subject, by gender, by disease status, and by age.

100
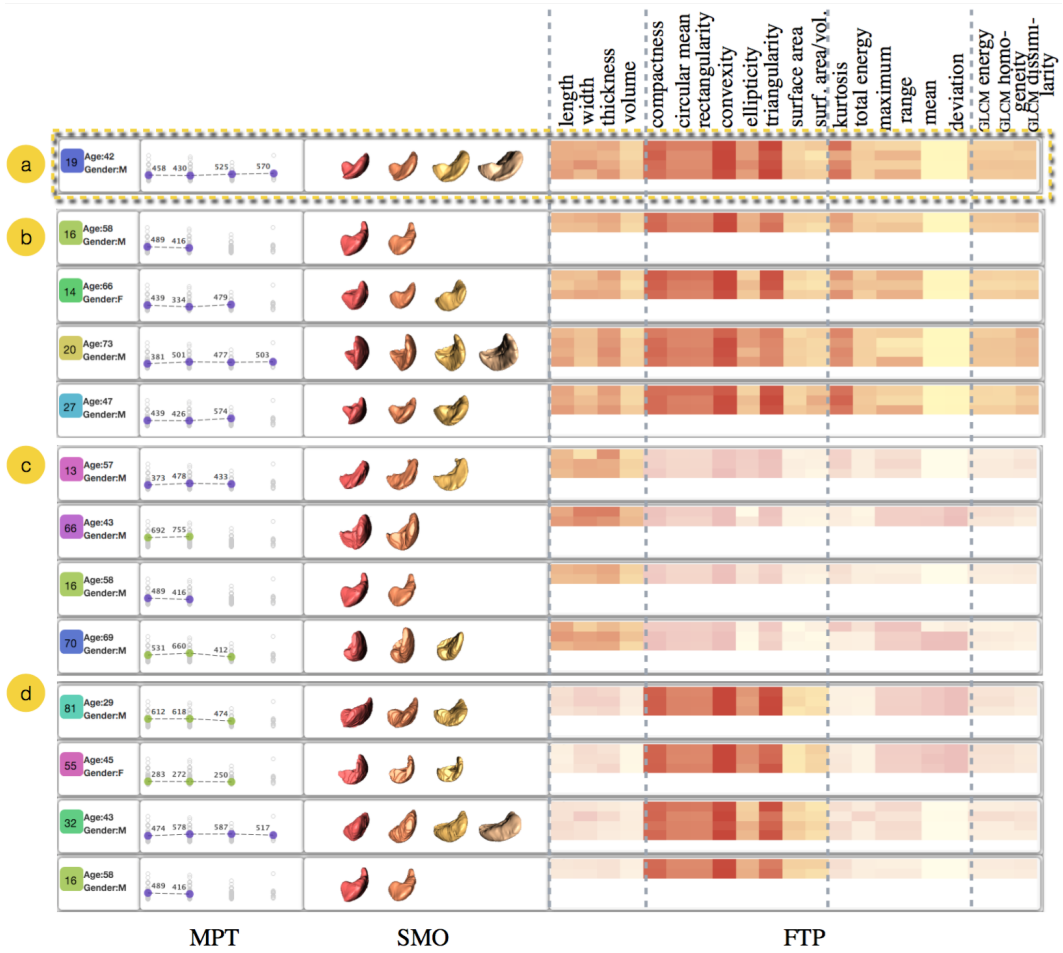
Figure 32: The result of similarity comparison for subject #19 (a): (b) based on all features; (c) measurement features only; (d) shape features only.
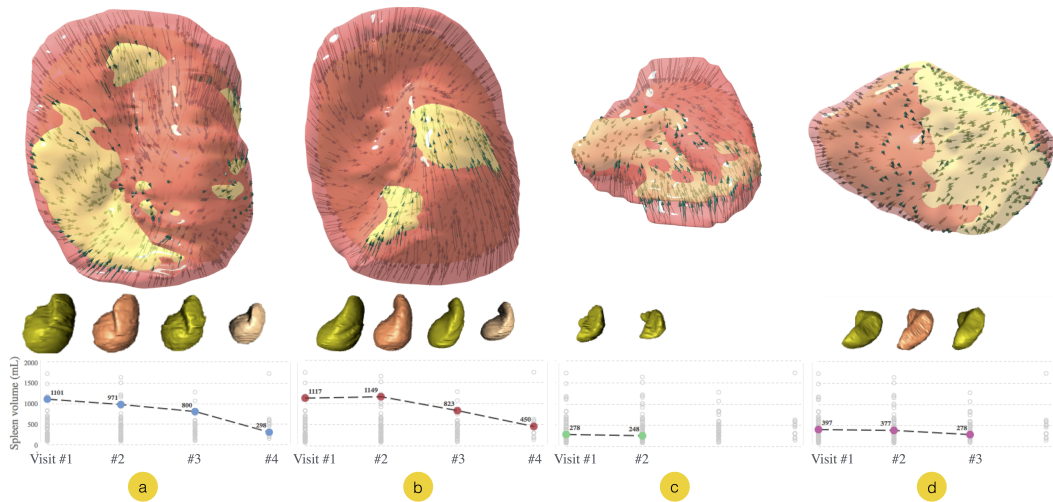
Figure 33: An example of OD, SMO, and MPT views for 4 subjects. Case (b) was found to be the most similar to case (a) in the sick subject data. Both show the same decrease in organ volume. Cases (c) and (d) were found to be most similar among healthy subjects. While the spleen (c) decreased in volume, (d) remained almost unchanged.

9.3.5 Feature Distribution Overview (FDO)

Histograms are common to visualize density distribution in the data. Feature visualization of multi-timepoint data should provide an understanding of distribution of average feature values and relative changes for each subject. The number of histogram bins and their size have to be chosen carefully to depict the overall density. While a basic histogram can provide an answer with respect to the overall distribution of values, it does not identify individual entities that contribute to each bin.

With the above ideas in mind, we construct a custom feature histogram view (Figure 31) that shows individual entity contributions. In our histogram the bin width is computed based on the Freedman-Diaconis rule [52]. For a dataset of 68 subjects, this resulted in up to 20-40 values per each bin for a given feature. The relatively small bin size allows us to display each entity individually as a rectangle of the respective width, and with a height computed based on the chart dimensions. The chart's width and height are equal for all feature types and are computed by the application based on the screen size. In our feature histogram, rectangles ▭ can be colored using one

of the color schemes: individually for each subject, disease status (healthy - ▭ , sick - ▭), gender (male - ▭ , female - ▭), or based on age (▭ to ▭). *AnaFe* allows for color selection based on one of these specified color schemes and switching between two summarized types of feature distribution visualization as required. The features can be summarized by average value per subject and average relative change between consecutive timepoints.

The FDO view is one of the key visualization components of the system. It provides a global overview without eliminating individual per subject values. Globally, this view provides insight into the distribution of features. Based on interaction (hovering over a rectangle), the user can investigate into which range a given subject belongs for each feature (the corresponding subject is highlighted in all views). In clinical research, the expert can specify population groups to be included in the similarity comparison. For instance, she might want to compare all the "long" spleens of sick subjects that were undergoing treatment. By brushing over the "length" FDO chart and over the "disease status" DO chart, the population is filtered and will be used for comparison.

