

Data Reconstruction and Visualization Techniques for Forensic Pathology

Alexander Ehlert, Zein Salah, Dirk Bartz

Visual Computing for Medicine Group[†]
University of Tübingen, Germany

Abstract

Forensic pathology is largely concerned with the determination of the cause and manner of deaths after accidents, or other circumstances in criminal investigations. A major task in that process is the documentation of surface injuries, which is traditionally done by drawing sketches, photography, or more recently by photogrammetry to generate a three-dimensional digital lesion cartography of the body surface.

In this paper, we describe a semi-automatic processing pipeline how data from 3D photogrammetry is combined and used to generate a visual surface representation of accident victims. In that course, a number of steps are performed to provide a high-quality interactive, point-based visualization of the acquired data, which can be used in a more routine way than previous forensic surface methods.

Keywords: *Scanned Data, Color and Intensity Matching, Geometry Matching, Point-based Rendering*

Categories and Subject Descriptors (according to ACM CCS): I.4.1 [Digitization and Image Capture]: Scanning, Imaging geometry; I.4.8 [Scene Analysis]: Color; I.3.7 [Three-Dimensional Graphics and Realism]: Color, shading, shadowing, and texture; I.3.8 [Application]: Virtual Medicine; J.3 [Life and Medical Sciences]: Forensic Pathology

1. Introduction

Forensic pathology is largely concerned with the determination of the cause and manner of deaths after accidents, or other circumstances in criminal investigations. A major task in this process is the documentation of surface injuries, which is traditionally done by drawing sketches, photography, or more recently by photogrammetry to generate a three-dimensional digital lesion cartography of the body surface. At the Institute of Forensic Medicine at the University of Tübingen, a quite labor-intensive process has been established, which employs off-the-shelf software for image processing and geometry extraction, and rendering [SWS00].

However, this process resulted in relatively poor quality and, due to the geometric complexity, low rendering performance. Even more important, the whole procedure took several days for processing the scanned data. Therefore, it is not a suitable approach for routine application.

In this paper, we describe an application how data from 3D photogrammetry is processed and integrated to generate a forensic surface representation of accident victims. In that course, a number of processing steps are performed to provide a high-quality interactive, point-based visualization of the acquired data, which are shown in Figure 1. Specifically, the scanned data points must be edited to remove parts that do not belong to the body (eg., the post-mortem table), and smoothed afterwards. After an initial color correction, the lighting and coloring information are matched between the data points from various scans and the colored geometry is registered and fused into one model.

In the remainder of the paper, we will briefly discuss some related work in Section 2 and present details of our approach in Section 3. Thereafter, we will present our results (Section 4) and conclude in Section 5.

[†] {ehlert,salah,bartz}@gris.uni-tuebingen.de

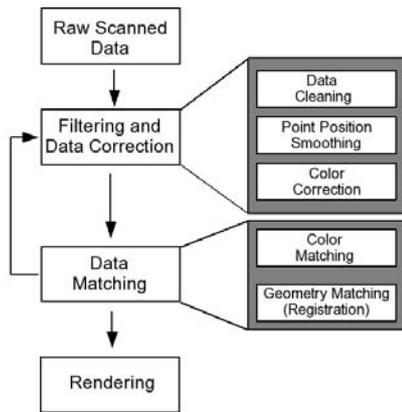


Figure 1: Processing pipeline: From scanned raw data to a point-based visualization.

2. Related Work

Reconstruction and rendering of scanned objects is currently one of the most investigated research topics in vision and computer graphics. The probably most prominent in recent years are the IBM Pietà project [IBM05, RBMT98, BRM*02] and Stanford's Digital Michelangelo project [LPR*00]. In the process from model acquisition to 3D rendering [BR02], many issues need to be addressed. Since most objects are too large to fit into the working space of the 3D scanners, several scans are stitched together using 3D registration algorithms [BM92]. The acquired textures must be adapted, since lighting conditions and camera equipment can be different, resulting in different color casts and brightness [BMR01, RB99]. After the geometric and texture data are corrected and matched, a polygonal surface can be constructed [BMR*99, RBMT98] or a point-based rendering avenue can be taken [RL00, PvBZM00].

In contrast to the scanning of objects, Xu et al. describe the acquisition and rendering of scanned out-door scenes [XC04]. They also generate a point-based representation from the scanned data points. While the IBM Pietà, the Digital Michelangelo Project, and our forensic science application aim at a realistic representation of the scanned data, Xu et al. provide a stylized rendering approach.

