Eurographics/ IEEE-VGTC Symposium on Visualization (2007)
Ken Museth, Torsten Möller, and Anders Ynnerman (Editors)

Story Telling for Presentation in Volume Visualization

Michael Wohlfart† and Helwig Hauser‡

VRVis Research Center, Vienna, Austria

Abstract

In this paper we present a novel approach to volume visualization for presentation purposes that improves both the comprehensibility and credibility of the intended visualization message. Therefore, we combine selected aspects from storytelling as well as from interactive volume visualization to create a guided but at the same time interactive visualization presentation approach. To ease the observer's access to a presented visualization result we not only communicate the result itself, but also deliver its creational process in the form of an annotated visualization animation, which we call a visualization story. Additionally, we enable variable means of interactivity during story playback. The story observers may just watch the presentation passively, but they are also allowed to reinvestigate the visualization independently from story guidance, offering the ability to verify, confirm, or even disapprove the presented visualization message. For demonstration purposes, we developed a prototype application that provides tools to author, edit, and watch visualization stories. We demonstrate the potential of our approach on the basis of medical visualization examples.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

1. Introduction

While in the past years the focus of the visualization community was more directed to visual exploration and analysis, growing effort is being put into visualization for presentation purposes recently. According to Thomas and Cook [TC06], production, presentation, and dissemination of analysis results are often the most time-consuming parts of data analysis. Also the time spent on the visualization user side increasingly becomes a crucial factor – in many application domains, e.g., in medical applications or when monitoring streaming data, it is very important that visualization users can come to their conclusions as fast as possible, usually very little time is available for free and long-lasting interaction. As a consequence, the development of new means to facilitate the production, dissemination, and presentation of visualization results has to be an integral part of future research in the visualization community.

When considering the communication of results from data exploration/analysis through a visualization presentation, two key aspects can be identified to assure that the presented information is well understood by the designated audience. First, the visualization message needs to be communicated *as clear as possible* in order to make it *comprehensible* for the receiver. Due to personal, social, or cultural contexts of the audience, this poses a non-trivial problem. We know, e.g., that the interpretation of colors varies across the population. Also, the more information a visualization result actually has to communicate, the more challenging it gets on the user side to fully understand the visualization message. This is especially true in 3D visualization where 3D renderings easily get complex and difficult to understand (without any interactive means, e.g., to rotate the visualization). Second, the visualization message will be stronger if it is reproducible by the receivers (e.g., in education). Therefore the observer has to be given means to reinvestigate and verify the presented visualization results. Thereby, observers can approve the presentation, leading to *more trust* into the presented visualization message.

† Wohlfart@VRVis.at
‡ Hauser@VRVis.at

1.1. The Basic Idea

Based on these considerations, our approach introduces a twofold improvement to visualization for presentation purposes.

Presentation Comprehensibility. In order to improve the comprehensibility of the presented visualization result we not only provide the result itself. Additionally, we also convey its buildup process, modeled after the paradigm of a story. The art of storytelling has long been used by humans to hand over information, education, and experience from one generation to the next. The potential of stories to comprehensibly convey great quantities of information in relatively few words is what we want to exploit in our approach. Traditional storytelling is based upon the concept of story plot – the general and sequential outline of the basic events in a story. By interpreting the user interaction events of an interactive data exploration session as the basis for such a story plot, we arrive at a simple, understandable, but very compelling metaphor for the generation process of a visualization result–the visualization story.

Presentation Credibility. The second innovation aims on improving the credibility of the presented visualization result. Therefore we propose a presentation scheme, which is guided through the visualization's generation process in the form of a story plot, but at the same time also is sensitive to user interaction. These concurrent influences on the visualization form a seemingly paradox situation, which is also known as the *narrative paradox* in the field of virtual storytelling [Ayl99]. In the context of visualization stories, the resolution of this paradox offers a great opportunity to improve on the credibility of the presented visualization result. This can be achieved by either alternating or splitting the control over the visualization between the observer and the story plot. This means, that the visualization is either (1) fully controlled by the story plot or the observer or (2) partly controlled by the observer and the story plot at the same time (so that, e.g., viewing is controlled by the observer, all other parameters are story-controlled). Both of these semi-flexible user interaction schemes, which we offer in our approach, give the presentation observer the ability to leave the story plot, reinvestigate some aspects of the dataset, and return to the story plot afterwards. This reinvestigation by the story consumers leads to additional insight into the data as well as to trust into the presented visualization message, which is significantly more difficult by only passively watching a presentation.

