DOI:10.1145/2950040

A full-body virtual autopsy of an ancient Egyptian mummy showed visitors he was likely murdered.

BY ANDERS YNNERMAN, THOMAS RYDELL, DANIEL ANTOINE, DAVID HUGHES, ANDERS PERSSON, AND PATRIC LJUNG

Interactive Visualization of 3D Scanned Mummies at Public Venues

BY COMBINING VISUALIZATION techniques with interactive multi-touch tables and intuitive user interfaces, visitors to museums and science centers can conduct self-guided tours of large volumetric image data. In an interactive learning experience, visitors become the explorers of otherwise invisible interiors of unique artifacts and subjects. Here, we take as our starting point the state of the art in scanning technologies, then discuss the latest research on high-quality interactive volume rendering and how it can be tailored to meet the specific demands

» key insights

- Volumetric visualization, together with CT scanning, creates opportunities for science communication, as visitors to public venues interactively explore the hidden interiors of rare museum artifacts and biological remains.
- Curating data and producing stories are bottlenecks for curators, interpreters, and exhibition designers; software solutions for supporting this workflow are lacking.
- Social context and engagement at public venues is available through touch tables with high resolution, image quality, and graphics performance together with user-interface design, balancing free exploration and guided storytelling.

of public venues. We then describe our approach to the creation of interactive stories and the design principles on which they are based and interaction with domain experts. The article is based on experience from several application domains but uses a 2012 public installation of an ancient mummy at the British Museum as its primary example. We also present the results of an evaluation of the installation showing the utility of the developed solutions.

Visitors walk into Gallery 64, the Early Egypt Gallery at the British Museum, eager to see and learn about one of the most famous and oldest

mummies in the collection. Known as the Gebelein Man, he was buried in a crouched position in a shallow grave during the late pre-dynastic period at the site of Gebelein in Upper Egypt. Unlike later Egyptian mummies, he was not artificially embalmed and his body was naturally mummified by the arid environment and direct contact with the hot dry sand more than 5,500 years ago. The mummy is safely protected behind glass and within a carefully controlled environment, beyond physical reach (see Figure 1). As visitors continue through the gallery a large touch-table display draws their

attention. Gathered around the table are several other visitors gazing down at the glass screen.

The same mummy is shown on the surface of the table, but, as a visitor moves a virtual slider on the table, the muscles, organs, and skeleton reveal themselves as the skin is gradually peeled away. Another visitor turns the mummy around to explore the other side. Someone touches the information icon near the left shoulder blade to discover that the cut on his back, as well as the damaged underlying shoulder blade and fourth rib, are the result of a single penetrating wound. The his-

Figure 1. The Gebelein Man (identified as human mummy EA 32751) is a pre-dynastic natural mummy and thus not prepared for artificial mummification, as later tradition directed. After being on display at the British Museum for more than 100 years, it was carefully transported to the Bupa Cromwell Hospital, London, U.K., for CT scanning.

Photo courtesy of the Trustees of the British Museum.

Figure 2. Visitors to the Mediterranean Museum in Stockholm can explore and virtually touch the coffins, wrapping, and body of the Neswaiu mummy, an Egyptian priest from Thebes, circa 300 B.C. This realistic tangible interaction is made possible through the integration of a range of technical developments, including high-resolution CT X-ray scans, laser scanning, photogrammetry, and modern GPUs.

Photo courtesy of Ove Kaneberg/Medelhavsmuseet.

tory of the Gebelein Man unfolds for the visitor in the same way researchers have used visualization to establish the cause of death; the Gebelein Man was murdered, with a metal blade the most likely weapon.²

The technology allowing visitors to the British Museum to explore the Gebelein Man is an image-generation technique called “volume rendering” that is efficiently executed on graphics processing units (GPUs). A stack of thousands of 2D images, as generated by modern computed tomography (CT) X-ray scanners, is processed in parallel to interactively create images. During the past decade, various museums have begun to scan human remains and artifacts from their collections. Not only are mummies scanned, so too are meteorites, bee flies, gecko lizards embedded in amber, and much more in 3D; see a 2013 report from the Smithsonian Institution.⁴