There are also different approaches for the correction of lighting and coloring. Paris et al. suggested a simple model that aims at a compromise between highly realistic skin textures and real-time rendering [PSQ03]. They achieve their goal by computing a skin reflectance map that encodes the lighting-independent parts into a texture. Combined with the standard Phong-based lighting model, this generates a quite realistic looking skin surface. However, this approach significantly modifies the original data and is thus less suited for forensic documentation.

Reinhard et al. suggested a color transfer approach

[RAGS01] that is quite similar to the one that we developed. In contrast to our simple RGB color space principal component analysis (PCA), they perform a color space conversion into the XYZ color space, before they perform the PCA.

Other methods of scanning have also been used in the context of forensic pathology. Thali et al. used post-mortem data from a multi-slice Computed Tomography (CT) scanner and Magnet Resonance Imaging (MRI) to examine the cause of death [TYS*02]. Preim et al. additionally provided segmentation and analysis tools to estimate the volume of certain body organs to aid at the determination of the sudden infant death syndrome [PCH*05]. A combination of local photogrammetry with radiological data to also capture the texture and color of lesions was also presented by Thali et al. [TBW*03].

Also related is the visual human project [Nat05], which provides scans from CT, MRI, and digitized anatomical cryosections of a deceased female and male human being.

3. Processing Data From Multiple Scans

In this section, we describe the different steps from data acquisition to the rendering of the scanned data. These steps include the cleaning of the scanned data from scanning artifacts and background objects, position smoothing, brightness and color correction, and finally the combination of the different scans into a single point set of the model (Fig. 1). Since the number of data points is quite large, we opted for a point-based representation that does not rely on a polygonal mesh with all its full neighborhood information.

3.1. Data Acquisition and Data Structure

Data acquisition is based on a scanner that consists of a line projector with a CCD camera. The 3D point geometry of the subject is acquired based on a coded light approach [Wah84, Wah86]. The color information is captured by the CCD camera and stored into an image texture.

The scanning process itself takes place at the Institute of Forensic Medicine and starts with the positioning of the objects on a post-mortem table. A vacuum mattress is used to stabilize the objects on the table. This, however, also complicates scanning, since the lateral parts of the bodies are less visible to the scanner. Since the limited scanning area of the scanner (up to $1.0m \times 1.0m$) does not allow the acquisition of the full body geometry within one scan, up to 28 (sometimes 40) scans are conducted from the front and the behind, where each scan acquired about 40,000 data points. In order to provide an optimal representation of features (e.g., lesions), the respective scan must ensure that the individual feature is well covered and in the center of the scanned region. For registration purposes, markers may be attached to the scanned object. Note that the scans of the front and back are combined into two different point-set models.

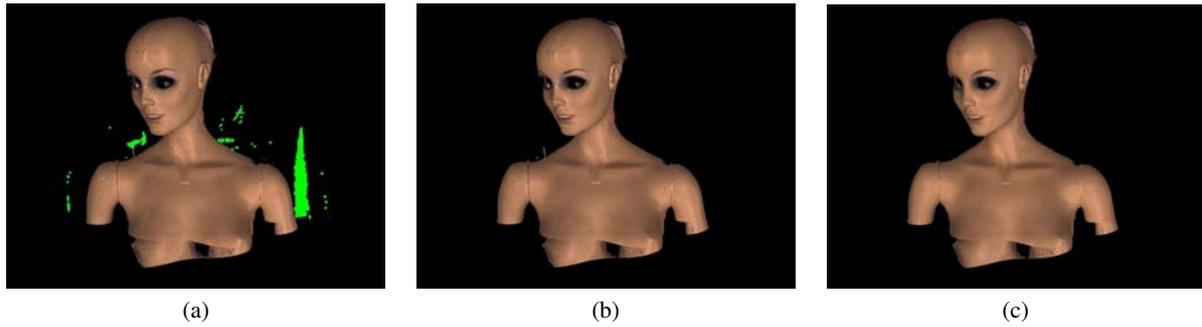


Figure 2: Filtering and data correction of the Mannequin dataset: (a) Data points which do not belong to the scanned subject (green) are selected and removed. (b) The positions of the remaining data points are smoothed. (c) Color values of the data points are corrected based on a median filter. A close-up can be seen in Figure 4.

The point sets acquired by the scanner can be quite large. As mentioned above, up to 28 datasets of approximately 11 million points – with approximately 8 million points remaining after various geometric cleaning operations – need to be processed. Since the large number of points would introduce considerable processing costs in terms of required memory (for connectivity and neighborhood information) if we reconstruct a polygonal isosurface, we use a point-based representation for processing and rendering. Several filtering operations, however, do require neighborhood information. In those cases, we employ a lazy neighborhood estimation scheme, based on a k-nearest neighbor search in an octree.

