# Advanced Storytelling for Volume Visualization

**Andreas E. Lie**
**andreas.lie@student.uib.no**
**andreas.eknes.lie@gmail.com**
November 2009

Visualization Group
Department of Informatics
University of Bergen

# Advanced Storytelling for Volume Visualization


**Andreas E. Lie**
**andreas.lie@student.uib.no**
**andreas.eknes.lie@gmail.com**
November 2009

Supervised by Helwig Hauser

Visualization Group
Department of Informatics
University of Bergen

http://www.ii.uib.no/vis/teaching/thesis/2009-lie/

# Preface

Over the years, Storytelling has proven to be a very effective way of both retaining and transferring knowledge between humans. We find stories compelling, intriguing and stimulating, and these traits make storytelling as versatile as it is. Storytelling has manifested itself in oral and visual form, both proven useful in their way.

Today researchers struggle with presenting their results from the vast amounts of data collected, often because of the complex visualizations created, and the implicit difficulty of presenting these visualizations to both non-professionals and indeed professionals as well. In this thesis we show that the art of storytelling can aid in presenting complex visualizations, and that stories can be presented in a way that is both comprehensible and credible to the audience.

Storytelling as a form of communicating can help bridging the communication gap in a variety of situations. We aim at improving daily medical communication, exemplified by doctor to doctor communication, doctor to patient communication and medical documentation. Our visualization stories can also be used in an asynchronous form, as a collaboration tool between multiple participants. This can be achieved by exporting the stories and its structure, accompanied with the original data. Any partners can then recreate the stories, edit them and continue this refinement process of collaboration.

Our solution allows the creation of advanced visualization stories, that support recording, editing and playback at different levels of interactivity. The stories also include textual labels, annotations and recorded oral comments to aid the audience in interpreting the visualization results.

# Contents

# 1 Introduction

*There have been great societies that did not use the wheel, but there have been no societies that did not tell stories.*

- Ursula K. LeGuin

Everyday data is collected for researchers and other personnel to use. Time is of essence, and accuracy is often required to achieve the needed results and information. There are numerous types of data — they come in all shapes and sizes, and can represent share transactions from the stock market to weather prediction data. The knowledge about this data, what information lies hidden in the data is often not revealed until properly explored. What is often common for all types of data, is that at some point someone has to inspect the data, look after some obvious interesting features or the more not so obvious relations between the data attributes contained within the data. He/she will finally present the findings — the knowledge gained from this scrutiny, to an audience.

Nowadays data volumes is at a magnitude that make it impossible or even meaningless to inspect every data value. Often data attributes must be seen in context to one another to grasp the underlying meaning of the data. Data is often collected over a series of timesteps, by for instance measuring the same spatial positions at different stages of time. Visualization is a way of unveiling the information entangled in such data. It is often said that visualization strives to gain insight into data. By using visualization, users can achieve an overview of the data, and produce a specification how the data should be organized and how the results of the presented in the most meaningful way.

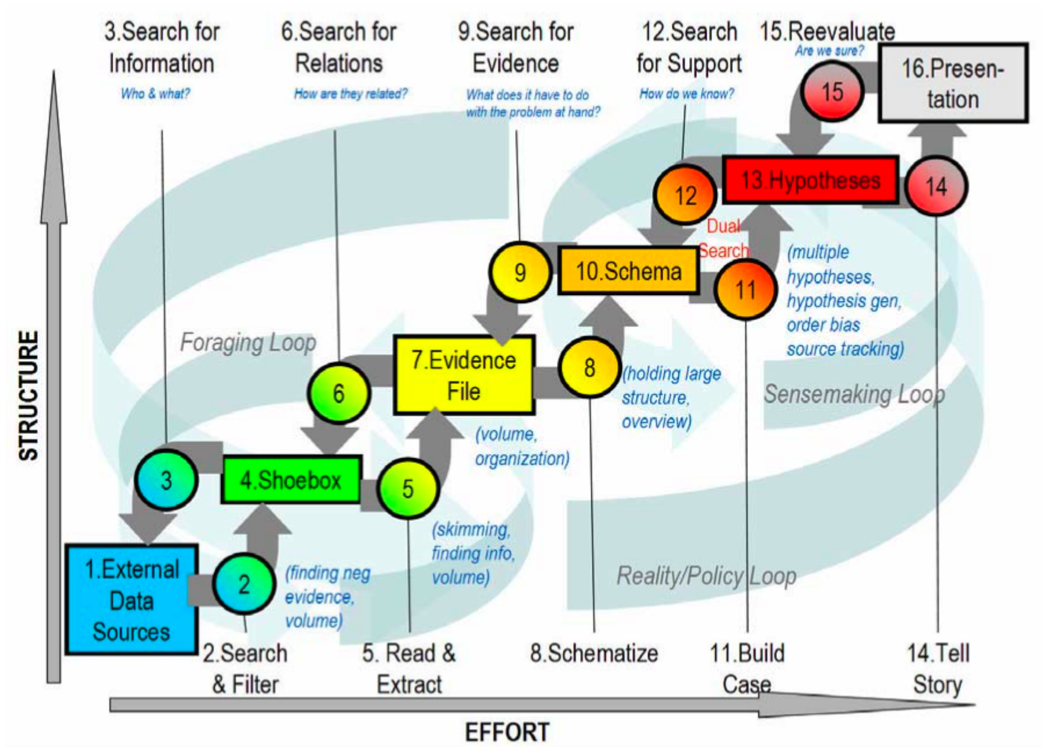


Figure 1.1: *The sensemaking process described by Pirolli & Card [PC05]. The Exploration process within visualization is analogous to the foraging loop, e.g. collecting evidence in a shoebox, while analysis is the consideration of this evidence. Ultimately any hypothesis or evidence found must be presented in one way or another.*

## 1.1 Visualization

Over the years, visualization has proven to be an efficient way of looking at and presenting sophisticated data. Visualization tools and techniques provide an overview of the data, make selections based on data of interest, and produce visualization results that present the data in a meaningful and understandable way. Within the field of visualization, historically there have been formed two separate communities who have for the same goals, but differ in their approaches and what they aim to achieve. These communities are the Scientific Visualization (SciVis) and Information Visualization (InfoVis). We will briefly describe the similarities and differences in the following paragraphs.

**SciVis**

This domain focuses on doing visualizations for interpreting and presenting research simulations, medical data, flowing patterns etc. The SciVis community strives to continually enhance visualization techniques to be able to explore and analyze data in better ways, and find ways to present those results with more impact on the audience. The data analyzed and visualized in SciVis is often time-dependent, i.e. the same spatial positions are measured numerous times according to a specified delta of time, which enable analysts to examine events over time. The data is often, but not always, grid-based, and often contains multiple attributes per spatial position. For example weather data from a region (over time) could consist of temperature, pressure and wind velocity, all sampled over a number of hours or days. Common visualizations done in the SciVis community range from simple graphs for statistics to more advanced 2D and 3D models devised from medical tomography data.

**InfoVis**

InfoVis is according to D. Keim [KMSZ06], the communication of abstract data through the use of interactive visual interfaces. Information Visualization (InfoVis) often entangles a variety of data and try to help with reasoning upon this data. Such data seldom have the spatial characteristics present in data visualized in SciVis. Common visualizations done in InfoVis comprise of statistical relations, presented as graphs, 2D plots to n-dimensional relational data viewed as parallel coordinates. Viegas & Wattenberg wrote a call to action [VW06] after the September 11 2001 incident that shocked the world. They argued the need for an asynchronous way of collaboration towards finding cues and evidence that might help avoid such situations. Better collaboration amongst partners is sought after in any domain, and indeed fields within visualization could benefit from better collaboration schemes.

In this thesis we focus mainly on SciVis, although much information contained here will also be readily compatible and no less true in InfoVis.

The visualization process can be roughly divided into three groups, according to Schumann & Müller [SM00], namely Exploration, Analysis and Presentation. These three distinct groups form a sort of work pipeline for visualizations.

**Exploration**

The exploration process is where a single user searches through the data at hand, without any prior knowledge about the data. The user is, in simple terms, in search for something meaningful, useful or otherwise surprising. The user produces a hypothesis based on the data explored, which later can be scrutinized further for either acceptance or rejection. According to Pirolli & Card this process can be

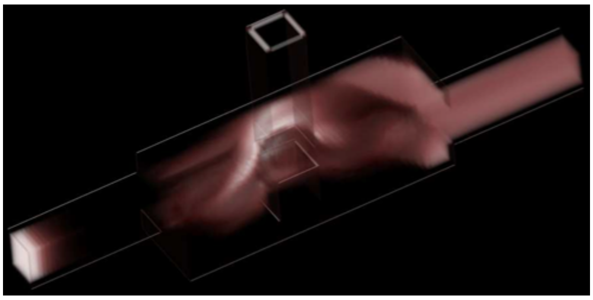

(a) *Flow patterns and temperature shown in a simulation dataset*

(b) *Chess set visualized according to DOI specification*

Figure 1.2: *Examples of DOI usage. Left side image visualizes the T-Junction dataset in the SimVis framework [PKH04]. The right side image depicts a chess set according to DOI specification. Pieces in focus have a high DOI value, while the blurred ones are kept as context data. Image courtesy of Hauser [Hau05].*

compared to collecting clips in a shoebox. Visualization users look after valuable bits and save them for further scrutiny later on. Figure 1.1 depict this process.

A way to cope with vast amounts of data is to categorize the data according to its importance. Focus+context, as described by Hauser [Hau05], is a very useful way of specifying how data should be visualized. If we elect to strip away all data not fully compliant to some certain criteria, the remaining data will often loose its inherent meaning. While we in many cases often look at single data items or specific numbers, much valuable information is often identified by inspecting the interrelated aspects of the data. Removing nearby data will make such inspection impossible. Focus+context visualization solves this by allowing the user to specify a DOI (degree of interest) value to each data attribute. This DOI value is used when deciding how data should be visualized. DOI values of 1 represents the data that we aim to focus on, while 0 means the opposite, data that are of no interest to us. By allowing a gradual decline (or rise) some data will be assigned a fracion of DOI. This specification will allow interrelated inspection, as well as spatial aiding when looking at 2D or 3D data. The figure 1.2 (right image), visualizes a chess set according to a DOI specification. The pieces that are in focus with respect to the camera, have a high degree of interest. Other pieces that contain a lower degree of interest appear blurred, but are kept in the visualization as context to the interesting pieces.

**Analysis**
When inspecting scientific data, researchers often have some ideas or correlations they want to inspect. They may already know some facts, certain criteria, that could indicate a specific event or situation. A technique named brushing is very suited to perform such analysis in a visualization environment. Brushing enhances regular focus+context visualization by enabling indirect selection of data based on other simple selections. These simple selections can be concatenated by logic to create more advanced selections. The data used in such analysis tasks are not limited to the usual 3D, or 4D if time is also taken into consideration, but can easily cope with data containing 10—40 attributes. An environment that can handle such vast amounts of data is the SimVis framework [PKH04]. Another way of identifying hidden relations or nuggets of information when looking at multidimensional data (more than 4 dimensions) is to use parallel coordinates. When visualizing with parallel coordinates, every attribute is plotted on the x axis of a diagram, while the y-axis represent the value or categorization according to that attribute. Lines are drawn between every point according to its specified value. See figure 1.3 for an example of such a visualization.

Another well known technique for visualization scientific data is direct volume rendering (DVR) [Lev88]. Volume rendering can be used on volume data that contain physical properties, such as density, to create stunning visualizations. Virtual rays are shot from the position of the camera, through the data, accumulating color and opacity along the way. An example of a volume rendering can be seen in Figure 1.3 where a stag beetle has been processed in a computed tomograph (CT) and visualized using a raycaster [BG05]. Volume rendering is a natural step from the 2D slice images used by medical personnel nowadays. By looking at such data in 3D, spatial relations are much easier to interpret. We will go into more detail regarding DVR later on.

**Presentation**
When data has been thoroughly explored and analyzed, there is an underlying need for presenting the data. Often the visualizations are of such nature that it is non-trivial to interpret them without any guidance in form of comments, labels or annotations. Care must be taken when presenting visualizations to avoid misinterpretations of the data being presented. We will go into further detail regarding presentation in the following section.

## 1.2 Visualization for presentation

Visualizations ultimately rely on the quality of the presentations to properly transfer the knowledge content contained in the visualizations. An expert could easily

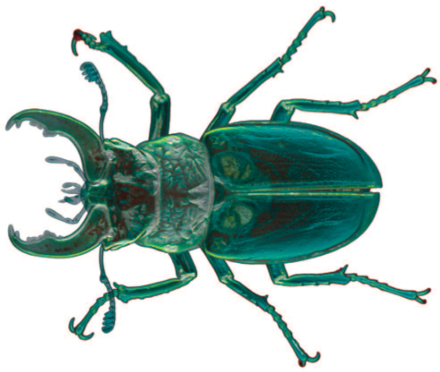

(a) *Direct volume rendering with semi-transparent regions, image courtesy of Bruckner et al. [BGK06].*

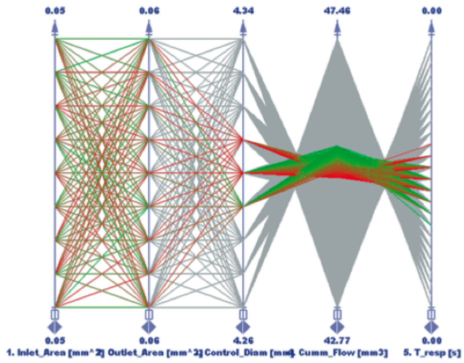


(b) *Parallel coordinates for n-dimensional data, image courtesy of Matkovic et al. [MJJ$^+$05].*

Figure 1.3: *A variety of visualizations*

explain and answer to questions regarding a visualization, but is very seldom that the creators of the visualizations themselves are present when having the visualizations presented to an audience. A well known alternative is through proper use of visualization aids, examples include color legends and scales, textual labels, annotations and commentaries.

According to Thomas & Cook [TC06], the most consuming parts of the visualization process, which turn data into presentation form, is the production, presentation and dissemination of visualizations.

**Production** refers to the process of producing the visualizations, i.e. looping through the visualization pipeline stages of exploration and analysis of data to a degree were results are interesting, intriguing and presentable. It is common knowledge that visually pleasing results triggers the audience in a better way than their counterparts.

**Presentation** is the presentation of the visualization results achieved from the production stage. Visualizations are made comprehensible, aesthetically pleasing and enriched with visualization aids like textual labels and annotations. It is vital that the presented visualizations are such that the message intended for the audience contains no ambiguity and is fully believable.

**Dissemination** is the stage were, according to Thomas & Cook [TC06], visualizations should be spread widely to allow sowing of seeds, in the metaphorical sense. Collaboration between analysts and visualization experts is key for proper dissemination. While much effort have been done to improve the foraging loop, the exploration and analysis in the visualization pipeline, there is still much potential to be realized in the sensemaking loop where presentation and collaboration takes place.

Although most researchers agree that presentation forms for visualizations are an obviously vital part for presenting the visualizations and the knowledge contained in them, little research have been done in that field the last years compared to exploration and analysis. The visualization community is urged by Thomas & Cook [TC05] to produce a more collaborative presentation scheme, with little or no ambiguity when presenting visualizations.

## 1.3 Storytelling

Storytelling is a special type of communication minded form. In storytelling, content (a story) is presented by a narrator to its audience. Storytelling has been around for as long as we know. It has been a crucial part of the preservation of knowledge, and as a learning aid in both old and modern time. The content of a story, can vary between fact and fiction, all depending on the narrators goal for the story. One such example are children that often experience storytelling at a very young age. In such a situation, the intent of the presenter or narrator, often can be morale, ethics or learning. One special feat regarding storytelling is that the audience emotionally attaches to the stories. Emotionally bound information is easier to remember, and thus a vital part of why storytelling have had such success. Nowadays storytelling is still used in various fields of work, ranging from industries to the medical domain. Everybody can take advantage of storytelling.

Storytelling has been documented by Orr [Orr86] in use at the Xerox corp amongst repairmen. They struggle with malfunctioning copiers that give little or no hint toward what the problem might be. To aid in this tedious task, they incorporate storytelling to help them in not only deducing what the problem might be, but also help them to remember what steps to be taken toward solving such a problem. Often they can be heard discussing what the problems with a specific copier might be sound like

> "Once, i was working on a copier similar to this model, where the ....".

The stories serve them also as a means of transfer for valuable knowledge, and can be used as a means of communication as Glassner [Gla99] states.

Another field where storytelling often is used is in the medical domain. Although most medical personnel can be considered medical expertise, they are certainly not experts in every field within the medical domain. T. Koch [Koc98] found that if nurses wrote their journals in a more similar fashion as stories, it would be easier for the next nurse, doctor or other to read and understand the emotional as well as the physical situation of their patient. Stories as such help bridging the gap between different (although similar) professions of doctors, and between doctors and patients.

## 1.4 Achieving better presentations

Researchers have come up with numerous ways to explore, analyze and present data. One common aspect concerning these tools, is that they were made by experts for experts to use. To be able to interpret data in such an environment, it is required that you have some prior knowledge about the data, the environment which you are about to utilize, and what you should be looking after. Although visualization is a fantastic way to show interesting features and/or trends in datasets, proper presentation of those visualization is crucial to certify that the audience fully comprehends the intended message.

One very interesting and elegant way of achieving better presenting visualizations made is to borrow the strengths from already existing presentation techniques. The oldest and maybe one of the more important presentation techniques is storytelling. Storytelling has been around for many many years, and its versatility has proven to be effective in many situations. Storytelling is capable of teaching practical tasks, ethics and moral, and also retaining history.
Taylor et al. [KST01] used learning material in animations in their studies. They could prove that the students viewing the animations containing descriptive examples had a much easier time to cope with learning the algorithms presented. Everybody has heard the phrase; a picture is worth a thousand words, and according to Gershon & Page [GP01], a story is should be worth a thousand pictures. By using the emotion and commitment found in storytelling, presentations of visualizations could really be taken to a higher level.

# 2 State of the Art

*Storytelling is an ancient and honorable act. An essential role to play in the community or tribe. It's one that I embrace wholeheartedly and have been fortunate enough to be rewarded for.*

- Russell Banks

## 2.1 Storytelling

In computer science, storytelling can be tightly compared to virtual reality (VR). Virtual reality is a fictional reality that exist only in computers, generated by computer software. Such virtual realities, can be perceived though as quite realistic. Realistic looking graphics, artificial intelligence and sound, will help heighten the user experience of such environment. These virtual realities are often modeled with a very specific purpose, such as having its user experiencing dangerous situations — like a flight simulator, without the risk of being harmed. Virtual environments can be thought of as stories that the audience are allowed to participate and interact in.

Interaction is an important feature of virtual realities. By granting the audience interaction possibilities, the whole environment is easier to grasp. According to A. Glassner [Gla99], to achieve a successful and believable virtual environments, the audience must be allowed to interact or participate as much as possible.

He also states that there is a bond between the narrator and the audience, a story contract. In this contract, the narrators responsibilities are the integrity of the story, and that there exist a sequence of plots that the characters involved in the story will follow. The audiences part of the contract is that they must grant the narrator to manipulate him/her intellectually, emotically and spiritually. This manipulation is more commonly known as *the willing suspension of disbelief.*

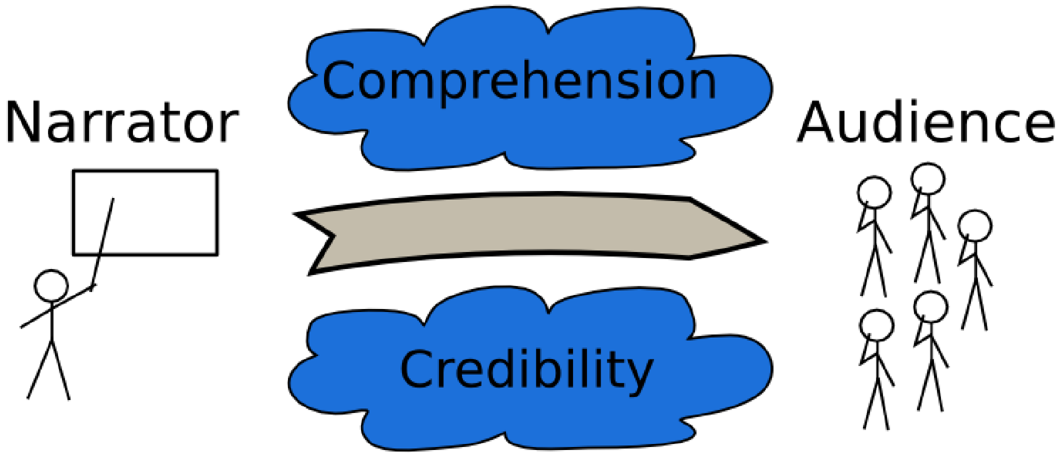


Figure 2.1: *The success of a story, how well the knowledge content is being transferred, is not only dependent on how well the narrator presents the content, but also at which degree the audience is willing to accept what is being told.*

The success of a particular story rests ultimately on two very important things, that the story is comprehensible, meaning that the audience fully understands what is being told to them, and that the narrator presents the story in such a way that is credible, i.e. that the audience accepts the story as it is being told. In figure 2.1, the process of a story being told is depicted. We will go into further detail regarding how to cope with comprehension and nurture credibility in section 3.2.3.

## 2.2 Presentation methodologies

Every day we are presented data and information in various ways. The simplest of these presentations use the form of still images. It is often said that an image is worth a thousand words, and while that analogy can be considered true, still images have certain constraints and drawbacks. According to Gershon & Page [GP01], still images are susceptible to uncertainties, and might require additional guidance in either written or oral form to clear them up. This includes labels, annotations and drawings of different kinds.

In the medical domain, doctors often use 2D slice based visualizations to examine and give diagnosis to patients. The 2D images used are often raw data from the acquisition devices, e.g. CT, MRI, used at hospitals. These devices sample their data from regions and output them in a slice based manner. An example of such 2D slice data can be seen in Figure 2.2. Such slice based medical data is considered difficult to interpret and fully understand for non-professionals, and

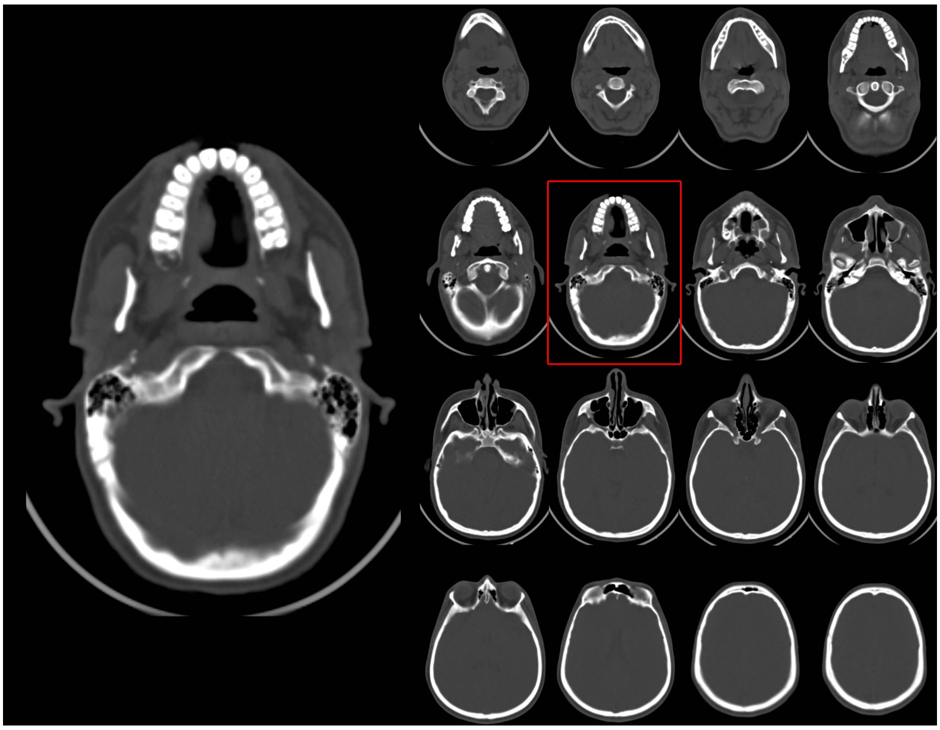

Figure 2.2: *An example set of 2D medical slice based data. Fully understanding the spatial relation is hard even though these slices are orthogonal to the main axis of the subject.*

are mostly interpreted by doctors.

