

# Bimodal Visualization of Industrial X-Ray and Neutron Computed Tomography Data

Xuan Huang, Haichao Miao, Hyojin Kim, Andrew Townsend, Kyle Chambley, Joseph Tringe, Valerio Pascucci, Peer-Timo Bremer

**Abstract**—Advanced manufacturing creates increasingly complex objects with material compositions that are often difficult to characterize by a single modality. Our collaborating domain scientists are going beyond traditional methods by employing both X-ray and neutron computed tomography to obtain complementary representations expected to better resolve material boundaries. However, the use of two modalities creates its own challenges for visualization, requiring either complex adjustments of bimodal transfer functions or the need for multiple views. Together with experts in nondestructive evaluation, we designed a novel interactive bimodal visualization approach to create a combined view of the co-registered X-ray and neutron acquisitions of industrial objects. Using an automatic topological segmentation of the bivariate histogram of X-ray and neutron values as a starting point, the system provides a simple yet effective interface to easily create, explore, and adjust a bimodal visualization. We propose a widget with simple brushing interactions that enables the user to quickly correct the segmented histogram results. Our semiautomated system enables domain experts to intuitively explore large bimodal datasets without the need for either advanced segmentation algorithms or knowledge of visualization techniques. We demonstrate our approach using synthetic examples, industrial phantom objects created to stress bimodal scanning techniques, and real-world objects, and we discuss expert feedback.

**Index Terms**—Multivariate Visualization, Visualization Application, Image Segmentation, Volume Visualization

## 1 INTRODUCTION

ADVANCED manufacturing techniques enable creation of increasingly complex parts and assemblies with internal structures, multiple materials, custom-designed impurities, and a host of other features. However, to confidently utilize these capabilities in engines, buildings, or industrial facilities, resulting parts must be inspected and ultimately certified. Furthermore, many of the more advanced manufacturing techniques are rapidly evolving, and the inspection of as-built parts becomes key to improving the process. The challenge is that the additional degrees of freedom provided by additive manufacturing, i.e., 3D printing, often create internal features that cannot be observed with traditional metrology tools such as surface probes. Instead, the field has moved to high-resolution X-ray computed tomography (CT) scans to enable volumetric analysis. Although the X-ray scan provides the ability to observe internal structures, it creates new challenges in managing, analyzing, and visualizing large-scale volumes. Members of our team are at the forefront of yet an additional challenge in this space where often a single scanning modality is

insufficient.

As combinations of metals, plastic, and shielded parts are common in critical applications, developing approaches that can effectively characterize such objects is of significant interest to our collaborators. To address the limitations of one modality, additional modalities are used, and here we are primarily concerned with the addition of neutron-based tomography. Unlike the high-energy photons of an X-ray beam, neutrons interact only weakly with many materials [1]. As a result, neutron tomography does not suffer from the shielding effects of high-density materials. Furthermore, unlike X-rays, neutrons do interact relatively strongly with organic materials, such as plastics. Consequently, a neutron image can provide insights into materials invisible to X-rays as well as information about otherwise shielded parts. However, due to the limited power of available neutron sources, the spatial resolution of scans is lower overall in comparison to X-ray CT.

Users are now faced with two large-size volumes corresponding to X-ray and neutron attenuation that must be analyzed jointly. The ultimate goal for most industrial parts is a segmentation into materials for subsequent analysis. However, the large volumes, bimodal data, and common reconstruction artifacts, such as streaks or beam hardening, make any automatic segmentation challenging, and we currently have no accepted solution.

To facilitate the process of initial data understanding, our goal is to create an exploratory visualization that can quickly provide an overview of an object, search for obvious flaws, and plan more detailed analysis steps. The result serves as a solid basis for material scientists at an early stage of analysis, allowing them to make quick judgments on the scanning quality and the overall structure of the datasets

- Xuan Huang and Valerio Pascucci are with the SCI Institute, University of Utah.  
E-mail: [xuanhuang@sci.utah.edu](mailto:xuanhuang@sci.utah.edu) [pascucci@sci.utah.edu](mailto:pascucci@sci.utah.edu)
- Haichao Miao and Hyojin Kim are with the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory.  
E-mail: [miao1@llnl.gov](mailto:miao1@llnl.gov) [kim63@llnl.gov](mailto:kim63@llnl.gov)
- Andrew Townsend, Kyle Chambley and Joseph Tringe are with the Nondestructive Characterization Institute, Lawrence Livermore National Laboratory.  
E-mail: [townsend10@llnl.gov](mailto:townsend10@llnl.gov) [chambley1@llnl.gov](mailto:chambley1@llnl.gov) [tringe2@llnl.gov](mailto:tringe2@llnl.gov)
- Peer-Timo Bremer is with the SCI Institute, University of Utah and the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory  
E-mail: [bremer5@llnl.gov](mailto:bremer5@llnl.gov)

before investigating specialized analysis tasks.

Existing tools in the visualization community are largely based on linked view systems providing matching views of both channels or on complex multimodal transfer functions. The former requires users to mentally assemble different 2D images, which can be difficult especially for complex 3D assemblies. The latter implies significant training and familiarity with the underlying visualization concepts that our typical user will not have. Since neutron imaging remains an experimental capability, few dedicated solutions to these challenges exist, and current approaches to visualize bimodal data in manufacturing are largely limited to a side-by-side view of orthogonal slices extracted through custom scripts, i.e., in matlab, which are not adequate for a meaningful exploration.

Here we present a new approach jointly designed by a team of nondestructive evaluation experts and visualization researchers to interactively explore bimodal scans of industrial objects at an early stage of analysis. The system combines a semiautomatic topological segmentation of the bivariate histogram with an interactive approach to adjust the visual mapping from segments to visual properties and to explore the detected materials.

The topological segmentation potentially coupled with localized thresholds using the relevance metric [2], [3] provides a scalable preprocessing step, which creates an easy-to-explore space of hierarchical segmentations that the user can navigate by simply varying the number of segments shown. The direct link to the joint histogram enables domain experts to quickly form hypotheses on the type of material that is segmented, and simple interactions such as assigning colors to materials and adjusting per-segment opacities provide a flexible interface to investigate hypotheses on the material composition.

We identify our main contributions as follows:

- We propose the application of topological segmentation of the bivariate histogram for the exploration of multimaterial objects that are scanned with x-ray and neutron CT. This approach is enhanced by a relevance-based hierarchy, is fast to compute, and provides easy material decomposition overview of the multimaterial industrial objects.
- A combined view of co-registered volumes integrating the advantages from both modalities, with straightforward color-material mapping and real-time rendering using Intel's high-performance OSPRay engine [4]. (Section 5.1)
- An interactive widget for the straightforward editing of the topological segmentation results allows us to adjust the number, color, and visibility of segments to avoid complicated transfer function design for each modality. We utilize the 2D space of the segmented bivariate histogram as an interface for merging, editing, and creating material segments.
- A thorough evaluation of two simulated datasets and three real-world acquisitions. We present three use cases proposed by experts from diverse backgrounds in the nondestructive evaluation field, ranging from CT reconstruction, analysis, and evaluation, to experimental validation. The result covers both objective compar-

ison and subjective comments on the effectiveness of our system. (Section 6)

## 2 RELATED WORK

In this section, we look at previous work in bimodal data visualization with a focus on volumetric data. In particular, we conduct a literature search of visualization for industrial CT. We also include approaches that focus on topology-based segmentation techniques, examine the use of transfer function design within the context, and discuss existing bimodal visualization systems with similar case studies.

### 2.1 Bimodal Data Representation and Segmentation

The terms bimodal and bivariate have been used interchangeably to describe a dataset with more than one value at any sample point. One of the main reasons a bimodal dataset is desired is that different modalities provide complementary information. The feature of interest is usually defined as a combination of values across different modalities. Thus, multimodality often causes problems in user interpretation, and an effective data representation is crucial for further analysis. The survey by Lawonn et al. [5] examines various visualization techniques for CT data, but the solutions are spread across different medical subdomains, and there is no integrated framework for commonly used techniques.

Data segmentation is considered an important basic operation. Topological methods have been introduced to offer a mathematical abstraction [6], and Jankowai and Hotz developed an algorithmic solution to segment bivariate datasets through isosurfaces [7] for important features. These existing topology-based or isosurface-based approaches reduce the complexity in the target dataset by highlighting certain subdomains, but they pay little attention to a comprehensive initial overview.

The potential of bivariate histogram segmentation has already been explored by LaManna et al. [8]. The authors introduced a bivariate histogram method over two modalities to leverage material contrast in volumes [8]. Although LaManna et al. 's work supports our decision to employ histogram segmentation, their approach remains a manual process of placing polygons inside the histogram that outputs labeled binary volumes, but requires a typical active time of 15-30 minutes for visualization output. Our work takes this idea further and includes a visualization system with semiautomated segmentation and interactive mapping of segments to visual properties.

Building on top of Kniss's [9] multidimensional transfer function work, our segmentation method is also similar to the work of Wang et al. [10] with automated Morse-complex based segmentation on the value vs gradient magnitude feature space. Although this system uses a similar automated pipeline, the underlying 2D space asks for a visualization goal that does not work smoothly with Morse-complexes. With gradient magnitude as the second field, the materials to be identified do not correspond to the visually distinct arch regions, but rather circular structures at the bottom. Therefore, although the system overall provides an automated Morse-complex based solution like ours, the resulting

segmentation is an indication but not a direct representation of different materials. In our system, the resulting segments align well with the target decomposition piece in the 2D space, and modifications can be made to match the material decomposition process more exactly.

## 2.2 Rendering of Bimodal Datasets

**2D Projection:** Since the final outputs of the visualization system are images, multiple modalities also have to be integrated into these final images in a reasonable way, such that the user is able to extract desired information. One of the most straightforward multichannel blending schemes is through Maximum Intensity Projection (MIP). Fishman et al. have discussed the advantage of volume rendering versus MIP in CT angiography [11] and concluded that MIP is best for fast demonstration or enhancement of certain organs, but lacks the ability to provide a clear definition of all structures such as soft tissues and muscles.

On the other hand, volume rendering is also a standard solution due to its ability to produce a better visualization result in general. However, volume-rendering-based visualization usually relies on user mastery of rendering parameters and interpretation to reveal the feature of interests. Intermediate structures, such as isosurfaces or additional data analysis panels, are often necessary during the process and thus require additional mental effort.