Figure 1 shows still images from a visualization story example, which is based on a CT-dataset of a human head that has suffered a tripod fracture. According to the predefined story plot, the observer is provided with visual guidance through the dataset as well as with annotational information on specific data parts of interest.

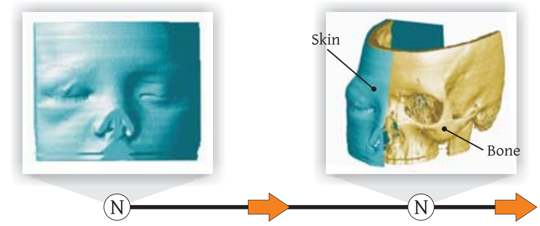

(a) **Overview.** The story is based on data from a CT-scan of a human head (left image). Applying partial clipping to the visualization, the two data representations contained in the presentation are revealed, one for the skin and one for the bone (right image).

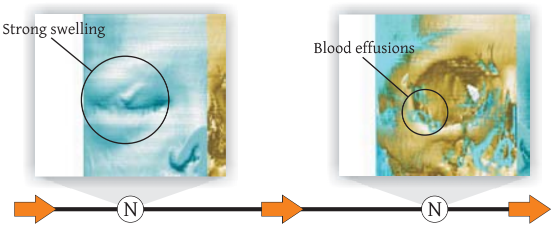

(b) **Zoom and filter.** Zooming in on the right eye, a strong swelling becomes visible (left image). Reducing the visual prominence of the skin (contour) exposes some blood effusions in the swollen region.

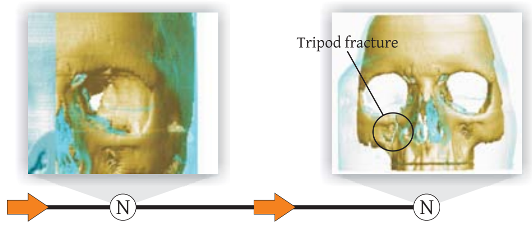

(c) **Visualization result.** For comparison to the swollen right eye, the non-injured left eye region is presented in more detail (left image). Finally, the cause of the swelling and the effusions is shown to be a tripod fracture (right image).

Figure 1: ***Visualization story example.*** This figure shows an example of what we understand as a visualization story. The story plot here is based on Shneiderman's information seeking mantra [Shn96], providing an overview first, then zoom and filter of specific objects and finally details on demand implicitly through interactive story approval.

1.2. Related Work

Although using the metaphor of a story in a visualization presentation context is a fairly novel approach, visualization has been used for presentation purposes in different ways for a long time already. The easiest (and also most unspectacular) possibility to present visualization results is by using images. But according to Gershon and Page [GP01], single images are also susceptible to uncertainties (e.g., occlusion) and might require some explanation or data reinvestigation for clarification.

In order to integrate more information into a visualization presentation, the use of animations is a useful choice. However, animations also have their issues. Setting up a meaningful animation is a complicated and time-consuming task, which is often not feasible in an application environment (e.g., medical presentations in hospitals). As a possible solution, Iserhardt-Bauer et al. [IBHE*01] introduced an automated web-service for standardized animation generation in the context of a specific medical application. This client-server approach converts a provided medical dataset (in addition to some parameters) into a complete set of standardized and reproducible animation sequences, handling the tasks of segmentation, rendering, as well as video generation. Though this is a straightforward approach for animation generation, it is highly specialized on a specific medical task, and leaves little room for changes through the animation's observer as well as its creator.

Early systems for the automatic generation of animation sequences, like ESPLANADE [KF93] are based on the hierarchical decomposition of predefined communication goals. A downside of these systems is their inflexibility regarding small animation changes. More flexible approaches use scripting to define the animation sequences. In his paper [But97], Butz introduces an approach for automatic generation of 3D animations, based on a high-level script language. This language is used to express specific communicative goals for the animation, which are translated into a low-level animation language according to a specific grammar. Another, more recent approach to script-based animation generation, that uses a simplified script-language, is presented by Mühler et al. [MBP06]. Their high-level language consists of a small set of high-level commands (e.g., the command `sceneOverview` gives a case dependent overview by rotating and moving around in the scene). Additionally, lower-level commands are provided, that allow for more control over specific parts of the visualization. Another possibility to embed more information into the visualization presentation is the use of textual annotations, as described by Hartmann et al. [HPS99, HGAS05]. In addition to the obvious additional textual information on the visualization, annotations are also useful to direct observer attention to particular areas of interest in the visualization.