The main problem with such 2D images is that they only show exactly what was sampled at that specific 2D plane. It is hard to interpret spatial relations meaning that it is hard to deduce where is the plane sampled or and along which axis is the plane located just by looking at the images. Without any spatial reference, it is hard to fully understand what the visualization at hand show, and how to analyze that information becomes futile.

Tietjen et al. [TMS+06] enhances such slice-based visualizations by introducing LiftChart widgets alongside slice-based visualizations that aid users in interpreting spatial relations. In figure 2.3 such a LiftChart widget is shown to the left, giving users valuable information such as the extent of the structures shown in the slice, the actual slice position itself, and the maximum extent (described by halos in the 2D slice) of the selected interesting structures.

Data with an inherent spatial relation, such as the 2D slices, are often the source for visualization techniques to create stunning 3D visualizations. Two very prominent techniques that result in 3D models are the raycasting as decribed by

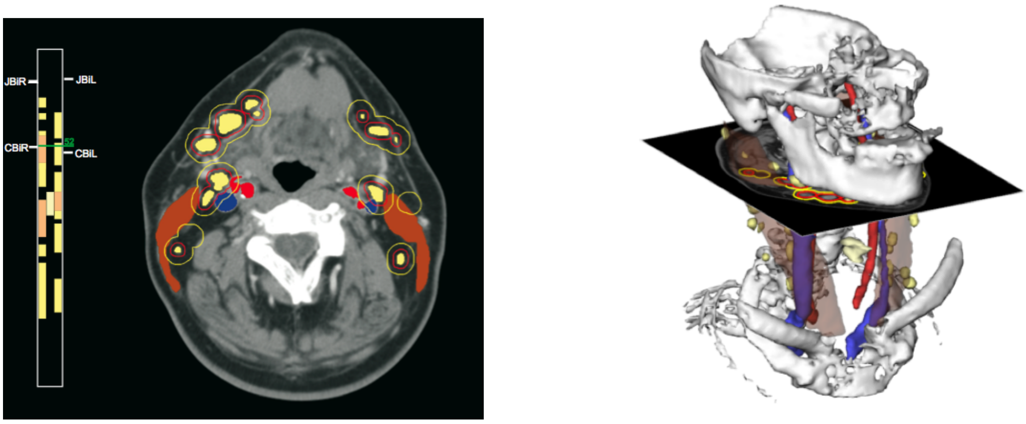


(a) **Liftcharts alongside a colorcoded slice image showing possible pathologic structures.**

(b) **3D visualization used in conjunction with the 2D slice in to aid in the interpretation of spatial relations**

Figure 2.3: *Slice images enriched and supported by adding volume rendering, color coding and liftcharts. Image courtesy Tietjen et al. [TMS+06].*

Levoy [Lev88] or marching cubes as described by Lorensen & Cline [LC87]. The resulting 3D visualizations often help with respect to occlusion by focussing on interesting areas. In figure 2.3b a 3D visualization is created from the slice-data where the focus has been set on the bones or higher density data. As seen in the figure, combining slices and 3D visualizations is a very powerful technique, since is is easy to understand the spatial relation of the structures described in the slice visualizations when you have such an overview 3D visualization alongside the original data to support optimal conditions for comprehension. The subject of comprehending data being presented will be discussed in section 3.2.3.

An example of a series of DVR visualizations used to present scientific data, can be seen in figure 2.4. The visualizations shown in the horizontal rows are rotated versions of the data being presented. In the vertical columns the variation is the amount of structures shown in the visualizations. By allowing the audience the option to vary between the different visualizations, it is easier to perceive the relation between the different anatomy structures, in addition to the fact that a broader audience can be targeted for such presentations as well (some might be interested in the surface, others in the different anatomy).

These 3D visualizations are of great help to the audience with respect to spatial references. Takahashi et al. [TFTN05] found methods for finding optimal viewpoints with respect to 3D models. By calculating optimal viewpoints for volumetric data, the usefulness of a visualization might be increased since it might be an optimal viewpoint for an increased number of viewers.

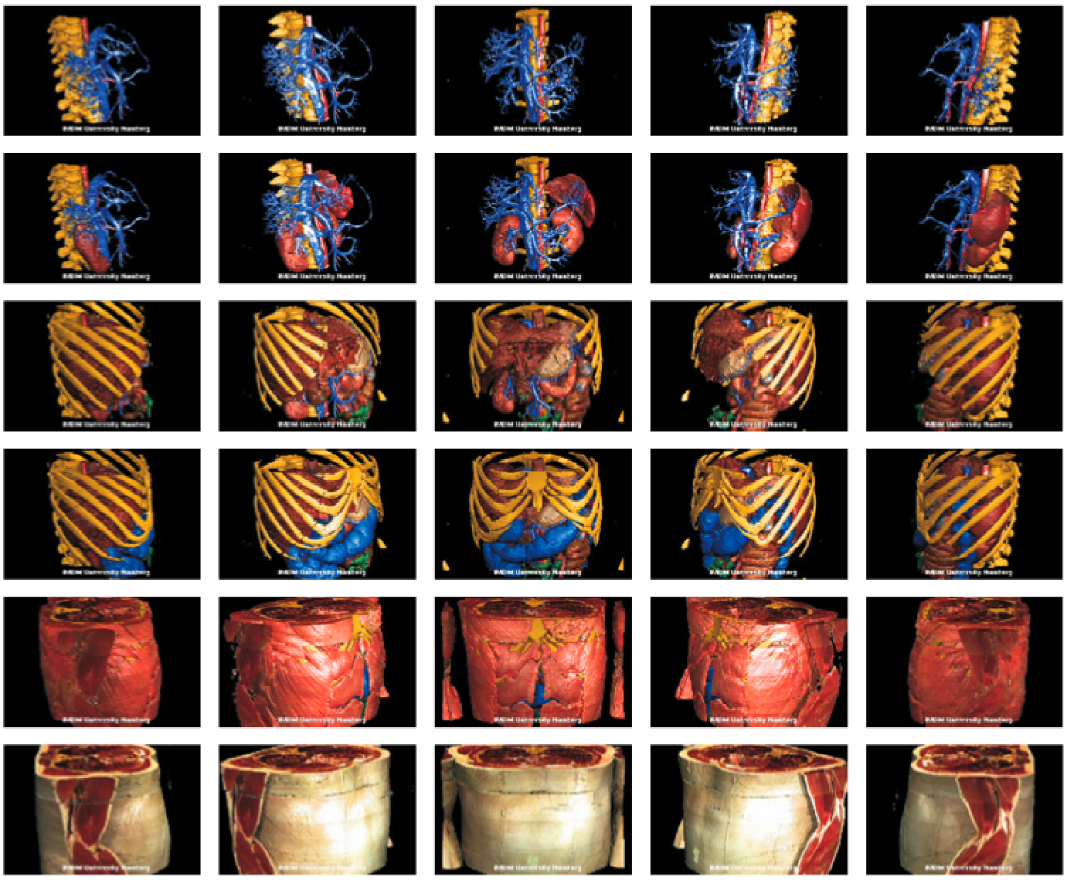

Figure 2.4: *A quadrant of images. Horizontal images variations are governed by rotation, while the vertical images emphasize the amount of information (structures) shown. Image courtesy of Schubert et al. [SPP+99].*

Producing visualizations that are more comprehensible, that have enriched content (for example liftcharts) and might include better viewpoints is according to Van Wijk [vW05] a way to increase the value of a visualization. He describes the value of a visualizations according to a set of parameters which includes its reusability, the amount of content shown, how many users can fully appreciate the visualization at hand, etc.

A natural advance with respect to presentation methodologies when having images of 2D/3D visualizations is animation. Animation is in fact just a series of slightly different images shown iteratively on the screen. Butz [But97] presented a system named CATHI that generated animation clips based on pre written user input. This input is considered a scripting language, a language where the user specified via different commands or rules how the view should be setup, where

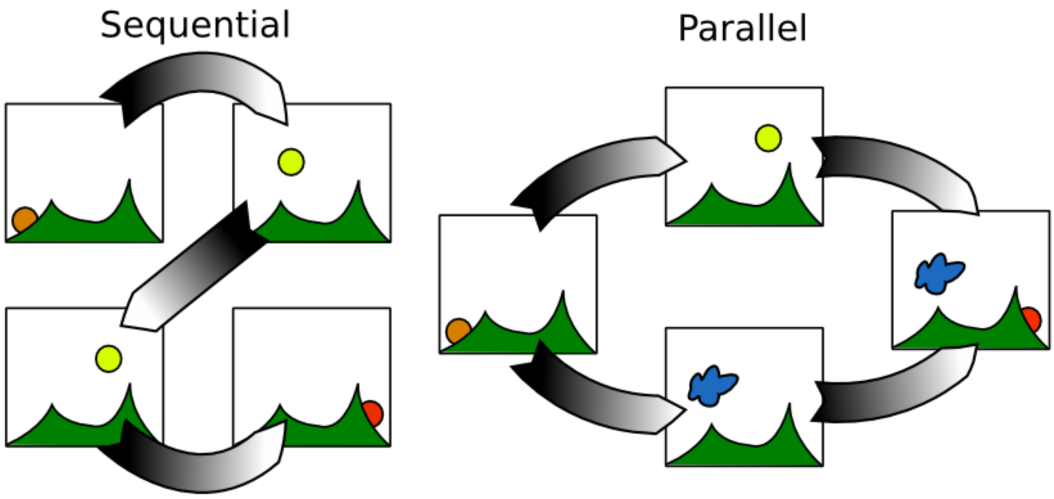


Figure 2.5: *Sequential vs Parallel sequences. Parallel sequencing of visualization events allow multiple changes to occur at the same time.*

lighting sources be placed, how object motions should be done etc. After the user has written the actions forming a short story, the system would in turn generate an animation based on the descriptive information given. The script syntax of CATHI is tightly correlated to terms and commands used by program developers within the computer graphics domain. An example of such command is;

(translate-object 1.0 7.0 :object "shaft-1" :startposition(71.0 79.5 30.0))

In this example the intended object named shaft-1, would be translated from its starting position, to a specified position XYZ (71.0 79.5 30.0), starting after one second, lasting a timespan of 6 seconds creating a gradual translation. One of the main advantages of this scripting language by Butz, was the possibility of doing actions in parallel. Instead of performing all the actions in a sequential order, for example move an object, do a certain rotation and finally scale the object for closer inspection, the system permitted parallelization of such actions, doing a gradual motion consisting of, as stated above, moving while rotating and scaling the object. The start and end of the sequential and parallel actions would still be the same. An example showing parallel events can be reviewed in Figure 2.5. Such parallelization of actions not only save time, but will make the audience more aware and receptive to its purpose because they immediately will focus more on the simple sequential stages of the animation, and consider the parallel actions as context or tension buildup.

An easier and thus more useful scripting language, with respect to its intended users, was proposed by Muehler et al. [MBP06]. They mapped high level commands to single or groups of low level commands, achieving a more natural way

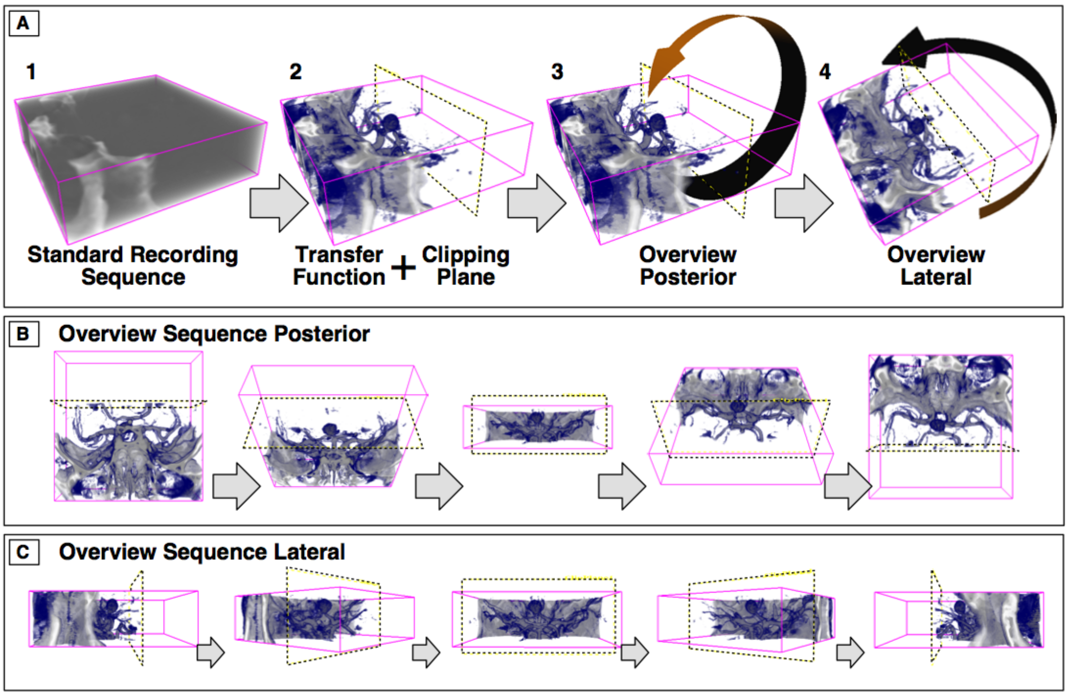


Figure 2.6: *Standardized camera paths used to allow quick and easy exploration and overview of complex 3D datasets. Image courtesy of Iserhardt-Bauer et al. [IBHE+01].*

of expressing the actions. Such animation scripts are often tailored to a specific dataset, making the animation worthless if the user exchanges the dataset with another one. Muehler et al. made their animation scripts reusable due to the fact that they are written in a high level language. Therefore it would require a lot less work to adapt existing animation sequences to fit any dataset. In their medical education software, the LIVERSURGERYTRAINER, they have used such pre-made animation scripts to better convey the lessons to the intended audience.

Iserhardt-Bauer et al. [IBHE+01], designed a system for an automatic 3D documentation running as a web service. This service offers the users automatic generation of video sequences, presenting data using standardized camera paths. By using pre-generated structures such as pre-defined camerapaths the time usually needed for generating visualizations for simple exploration will be significantly reduced . Figure 2.6 show the pre-defined standardized camera paths used when producing the video sequences through their webservice.

Zooming can also be used to emphasize important structures of regions when visualizing 3D datasets. Preim et al. use in their system ZOOMILLUSTRATOR [PRS97] fisheye zooming to focus on interesting structures.

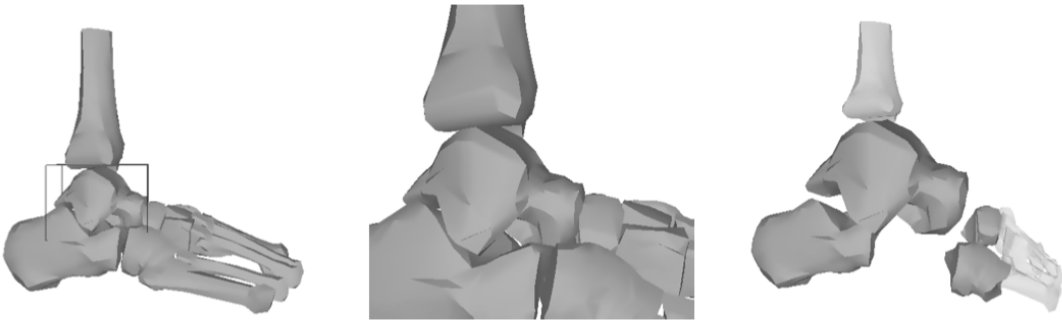

Figure 2.7: *Original 3D model shown left. The middle image shows conventional full-zoom, the right image shows 3D fisheye zoom as used in ZOOMILLUSTRATOR. Image courtesy of Preim et al. [PRS97].*

The ZOOMILLUSTRATOR effectively combines 3D models with textual labels. These labels allow users to interact with the system, selecting which parts of the visualization should be in focus and which parts to be shown as context data. This is achieved by clicking on the textual labels. The focused regions would then be emphasized by the wide-angle fisheye zoom, while the context regions would be made less prominent. This type of zoom focus can be seen in Figure 2.7, where conventional zoom is in the middle, and fisheye zoom in used in the right image. The motivation for combining 3D models with interactive textual labels was to produce a more interactive environment than regular textbooks could.

A research area where fisheye-zooming is used extensively is virtual endoscopy. Regular endoscopy involves a patient being examined by the use of an endoscope. Such examinations can be very uncomfortable for patients and an interesting alternative, virtual endoscopy, has come up. In virtual endoscopy, the patient is examined using data from CT, MRI, or another similar acquisition device. The resulting 3D datasets can then be explored and visualized on a computer instead of using the fiber and camera endoscope. Tiede et al. [TvSGSH02] uses panorama views to avoid the distortion caused by such extreme wide-angle views.

## 2.3 Labelling in visualization

The use of color is an essential part in visualization, and examples of the use of color include the depiction of importance, similarities or reflecting underlying data values. Interpreting the contextual meaning of color is difficult without an adjoining colormap. These maps visually delimit the range of colors used, and add descriptions to allow interpretation. The colors used in such maps should be chosen with care. So called rainbow color legends have long been discouraged as

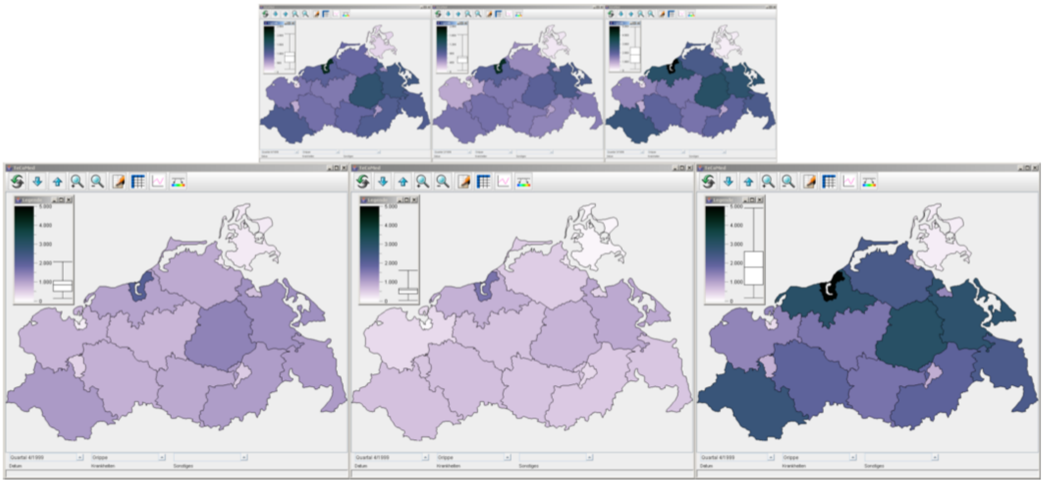

Figure 2.8: *The figure depicts three timesteps of quantitative disease data. The upper maps results from a color coding which is misleading when visualizations are compared. They also fail to emphasize the differences in the visualizations. The lower visualizations have a common color legend, which allow correct comparison of the three visualizations. Adjoining Box-whisker plots show the statistical distributions to aid interpretation of the colors and related values. Image courtesy of Schultze-Wollgast et al. [SWTS05].*

color mappings due to their misleading nature as described by Rogowitz and Treinish [RTB96]. They also point out several pitfalls to avoid when using color scales in visualization and describe a set of rules to achieve better visualizations called PRAVDA (Perceptual Rule-based Architecture for Visualizing Data Accurately). Color legends or colormaps should also be adjusted according to their statistical properties according to Schulze-Wollgast et al. [SWTS05]. Figure 2.8 shows how such distrubutions can contribute to achieve better visualization images.

Scientific diagrams and learning material are often of complex nature. They often require explanation and pinpoints to assist user interpretation and guidance. Denisovich clarifies in his work [Den05], how visualizations could be enriched by handdrawn marks. Annotations as such are vital when presenting visualizations, but should be used with caution with respect to occlusion. Textual labels have been used by illustrators for a long time to reduce complexity and aid spatial awareness in similar graphics. Hartmann et al. have done comprehensive studies on how to incorporate and effectively use labels in visualizations. Labels can be either of internal or external character. See the work of Hartmann et al [HGAS05] and Götzelmann et al [GAHS05] for metrics controlling label placement and discussion on placement algorithms for both internal and external labels. Internal labels can

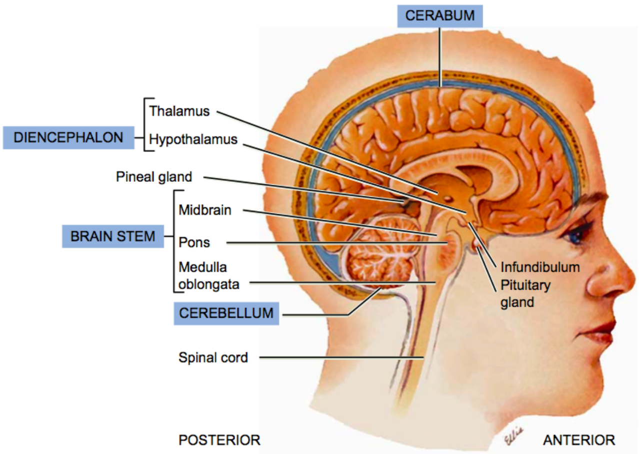


(a) **Contextual label grouping.**

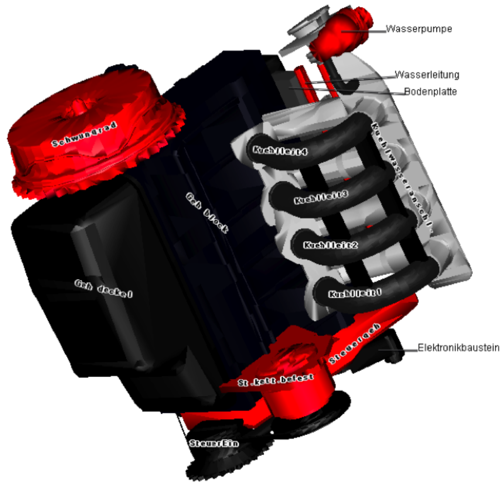


(b) **Internal and external labelling**

Figure 2.9: *Labels are contextually grouped in figure 2.9a. This eases interpretation of labels, by allowing meta-labels. Figure 2.9b depict internal and external labels. Image courtesy by Götzelmann et al. [GHS06b] [GAHS05].*

be seen in figure 2.9b, where an engine is depicted.

Internal labels are often placed on homogenous surfaces and is aligned along the main axis of the object it relates to, to minimize the amount of occlusion. Another feature of internal labels, is that their curvature can give helpful visual cues on how an object's geometry is presented in the work of Maass & Döllner [MD07]. External labels often consist of three parts. The label itself containing some descriptive information, the label target, a spatial position often marked by a small circle or square, and a line connecting them often described as anchor line. External labels are dependent on either the use of anchor lines to pair the label with its spatially counterpart, or use color coding (same color on label and circle/square) to have a correlation between the label and its spatial position. Labels should be placed in such manner that they minimize occlusion. E.g. that their anchor lines occlude as least as possible of the visualization at hand, or each other. And that the label anchor lines do not cross other anchor lines, that would be a source of confusion to the viewer.

Labels should also be grouped according to their contextual meaning as described by Götzelmann et al. [GHS06b]. This enables meta-labels, which are labels about labels, which eases interpretation of labels. See figure 2.9 for meta-labels and labels grouped according to their relation.

Labels are usually placed according to a bounding box or silhouette that circumferences the graphical regions. Figure 2.10a depicts a silhouette of control

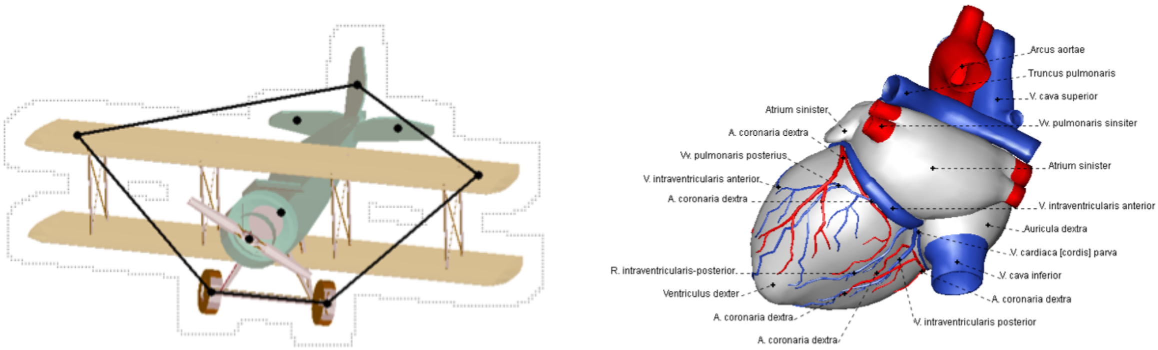


(a) **Silhouette used to place labels. Image courtesy by Götzelmann et al. [GHS06a].**

(b) **Labels placed by silhouette layout. Image courtesy by Ali et al. [AHS05].**

Figure 2.10: *In figure 2.10a the convex hull of a set of control points for the object is measured to determine where the object's silhouette is defined. This is used to place labels esthetically and efficiently. Figure 2.10b depicts the outcome of such silhouette label layout.*

points used for label placement. In figure 2.10b labels have been placed around a heart according to its silhouette.