**Data Fusion:** Another obvious direction is to combine all channels so that the standard scalar field volume rendering pipeline can be applied naturally. Therefore, the image fusion technique is widely used by industrial CT solutions to form a single scalar field [12] or to resolve a bimodal image [13]. Various metrics are discussed to blend multiple scalar fields [14], as well as the use of different intermixing schemes for direct volume rendering on GPU-based devices [15]. Redesigning the color blending function also helps eliminate false additive colors [16], [17], but qualitative analysis in the color space remains an open problem due to challenges in human perception. In general, although the technique itself is valuable for producing high-quality fused datasets for specialized analysis conditions, fusion is performed as an extra merge operation outside the visualization system. Thus, such techniques raise the risk of information loss or bias and allow little flexibility in overview exploration.

**Direct Volume Rendering:** Although 3D data rendering itself is well supported in modern graphics pipeline, multimodality poses additional challenges when targeting an interactive system. Due to the bimodal nature, the data size will double for bivariate datasets, and the multimodalities are considered to be only different RGBA channels that still require fusion. Schubert and Scholl [15] provide a data fusion solution with CUDA and GPU with interactive rate. Ghosh proposes hardware-assisted volume rendering [18] to speed up the process by efficient parallelization of multiple graphics hardware boards. Both techniques suffer from the inherent volume slicing artifacts, and a dual node PC cluster is used in the latter case because a single machine implementation does not perform well.

The visualization system in this paper uses direct volume rendering with a CPU-based rendering engine [4]. The

embedded high-performance ray-tracing method allows interactive rate on desktop machines for reasonable sizes scientific data and produces higher quality images.

## 2.3 Transfer Function and Interface Design

Transfer function design directly affects user interaction with the goal of balancing generalization and precise user manipulation. Automated clustering-like solutions are well explored in 2D transfer function design for scalar field volumes [10], [19], but the solution for multivariate datasets is not straightforward with dense value pairs, as in our cases. To overcome the difficulties in single techniques, interpanel linking is widely adopted to reflect multiple changes interactively [20].

Kniss designed a multidimensional transfer function interface with dual-domain and classification widgets [9] that is much more complicated than the common scalar field transfer function widgets. Illustration-based approaches establish indirect mapping from multidimensional attribute space to color space through formulations, such as sets, numeric operations, and semantics [21], [22], [23]. Similarly, the idea of fiber surfaces introduces new representations of multivariate datasets as shown in the work of Athawale et al. [24]. Whereas these methods provide expressive visual abstractions for bimodal datasets, the interpolated visual style elements may introduce ambiguity in material segmentation. SegMo [25] provides a straightforward Morse-complex based segmentation solution that is voxel-level precise but requires a heavy topological structure in volume space. In the direction of more specialized transfer functions, Sereda [26] discussed a semiautomatic hierarchical method based on the concept of the LH space boundary. The potential of dimension projection and parallel coordinates [27] is also explored. Lu and Shen also looked into using subspaces to reduce visual elements [28]. However, the newly introduced ideas and complex panels also pose challenges to users who are usually not in the field of visualization.

Finally, Tzeng and Ma [29] showed a cluster-space visual interface for preprocessed classified volume. We took a similar approach by including views of individual segments, but we provided a more flexible environment for users to make adjustments to the segmentation result interactively with the rendering.

This seminal work takes the segmentation from bivariate histogram, which is easily translated into color mapping on the bivariate volume attributes while hiding the complexity of rendering parameters. We also include a brushing and linking [30] functionality to enhance the connection between the histogram and volume rendering interface.

## 2.4 Bimodal Visualization Systems for Industrial CT

A survey by Heinzl et al. [31] examined the field of material science visualization and proposed several high-level challenges. Integrated visualization systems are often required for generated material science data to investigate features of interest, and a tailored feature extracting pipeline also has to be embedded for further analysis. Schiwarth et al. [32] described a workflow for a series of fiber-reinforced polymer X-ray computed tomography datasets. To show the analysis result for all features, such as fibers or voids,

the system uses a combination of 3D rendering and 2D multivariate scalar plot matrix. ImNDT [33] pushed the interactivity of exploring multivariate material data into an immersive virtual reality experience, by showing an overview with a grid-like displacement system so that the user can subdivide the domain spatially.

Our work avoids the high-dimensional charts or any additional data interpretation UI panel in that the bivariate histogram serves both as the result of segmentation and as color-mapping on the transfer function interface. Together with a multi-resolution data loading mechanism, by using an high-quality ray-tracing engine, we utilized the built-in state-of-the-art volume rendering data structures and algorithms that are parallelized and optimized for modern instruction sets, and thus enables an interactive rate, higher fidelity volume rendering for all datasets we use (see Table 1). Moreover, both the segmentation and the rendering components are designed in a modularized way such that other algorithms can be easily swapped into the pipeline for even greater flexibility in application development.

### 3 BACKGROUND

Our domain experts are scientists focusing on the development of novel methods for the characterization of advanced materials and industrial objects. They have carefully constructed a range of phantom datasets to stress test their approach, which we use in this work. The phantoms are multimaterial objects, containing small features that stress the single modality, shown in Table 1. Furthermore, as manufacturing these test objects and acquiring data with different parameters are time-consuming and expensive, our experts rely on physics-based simulations [34] to create additional data as the scanning systems are being developed. Physics-based simulations have the additional advantage that different setups can be easily explored, such as varying X-ray and neutron energies to plan for hardware changes or future systems.

X-ray imaging, although appreciated for its high spatial resolution, struggles with high-density metals, such as tungsten. As a result, important features could be hidden behind these materials, leaving this modality vulnerable to overlooking crucial defects. Neutron imaging, on the other hand, has a similar resolution everywhere, but a lower resolution overall [35]. The combination of these two modalities promises a vast improvement in the ability to detect features that are otherwise difficult to identify.

The complementary nature of these two modalities motivated our collaborators to construct a scanner that is able to acquire X-ray and neutron volumes simultaneously, where the X-ray and neutron sources are offset by a fixed angle. As a result, the data will not require computational registration. Our collaborators primarily try to identify the materials and their boundaries. Currently, this task relies on a side-by-side comparison of orthogonal slices from the two modalities extracted through scripts or simple interfaces.

#### 3.1 Bivariate X-Ray and Neutron Data

We use seven X-ray and neutron CT datasets: JH2A, JH2B, XR05A, XR05B, Synthetic Cylinder, battery and meteorite.

TABLE 1: The datasets used in this work. The phantoms are created by our collaborating experts to test acquisition, analysis, and characterization of materials using X-ray and neutron imaging. The phantoms are created using a physics-based simulation [34]

Phantom object	Type	Size
JH2A	Physics-Based Simulation	2x128x128x128
JH2B	Physics-Based Simulation	2x1000x1000x128
XR05A	Physics-Based Simulation	2x512x512x512
XR05B	Acquisition	2x256x256x256
Synthetic Cylinder	Simulation	2x512x512x512
Battery	Acquisition	2x2159x2159x2559
Meteorite	Acquisition	2x1999x1999x1549

JH2A, JH2B, and XR05A are physics-based simulations of nested cylinders with a small volume of material inclusions in the center of the object. Whereas JH2A and XR05A contain similar rod-shaped inclusions in the center with varying sizes, the JH2B dataset contains even smaller inclusions.

The Synthetic Cylinder is created without physics-based simulation to clearly demonstrate the advantage of bivariate histogram segmentation, and all its materials are concentric cylinders in space. All datasets have matching sizes on both modalities and are co-registered. All histograms and thus the following segmentations are computed on full resolution. The last two datasets, Battery and Meteorite, are downsampled to 1/4 on each axis in the rendering process for desktop performance (Section 4). To demonstrate our interactive approach to actual acquisition, we also include XR05B and battery, which are scanned from constructed objects, and a meteorite dataset that we analyze without any previous knowledge.

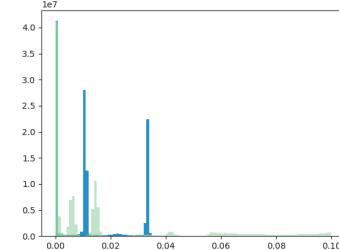


Fig. 1: The 1D histogram of the X-ray (blue) and neutron channel (green) of JH2B, which has five materials. The distinct peaks differ greatly, showing that one modality is inherently insufficient to capture all materials. In the X-ray, we can see three distinct peaks, whereas the neutron has three lower density peaks and a wider range of values that are spread out without any significant additional peak.

In Figure 1, we can see two histograms of the two modalities of the JH2B dataset, which consists of five materials. The X-ray and neutron clearly react differently to the materials, resulting in distinct peaks over the 1D histogram. As demonstrated by the number of peaks, neither single modality can capture all the material in the phantom object. The X-ray provides an overall high-quality capture while losing precision at the interior after hitting high-density metal. The neutron behaves consistently with less detail, but is able to capture the internal structure. We demonstrate our

method with JH2A and JH2B in section 4 and section 5. All datasets are discussed in more detail in section 6 with the domain experts.

## 4 MULTIMODAL VISUALIZATION OF X-RAY AND NEUTRON CT

We work with domain experts in nondestructive evaluation and design the visualization system with the goal of an efficient multimodal data analysis. The main idea is to provide a quick overview of features from both X-ray and neutron CT for the scientists without going into technical details in the field of visualization. Although experts traditionally inspect the difference in density values, their overall goal is to understand the material compositions. Hence, our work offers both an effective overview with minimum user interaction as well as the flexible inspection of individual materials.

The pipeline is summarized as follows (Figure 2): The material identification is semi-automated. Our approach first creates the bivariate histogram and a topological hierarchy. The user can specify the number of segments, which are then assigned color and opacity. For analysis, the user can toggle the visibility of each segment through the interaction panel. Since some datasets are too large to fit into memory, we compute the histogram on the full-resolution volume but utilize the multiresolution approach by Kumar et al. [36] and the OpenVisus [37] framework to downsample the data for rendering. Our multimodal renderer looks up the color based on the localization of each value pair in the histogram and visualizes each segment with the respective color.

With a Morse-complex based segmentation on the bivariate histogram formed by both modalities, we enable the user to oversegment the bivariate histogram from the two volumes into regions that roughly correspond to object materials. The outputs from this step are 2D segmentation images with color assignment of various number of segments. All images will then be taken as input into the renderer as lookup tables. Material colors can be adjusted by modifying the segmentation image on a simple interface.

### 4.1 Domain Goals in Nondestructive Evaluation

We conducted our work as part of a large project aimed at developing new methods for the nondestructive evaluation of industrial objects. An important component is the use of multiple modalities that produce complementary images of the same objects because, as shown in Figure 1, a single modality is insufficient to capture all materials in the object. Over the course of a year, we met weekly to discuss challenges and refine goals for the multimodal analysis. We were tasked with developing an effective method to examine X-ray and neutron data simultaneously without the need to mentally match multiple views of slices, so that the scientists are able to gain initial yet comprehensive knowledge of the acquired datasets to understand material compositions.