An interesting approach to interactive visualization presentation is the use of Quicktime VR (QTVR) objects (e.g., Tiede et al. [TvSGSH02]). QTVR objects are basically a set of images, with specific scene parameters slightly varied from image to image. The images are packaged and viewed inside a media player, mapping specific mouse interactions to the scene parameters that were varied during the pre-rendering step and therefore mimicking true interactivity. However, one important limitation of QTVR is that it is impossible to sample the multi-dimensional parameter space of a visualization (viewing, iso-values, lighting, etc.) finely enough, which would be necessary for a thorough reinvestigation of the visualization during presentation – usually

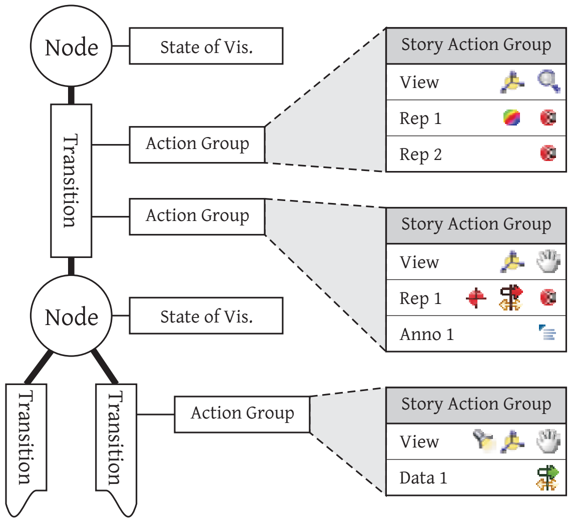

Figure 2: ***The visualization story model.*** This figure shows the elements of our story model and how they relate to each other to form a short sample story. Story nodes are connected through story transitions. Story transitions consist of one or multiple sequential story action groups. Story action atoms store the atomic visualization changes.

only one or two parameters are varied. Another problem with QTVR objects is that they require a lot of production effort, which is not desirable in real-life applications.

Another research field related to our work is the field of digital or virtual storytelling [Gla01, Per05]. The scope of virtual storytelling includes interactive narrative, virtual characters as well as drama and emotion. Therefore, we see our work only loosely tied to virtual storytelling, since we mainly focus on a tool to create as well as watch and interact with stories in a visualization context.

2. Visualization Stories

What we consider to be a *visualization story* is a set of linked *story nodes* which represent important stages (milestones) of the visualization build-up process. Figure 2 shows the composition of our story model on the basis of a short story example. The particular visualization scene, which is central to the story node, naturally is also annotated. Several types of annotations can be useful in this context, like feature-specific annotations (arrows plus text) or caption-like annotations (textual or verbal, which is quite common in medical diagnosis). Figure 3 shows possible annotation styles.

The individual story nodes are connected through *story transitions*, which account for the smooth visual metamorphosis between story nodes. Transition discontinuities as well as abrupt visual changes during story playback can disturb the observer's mental map, resulting in an unnecessary

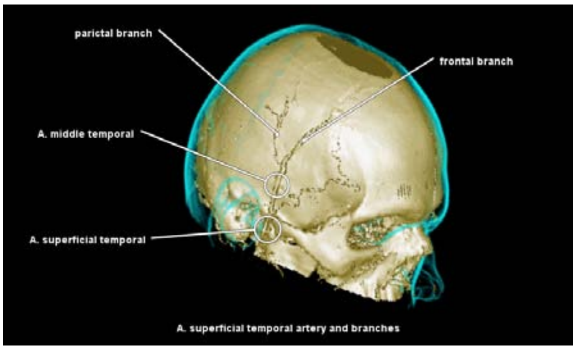

Figure 3: ***Annotation possibilities.*** This figure shows annotation styles available in our prototype. General text annotation (bottom). Feature-specific annotation (top).

interruption of story guidance. Therefore, seamless changes, which allow for a continuous adaptation of the mental map, are helpful to improve story comprehension. Story transitions are usually comprised of several sub-transitions (called *story action groups*), which are temporarily aligned one after the other. Story action groups contain a variable number of atomic changes to the visualization (called *story action atoms*), which are simultaneously executed during story playback. Additionally, a timing parameter specifies the playback duration for every story action group.