This article describes the underlying research and development of algorithms and technologies that have become the basis for a software solution for multi-touch tables. It also describes how the workflow of scanning, curating, and integrating the data into the overall creation of stories for the public can lead to engaging installations at museums around the world. Our work is an example of how computing technologies, especially in computer graphics and visualization, allow a general audience to interact with and explore the very same data that only scientists and domain experts have previously been able to study. Figure 2 shows children exploring the Neswaiu mummy at the Mediterranean Museum in Stockholm.

State-of-the-Art Scanning

The first commercial CT machinery was invented by Sir Godfrey Hounsfield (who was co-winner of the 1979 Nobel Prize for Physiology or Medicine) and made public in 1972. A tomography image of 80×80 pixels initially took 2.5 hours to compute, image-acquisition time not included. Fast-forwarding through four decades of development of detector materials, computer hardware, and image-reconstruction algorithms, a modern CT scanner can scan an entire body in fewer than five seconds, generating up to 10,000 slices, with down-to-0.3mm slice thickness.

The latest technology enables scanning of the body with multiple simultaneous energy levels. Multiple energy levels enable better tissue separation and, at the same time, give rise to new challenges related to visualization of this new type of data. Also, the radiation dose has been reduced approximately 70-fold over the past 10 years. A heart scan today results in less than one month worth of average background radiation.

The increased availability of CT scanners in modern hospitals and radiology clinics has provided an opportunity to scan the fragile mummies and learn about and reveal their interior content and structure without having to open or damage them. Not even the wrapping has to be touched. The first report on CT scanning of mummies was presented in 1979 by Harwood-Nash.⁷ Scanning protocols for mummies require custom settings, as the body is completely dehydrated and regular CT protocols assume naturally hydrated tissues, as discussed by Conlogue.⁵ Furthermore, standard protocols aim to minimize the radiation dose, and parameters (such as slice thickness) are set to the largest value that still yields sufficient image quality for clinical diagnosis. The scanning settings need a custom configuration to yield the best image quality. For mummies, the radiation dose is generally viewed as less of an issue, but some scholars have begun to raise concerns about DNA degradation.⁶

The presence of metallic objects (such as amulets) can pose a challenge to CT scanning, as these objects potentially block X-ray radiation from reaching the sensor array, and the signal reconstruction suffers from artifacts. Such artifacts can now be reduced with state-of-the-art reconstruction algorithms but may also be reduced with modified scanning protocols and multi-energy CT. As museums around the world are increasingly digitizing their collections, a shift toward dedicated CT scanners (such as industrial and micro-CT scanners) is a possible development.

Rendering What Hides Inside

The data generated from a CT scan after reconstruction is simply a stack of tomographic images, where the pix-

Historical Notes on Virtual Autopsy

What would eventually become the Inside Explorer visualization table began in 2004, when the Research Group in Scientific Visualization at the Department of Science and Technology, Campus Norrköping, Linköping University, Sweden, began a collaboration with the Center for Medical Imaging and Visualization at Linköping University on a newly started project on virtual autopsies. The need to scan and visualize an entire cadaver before the forensic autopsy posed technical challenges. The most critical were how to deal with the large amounts of data and how to interactively view it when the available memory in a typical workstation was not enough to hold the data of an entire body. This challenge drove our thinking and development of fundamental data-handling components and rendering techniques to support full-body virtual autopsy visualization.¹² The resulting software laid the groundwork for the initial versions of the Inside Explorer Visualization table. It essentially addressed the challenges by adjusting the local level of detail based on the perceived visual impact on the final image and adapting image quality at runtime during interactive exploration. Much research has been devoted over the past 20 years to developing and refining DVR, as well as managing increasingly large and complex datasets, as surveyed by Beyer et al.³ in 2015.