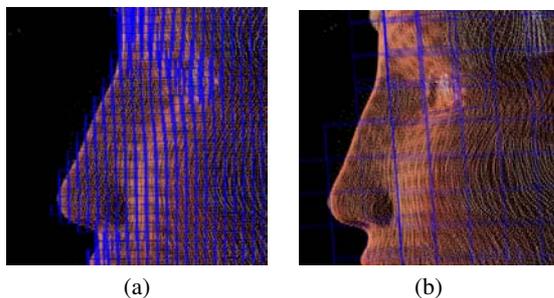


Figure 3: Octree leaf block resolutions in left cheek area: (a) Blocks with 5mm resolution, (b) blocks with 20mm resolution.

After data acquisition, the data points and all their attributes (position, normal, color) are temporarily stored in array A. Successively, the indices of all these data points are sorted into the octree, whereas an axis-aligned bounding box of these data points represents the respective octree leaf block (Fig. 3). If the number of data points in an octree block exceeds a certain threshold T , this octree block is split into its child blocks. Empty blocks, however, are skipped.

Once the octree and its data point indices are created, the

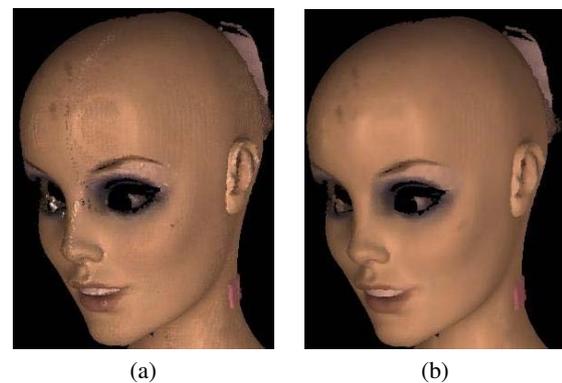


Figure 4: Filtering and data correction of the head of the Mannequin dataset: (a) The positions of the data points are smoothed. (b) Color values of the data points are corrected based on a median filter.

leaf blocks of the octree are sequentially traversed and all indexed data points are copied into a new array B. At the same time, the respective indices are updated to array B. Array A is afterwards deleted, since it is no longer needed. After this operation, the data points of a local neighborhood can be accessed in a more cache sensitive fashion, which in turn speeds up the following filter operations.

3.2. Filtering and Data Correction

The first processing step consists of several cleaning operations, since systematic and statistical measurement errors generate invalid data points. Furthermore, structures like the post-mortem table are part of the scanned data, but need to be removed before further processing (see Fig. 2a). Therefore, several interactive and automatic editing operations are implemented.

All editing operations are performed on the level of the

octree leaf blocks, hence data points cannot be modified individually. However, a typical leaf block resolution of 5mm in each direction usually provides sufficient granularity for editing purposes. Figure 3 shows two different spatial leaf block resolutions that contain a different number of data points.

Interactive editing can remove blocks that are either specifically selected (picking), selected by a drawn frame (frame selection), or by successively selecting all leaf blocks in the ϵ -proximity of already selected leaf blocks (region growing). Automatic editing provides tools which remove leaf blocks in cuboid area of the data space (cuboid removal). Finally, the minimum filter focuses on outlier data points by removing all leaf blocks that contain less than a specified number of data points.

After editing, the noise on the geometric surfaces is smoothed by fitting a tangential plane P to the neighbors of every remaining data point. The respective data point is then projected onto this plane [Lev03]. Figure 2b shows the result of this geometric smoothing process, whereas an enlarged area of the head is shown in Figure 4a.

In the final step of the initial data processing, a simple color cleaning is performed by applying a median filter to the color values. Data points that have a color value outside the range specified by the median filter are removed. Since only a relatively small number of data points is removed, the overall sharpness of the acquired data is maintained (see Fig. 2c and the enlarged head area in Fig. 4b). This last step, however, has a clear trade-off between maintaining the full color fidelity of the data (Fig. 4a) and the corrections of potential scanning errors. In particular small lesions might experience changes of their color appearance. This can be an issue for the documentation purpose if not all lesions were detected during the postmortem examination.

3.3. Data Matching

The second major processing stage is concerned with data matching. Specifically, the colors of the remaining data points need to be corrected to adapt for different lighting conditions at data acquisition. Afterwards, the data points of the different scans (up to 28 per body orientation) are registered and combined into two dataset, one for each body orientation.

In an ideal scanning setup, a sufficient number of diffuse light sources would ensure uniform lighting conditions for every scan. Furthermore, a calibration of the white points of different CCD cameras would reduce differences in coloring. However, the scanned data of our project did not benefit from such an ideal setting, even though the postmortem room allows a much more controlled scanning environment. Instead, the data from different scans vary substantially in lighting and in color cast, eg., one scan showed a rather yellowish coloring (Fig. 5b).