Story layout

Based on this simple visualization story model, several distinct story layouts can be envisioned. One useful guideline for a story layout in general can be derived from Ben Shneiderman's *information seeking mantra* [Shn96]. Such a visualization story gives an *overview* at first to introduce the story observer to the dataset (e.g., through a full rotation around the visualization or an outside-in view with slight rotation to improve 3D perception). *Zoom and filter* forms the second step of this general story layout, guiding the observer's attention to specific areas of interest. Here focus+context visualization plays an important role, since it offers means for concurrently focusing on specific details in the visualization while at the same time retaining an overview of the rest of the visualization to preserve the observer's mental map [Hau05]. The third step as described by Shneiderman, *details on demand*, is inherent in visualization stories through observer interaction. Nevertheless, it is often very useful to explicitly provide details on specific parts of the visualization as suggested by Shneiderman. Therefore our model offers means to integrate multiple story paths into the visualization story.

Comparative story layout. The comparison of distinct data parts forms a visualization scheme, which is helpful in many visualization application areas. In general, comparison is used to show the similarities and differences between two distinct data parts. These parts might originate from either (1) a single quasi-symmetrical dataset (e.g., the right injured eye vs. the left healthy eye from Figure 1) or (2) multiple datasets (e.g., from a pre- and post-OP scenario). A comparative story layout gives an overview of the original dataset at first and then concentrates on the specific data part that will be compared (e.g., by zoom and filter). The next step is a blend between the first and the second dataset (which must have been registered in advance). Finally, the specific differences between the two compared data parts are presented step by step.

Iterative story layout. Another story layout emerges, when a large number of similar data features have to be presented to the observer (e.g., lymph nodes in a neck dissection planning [KTH*05]). Here, every single feature is presented by basically the same visualization changes (e.g., zoom in on feature→ emphasize feature → rotate → zoom out). An iterative story layout summarizes all of these visualization changes by executing the same visualization changes for every feature (i.e., `for each` *feature* `do` *visualization changes*).

While this consideration on story layouts is far from complete, it makes clear that several visualization and presentation tasks can be simplified to specific story layouts. In our current prototype application, stories are generated solely by the story creator without building upon such basic story layouts. While this approach gives maximum flexibility to the creator, it also makes story generation more difficult. Therefore, we plan to extend our prototype to support story generation based on story layouts in our future research.

3. Story Telling and Interaction

Traditional approaches to visual presentation, which rely on images or animations, suffer from two major drawbacks. These are (i) the difficulty to recognize the exact spatial positions of and the relations between the visualized objects and (ii) possible occlusions of relevant parts. The integration of interactive elements into visualization presentation can solve these issues by providing means to reinvestigate the visualization interactively during presentation, leading to a better understanding of size, shape and positions of the visualization objects. It is one central experience from our work with volume visualization users that the ability to interactively adapt the viewing parameters (e.g., by rotating a visualization slightly to the left and the right) greatly improves the 3D perception of the visualization result. In order to provide these opportunities to visualization users, special care has to be taken to avoid specific problems of interactivity in a presentation context. Interactive exploration always exposes its users to the risk of losing sight of specific objects or orientation as a whole, due to either a lack of user experience or insufficient interaction techniques, as described by Bade et al. [BRP05]. To avoid these pitfalls, our approach provides interaction schemes, which distribute visualization control between the storytelling application and its user similar to the multiple levels of control proposed by Zeltzer et al. [Zel91]. The possible interaction schemes, which mark the key component of our approach, result from a blending between total story guidance and free user interaction. Four different kinds of story consumption are considered.

- *Passive story consumption*. Totally guided, similar to animation
- *Story playback with interactive approval*. Interactive excursions from the story for deep insight
- *Semi-interactive story playback*. Interaction simultaneous to guidance, extent of guidance determined by user
- *Total separation from story*. Exiting the story to totally free and interactive data exploration/analysis

While the main focus of visualization presentation through stories is directed to the blended interaction schemes, the traditional non-interactive or non-guided ones are still possible and useful. Figure 4 shows instances of the different interaction schemes as paths through the story traversal/story distance domain of a short story example. The default way of consuming a story is comparable to watching a non-interactive visualization animation (yellow path). The visual presentation traverses through the story plot, providing smooth animations during transitions. In story nodes, the traversal comes to a halt, giving the audience the chance to visually adapt to the visualization and conceive the presented visual and textual information. All visualization parameters are controlled by the story plot. For the interactive exploration of the story visualization, we provide means to (partly or fully) decouple from the predefined story path. Therefore, we at first categorize all possible visualization interactions according to their impact on the visualization scene.