In the virtual autopsy project, we put each forensic case with casualties through a full-body CT scan prior to the traditional clinical autopsy of the cadaver. Since inception of the virtual autopsy project in 2003, approximately 800 cases have been processed at the Center for Medical Imaging and Visualization, and all cadavers now routinely undergo postmortem full-body CT imaging when there is suspicion of murder. A large body of knowledge has been established, and various studies have been conducted for specific types of autopsy examinations, including classification of dental fillings. This work has been beneficial, if not pivotal, for the subsequent scanning of mummified cadavers. The knowledge gained from scanning and visualizing cadavers could be translated and adapted into scanning and visualizing mummies and other interesting artifacts at the world's premier museums.

Figure 3. The volume-rendering integral. On a modern GPU, thousands of these ray integrals can be computed in parallel, enabling full HD resolution rendering at 30Hz–60Hz.

els within each image slice become samples, or voxels, in a 3D volume. The rendering of images of the volumetric dataset on a computer screen is known as volume rendering,¹³ or direct volume rendering (DVR), as the image is generated directly from the data without any intermediate (polygonal) representation. The most common approach to DVR is raycasting, where the stack of images is traversed by simulated rays that collect the color contribution, $g(s)$, from each sampled point as the rays traverse the stack (see

Figure 3). This approach represents an excellent case of a massively data parallel algorithm that fits modern GPU architectures, with thousands of parallel processors, that execute threads of ray casting programs. Each ray estimates how simulated light is reflected and attenuated as the ray propagates from the user's eye through the volume,^a and the final footprint of

^a This follows the principle of reverse ray tracing and allows for several optimizations of the computation.

Figure 4. The four different TFs designed for the Gebelein Man are shown together in this blend of four screenshots, where angled lines separate the different settings, from left to right: skin, soft tissue, soft tissue and skeleton, and skeleton. As the user slides the bottom bar, the TF changes gradually in the volume as well, unveiling the next layer.

Figure 5. Including complex shadows and lighting effects can provide superior cues for depth, structure, and texture of the data. Interactive performance requires simplifications and approximations in the rendering algorithms. This image shows a state-of-the-art interactive volume illumination and rendering of the Gebelein Man mummy based on the photon mapping algorithm.⁹

Dataset courtesy of the Trustees of the British Museum; images by Daniel Jönsson.

Figure 6. Illustration of the simplified user interaction with the touch table: rotations are restricted to the horizontal axis; clipping planes can be used from top and bottom of the screen; zooming is achieved with the familiar two-finger pinch/stretch; a simple slider selects presets of the TF; and embedded points of interest are displayed according to the current context.

Illustration by Martin Falk.

color and opacity, $\mathbf{I}$, is then deposited onto a pixel on the screen

$$\mathbf{I} = \int_a^b \mathbf{g}(s) e^{-\int_a^s \tau(u) du} ds$$

Computing such a ray for every pixel on the screen showing the image, and doing so at least 30 times per second, results in an interactive visualization of the interior of the scanned body or artifact. However, what is revealed depends on the user's settings (such as viewpoint and orientation) and, most important, how the X-ray absorption values are translated into colors and opacities. The function that describes this translation is referred to as the “transfer function” (TF), essentially included in $\mathbf{g}(s)$ for the emission/absorption of light contribution and $\tau(u)$ for the density-causing attenuation. The TF is one of the most central concepts in volume rendering. Here, it controls whether bones, muscles, fat, or various soft tissues will be visible or not, and to what degree (see Figure 4). As different tissues have varying degrees of overlapping ranges in the data domain, the design of TFs is also one of the major algorithmic challenges, as discussed in Ljung et al.¹¹

For a realistic high-quality final result, the ray-casting process must consider how light sources affect every sample and how samples globally affect each other in a participating medium. This multiple scattering effect is referred to as “volumetric illumination” and is technically embedded within $\mathbf{g}(s)$ and can thus become highly computationally demanding to compute due to the recursive multiple scattering events occurring in reality. The implementation must therefore exploit several simplifications and approximations to maintain an interactive experience. There are numerous approaches to volumetric illumination, ranging from effective methods based on approximations of light transport to more physically based methods with much higher computational demand; see Figure 5 for an example of a more complex lighting technique. The methods used in the work presented here were all discussed through a 2014 survey on volumetric illumination techniques for interactive DVR.¹⁰