9.3.6   Feature Progression over Time (FPT)

FPT is a heatmap visualization of feature vectors for each subject's visit (Figures 30 and 32 (3)). The FPT view is not immediately visible and can be shown on demand through selection on the control panel. The row number of within the chart indicates the visit, and the cell is a value of the feature at a given point in time. Cell color is determined based on the normalized value of the feature (the data value is mapped to interval from 0 to 1). The darkest red ▭ identifies the highest value of the feature (normalized to 1) and yellow ▭ indicates lowest value (normalized to 0). The original value can be viewed by hovering over the cell. Through this color coding, the FPT view shows the change of each individual feature for the subject over time in a single view. For instance, measurements such as volume, length, width, and thickness of the organ (shown in the first group in Figure 32) and quantitative radiomics shape features can be surveyed in parallel across a treatment regimen.

9.3.7 Visual Queries and Interaction

Our visual query interface is comprised of multiple feature distribution histogram charts. We have identified 4 groups of imaging-derived features. In this application we currently use up to 8 features in each group with the possibility of adding more features in the future.

Query construction. The feature query interface allows the user to specify three types of information required for the comparison of time-series similarity: inclusion/exclusion of each feature, weight of feature contribution, and range of feature interval. The user can *select a feature* (include/exclude) via a checkbox in the legend in order for the feature to be included. The *feature weight* can be specified via a slider located next to each feature histogram. *Feature range selection*, brushing selection over a block of the histogram, allows the user to select a range of values for each feature. If a histogram bin is selected, all sets within that range will be included in the similarity comparison. Mouse hover *links* all feature values for a given subject on all histograms.

Query results. A similarity query compares the currently selected study (displayed first) with every other multi-timepoint study available in the dataset. If a value of a feature falls outside the range selected by brushing, and leaves only a single timepoint per subject, the subject will be compared based on that one point. After the user completes the feature selection, she can request the similarity results. Results are displayed in all table/list views: subject information, MPT, and SMO.

9.3.8 Implementation

The *AnaFe* front-end is implemented using standards-compliant Javascript, CSS, HTML5, and WebGL and utilizing `Bootstrap`[9] layouts. The core application scaffolding is driven by `Angular.js`[10]. The visualizations are generated using `three.js`[11] and `Highcharts.js`[12]. The core web application functionality is implemented on top of `Node.js`[13], with `Sails.js`[14]. Data

---

[9]http://getbootstrap.com

[10]https://angularjs.org

[11]http://threejs.org

[12]http://www.highcharts.com

[13]https://nodejs.org/

[14]http://sailsjs.org

pre-processing and derivation of all imaging features is handled in `Python` utilizing `scikit-image`[15].

9.4 Application

In this section, first, we describe the data collection used. Then we present two usage scenarios, showcasing our application in the process of analyzing a collection of splenic imaging datasets, and obtain related domain expert feedback. For both of the scenarios we show our system by providing an overview of the data, allowing the user to zoom, filter, and examine example results and details on demand. Finally, we present an example of *AnaFe*'s extensibility to other organs and medical imaging types.

9.4.1 Spleen Data

The data used in our study was collected retrospectively over a period of three years. This data includes two groups of subjects: those with *no known causes of splenomegaly or diseases* affecting the spleen (healthy), and those *with lymphoma or leukemia*, diseases known to affect the organ size, that were undergoing treatment (sick). Only subjects with two or more abdominal CT scans over a three year period were included. This study was approved by our IRB. The DICOM image data was de-identified and stored in a secure off-line database.

A total of 33 healthy subjects with up to 3 follow-ups over the given time period were included, resulting in 79 healthy datasets. For "sick" subjects, we included data with up to 4 follow-up scans over the course of treatment / disease observation, resulting in 35 subjects with a total of 110 datasets. The data included 24 female subjects and 44 male.

For each study, manual tracing of the splenic contour was performed on each axial slice in a blinded fashion under the supervision of a fellow-trained attending radiologist, providing a ground truth segmentation. This work was performed using the Alice software[16].

While the time between the follow up scans varied from one week to over a year, it was noted by our collaborators that only the difference between measurements was important, and the time between visits was not of high significance, as explained earlier.

---

[15]http://scikit-image.org

[16]Paraxel Informatics; Waltham, MA

9.4.2 Case 1: Similarity Comparison

In this section we demonstrate the similarity comparison feature for multiple datasets and features. We focus on the usability and interactivity of the system by providing step-by-step scenarios of usage.

Data selection. First, the user should select a multi-subject study and select a set for comparison. By default, all of the study parameters and imaging-derived features will be included in the search. *AnaFe* provides full customization of the query parameters to the user.

Visual queries. The search for similar shapes can be narrowed down through *AnaFe*'s visual query interface in the FDO view. Based on the selection of search parameters (types of measurements and features), some timepoints will be excluded from comparison. Once the query selection is complete, the entire time series for a particular subject is compared to that of all other subjects. The similarity is computed in real-time (for 189 spleens), and the result is displayed immediately in the list/table views. Initially, the preview of temporal sets in the 3D SMO view provides immediate feedback to the user.

Example of findings. Figure 32(a) shows an example of the four most similar sets as compared to subject #19. First, similarity comparison was performed using all available features (Figure 32(b)), next based on diameters and volume of the organ (c), and finally based only on shape features (c). The color of the marker in the MPT scatterplot indicates disease status of the subject (with green being healthy). In cases (c) and (d) the system has identified healthy subjects #66, #70, #81, and #55 as similar. In the first case, with all features selected, the system has identified only diseased subjects. Similarity comparison in this case did not include demographics information and was based solely on imaging features.

Zooming-in on the result. Next, the user can select the organs by double clicking on them for zoomed-in investigation in the 3D OD view. Figure 33 shows rendering of temporal scans for four subjects. All measurements (volume, width, length, thickness) were used for comparison query. 3D meshes of the organ are all registered sequentially. The reviewer can inspect the direction of change. For example, for sick subjects (a) and (b) in Figure 33, change happens in all dimensions. In case (c), some change has occurred in the healthy subject as well. For such a small organ volume, this change may indicate the variation in hydration status. However, for most of the healthy subjects found in our data, there is no change between consecutive

visits.

9.4.3 Case 2: Interactive Feature Exploration

Radiomics uses medical imaging features for the prediction of diagnostic outcomes [34, 113]. In the examined applications these features provide a non-invasive way of quantifying and monitoring tumors. We use this comprehensive quantification of imaging to find similarities in progression of disease and response to treatment in the spleen data. However, there haven't been any previous attempts at visualizing these features for detailed user examination.

Feature overview. *AnaFe* provides an overview of all features in the FDO view, thus creating a second scenario for the view usability. From this view, we can see that, as expected, the mean and standard deviation of intensity communicate homogeneous density properties of the organ on the CT scan. Thus, these features might not be discriminative enough to show differences between subjects.

Feature selection. The user can unselect these features from the similarity comparison, and the corresponding columns will be dimmed out on the FPT heatmap (Figure 30). The user can inspect relationships between features of different subjects by brushing the selection over the FDO histograms.