- *Viewing interaction*. This category encompasses interactions which affect global visualization and viewing parameters. This means in particular, that interactions of this category do *not* influence *what* parts of the data are shown or *how* they are represented. The most important use of this interaction category obviously is to look at the 3D visualization from different sides (possibly under varying lighting conditions).
- *Representation interaction*. Interactions regarding the representation parameters of specific data parts form the second interaction category. Whereas the visualization question of *what* data parts to display is still not influenced, we here allow alterations of the way these data parts are visually represented (e.g., the users would alter the opacity of a visualization object to peer inside or change the visualization style for a particular object to improve on occlusion issues in the visualization).
- *Content interaction*. The third category, that brings the story consumer farthest away from the story path, subsumes interactions affecting the data parts themselves. This category contains, e.g., an iso-value change of a surface representation or the addition of a new object in the visualization which originally was not part of the visualization story.

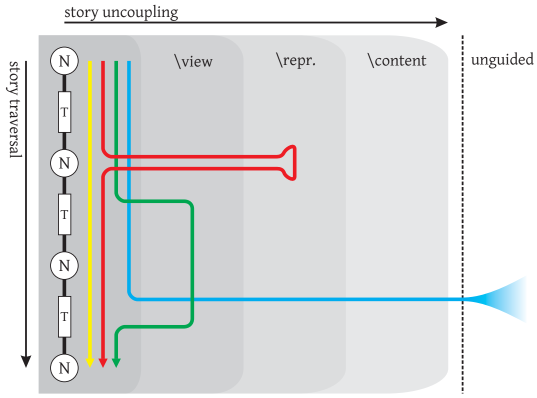

Figure 4: ***Storytelling and interaction*** This figure shows instances of the possible interaction schemes. Non-interactive story consumption (yellow). Story playback with interactive approval during a halt at a story node affecting representation parameters (red). Interactive visualization during story playback affecting viewing only (green). Total uncoupling from story guidance (blue).

By assigning the control over these categories (and their respective interactions) to either the story or the observer a blending between story guidance and interactive exploration can be achieved.

As soon as some part of the story especially attracts the attention of the story observers (through, e.g., interest or disbelief), they are able to pause playback and explore the current visualization interactively. Depending on the type of interactions performed by the observers, the according interaction categories are decoupled from the story path and control over them is assigned to the observers. After finishing the interactive investigation, all interaction categories are recoupled to the story path again, bringing the visualization state back to the point where the story decoupling took place. We call this particular form of interaction *story playback with interactive approval* (red path). After decoupling from the story path and interactively reinvestigating the visualization, the observers are not forced to (fully) return to the starting point of their investigation. By restarting story playback while some interaction categories are still controlled by the story observers, only story actions concerning story-controlled interaction categories are executed during playback. Thus, story playback is directly combined with interactivity, resulting, e.g., in the playback of a story with user controlled viewing parameters. We call this particular form of interaction *semi-interactive story playback* (green path). As a consequence of letting the story consumer leave the story path to a certain extent, a total disconnection from story guidance is possible. Thus, the users are allowed to alter the scene in any possible way, turning the storytelling application into an interactive exploration/analysis tool (blue path).

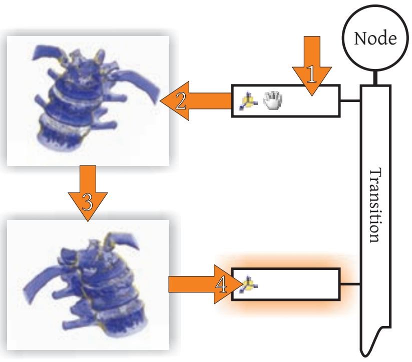

Figure 5: ***Interactive story authoring.*** This figure shows, how interactive story authoring works. The user navigates to a position in the story (1). The visualization is changed accordingly (2). The user changes the viewing angle (3). User interaction is recorded and incorporated into the story as a new story action group (4).

4. Story Authoring and Representation

While improvements to visualization presentation make up one important challenge for future research in visualization, another one is the integration of interactive visualization and the production of analysis results. According to Thomas [TC06], *production* is the creation of materials to summarize the results of an analytical effort. In order to reduce the time required to prepare analysis results for presentation as visualization stories, we present a technique that provides means to *generate* stories in the first place, as well as *refine and edit* already existing stories.

Story generation. As already stated, visualization stories represent the creational process of a visualization result. In order to collect this information during interactive data exploration/analysis, we propose an approach that records all user interactions into a sequential user interaction path, producing a first draft version of the desired story. This recording process starts like any interactive visualization method with the creation and interactive adjustment of one or multiple data representations. As soon as an interesting visualization is achieved, the story author enables story recording, resulting in the addition of the current visualization state to the (previously) empty story. From this point on, all user interactions are inserted into the story serially as story action atoms encapsulated into story action groups. An illustration of this recording process is shown in Figure 5.