Touching Data: Creating an Interactive Narrative

To bring volume rendering into the galleries of public venues we used several building blocks to design and implement a prototype of a visualization table for volumetric data termed “Inside Explorer” (see the sidebar “Historical Notes on Virtual Autopsy”). As starting points we used: state-of-the-art DVR software based on our research and development for high-quality visualization of large-scale volume data; large display screens with multi-touch interfaces; interaction design targeted at minimizing visual clutter and hiding of advanced concepts (such as the TF); and challenging large-scale volumetric data obtained from the scanning of humans and animals, as well as various artifacts.

It should be noted that interactive DVR has rarely been deployed at public venues to enable a general audience to explore scientific content. It is primarily a tool for research and limited professional use, requiring complex settings and domain-specific knowledge. The freedom and flexibility provided in professional medical workstations and experimental research software pose interesting challenges. For example, see Yu et al.¹⁵ and Jian et al.⁸ who both use multi-touch tables but primarily target expert users. These tools allow free exploration desired by expert domain users. However, for the general public of novice users, built-in guidance, simplified visual representations, and interaction are all necessary for efficiently supporting interactive narrative exploration. Here, we provide several insights gained with respect to the constraints and demands imposed on visualization at public venues. Our discussion focuses on four critical areas.

Performance. In a public installation, all aspects of performance must be considered in order to optimize the user experience. Such installations require fast loading times, as case-specific data is being loaded at runtime, and fast rendering times, with at least 30 frames per second at consistent high image quality. There are several reasons behind this requirement. In contrast with the medical- and scientific-visualization domains, where accuracy is more important than fast rendering and responsiveness, users at public

A design constraint is the limitation of interactions that cause rapid context switches and dramatic scene changes.

venues are eager to get started, not necessarily technically knowledgeable, and can potentially lose patience and interest. They are also not aware of the demanding nature of the datasets that are being displayed and often compare performance with the latest computer games. However, the key realization for software developers with respect to performance is the extreme sensitivity of human perception to large-scale touch interfaces and lagging visual updates. While the system is capable of creating the notion of touching a real artifact, that experience is lost as soon as users perceive lag in the interaction or rendering is perceived.

The Gebelein Man dataset has a size of $512 \times 512 \times 1,834$ voxels, 16-bit integers, and rendering image size of $1,920 \times 1,080$ (full HD). To ensure high frame rates with sufficient image quality, the renderer employs a set of optimizations: generation of a tight bounding surface for ray entry points and skipping empty blocks of data; termination of ray integration when opacity is saturated along the ray; and use of fast simplified shadowing using a “planesweep” technique that diffuses shadows from the light source and creates a shadow volume.¹⁴ The frame rates achieved range from 32Hz to 60Hz for the four different TFs used, as in Figure 4.

Robustness. From both a hardware and a software perspective a public environment poses particular robustness challenges in terms of number and type of user, hours of operation, and physical handling. The software must be almost as robust as high-availability systems and not provide users any opportunity to bypass the user interface and access lower layers of data or software. During hours of operation, the equipment must be able to endure handling by a wide range of visitors and sustain substantial wear and tear, as well as intentional damage caused by some of the visitors. Our experience is that multi-touch interfaces are indeed very robust, with all integrated technology hidden behind tempered glass and difficult to damage, even in environments like the British Museum, the leading visitor attraction in the U.K., with more than six million visitors per year.

Interaction. At public venues, the threshold time for “walk up and use”

Figure 7. By providing interactive narrative exploration and information at key points of interest, museum visitors are free to explore the narrative at their own pace and in their own sequence. The broken bones in the shoulder of the Gebelein Man were not the only injury visitors discovered. The posterior part of the fourth rib was also damaged. This discovery revealed a 5,500-year-old murder.