Example of findings. As shown in Figure 30, organs with larger volumes also have larger craniocaudal length, which is a well known relationship for the spleen. However, relationships between other features for this organ are largely unknown. *AnaFe* creates a unique opportunity for the user to inspect the derived features simultaneously and select particular ones for further statistical testing.

9.4.4 Extensibility to Other Domains

AnaFe was developed with the goal of exploring changes in temporal imaging-derived features, which does not limit its application to CT spleen data. To show extensibility of our application, we use observations of prostate cancer on Magnetic Resonance Imaging (MRI). Prostate cancer is the most common malignancy and second most common cause of cancer-related mortality in men. Similar to splenic increase in size, enlargement in the prostate can be palpable and is used in primary staging. Thus, the size of the organ is one of the useful indicators that can be obtained from the segmented organ

imaging.

We use temporal prostate data from the Prostate MR Image Database[17]. Manual segmentation was available only for two data timepoints per subject, and demographics information was not available. Additionally, no disease status was available for these subjects. The study by domain experts is required to evaluate results of the analysis. Figure 34 shows an example of *AnaFe* with prostate data.

9.4.5 Expert feedback

Two of our collaborating radiologists inspected the application and its features. We have recorded their feedback and the most important points they made about the usability of the tool.

Linking measurements and 3D data over time: MPT, SMO, and OD views. The MPT view for craniocaudal length and volume (well known measurements) allows the researcher to compare their trends rapidly and gain immediate visual feedback. It becomes particularly powerful in combination with 3D rendering of the organ sequence (SMO). The overlay rendering of aligned mesh surfaces (3D OD view) provides a useful picture of how the organ has changed between patient visits. In the software that is currently used to perform medical research on the spleen, neither registration of the organ, nor 3D rendering of overlays are implemented. Thus, the MPT, SMO, and OD views, when used jointly, *provide the missing links for the current analytical workflow*, and allow for substantial time savings when comparing multiple organs. As a potential improvement, our collaborators indicated that some sort of on-mesh annotation of the areas of major change in shape or size within the 3D view would be very useful.

Feature distribution and characterization of the data: FDO view. The FDO view was selected as the most interesting feature due to its ability to analyze multiple shape features simultaneously. As was noted by our collaborators, *differentiation between sick and healthy organs* based on a set of features was unclear. The FDO chart provided visualization of distribution and ranges of each feature, as well as immediate visual feedback (based on the color of the category) if the separation between the values based on the category exists. For instance, individual color coding by subject has shown an *already well known correlation of the organ's craniocaudal*

---

[17]http://prostatemrimagedatabase.com

length and volume. That is for spleens with a large volume, the length is also large. For some of the shape features, differentiation between sick and healthy subjects was very visible. Using this knowledge, and the subject's diagnostic information (currently not available in the anonymized data), the clinician could investigate sources of such differences, for example, the effect of portal vein hypertension on the organ shape. This is an interesting direction for future research: analysis of statistical significance of shape features in predicting organ disease status, and predicting normality of diseased organs and abnormalacy of healthy ones.

In summary, the ability to visualize a large number of datasets simultaneously has significantly improved over routine comparison that has previously required context switching between different applications and was limited to comparison of only a few parameters in static charts (IS1). Most importantly, *AnaFe* has provided a unique look at multiple snapshots of spleen data and its imaging-derived features changing over time (IS2).

9.5   Conclusions

In this work, we described a visual analytics framework for the exploration of large collections of medical imaging datasets, with a focus on disease progression accompanied by progressive spleen enlargement. Through multiple linked views we allow the user (researcher or the radiologist) to interactively explore multiple imaging-derived features. Meanwhile, a 3D mesh view allows the user to examine changes in the organ more closely and see the relation of the traditionally used 1D parameters to the full organ view. Currently, our application supports the set of most common robust and reproducible imaging features. These features describe organ intensity, shape, texture, and measurements. Our application allows for observation of trends over time to determine similarity in disease progression and outcomes. We have currently explored only a limited set of radiomics features, and were already able to find their usability for the spleen. It is also in our interest to pursue a time-dependent comparison in the future with additional parameters from other tests considered (beyond CT imaging), which requires acquisition of more data and additional IRB approval.

Figure 34: An example of *AnaFe* use with prostate data. User interaction with the system is displayed in the average (top) and change FDO views (bottom). Hovering the mouse over the marker (subject #73) in the MPT view (a) links to all features for this subject in the FDO view. Selected subjects in the view (c) are highlighted across all FDO charts. Not selected subjects (b) are dimmed out in the FPT view.

10 Conclusions

10.1 Summary of Contributions

This dissertation has presented contributions in several areas of interactive visual analytics, both on mobile devices and in the form of the web-based tools. Through the number of the following contributions, we address the challenges of making data visualization, interaction, and analytics ubiquitous and informative.

In the first part of this dissertation, we introduce improvements and solutions to the problem of mobile visualization of 3D medical data. The size of the data coming from modern CT scanners presents a challenge for both transmission to and rendering on the mobile device. Consequently, we follow two distinct approaches to mobile visualization:

- Following an image-based approach to mobile visualization, we design a saliency-aware scalable compression scheme. The method assumes unique properties of medical data visualization based on the transfer function design. It allows for the progressive transmission of the data from initially lossy to fully lossless reconstruction.

- Following a streaming-based approach to mobile visualization, we implement a pipeline for remote rendering. In this framework, the rendering of the data happens entirely on the server. The rendering result is encoded into H.264 video stream and streamed to the client. The client application supports user interaction and render state command encoding which is transmitted back to the server and used to update the rendering.

- Additionally, we evaluate new approaches to existing problems of interaction with volumetric data on touchscreen devices. In particular, we focus on the target selection in 3D volumetric data. In this work, we evaluate three types of selection techniques that are currently unique to touchscreen devices: selection with a finger, selection with a stylus with a ray cast at a straight angle, and at an arbitrary angle based on stylus' tilt. We show that users benefit from using the stylus for 3D selection both with and without an angle.

- Finally, we explore the space of immersive analytics and study interactions with scientific data. Particularly, we focus on mid-air interactions

111

with scientific volumetric data in mixed reality. We propose a set of natural gestures based on positional hand tracking through head mounted display and rotational tracking through a custom built hardware.

In the second part of this work, we introduce a number of contributions to the analysis of splenic data. As spleens shrink and grow in response to disease progression and treatment, one must carefully measure the volume of the organ. Thus, we analyze existing approaches to estimation of the splenic volume and improve upon them. Most importantly, high variation in organ's shape and size lead us to employ visual analytics solution to further study variation of the organ and its changes.

- We analyze the correlation of several existing volumetric measurements to the time-varying splenic data. Unlike previous works, we consider several repeated measurements over time for each subject. We evaluate the correlation of changes in volume to changes in "splenic indices" based on one, two, or three orthogonal measurements of the organ.