For varying brightness in the different scans, we simply apply a histogram stretching based on the different white and black points of the cameras in order to map the brightness ranges into a common range. Furthermore, a gamma correction is applied for further improvements.

For the correction of the different color casts (see Fig. 5), we apply a principal component analysis (PCA). With this PCA, we achieve an almost orthogonal color space, where each component provides only a minimized correlation to the other components. Concurrently to our research (which started in early 1999), Reinhard et al. proposed a more complex method that also employs a PCA to generate a near-orthogonal color space [RAGS01]. Their approach, however, aims at a transformation from the RGB color space into a perception-driven $l\alpha\beta$ color space [RCC98].

The image sequence in Figure 5 shows how the colors in an image with a yellowish color cast (Fig. 5b) are corrected using the described method. Figure 5a shows the image with the correct (reference) color cast and Figure 5c shows the re-colored image of Figure 5b.

In the final color correction step, we correct lighting artifacts which are due to the scanning setup (see Fig. 6). Among these are shadows and surface highlights. We identify the respective points by examining all pixel values for very low (shadows) and very high pixel intensities, similar to Rushmeier et al. [RB99]. Using the surface normal of the specific data points – estimated based on its point neighborhood – and the estimated light position, we approximate the diffuse lighting for that point, remove its influence, and use instead an interpolated value from its (possibly already corrected) neighborhood pixels.

After color correction, we need to register the geometry of the different scans in order to achieve a combined model (see Fig. 8). Essentially, we employ a variation of the iterative closest point algorithm [BM92, RL01], where a random set of data points \vec{s}_i is selected from the reference dataset and a match point set \vec{t}_i is searched in the target dataset. \vec{t}_i are the data points in the target dataset which are located closest to the reference data point (according to the initial positioning). This matching process can be specified by solving the linear equation system $M \cdot \vec{s}_i = \vec{t}_i$. For a rigid registration, M has six degrees of freedom – three parameters for rotation and three parameters for translation – which are optimized to find an appropriate registration. We use the popular Gauss-Newton algorithm to minimize the distance S between reference points and target surface $P(\vec{t}_i)$, with $P(\vec{t}_i)$ as an approximation of a tangential plane in an ϵ -neighborhood around \vec{t}_i :

$$S = \sum_i \text{dist}(M \cdot \vec{s}_i, P(\vec{t}_i))^2 \quad (1)$$

To ensure meaningful registration results, reference and target datasets need to be roughly pre-positioned. The actual ICP algorithm will then compute the specific transformation.

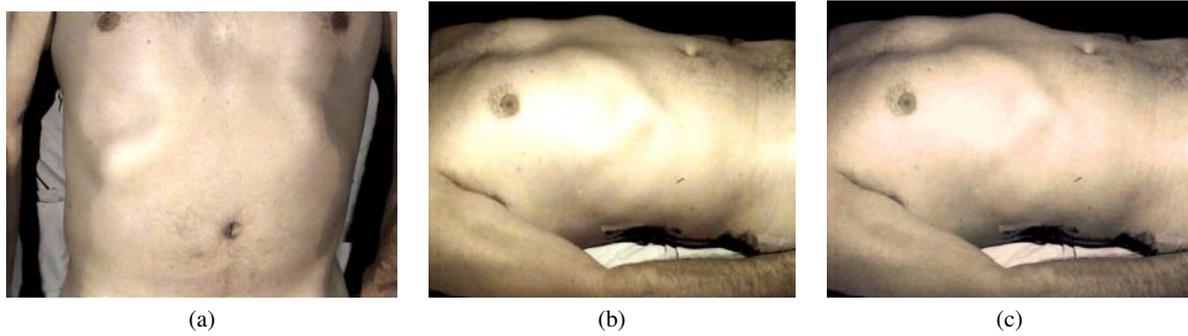


Figure 5: Color matching of the Accident Victim dataset: (a) Image with reference colors. (b) Image with yellow color cast that needs to be corrected. (c) Image with new matched colors.