Photo courtesy of Interactive Institute Swedish ICT / Interspectral.

Figure 8. Workflow of Inside Explorer, from initial data acquisition and curation in the collection phase, to exploration, where scientific discoveries and annotations are added to the data. Production designers then provide their interpretations in the production stage, further enhancing the data with stories and visual design. User feedback and interaction patterns are then collected during exhibition.

Illustration by Martin Falk.

must be close to zero. In a matter of a few seconds the user must be able to grasp how to interact, how to rotate the objects without losing control, zoom in and out, and change rendering parameters to reveal the data of interest. During the course of the project's development, touch interfaces have rapidly become commonplace, and the potential user base spans all of modern society. This combination is beneficial since the same user interface must accommodate young

children, seniors, and people from varying cultural backgrounds and physical ability. In view of this universal public expectation, we chose multi-touch surfaces as our primary interaction methodology at an early stage of the project. Our ambition with the user interface design was to enable intuitive and explorative interaction while still allowing users to maintain a mental connection with the corresponding real artifact or subject (see

Figure 6). This has led us to apply design principles that can be characterized by the following parameters:

Object focus. The main focus of the representation on the screen should always be the rendered object, thus placing serious constraints on the choice of interaction paradigms and the implementation;

Judicious interaction constraints and micro-level interaction freedom. Interaction must be perceived as free but still constrained to noninvasively guide the user. We thus chose to constrain global interaction but increase freedom as points of interest are reached;

Minimalistic icons. Icons should be sparse and blend into the scene so as not to interfere with the object focus; they should also be associated with interaction points on the touch surface; and

Limiting multiuser interaction. When multiple eager and uninitiated users are interacting simultaneously, registration of unwanted interaction can occur frequently; a design constraint is thus the limitation of interactions that cause rapid context switches and dramatic scene changes. An example of the application of our principles is reflected in the TF interface design. In general, such design is a challenging task and, as noted earlier, is a major topic in the visualization research community. To hide the complexity of TF interaction, we use a range of predefined static TF settings from which the user can choose. By arranging these fixed TF settings along a slider it becomes intuitive and simple for users to peel away layer by layer of the content. The metaphor also maps easily onto polygonal geometries like coffins and chests.

A common technique used in volume rendering is known as "clipping planes" to make cutouts and reveal the interior of an object. From the user perspective, the planes are depicted as a pair of virtual scissors that enable virtual dissection of the subject/object. It is a simple yet powerful technique that is both informative and visually exciting when the user is able to discover content inside the skull or underneath the skin. To minimize usage complexity, we limit the interaction to a few fixed clipping directions with reduced flexibility. This limited interaction allows the user to explore while still making sure that focus and context are not lost due to excessive clipping and an

overly complex user interface. For ease of use, rotation is limited to a single axis, and when zooming and clipping the volume, that axis is always automatically centered on the screen.

Narrative. While adhering to the design principles described in the earlier section on interaction, the system must be able to support narratives, sometimes at multiple levels—the narrative of the case being shown, as well as the narrative of the scientific discovery and the user’s own experience of discovering the content.

At the heart of the interactive narrative exploration are the results of the work conducted after the capture and reconstruction of an object or body when curators analyzed and annotated the volumetric data. They can make

discoveries that need to be included, sometimes leading to the definition of a region or point of interest and a description of what findings have been made. The curators provide the fundamental scientific attribution to the data, documenting it such that exhibit producers can tailor the narrative and adapt the content to the public audience; see Figure 7 for an example. It is important for the software developer to realize that the narratives are implicitly embedded in the whole application, from visual settings to enabled interaction modes. For explicit information, we employ the notion of volumetric information hot spots that provide contextualized information about the scanned artifacts and subjects. During exploration, information spots

appear in the different layers, providing the user with contextual information, as well as details about particular findings. This information enables guided exploration through the datasets, as hints to other points of interest can be given.