Story editing. While story authors are allowed to produce large serial sequences of recorded interactions, they are not forced to record the whole story in one coherent recording step. When a first story draft has been recorded, the author might want to have a look at the recorded sequence. Therefore, navigational controls similar to a standard media player are provided. If desired, additional recordings can be made, which are incorporated into the story in a way that depends on the current story traversal position. If the current position is a story action group (belonging to a story transition), the recorded interaction is inserted at this location. If the current traversal position is a story node, a new story transition (containing the recorded interaction) is appended to the story node, resulting in a story branch.

In addition to images and animations, telling visualization stories also relies on additional textual information. In the story editing step, we provide two distinct ways to enrich specific story elements (i.e., story nodes or transitions) with text. On the one hand, general story text can be overlaid to provide the basic story content. On the other hand, specific parts of interest in the visualization can be annotated through labels. While other approaches to labeling try to extract optimal label placements from the data [AHS05], our approach is based on an interactive labeling scheme. Therefore, labels do not rely on pre-segmented data (i.e., an arbitrary number of labels can be added to the visualization), but are still anchored in 3D space. This is achieved by a special on-view mouse interaction. A first mouse-click shoots a picking ray into the scene delivering the 3D position of the label as the first intersection between any scene object and the picking ray. The following mouse-drag determines the size of the annotation marker (e.g., a circle). Finally, a second click specifies the relative 2D screen position of the label's text, which remains unchanged during story playback.

When authoring a story, all visualization interactions are serially integrated into the story as story action atoms contained in story action groups. In order to parallelize some atomic scene changes, story action groups can be merged together. This is achieved by joining all story action atoms (i.e., atomic visualization changes) from two successive story action groups into one single group. The resulting animations are not only nice to look at, but also reduce overall playback time.

5. Implementation

In order to demonstrate the potential of our visualization story approach, we implemented a prototype application based on the Real-Time Volume Rendering Java Library (RTVR) [MH01]. The main parts of the implementation process were directed to improvements of the underlying RTVR renderer, the implementation of the story model, as well as the integration of story actions and user interactions. Figure 6 shows the basic layout of the application's UI and describes its most important parts.

For the purposes of storing and distributing visualization stories, we developed an XML scheme, that resembles the story model defined in Section 2. It is designed in a way that the general visualization story outline is kept as application independent as possible. Additionally, individual applications are allowed to also incorporate more application-specific visualization changes into the story description.

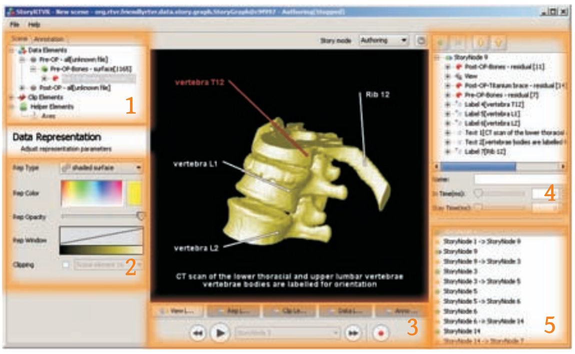

Figure 6: ***The storytelling prototype application*** The scene tree view is used to import data, derive data subsets and create visual data representations (1). Below are the parameter controls for the active data representation (2). The navigational controls influence story playback as well as story coupling (3). The story tree view shows the atomic visualization changes at the current story position (4). The story outline view gives an overview of the story plot (5).

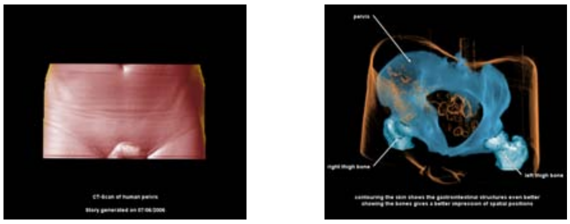

(a) **Overview.** The story is based on a CT-dataset of a human pelvis (left). Reducing the visual prominence of the skin (contour) reveals the bones and parts of the gastro-intestinal tract (right).

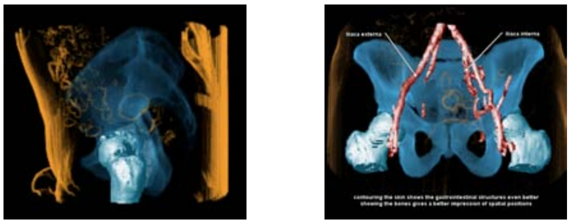

(b) **Interactive approval.** To better conceive the spatial relation between the visualization objects, the observer leaves the story path to interactively adjust the viewpoint (left). After returning to the story, the main blood vessels become visible (right).