- Using conformal welding heuristic, we determine the optimal plane where two orthogonal measurements of the organ can be taken to approximate the volume of the spleen. t

- *AnaFe* is a versatile tool for exploration of splenic size and morphology. It offers insights into the overall distribution of robustly extracted and reproducible quantitative imaging features and their changes within the population, and also enables detailed analysis of individual cases. It performs similarity comparison of temporal series of one subject to all other series in both sick and healthy groups.

10.2 Future Work

Based on the advancements made in this thesis, we propose a long and short term work plan. Here we outline the next steps to continue on the latest projects, as well as an overall overview of the field.

10.2.1 Short-term Guidance

We propose two possible directions to continue work on interaction with mobile visualization targeted at touchscreen devices. First, our preliminary study of three types of interaction provided some initial information on angle and depth parameters for 3D target acquisition. As a next step, one can complete the formal study of model interaction for volumetric target selection. Secondly, in this work we only focused one type of interaction with volumetric data on touchscreen devices. Hence, there is an opportunity to research novel approaches for high quality, informative, and user-friendly interaction that takes advantage of novel input types. Such work can address both a design of a traditional transfer function and introduce novel methods of volumetric navigation and exploration. For example, the design of a good transfer function is a complicated process and might require user's knowledge about the mapping properties of the function. However, a good design is essential for effective visualization. A stroke-based transfer function has been previously proposed for a traditional desktop setup with a mouse input device. On the touchscreen devices one accommodate for varying width and depth of stroke based on the user's touch. Optimal visualization result is highly dependent on the ability of the user to adjust the viewing options of the data. Additional methods of interaction with volumetric data, for example, clipping plane positioning, volume of interest selection can be reintroduced to take advantage of innovations in the hardware of the touchscreen devices.

Rapidly emerging space of mixed reality provides a broad spectrum of opportunities for immediate research work as well, and is much less explored compared to the touchscreen devices. Interaction with volumetric data in the mixed reality setting is drastically different from the "traditional" projection rendering onto the 2D display. While added dimensionality removes the need of mapping of 3D interactions to the 2D input, it faces the challenge of user perception of such interaction. In out work, we have focused only on the quantitative evaluation of the building blocks of mid-air interactions: simplest gestures comprised of hand position and rotation. However, user

understanding and perception of mid-air direct touch for all of the above-mentioned types of interaction with volumetric data needs to be carefully studied.

Augmented reality applications allow easy participation in collaborative data exploration. Augmented and mixed reality HMDs create personalized and private views of the data for each user. Custom views for medical data can mean a number things: personal rendering modes for each user, different multi-modal visualization settings, varying data privacy settings. But what does this mean for interaction with volumetric data in collaborative setting? While the data may be visualized per user in a certain rendering mode, user interaction should be visualized for all collaborators independent of the viewing setting. The study of consistent and simple visual feedback metaphors for volumetric data exploration in collaborative setting is one of the short-term projects that can be built upon our presented work.

10.2.2 Long-term Guidance

In this dissertation, we have addressed one of the key challenges when working with multidimensional medical data: size of the data that needs to be rendered on the mobile device. Similarly, in the immersive environment, self-contained HMDs are limited to on-device computation. Hence, the bottleneck of rendering on the HMD itself is a challenge that is exacerbated by the computation overhead for communication with other HMDs in the collaborative setting. Such limitation can be addressed in couple of ways. First of all, it is important to highlight that remote volumetric rendering remains needed until significant improvements in mobile hardware provide sufficient resources for highly demanding volumetric data rendering. Currently, in collaborative setting, a fully remote rendering framework would require a server to produce multiple custom per user views. Alternatively, distributed rendering frameworks can offload the expensive parts of rendering computation, such as lighting and shading, to a server or rendering cluster, thus making a single computation instance available for all users. At the same time, personal HMDs can do a camera pass individually for each user based on head orientation.

Secondly, one can take advantage of advanced compression schemes for immersive visualization. Similar to our work, one can immediately take advantage of saliency-aware compression schemes and ensure that lossless requirement is satisfied in the immersive setting. Compression schemes for

medical imaging data can undergo a significant upgrade by utilizing deep learning from large imaging databases. Currently, such research is constrained by data privacy regulations. While providing patient protection, such regulations significantly stifle advancements in medical imaging analytics. Similarly to the open source software, one can imagine finer guidelines towards open clinical research data, including imaging, that accelerate the research.

Interaction research work in this dissertation has been limited by readily available or custom assembled hardware from existing sensors. Similar works in the field of immersive visualization and analytics focus on tactile and tangible interfaces for data exploration. Similar to these interfaces, mid-air interaction needs to adapt to a variety of user behaviors. One can easily make enhancements to the proposed volumetric interaction framework, for example, using muscle sensors or gaze tracking devices. However, muscle sensors have proven to be hard to calibrate for a large group of users. This brings us to a key challenge of natural user interfaces for immersive environments: creating simple and intuitive methods of interacting with the data that are robust and uniform across a large number of users. Standardization of interaction methods based on the natural user input is one of the requirements for a user-friendly mixed reality world.

References

[1] *Intel Media SDK* . http:// software.intel.com/en-us/vcsource/tools/media-sdk.

[2] *LibVLC library.* http://libvlcnet.sourceforge.net/.

[3] *LIVE555 library.* http://www.live555.com/liveMedia/.

[4] *NVCUVENC and NVENC video encoders, NVUVID.*

[5] *NVIDIA OpenGL SDK 10 Code Samples: Render to 3D Texture.* http://developer.download.nvidia.com/SDK/10/opengl/samples.html.

[6] *Unity Volume-Raytracing Sample Code.* https://github.com/gillesferrand/Unity-RayTracing.

[7] P. Angelelli, S. Oeltze, J. Haasz, C. Turkay, E. Hodneland, A. Lundervold, A. J. Lundervol, H. Hause, and B. Preim. Interactive visual analysis of heterogeneous cohort study data. *IEEE Computer Graphics and Applications*, 1(5):70–82, 2014.

[8] F. Argelaguet and C. Andujar. Visual feedback techniques for virtual pointing on stereoscopic displays. *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*, pages 163–170, 2009.

[9] F. Argelaguet and C. Andujar. A survey of 3d object selection techniques for virtual environments. *Computers & Graphics*, 37(3):121–136, 2013.

[10] R. Azuma, Y. Baillot, R. Behringer, S. Feiner, S. Julier, and B. MacIntyre. Recent advances in augmented reality. *IEEE Computer Graphics and Applications*, 21(6):34–47, 2001.

[11] B. Bach, R. Dachselt, S. Carpendale, T. Dwyer, C. Collins, and B. Lee. Immersive analytics: exploring future interaction and visualization technologies for data analytics. *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*, pages 529–533, 2016.

[12] R. Ballagas, J. Borchers, M. Rohs, and J. G. Sheridan. The smart phone: a ubiquitous input device. *IEEE Pervasive Computing*, 5(1):70–77, 2006.

[13] J. Barcroft and J. G. Stephens. Observations upon the size of the spleen. *The Journal of Physiology*, 64(1):1–22, 1927.

[14] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Hybrid tactile/tangible interaction for 3d data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):881–890, 2017.