Metagraphics such as labels are easy to place on still images. Nowadays animation and rotating 3D graphics are more and more commonly seen. Calculating the label placement with respect to each frame is possible, but will very likely result in very jittery labels when images change. Calculating label positioning metrics based on several images, allows labels to be positioned such that they need not move much, and move as if they were spring based, which Ali et al. [AHS05] described. Animation sequences are especially suited for such calculations, where the images appear in fixed order. The labeled objects spatial positions are measured in all images, which is then used to calculate the optimal placement and eventual movement for the labels, as presented in the work of Götzelmann et al. [GHS07].

One very limiting aspect of textual labels and annotations is the available screen space for use. Once the amount of space used becomes high, the textual labels appear as visual clutter. Another challenge with respect to the visual clutter, is the inherent limit of text the labels can display. Using a variety of different modalities have been proposed by Schär and Krueger [SK00]. Using both visual and auditory modalities, valuable screen space can be saved by allowing verbally recorded comments and discussions to be incorporated alongside the visualizations. In their work, they state that *Voice guides the attention and gives value to the*

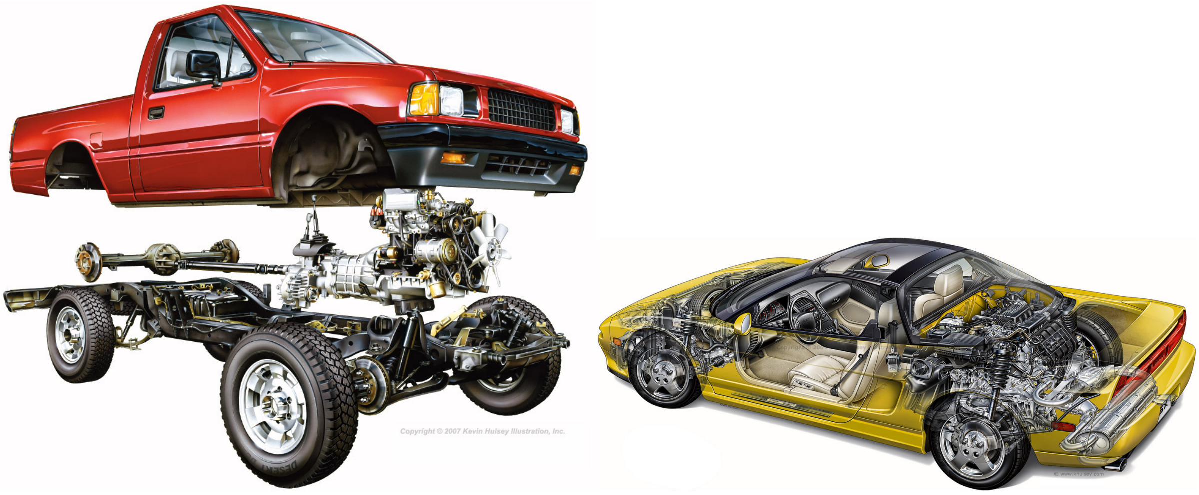

(a) **Exploded view** (b) **Ghosting technique**

Figure 2.11: *Two different techniques applied to avoid occlusion of interesting structures. In 2.11a the chassis of the truck is separated and displaced from the engine, to allow the engine full scrutiny. 2.11b displays an illustrative technique where uninteresting parts are made semi-transparent to allow underlying structures to become visible. Image courtesy of K.Hulsey [Hul].*

*learning content*, which is very similar to textual labels. They also point out that overloading one sense modality can cause reduced attention, thus, a balance between visual and auditory information reduces the cognitive load.

## 2.4 Illustrative visualization for presentation

One of the hardest problems regarding visualizations is how to cope with occlusion. Occlusion occurs when the object or area of interest is fully or partially covered by other structures. There are several different techniques that reduces the problem concerning occluding structures, or at least minimizes the information loss caused by the occlusion. The most simple way of dealing with occlusion would be to remove or disregard any data in front of the data in focus, with respect to the line of sight. This technique is referred to as view-aligned clipping or slicing. Removal of data can cause loss of context or make visualizations more difficult to interpret.

More advanced slicing is proposed by Weiskopf et al. [WEE03], where arbitrary objects can be used to clip volumetric data. To avoid loss of contextual information, volumetric objects can be separated from each other, or virtually exploded as Bruckner [Bru06] presented. This allow otherwise occluded structures to be seen along side with its contextual shapes. Figure 2.11a is an example of an exploded

(a) **Context**

(b) **Gecko**

Figure 2.12: *Figure 2.12a is a visualization where lighting-driven feature classification is used to illustrate the visible human male dataset. Image courtesy by Bruckner et al. [BGK06]. In figure 2.12b a cone is calculated from the focused region which alter the transparency with respect to the view direction. Image courtesy by Viola & Kanitsar [VK05].*

view, where the chassis of the truck is separated and displaced from the engine, allowing the engine to be scrutinized.

By making cutouts of the data with respect to the occluding structures, better cutaways can be achieved. Li et al. present interactive cutaways [LRA+07], where the occluding objects are cut to make focus regions visible, and at the same time preserve some context from the occluding structures. In figure 2.4 such interactive cutaways are done with respect to different focus parts.

In educational books, peel away structures are often used to allow users to examine what lies underneath. The user would bend a piece of paper, depicting the outer structures of an object, away to reveal structures underneath. Within the field of visualization, Birkeland [Bir08] uses similar techniques known as view-dependent peel-away to visualize volumetric data.

Instead of warping the volumetric data with explosion, slicing and peel-aways, changes to the opacity determines how visible the structures are. This ghosting of un-interesting otherwise occluding regions allows the focus to be visible, while maintaining spatial comprehension by keeping most of the surrounding data visible and in place. The ghosting technique is frequently used by illustrators, and can be seen in figure 2.11b where the chassis of the car is made transparent to

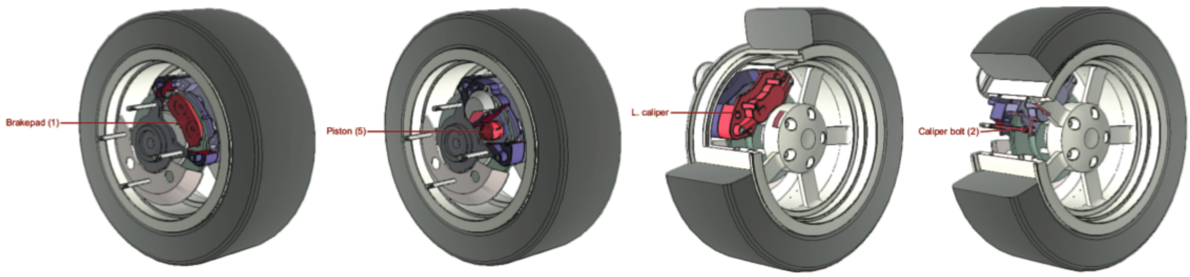


Figure 2.13: *Cutaways are done perpendicular to the occluding objects primary axis and with respect to nearby parts. This allows the user to easily mentally complete the structures removed, and to comprehend how the spatial relations are. Image courtesy by Li et al. [LRA+07].*

allow the engine and internal parts to be visible. Notice that we still concieve that the chassis is intact and how it is formed by interpreting the telltale signs from the illustration. In the work by Voila et al. [VK05], a cone is projected from the focus regions toward the viewer, where the content inside is made transparent. This conserves spatial understanding and allow the user to see the structure of the semi-transparent content. Figure 2.12b displays the Leopard Gecko dataset where the focus region is set to the geckos intestines. A similar method, namely Context preserving rendering, which can be seen in Figure 2.12 is presented by Bruckner et al. [BGK06]. Here they use a lighting-driven feature classification to define how opacity should be handled. Haidacher [Hai07] presented a novel way of specifying importance in multimodal volumetric rendering. His work combines several modalities and tweaks visibility by granting more importance to one of the modalities than the others. Bruckner & Gröller [BG07] enhance depth perception and makes objects more prominent by adding flexible halos in their volume rendering visualizations. By using halos, depth relations are also easier to perceive and objects become more noticeable for our visual system.

## 2.5 History Mechanisms

Creating visualizations is a lengthy process consisting of fine tuning numerous parameters. In the process the user might iterate through many variations of visualizations, applying more or less the same changes to the them. According to the work of Pirolli and Card [PC05], this process of identifying interesting nuggets of information is known as the foraging loop, se figure 1.1. The user searches for nuggets within the data, an exploration, which later can be presented and vali-

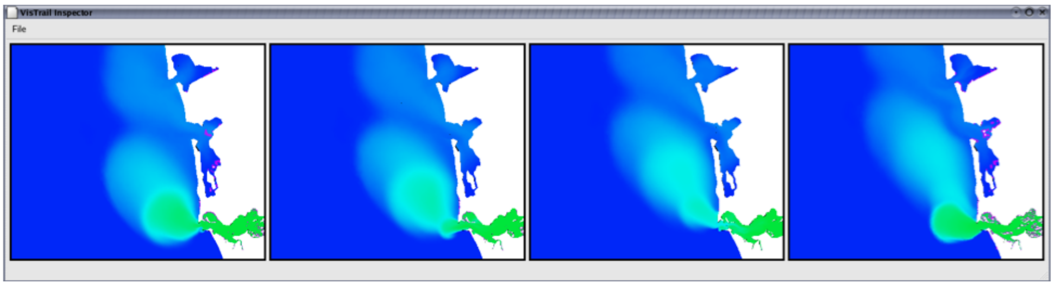


Figure 2.14: *Different timesteps of the Columbia river forecast. Visualizations contain only small subset of difference, and computational effort can be saved. Image courtesy by Bavoil et al. [BCS+05].*

dated in the sensemaking loop. When deriving visualizations from the data, often the user performs similar changes to the dataset for a long time, before changing a small subset of parameters altering the visualization in the end. Much redundant work is performed in these instances. Bavoil et al. solve this issue in their work [BCS+05], by creating visualization trails, or vistrails. These visualization trails contain the needed information to build a visualization up to a specified timestamp, making the process more effective by avoiding a complete recalculation of attributes for each visualization made. Using such visualization trails is very effective when considering time-varying datasets like meterology. Figure 2.14 shows forecast images from the Columbia river, at different stages of time. These visualizations are very similar, and the framework saves much computational effort by creating a pipeline specification it can reuse for creating visualizations with containing only small differences.

According to Schneiderman [Shn96], a common technique used to create visualizations is "trial and error". By keeping history of actions to support undo, replay and progressive refinement, this technique can be avoided. These action and history mechanisms are also described in the works of Kreuseler et al. [KNS04]. Kreuseler et al. define a datastructure they call a history tree which all history functionality is described in. This tree stores the visualization, and its branches, making it possible to explore and present the different visualizations at a later stage.

## 2.6 Interviews with professional medical personnel

During our research we contacted and conducted interviews with three medical professionals to gain some insight how storytelling could be used in their daily work. The three participants were

Professor Arvid Lundervold, institute of biomedicine.
Professor II Odd Helge Gilja, section of gastroenterology.
Associate Professor II Martin Biermann, the section of radiology.
Although some of their daily work include examinations, most of the time are still spent on analyzing and interpreting the results, and presenting these results to fellow doctors or to patients. We were interested to learn more on their thoughts on using storytelling and visualization stories as a tool to enhance and increase effectivity of medical communication. Among the questions we queried the participants were;

*How could your work benefit from storytelling?*
*How much time are you willing to devote?*
*Do you consider storytelling worth the investment of time?*
*Do you see any use for template stories? E.g comparative, overview & zoom etc.*
(The complete interview questionnaire can be found in appendix A).

**A compiled summary of the interviews follow.**
Storytelling is undeniably very suited for teaching and learning. Making visualization stories presenting material and relations is both pedagogic and very entertaining. It could also very well be used for the medical epikrisis, which is a documentation containing medical history, findings, procedures planned, diagnosis and treatment. These epikrisis can take much time to write, and usually many doctors and medical personnel will read through this documentation over time.
It is a increasing trend that patients have the right to read, and would like to review their own patient history. Visualization stories and storytelling can be of great help when reviewing such material, explaining complicated relations and having the possibility to pause and halt the stories to avoid being overwhelmed by new information.
New modalities with respect to radiology, e.g. PET, PET-CT, etc., can be hard to interpret without co-references. Spatial positions often hard to interpret, data is often complex and require training to interpret correctly. Stories can help present such data, adjoined by labels, annotations and oral commentaries to support comprehension.

Visualization stories for specific trauma or pathologies are often very similar. Stories might require only small changes to fit the need for other patients, thus templating and story re-use is very probable.
At hospitals, it is common that case demonstrations are held several times a week. These demonstrations often have an audience of 40 persons or more. In these situations the ratio for investment of time into stories could be quite high. These visualization stories could enhance such demonstrations while also encourage discussion and can make adjustments and new investigations on the fly.
Comparative visualizations can be very helpful toward analysis and learning. Viewing comparable data side by side, such as young vs old, healthy vs sick, male vs female and sickness spreading over time eases the learning curve and is in general very helpful for analysis.
Storytelling undoubtedly helps the audience to look at data and visualizations in the right context, and is very user friendly as such.

# 3 Storytelling

*Visual storytelling of one kind or another has been around since cavemen were drawing on the walls.*

- Frank Darabont

A story can be defined as a quantity of knowledge that is transfered from one person to another. There are no requirements toward what the story's content should be, and for what purpose the story is told. The content can always be categorized as either fact or fiction.
During the history, stories have been used for many purposes. Some stories aim at teaching the audience ethics or morals, manageable by both facts and fictional stories, others aim at the transfer of knowledge. The latter type of stories were responsible for maintaining history from generation to generation before more permanent techniques as painting and writing became available.

Stories exist in many forms, some are vocally transfered, others are in written form. Neither form excludes the other. One of the earliest forms of known written or visual stories can be seen in figure 3.1. This cavepainting is approximately 6000 years old, and depicts a tribe together with a herd of elephants.

Central in a story is the plot. The plot of a story governs how a story will unfold. This plot can be understood as a series of predetermined steps the story must take towards the goal. These predetermined steps are intertwined, they are tightly connected to each other. From one step to the next there is a gentle change of the story, such that the audience can perceive and fully understand the changes being done. These predetermined steps are the important states during the story, without them the story would loose its inherent meaning.

A vital part of a story is also the narrator and its audience. Without them, the life of a story would come to an end. A narrator is the teller of the story. He (she) presents the story in the most comprehensible way possible to an audience. The

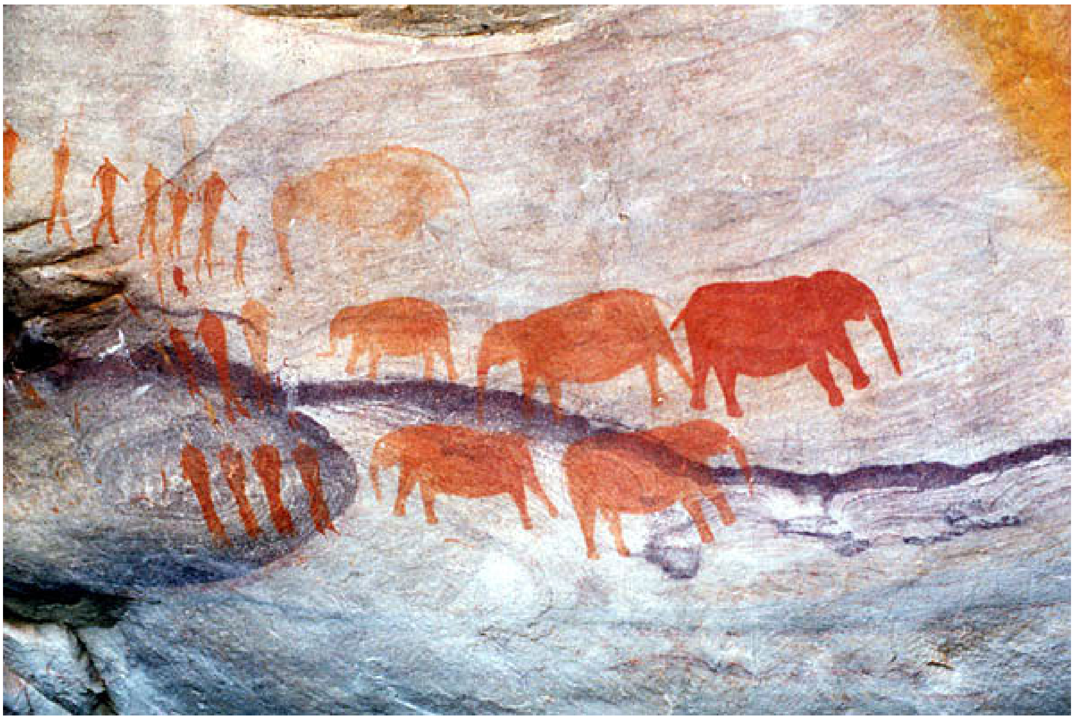

Figure 3.1: *Cavepainting found in the Cederberg mountains near South Africa. Estimated to be a 6000 year old visual story, it depicts a tribe with spears nearby a herd of elephants. Image courtesy of Ventureforth [Ven].*

job of the audience is to allow the narrator to manipulate him/her intellectually, emotically and spiritually. If the audience do not believe the story, find the information presented credible, the transfer of knowledge during the story has failed. Care must be taken to allow a story to be both comprehensible (understandable) and credible (believable).

## 3.1 Storytelling in different fields

Storytelling is also present in many different fields and applications used today. The potent characteristics of storytelling includes not only the transfer of information from the narrator to the intended audience, but also the fact that stories can be used as an aid for collaboration between individuals as well. According to Gershon & Page [GP01], there is a need for knowledge transfer in every community, and storytelling can be a very powerful asset to aid in information collaboration of any sort.

The information visualization genre seeks to use efficient information sharing

& collaboration strategies to achieve effective and valuable visualizations. Viegas & Wattenberg argue in their work [VW06] for better ways to collaborate between different parties when analysing and presenting visualizations, and point out that old fashioned ways are still very often in use. Researchers often work offline by themselves and share too little of their findings during the exploration and analysis phase. Thomas & Cook also state [TC05] that a narrative story would be of great importance to support evidences found in visualizations. This fits perfectly with the model of asynchronous collaboration, which would enable better collaboration with less amount of misinterpretations.

Yang et al. [YXRW07] also point out that researchers often work offline during most of the visualization process. What if the visualization or results one was working on already was finished, or helpful hints could be found in a similar visualization made from the same data? Yang et al. provided a nugget pool, a collection of findings, important and valuable information such that users could share their results at an earlier stage than its final result.

In that context, Jern et al. [JRAY08] present a collaboration tool that could export visualization images adjoined by textual explanations to the HTML format. They also added the data used in the visualizations in a suitable XML format to allow collaboration between different partners. This asynchronous way of working on improving visualizations or indeed deriving new visualizations based on previous results is much sought after.

Shrinivasan & van Wijk [SvW08] sought to enhance the reasoning process for the audience by defining three different views in their visualization system. In the knowledge view, they allowed saving of milestones and providing an overview of existing milestones. In the navigation view, recording of actions, which steps were taken toward creation of an distinct visualization was recorded, and finally data view where the actual visualization was done. The system was capable of loading previously saved states, and would thus allow easy comparison of visualizations made.

By allowing an increased number of people to examine the visualizations or indeed the data itself, the process of data exploration and data analysation can be greatly sped up. By allowing a certain interactivity to the data presentations, people can refine the presentations, or even dig deeper in the data attached to the presentations. There exist solutions online, exemplified by Manyeyes [Eye] and Gapminder [Gap], that allow users to upload their own data, create stunning visualizations, and share them with the public. Others may now explore, analyze or have the visualizations presented to them. Examples of such collaboration opportunities freely available on the world wide web can be seen in figure 3.2.

Another interesting environment is the Tableu database visualization system by Heer et al. [HMSA08]. They introduce history mechanisms to record and cre-

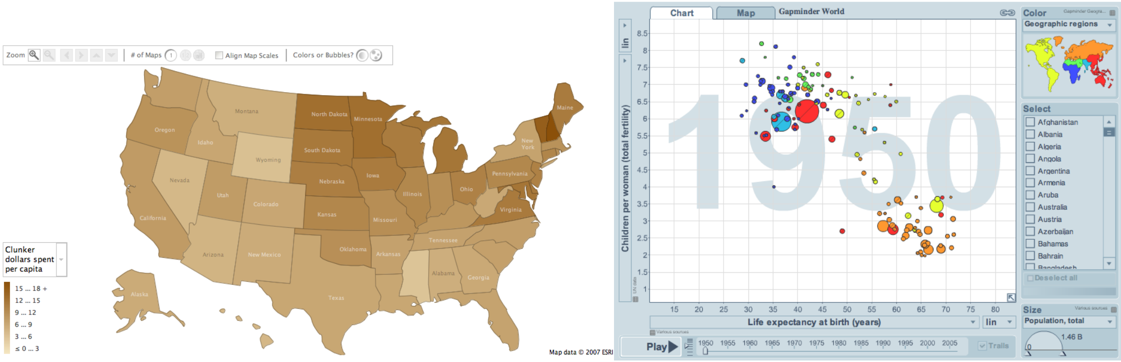


(a) **Many-Eyes** (b) **Gapminder**

Figure 3.2: *Many-Eyes is property of IBM [Eye], Gapminder property of Google Inc. [Gap].*

ate visual interaction histories that aims to support analysis, communication and evaluation of data within the field of Information visualization. These histories can be exported to support for collaboration with different partners as well.

Working with time-dependent data, Lu & Shen [LS08] presented a system were the audience could interactively choose which time ranges to visualize and provided some controls to tweak their visualizations also. These interactive storyboards were able to playback the data after user input, and would be presented as stories for visual exploration since the audience could refine their playback attributes easily.

The LiverSurgeryTrainer by Cordes et al. [CMO+07], is a visualization environment that includes a scenario training system for surgery of the liver. It allows doctors and students to participate and learn about surgery in an non-harmful virtual environment. The framework has a predefined set of stories that can aid the users to learn about its related content.

## 3.2 Storytelling in Volume Visualization

Storytelling has proven repeatedly to be a very potent way of retaining, transferring and absorbing knowledge. The field of visualization is ultimately all about finding useful information, and presenting it in the best possible manner. Seeing storytelling as a motivation, Wohlfart & Hauser [WH07] adopted the presentation form to suit stories made in volume visualization. While visualizations often rely on some expertise in the audience or careful commenting, storytelling would allow a much broader audience due to the fact that series of images could be presented. The relation between the visualizations would also be easier to grasp when proper narration and context information is available for the audience.