From the above-mentioned considerations and description of challenges, we identify the following goals in terms of visualization and analysis:

- Show an efficient overview of materials in industrial objects that are simultaneously scanned with X-ray and neutron CT.

- Avoid labor-intensive operations such as adjusting transfer functions or manual segmentation.
- Leverage existing views that are already familiar to the experts.
- Enable interactive exploration and the ability to see the shapes of individual material regions.
- Keep user interactions at a minimum while providing flexibility in segmentation and visibility parameters updates.

### 4.2 Bivariate Histogram Generation

1D histograms are commonly employed by the experts to understand the material compositions without spatial information. The bivariate histogram has already been shown to be effective for identifying distinct materials from X-ray and neutron data [8]. Another advantage in utilizing the value frequency in the bivariate histogram for segmentation and the inherent binning of the values is that it also accounts for the previously mentioned artefacts in the acquisition process. If the value pairs are directly taken as input, the subsequent morse-complex segmentation would result in false positive segments as a result. Finally, we also consider the familiarity of our collaborators with bivariate histograms for material identification and this approach enables us for a seamless adoption to their existing workflow.

The bivariate histogram provides the underlying data for the segmentation. The materials appear as connected pixel areas formed by different value combinations of the two modalities, but not all materials will necessarily result in high peaks. As a 2D domain, the histogram consists of, for example, 100x100 bins that provide us with a quick view of material clusters. The peak in the histogram will likely correspond to a material type, as in the 1D case. We have seen in Figure 1 that each 1D histogram has fewer material peaks than the actual number of materials. It is easy to see in Figure 3 that the X-ray 1D histogram on the horizontal axis has brown and light blue materials as similar values, which make those indistinguishable on purely X-ray, whereas the neutron 1D histogram on the vertical axis will not separate green, orange, and brown materials. However, in the bivariate histogram, we can see the distinct material peaks in the 2D domain. Figure 3d shows five clusters that, in a simplified assumption, correspond to the materials. The background peak on the bottom left is naturally the highest.

### 4.3 Morse-Complex Segmentation

#### 4.3.1 The Base Algorithm

Similar to other approaches that have used topology-related approaches to segment bivariate histograms [19], [38], [39], the key insight is that, at least in our use case, such histograms are formed through a superposition of blurred Dirac functions, one for each identifiable material, including the surrounding air, i.e., the background. In a perfectly resolved scan, with uniform materials, and without noise or artifacts, each material would be identified through a unique neutron-X-ray pair of values and be confined to a single bucket of the histogram. In practice, noise in both the X-ray and neutron sources as well as in the detector, assumptions in the underconstrained reconstruction problem, material variations, partial volume artifacts, etc., cause

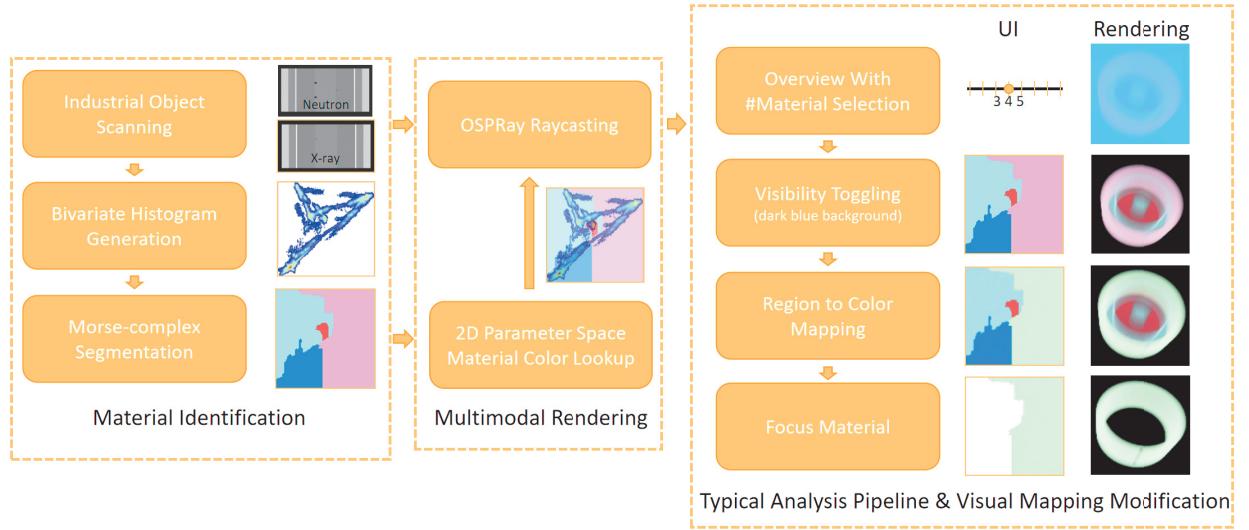


Fig. 2: Pipeline overview. We load the X-ray and neutron data at full resolution, generate the bivariate histogram, and compute the Morse-complex segmentation. Both the segmentation and the dataset itself are set as input for the renderer. During the ray casting process, when a voxel is queried it will be sampled for both modalities, resulting in a pair of sampled values. The 2D vector corresponds to the 2D position in the bivariate histogram. The RGB values are determined by the same coordinate lookup in the segmentation image, which is then modifiable through the application's built-in GUI.

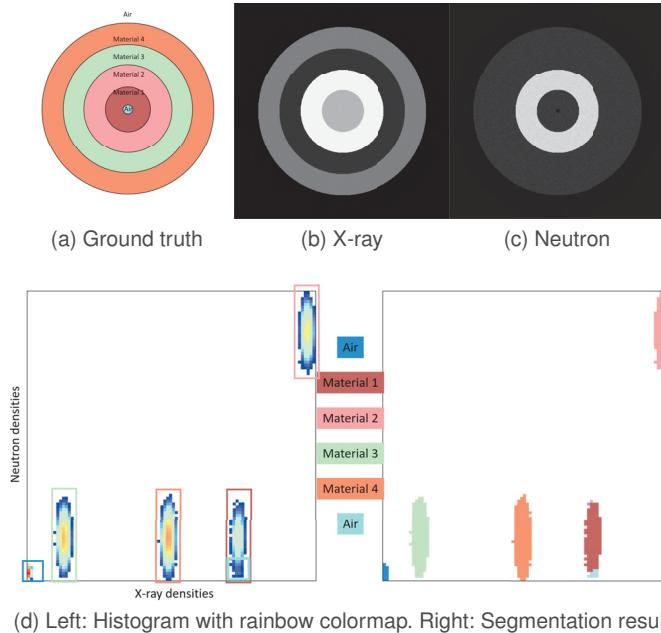


Fig. 3: A Test example: The Simulated Synthetic Cylinder that contains four concentric cylinders of different combinations of values from the X-ray and the neutron channel. (a) Shows the ground truth, (b) and (c) show a slice of the two channels respectively, with X-ray missing the innermost hole and Neutron falsely mixing orange and green materials. (d) Shows that the materials and backgrounds are well separated into 4+2 clusters on the bivariate histogram on the left. The information is successfully captured in the Morse-complex segmentation on the right.

these Dirac peaks to be blurred into a more continuous response. Computing the Morse-complex, also called the watershed, of this response function identifies all local peaks

and their corresponding region of influence and thus splits the contributions of the various materials as well as possible without assuming further knowledge of the system, i.e., the number and types of materials present or the total amount of material in an object.

Here, we use a simple steepest gradient style algorithm [40] to identify all local maxima and compute an initial segmentation. More specifically, by interpreting the histogram densities as function values on a quadrilateral grid, each vertex computes its steepest ascending neighbor and breaks ties through standard simulation of simplicity. Vertices with no ascending neighbors are labeled as maxima, and the corresponding segment is defined by its *stable manifold* - the collection of vertices whose steepest gradient path will end at the given maximum. Finally, saddles are identified as the highest vertices whose immediate neighbors belong to more than one segment. As apparent in many of the images, this approach suffers from some mesh imprinting, especially in areas of low overall density, which could likely be improved with more advanced techniques, such as the accurate Morse-Smale complexes of Gyulassy et al. [41]. However, as discussed in more detail below, the areas of low density in the joint histogram by definition have minimal effect on the final image as very few voxels are involved. In regions of high density, we found the steepest gradient algorithm quite reliable even when dealing with various sources of noise. Note also that the algorithm described above does not identify the full Morse complex as defined in traditional discrete Morse theory [42], [43]. For example, we avoid the complexity of identifying saddles as edges between vertices, do not explicitly represent minima or the potentially degenerate one manifolds from saddle to minima, and will not detect *strangulations* (saddles connected to the same maximum twice). However, for a simple peak-finding approach, the above algorithm is equivalent to more sophisticated solutions and trivial to implement.

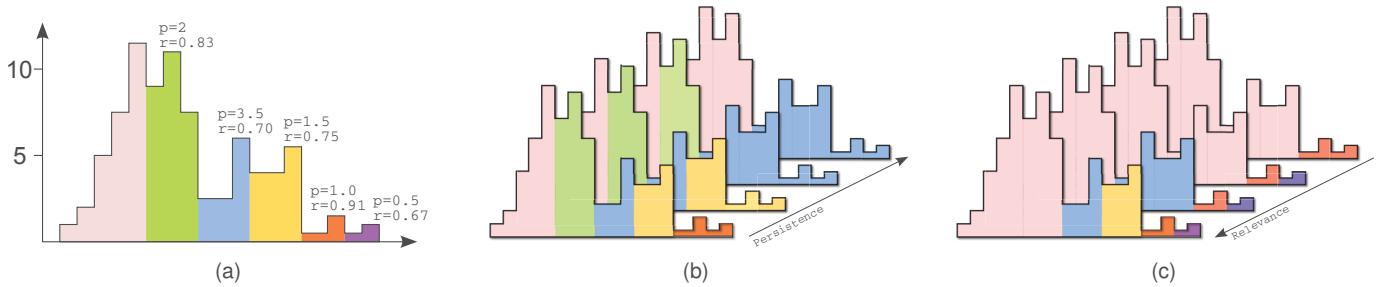


Fig. 4: Persistence vs. relevance simplification in a 1D histogram. (a) Original histogram segmented into modes indicated by color with the relevant persistence and relevance values of the local maxima indicated. (b) Hierarchical simplification by increasing persistence, which prioritizes high value peaks corresponding to common values. (c) Hierarchical simplification by decreasing relevance, which prioritizes distinct peaks when compared to their local neighborhood and preserves rare values.