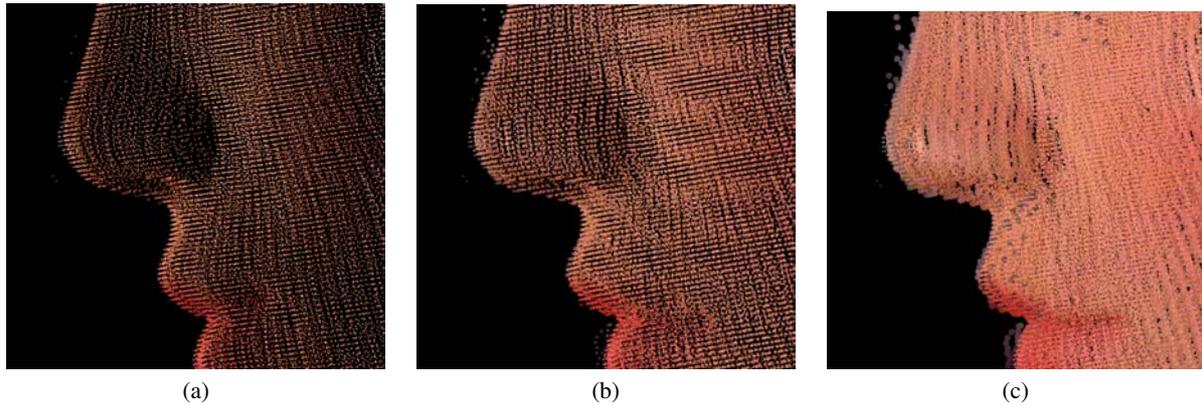


Figure 7: Nose region of Mannequin with different QSplat sizes: (a) Small QSplats, (b) medium QSplats, and (c) large QSplats.

3.4. Rendering

After cleaning and matching of the data scans, the combined dataset is rendered. Since we opted for a point representation that only maintains rough neighborhood information for every data point, we used a point-based rendering method to visualize the combined dataset.

Each cleaned point model (two for each body) contains approximately eight millions points, sorted in the octree. During the rendering stage, the octree is traversed and the leaf blocks are depth-sorted. For the actual OpenGL rendering, we follow the QSplat approach [RL00] in using anti-aliased OpenGL points as rendering primitive (Fig. 7). Hence, all data points associated with the respective leaf blocks are passed as colored GL_POINTS using the standard OpenGL lighting and shading model, weighted with the matched colors. Note that the data points within a leaf blocks are not depth-sorted. However, the resulting visual defects are negligible.

Anti-aliased rendering is also improved by attenuating the brightness of the OpenGL points (posing as QSplats) by their size, transparency, and distance from the view point,

using the point parameter extension of the OpenGL ARB (ARB_point_parameters). If that extension is not available, the size of the OpenGL points is manipulated accordingly.

4. Results

We demonstrate results of the developed techniques on two different datasets, where each scanning subject was in dorsal and in a face-down position. The first dataset (Figs. 2, 4, and 8) is from a scan of a Mannequin, which consists of 28 partial scans for each body positioning. Figure 8 shows the results from the face-down (left) and dorsal position (right). In the middle, four partial scans from the dorsal position are combined into one model.

The second dataset is from the scan of an Accident Victim, which is also composed of 28 scans per body position. Figure 5ab shows two photographs of the victim laying on the post-mortem table, and the color correction of the second photograph (Figs. 5c). Figure 6 shows the light correction of the reconstructed model from Figure 5a.

Most manual work of the whole processing pipeline after the scanning itself is spent to a smaller extent by the clean-

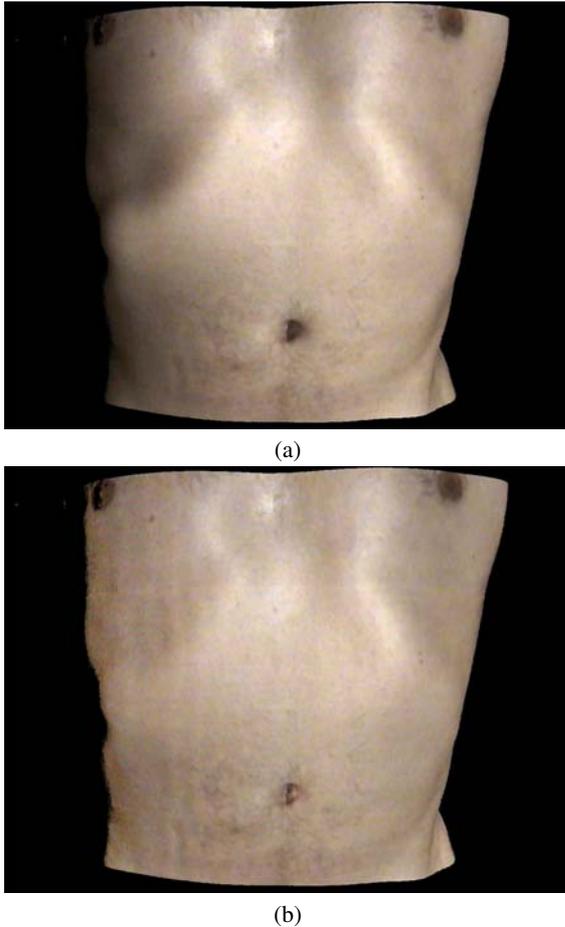


Figure 6: Lighting correction with an image of the Accident Victim dataset: (a) Uncorrected image, (b) corrected image.