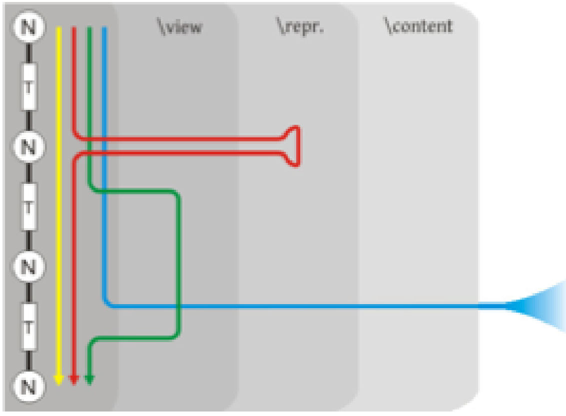


Figure 3.3: *The four interaction schemes defined by Wohlfart & Hauser [WH07]. Yellow line; regular playback, the red line desribes interactive approval, green represents semi-automatic storytelling and blue is full detachment. The distance along the x-axis represents the amount of interactivity at a given stage in the story.*

Wohlfart & Hauser defined a four level interaction scheme that would allow the audience to take part in the story being presented. These interaction levels are defined as playback, interactive approval, semi-automatic storytelling and full detachment. The default scheme would be playback, were the story would be presented to the audience in the precise manner that it was intended to be viewed. The interactive approval scheme would allow the audience to pause the story, make changes to the story at that given state, for instance changing the viewpoint to allow inspection or confirmation of the data from another angle, then returning to the original storyplot again. The next level of interaction, the semi-automatic mode, allows the audience to make a permanent change to one or several parameters indirectly controlling the story, and keep this parameter with its new value for the remainder of the story. A depiction of the impact these levels of interaction have on the story, and what freedom they allow with respect to the story plot, can be seen in Figure 3.3. An example of this type of interaction could be viewing a CT volume of a body, and making a change to the transfer function causing the skin represented in the volume to be transparent. This would reveal the underlying structures, i.e. bones and skull for the rest of the story. The last scheme is when the audience choose to disembark from the story, taking full control of the story themselves. In this scheme, the story is completely disregarded at a certain point, and the audience freely can do whatever changes wanted.

### 3.2.1 Storyplot

The plot of a story contains highlights, important scenes, which follow one another in a predetermined fashion. These highlights are considered the most important events in the story. The story unfolds itself by visiting these highlights one after the other, and gradually twist the story toward the next highlight to create an seamless flow connecting the events. In theatre and movie production it is common to sketch graphical representations of the important scenes, turning points and climax at an early stage during production. A sketch artist often draw a simple sketch frame to represent each highlight or milestone of the story, to allow actors and producers an overview of the plot.

It is vital to the production of stories to identify and specify milestones. Although the milestones depict a breakdown of important events included in the plot, they are limited to describe only the state of the movie at that very instant. So with respect to the story, there are large gaps between the milestones themselves.

This representation of the plot in milestone form is commonly used in comics. When interpreting, i.e. reading comics, the audience assumes a relation between the milestone frames drawn, and mentally completes the story themselves. The milestones often depict situations the audience have some familiarity with, and this familiarity aids the audience to make an educated guess toward how the transition between the milestones is supposed to be. The possibility of having the audience assuming logic relations between milestones is only possible when describing milestones that include situations that the audience have prior experience of. Whenever stories include more complex milestones, the relations between these must be described either visually or orally explicitly.

Between the milestones there must exist transitions of states, that lead the story from a milestone to the next. These transitions can be considered the glue that holds a story together, a potent set of cues that guides the story in the right direction. In Figure 3.4 an example story is depicted. The resulting story would comprise of a gentle flow of small changes, producing an animation like effect. The important visualization states of the story, the milestones, would be visited exactly once. The transitions between the milestones would handle the task of bringing the story from milestone to milestone.

Nuij & van Wijk [vWN04] show that when presenting a context change from one state to another, it is important to facilitate for transitions that aid the audience in interpreting and fully understanding the extent of the change. They exemplified that by a scenario where an audience were presented mapdata, and the current location were zoomed in on New York. When the focus was shifted to Los Angeles, it would be more comprehensible to zoom out to achieve an overview, then zoom in again on the new city, as opposed to in zoom panning along the continent.

The transitions can also be tailored to give the audience further information about the upcoming milestones in the plot. If a story progresses with much change, we can assume that the changes are of less importance but still keeping the audience focused and not be confused by the changes being done. In the opposite situation of much change, less change during a transition signify that the audience should stay focused on the particular changes being done and thus making a crescent tension in the story. The work by Beyer & Hassan [BH06] states that animated

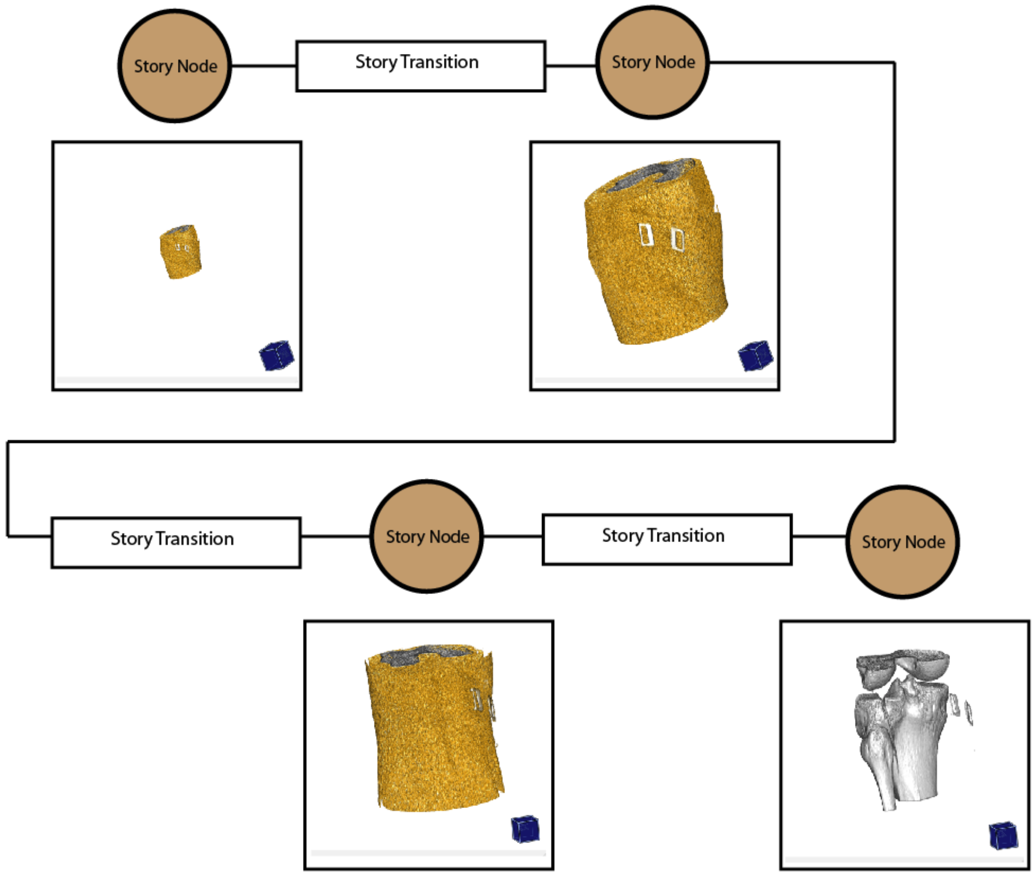


Figure 3.4: *An example visualization story. The story is created by defining milestone visualization states that the story must visit in a predefined sequence. The transitions guide the story from milestone to milestone by gradually changing the visualization parameters, visiting each milestone exactly once. The resulting experience is an animation like flow of visualizations.*

storyboards, i.e. animations containing milestones with proper transitions, would be able to emphasize and communicate interesting events that single images or movies containing single distinct images could.

### 3.2.2 Narrative paradox

During regular storytelling, the most common situation is that the narrator presents the story to an audience from start to end without interruptions of any kind. In contrast to regular storytelling, stories contained in a virtual environment often allow or even encourage the audience to participate in the story itself. This participation or interaction can cause a conflict between the narrator's original story content, the intended message, and how the story is perceived when the audience is allowed to make changes to the story plot. This situation is known as the narrative paradox in virtual storytelling. Louchart & Aylett [LA03] presented the emergent narrative in their work, which is a way to cope with this paradox. The narrative paradox arises when the audience is allowed to alter the pre-defined story.

The main reason for allowing the audience to alter the story is to allow them more freedom with respect to the predefined story at hand. By allowing the audience more freedom, the story becomes more credible.

This interaction can be used to affect the story to a whole lot, depending on how much change and interaction the audience wants. An example of interaction that is common that audiences in virtual environments take advantage of is approval. The audience pauses the story, and makes some changes to check, or re-check some feature. This approval is important for the audience, since it allows them to inspect the story at any given point to support confirmal or denial of facts. When any interaction is over, the audience might want the story to continue as is was originally intended. The changes done since the story was paused is discarded, and the story continues in its original form.

Another situation that could arise, is for instance when viewing a dataset from a specified viewpoint, the audience might want to see the dataset from a an angle that reveals different characteristics. After pausing the story, the newly defined viewpoint can be taken into account when the story is set to continue. This conflict makes a bigger impact on the original story, but does not necessary change the content contained in the story. When continuing, the new specifications changes the story in the way that the audience sees the story from another viewpoint, and all changes done to the story will be observed from this newly specified viewpoint.

Although these examples clearly changes the stories predetermined setup, this interaction makes the stories much more flexible with respect to the amount or size of the audience that might find an specific story of any interest. Stories can be used over and over again, since they can be tailored and rechecked during runtime to fit the audience better.

### 3.2.3 Achieving successful stories using visualizations

The success of storytelling relies heavily on the impact the story and its contents has on the audience. The narrator, the presenter of the story, needs to achieve two important criterions to successfully transfer the knowledge content to the audience. These criterions is that the story must be ***1) comprehensible***, and ***2) credible.***

A comprehensible story requires a minimization of information loss between the narrator and the audience. The content should be as rich as possible, with respect to information, while being both clear and concise. If any content would be misinterpreted by the audience, the narrator has failed in successfully telling the story.

The other requirement for having a successful story is that the story content must also be credible. The narrator must present the story content in such a way that the audiences believes what they hear/see.

These two criteria are also emphasized by A. Glassner in his work [Gla99]. He states that there is a bond between the narrator and the audience, a story contract. In this contract, the narrators responsibilities are the integrity of the story — ensuring the comprehensibility of the story, and that there exist a sequence of plots that the characters involved in the story will follow. The audiences part of the contract is that they must grant the narrator to manipulate him/her intellectually, emotically and spiritually. This allowance of manipulation is a vital part of ensuring that the story is credible. If the audience is not willing to be moved, i.e. that they are focused on disproving the content of the story, the narrator has a nearly non-existing chance of telling the story and transfering the content.

There are several techniques that aid comprehension in visualizations. An easy, but very powerful method is to use redundant coding to ensure that the audience notices the importance of a certain feature. When using redundant coding, the visualization contains multiple mappings of attributes to a certain subset of data. This ensures that there are enough visual cues to make the selected data to stand out, metaphorically speaking, being emphasized so that the audience understands its importance. Examples of visualization cues that can be used for redundant coding is, for example, using color in addition to shape. Another example of redundant coding could be using both hue and saturation to emphasize the data.

In visualization, it is very common to use color or grayscale as an agent for describing some quantity or grouping of data. When interpreting such visualizations, it is impossible to compare data to one another without having proper color legends or color maps adjoining the visualizations. Color legends or colormaps should also be adjusted to their statistical properties to allow for a proper differentiation of the bulk of the data according to Schulze-Wollgast et al. [SWTS05]. By doing so, color representations will increase its usefulness.

Visualization often include mathematical transformations that affect the 2D or 3D content such as scaling, rotation and translation. These transformations alter the visualizations, for instance zooming onto a specific position of an object. Without a gradual motion describing this zooming from the overview position to the final close-up, the audience will loose the context and any spatial relation with respect to the original position. When doing such translations, Shoemake [Sho85] described spherical line interpolation that would generate gradual camera motions well suited to help viewers retaining the spatial relation when using transformations in computer graphics.

Another often used solution to aid comprehension in visualization, is to use focus+context techniques as described by Hauser [Hau05]. Focus+context visualization techniques effectively solve two problems for visualization, how to be able to differentiate between attributes to emphasize focus-data (and de-emphasize uninteresting data), while at the same time keep data to support as a spatial reference. The user specifies how important every data attribute is by giving it a DOI-value (degree-of-interest). The most important attributes are given the value 1, and the least important attributes are assigned the DOI of 0. This gradual value describing every attributes interest is a key factor when composing the visualization. By using these DOI-values when inspecting the data, the visualization can be created by emphasizing the focus-data, while keeping important surrounding data as context. Figure 1.2 is such an example of f+c.

Labels are known from textbooks and posters to aid as descriptive tags. Such labels are often placed in the surrounding of the graphics as a silhouette occluding on as little amount of the data possible. Labels and annotations enrich visualizations by adding descriptive text to the images. This help interpreting the visualizations, and can also be used to prepare the audience for upcoming changes and/or comparable results. Oral comments are according to Krueger and Schär [SK00] an underrated way of adding content. Not only will it allow the audience to fully focus the on the visualization while listening to the oral comments, but Kruger & Schär also show that using a variety of modalities will minimize the risk of overloading one sense.

To achieve a successful story, the two criteria, comprehensibility and credibility had to be fulfilled. While we have described several methods that aid the narrator in reaching the first criterion, little have been discussed with respect to the latter. Credibility of the visualization story can be improved by several means. By just allowing the audience to experience the story again, have it being retold (or replayed) will allow the audience to be better prepared for the story and thus focus on the important parts (since they know more or less what will happen). Another very strong mean for ensuring credibility is to allow the audience to halt and inspect the story. This is the equivalent to asking questions to the narrator, and receiving replies. This introduces the strongest mean available, the option to allow the audience to alter the story. This implies acceptance that the narrative paradox as describes in part 3.2.2 will arise. By allowing the audience to alter the story, different viewpoints can be selected, and focus can be shifted to make other data more visible (for instance making skin transparent to reveal bones underneath).

## 3.3 Goals for this thesis

In this thesis we aim at exploiting storytelling qualities to enhance visualization presentation. We would like to clarify and exemplify different types of stories, and how these stories could benefit different types of daily medical communication.

Previous work has been done on storytelling for presentation by Wohlfart & Hauser [WH07] that functions well with a defined range of visualization attributes. We aim to improve this storytelling model to allow an arbitrary amount of attributes to be covered in a visualization story. This will allow newly devised plugins for our environment to be used in the visualization stories without any changes needed.

We aim to incorporate storytelling features in our visualization framework, Volumeshop, that allows recording, editing and playback of visualization stories. In addition to creation of stories, we would also give the audience the ability to participate in the stories through interactivity. The audience of the visualization stories should be able to control playback, do interactive re-checking, and to make changes that would be reflected in the visualization stories right away. The visualization framework should present an understandable overview, similar to a storyboard, of the recorded visualization story.

Labelling and annotations have been a crucial part of presentations of visualizations to allow the audience to better comprehend the message encapsuled in the visualization. We aim at giving the narrator of the story the ability to insert labels and annotations into the visualization stories created in our visualization environment, to support comprehension and clarification. As visualization stories might raise additional questions regarding different medical pathologies, we aim to enhance the visualization stories with direct linking to medical dictionaries and other content found on the internet. This type of content enrichment is an optional part of stories, and will not be forced on the audience. We also aim to include the possibility to record oral comments during the creation of the visualization stories to and to have the recorded verbal comments enhance and clarify the visualizations during playback. According to Schär & Krueger [SK00] relying on different types of modalities (oral and visual) will reduce the chances for overloading of one senses.

To support story longevity and story collaboration, we would allow to import and export visualization stories based on an XML model that would support any visualization system. The structure of the XML will be discussed at a later stage in this thesis.

Time is of essence when working with visualizations, and we propose a set of story types that general enough that they fit many stories resulting in a high degree of re-usability. We propose to make templated versions of these stories to save time during the story creation process.

# 4 Storytelling in Volume Visualization

> *Stories are how we learn. The progenitors of the world's religions understood this, handing down our great myths and legends from generation to generation.*
>
> \- Bill Mooney and David Holt

This chapter aims to thoroughly explain and exemplify visualization stories similar to the visualization story seen in the visualization story example in Figure 3.4. We show how visualization stories are created, edited and interactive playback can be achieved within our visualization environment, Volumeshop, and go into detail regarding how visualization stories are structured and the parts of our story model.

The framework Wohlfart & Hauser [WH07] used, RTVR, was limited to a fixed set of visualization parameters, meaning that the story functionality would have to be altered to cope with any new visualization parameters introduced to the framework. In our work, we show how visualization stories can exceed this limitation and cope with an in principle arbitrary selection of parameters. At last we discuss how labels, annotations and oral commentaries can enhance and support comprehension when interpreting visualizations.

## 4.1 An enhanced visualization story model

The most simple and straightforward approach of presenting a story, is to present one change at a time. Using such an approach generates a long animation sequence for the audience to view. There are several distinct drawbacks using such an approach; **1)** there is no possibility to distinguish the parts of the visualization story that represent the peaks (the important informative visualizations) from the transitional visualizations (the middle steps towards the next peak). Meaning that every change is treated as important as the other, in contrast to a plot where

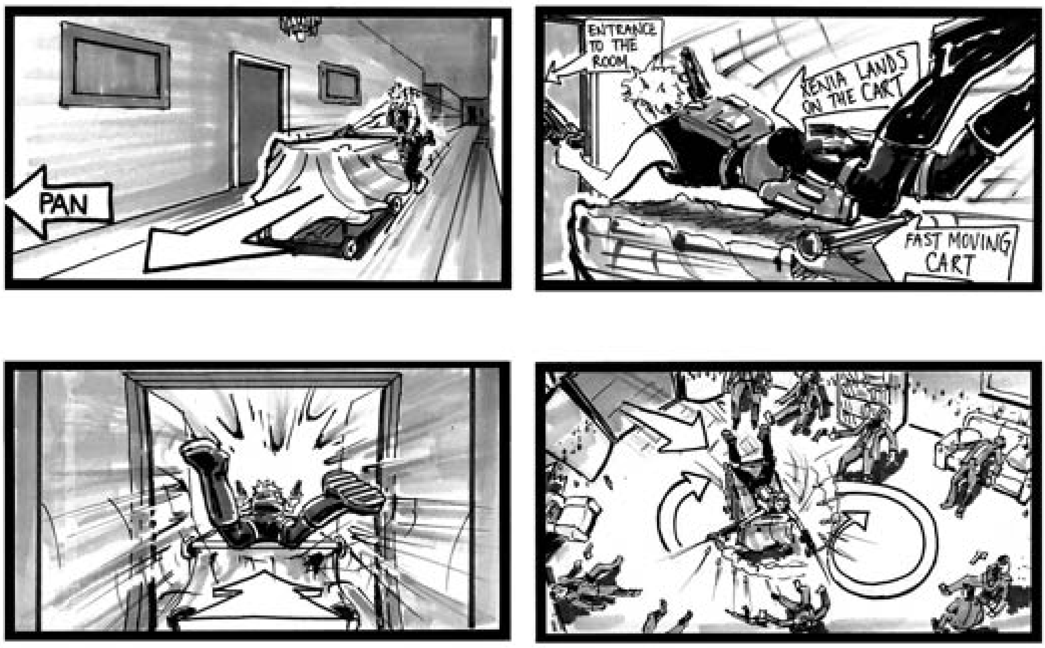


Figure 4.1: *Sample Storyboard from a film. Metagraphics has been added to the scene frames to allow better interpretation of the scenes. Image courtesy of Storyboards Inc. [Sto].*

it is only natural to differentiate between these changes. **2)** the long animation sequence leads to long waiting times, and the audience can loose interest or focus with respect to the important parts of the visualization story. A better story model and presentation scheme is needed. We examined more traditional presentation forms like movies and cartoons for clues, and have derived a visualization story model that suits and enhances virtual storytelling. In the next paragraphs we will elaborate on this new story model.

**Storyboards**

Movies makers and cartoonists have for long used a technique named storyboarding in their respective professions. A storyboard is essentially a sequence of images depicting important events in a story. Storyboards can be considered as a scene by scene breakdown, rather than a frame by frame breakdown. When drawing storyboards, the artists focuses on the most important events in a story and often includes captions, comments and directives to every frame to explain how the scene is intended to by interpreted or produced. This helps readers, directors and cinematographers to properly visualize the scenes. An example storyboard can be seen in Figure 4.1 where metagraphics has been added by the illustrator to emphasize how the specified scene was intended to be. Such storyboards are also present in other fields of work as well, exemplified by graphical user interface design where the storyboards not necessarily are a sequence of crescendos in a plot, but could indicate, as Benyon et al. [BTT05] use storyboards, scenarios for different user options accompanied with metagraphics explaining how different types of interaction inside the GUI will result.

A visualization story, as we see it, is composed of a series of visualization states. Each visualization state describes an important scene in the visualization story. These visualization states are represented in our story model as story nodes or story milestones. A set of story nodes can be compared to storyboard as they both describe necessary scenes in their respective context. Each story node contains a complete specification of visualization parameters to be able to fully recreate that specific visualization.

### Keyframing

While storyboards are excellent scene breakdowns, that give a nice overview of the story and plot, the natural flow that animations and movies provide are lacking. We also miss important information regarding the transitions made between the scenes since we have to make assumptions how the flow and change between each scene is supposed to be. Movies solve these gaps by filming the scenes and fitting them together in subsequent steps. In computer animation production, a common technique for producing animations is keyframing. See the works of Melikhov et al. [MTS+04] for more information regarding this technique. With keyframing, certain frames are selected as keyframes, and these frames are guaranteed to be visited during in the resulting animation. This technique solves the problem of disjunct storyboards, frames without any complete transition, by allowing a specification of timing and interpolation schemes between such frames. The result is very similar to what producers of cinematic movies end with after the clipping process; smooth transitions between the scenes. Keyframing software often allows different parameters for different attributes, meaning it would

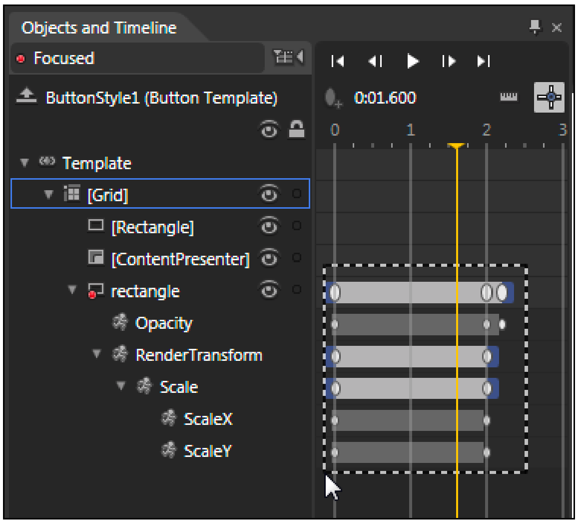


Figure 4.2: *Keyframe editing in the Express Blend environment. Image courtesy of Microsoft Express Blend*

be possible to have different interpolation schemes (e.g. linear, b-spline, Bezier) for different objects in an animation, over different amounts of time. An example of keyframe editing can be seen in Figure 4.2 where several parameters for an animation are tweaked along the timeline.

In a visualization story, there must exist a connection, a gentle transition that moves the story in a purposeful and gradual manner towards the visualization state contained in the subsequent story node. We specifically want to avoid abrupt visual changes that can cause confusion and disrupt the mental map and comprehension for the audience. This is the task of the story transition, to guide a viewer gracefully from story node to story node through the whole story.

**Elements in the Visualization Story**

The storyboard approach is a simple yet powerful technique that coincides well with storytelling, by describing a series of scenes presented in a linear fashion. A drawback of the original storyboarding technique is, as stated above, the lack of structure and information in between the storyboard frames.

In our visualization story model we take advantage of the flexibility the keyframing technique provides, in form of independently setting timing and interpolation schemes between the keyframes. The keyframing process from computer animation is notoriously known for the tedious work it requires, we specifically want to avoid this overload, and aim to automatize the creation of these transitions.

Combining these two, we end up with a versatile story model that supports an iterative build-up process, i.e. it allows recording while the user interacts with the visualization environment, while it also supports complex transitions that provide narrative cues for the audience.

We go through the different elements of our story model for visualization stories in the following paragraphs. In our framework we have chosen to separate the story structure into three different parts, namely *story action groups*, *story actions* and *story milestones* as described below.

**Story Milestones**

The story milestones are the essential visualization states in our framework, and would if printed, comprise as a storyboard for the story. The milestones contain complete information of the visualization state, to enable fully recreation of that visualization, in addition to a timing parameter to control how long the specified milestone should be visited during regular story playback. By default, the story nodes will be visited in a iterative, linear manner, one after the other in a specified sequence. Every story node is indirectly connected to exactly two other story nodes (one before, one after) with the exception of the beginning and end node, which both is connected to only one story node.