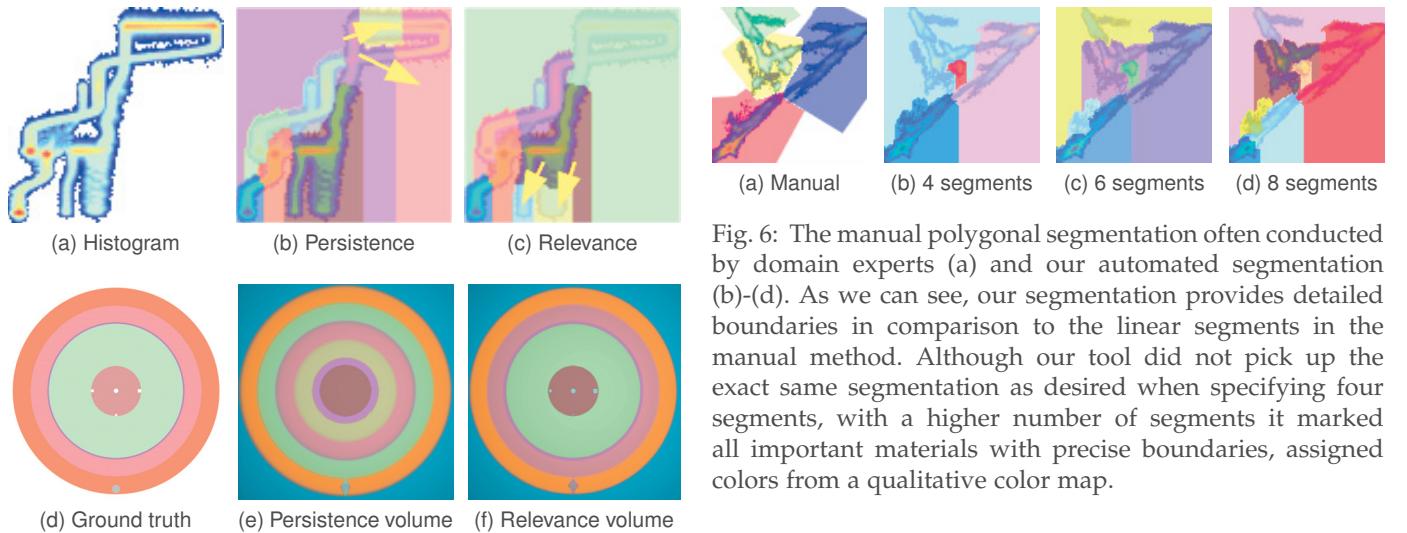


Fig. 5: JH2B bivariate histogram (a) segmented with persistence (b) and relevance (c). With persistence, the purple color segment in (b) tries to further split the highest value peak into yellow and pink segments indicated by arrows, whereas the relevance imposes a better segment-material correlation by identifying a possible new region around the lower value area (brown, light blue, and yellow in (c)). (d)-(f) shows the ground truth illustration and the volume result of segmenting by persistence vs by relevance. The persistence scheme tends to segment out the bigger volume with smooth value variation, whereas relevance prioritized the smaller pieces with more distinct value peaks that are more likely corresponding to individual materials.

#### 4.3.2 Problem With Persistence Simplification

The next challenge is that discrete approaches like the one described above are well known to identify a large number of spurious maxima that must be discarded for the segmentation to become meaningful. The most common approach to simplify Morse complexes is driven by *persistence* as proposed in Edelsbrunner et al. [44] for histograms. Other approaches include weighted sums of persistence and spatial distances [45] or volumes of segments [46]. The persistence in this context can be thought of as the height of a peak,

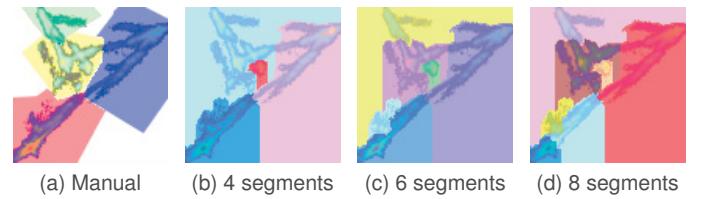


Fig. 6: The manual polygonal segmentation often conducted by domain experts (a) and our automated segmentation (b)-(d). As we can see, our segmentation provides detailed boundaries in comparison to the linear segments in the manual method. Although our tool did not pick up the exact same segmentation as desired when specifying four segments, with a higher number of segments it marked all important materials with precise boundaries, assigned colors from a qualitative color map.

i.e., the difference in density between the maximum and the highest neighboring saddle. To simplify the segmentation, one iteratively identifies the remaining maximum with lowest persistence and merges the corresponding segment with the segment on the other side of the corresponding saddle. However, although intuitive for terrains, persistence is not an ideal metric for bivariate histograms. By definition, the height in a histogram indicates the number of voxels in the original data that share the corresponding value pair. Consequently, how large a particular peak is does not necessarily correlate with its importance, but only with the quantity of the material present. In fact, by far the largest peak is typically associated with the background material, air in our case. Therefore, using the absolute height and even the absolute difference in height between a peak and its saddle as a measure of importance will significantly overvalue common materials and ignore rare combinations, i.e., small inclusions. Such a result is the opposite of the desired effect since the most common materials in an object are often the least interesting.

#### 4.3.3 Modification for Our Cases: The Relevance Metric

We propose a different metric based on relative persistence, also called *relevance* [2], [3]. Figure 5c. Consider two neighboring segments capped by maxima  $m_1$  and  $m_2$  separated by the saddle  $s$  with  $f(m_1) < f(m_2)$ . The persistence of

the  $m_1 - s$  pair is defined as  $f(m_1) - f(s)$ . Its relevance is defined as  $1 - (f(m_1) - f(s))/f(m_2)$  as the global minimum of our histograms is assumed to be 0. Conceptually, relevance acts more like a local signal to noise ratio. In a high-density region of the histogram, even a comparatively small amount of noise may cause large density variations. On the contrary, in a low-density region the same amount of noise will cause much smaller variations. Persistence measures absolute variation whereas relevance measures local variation. As a result, simplifying according to relevance better preserves low peaks corresponding to the rare materials. Figure 4 illustrates the differences between a persistence- and a relevance-based segmentation for a 1D example. We also impose a small absolute lower bound on the density and will not consider maxima with less than a certain number of samples. This action is necessary to prevent numerical noise from being amplified in the relevance computation in regions of minimal density. The final result is a simple-to-implement and fully automatic hierarchical segmentation, which identifies materials according to how (relatively) distinct they appear in the joint histogram. The result corresponds well to how our experts would manually perform the segmentation and avoids common problems when colormapping histograms of large ranges of densities.

We choose relevance because favoring smaller features instead of further subdividing higher peak gives a better segmentation for our datasets.

## 5 HISTOGRAM-BASED BIMODAL RENDERER

Based on the segments from the histogram segmentation, we can now render the two volumes. We developed a bimodal renderer that loads the segmentation as 2D texture. Together with the segmentation images, the bimodal CT datasets are passed to the renderer as raw format. To reach the goal of a joint view of both X-ray and neutrons CT images for domain analysis, the two channels are composited in the renderer in a user-specified configuration. The options cover a wide range of operations, from choosing the number of segments to setting the color or visibility of individual segments. In this section, we describe the details of our implementation. The entire project is open-source and can be found at: <https://github.com/xuanhuang1/multiChannelOSPMModule>.

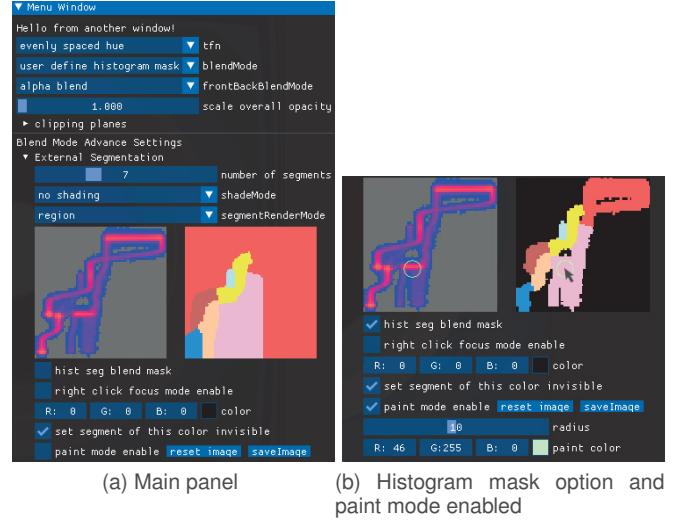
### 5.1 Co-Registered X-Ray and Neutron Volume Renderer

We used the open-source ray tracing engine OSPRay [4], and modified the built-in Scivis renderer for a true bimodal ray-casting rendering pipeline. Combined with the histogram segmentation, our visualization encodes original information from both channels.

The transfer function is defined as a color lookup from the segmentation images. For each voxel sampled, the density value pair from the X-ray and neutron data determines its 2D coordinate in the color segmentation image, which is passed to the renderer as a 2D texture. Then the corresponding pixel on the texture is sampled, and its RGB color is assigned to the voxel. Opacity is set as uniform value for all segments. The resulting renderer is a modified ray-caster

with a specialized transfer function mapping. We chose our current pipeline for interactivity and a straightforward rendering result to users.

## 5.2 Interactions



(a) Main panel

(b) Histogram mask option and paint mode enabled

Fig. 7: Graphics user interface. The main menu in (a) lists user operations to set shading options and the number of segments. By clicking on the segmentation image, the user is able to set color and visibility for each segment. As shown in (b), a histogram over segmentation mask option is available for visual correlation. A paint mode widget (enabled here on the right) is located at the bottom for further manual segmentation editing. Extra utilities include global opacity scalar and clipping planes. The entire interface consists of only commonly seen elements and simple operations.

In order to reduce the amount of interactions, the segmentation process only requires the user to select the number of segments. When the volumes do not contain any noise, the number of segments correspond exactly to the number of materials plus the background. However, as X-ray and neutron scans will not necessarily capture all materials as the same density value pairs, as described in subsection 4.3, the number of segments will also depend on the quality of the scans, the material, and the location in the volume. By providing the flexibility to adjust the number of segments, we allow the user to explore different segmentations without the need to manually draw polygons in the histogram, or segment the volume. The bimodal renderer provides a quick overview of the segmentation, as shown in Figure 2. The main features of interaction are listed below:

- To resolve the inherent visual occlusion problem with volume visualization, we provide the user with two ways to toggle the visibility for all available segments: all visible except for the selected segment or none visible except for the selected segment.
- Further flexibility and support for exploration is provided through the color assignment for each segment.
- An overall opacity scalar value slider for quick transparent overview or solid material observation.