Lessons Learned

One key to the successful installation of Gebelein Man is careful iterative interaction between the museums and the development team of programmers, modelers, and interaction designers. For example, in the design process for the Gebelein Man, we closely collaborated with the British Museum curatorial teams and the Interpretation and Learning Departments, as well as access managers; see Figure 8 for

Figure 9. The digitization of Neswaiu included multiple CT scans of the body as well as photogrammetry and laser scanning of the wooden sarcophagi in which he was buried. On the visualization table, both the textured polygonal models and the volume-rendered remains of the mummy are shown as different layers that can be peeled away, one at a time.

Left photo courtesy of Ove Kaneberg/Medelhavsmuseet;
right image courtesy of Interactive Institute Swedish ICT/Interspectral.

Figure 10. Reproducing hidden artifacts is another benefit of the process of digitizing museum collections. Here, a falcon amulet from the Neswaiu mummy in Stockholm is shown. It was produced through inexpensive 3D printing of a plastic model used to create a mold in which a gold-plated bronze copy was cast.

Image and photo courtesy of Interactive Institute Swedish ICT/Interspectral.

a schematic of the adopted workflow. The iterative approach was based on frequent videoconferencing, as well as regular on-site meetings and design reviews, leading to refinement of the software, as well as the physical exhibit design, to improve aspects of the interactive public experience (such as accessibility to visitors in wheelchairs).

The final Gebelein Man exhibit in the Early Egypt Gallery at the British Museum was subject to an evaluation by the museum's Interpretation Department prior to, as well as after, installation of the Inside Explorer visualization table to determine its contribution to the visitor experience.¹ In the pre-table evaluation, we tracked 114 visitors and interviewed more than 50 of them. In the post-table evaluation, we tracked 133 visitors and interviewed 95 of them. Here are some of the main findings:

Time spent in the gallery. Prior to table installation, the median time spent in the gallery was two minutes. After table installation, visitor dwell time in the gallery increased by 40%, and visitors made an average of three stops (instead of two) in the gallery. An important finding was that the fraction of visitors stopping to view other displays in the gallery increased from 3% to 20%;

Visitor attention to original display. We were initially concerned that the virtual visual exploration table would divert interest from the physical artifacts and mummy in the gallery, but the opposite occurred, where the fraction of visitors viewing the physical display increased from 59% to 83%. The total visitor attention time given to Gebelein Man was shared between the table and the mummy, leading to a decrease in time spent observing the physical subject alone; and

Interaction with the table. The table saw heavy visitor interaction and was in use 95% of the evaluation time. Of the visitors to the gallery, 83% stopped at the Gebelein Man display, 59% stopped at the Inside Explorer table, and 36% engaged with the table. The average interaction time was two minutes, nine seconds. In our interviews, the display was scored nine out of 10 on average for ease of use and enjoyment and an average of 8.5 for informativeness. Approximately 60% of table users discovered the cause of death.

The amulet of Neswaiu, still present within the mummy wrappings and behind glass, can now be physically reproduced in high-resolution detail through a process of digital virtual extraction.

Overall, the visualization table had a positive effect on the display of Gebelein Man and the gallery as a whole. It increased the attention-grabbing power of the display by 24% over surveys before table installation. Our visitor interviews revealed that the Inside Explorer visualization table had contributed to visitors' understanding of the mummy, with many returning to observe it after having discovered information on his age, preservation, and cause of death. The Inside Explorer table also appears to have had a positive effect on the other physical displays and cases in the gallery. Based on these conclusions, the British Museum decided to turn the temporary installation into a permanent exhibit.

Even though these lessons learned are described primarily in the context of our work with the British Museum, we see similar results in many other installations worldwide and regard our work as only the beginning of the exciting possibilities promised by interactive visualization at public venues.

The Future of Science Visualizations at Public Venues

The Inside Explorer visualization table described here is an example of an ongoing trend that is fundamentally changing the visitor experience in museums and science centers.⁴ Instead of passively absorbing information and studying static objects or subjects on display, visitors are allowed to engage with digital objects and explore the secrets hidden within an object in a manner similar to the way researchers would process and explore the real object and its digital representation. We curated and annotated the datasets and provided additional information to furnish a complete story for the user experience with the table.