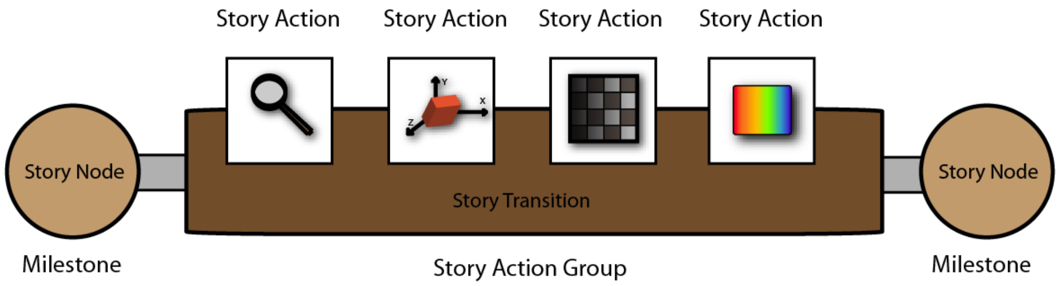


Figure 4.3: *An simple example of the story structure. The story visits the first milestone, before the upcoming transition gradually changes the visualization into the next milestone. This result in an animation like experience. The transition in this example consist of four different story actions, all done in parallel when the transition is executed.*

### Story Action Groups

After having defined the milestones as the placeholder for the story nodes (according to the story model), we define the story action group as a container object for story specific atomic actions. The Story Action Group enables the story to treat each transition individually similar to the keyframing technique. It stores along with the atomic actions also a timing parameter, to control how long the contained actions should last. The groups are relative to its neighbouring milestones, meaning that they only consist of changes to visualization parameters that should be done during playback. The groups are calculated automatically during recording, but it is possible to indirectly edit the content and directly change the timing parameter. We will go in-depth with respect to editing of such groups later on. By allowing the groups to accumulate more than one visualization parameter change (it can be only one as well) the possibility of parallel sequencing, as Butz described in his work [But97], is enabled. An example of a parallel sequence would be a story that rotated the datavolume while at the same time changed the opacity of the skin represented in that data.

### Story Actions

Finally in we come to the story actions, which in essence describe the actual transition content. Any change done from one visualization state to another, requires a minimum of one story action. The story actions represent atomic visualization parameter changes and examples include a rotation of the data volume, change of camera fov (field of view), change of transfer function (opacity levels and color mapping) or sampling distance for the raycasting equation. The story actions know a starting value and an corresponding end point for those that value. Since the story action groups iterate over the story actions and force them to complete the transition from start to end point over its specified amount of time, it is a lot easier to handle the story actions, and not necessary with additional timing at that level of the model. Figure 4.3 show a simple story structure consisting of two milestones and the transition between them.

Summing up, the transitions are relative states with respect to the nodes, which in turn are the absolute states of the story. As we will describe in more detail later on, the stories can be edited after creation, both deleting and inserting nodes in arbitrary sequence positions. The story automatically adapts to the new nodes and creates the needed transitions to properly guide the story as intended. During playback, the presented stories feels like real animations, with complete transitions moving the story from the beginning visiting each visualization state defined in turn towards the end. Although the stories presented here are modeled in a linear fashion, any branching story can be modeled as a longer linear story.

## 4.2 Stories with arbitrary sets of changing parameters

In previous work done on visualization stories by Wohlfart & Hauser [WH07], the stories could only cope with a predefined set of changing visualization attributes when composing and playing visualization stories. In this project we realized an approach that supports an arbitrary set of visualization parameters to change when integrating our story model in the visualization framework.

All previous approaches regarding the capture of visualization parameter changes relied on iterating over a set of known parameters and comparing those against the previous state to determine whether there was a change or not. This implied that the framework responsible for the story recording process would have to be rewritten or adapted once new visualization parameters were to be introduced in the stories. The longevity of such a solution is undeniably limited. We sought to avoid this limitation by making a framework for creating, editing and playing stories that could adapt to a more flexible number of visualization parameters. In order to achieve such a solution, two criteria in the changing framework had to be met; **1)** a predefined set of variable types would have to exist that allow type checking, i.e. inspection of what kind of type any given parameter is, and **2)** a way to extract and insert a given visualization parameter into the framework.

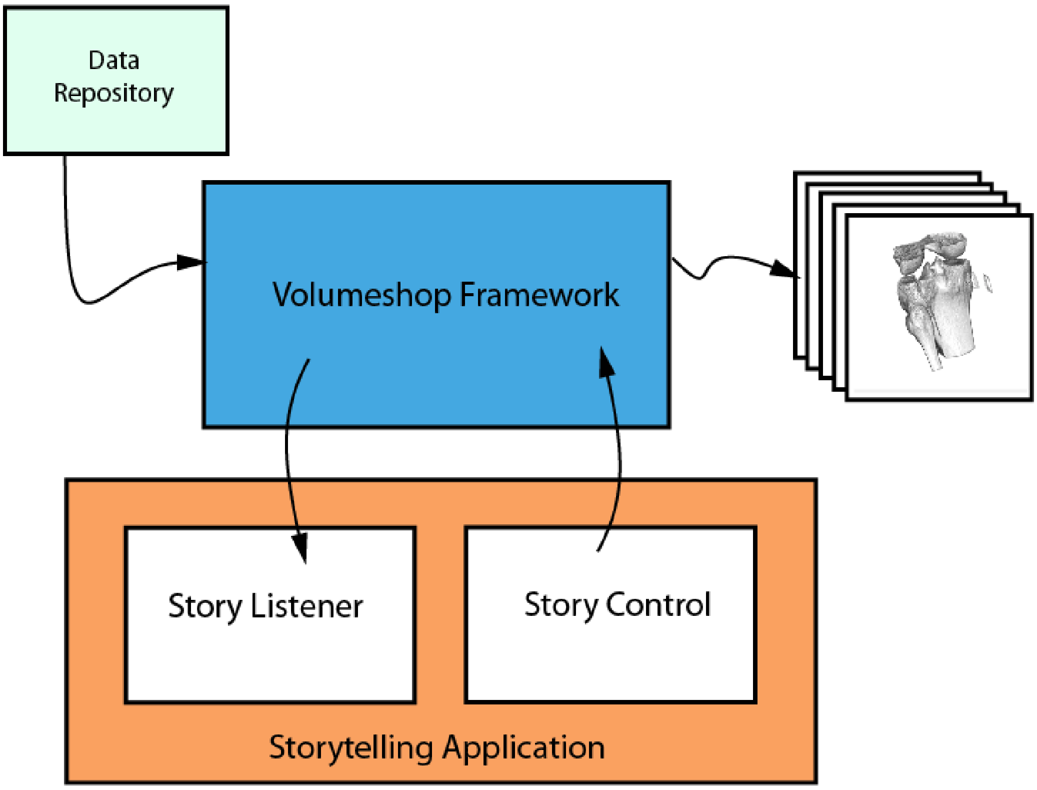


Figure 4.4: The storytelling application collects information via the listener object, while at the same time it needs to be able to inject visualization parameters into the visualization framework to successfully handle story recording and story playback.

The first criterion ensures the correct interpretation of all changing visualization parameters, by allowing the program to check which kind of visualization parameter is in question. Understandably regular visualization parameters like a matrix (for defining camera, rotation etc.), a vector (position of light), simple types like floats and integers would have to be treated very differently. By having a predefined set of such parameter types, the rules for correctly interpreting the changes done in the stories can be tailored.

The second criterion required for the handling of visualization stories is that it must be possible for the story control to both read from and write (make changes) to the visualization parameters in the environment. Reading from these parameters enables comparison of values to create the stories, and writing values back to the environment enable animation like playback since the story control effectively determines which visualization is set up.

In Figure 4.4 we show how we achieved storytelling capabilities within our selected visualization framework, the Volumeshop environment by Stefan Bruckner et al. [BG05]. Volumeshop accesses volumetric data, and presents visualization results based on a number of visualization parameters. The framework uses variants, a container class for all visualization parameters used, that also provides information regarding which type of visualization parameter it contains. This fulfills the first criterion for arbitrary sets of visualization parameters. Secondly, it allows attaching of a listener object that polls information with respect to what changes have happened in the visualization process, but also complete visualization states whenever needed. This enables us to fully create visualization stories based on the information at hand. Lastly we can use the stories created to enable interactive playback in the environment. Our basis for creating and playing visualization stories with arbitrary parameters is then complete.

## 4.3 Interaction in Visualization Stories

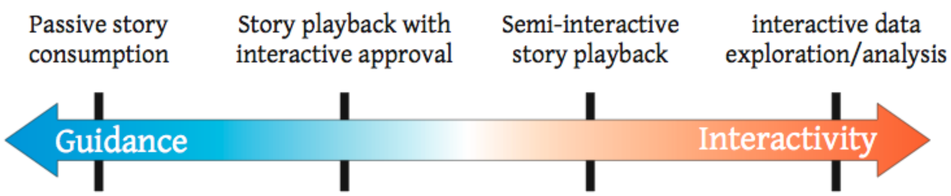


Figure 4.5: Interaction schemes with respect to the level of story control that the narrator provides. Image courtesy of Wohlfart & Hauser [WH07].

Interaction is a key feature of visualization stories. As described by Wohlfart & Hauser [WH07], traditional approaches to presentation of visualizations which rely on simple animations or sequences of images have two major problems. **1)** It is difficult to correctly interpret and conceive spatial relations in visualizations, and **2)** important structures can be hard to perceive due to occlusion and lack of optimal viewpoints. We have stated earlier in section 3.2.3, that the support for comprehension and credibility are fundamental in visualization stories. Interaction as such would allow the audience to, for example, halt the story while rechecking the data from another view angle. We have also discussed in section 3.2.2 the paradox that arises when the audience is allowed to participate and make changes to a story. In the domain of virtual storytelling, Louchart & Aylett [LA03] presented a solution to cope with this dilemma. In volumetric storytelling, Wohlfart & Hauser presented four interaction schemes that gradually shifted the story control from the narrator toward the audience. In figure 4.5 we can identify these four interaction schemes with respect to how much guidance and freedom the audience are governed with. We will describe these four interaction schemes in the following paragraphs as we have realized them in our visualization framework.

**Passive Playback**

In the passive playback scheme, the audience follows the story without any degree of interaction done. The story progresses in a natural way, from visualization state, e.g. milestone, to the next, with gradual comprehensible transitions in between. This resembles a VCR-like playback, which can be paused and started at will, and that also can be replayed over again if the audience feel they were overwhelmed by information.

**Interactive Approval**

While watching passive playback, there is a certain probability that the audience finds features visualized in the visualization stories intriguing. If the audience takes note of an interesting feature, or something that for any other matter attracts their attention, they possibly want to inspect this feature further. The interactive approval scheme allows the audience to halt and immediately diverge from the story in the way that they can interactively investigate any feature while in the midst of the playback, and changing any needed visualization parameter to do so. This interactive reinvestigation is crucial to support optimal credibility in visualization stories. When the audience is finished with their interactive investigation, the story gradually returns to the point in the story where it was paused, to avoid disturbing the mental map of the visualization story. The story continues then again according to the predefined structures. Figure 4.6 show on the left side the interactive approval interaction scheme. The story is paused while the audience diverges from the story to examine or recheck the data. The story then returns to the point where it was paused, and continues as planned.

**Semi-Automatic Playback**

Creating the perfect visualization story is an iterative process and it is likely that the stories produced would affect people in a varied way. Some might be more interested in certain specific regions of the data, while others might opt for different viewpoints for the visualization story. The semi-automatic scheme enables the audience to pause the story and change visualization parameters to their liking. An example of this, is as mentioned above, changing the view-direction for the camera to achieve a different view for some reason. The story would then disregard any visualization parameter there was issued a change to during the time this interaction scheme was selected, meaning that whenever playback is restored

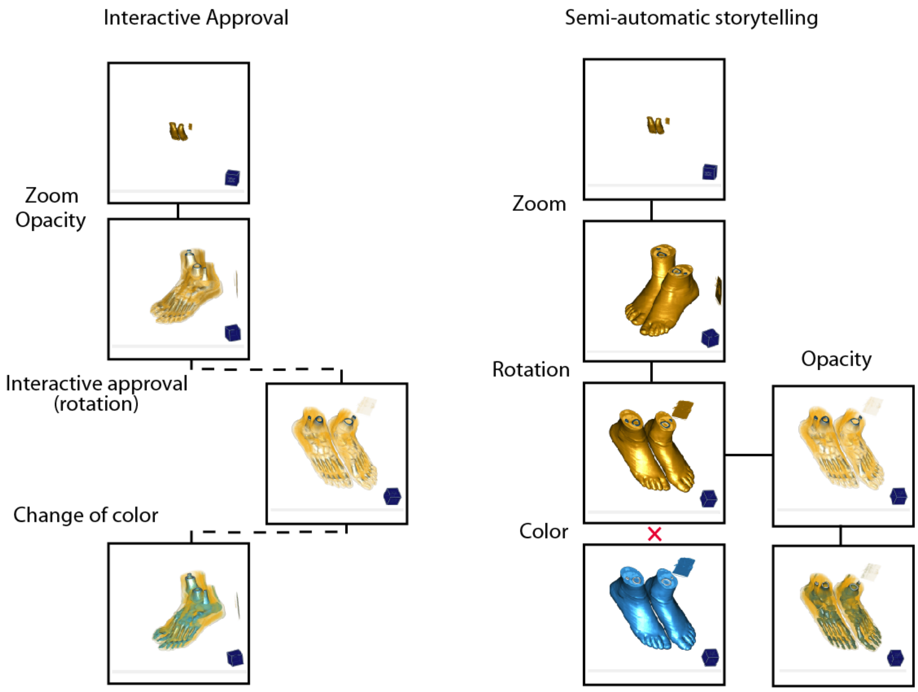


Figure 4.6: *The interaction schemes Interactive approval and Semi-automatic storytelling. In interactive approval, the story is paused when the audience want to recheck or investigate the data, and then returns back to the story. In semi-automatic storytelling, the audience issues a change to the story, which the story takes into consideration. The story progresses, but uses any newly defined changes as specified by the audience. Dataset courtesy by Casimage [Cas].*

and the story goes on, it will disregard any change done to the newly tweaked visualization parameters and use the newly defined values for those parameters instead. This way, stories can easily be refitted to suit a broader audience very effectively. Such visualization stories according to Van Wijk [vW05] have a high visualization value, since its reusability is quite high, and the amount of work involved in achieving that would be quite low. If the user opt for changing back to passive playback instead of the semi-automatic playback, control with respect to the story would be restored again. Figure 4.6 show on the right side how the structure of such an interaction scheme will change the original story.

### Full Detachment

This interaction scheme allows the audience complete freedom with respect to interaction. In this scheme, the audience chooses to disregard the story as a whole,

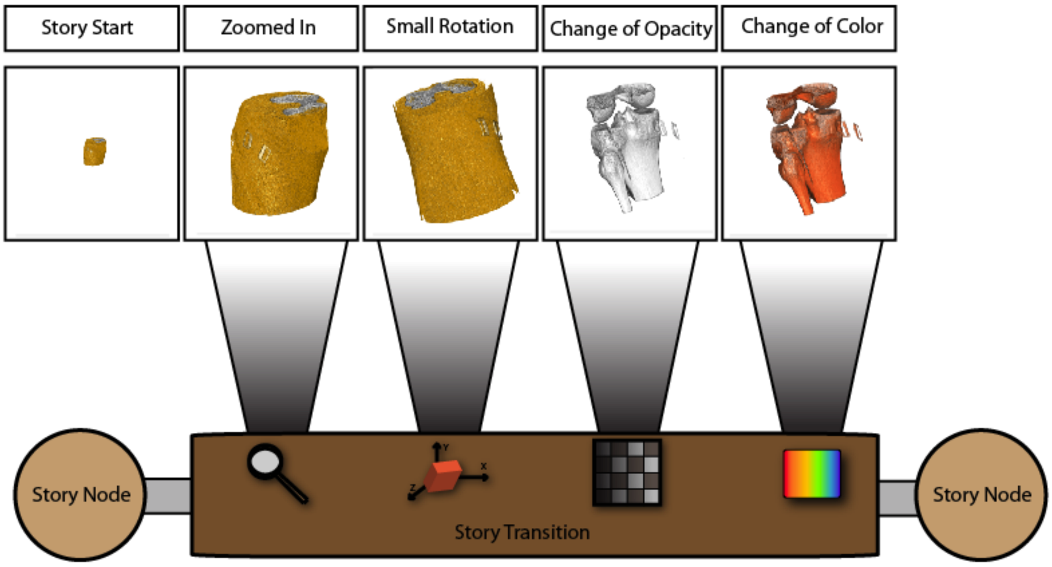


Figure 4.7: An example recording of a story milestone and its inherent transition. The data is zoomed onto, rotated slightly, a change of opacity is done do reveal the broken bone and finally the colormapping was changed to emphasize the cracked tibia.

and can explore and analyze the visualizations at free will. Though if for any reason the audience wants to continue the story, the framework would allow so. This type of interaction allows for asynchronous collaboration since the audience can now freely explore the data to possibly identify new interesting features that could be included in the visualization stories.

## 4.4 Recording, playing and editing stories

In the previous sections we have described how visualization stories are structured, and different interactive possibilities the audience has when the stories are played back. Without tools for creating such stories, the extent of storytelling would be scarce. In this section we will cover how such visualization stories can be recorded, how they can be played back and how they can be edited as well.

Groth & Streefkerk presented in their work [GS06], a model for recording interactions done during the exploration and analysis phases of visualization. We have in our framework used a similar approach that records all interaction done since the previous milestone, e.g story node was defined.

**Story Recording**

When creating visualization stories, it is common to start the with an empty visualization. The user loads the data in question into the framework and starts to specify visualization parameters to either explore, analyze or present already known visualization states. The framework will during interaction, keep a log of all visualization parameters changed since the last milestone was set, to be able to generate proper transitions, e.g. comparing to find the amount of change, without testing all visualization parameters for changes. Once a new milestone is to be defined, the framework will start comparing the visualization parameters according to its log, to find the parameters that actually was changed since last milestone. There exist a case where a parameter was changed, and then changed back to its original state. Such changes will be omitted since they do not actually change anything in the visualization. In the case where the user starts with an empty visualization, every change will be inserted into the story action groups. In the case where only single parameter changes between two milestones, the playback would resemble sequential playback of changes. In all other cases where multiple changes are being done between two milestones, parallelization of changes, similar to the work of Butz [But97], are being done. In Figure 4.7 the process of creating story nodes is depicted. Multiple visualization parameters are gathered in one story transition dependant on the interaction and changes done between to consecutive milestones. It then defines the milestone and the transition related to the newly defined milestone and its predecessor, and then links the transition to them both. Once the transition is created, the log is emptied, and the framework starts listening for further interaction and changes for future transitions.

**Story Playback**

At any time during interaction, the story can be subject to playback. The authoring process halts in playback mode and the visualization story can be played back according to any of the four interaction schemes defined earlier; namely passive playback, interactive approval, semi-automatic playback and full detachment. The user also have the ability to manually step from milestone to milestone back and forth in the story (also jump beginning and end) much similar though the use of the VCR-like interface included. Once playback is done, either because the story reaches its end or because the stop button has been pressed, the framework returns to the recording mode and is ready to record or edit further.

**Story Editing**

The visualization process consist of a possibly large amount of parameter tweaking and overall refinement. Users go back and forth over the visualizations, changing small details, adding labels and captions, removing unnatural milestones, etc.

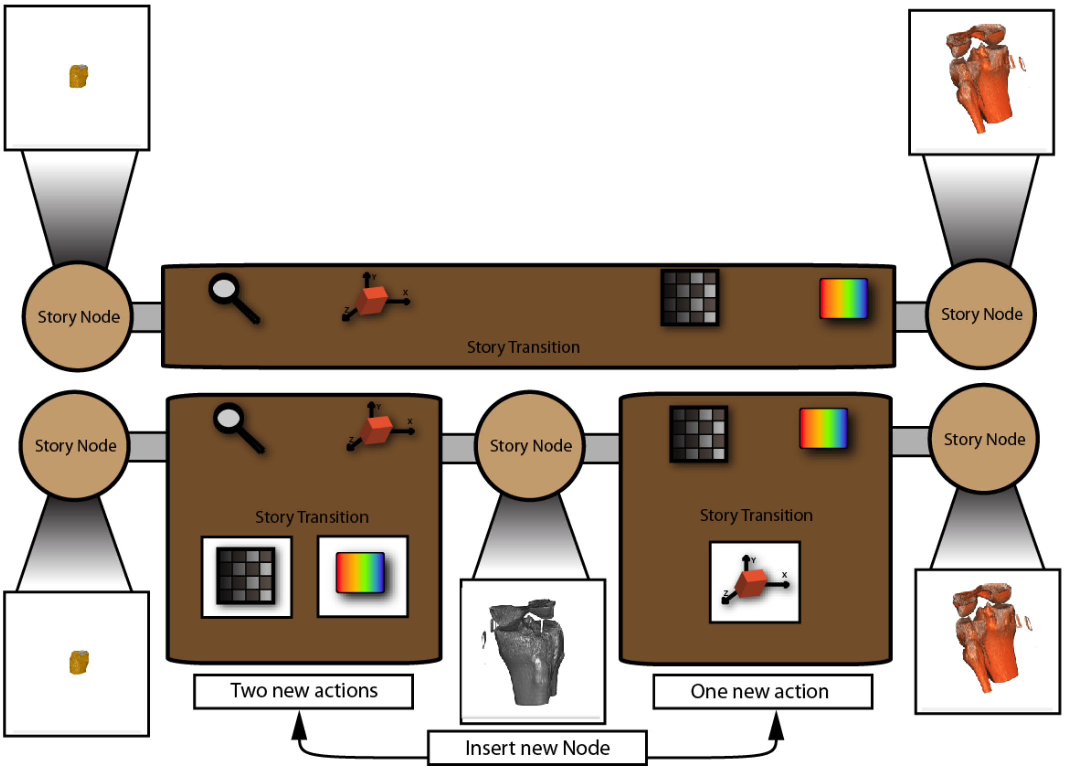


Figure 4.8: A new story node is inserted in the midst of the visualization story. A new story action group has been inserted, and additional actions have been added to the previous one to facilitate a proper gradual flow of changes when playing the stories. Upper row indicate before, lower row is after the changes are done.

Editing is as such a very important step in the creation of the visualization story. Another important feature as a result of editing possibilities, is the collaborative refinement of stories and visualizations between partners. Users can now create visualization stories, and send them to their partners for viewing, reviewing and iterative collaboration. This is very well aligned with the arguments by Viegas & Wattenberg argues in their work [VW06], who seeks better ways to collaborate between different parties with respect to visualization.

As mentioned earlier, there exist timing parameters for both the story milestones (nodes) and story action groups (transitions). These timing parameters can be changed freely to suit the story as the narrator want. We will describe two ways to alter stories in the following paragraphs, namely to insertion and removal of milestones at arbitrary positions in the visualization stories.

**Story Node Insertion**

As we mentioned in the previous paragraph, creation of visualizations and indeed also visualization stories is a refinement process. After reviewing the visualization story created, the user might want to change certain aspects of it, for instance introducing new milestones in the midst of the story. By stepping to the milestone before the intended change with the stepping controls, recording as insertion of milestones is possible. The user changes the visualization parameters to his/her liking, and opts for a new milestone. The framework will create a new milestone with the already logged changes. It will then calculate two new transitions (as opposed to normally only one), both dependant on the neighbouring milestones and the newly created milestone. The previously defined story action groups (transition) that connected the neighbouring milestones is removed, and the new transitions and milestones takes its place. See Figure 4.8 for an example of inserting a milestone in the midst of the story. This insertion (and as we will see, also removal) is very intuitive for the user, since the user only has to traverse to the designated point in the story to make the edit. The framework is fully aware of the position in the story the user currently is at, and will automatically solve the issue with incomplete transitions as specified above.

**Story Node Removal**

Story Node removal is as important as its opposite, since it enables visualization users to undo any milestones they previously have set. The user selects the milestone he/she wants to get rid of, and marks that for deletion. The framework responds by purging the story node, inclusive all visualization parameters contained in that node. Also, the forward and backward facing (if the node was not at beginning or the end) is also removed and replaced by a single calculated transition based on the new neighbouring milestones. The Figure 4.9 describes how the framework handles story node removal in an automatic way.

This concludes the story editing possibilities. We have opted and aimed for powerful, yet simple ways of recording, editing and playing our visualization stories, while at the same time making visualization stories possible for non-experts to create, edit and play.