- A segmentation editing panel that allows free color painting directly on the histogram segmentation result, acting as a manual refinement tool on top of an automatic segmentation. This widget interactively updates the volume rendering mapping, which enables user interaction with small 2D areas in bivariate space plus the the precision and convenience of an automated process. Figure 7 shows individual components in detail.
- Since the 2D space pixel color is simply the segmentation ID for each material, the resulting image from editing encodes the desired segmentation for this visualization session and can then be saved as a simple png file for further analysis tasks.

## 6 RESULTS AND EXPERT FEEDBACK

We demonstrate our approach on four datasets that our domain scientists use. From the weekly meetings with them, we derived the main case studies that we describe in the context of each individual dataset. We collected feedback from the four domain experts by demonstrating the visualization at regular meetings and, in addition, two of the experts used our system directly to study the datasets described in Table 1. All domain experts are senior scientists that are focused on nondestructive evaluation with more than 10 years of experience. They lead groups or projects for the characterization of materials in advanced manufacturing. From these informal sessions, we collected the experts' informal feedback and formed the following four case studies. We also describe our observations on their usage, what they found easy to understand and what was difficult. We summarize their responses at the end of all user case analyses in subsection 6.5. The users tested the approach on a laptop with an Intel i9-1088H @ 2.4 GHz processor, 64 GB RAM, and a Nvidia Quadro RTX 3000 (6 GB RAM). The application maintained at least around 10 fps on all case studies.

### 6.1 Case Study 1 JH2B and XR05A: Semiautomated Fast and High-Precision Material Segmentation

In this section, we perform a qualitative, material-to-material analysis to demonstrate the effectiveness of the segmentation system. We present our result with two datasets, JH2B and XR05A. The segmentation method is also compared to the cluster-based k-means segmentation as well as Wang et al. 's gradient magnitude based approach.

In the JH2B dataset, the number of materials in the structure is known, but the number of X-ray and neutron value pairs for a material may be different depending on the other materials in the object due to the scanning artifacts described earlier. In particular, this dataset has long, thin cylindrical rods in the center of the object that often require closer attention. Here, we first demonstrate the histogram-based approach to identify material peaks. As shown in Figure 6, the bivariate histogram already indicates several high peaks that our experts would associate with the corresponding materials. They also confirmed that they would draw a similar segmentation manually, but Morse-complex segmentation is able to identify a detailed boundary that would require the placement of polygons with many sides, which would be tedious to do by hand.

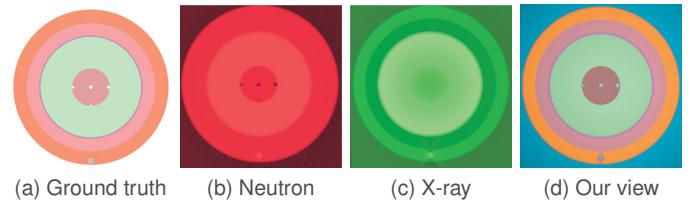


Fig. 8: JH2B ground truth, neutron (red), X-ray (green), and the bimodal rendered result. The same slice of JH2B is shown. The structure consists of aluminum (orange), HDPE outside (pink), stainless steel (purple), tungsten (green), and HDPE inside (brown). Note that tungsten not only blocks X-ray from depicting the inside of HDPE, but also results in a wide range of X-ray values due to absorption. Neutron preserves the interior rods but does not distinguish aluminum and HDPE. Our final view successfully combines the desired features from two modalities.

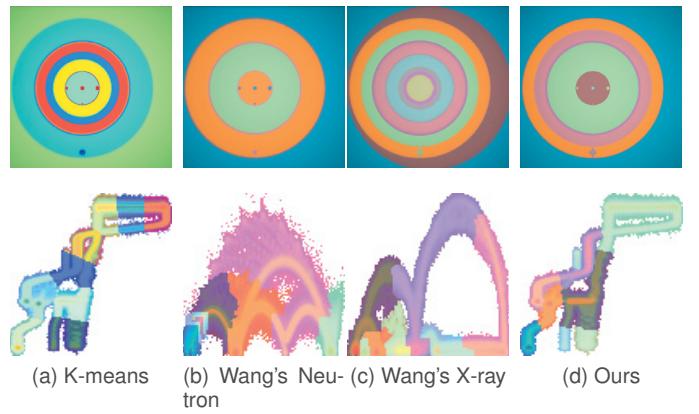


Fig. 9: JH2B 2D histogram segmentation methods comparison. Column (a) shows the k-means clustering with the histogram as scatter plot, which doesn't conform to material shapes at boundaries. (b) and (c) Shows Wang et al. 's method that uses gradient magnitude as the second axis on histogram for each modality and apply the Morse-complex segmentation individually. Note how the Morse-complex segmentation does not pick up individual arcs easily, and reflects the drawback of each modality as seen in Figure 8 (b) and (c). Column (d) shows a more desired result from our method.

The expert is interested in a closer look at the dataset to see material boundaries and any small features in the volume, and therefore the volume is clamped onto the slice-like view that experts are familiar with. We will also compare our result with the phantom images produced by the existing traditional slice-view method.

As we can see in Figure 8, the X-ray (green) fades undesirably with a lower precision at the center but is of higher quality overall, whereas the neutron (red) remains consistent within the materials but has lower accuracy in terms of reconstruction. Our result includes both the holes on the inner ring from the neutron and high-contrast boundaries from the X-ray, generating a result very similar to the ground truth.

To demonstrate the effect of different segmentation

methods we applied three 2D histogram segmentation on the JH2B dataset expecting nine resulting segments, as shown in Figure 9. Comparing to the ground truth illustration Figure 8 (a), the clustering-based method groups the 2D points by distance and does not catch precise shape boundaries as in the topological-based method, where the segment boundary are intentionally defined as along the local minima.

We also compare our segmentation to the method of Wang et al. [10], which performs the Morse-complex segmentation on the bivariate histogram of the value and gradient magnitude 9 column (b) and (c). whereas their Morse-complex based method segments the histogram more toward distinct peaks, the material identification of these composed histograms mainly relies on recognizing arc motifs that correspond to material boundaries. This misalignment between the nature of the Morse-complex based method and the segmentation goal results in inefficient material identification. For example, in column (b) on the bottom we can clearly see four distinct curves, but the orange and pink colors actually split and merge part of these structure randomly. Besides, our key goal was to provide a joint view of the two modalities, and using the their method would still require us to mentally assembling two complementary visualizations that each lacks different pieces of information, as seen in Figure 8 (b) and (c). Our method in Figure 9 column (d) demonstrates a more efficient 2D histogram segmentation that prioritize key materials and is straightforward to interpret.

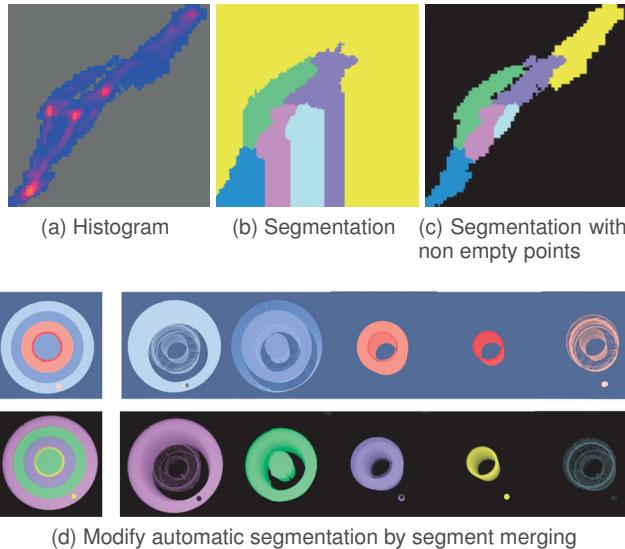


Fig. 10: A five-component segmentation with XR05 (background set to transparent). This is a slightly different variation of the JH2B dataset. (a) Shows the histogram segmentation. The top row in (d) is the manually labeled volume, which can be viewed as the ground truth of the segmentation. The second row in (d) is our result. We have reached a close approximation with only five segments. The small outer rod piece gets blended into other segments due to the value inaccuracy in the volume data itself and thus shows no distinct peak in the histogram, but can be easily identified spatially.

The XR05A is another simulated dataset with a higher quality X-ray. This dataset is expected to produce high-quality segmentation, and we validate our results with a labeled volume that describes the desired segmentation.(Figure 10).

We are aware that the segmentation does not group components precisely, in that we can see the tube on the side appears in two segments. Not all material is segmented into its own segment because of the inconsistency of the X-ray, resulting from the data acquisition step itself and thus introducing noise into the bivariate histogram compared to the ground truth. However, the final visual result could be improved by oversegmenting and merging, as shown in the next case study.

## 6.2 Case Study 2 JH2B: Exploratory Material Identification via Segmentation Adjustment

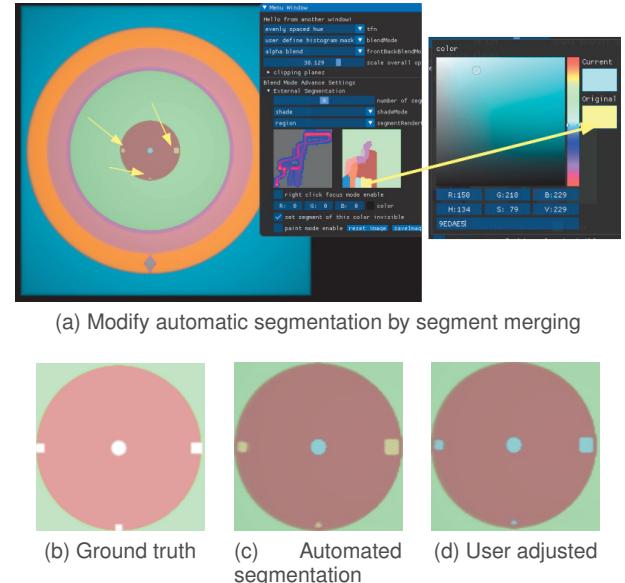


Fig. 11: JH2B segment merging: (a) shows the process of changing color with the color picker from yellow to light blue, so that the two segments get merged, and (b)-(d) are comparisons to the ground truth.

One of the main goals of the visualization task is to quickly perform exploratory analysis of the structure to see if important features can be easily spotted. These exploratory tasks are achieved by adjusting the number of segments and setting the opacity. In this case study, we first demonstrate the interactive adjustment of the segmentation results with the JH2B dataset.