[15] A. S. Bezerra, G. D'Ippolito, S. Faintuch, J. Szejnfeld, and M. Ahmed. Determination of splenomegaly by CT: Is There a Place for a Single Measurement? *American Journal of Roentgenology*, 184(5):1510–1513, 2005.

[16] M. Billinghurst and H. Kato. Collaborative augmented reality. *Communications of the ACM*, 45(7):64–70, 2002.

[17] J. Blaas, C. P. Botha, and F. H. Post. Interactive visualization of multi-field medical data using linked physical and feature-space views. *EuroVis*, pages 123–130, 2007.

[18] D. Bowman, E. Kruijff, J. J. LaViola Jr, and I. Poupyrev. *3D User Interfaces: Theory and Practice, CourseSmart eTextbook*. Addison-Wesley, 2004.

[19] G. Bruder, F. Steinicke, and W. Stürzlinger. Effects of visual conflicts on 3d selection task performance in stereoscopic display environments. *IEEE Symposium on 3D User Interfaces*, pages 115–118, 2013.

[20] G. Bruder, F. Steinicke, and W. Sturzlinger. To touch or not to touch?: comparing 2D touch and 3D mid-air interaction on stereoscopic tabletop surfaces. *Proceedings of the 1st Symposium on Spatial User Interaction*, pages 9–16, 2013.

[21] V. Buchmann, S. Violich, M. Billinghurst, and A. Cockburn. Fingartips: gesture based direct manipulation in augmented reality. *Proceedings of the 2nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*, pages 212–221, 2004.

[22] S. Busking, C. P. Botha, and F. H. Post. Dynamic multi-view exploration of shape spaces. *Comput Graph Forum*, 29(3):973–982, Aug. 2010.

[23] C. Butson, G. Tamm, S. Jain, T. Fogal, and J. Kruger. Evaluation of interactive visualization on mobile computing platforms for selection of deep brain stimulation parameters. *IEEE Transactions on Visualization and Computer Graphics*, 19(1):108–117, 2013.

[24] J. J. Caban, P. Rheingans, and T. Yoo. An evaluation of visualization techniques to illustrate statistical deformation models. *Computer Graphics Forum*, 30(3):821–830, 2011.

[25] B. Cabral, N. Cam, and J. Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. *Proceedings Symposium on Volume Visualization*, pages 91–98, 1994.

[26] L. Campoalegre, P. Brunet, and I. Navazo. Interactive visualization of medical volume models in mobile devices. *Personal and Ubiquitous Computing*, pages 1–12, 2012.

[27] L. Campoalegre, P. Brunet, and I. Navazo. Interactive visualization of medical volume models in mobile devices. *Personal and Ubiquitous Computing*, 17(7):1503–1514, 2013.

[28] J. Carmigniani, B. Furht, M. Anisetti, P. Ceravolo, E. Damiani, and M. Ivkovic. Augmented reality technologies, systems and applications. *Multimedia Tools and Applications*, 51(1):341–377, 2011.

[29] J. L. Carrasco, B. R. Phillips, J. Puig-Martinez, T. S. King, and V. M. Chinchilli. Estimation of the concordance correlation coefficient for repeated measures using SAS and R. *Comput Methods Programs Biomed*, 109(3):293–304, 2013.

[30] L.-W. Chan, H.-S. Kao, M. Y. Chen, M.-S. Lee, J. Hsu, and Y.-P. Hung. Touching the void: direct-touch interaction for intangible displays. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2625–2634, 2010.

[31] P. N. Chaware, S. M. Belsare, Y. R. Kulkarni, S. V. Pandit, and J. M. Ughade. The Morphological Variations of the Human Spleen. *Journal of Clinical and Diagnostic Research*, pages 159–162, 2012.

[32] B. D. Cheson, R. I. Fisher, S. F. Barrington, F. Cavalli, L. H. Schwartz, E. Zucca, and T. A. Lister. Recommendations for initial evaluation, staging, and response assessment of Hodgkin and Non-Hodgkin Lymphoma: The Lugano classification. *Journal of Clinical Oncology*, 32(27):3059–3067, 2014.

[33] B. Chow, P. Lu, and L. Ni. Hamilton's Ricci Flow. *American Mathematical Soc.*, 77, 2006.

[34] C. Christin, H. C. J. Hoefsloot, A. K. Smilde, B. Hoekman, F. Suits, R. Bischoff, and P. Horvatovich. A critical assessment of feature selection methods for biomarker discovery in clinical proteomics. *Molecular & Cellular Proteomics*, 12(1):263–276, 2013.

[35] K. K. Chui, J. B. Wenger, S. A. Cohen, and E. N. Naumova. Visual analytics for epidemiologists: understanding the interactions between age, time, and disease with multi-panel graphs. *PloS ONE*, 6(2):e14683, 2011.

[36] D. Coffey, N. Malbraaten, T. B. Le, I. Borazjani, F. Sotiropoulos, A. G. Erdman, and D. F. Keefe. Interactive slice WIM: Navigating and interrogating volume data sets using a multisurface, multitouch VR interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626, 2012.

[37] A. Cohé, F. Dècle, and M. Hachet. tBox: a 3D transformation widget designed for touch-screens. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3005–3008, 2011.

[38] A. Colley, J. Häkkilä, J. Schöning, F. Daiber, F. Steinicke, and A. Krüger. Touch the 3rd dimension! Understanding stereoscopic 3D touchscreen interaction. *Australian Computer-Human Interaction Conference*, pages 47–67, 2013.

[39] L. Cools, M. Osteraux, L. Divano, and L. Jeanmart. Prediction of Splenic Volume by a Simple CT Measurement: A Statistical Study. *Journal of Computer Assisted Tomography*, 7(3):426–430, 1983.

[40] L. D. Cutler, B. Fröhlich, and P. Hanrahan. Two-handed direct manipulation on the responsive workbench. *Proceedings of the Symposium on Interactive 3D Graphics*, pages 107–114, 1997.

[41] E. De Guzman, F. Ho-Ching, T. Matthews, T. Rattenbury, M. Back, and S. Harrison. Eewww!!: Tangible interfaces for navigating into the human body. *CHI Extended Abstracts on Human Factors in Computing Systems*, pages 806–807, 2003.

[42] R. M. Dickey, N. Srikishen, L. I. Lipshultz, P. E. Spiess, R. E. Carrion, and T. S. Hakky. Augmented reality assisted surgery: a urologic training tool. *Asian Journal of Andrology*, 18(5):732, 2016.

[43] B. Diedenhofen and J. Musch. cocor: A comprehensive solution for the statistical comparison of correlations. *PloS ONE*, 10(4):e0121945, 2015.

[44] K. Dmitriev, I. Gutenko, S. Nadeem, and A. Kaufman. Pancreas and cyst segmentation. *SPIE Medical Imaging*, pages 97842C–97842C, 2016.

[45] J. A. Fails, A. Karlson, L. Shahamat, and B. Shneiderman. A visual interface for multivariate temporal data: finding patterns of events across multiple histories. *IEEE Symposium On Visual Analytics Science And Technology*, pages 167–174, 2006.

[46] Y. Fang, Z. Chen, W. Lin, and C.-W. Lin. Saliency detection in the compressed domain for adaptive image retargeting. *IEEE Transactions on Image Processing*, 21(9):3888–3901, 2012.