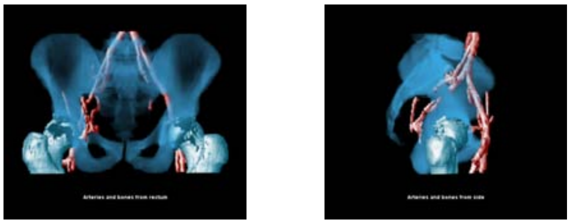

(c) **Zoom and filter.** Animation to the back and the side makes the relative spatial positions of bones and arteries clearer.

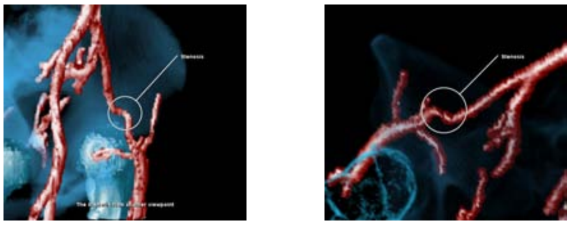

(d) **Visualization result.** A zoom animation guides observer attention to a specific point in the visualization, showing an anomaly. The narrowed width of the visible artery indicates a stenosis.

Figure 7: ***Visualization story example.*** A sample visualization story from our application prototype, based on a human pelvis CT-dataset.

6. Sample Story and Applications

In Figure 7 we present a sample visualization story based on a CT-dataset from a human pelvis. The underlying visualization consists of 4 elements, representing the skin, the thigh bones, the pelvic bone and the arteries in the region. The goal of this story is to (1) guide the observers through the visualization (2) put the contained visual representations into context to each other and finally (3) introduce them to an important feature in the data, a partly stenosis of the artery iliaca externa on the left side.

Since the production, presentation, and dissemination of results is of particular interest for visualization in general, application areas for our approach are widespread. Although in this paper, we mainly focus on the medical domain (i.e., medical diagnosis, presentation and education), an application of our approach is thinkable in many other areas (e.g., presentation of flow visualization data, drilling for oil). In general, visualization stories are useful wherever analysis results are gathered through visualization and need to be further distributed.

For additional story examples as well as further information on the visualization story project, see http://www.VRVis.at/vis/research/vis-stories/.

7. Conclusions and Future Work

In this paper, we introduced a new method for visualization presentation, which was inspired by the research fields of volume visualization and virtual storytelling. This new approach proposes the integration of storytelling into visualization for presentation purposes. Therefore a separation into a story authoring and a story telling step is necessary. In the story authoring step, users can explore a dataset, extract interesting information and arrange it in the form of a story. In the story telling step the users watch the story passively but are also allowed to interact with the story to a certain degree.

The application of our new data presentation approach leads to better understanding of the underlying data as well as information. The possibility of interaction, the ability to look at things from different angles and the feeling to still have some control over what is happening could pose a significant improvement to data presentation. Nevertheless, a user study is needed to shed more light on the possible improvements to presentation comprehensibility and credibility that visualization stories offer, and should therefore be conducted in the future. Additionally, we want to improve our prototype application for visualization stories by several means. Due to the fact that RTVR offers the possibility of web-based interactive volume visualization, an extension of our visualization stories approach to tell stories over the Internet is planned. Another web-based improvement would be the export of a visualization story as a webpage, featuring images of the visualized scene in the story nodes as well as the annotational information. A further extension would be the integration of verbal annotations into the prototype application. While textual annotations have already been implemented, verbal annotations, which are quite common e.g., in medical diagnosis, would improve on the applicability of our application.

8. Acknowledgements

This work has been carried out as part of the strategic research on visualization at the VRVis Research Center in Vienna, Austria (http://www.VRVis.at/vis/), which is partly funded by an Austrian research program called Kplus. The authors would like to thank Lukas Mroz for the support of their work. Additional thanks go out to Stefan Bruckner, and Ivan Viola, who contributed to this work by providing helpful thoughts and ideas. Furthermore, the valuable input of Knut Hartmann as well as Bernhard Preim needs to be mentioned.

References

[AHS05] ALI K., HARTMANN K., STROTHOTTE T.: Label Layout for Interactive 3D Illustrations. In *WSCG '05: Journal of the 13-th Intl. Conf. in Central Europe on CG, Vis. and Computer Vision* (2005), pp. 1–8.