Based on the work presented here, we now intend to take another step, aiming to develop a solution that addresses the bottlenecks in the process to enable rapid and efficient conversion of a scanned subject into a compelling visitor experience. The first bottleneck is in the deeper initial exploration of scans of the subject to facilitate discovery. Deep technical knowledge of the process should not be required, as it is today, making it possible to address

a wider group of domain experts, as well as non-experts. The second main production bottleneck is the process of closely linking any scientific or curatorial findings with the process of interpretation and exhibit production. The envisioned solution will thus support the whole chain, from scientific data collection and documentation of scientific discoveries, design, and interpretation of the information into non-linear stories, to production of end-user applications for exhibits. Figure 8 outlines how we refined, curated, and annotated the data, though the original data is the basis of every step in the process of developing the visitor experience in the gallery.

We have also taken initial steps to explore other opportunities offered by the digitization of museum artifacts, including laser scanning and photogrammetry, employed for the mummy of Neswaiu, as in Figure 9, and also 3D printing technologies of the scanned data. For example, the amulet of Neswaiu, still present within the mummy wrappings and behind glass, can now be physically reproduced in high-resolution detail through a process of digital virtual extraction (see Figure 10). Visitors can clasp and feel it and explore it physically, even though the mummy has never been unwrapped.

What does this development mean for the future of science visualization at public venues? One important aspect of sharing visualizations with the public may be authenticity of the science and resulting data. The public can interact with and explore the real underlying data, taking the visitor experience to a new level of engagement. The scientific story can be embedded in the interactive narrative on the table using points of scientific interest and told through the framing of the table installation with additional information about the process, from scan to installation.

Conclusion

Our findings demonstrate that science visualization has the potential to narrow the gap between the general public and research, as it allows scientists and curators to share the methods used to interpret and analyze the collections with visitors (such as determination of the age and sex of a mummy). Bringing

the original research data to the public and providing tools enabling learning and exploration is an exciting and challenging scientific adventure that creates many new research opportunities involving not only computer science as enabling technology but also learning and communication, social science studies, and domain-specific areas from which the objects of study originate.

Acknowledgments

The authors wish to thank the Trustees of the British Museum. The Gebelein Man dataset is courtesy of the Trustees of the British Museum. We also wish to acknowledge the help and input from the Department of Ancient Egypt and Sudan at the British Museum. We also thank Medelhavsmuseet/the Museum of Mediterranean and Near Eastern Antiquities in Stockholm, Sweden. We also thank the team from Interactive Institute Swedish ICT/Interspectral, specifically David Karlsson, Claes Ericson, Karl Lindberg, and Kristofer Jansson. Finally, we thank the Norrköping Visualization Center C and the Center for Medical Image Science and Visualization, Linköping.

This work is supported by Swedish eScience Research Center, Excellence Center at Linköping-Lund in Information Technology, Knut and Alice Wallenberg Foundation, Foundation for Strategic Research, The Knowledge Foundation, Swedish Research Council, and Vinnova.