## 4.5 Visual and oral visualization aids

Visualizations alone are often non-trivial to interpret. To optimally support comprehension when viewing (and indeed interacting) with our visualization stories, we have included some visual and oral aids to guide the audience. Textual labels and annotations can be added as story actions, and aid the audience with interpreting results, relations and prepare them for any upcoming changes in the stories

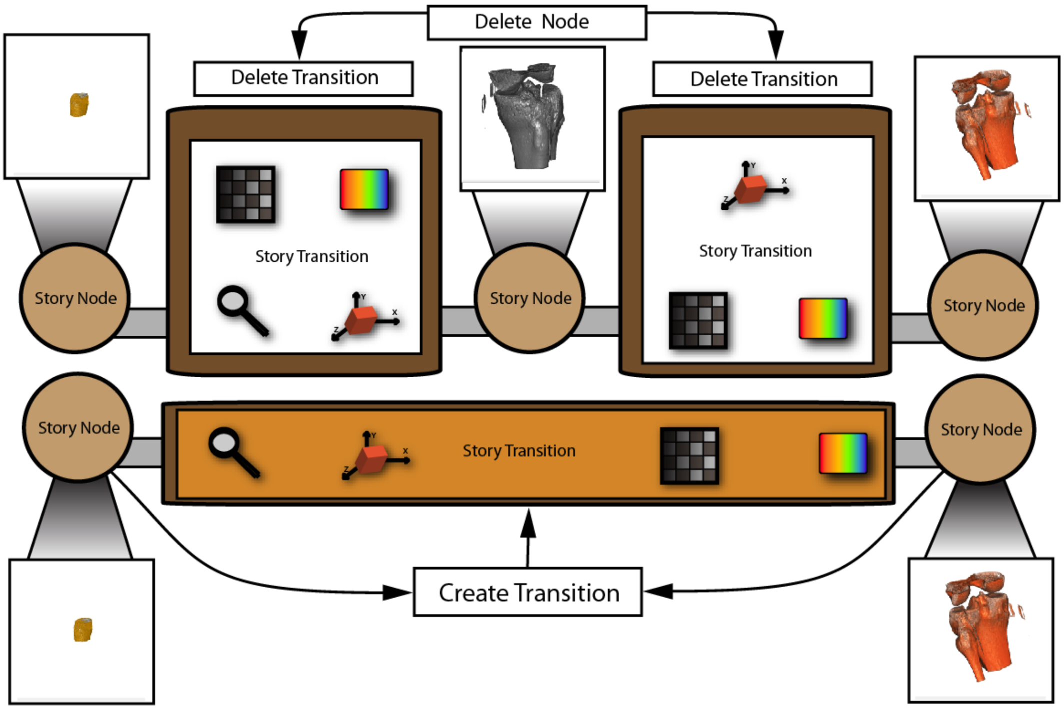


Figure 4.9: The middle node was marked for deletetion, and thus both story action groups (transitions) that was connected to the deleted story node is removed. The framework then creates a new story action group (transition) by inspecting the two neighbouring story nodes and find the differences between them. Upper row indicate before deletion, lower row is after the deletion and creation of new transition is done.

as well. Such story actions would be handled in the same manner as other story actions, meaning that they after a specific amount of time would be included in the visualization, and possibly be removed again depending on the story at hand.

The labels are external labels, meaning that they are positioned outside the convex hull of the visualization. They expand according to the amount of characters (text) written, and point toward a spatial location with their anchor line. The textual labels work in the 3D domain, meaning that the anchor of the labels (where the label point toward) points at the same spatial position with respect to the volume if the volume was to be transformed in any way (zoom, rotation etc.).

Annotations are always positioned below the visualizations, since it is common to place written text under figures or images. The annotations we included can be used to link to web pages also. The visualization stories can contain related material, for example a link to a medical dictionary or lexicon. This content can be activized by clicking on the annotation during the visualization presentation. The stories support multiple weblinks simultaneously. An example showing the use of labels and annotations can be reviewed in Figure 4.10.

To avoid overloading one sense (the visual sense), we have included the possibility to record and playback oral commentaries in the visualization stories also. According to Krueger and Schär [SK00], using a variety of modalities when presenting information enhances the focus of the audience. The oral commentaries are, similar to the labels and annotations, included in the story as story actions.

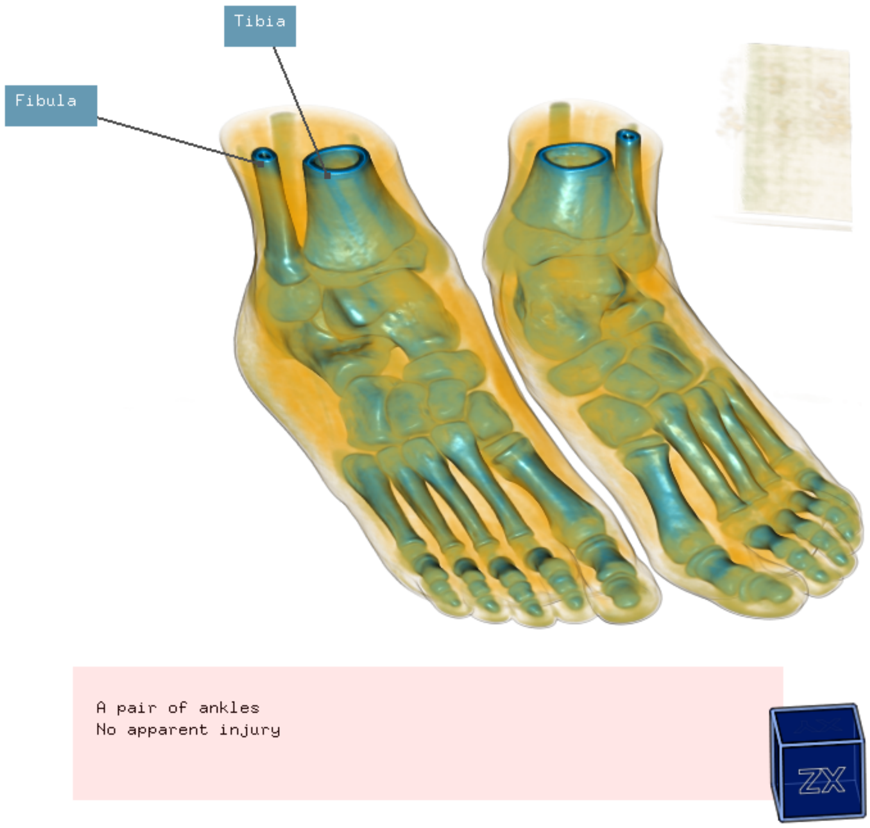


Figure 4.10: Labels added to aid audience in interpreting the visualization. Labels work in 3D, meaning that they will point to the designated location regardless of any transformations done to the data set. Dataset courtesy by Casimage [Cas].

## 4.6 Story templates

As we know, creating visualizations can be time consuming, and inherently the creation of whole visualization stories will demand its time. Time is of the essence, especially in the medical community where, for example, a doctor has a limited window of opportunity for documentation (or creation of stories) with respect to his/hers patients. The idea behind the story template is to increase the re-usability of stories, and present some generalized pre-defined storiy types that the creators of visualization stories can use to save time when creating their own visualization stories. We will present 5 such story templates, that will make story creation more effective and also pose as ideas for further creation of the visualization stories.

**Overview & zoom**

A general, but very powerful story template relies on Ben Shneiderman's information seeking mantra [Shn96] *Overview first, zoom and filter, details on demand.* The mantra states that first and foremost, the visualization needs to give the audience an overview of the situation (data) to allow the audience to complete their mental map of the visualization. This can be achieved by zooming out to the point where the audience can get an overview of the data, and perform a 360 degree rotation of the data (or indeed moving the camera around the data instead).

Secondly zoom and filter states that the visualization should guide the audience toward interesting features by simplifying the data at hand. This can be achieved by limiting the visible structures by tweaking the transfer function, or through focus+context techniques by specifying the importance of the data.

Last the mantra states that details on demand should be given. By adding annotations, labels and oral comments, details can be explicitly defined and presented. The story template as such focusses on presenting the data in a way that allows the audience to have a complete overview first, then filtering the data to the point where any interesting regions or similar can be identified, and lastly to focus on these regions while adding content to aid in presenting these findings.

**Comparative stories**

Comparison is a natural way of presenting likenesses or diverging aspects of similar objects. A story template for visual comparison of the data would be a comprehensible way for presenting the data for an audience. Inspecting and fully understanding a medical trauma will only be completely comprehensible when the audience is presented healthy tissue/organs for a visual comparison.

These comparative stories can focus on several interesting differences within the medical domain. Examples for comparable data includes, healthy versus sick, young versus old, left eye versus right eye (or any other part for that matter) or pre-operation versus post-operation.

Comparative stories can be presented in several ways, the most obvious way is to visually present the two opposing data sets side by side, but they could also be presented one on top of the other to reveal anomalies or any other differences with respect to the data.

**Iterative stories**

Whenever working with multiple similar pathologies, exemplified by multiple lymph node metastasis, an iterative scheme type of story would be interesting. The story would present the metastasis one at the time, in a sequential manner to allow inspection and visual comparison between the nodes.

The story can devise a scheme that is executed for every lymph node containing a specified set of visualization changes,
i.e. *for each* **medical pathology** *do.* The story would then execute this instruction set for every node, presenting them in the same matter for an easy comparison.

**Copy & Paste**

Whenever similar trauma or common medical pathology is to be visualized much time can be saved by making a template that would fit by just switching in the data at hand. Many simpler cases of trauma and pathologies raise the same can very easily be presented the same way, and will probably raise the same types of questions. The visualization stories can then link to medical lexicons available on the internet, or indeed any other available information. This will answer most questions regarding the pathology at hand.

These copy & paste templates work best when dealing with trauma that occur frequently. Van Wijk [vW05] states that the value of a certain visualization increases when the re-usability is high, and the amount of work for generating this visualization is low.

**Varying representations**

There exists many different ways of visualizing 2D or 3D data. In many cases we want to use multiple of these visualization techniques to unveil the information inside the data. A story template consisting of presenting the data by different techniques, either side by side (with differing visualization techniques), in a sequential manner and changing techniques after a period of time, or by fusing images created using different types of visualization techniques on top of each other.

An example of such varying representations could be using regular direct volume rendering (DVR) in addition with maximum intensity projection (MIP).

# 5 Demonstration stories

> *Stories have power. They delight, enchant, touch, teach, recall, inspire, motivate, challenge. They help us understand. They imprint a picture on our minds.*
>
> - Janet Litherland

In this chapter we present two example stories that show the potential of volumetric storytelling within the medical domain. These two stories we will present is the Ischemic Stroke story and the Metastatic Lesions story.

The Ischemic Stroke story the patient is the subject of an ischemic stroke. Blood supply was restricted, probably due to a constriction or blockage of the blood vessels, causing the organ (in this case the brain) a shortage of oxygen and nutrients. The patient was examined shortly after the incident, and it was indicated that a large region of the brain, more specifically in the left hemisphere of the brain, was classified as tissue at risk. Such tissue at risk is tissue that will become dysfunctional or die if not treated properly. To readily identify such affected tissue, a contrast agent is injected into the patient while he/she is undergoing the examination in the MR. Most tissue is salvageable if treatment is started at the earliest possible stage in time. Available treatment was issued to the patient, the follow-up examinations were done in the following hours to ensure that the treatment was working and hopefully salvaging as much of the brain functionality and tissue as possible. Dataset courtesy by Haukeland university hospital.

The Metastatic Lesions story is a story of a patient with malignant tumor in multiple organs. The patient has undergone different examinations, and the data from these examinations has been concatenated to allow better interpretation by allowing data from both examinations to be compared for any spatial position.

The examinations done for this story is a CT-scan, and a PET-scan. The CT scan is for the story's purpose used as a spatial reference, since is easy to identify the different organs, skeletal structures etc. from this data, while the PET-scan shows the uptake of an injected contrast agent to identify the suspicious regions that could be identified as malignant tumor. By interpreting the PET-data the extent of the tumor can be identified and compared with previous results, thus revealing whether the sickness has spread or not. Dataset courtesy by Casimage [Cas].

The stories can both be categorized as doctor to patient type stories. Both their content is from the medical domain, and the stories will be explained in a simple manner to assist in clarifying the situation of these patients.

## 5.1 Ischemic Stroke Story

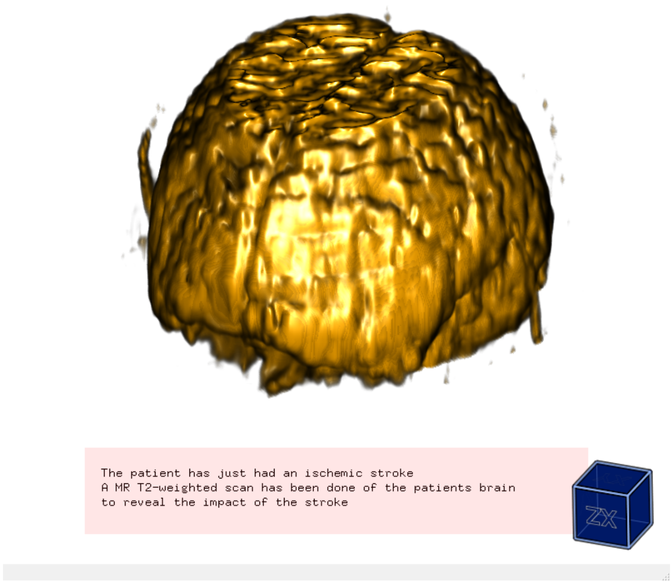


Figure 5.1: The story starts with a volume rendering of the patients brain. This volume is shown to give the audience a spatial reference to the story. Volume renderings are usually much more detailed (more slices across the volume is sampled), but here the focus of the examination is to sample the same spatial positions over time, e.g. time dependent data (to measure the uptake with respect to the contrast agent used).

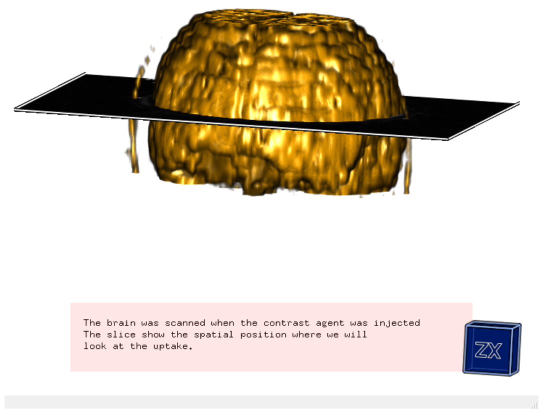


Figure 5.2: We rotate the volume and introduce a slice of data in the xy-plane of the data volume to give the audience a clear understanding where in the brain we are situated (as a spatial reference). During the examination of the patient, a contrast agent was injected to reveal how well the brain responds to the agent. Regions that does not respond to the agent, will have to be scrutinized further.

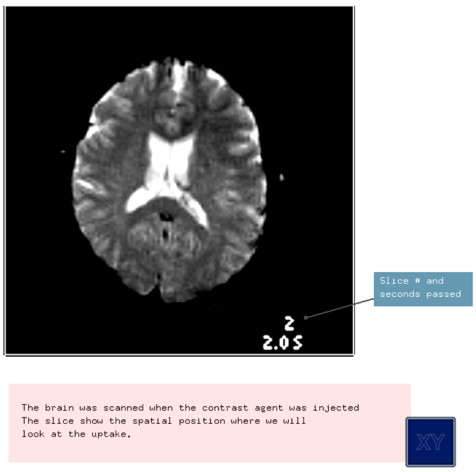


Figure 5.3: The volume rendering gradually disappears to allow better inspection of the slice data. The slice is transformed, rotated and zoomed upon, such that a clear orthogonal view is achieved. The following slices originates from the first examination done on the patient. Doctors suspect tissue at risk will be identified. No contrast agent has been injected yet.

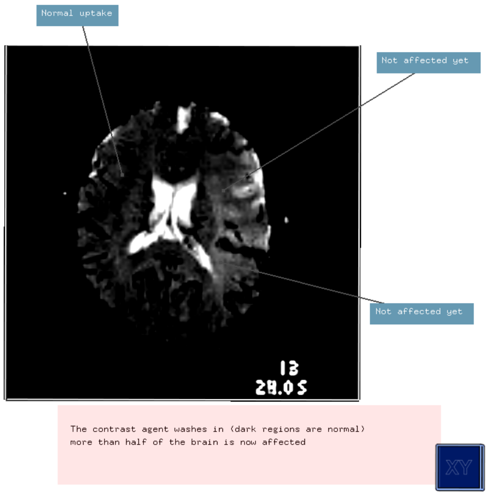


Figure 5.4: We iterate over the slices along the timeline and can readily identify the contrast agent affecting the brain on the left side of the image. In this data the dark regions represent the contrast agent taking effect with respect to the brain. We look for regions that remain unaffected when the contrast agent is injected.

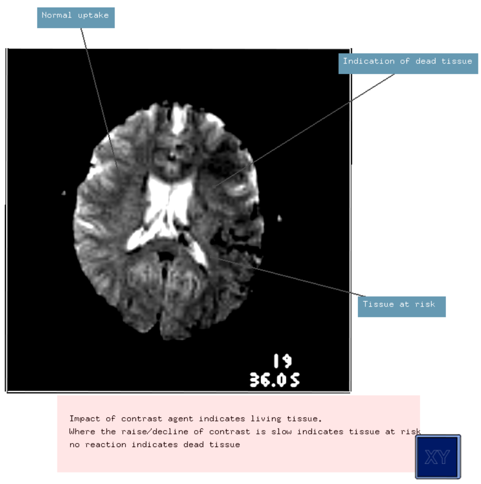


Figure 5.5: The contrast agent is starting to affect the right side of the brain as well. We can notice that there exist a region on the right side of the image where there is no reaction (or little) to the contrast agent. This is the tissue at risk. Above that region, is a region that is not affected at all, this is the dead tissue which we cannot do anything to save.

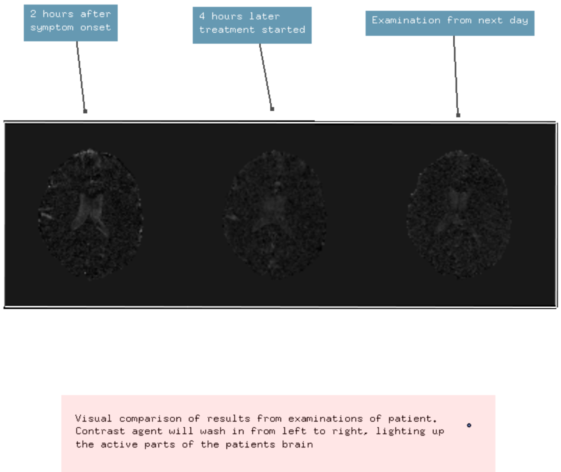


Figure 5.6: We issue thrombolytical treatment to the patient to salvage as much as possible of the identified tissue at risk. In this visualization we show slicedata from three different stages of treatment. The first slice is before the treatment, approx. 2 hours after symptom onset. The second slice show the patient after the treatment was started, approx 4 hours later. The last slice show the patient the next day after treatment.

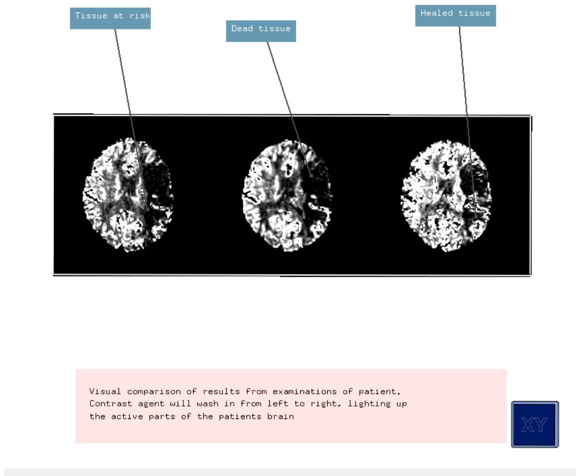


Figure 5.7: In these visualizations, the contrast agent uptake is shown as bright regions (in contrast to the previous images where they were dark). We can se the contrast agent uptake is highly active in the left hemisphere of the brain, and that it starts to affect the right hemisphere as well.

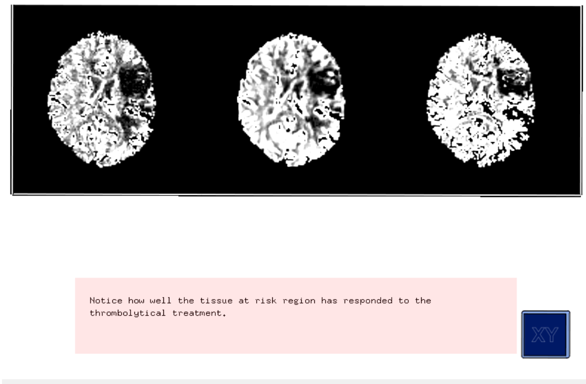


Figure 5.8: We compare the tissue at risk regions from the different visualizations, and can clearly identify that the uptake in the latter visualizations are much higher than before the thrombolytical treatment was started. This indicates successful treatment.

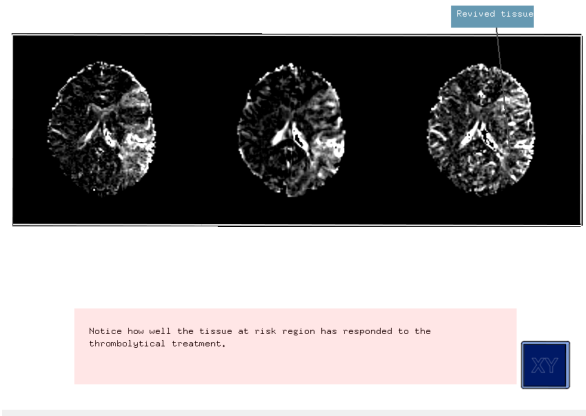


Figure 5.9: The contrast agent washes out from the right hemisphere, and we can see that the tissue at risk regions in the latter visualizations returns to normal much faster than the pre-treatment visualization. We can conclude that the treatment has helped revive and rescue important tissue, and thus brain activity, that would have otherwise been lost.

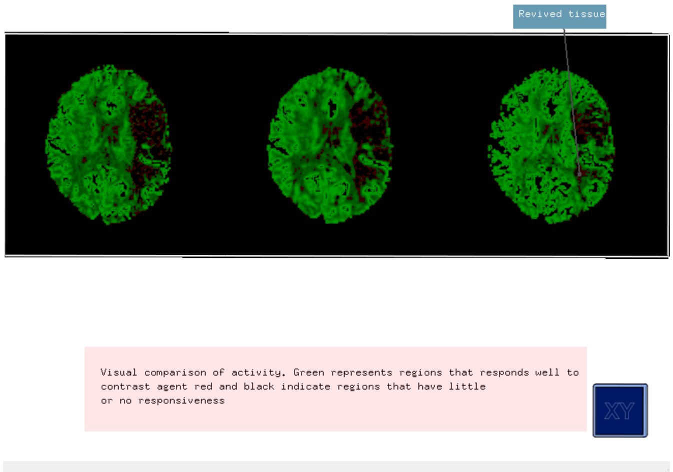


Figure 5.10: To support a better visual comparison of the three different examinations, the color mapping has been altered. The green regions depict the healthy regions within the patients brain, while the red indicate tissue at risk and dead tissue. It is evident that the thromobolytical treatment had a positive effect on the patient.

## 5.2 Metastatic Lesions Story


Figure 5.11: In this story we want to examine the extent of tumor in the patient. The patient has undergone CT and PET scans to reveal unnatural uptake of the contrast agent injected.


Figure 5.12: We show the CT and PET simultaneously to allow the audience to easily comprehend the spatial position of the PET data. We can readily identify several regions with high uptake.

Figure 5.13: Interactive approval is performed, we view the data from the left side of the patient. After approval, we return to the story.

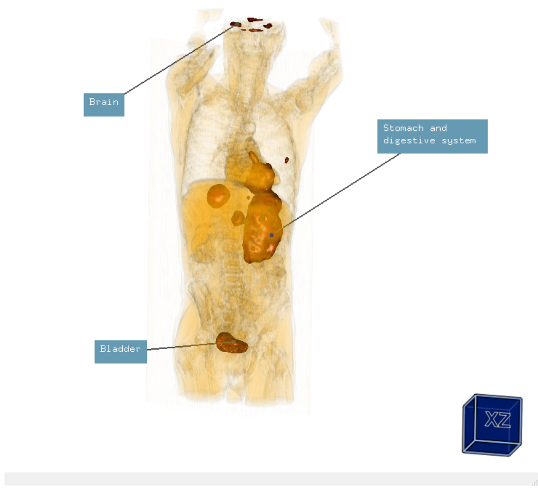


Figure 5.14: Not all regions with high uptake are to be considered malignant, it is normal to have a high uptake in the bladder, brain and digestive system. We use labels to aid the audience in interpreting the data.