[47] Z. Fang, T. Möller, G. Hamarneh, and A. Celler. Visualization and exploration of time-varying medical image data sets. *ACM Proceedings of Graphics Interface*, pages 281–288, 2007.

[48] H. M. Fenlon, D. P. Nunes, P. C. Schroy, M. A. Barish, P. D. Clarke, and J. T. Ferrucci. A comparison of virtual and conventional colonoscopy for the detection of colorectal polyps. *New England Journal of Medicine*, 341(20):1496–1503, 1999.

[49] P. M. Fitts. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology*, 47(6):381, 1954.

[50] D. Fleischmann and D. C. Miller. Clinical 3D and 4D Imaging of the Thoracic Aorta. *Diseases of the Heart, Chest Breast Diagnostic Imaging and Interventional Techniques*, pages 119–130, Feb. 2007.

[51] J. D. Foley, V. L. Wallace, and P. Chan. The human factors of computer graphics interaction techniques. *IEEE Computer Graphics and Applications*, 4(11):13–48, 1984.

[52] D. Freedman and P. Diaconis. On the histogram as a density estimator: L2 theory. *Probab. Theory Related Fields*, 57(4):453–476, 1981.

[53] C.-W. Fu, W.-B. Goh, and J. A. Ng. Multi-touch techniques for exploring large-scale 3d astrophysical simulations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2213–2222, 2010.

[54] F. P. Gardiner and N. Lakic. Quasiconformal Teichmüller theory. (76), 2000.

[55] E. M. Geraghty, J. M. Boone, J. P. McGahan, and K. Jain. Normal organ volume assessment from abdominal CT. *Abdominal Imaging*, 29(4):1–11, 2004.

[56] E. Gobbetti, P. Pili, A. Zorcolo, and M. Tuveri. Interactive Virtual Angioscopy. *IEEE Visualization*, pages 435–438, 1998.

[57] D. Gotz and H. Stavropoulos. Decisionflow: visual analytics for high-dimensional temporal event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1783–1792, 2014.

[58] D. L. Gresh, B. E. Rogowitz, R. L. Winslow, D. F. Scollan, and C. K. Yung. WEAVE: A system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. *IEEE Visualization*, pages 489–492, 2000.

[59] T. Grossman and R. Balakrishnan. Pointing at trivariate targets in 3d environments. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 447–454, 2004.

[60] T. Grossman and R. Balakrishnan. The design and evaluation of selection techniques for 3d volumetric displays. *Proceedings of the 19th annual ACM Symposium on User Interface Software and Technology*, pages 3–12, 2006.

[61] S. Gu, R. Gupta, and I. Kyprianou. Computational high-resolution heart phantoms for medical imaging and dosimetry simulations. *Physics in Medicine and Biology*, 56(18):5845, 2011.

[62] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(8):949–958, 2004.

[63] Y. Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of motor behavior*, 19(4):486–517, 1987.

[64] I. Gutenko, K. Dmitriev, A. E. Kaufman, and M. A. Barish. Anafe: Visual analytics of image-derived temporal features–focusing on the spleen. *IEEE Transactions on Visualization & Computer Graphics*, (1):171–180, 2017.

[65] I. Gutenko, H. Peng, X. Gu, M. Barish, and A. Kaufman. Maximal area and conformal welding heuristics for optimal slice selection in splenic volume estimation. *SPIE Medical Imaging*, pages 97853V–97853V, 2016.

[66] I. Gutenko, H. Peng, X. Gu, M. Barish, and A. Kaufman. Maximal area and conformal welding heuristics for optimal slice selection in splenic volume estimation. *Proceedings SPIE Medical Imaging*, 9785:97853V–97853V–8, 2016.

[67] I. Gutenko, K. Petkov, C. Papadopoulos, X. Zhao, J. H. Park, A. Kaufman, and R. Cha. Remote volume rendering pipeline for mhealth applications. *SPIE Medical Imaging*, pages 903904–903904, 2014.

[68] J. G. Hagedorn, J. P. Dunkers, S. G. Satterfield, A. P. Peskin, J. T. Kelso, and J. E. Terrill. Measurement tools for the immersive visualization environment: Steps toward the virtual laboratory. *Journal of Research of the National Institute of Standards and Technology*, 112(5):257, 2007.

[69] R. S. Hamilton. The Ricci flow on surfaces. *Contemp. Math*, 71(1):237–261, 1988.

[70] S. G. Hart and L. E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in Psychology*, 52:139–183, 1988.

[71] P. Henrici. Applied and computational complex analysis. 3, 1988.

[72] M. Hermann, A. C. Schunke, T. Schultz, and R. Klein. A visual analytics approach to study anatomic covariation. *IEEE Pacific Visualization Symposium*, pages 161–168, 2014.

[73] J. D. Hincapié-Ramos, K. Özacar, P. P. Irani, and Y. Kitamura. GyroWand: an approach to IMU-based raycasting for augmented reality. *IEEE Computer Graphics and Applications*, 36(2):90–96, 2016.

[74] L. Hong, S. Muraki, A. Kaufman, D. Bartz, and T. He. Virtual voyage: Interactive navigation in the human colon. *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 27–34, 1997.

[75] W. Hürst and K. Vriens. Multimodal feedback for finger-based interaction in mobile augmented reality. *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 302–306, 2016.

[76] M. Hussain, K. Hassan, B. Yadav, and N. Usman. Anatomical Variations of Spleen in North Indian Population and its Clinical Significance. *Innovative Journal of Medical and Health Science*, 3(4):1–3, 2013.

[77] T. Isenberg. Interactive exploration of three-dimensional scientific visualizations on large display surfaces. *Collaboration Meets Interactive Spaces*, pages 97–123, 2016.

[78] T. Isenberg and M. Hancock. Gestures vs. Postures: "Gestural" Touch Interaction in 3D Environments. *Proceedings of The 3rd Dimension of CHI Workshop: Touching and Designing 3D User Interfaces*, pages 53–61, 2012.

[79] P. Issartel, F. Guéniat, and M. Ammi. Slicing techniques for handheld augmented reality. *IEEE Symposium on 3D User Interfaces*, pages 39–42, 2014.

[80] P. Issartel, F. Guéniat, T. Isenberg, and M. Ammi. Analysis of locally coupled 3D manipulation mobile device motion. *arXiv preprint arXiv:1603.07462*, 2016.

[81] S.-J. Jeong and A. E. Kaufman. Interactive wireless virtual colonoscopy. *The Visual Computer*, 23(8):545–557, 2007.

[82] M. Jin, J. Kim, F. Luo, and X. Gu. Discrete surface Ricci flow. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1030–1043, 2008.

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

[84] B.-J. Kim, Z. Xiong, and W. Pearlman. Low bit-rate scalable video coding with 3-D set partitioning in hierarchical trees (3-D SPIHT). *IEEE Transactions on Circuits and Systems for Video Technology*, 10(8):1374–1387, 2000.

[85] M. Kim and J. Y. Lee. Touch and hand gesture-based interactions for directly manipulating 3d virtual objects in mobile augmented reality. *Multimedia Tools and Applications*, 75(23):16529–16550, 2016.