References

1. Antoine, D. *Evaluation of the Virtual Autopsy Table (Room 64: Early Egypt Gallery)*. Internal Report, Supplemental Material, Interpretation Department, The British Museum, London, U.K., Feb. 2016.
2. Antoine, D. and Ambers, J. The scientific analysis of human remains from the British Museum collection: Research potential and examples from the Nile Valley. In *Regarding the Dead: Human Remains in the British Museum*, A. Fletcher, D. Antoine, and J. Hill, Eds. The British Museum, London, U.K., 2014, 20–30; https://www.britishmuseum.org/PDF/Regarding-the-Dead_02102015.pdf
3. Beyer, J., Hadwiger, M., and Pfister, H. State of the art in GPU-based large-scale volume visualization. *Computer Graphics Forum* 34, 8 (Dec. 2015), 13–37.
4. Clough, G.W. *Best of Both Worlds—Museums, Libraries, and Archives in a Digital Age*. e-book. Smithsonian Institution, Washington, D.C., 2013; <http://www.si.edu/bestofbothworlds>
5. Conlogue, G. Considered limitations and possible applications of computed tomography in mummy research. *The Anatomical Record* 298, 6 (June 2015), 1088–98.
6. Grieshaber, G.M., Osborne, D.L., Doubleday, A.F., and Kaestle, F.A. A pilot study into the effects of X-ray and computed tomography exposure on the amplification of DNA from bone. *Journal of Archaeological Science* 35, 3 (Mar. 2008), 681–687.
7. Harwood-Nash, D. Computed tomography of ancient

Egyptian mummies. *Journal of Computer Assisted Tomography* 3, 6 (Dec. 1979), 768–773.

8. Jian, W., Tu, H.-W., Hua Han, X., Tateyama, T., and Chen, Y.-W. A preliminary study on multi-touch-based medical image analysis and visualization system. In *Proceedings of the Sixth International Conference on Biomedical Engineering and Informatics (BMEI 2013)* (Hangzhou, China, Dec. 16–18). IEEE, 2013, 797–801.
9. Jönsson, D., Kronander, J., Ropinski, T., and Ynnerman, A. Historygrams: Enabling interactive global illumination in direct volume rendering using photon mapping. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2364–2371.
10. Jönsson, D., Sundén, E., Ynnerman, A., and Ropinski, T. A survey of volumetric illumination techniques for interactive volume rendering. *Computer Graphics Forum* 33, 1 (Feb. 2014), 27–51.
11. Ljung, P., Krüger, J., Gröller, E., Hadwiger, M., Hansen, C.D., and Ynnerman, A. State of the art in transfer functions for direct volume rendering. *Computer Graphics Forum* 35, 3 (June 2016).
12. Ljung, P., Winskog, C., Persson, A., Lundström, C., and Ynnerman, A. Full-body virtual autopsies using a state-of-the-art volume-rendering pipeline. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (Sept.–Oct. 2006), 869–876.
13. Max, N. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 1, 2 (June 1995), 99–108.
14. Sundén, E., Ynnerman, A., and Ropinski, T. Image plane sweep volume illumination. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2125–2134.
15. Yu, L., Svetachov, P., Isenberg, P., Everts, M.H., and Isenberg, T. FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 1613–1622.

Anders Ynnerman (anders.ynnerman@liu.se) is a professor focusing on scientific visualization and head of the Media and Information Technology Division of Linköping University, Linköping, Sweden, and director of the Norrköping Visualization Center C, Norrköping, Sweden.

Thomas Rydell (thomas.rydell@interspectral.com) is the CEO and co-founder of Interspectral AB, Norrköping, Sweden, and formerly studio director of the Interactive Institute Swedish ICT, Norrköping, Sweden.

Daniel Antoine (dantoine@thebritishmuseum.ac.uk) is the curator of physical anthropology at The British Museum and an honorary senior research fellow in the Institute of Archaeology, University College London, U.K.

David Hughes (david.hughes@interspectral.com) is a technical advisor and strategic consultant at Interspectral AB, Norrköping, Sweden, and previously held the same position at the Interactive Institute Swedish ICT, Norrköping, Sweden.

Anders Persson (anders.s.persson@liu.se) is a professor of radiology at Linköping University, Linköping, Sweden, and the director of the Center for Medical Image Science and Visualization, Linköping, Sweden.

Patric Ljung (patric.ljung@liu.se) is a senior lecturer in immersive visualization at Linköping University, Linköping, Sweden, and research coordinator at Norrköping Visualization Center C, Norrköping, Sweden.

Copyright held by the authors.

Watch the authors discuss their work in this exclusive *Communications* video. <http://cacm.acm.org/videos/interactive-visualization-of-3d-scanned-mummies-at-public-venues>