The users in this experiment first applied clipping planes to cut along the xy-plane. Once a location of interest was reached, the number of segments was set to a higher number for precise segmentation. As shown in Figure 11c, the histogram is initially oversegmented until all parts receive their respective segments. Here, the user decided to merge the light blue and light green segments by assigning them the same color, since they are made of the same material. This case study also involves the greatest amount of interactions, but the process is still intuitive as all the widgets are widely used in common scalar field visualization interfaces. We

demonstrate here the ability of the interactions to easily correct the segmentation. The result after adjustment is shown in Figure 11c, which can be compared with the ground truth in Figure 8.

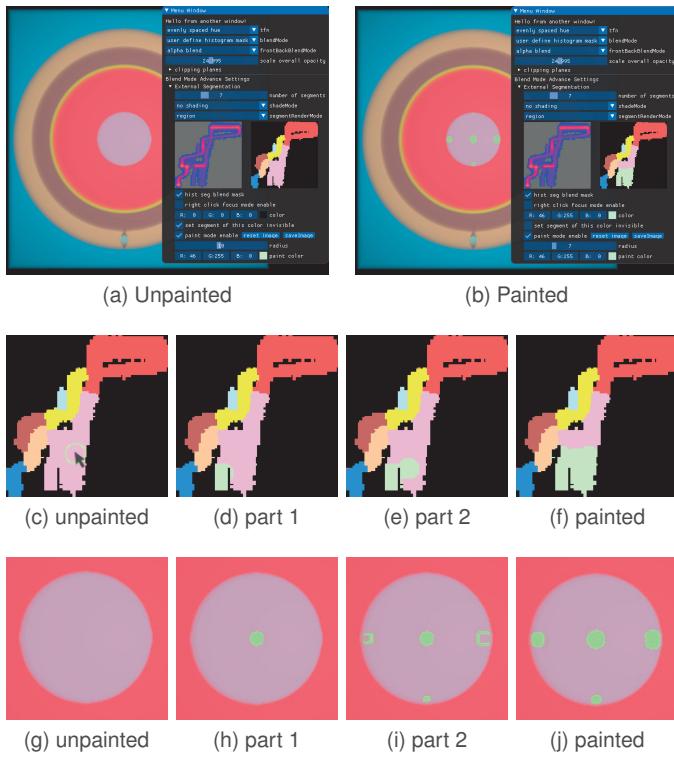


Fig. 12: JH2B user-adjusted. Here we show an example of an automated segmentation with manual modification. The process of painting is shown in the two bottom rows.

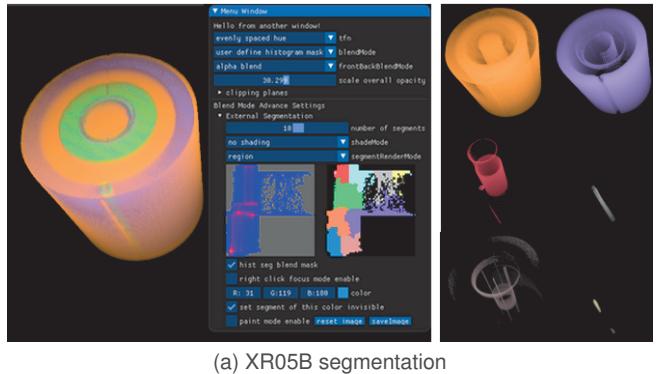
An alternative to oversegmentation with merging is to add the missing material on the histogram with a smaller number of segments. As shown in Figure 12, with the case of seven segments, the structure of the outer layers is already clear, but the interior holes are still missing. By looking at the bivariate histogram, the domain scientists suspect that the finger-shape areas in the bottom center likely form another distinct peak and thus represent a new material, but have not yet been segmented out from its pink surroundings. Therefore, in this situation, the user can utilize the painting tool to test the assumption. When the paint mode is enabled, the user selected a color to paint with mouse at the desired region, and the volume was updated interactively as seen with column pairs (c)-(j). Figure 12.

### 6.3 Case Study 3 XR05B and Battery. Bivariate Data Validation with Manufactured Scanning Objects

In this section, we present the ability to perform fast segment-based visualization on two real-world scanning datasets, XR05B and Battery. For these two real-world datasets, although the scientists are familiar with the structure of the original constructed object, the scanning and registration quality remains unclear, and thus our work fits well with the need for a quick check of previous steps. We compare the process with the slice-based views that are

commonly used, which is also the only easily accessible validation tool that we are aware of for the bivariate datasets.

The XR05B object has been deliberately constructed to stress the scanning process, and therefore the user of this case would like to have a fast validation of the data quality. Because the scientists rely on slicing for general initial observation, the XR05B testing object is designed to be identical along z-axis, similar to the JH2s. Based on the knowledge of such underlying structures, the scientist was able to use the automated segmentation result to spot important visual clues of potential acquisition defects. Figure 13.



(a) XR05B segmentation

Fig. 13: XR05B with the collection of segments on the right. Segmentation results in partially missing pieces.

To validate the quality of this acquired bivariate dataset, the user first noticed the bleeding effects on the histogram (blue scattered points in Figure 13) spreading out straight across the axis. Those horizontal thin lines were less often seen in the actual density value pairs, but were more likely resulting from scanning or registration artifacts. Therefore, after identifying several tilted or partially missing components that corresponded to the bleeding area on the histogram. Figure 13, the user was then able to confirm that there is a problem from a previous process, potentially a registration misalignment. This dataset would likely be unsuitable for further analysis, and the subtle misalignment would be difficult to diagnose by viewing the two channels side by side. Our system has successfully served as a fast visual system to observe the data quality resulting from the acquisition step.

For most real-world datasets, the data distribution is not as convenient as for JH2B or XR05B. Battery is another real-world scanned object that is more complicated to decipher. The traditional data interpretation process requires multiple slices at exact locations for both modalities to identify all structures, whereas with our system it is straightforward in the overview Figure 14.

### 6.4 Case Study 4 Meteorite: Investigating Distinct Material Quantities in an Unknown Object

The experts working with meteorite had little knowledge of this dataset at the time of viewing. The object is scanned to inspect the inside materials. In this case study, the analysis is exploratory and the main goal is to obtain a comprehensive first-step view of the interior structure. The feature of interest for this dataset is the iron-nickel inclusions.

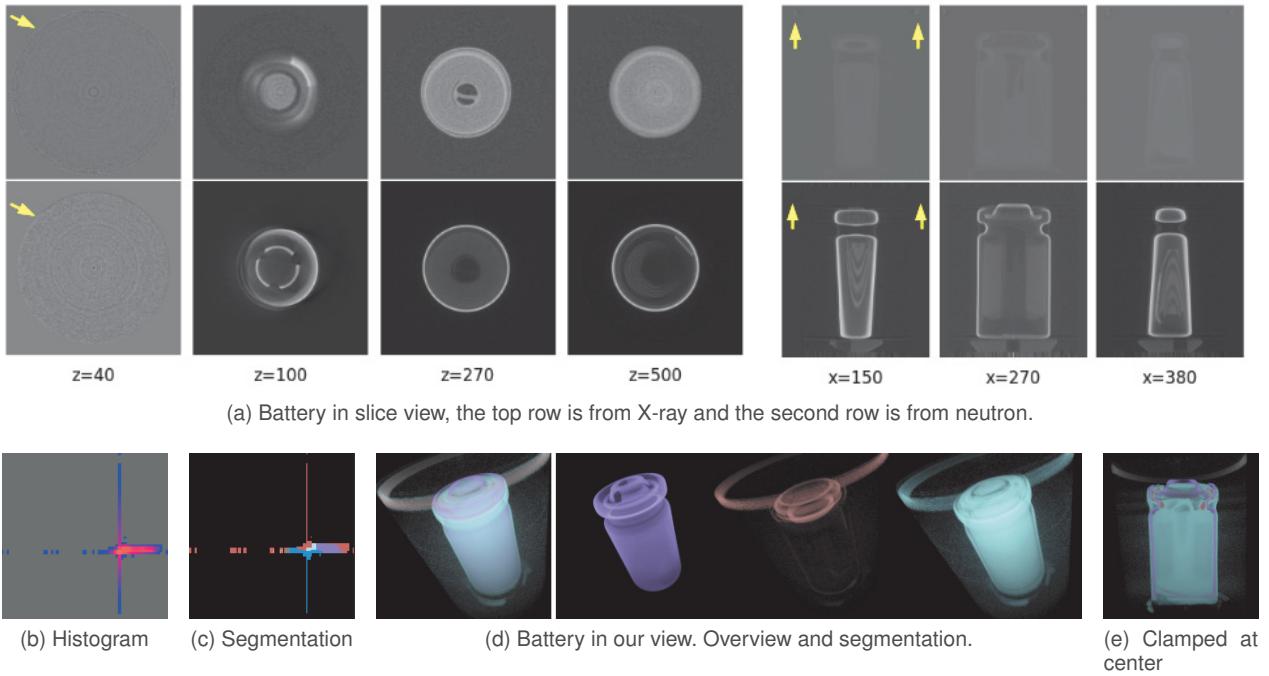


Fig. 14: Battery: (a) shows a traditional slice based view of the dataset. Our work (b) shows a more intuitive overview of the datasets, and a clamp operation as in (c) enables a close look of the segmentation similar to slicing for a familiar view. The bimodal visualization system provides a good estimation of the structure through automated segmentation while avoiding dealing with a stack of 2D slices in both modalities. Note that our work also clearly captures the upper ring, which is blurry in X-ray and almost missing in neutron.

The assumption is that this alloy lies in the meteoroid and can be well detected by the scanning devices, as the substance has different density values from the surrounding rocks. However, since iron–nickel is a small feature scattered across the entire object with no uniform shape, this dataset is considered an extremely labor-intensive case for traditional material identification techniques.

the segmentation, the user was then able to loop through each individual segment easily and got a good initial understanding of the dataset itself. The users stated that this is a highly relevant use case for our method, as they can use our tool to clearly view the various inclusions in the object with just a few clicks. This tool helps to make a more informed decision about further nondestructive evaluation techniques.

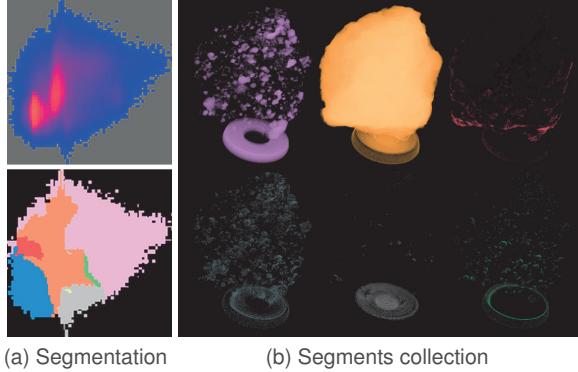


Fig. 15: Meteorite for case study 4. Showing collection of individual segments. The first segment in purple corresponds to the domain experts' expectation of the iron–nickel alloy. The rest are different parts of rock.