[86] S. H. Kim, J. M. Lee, J. Y. Choi, K.-S. Suh, N.-J. Yi, J. K. Han, and B. I. Choi. Changes of Portosystemic Collaterals and Splenic Volume on CT After Liver Transplantation and Factors Influencing Those Changes. *American Journal of Roentgenology*, 191(1):W8–W16, 2008.

[87] T. Klein, F. Guéniat, L. Pastur, F. Vernier, and T. Isenberg. A design study of direct-touch interaction for exploratory 3d scientific visualization. *Computer Graphics Forum*, 31(3pt3):1225–1234, 2012.

[88] P. Klemm, S. Oeltze-Jafra, K. Lawonn, K. Hegenscheid, H. Volzke, and B. Preim. Interactive visual analysis of image-centric cohort study

data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1673–1682, 2015.

[89] R. Kopper, D. A. Bowman, M. G. Silva, and R. P. McMahan. A Human Motor Behavior Model For Distal Pointing Tasks. *International Journal of Human-Computer Studies*, 68(10):603–615, 2010.

[90] J. Krause, A. Perer, and H. Stavropoulos. Supporting iterative cohort construction with visual temporal queries. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):91–100, 2016.

[91] J. Kruger and R. Westermann. Acceleration techniques for gpu-based volume rendering. *Proceedings of IEEE Visualization*, pages 187–192, 2003.

[92] B. Laha, D. A. Bowman, D. H. Laidlaw, and J. J. Socha. A classification of user tasks in visual analysis of volume data. *IEEE Scientific Visualization Conference*, pages 1–8, 2015.

[93] B. Laha, K. Sensharma, J. D. Schiffbauer, and D. A. Bowman. Effects of immersion on visual analysis of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):597–606, 2012.

[94] P. M. Lamb, A. Lund, R. R. Kanagasabay, A. Martin, J. A. W. Webb, and R. H. Reznek. Spleen size: how well do linear ultrasound measurements correlate with three-dimensional CT volume assessments? *The British Journal of Radiology*, 75:573–577, 2002.

[95] F. Lamberti and A. Sanna. A streaming-based solution for remote visualization of 3d graphics on mobile devices. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):247–260, 2007.

[96] M. Levoy. Display of surfaces from volume data. *IEEE Computer graphics and Applications*, 8(3):29–37, 1988.

[97] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum (Proc. SGP 08)*, 27(5), July 2008.

[98] L. I.-K. Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pages 255–268, 1989.

[99] D. López, L. Oehlberg, C. Doger, and T. Isenberg. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3d data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(5):1616–1629, 2016.

[100] W. E. Lorensen and H. E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.

[101] C. Lundstrom, T. Rydell, C. Forsell, A. Persson, and A. Ynnerman. Multi-touch table system for medical visualization: Application to orthopedic surgery planning. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):1775–1784, 2011.

[102] J. B. Madsen and R. Stenholt. How wrong can you be: Perception of static orientation errors in mixed reality. *IEEE Symposium on 3D User Interfaces*, pages 83–90, 2014.

[103] M. Malik and M. E. Groller. Feature peeling. *Proceedings of Graphics Interface*, pages 273–280, 2007.

[104] S. Malik, F. Du, M. Monroe, E. Onukwugha, C. Plaisant, and B. Shneiderman. Cohort comparison of event sequences with balanced integration of visual analytics and statistics. *Proceedings 20th International Conference on Intelligent User Interfaces*, pages 38–49, 2015.

[105] N. A. Michels. The Variational Anatomy of the spleen and splenic artery. *American Journal of Anatomy*, 70(1):21–72, 1942.

[106] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman. Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2227–2236, 2013.

[107] J. Noguera and J. Jiménez. Visualization of very large 3d volumes on mobile devices and webgl. *20th WSCG International Conference on Computer Graphics, Visualization and Computer Vision*, 2012.

[108] J. M. Noguera, J.-R. Jiménez, C. J. Ogáyar, and R. J. Segura. Volume rendering strategies on mobile devices. *GRAPP/IVAPP*, pages 447–452, 2012.

[109] R. A. O'Reilly. Splenomegaly in 2,505 patients at a large university medical center from 1913 to 1995. 1963 to 1995: 449 patients. *The Western Journal of Medicine*, 169(2):88, 1998.

[110] S. Owada, F. Nielsen, and T. Igarashi. Volume catcher. *Proceedings of the 2005 Symposium on Interactive 3D Graphics and Games*, pages 111–116, 2005.

[111] C. Papadopoulos, I. Gutenko, and A. Kaufman. Veevvie: Visual explorer for empirical visualization, vr and interaction experiments. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):111–120, 2016.

[112] C. Papadopoulos, S. Mirhosseini, I. Gutenko, K. Petkov, A. Kaufman, and B. Laha. Scalability limits of large immersive high-resolution displays. *IEEE Virtual Reality*, pages 11–18, 2015.

[113] C. Parmar, E. Rios Velazquez, R. Leijenaar, M. Jermoumi, S. Carvalho, R. H. Mak, S. Mitra, B. U. Shankar, R. Kikinis, B. Haibe-Kains, P. Lambin, and H. J. W. L. Aerts. Robust radiomics feature quantification using semiautomatic volumetric segmentation. *PLoS ONE*, 9(7), July 2014.

[114] K. Pearson and L. N. G. Filon. Mathematical contributions to the theory of evolution. iv. on the probable errors of frequency constants and on the influence of random selection on variation and correlation. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 191:229–311, 1898.

[115] H. Peng, Z. Ruan, F. Long, J. H. Simpson, and E. W. Myers. V3D Enables Real-time 3D Visualization and Quantitative Analysis of Large-scale Biological Image Data Sets. *Nature Biotechnology*, 28(4):348–353, 2010.

[116] H. Peng, X. Wang, Y. Duan, S. H. Frey, and X. Gu. Brain morphometry on congenital hand deformities based on Teichmüller space theory. *Computer-Aided Design*, 58:84–91, 2015.

[117] A. Perer and D. Gotz. Data-driven exploration of care plans for patients. *CHI Extended Abstracts on Human Factors in Computing Systems*, pages 439–444, 2013.

[118] P. Prassopoulos, M. Daskalogiannaki, M. Raissaki, A. Hatjidakis, and N. Gourtsoyiannis. Determination of normal splenic volume on computed tomographyin relation to age, gender and body habitus. *European Radiology*, 7(2):246–248, 1997.

[119] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. http://www.R-project.org/.

[120] D. S. Raicu. Mining knowledge in computer tomography image databases. In *Multimedia Data Mining and Knowledge Discovery*, pages 487–508. Springer, 2007.

[121] R. G. Raidou, U. A. Van Der Heide, C. V. Dinh, G. Ghobadi, J. F. Kallehauge, M. Breeuwer, and A. Vilanova. Visual analytics for the exploration of tumor tissue characterization. *Computer Graphics Forum*, 34(3):11–20, Apr. 2015.