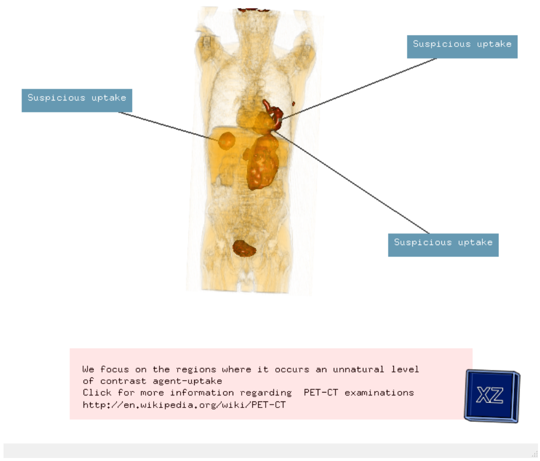


Figure 5.15: We disregard the regions we know normally have a high degree of contrast agent uptake, and focus on the remaining regions found in the PET examination.

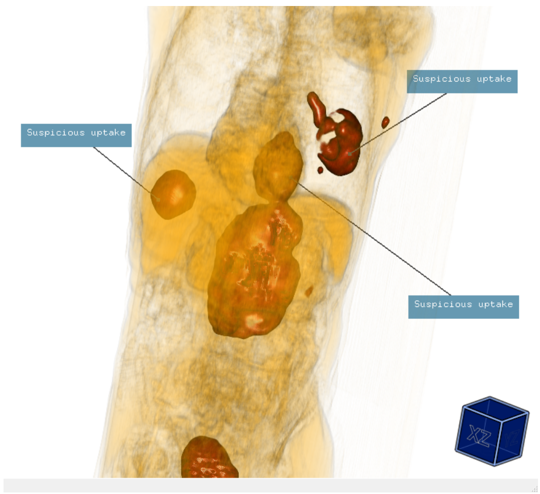


Figure 5.16: We rotate the volume to allow inspection from another interesting angle. We get a better notion of the extent of the tumor by inspecting from different angles.


Figure 5.17: To better indicate the extent of the pathology, and to determine to what degree the liver is affected, we introduce a slice along the xy-plane. This slice can be adjusted along the patients main axis (z-axis). It seems highly probable that the liver and lungs contain metastatic lesions.

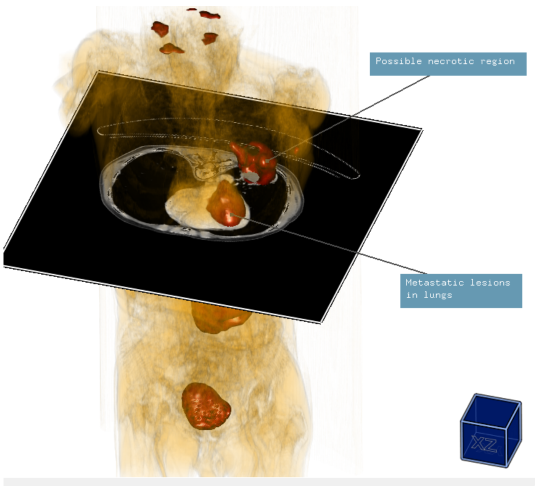


Figure 5.18: We use the slice data in conjunction with the CT and PET data. We can readily identify metastatic lesions in both lungs and liver.

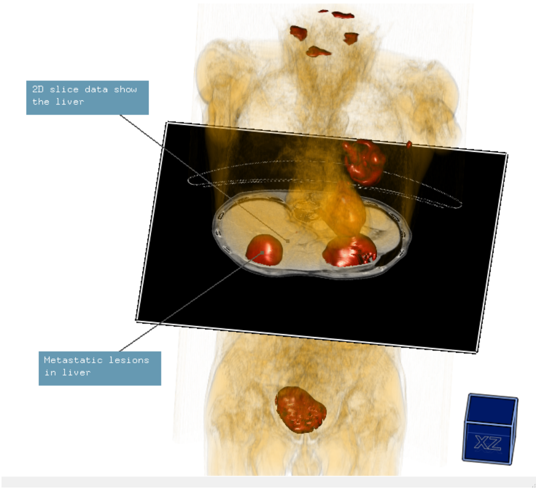


Figure 5.19: We investigate the region around the liver to be able to determine the extent of the tumor.

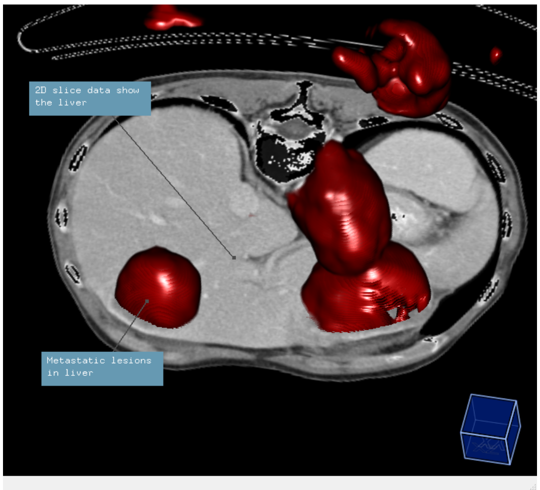


Figure 5.20: We remove the CT data to avoid further occlusion of the data, and rotate the slice data orthogonal to the view direction. We can now easily see the extent of the metastatic lesions.

# 6 Realization in VolumeShop

*I know only one thing about the technologies that await us in the future: We will find ways to tell stories with them.*

- Jason Ohler

This chapter describes how we implemented storytelling functionality into the visualization environment named Volumeshop. This implementation was done mostly using the C++ programming language, but also includes some small snippets of c-style code written as GLSL shaders. The choice of the programming language was implicit with respect to the already existing environment, but could have been another languages as well. Wohlfart & Hauser [WH07] implemented a similar functionality in a Java based visualization environment named RTVR made by Mroz et al. [MH01]. The plugin that we created to suit the Volumeshop environment can not only create, edit and play stories, but also is capable of handling a large variety of plugins and thereby also a great number of visualization parameters. Secondly, it is not limited to a fixed set of functionality as previous solutions were.

## 6.1 Volumeshop

Volumeshop is a framework for the creation of volume visualizations by Stefan Bruckner et al. [BG05]. Volumeshop comes with a magnificent volumetric raycaster that creates stunning visualizations at interactive framerates by taking advantage of shader units present in consumer level graphics hardware. The term "interactive" in the context of volume rendering refers to the immediate or instant creation of visualizations of the respective data, and should not be confused with interaction as discussed with respect to storytelling. A screenshot of the Volumeshop software can be seen in Figure 6.1. Most visualization frameworks

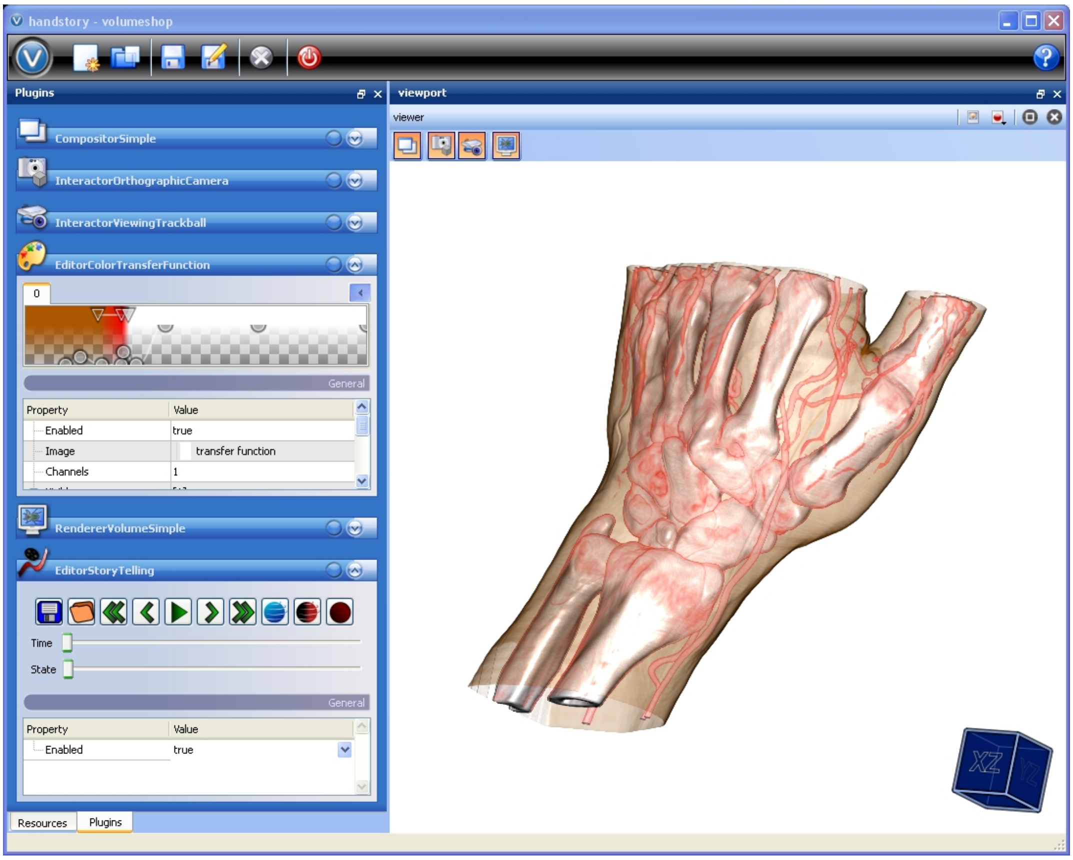


Figure 6.1: *The Volumeshop environment for volume visualization. The figure shows the hand dataset being visualized with differing opacity and color for skin, vessels and bone. Image courtesy of Stefan Bruckner et al. [BG05].*

contain a fixed pipeline for visualizations, meaning that they have a fixed set of functionality with respect to the handling of data and creation of visualizations. Volumeshop is designed to allow its pipeline setup to be expanded and enhanced with new plugin software. In short, any part of the pipeline can be changed and tweaked by adding new plugins to the framework. This allows for a continuous growth of visualization styles and techniques, and constant adaption to new types of data as well. This is an optimal environment for visualization stories because of its longevity due to the plugin based pipeline approach.

## 6.2 Storytelling Plugin

We have created a storytelling plugin named EditorStoryTelling for the Volumeshop environment. The plugin enables recording, editing and interactive playback of visualization states.

To be able to sufficiently overwatch (read and write) a large variety of visualization parameters, and thereby avoid the fixed functionality many visualization environments offer, the plugin must communicate with the program core. To be able to handle parameters that are not known a priori, the plugin and the environment, as stated earlier in Section 4.2, must successfully meet two criterion of two-way communication and variables accompanied with information regarding which type of variable it is.

In the volumeshop environment, every plugin keeps a record of every visualization parameters it has, all present in the QT graphical user interface it is required to have. We are able to access these lists of visualization parameters, and therefore can check whether there have been any changes to the visualization parameters and the amount of change done. Through the same interface, we can inject our own values for those visualization parameters, effectively controlling the environment as we want. The Volumeshop also provides a container class for every visualization parameter named Variant. These variants contain, in addition to the parameter itself in form of float, int, string, matrix, vector etc., information on what parameter it actually is. By using these variants, we can define rules for handling the different types of variants. E.g. it makes sense to treat a matrix and a float differently with respect to how to interpret and interpolate those values.

During interactivity the plugin keeps tally of which visualization parameters where changed since the last defined visualization milestone. Whenever a new milestone (key visualization in the story) is defined, the plugin automatically discovers the extent of the parameter changes by inspecting current and previous state, and calculates how those changes should be incorporated into the visualization story. By keeping track of the visualization parameters changed, we avoid the computationally expensive alternative of comparing all parameters available whenever a milestone is being produced.

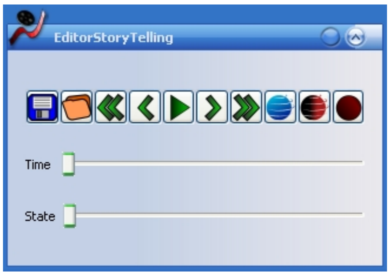


Figure 6.2: *Interface for the EditorStoryTelling plugin. Buttons are fashioned to resemble well known interfaces for media*

The story functionality the plugin offers can be accessed through the use of the GUI (graphical user interface) for the plugin. The GUI can be seen in Figure 6.2. A well known set of buttons are readily identifiable for controlling the story. The record button, located to the right, triggers the saving of new milestones in the story. By allowing the user to specify when the recording of visualization states should be done we allow for a more relaxed recording approach, since the environment lets the user control when the actual recording of the changes should be done. This is in contrast to active recording, where every user interaction usually is recorded, even though the accumulated changes might result in the same visualization state.

The user can now freely explore the data as he/she wishes without any concern with respect to the story, and record only whenever an interesting visualization states is found. A regular set of stepping buttons is also included, allowing the user to review previous (and upcoming) milestones in the story. Interactive sliders are included to allow the user to have an overview of where in the story we currently are (timeslider showing relative milestone compared to the total amount of milestones) and at what interactive state we currently are in during playback (stateslider showing the four interactive states defined by Wohlfart & Hauser [WH07]).

We included two convenience functions with respect to story overview, namely *world rotation* and *overview rotation*. These functions adds overview rotations to the story at the press of a button. The buttons issuing these different rotations, can be seen in Figure 6.2 where the blue button represent the world rotation, and the red button issue the overview rotation. The difference between these rotations, is that the world rotations issues a 360-degree rotation of the volume relative to the transformation already done, while the overview rotation tilts the volume gently 90-degrees to the left at first, then tilts the volume -180-degrees such that it is rotated to the right, before it returns to its original position.

Whenever a milestone is recorded, the plugin updates its own storytelling board, visible in figure 6.3. In the shown simple example story, the object is rotated and zoomed onto. A change of opacity is done to reveal the underlying structure, i.e., the bones. Finally the object is rotated to another view angle, and another visualization technique is applied while the opacity is tweaked further. The resulting visualization story is a gentle flow of events from the beginning to the end.

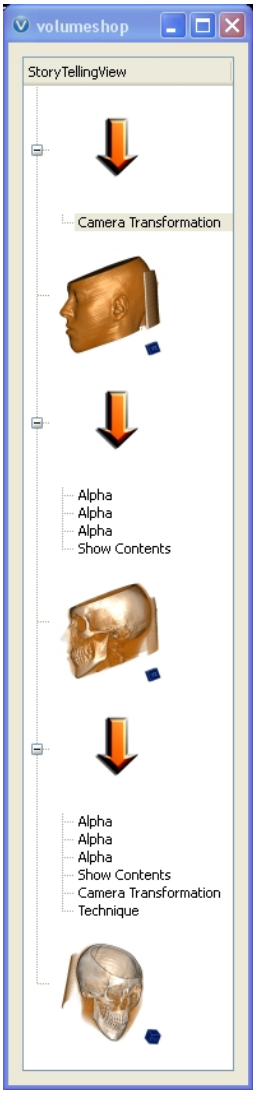


Figure 6.3: *Storyboard for a visualization story in the Volumeshop environment.*

In the storytelling board-pane, every milestone is identifiable with its own thumbnail image (of that specific visualization). Between the milestones, the transitions are presented with textual representations to enable support for interpreting the changes done in that specific transition. Both milestones and transitions found in this external storytelling board-pane are interactive, meaning that by right-clicking on a respective item, operations like timing changes (running time for transitions and waiting time for milestones) and deletion can be done. The intent for the storytelling board is to allow the user an overview of the story, while at the same time allow interaction with the story elements.

## 6.3 Handling interactive playback

Since the visualization environment is designed to produce results as fast as possible, it was apparent that controlling the environment would have to be done in a separate processor thread. Through using this thread, we gain the ability to steer the changes of the visualization ourselves. We induce a small delta of change and issue a redraw to the environment itself to achieve this. Since the framework takes advantage of the parallelism of the shader units, it responds fast enough that we readily can produce multiple images a second and thereby achieving animation like behaviour.

Through this thread, interaction can also be introduced. We can pause, make changes or diverge from the already existing stories easily by just do appropriate changes to the environment from the thread. It is apparent that playback of the visualization stories must follow some specification. The algorithm we designed to readily playback any story in a visualization environment is described in Algorithm 1. The algorithm visits every StoryActionGroup sequentially from the beginning to the end of the story and performs the containing StoryActions accordingly. When executing the StoryActions in a StoryActionGroup, each $StoryAction_{\Delta}$ is calculated dependant on the range of values the StoryAction should reproduce. This delta is iteratively added, for all StoryActions, while repainting of the visualization is being done. The output of the algorithm produces the animating behaviour sought after. It should be noted that not all visualization parameters are linearly interpolated this way, such an approach would make sense for variables like floats, ints and reals, but would for example be handled differently for matrices, vectors and strings. Nevertheless, the general playback model is done according to the algorithm explained.

**Algorithm 1**: The playback algorithm for iterating over the StoryActionGroups contained in the story and performing delta changes according to the StoryActions present in the StoryActionGroups.

**Input**: Story
**foreach** *Story Action Group in Story* **do**
  Steps $StoryActionGroup_{steps}$;
  Timer $T_{current} = 0, T_{end} = 1$;
  **while** $T_{current}$ != $T_{end}$ **do**
    **forall** *Story Actions in Story Action Group* **do**
      $StoryAction_{\Delta} = (StoryAction_{end} - StoryAction_{start})/StoryActionGroup_{steps}$;
      Perform $StoryAction_{\Delta} * T_{current}$;
    **end**
    $T_{current} = T_{current} + (1/StoryActionGroup_{steps})$;
  **end**
  Sleep at milestone for time $T_{sleep}$;
**end**
**Output**: Visualization Story

## 6.4 Sharing stories through XML

Creating visualization stories, and presenting them to an audience is undoubtedly a powerful way of communicating the information contained in the data. Being able to save and share those visualization stories would enable partners to asynchronously collaborate on the data and the visualization stories as well. By saving the visualization stories as XML (extensible markup language) accompanied with the data, the stories can with ease be sent between partners or published to a website, thus the foundation for asynchronous collaboration is made possible.

The structure of the XML template can be seen in Figure 6.4. The first part of the XML diagram shows the actual story elements, while the last part includes the framework specific state saves. By separating the story part and the framework dependant part that way in the XML output, this XML template can and will suit different visualization environments with a minimum amount of change.

The StoryActionGroups are listed in correct sequence, with id-number and information regarding the number of actions contained. Each StoryAction includes information regarding which plugin the visualization parameter resides from, what kind of parameter it is and the specific name given that parameter. Every StoryAction also has two Variants (framework specific types), one defining the starting value for the visualization parameter, and the other for defining the end value

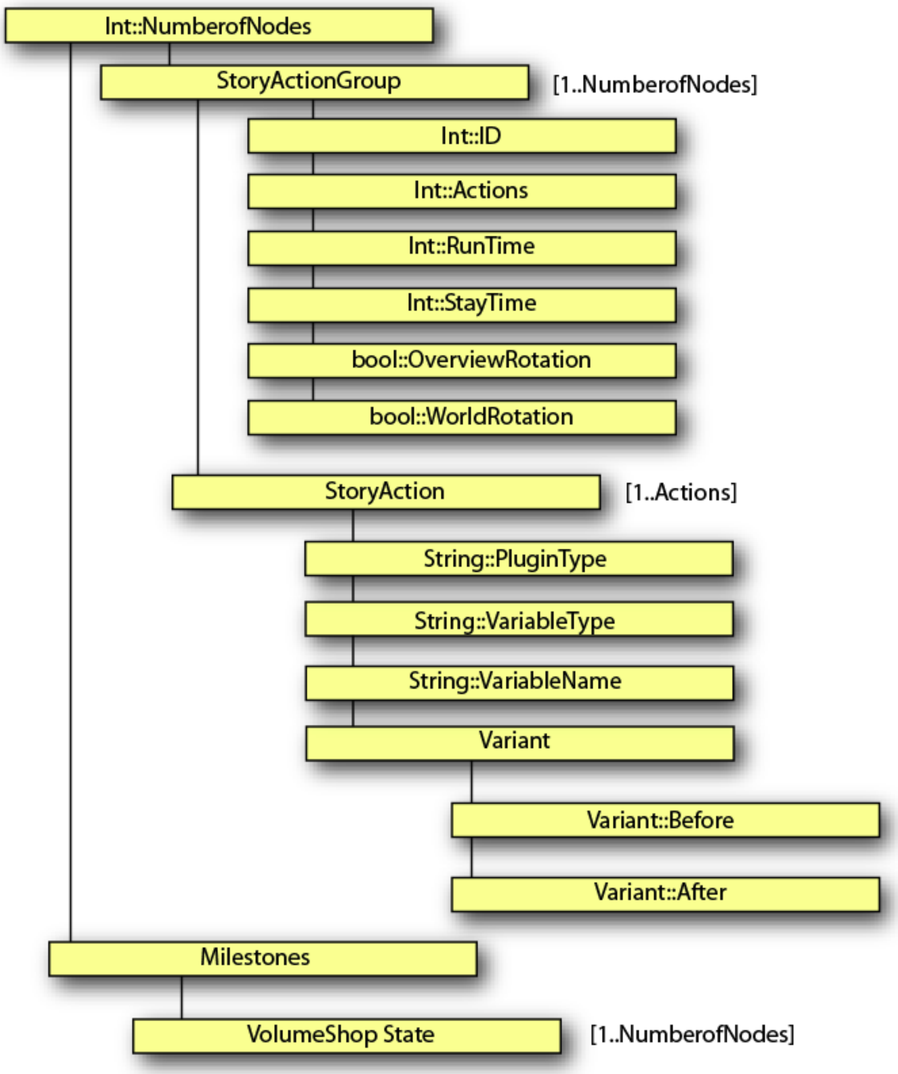


Figure 6.4: *An overview of the XML layout in context when saving and loading visualization stories.*

for the same parameter. The StoryActionGroups also includes timing information regarding length of transitions and wait time for the adjoining milestones.

## 6.5 Annotations, labels and commentary-plugin

Visualizations are often complex by nature, and require careful comments or descriptions to aid correct interpretation. Many visualization environments already have included textual labels and annotation possibilities to enrich their visualizations properly. We have made a plugin that not only can place labels with anchorpoints in 3D, but also is capable of inserting annotations to the visualizations and maybe most importantly is capable of recording and playing back oral comments to the visualizations. By using a variety of modalities when enriching the visualizations with information, there is significantly less chance of overloading one sense. The plugin enables the use of labels, annotations and oral commentary within the visualization stories as well. All functionality is contained in a separate plugin, and while it can be recorded by the storytelling plugin described earlier, it has no notion or knowledge of that process.

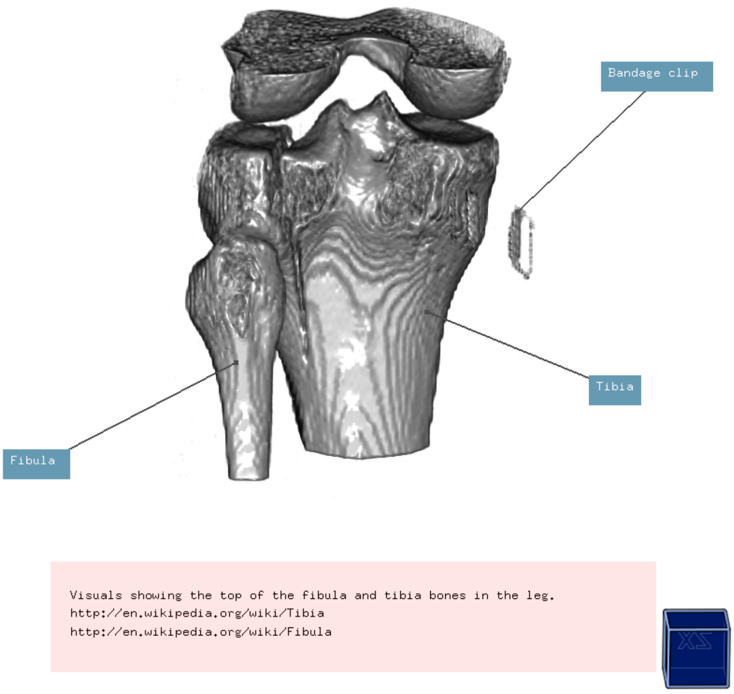


Figure 6.5: *A visualization enriched with annotation and labels. The annotation pane also works as a information gateway to external sources, e.g. hyperlinks.*

**Annotations**
Annotations enrich visualization by adding a frame containing text regarding the

visualization content. Although the annotation is simple, and contain text only, it is a prominent way to introduce general information about the data, or upcoming changes in a story. Such textual information can aid the audience in interpreting the scenes, and indeed the story as a whole. In the plugin the annotation is drawn in screen space, meaning that it is only in 2D, and is blended on top of the visualization itself. The annotation pane also functions as a information gateway to external sources if the creator of the visualization story has included hyperlinks to related material. By clicking the annotation pane during parts of the visualization story where hyperlinks are defined, the system opens a web browser and tries to connect to those pages. An example annotation can be seen in Figure 6.5.