With our visualization system, the analysis started by removing the background segment. After changing the number of segments a few times with a slider, the user observed that the majority of this object is rock, and thus further segmentation is not necessary. When satisfied with

## 6.5 Expert Feedback

The four experts in nondestructive evaluation also provided qualitative feedback on the current prototype of our approach. Three experts were involved in the design of our approach and contributed valuable data, regular discussions, and testing of our system, and hence they are also included in the author list. Their expertise ranges from CT reconstruction, analysis, and evaluation, to experimental validation.

The fourth expert who had previously focused on finding appropriate methods for visualization of X-ray and neutron data was not involved in the design process of our approach. He provided valuable feedback, but is not included in the author list. For these four experimental sessions, two experts were able to directly use the developed system; one tested the system remotely via screen sharing, and the other was shown the results directly and provided feedback. A third expert was able to use the manually painted segmentation image afterwards to extract surfaces for further analysis.

We have received generally positive feedback from our target users and grouped them into four categories, as

discussed below. After a short demo from us, the experts were able to understand the theoretical background of the segmentation for data interpretation. They were able to identify materials through simple view manipulation and segment color toggling.

#### 6.5.1 Efficient System for Material Identification

One aspect that was praised by all experts is the ability to easily select the material regions from the bivariate histogram without the need to draw the exact material boundaries themselves. One expert had previously worked on manual segmentation of the bivariate histogram and was very impressed with the ability to automatize this process. Not only does our system automatize the drawing of polygons, but the detected boundaries can also be much more detailed than manual drawing can achieve. According to him, this manual segmentation was the most laborious part of the previous analysis procedure and our automated segmentation saves around 20 minutes of manually placing polygons for each dataset. Furthermore, the experts appreciated being able to simply change the number of segments.

The bimodal visualization has also been reported to be helpful in highlighting an overview of all material boundaries, which is difficult to obtain with the traditional slicing method. However, for a detailed examination, the experts preferred to look only at smaller slabs, using clipping planes, to avoid visual occlusion issues. One expert compared our work to the fusion of X-rays and neutrons but preferred our semiautomated method to change the combined view and the number of segments. With this method, artifacts and errors in the acquisition as well as imperfect automated fusion algorithms can be adjusted for.

#### 6.5.2 Self-Explanatory Concept and Familiar View

The experts were very satisfied with the simplicity of the visualization and the incorporation of familiar views, such as the bivariate histogram and the intuitive histogram-to-volume color mapping. The visual result did not introduce any "black boxes", such as in the deep learning approach, that would decrease confidence in the interpretation of the results. Instead, we use well-established bivariate histograms and combine them with a segmentation that can be simply understood by the user (using hills and valleys). The experts immediately understood that the Morse-complex segmentation is related to the common watershed solution, but preferred our hierarchical method due to the ability to simply adjust the number of segments.

#### 6.5.3 Minimum Data Manipulation

A senior member of the nondestructive evaluation team appreciated the fact that our method did not modify the datasets as typical data analysis techniques do. The result from our automated segmentation is not only flexible enough to play around in real-time, but it is also not "corrected" to answer any specific question. Rather, the final combined view remains as a direct representation of the raw input. For example, the material closer to the center easily gets oversegmented due to X-ray decay. The experts commented that although it may take some thinking to see the reason behind this tendency, the result revealed

the underlying physics of how the beams behave for these cylindrical objects. Thus, the visualization system itself was considered a more faithful representation of the dataset, and thus is well suited for general purpose analysis.

#### 6.5.4 Further Expectations

The scientists have expressed an overall interest in using the system further with more datasets because they believe this tool can be very helpful in general. One of them stated he would like to see further integration of other data analysis toolkits to make full use of the visualization, for example, validating the boundaries that are shown in the view.

The main negative comments about the visualization come from two aspects, the quality of automated segmentation and the lack of shading. The final segmentation is not as accurate as a data-driven approach. However, experts did acknowledge that our histogram-based segmentation approach is not a 'black-box', and that the purpose of our system is not to replace specialized analysis tools but to give an overview. The ability to manually adjust the segmentation results enables them to account for inaccuracies caused by artefacts. The fact that we encourage the user to oversegment first is also unintuitive at first. In the end, the scientists commented that this is something they need to get used to, as the number of material peaks in the histogram does not correspond to the number of actual materials, but different density clusters in the data.

The complaints of no shading come from an expectation of creating underlying geometries. For instance, the scientists often have some knowledge of the object and are aware that there may be surface-like materials to be extracted for further analysis, which is beyond the topic of visual representation here. The segmented volumes are also not always clean enough to form crisp geometry structures like surfaces, and it would be time consuming to compute all the potential geometry shapes in 3D space with different number of segments. The process could be a one-time computation but will defeat the purpose of our method as being fast to use. The experts wanted surface-like shading with thinner materials of high opacity to enhance their spatial perception of such structures. It is possible to create a customized surface segment and extract the corresponding data points later on, but the quality remains unclear since the segment is manually defined in histogram space.

## 7 DISCUSSION AND LIMITATIONS

After the system testing sessions and discussions, the experts confirmed that this research is already showing great potential to characterize multimaterial industrial objects. Although our approach is also ready to be applied to X-ray and ultrasound data, one key challenge is that these two modalities must be registered first, and we seek to extend this work to a wider range of multivariate volume datasets. In the future, we also aim to explore the possibility to edit and correct the segmentation, allowing users full flexibility to create a segmentation that are ground truth quality.

Our visualization currently does not support integrated multiresolution data loading. The automated multiresolution loading process will open up the possibility of straight-

forward large-scale data analysis. The segmentation quality can also be improved by adopting ideas from higher dimensions and calculating high-dimension distances for more complicated underlying topological structures. Other distance measures (geodesic distance for example) and more advanced topological algorithms should also be taken into consideration.

Besides direct volume rendering, it will also be preferable to enable mesh geometry extraction and surface rendering. We would also like to dig further to adapt the properties of X-ray and neutron for precise material identification, for instance to favor the high-resolution X-ray before it hits any blocking material, and bias toward the neutron afterwards. Such physics-based adjustments can be applied to aid further data analysis, and with further modification, the system can be easily adapted to other types of bimodal medical data.

## 8 CONCLUSION

With the ability to rapidly inspect a joint view of X-ray and neutron data, our collaborating expert scientists can now quickly investigate the material composition of industrial objects. This additional visualization capability is essential for the development of advanced material characterization in nondestructive evaluation. We have successfully developed a bimodal visualization system that provides an efficient overview of both modalities and is easy to navigate. By incorporating the well-known bivariate histogram, we ensure that the proposed approach can be easily understood by expert scientists, which is confirmed by expert feedback.

## ACKNOWLEDGMENTS

The authors wish to thank the reviewers as well as Jacob LaManna at NIST, and Allan Treiman at the Lunar and Planetary Institute for providing the sample datasets and the helpful feedback. NSF OAC awards 2127548, 1941085, 2138811 NSF CMMI awards 1629660, DoE award DE- FE0031880, and the Intel Graphics and Visualization Institute of XeLLENCE, and oneAPI Center of Excellence.

## REFERENCES

- [1] J. Banhart, *Advanced Tomographic Methods in Materials Research and Engineering*. Oxford University Press, 03 2008. [Online]. Available: <https://doi.org/10.1093/acprof:oso/9780199213245.001.0001>
- [2] A. Mascarenhas, R. W. Grout, P.-T. Bremer, E. R. Hawkes, V. Pas- cucci, and J. Chen, *Topological feature extraction for comparison of terascale combustion simulation data*, ser. Mathematics and Visualization. Springer, 2011, pp. 229–240.
- [3] P.-T. Bremer, A. Gruber, J. Bennett, A. Gyulassy, H. Kolla, J. Chen, and R. Grout, “Identifying turbulent structures through topological segmentation,” *Com. in App. Math. and Comp. Sci.*, vol. 11, no. 1, pp. 37–53, 2016.
- [4] I. Wald, G. Johnson, J. Amstutz, C. Brownlee, A. Knoll, J. Jeffers, J. Gunther, and P. Navratil, “Ospray – a cpu ray tracing framework for scientific visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, pp. 1–1, 01 2016.
- [5] K. Lawonn, N. Smit, K. Bühler, and B. Preim, “A survey on multimodal medical data visualization,” *Computer Graphics Forum*, vol. 37, 10 2017.
- [6] C. Correa, P. Lindstrom, and P.-T. Bremer, “Topological spines: A structure-preserving visual representation of scalar fields,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, pp. 1842–51, 12 2011.
- [7] J. Jankowai and I. Hotz, “Feature level-sets: Generalizing iso-surfaces to multi-variate data,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 2, pp. 1308–1319, 2020.
- [8] J. LaManna, J. Chen, S. Althaus, Y. Liu, D. Hussey, and D. Jacobson, “Bivariate histogram segmentation of simultaneous neutron and x-ray tomography for improved compositional and structural determination of source rock shales,” *Microscopy and Microanalysis*, vol. 26, pp. 1–3, 07 2020.
- [9] J. Kniss, G. Kindlmann, and C. Hansen, “Multidimensional transfer functions for volume rendering,” *Visualization and Computer Graphics, IEEE Transactions on*, vol. 8, pp. 270–285, 08 2002.
- [10] Y. Wang, J. Zhang, D. Lehmann, H. Theisel, and X. Chi, “Automating transfer function design with valley cell-based clustering of 2d density plots,” *Computer Graphics Forum (In Proc. EuroVis)*, 01 2012.
- [11] E. Fishman, D. Ney, D. Heath, F. Corl, K. Horton, and P. Johnson, “Volume rendering versus maximum intensity projection in ct angiography: What works best, when, and why?,” *Radiographics: a review publication of the Radiological Society of North America, Inc*, vol. 26, pp. 905–22, 05 2006.
- [12] C. Heinzl, J. Kastner, and E. Gröller, “Surface extraction from multi-material components for metrology using dual energy ct,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1520–1527, 2007.
- [13] D. Mueller, A. J. Maeder, and P. J. O’Shea, “The generalised image fusion toolkit (gift).” Kitware, Inc., July 2006, pp. 1–16, special issue: MICCAI 2006 Workshop on Open Science. [Online]. Available: <https://eprints.qut.edu.au/6454/>
- [14] W. Cai and G. Sakas, “Data intermixing and multi-volume rendering,” *Computer Graphics Forum*, vol. 18, 1999.
- [15] N. Schubert and I. Scholl, “Comparing gpu-based multi-volume ray casting techniques,” *Computer Science - Research and Development*, vol. 26, pp. 39–50, 02 2011.
- [16] C. Morris and D. Ebert, “Direct volume rendering of photographic volumes using multi-dimensional color-based transfer functions,” *Proceedings of the Symposium on Data Visualisation 2002*, pp. 115–ff, 01 2002.
- [17] J. Chuang, D. Weiskopf, and T. Möller, “Hue-preserving color blending,” *IEEE transactions on visualization and computer graphics*, vol. 15, pp. 1275–82, 11 2009.
- [18] A. Ghosh, P. Prabhu, A. Kaufman, and K. Mueller, “Hardware assisted multichannel volume rendering.” vol. 2003, 01 2003, pp. 2–7.
- [19] L. Cai, B. Nguyen, C.-K. Chui, and S. Ong, “A two-level clustering approach for multidimensional transfer function specification in volume visualization,” *The Visual Computer*, vol. 33, 02 2017.
- [20] A. Buja, J. A. McDonald, J. Michalak, and W. Stuetzle, “Interactive data visualization using focusing and linking,” *Proceeding Visualization ’91*, pp. 156–163, 1991.
- [21] P. McCormick, J. Inman, J. Ahrens, C. Hansen, and G. Roth, “Scout: a hardware-accelerated system for quantitatively driven visualization and analysis,” in *IEEE Visualization 2004*, 2004, pp. 171–178.
- [22] J. Woodring and H.-W. Shen, “Multi-variate, time varying, and comparative visualization with contextual cues,” *IEEE transactions on visualization and computer graphics*, vol. 12, pp. 909–16, 09 2006.
- [23] P. Rautek, S. Bruckner, and E. Gröller, “Semantic layers for illustrative volume rendering,” *IEEE transactions on visualization and computer graphics*, vol. 13, pp. 1336–43, 12 2007.
- [24] T. M. Athawale, C. R. Johnson, S. Sane, and D. Pugmire, “Fiber uncertainty visualization for bivariate data with parametric and nonparametric noise models,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 613–623, 2023.
- [25] Y. Nagai, Y. Ohtake, and H. Suzuki, “Segmo: Ct volume segmentation using a multi-level morse complex,” *Computer-Aided Design*, vol. 107, pp. 23–36, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010448518301945>
- [26] P. Sereda, A. Vilanova, and F. Gerritsen, “Automating transfer function design for volume rendering using hierarchical clustering of material boundaries.” vol. 2006, 01 2006, pp. 243–250.
- [27] H. Guo, H. Xiao, and X. Yuan, “Scalable multivariate volume visualization and analysis based on dimension projection and parallel coordinates,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 9, pp. 1397–1410, 2012.
- [28] K. Lu and H.-W. Shen, “Multivariate volumetric data analysis and visualization through bottom-up subspace exploration,” in *2017 IEEE Pacific Visualization Symposium (PacificVis)*, 2017, pp. 141–150.