ing the dataset (step one of the filtering and data correction). This in particular involves the removal of the also scanned parts of the postmortem table, and of the vacuum mattress, although most geometry is removed by the minimum editing filter. To the larger extent, manual work is necessary for the pre-positioning of the individual scans to register them into the combined model, if the attached markers cannot be used. The required precision of the pre-positioning depends on the overlap between the different scans and if distinctive features are covered by the scans; largely flat areas require a significant more precise pre-positioning than curved areas. As mentioned in Section 3.1, lateral and possible incomplete scans require extra pre-positioning effort.

On a standard PC equipped 1GB of main memory, a 2.6GHz Intel P4 CPU, and an ATI Radeon 9700 graphics subsystem, rendering achieves about 10-100 fps, depending on the size of the dataset (number of data points) and the selected resolution. Data pre-processing took several min-

utes (usually between 10-40mins, depending on manual tuning requirements), whereas the geometry matching stage required manual pre-positioning of the partial geometries.

5. Conclusion and Future Work

In this paper, we presented an application of various data reconstruction and visualization techniques to generate a colored surface model of dead human beings for forensic pathology. The data was acquired through an optical scanner with a CCD camera, line projector and employing the coded-light approach. It was cleaned and thereafter matched in brightness and color. Finally, its geometry was matched into a combined model. The generated models were then rendered at interactive framerates using a point-based rendering approach.

The goal of this application is to provide additional means for the documentation of surface injuries inflicted by fall, shock, or by a blow. Note that it is meant as a supplement to the traditional documentation by sketches and photography, not for a full body data acquisition like with the visible human datasets. Beyond the mere documentation, the acquired data is also intended for the analysis of the mechanism of accidents, and hence for improving the safety of vehicles.

Acknowledgments

This work has been supported by project “Digital Wounded”, funded by the State of Baden-Württemberg, the Institute of Forensic Pathology and Department for Neurosurgery of the University Hospital Tübingen, and by DFG project VIRTUE. The datasets are courtesy of the Institute of Forensic Medicine of the University of Tübingen.

References

- [BM92] BESL P., MCKAY N.: A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 2 (1992), 239–256. 2, 4
- [BMR*99] BERNARDINI F., MITTLEMAN J., RUSHMEIER H., SILVA C., TAUBIN G.: The Ball-Pivoting Algorithm for Surface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 5, 4 (1999), 349–359. 2
- [BMR01] BERNARDINI F., MARTIN I., RUSHMEIER H.: High-Quality Texture Reconstruction from Multiple Scans. *IEEE Transactions on Visualization and Computer Graphics* 7, 4 (2001), 318–332. 2
- [BR02] BERNARDINI F., RUSHMEIER H.: The 3D Model Acquisition Pipeline. *Computer Graphics Forum* 21, 2 (2002), 149–172. 2
- [BRM*02] BERNARDINI F., RUSHMEIER H., MARTIN I., MITTLEMAN J., TAUBIN G.: Building a Digital Model of Michelangelo’s Florentine Pietà. *IEEE Computer Graphics and Applications* 22, 1 (2002), 59–67. 2