**Labels**

The labels textual part resembles the annotations in the way that they are drawn in screen space. The labels anchor point, where the label points towards, is defined in 3D. This enables the visualization to change, i.e., a rotation of the data or change of camera position, while the label still points to the predetermined location. To achieve this 3D mapping whenever a label is created, a ray is shot into the data based on the where the mouse was clicked on the dataset. The ray traverses the data until it reveals a point which is non-transparent and recalculates the original position of that data by inversing the rotations already applied to the data. This way, we can guarantee that any label that points at the dataset itself, i.e. a visible point, will continue to point at that position regardless of any upcoming changes to the data. In the case where the user clicked outside the dataset, meaning that the ray could not hit any point, the system will use the screen space coordinates instead. An example of labels in use is available in Figure 6.5.

**Oral Comments**

Oral commentaries are made possible in the plugin by utilizing the already existing WINMM-API in Microsoft Windows. This API allow both recording and playback of soundfiles without the need of installing extra software or utilities. The recording is done during visualization story creation, and the playback is related to the milestones in the story. Whenever a milestone is encountered that contains recorded material, it will playback that material for the audience to hear. To active recording during story creation, the user must press a certain key combination on the keyboard to issue the actual vocal recording. When these keys are released, recording stops and the systems properly intertwines the recorded vocal commentary into the visualization story. During playback of the visualization stories the vocal commentary is played back automatically.

# 7 Summary

*A story is the vehicle we use to make sense of our lives in a world that often defies logic.*

- Jim Trelease

Nowadays data volumes are at a magnitude that make it impossible or even meaningless to inspect every data value. Visualization is a way of unveiling the information entangled in such data. One of the limiting aspects of visualizations, is the fact that visualizations are bound by the number of pixels (in case of a computer screen) or dots (in case of a printout) that can be used to represent, and ultimately state facts regarding the data.

A natural progress from showing single visualization images is presenting multiple visualization images in an animation like manner. By looking at multiple visualizations, the audience can more easily retain their mental map of the visualization changes, while following the gradual changes an animation sequence provides. Regular animations themselves have no interaction with respect to the content other than the inherent support of play and pause during playback.

We strive to find a presentation form that allows playback, but also provides a rich set of interaction possibilities. Storytelling is an old and versatile presentation form that has proven its usefulness, effectiveness and versatility over the last couple of millennia. By using the emotion and commitment present in storytelling, presentations of visualizations could really be taken to a higher level.

The success of a story, how well the knowledge content is being transferred, is not only dependant on how well the narrator presents the content, but also at which degree the audience is willing to accept what is being told. Stories relies on two critical criterions, that the stories are comprehensible, and that they are credible. Stories being comprehensible implies that the the audience understands the spatial relations of the data as well as the interrelated aspects of the data being presented, while credibility (does the audience actually believe what they see?) can be supported by allowing interaction in the presentations. This interaction result in a paradox named the narrative paradox, which arises when the narrator gives the audience the ability to influence the story themselves.

The visualization stories are structured with strong resemblances to storyboarding from the movie industry, and keyframing techniques from computer animation. Visualization stories define the important scenes in the stories as a sequential series of states that should be visited in correct order. Between these important scenes, an interpolation scheme makes gradual changes from one milestone in the story to the next, resulting in a gentle flow of changes the audience easily can interpret and cope with.

Visualization stories have qualities that currently few visualization framework possesses. Most frameworks can only change one visualization parameter at the time in a sequentially manner, while visualization stories can affect multiple visualization parameters in parallel. This is achieved by creating transitions between the key visualization states, that comprise of multiple changes which would occur simultaneously. The parallelism of events is also a strong visual cue for the audience, since multiple changes are interpreted as context changes, i.e. changes that are needed but not the focus of the story. Whenever transitions comprise of single gradual changes, the audience readily becomes more alert and interprets the change as step toward a climax of the story.

Our framework (Volumeshop) now supports an advanced story model that can handle arbitrary sets of visualization parameters. This is a huge advancement since until now visualization stories would only copy with a fixed set of such parameters. The visualization stories created in the Volumeshop environment are only limited by the visualization techniques supplied by the plugins used in the visualization pipeline. Since these plugins can be updated without tinkering with the storytelling model, the are no obvious limitations with respect to what the visualization stories can cope with.

We have enriched visualization stories by adding the possibility to use textual labels, annotations and oral commentaries to emphasize, clarify and distinguish content used in the stories. Labels and annotations are known from books and scientific diagrams, for example for explaining parts of the human anatomy or the cities in a country. Oral commentaries are seldom used (recorded and played back) whenever using computers, but their degree of personal touch and impact on presentations are undoubted. By using a variety of modalities (visual and aural), we significantly reduce the risk of overloading one sense. These tools can also easily be used to distinguish and focus on certain parts of the story, since the impact of visual and aural aids differ strongly. Through our story framework the annotation pane also works as an information gateway to external sources, i.e. hyperlinks. These hyperlinks will be opened automatically whenever the audience activates that functionality during a story.

Visualization Stories can be used as a collaboration tool between multiple partners. Since the stories can be edited, saved and exported quite easily, a possibility of asynchronous collaboration is given. The stories can be exported to the XML format, making stories compact and not dependant on any proprietary file format. The XML layout for visualization stories is structured to be as adaptable to possible, such that any visualization environment can save, and use the structure defined.

# 8 Conclusions

*A good story cannot be devised; it has to be distilled.*
- Raymond Chandler

Storytelling is an old and versatile presentation form that shown its usefulness, effectiveness and versatility over the last couple of millennia. We show that producing presentations for visualization in a storytelling fashioned way, results in presentation material that is not only more compelling, but also more intriguing.

The use of visualization stories makes presentations more comprehensible, since they rely on easily interpretable gradual transitions and avoid abrupt changes. Since visualization stories give the audience an overview of the data, and pace the story changes at a rate the audience easily can keep up with, more focus will be on story material rather than story structure. By relying on a variety of modalities (visual and aural) when enriching the stories with textual labels, annotations and oral commentaries, stories avoid overloading one sense and can use the these different modalities to emphasize important information. Since visualization stories allow the audience to participate in the story, to do active reinvestigation during playback interactively and insert own preferred changes to the story itself, the story maintains and retains its credibility.

Visualization stories help bridging the communication gap in the medical community. Daily communication, including communication between doctor-to-doctor and doctor-to-patient, as well as documentation and teaching will profit from using visualization stories as their medium.

As a medium, visualization stories are effective as a collaboration tool. The stories can with ease be created, edited, exported and shared between multiple partners for asynchronous collaboration.

At a personal level, i firmly believe presentation schemes like visualization stories will prove themselves very useful in this information age we live in. Every day we expect more and more information to be available to us, ranging from tax reports to medical history and indeed the data from any examinations. The information, or data, becomes harder to interpret, and the need for proper explanation and visual aids become evident.

I have a feeling this is not the end of the story..

# 9 Acknowledgements

I would like to take the opportunity to thank a multitude of people for participating and sharing knowledge with respect to the development of the storytelling plugin, writing and research done in context of this thesis.

Thanks to Stefan Bruckner for allowing us to use the impressive volume raycaster framework that he has built, and to all the helpful aid he gave in various stages of the plugin development.

Thanks to our local medical professionals; Martin Biermann, Odd Helge Gilja and Arvid Lundervold for participating in interviews regarding the usefulness and impact storytelling can have in medical treatment, planning and care.

Also a thank to Steffen Öeltze and Arvid Lundervold for providing the follow-up study datasets for the ischemic stroke patient and providing support for correct interpretation of the data as well.

Last but surely not the least, i want to express special gratitude toward my supervisor, professor Helwig Hauser, for the countless hours of discussion and supervision of this project. Thank you.

# A
# Interview Questionnaire

Here are the questions presented to the various medical personnel during the interviews regarding the usefulness of storytelling with respect to direct usage in the medical community. Names of the interviewees and a summary of the interviews can be found in section 2.6.

- How do you see storytelling could help your work?
- How can we improve storytelling to fulfill your needs for a presentation form?
- How much time do you have available for post processing?
- How much of that time would you be willing to devote to story creation?
- What do you consider visualization stories main strengths?
- What do you consider visualization stories greatest weaknesses?
- Do you think storytelling would be worth the investment of your time?
- Which templates could fit your daily work? (Overview+zoom, Varying representations, iterative stories, comparative stories, copy & paste)
- Would oral (voice) commentary help explaining features, medical trauma, in 2D/3D visualizations?
- Will the possibility of viewing two or more visualizations at the same time aid your work? Compare pre-post etc.
- What similar tools (if any) are available for you to use in your daily work? Any advantage/disadvantage of these?

- Can visualization stories be a nice supplement for teaching, research or documentation?
- Could handdrawn marks and/or measurements be helpful in visualization stories?

# Bibliography


[AHS05] K. Ali, K. Hartmann, and T. Strothotte, *Label layout for interactive 3d illustrations*, The Journal of WSCG **13** (2005), 1–8.

[BCS+05] Louis Bavoil, Steven P. Callahan, Carlos Eduardo Scheidegger, Huy T. Vo, Patricia Crossno, Cláudio T. Silva, and Juliana Freire, *Vistrails: Enabling interactive multiple-view visualizations*, IEEE Visualization, IEEE Computer Society, 2005, p. 18.

[BG05] Stefan Bruckner and M. Eduard Gröller, *Volumeshop: An interactive system for direct volume illustration*, IEEE Visualization, IEEE Computer Society, 2005, p. 85.

[BG07] S. Bruckner and M.E. Groller, *Enhancing depth-perception with flexible volumetric halos*, Visualization and Computer Graphics, IEEE Transactions on **13** (2007), no. 6, 1344–1351.

[BGK06] Stefan Bruckner, Soren Grimm, and Armin Kanitsar, *Illustrative context-preserving exploration of volume data*, IEEE Transactions on Visualization and Computer Graphics **12** (2006), no. 6, 1559–1569, Member-Groller,, M. Eduard.

[BH06] Dirk Beyer and Ahmed E. Hassan, *Evolution storyboards: Visualization of software structure dynamics*, icpc **0** (2006), 248–251.

[Bir08] Åsmund Birkeland, *View-dependent peel-away visualization for volumetric data*, Nov 2008.

[Bru06] Stefan Bruckner, *Exploded views for volume data*, IEEE Transactions on Visualization and Computer Graphics **12** (2006), no. 5, 1077–1084, Member-Groller,, M. Eduard.

[BTT05] D. Benyon, P. Turner, and S. Turner, *Designing interactive systems*, Addison-Wesley, 2005.

[But97] Andreas Butz, *Anymation with CATHI*, Proceedings of AAAI/IAAI '97, AAAI Press, 1997, pp. 957–962.

[Cas] Casimage, *http://pubimage.hcuge.ch:8080*, 2009.

[CMO+07] Jeanette Cordes, Konrad Mühler, Karl J. Oldhafer, Gregor Stavrou, Christian Hillert, and Bernhard Preim, *Szenariobasierte entwicklung eines chirurgischen trainingssystems*, eLearning in der Medizin und Zahnmedizin (Proceedings 11. Workshop der GMDS AG) (Aachen), Shaker Verlag, 17.-18. Mai 2007, pp. 17–30.

[Den05] Ivan Denisovich, *Software support for annotation of visualized data using hand-drawn marks*, IV, IEEE Computer Society, 2005, pp. 807–813.

[Eye] IBM Many Eyes, *http://manyeyes.alphaworks.ibm.com/manyeyes/*, 2004.

[GAHS05] Timo Götzelmann, Kamran Ali, Knut Hartmann, and Thomas Strothotte, *Form follows function: Aesthetic interactive labels.*, Computational Aesthetics (László Neumann, Mateu Sbert Casasayas, Bruce Gooch, and Werner Purgathofer, eds.), Eurographics Association, 2005, pp. 193–200.

[Gap] Google Gapminder, *http://www.gapminder.org/*, 2006.

[GHS06a] Timo Götzelmann, Knut Hartmann, and Thomas Strothotte, *Agent-based annotation of interactive 3d visualizations.*, Smart Graphics (Andreas Butz, Brian Fisher, Antonio Kröger, and Patrick Olivier, eds.), Lecture Notes in Computer Science, vol. 4073, Springer, 2006, pp. 24–35.

[GHS06b] Timo Gtzelmann, Knut Hartmann, and Thomas Strothotte, *Contextual grouping of labels*, In SimVis, SCS, Publishing House, 2006, pp. 245–258.

[GHS07] Timo Götzelmann, Knut Hartmann, and Thomas Strothotte, *Annotation of animated 3D objects*, Simulation und Visualization 2007 (Magdeburg) (Thomas Schulze, Bernhard Preim, and Heidrun Schumann, eds.), SCS Publishing House e.V., 8.-9. März 2007, pp. 209–222.

[Gla99] Andrew S. Glassner, *Active storytelling*, Computer Graphics International, 1999, p. 2.

[GP01] Nahum Gershon and Ward Page, *What storytelling can do for information visualization*, Commun. ACM **44** (2001), no. 8, 31–37.

[GS06] Dennis P. Groth and Kristy Streefkerk, *Provenance and annotation for visual exploration systems*, IEEE Trans. Vis. Comput. Graph **12** (2006), no. 6, 1500–1510.

[Hai07] Martin Haidacher, *Importance-driven rendering in interventional imaging*, Master's thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, August 2007.

[Hau05] Helwig Hauser, *Generalizing focus+context visualization in scientific visualization: The visual extraction of knowledge from data*, ch. 4, pp. 305–327, Springer Berlin Heidelberg, 2005.

[HGAS05] Knut Hartmann, Timo Götzelmann, Kamran Ali, and Thomas Strothotte, *Metrics for functional and aesthetic label layouts*, Smart Graphics, 5th International Symposium, SG 2005, Frauenwörth Cloister, Germany, August 22-24, 2005, Proceedings (Andreas Butz, Brian D. Fisher, Antonio Krüger, and Patrick Olivier, eds.), Lecture Notes in Computer Science, vol. 3638, Springer, 2005, pp. 115–126.

[HMSA08] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala, *Graphical histories for visualization: Supporting analysis, communication, and evaluation*, Visualization and Computer Graphics, IEEE Transactions on **14** (2008), no. 6, 1189–1196.

[Hul] Kevin Hulsey, *http://www.khulsey.com/*, 2009.

[IBHE+01] S. Iserhardt-Bauer, P. Hastreiter, T. Ertl, K. Eberhardt, and B. Tomandl, *Case study: medical web service for the automatic 3D documentation for neuroradiological diagnosis*, Proceedings of the conference on Visualization'01 (2001), 425–428.

[JRAY08] Mikael Jern, Jakob Rogstadius, Tobias Åström, and Anders Ynnerman, *Visual analytics presentation tools applied in html documents*, IV '08: Proceedings of the 2008 12th International Conference Information Visualisation (Washington, DC, USA), IEEE Computer Society, 2008, pp. 200–207.

[KMSZ06] Daniel A. Keim, Florian Mansmann, Jörn Schneidewind, and Hartmut Ziegler, *Challenges in visual data analysis*, IV, IEEE Computer Society, 2006, pp. 9–16.

[KNS04] Matthias Kreuseler, Thomas Nocke, and Heidrun Schumann, *A history mechanism for visual data mining*, INFOVIS, IEEE Computer Society, 2004, pp. 49–56.

[Koc98] T. Koch, *Story telling: is it really research?*, Journal of Advanced Nursing **28** (1998), no. 6, 1182–1190.

[KST01] Kehoe, Stasko, and Taylor, *Rethinking the evaluation of algorithm animations as learning aids: An observational study*, IJHCS: International Journal of Human-Computer Studies **54** (2001).

[LA03] Sandy Louchart and Ruth Aylett, *Solving the narrative paradox in ves - lessons from rpgs.*, IVA (Thomas Rist, Ruth Aylett, Daniel Ballin, and Jeff Rickel, eds.), Lecture Notes in Computer Science, vol. 2792, Springer, 2003, pp. 244–248.

[LC87] W. Lorensen and H. Cline, *Marching cubes: A high resolution 3D surface construction algorithm*, Computer Graphics **21** (1987), no. 4, 163–169.

[Lev88] Marc Levoy, *Display of surfaces from volume data*, IEEE Computer Graphics and Applications **8** (1988), no. 3, 29–37.

[LRA+07] Wilmot Li, Lincoln Ritter, Maneesh Agrawala, Brian Curless, and David Salesin, *Interactive cutaway illustrations of complex 3d models*, ACM Trans. Graph. **26** (2007), no. 3.

[LS08] Aidong Lu and Han-Wei Shen, *Interactive storyboard for overall time-varying data visualization.*, PacificVis, IEEE, 2008, pp. 143–150.

[MBP06] Konrad Mhler, Ragnar Bade, and Bernhard Preim, *Adaptive script based animations for intervention planning.*, MICCAI (1) (Rasmus Larsen, Mads Nielsen, and Jon Sporring, eds.), Lecture Notes in Computer Science, vol. 4190, Springer, 2006, pp. 478–485.

[MD07] Stefan Maass and Jürgen Döllner, *Embedded labels for line features in interactive 3D virtual environments*, Proceedings of the 5th International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa, Afrigraph 2007, Grahamstown, South Africa, October 29-31, 2007 (Hannah Slay, Stephen N. Spencer, and Shaun Bangay, eds.), ACM, 2007, pp. 53–59.

[MH01] Lukas Mroz and Helwig Hauser, *Rtvr: a flexible java library for interactive volume rendering*, VIS '01: Proceedings of the conference on Visualization '01 (Washington, DC, USA), IEEE Computer Society, 2001, pp. 279–286.

[MJJ+05] Kresimir Matkovic, Mario Jelovic, Josip Juric, Zoltan Konyha, and Denis Gracanin, *Interactive visual analysis and exploration of injection systems simulations.*, IEEE Visualization, IEEE Computer Society, 2005, p. 50.

[MTS+04] Konstantin Melikhov, Feng Tian, Seah Hock Soon, Quan Chen, and Jie Qiu 0002, *Frame skeleton based auto-inbetweening in computer assisted cel animation.*, CW, IEEE Computer Society, 2004, pp. 216–223.

[Orr86] Julian E. Orr, *Narratives at work: story telling as cooperative diagnostic activity*, CSCW '86: Proceedings of the 1986 ACM conference on Computer-supported cooperative work (New York, NY, USA), ACM, 1986, pp. 62–72.

[PC05] P. Pirolli and S. Card, *The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis*, Proceedings of 2005 International Conference on Intelligence Analysis (2005), 2–4.

[PKH04] H. Piringer, R. Kosara, and H. Hauser, *Interactive focus+context visualization with linked 2d/3d scatterplots*, July 2004, pp. 49–60.

[PRS97] B. Preim, A. Raab, and T. Strothotte, *Coherent Zooming of Illustrations with 3D-Graphics and Text*, Proceedings of Graphics Interface **97** (1997), 105–113.

[RTB96] Bernice E. Rogowitz, Lloyd A. Treinish, and Steve Bryson, *How not to lie with visualization*, Comput. Phys. **10** (1996), no. 3, 268–273.

[Shn96] Ben Shneiderman, *The eyes have it: A task by data type taxonomy for information visualizations*, VL, 1996, pp. 336–343.

[Sho85] Ken Shoemake, *Animating rotation with quaternion curves*, SIGGRAPH Comput. Graph. **19** (1985), no. 3, 245–254.

[SK00] Sissel Guttormsen Schär and Helmut Krueger, *Using new learning technologies with multimedia*, IEEE MultiMedia **7** (2000), no. 3, 40–51.

[SM00] Heidrun Schumann and Wolfgang Mller, *Visualisierung - grundlagen und allgemeine methoden*, Springer Verlag, 2000.

[SPP+99] Rainer Schubert, Bernhard Pflesser, Andreas Pommert, Kay Priesmeyer, Martin Riemer, Thomas Schiemann, Ulf Tiede, Pascal Steiner, and Karl H. Höhne, *Interactive volume visualization using intelligent movies.*, In proceedings Medicine Meets Virtual Reality, 1999, pp. 321–327.

[Sto] Storyboardsinc, *http://www.storyboardsinc.com*, 2009.

[SvW08] Yedendra Babu Shrinivasan and Jarke J. van Wijk, *Supporting the analytical reasoning process in information visualization*, Proceedings of the 2008 Conference on Human Factors in Computing Systems, CHI 2008, 2008, Florence, Italy, April 5-10, 2008 (Mary Czerwinski, Arnold M. Lund, and Desney S. Tan, eds.), ACM, 2008, pp. 1237–1246.

[SWTS05] Petra Schulze-Wollgast, Christian Tominski, and Heidrun Schumann, *Enhancing visual exploration by appropriate color coding.*, WSCG (Full Papers), 2005, pp. 203–210.

[TC05] J. Thomas and K. Cook, *Illuminating the Path: The R&D Agenda for Visual Analytics*, 2005.

[TC06] JJ Thomas and KA Cook, *A visual analytics agenda*, Computer Graphics and Applications, IEEE **26** (2006), no. 1, 10–13.

[TFTN05] S. Takahashi, I. Fujishiro, Y. Takeshima, and T. Nishita, *A feature-driven approach to locating optimal viewpoints for volume visualization*, Oct. 2005, pp. 495–502.

[TMS+06] Christian Tietjen, Björn Meyer, Stefan Schlechtweg, Bernhard Preim, Ilka Hertel, and Gero Strauß, *Enhancing Slice-based Visualizations of Medical Volume Data*, Proceedings of the Eurographics / IEEE VGTC Symposium on Visualization (EUROVIS 2006, May 810, 2006, Lisbon, Portugal) (Dieter Fellner, Torsten Möller, and Stephen Spencer, eds.), 2006, pp. 123–130.

[TvSGSH02] Ulf Tiede, Norman von Sternberg-Gospos, Paul Steiner, and Karl Heinz Hhne, *Virtual endoscopy using cubic quicktime-vr panorama views.*, MICCAI (2) (Takeyoshi Dohi and Ron Kikinis, eds.), Lecture Notes in Computer Science, vol. 2489, Springer, 2002, pp. 186–192.

[Ven] Ventureforth, *http://www.ventureforth.co.za/trekking-hiking/cederberg/*, 2009.

[VK05] Ivan Viola and Armin Kanitsar, *Importance-driven feature enhancement in volume visualization*, IEEE Transactions on Visualization and Computer Graphics **11** (2005), no. 4, 408–418, Member-M. Eduard Groller.

[vW05] JJ van Wijk, *The value of visualization*, Visualization, 2005. VIS 05. IEEE (2005), 79–86.

[VW06] F. B. Viégas and M. Wattenberg, *Communication-minded visualization: A call to action.*, Tech. Report 4, 2006.

[vWN04] Jarke J. van Wijk and Wim A. A. Nuij, *A model for smooth viewing and navigation of large 2d information spaces.*, IEEE Trans. Vis. Comput. Graph. **10** (2004), no. 4, 447–458.

[WEE03] D. Weiskopf, K. Engel, and T. Ertl, *Interactive clipping techniques for texture-based volume visualization and volume shading*, IEEE Transactions on Visualization and Computer Graphics **9** (2003), no. 3, 298–312.

[WH07] Michael Wohlfart and Helwig Hauser, *Story telling for presentation in volume visualization.*, EuroVis (Ken Museth, Torsten Mller, and Anders Ynnerman, eds.), Eurographics Association, 2007, pp. 91–98.

[YXRW07] Di Yang, Zaixian Xie, Elke A. Rundensteiner, and Matthew O. Ward, *Managing discoveries in the visual analytics process.*, SIGKDD Explorations **9** (2007), no. 2, 22–29.