[Ayl99] AYLETT R.: Narrative in Virtual Environments - Towards Emergent Narrative. In *Papers from the AAAI Fall Symposion* (1999), AAAI Press, pp. 83–86.

[BRP05] BADE R., RITTER F., PREIM B.: Usability Comparison of Mouse-Based Interaction Techniques for Predictable 3D Rotation. In *SG '05: Proc. of the 5th Intl. Symposium on Smart Graphics* (2005), pp. 138–150.

[But97] BUTZ A.: Anymation with CATHI. In *Proc. of AAAI/IAAI '97* (1997), AAAI Press, pp. 957–962.

[Gla01] GLASSNER A. S.: Interactive Storytelling: People, Stories, and Games. In *ICVS '01: Proc. of the 1st Intl. Conf. on Virtual Storytelling: Using Virtual Reality Technologies for Storytelling* (London, UK, 2001), Johnson C., Moorhead R., Munzner T., Pfister H., Rheingans P., Yoo T. S., (Eds.), Springer-Verlag, pp. 51–60.

[GP01] GERSHON N. D., PAGE W.: What storytelling can do for information visualization. *Communications of the ACM 44*, 8 (2001), 31–37.

[Hau05] HAUSER H.: Generalizing Focus+Context Visualization. In *Scientific Vis.: The Visual Extraction of Knowledge from Data*, Bonneau, Ertl, Nielson, (Eds.), Mathematics+Visualization. Springer, 2005, pp. 305–327.

[HGAS05] HARTMANN K., GÖTZELMANN T., ALI K., STROTHOTTE T.: Metrics for Functional and Aesthetic Label Layouts. In *8. Konferenz Elektronische Sprachsignalverarbeitung* (2005), pp. 115–126.

[HPS99] HARTMANN K., PREIM B., STROTHOTTE T.: Describing Abstraction in Rendered Images through Figure Captions. *Electronic Transactions on Artificial Intelligence 3*, A (1999), 1–26.

[IBHE*01] ISERHARDT-BAUER S., HASTREITER P., ERTL T., EBERHARDT K., TOMANDL B.: Case Study: Medical Web Service for the Automatic 3D Documentation for Neuroradiological Diagnosis. In *VIS '01: Proc. of the 12th IEEE Conf. on Vis.* (2001), IEEE, pp. 425–428.

[KF93] KARP P., FEINER S. K.: Automated presentation planning of animation using task decomposition with heuristic reasoning. In *Graphics Interface 93* (1993), pp. 118–127.

[KTH*05] KRÜGER A., TIETJEN C., HINTZE J., PREIM B., HERTEL I., STRAUSS G.: Interactive Visualization for Neck Dissection Planning. In *EUROVIS '05: Proc. of the 7th Eurographics – IEEE Symposium on Vis.* (2005), pp. 295–302.

[MBP06] MÜHLER K., BADE R., PREIM B.: Adaptive script based animations for medical education and intervention planning. In *MICCAI '06 - Medical Image Computing and Computer-Assisted Intervention* (2006), Springer, pp. 478–485.

[MH01] MROZ L., HAUSER H.: RTVR - A Flexible Java Library for Interactive Volume Rendering. In *VIS '01: Proc. of the 12th IEEE Conf. on Vis.* (2001), pp. 279–286.

[Per05] PERLIN K.: Toward Interactive Narrative. In *ICVS '03: Proc. of the 3rd Intl. Conf. on Virtual Storytelling* (2005), Johnson C., Moorhead R., Munzner T., Pfister H., Rheingans P., Yoo T. S., (Eds.), pp. 135–147.

[Shn96] SHNEIDERMAN B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *VL '96: Proc. of the 1996 IEEE Symposium on Visual Languages* (College Park, Maryland 20742, USA, 1996), pp. 336–343.

[TC06] THOMAS J. J., COOK K. A.: A Visual Analytics Agenda. *IEEE Computer Graphics and Applications 26*, 1 (2006), 10–13.

[TvSGSH02] TIEDE U., VON STERNBERG-GOSPOS N., STEINER P., HÖHNE K. H.: Virtual Endoscopy Using Cubic QuickTime-VR Panorama Views. In *MICCAI '02: Medical Image Computing and Computer-Assisted Intervention* (2002), pp. 186–192.

[Zel91] ZELTZER D.: *Task-level graphical simulation: abstraction, representation, and control.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1991, ch. 1, pp. 3–33.