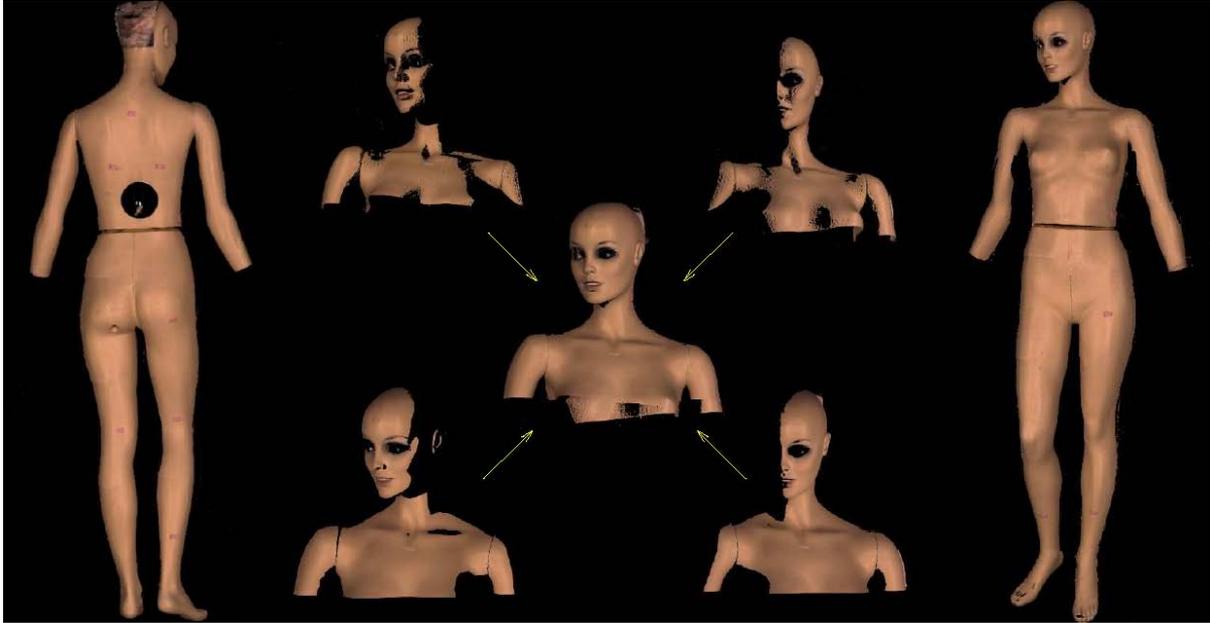


Figure 8: Registration of partial scans: In this example, four partial scans are registered into a combined model.

- [IBM05] IBM RESEARCH: Pietà Project. <http://www.research.ibm.com/pieta/index.html>, re-accessed 2005. 2
- [Lev03] LEVIN D.: Mesh-independent Surface Interpolation. In *Geometric Modeling for Scientific Visualization* (2003), Brunnert G., Hamann B., Müller H., Linsen L., (Eds.), Springer-Verlag, Heidelberg, pp. 37–49. 4
- [LPR*00] LEVOY M., PULLI K., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., CURLESS B., SHADE J., FULK D.: The Digital Michelangelo Project: 3D Scanning of Large Statues. In *Proc. of ACM SIGGRAPH* (2000), pp. 131–144. 2
- [Nat05] NATIONAL LIBRARY OF MEDICINE: The Visual Human Project Project. http://www.nlm.nih.gov/research/visible/visible_human.html, re-accessed 2005. 2
- [PCH*05] PREIM B., CORDES J., HEINRICHS T., JACHAU K., KRAUSE D.: Quantitative Bildanalyse und Visualisierung für die Analyse von post-mortem Datensätzen. In *Proc. of Bildverarbeitung für die Medizin* (2005), Springer, pp. 6–10. 2
- [PSQ03] PARIS S., SILLION F., QUAN L.: Lightweight Face Relighting. In *Proc. of Pacific Graphics* (2003), pp. 41–50. 2
- [PvBZM00] PFISTER H., VAN BAAR J., ZWICKER M., M.GROSS: Surfels: Surface Elements as Rendering Primitives. In *Proc. of ACM SIGGRAPH* (2000), pp. 335–342. 2
- [RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color Transfer Between Images. *IEEE Computer Graphics and Applications* 21, 5 (2001), 34–41. 2, 4
- [RB99] RUSHMEIER H., BERNARDINI F.: Computing Consistent Normals and Colors from Photometric Data. In *Proc. of International Conference on 3D Digital Imaging and Modeling* (1999), pp. 99–108. 2, 4
- [RBMT98] RUSHMEIER H., BERNARDINI F., MITTLEMAN J., TAUBIN G.: Acquiring Input for Rendering at Appropriate Levels Of Detail: Digitizing a Pietà. In *Rendering Techniques (Proc. of Eurographics Workshop on Rendering)* (1998), pp. 81–92. 2
- [RCC98] RUDERMAN D., CRONIN T., CHIAO C.: Statistics of Cone Responses to Natural Images: Implications for Visual Coding. *Journal of the Optical Society of America* 15, 8 (1998), 2036–2045. 4
- [RL00] RUSINKIEWICZ S., LEVOY M.: QSplat: A Multiresolution Point Rendering System for Large Meshes. In *Proc. of ACM SIGGRAPH* (2000), pp. 343–352. 2, 5
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient Variants of the ICP Algorithm. In *Proc. of International Conference on 3D Digital Imaging and Modeling* (2001), pp. 145–152. 4
- [SWS00] SUPKE J., WEHNER H., SZCZEPANIAK S.: Streifenlichttopometrie (SLT): A New Method for the

Three-dimensional Photorealistic Forensic Documentation in Colour. *Forensic Science International* 113, 1-3 (2000), 289–295. 1

[TBW*03] THALI M., BRAUN M., WIRTH J., VOCK P., DIRNHOFER R.: 3D Surface and Body Documentation in Forensic Medicine: 3-D/CAD Photogrammetry Merged with 3D Radiological Scanning. *Journal of Forensic Science* 48, 6 (2003), 1356–1365. 2