[122] C. Ratti, Y. Wang, B. Piper, H. Ishii, and A. Biderman. Phoxelspace: an interface for exploring volumetric data with physical voxels. *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, pages 289–296, 2004.

[123] B. Reitinger, D. Schmalstieg, A. Bornik, and R. Beichel. Spatial analysis tools for virtual reality-based surgical planning. *IEEE Symposium on 3D User Interfaces*, pages 37–44, 2006.

[124] P. Rezai, S. M. Tochetto, M. S. Galizia, and V. Yaghmai. Splenic volume model constructed from standardized one-dimensional MDCT measurements. *American Journal of Roentgenology*, 196(2):367–372, 2011.

[125] C. Rezk-Salama and A. Kolb. Opacity Peeling for Direct Volume Rendering. *Computer Graphics Forum*, 25(3):597–606, 2006.

[126] F. Robertson, P. Leander, and O. Ekberg. Radiology of the spleen. *European Radiology*, 11:80–95, 2001.

[127] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *IEEE Multimedia*, 13(3):24–30, 2006.

[128] I. Rössling, C. Cyrus, L. Dornheim, A. Boehm, and B. Preim. Fast and flexible distance measures for treatment planning. *International Journal of Computer Assisted Radiology and Surgery*, 5(6):633–646, 2010.

[129] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Signal Processing*, 26(1):43–49, 1978.

[130] J. Serrà and J. L. Arcos. An empirical evaluation of similarity measures for time series classification. *Knowledge-Based Systems*, 67:305–314, 2014.

[131] E. Sharon and D. Mumford. 2D-Shape analysis using conformal mapping. *International Journal of Computer Vision*, 70(1):55–75, 2006.

[132] S. Shen, G. L. DeNardo, A. Yuan, C. Hartmann-Siantar, R. T. ODonnell, and S. J. DeNardo. Splenic volume change and nodal tumor response in Non-Hodgkin's lymphoma patients after radioimmunotherapy using radiolabeled Lym-1 antibody. *Cancer Biotherapy and Radiopharmaceuticals*, 20(6), 2005.

[133] P. Song, W. B. Goh, C.-W. Fu, Q. Meng, and P.-A. Heng. WYSI-WYF: exploring and annotating volume data with a tangible handheld device. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1333–1342, 2011.

[134] V. Spitzer, M. J. Ackerman, A. L. Scherzinger, and D. Whitlock. The visible human male: a technical report. *Journal of the American Medical Informatics Association*, 3(2):118–130, 1996.

[135] M. D. Steenwijk, J. Milles, M. A. van Buchem, J. H. C. Reiber, and C. P. Botha. Integrated visual analysis for heterogeneous datasets in cohort studies. *IEEE VisWeek Workshop on Visual Analytics in Health Care*, pages 1–8, 2010.

[136] G. Taubin. A signal processing approach to fair surface design. *Proceedings 22nd Annual Conference on Computer Graphics and Interactive Techniques*, pages 351–358, 1995.

[137] D. Taubman. High performance scalable image compression with EBCOT. *IEEE Transactions on Image Processsing*, 9(7):1158–1170, 2000.

[138] R. J. Teather and W. Stuerzlinger. Pointing at 3d targets in a stereo head-tracked virtual environment. *IEEE Symposium on 3D User Interfaces*, pages 87–94, 2011.

[139] P. Tormene, T. Giorgino, S. Quaglini, and M. Stefanelli. Matching incomplete time series with dynamic time warping: an algorithm and an application to post-stroke rehabilitation. *Artificial Intelligence in Medicine*, 45(1):11–34, 2009.

[140] E. Tufte. Envisioning information. *Optometry & Vision Science*, 68(4):322–324, 1991.

[141] C. Turkay, A. Lundervold, A. J. Lundervold, and H. Hauser. Hypothesis generation by interactive visual exploration of heterogeneous medical data . *Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data*, pages 1–12, 2013.

[142] D. Valkov, A. Giesler, and K. H. Hinrichs. Imperceptible depth shifts for touch interaction with stereoscopic objects. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 227–236, 2014.

[143] D. Valkov, F. Steinicke, G. Bruder, K. H. Hinrichs, J. Schöning, F. Daiber, and A. Krüger. Touching floating objects in projection-based virtual reality environments. *EGVE/EuroVR/VEC*, pages 17–24, 2010.

[144] M. Wagner. *The Geometries of Visual Space*. Psychology Press, 2006.

[145] Y. Wang, X. Gu, T. F. Chan, P. M. Thompson, and S.-T. Yau. Brain surface conformal parameterization with algebraic functions. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 946–954, 2006.

[146] Y. Wang, L. M. Lui, X. Gu, K. M. Hayashi, T. F. Chan, A. W. Toga, P. M. Thompson, and S.-T. Yau. Brain surface conformal parameterization using Riemann surface structure. *IEEE Transactions on Medical Imaging*, 26(6):853–865, 2007.

[147] A. Wiebel, F. M. Vos, D. Foerster, and H.-C. Hege. WYSIWYP: What You See is What You Pick. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2236–2244, 2012.

[148] A. Wiebel, F. M. Vos, and H.-C. Hege. Perception-oriented Picking of Structures in Direct Volumetric Renderings. Technical report, Technical Report 11–45, ZIB, Berlin, Germany, 2011.

[149] K. Wongsuphasawat and D. Gotz. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2659–2668, 2012.

[150] K. Wongsuphasawat, J. A. Guerra Gómez, C. Plaisant, T. D. Wang, M. Taieb-Maimon, and B. Shneiderman. Lifeflow: visualizing an overview of event sequences. *Proc SIGCHI Conf Hum Factor Comput Syst*, pages 1747–1756, 2011.

[151] J. Xu, Z. Xiong, S. Li, and Y.-Q. Zhang. Three-dimensional embedded subband coding with optimized truncation (3-D ESCOT). *Applied and Computational Harmonic Analysis*, 10(3):290–315, 2001.

[152] E. M. Yetter, K. B. Acosta, M. C. Olson, and K. Blundell. Estimating splenic volume: sonographic measurements correlated with Helical CT determination . *American Journal of Roentgenology*, pages 1615–1620, 2003.

[153] D. Yim, G. N. Loison, F. H. Fard, E. Chan, A. McAllister, and F. Maurer. Gesture-driven interactions on a virtual hologram in mixed reality. *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces*, pages 55–61, 2016.

[154] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. Efficient structure-aware selection techniques for 3D point cloud visualizations with 2DOF input. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2245–2254, 2012.

[155] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. CAST: Effective and efficient user interaction for context-aware selection in 3D particle clouds. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):886–895, 2016.

[156] L. Yu, P. Svetachov, P. Isenberg, M. H. Everts, and T. Isenberg. Fi3d: Direct-touch interaction for the exploration of 3d scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1613–1622, 2010.

[157] Z. Zhang, B. Wang, F. Ahmed, R. Zhao, A. Viccellio, and K. Mueller. The Five W's for information visualization with application to health-care informatics. *IEEE Transactions on Visualization and Computer Graphics*, 19(11):1895–1910, 2013.

[158] G. Y. Zou. Toward using confidence intervals to compare correlations. *Psychological methods*, 12(4):399, 2007.