- [29] F.-Y. Tzeng and K.-L. Ma, "A Cluster-Space Visual Interface for Arbitrary Dimensional Classification of Volume Data," in *Eurographics / IEEE VGTC Symposium on Visualization*, O. Deussen, C. Hansen, D. Keim, and D. Saupe, Eds. The Eurographics Association, 2004.
- [30] D. Keim, "Information visualization and visual data mining," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 1–8, 2002.
- [31] C. Heinzl and S. Stappen, "Star: Visual computing in materials science," *Computer Graphics Forum*, vol. 36, pp. 647–666, 06 2017.
- [32] M. Schiwarth, J. Weissenböck, B. Plank, B. Fröhler, C. Heinzl, and J. Kastner, "Visual analysis of void and reinforcement characteristics in x-ray computed tomography dataset series of fiber-reinforced polymers," *IOP Conference Series: Materials Science and Engineering*, vol. 406, p. 012014, 09 2018.
- [33] A. Gall, E. Gröller, and C. Heinzl, "Imndt: Immersive workspace for the analysis of multidimensional material data from non-destructive testing," in *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology*, ser. VRST '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: <https://doi.org/10.1145/3489849.3489851>
- [34] K. M. Champlley, T. M. Willey, H. Kim, K. Bond, S. M. Glenn, J. A. Smith, J. S. Kallman, W. D. Brown, I. M. Seetho, L. Keene, S. G. Azevedo, L. D. McMichael, G. Overturf, and H. E. Martz, "Livermore tomography tools: Accurate, fast, and flexible software for tomographic science," *NDT & E International*, vol. 126, p. 102595, 2022.
- [35] P. Vontobel, E. H. Lehmann, R. Hassanein, and G. Frei, "Neutron tomography: Method and applications," *Physica B: Condensed Matter*, vol. 385–386, pp. 475–480, 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921452606011161>
- [36] S. Kumar, V. Pascucci, V. Vishwanath, P. Carns, M. Hereld, R. Latham, T. Peterka, M. E. Papka, and R. Ross, "Towards parallel access of multi-dimensional, multi-resolution scientific data," in *2010 5th Petascale Data Storage Workshop (PDSW '10)*, 2010, pp. 1–5.
- [37] "Visus wiki," [https://wiki.visus.org/index.php/Main\\_Page](https://wiki.visus.org/index.php/Main_Page), accessed: 2023-08-24.
- [38] N. Kotava and A. Knoll, "Morse–smale decomposition of multivariate transfer function space for separably-sampled volume rendering," *Computer Aided Geometric Design*, vol. 30, pp. 549–556, 07 2013.
- [39] E. Vazquez, R. Baldrich, J. Vazquez, and M. Vanrell, "Topological histogram reduction towards colour segmentation," in *Iberian Conference on Pattern Recognition and Image Analysis*. Springer, 2007, pp. 55–62.
- [40] V. Robins, P. J. Wood, and A. P. Sheppard, "Theory and algorithms for constructing discrete morse complexes from grayscale digital images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 8, pp. 1646–1658, 2011.
- [41] A. Gyulassy, P.-T. Bremer, and V. Pascucci, "Computing morse–smale complexes with accurate geometry," *IEEE transactions on visualization and computer graphics*, vol. 18, no. 12, pp. 2014–2022, 2012.
- [42] R. Forman, "Morse theory for cell complexes," *Advances in mathematics*, vol. 134, no. 1, pp. 90–145, 1998.
- [43] —, "A user's guide to discrete morse theory," *Sém. Lothar. Combin.*, vol. 48, 12 2001.
- [44] H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," in *Proceedings 41st annual symposium on foundations of computer science*. IEEE, 2000, pp. 454–463.
- [45] J. Reininghaus, N. Kotava, D. Guenther, J. Kasten, H. Hagen, and I. Hotz, "A scale space based persistence measure for critical points in 2d scalar fields," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2045–2052, 2011.
- [46] H. Carr, J. Snoeyink, and M. van de Panne, "Simplifying flexible isosurfaces using local geometric measures," in *IEEE Visualization '04*. IEEE Computer Society, 2004, pp. 497–504.

**Haichao Miao** is a computational scientist at the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory, United States. His research focuses on multifield, virtual reality and large-scale data visualization. He received his PhD in computer science from TU Wien, Austria, in 2019 and worked in the Molecular Diagnostics Department at the Austrian Institute of Technology, where he focused on the development of in silico methods for the design and visualization of nanostructures.

**Andrew Townsend** is an engineer in the Nondestructive Evaluation Group at the Lawrence Livermore National Laboratory (LLNL) in Livermore, CA. Current work includes X-ray radiography and computed tomography research, characterizing complex materials and structures. He was a quality manager in the aerospace industry for 20 years before transitioning to research. Dr. Townsend has authored over a dozen peer-reviewed journal papers. He holds a PhD in Mechanical Engineering from Huddersfield University, UK.

**Kyle Champlley**, PhD: Kyle joined LLNL in 2012 and is an applied mathematician who works in the Nondestructive Characterization Institute and is the lead for the Signal and Image Processing Research group. He develops algorithms and writes software for data processing and reconstruction of Computed Tomography (CT) data. Kyle is the primary developer for the Livermore Tomography Tools (LTT) software package. He currently serves as an independent contractor for two medical imaging startup companies, Imatrex and PET/X. Previously, he was a staff scientist at the General Electric Global Research center where he developed CT reconstruction algorithms for GE's Revolution CT system. Kyle received his PhD in Electrical Engineering from the University of Washington where he performed research in Positron Emission Tomography (PET).

**Joseph Tringe** leads a dynamic group of more than 40 engineers and technicians in the Lawrence Livermore National Laboratory (LLNL) Materials Engineering Division, focused on characterizing a wide range of complex materials, structures, devices and systems. Techniques of interest to the nondestructive evaluation (NDE) group include x-ray computed tomography (CT), laser and transducer-based ultrasound characterization, microwave interferometry and microwave interrogation of embedded sensors. Dr. Tringe has led the group since 2017. From 2003 to 2017, Dr. Tringe was a staff scientist in the Physics and Life Sciences Directorate where he synthesized and characterized devices and materials with novel electrical, optical and molecular transport properties. He also developed advanced methods for characterizing phenomena associated with energetic materials safety. Prior to joining LLNL in 2003, he led a group in the Air Force Research Laboratory (AFRL) Space Vehicles Directorate which created and characterized radiation- and defect-tolerant electronic materials, devices and circuits for satellites. Dr. Tringe has authored more than thirty-five peer-reviewed publications and five U.S. patents. He holds a PhD in Materials Engineering from Stanford University and a bachelor's degree in Physics from Harvard College.

**Valerio Pascucci** is the founding Director of the Center for Extreme Data Management Analysis and Visualization (CEDMAV) of the University of Utah. Valerio is also a Faculty of the Scientific Computing and Imaging Institute, a Professor of the School of Computing, University of Utah, and a Laboratory Fellow, of PNNL. Before joining the University of Utah, Valerio was the Data Analysis Group Leader of the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory, and Adjunct Professor of Computer Science at the University of California Davis. Valerio's research interests include Big Data management and analytics, progressive multi-resolution techniques in scientific visualization, discrete topology, geometric compression, computer graphics, computational geometry, geometric programming, and solid modeling. Valerio is the coauthor of more than one hundred refereed journal and conference papers and has been an Associate Editor of the IEEE Transactions on Visualization and Computer Graphics.

**Peer-Timo Bremer** holds a shared appointment at Lawrence Livermore National Laboratory's (LLNL's) Center for Applied Scientific Computing (CASC), focusing on large-scale data analysis and visualization, and at the University of Utah, serving as Associate Director for Research of the Center for Extreme Data Management Analysis and Visualization (CEDMAV). His research interests include large-scale machine learning, data analysis, visualization, medical image analysis, topology, volume modeling, and virtual reality.