[TYS*02] THALI M., YEN K., SCHWEITZER W., VOCK P., BOESCH C., OZDOBA C., SCHROTH G., ITH M., SONNENSCHN M., DOERNHOEFER T., SCHEURER E., PLATTNER T., DIRNHOFER R.: Virtopsy, a New Imaging Horizon in Forensic Pathology: Virtual Autopsy by Postmortem Multislice Computed Tomography (MSCT) and Magnetic Resonance Imaging (MRI) - a Feasibility Study. *Journal of Forensic Science* 48, 2 (2002), 386–403. 2

[Wah] WAHL F.: *A Coded Light Approach for 3-Dimensional (3D) Vision*. Tech. Rep. IBM RZ 1452 (52546), IBM Zurich Research Laboratory (1984). 2

[Wah86] WAHL F.: A Coded Light Approach for Depth Map Acquisition. In *Proc. of DAGM-Symposium* (1986), pp. 12–17. 2

[XC04] XU H., CHEN B.: Stylized Rendering of 3D Scanned Real World Environment. In *Proc. of Symposium on Non-Photorealistic Animation and Rendering (NPAR)* (2004), pp. 25–34. 2

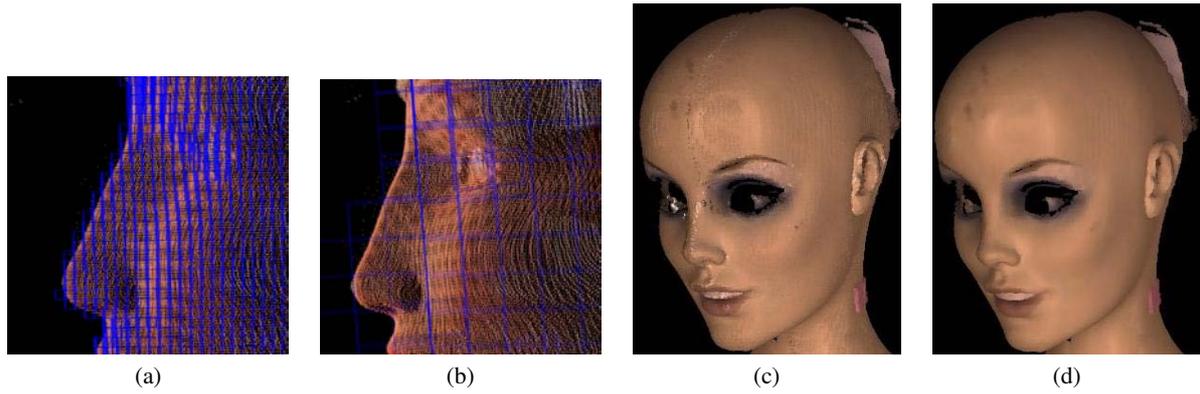


Figure 9: (a) Octree leaf block resolutions in left cheek area with 5mm resolution, and (b) 20mm resolution. (c) Filtering and data correction of the head of the Mannequin dataset: The positions of the data points are smoothed. (d) Color values of the data points are corrected based on a median filter.

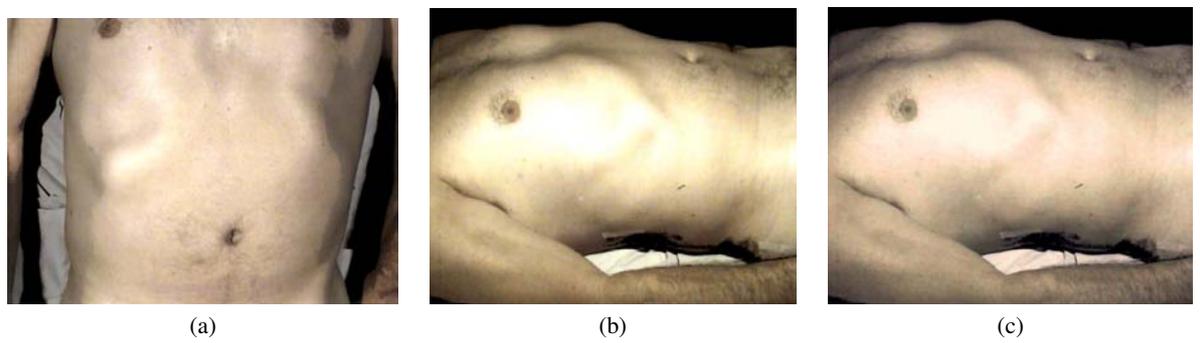


Figure 10: Color matching of the Accident Victim dataset: (a) Image with reference colors. (b) Image with yellow color cast that needs to be corrected. (c) Image with new matched colors.