

Slice and Dice: A Physicalization Workflow for Anatomical Edutainment

Renata G. Raidou,^{1,2} M. Eduard Gröller,^{2,3} Hsiang-Yun Wu²

¹ University of Groningen, the Netherlands, ² TU Wien, Austria, ³ VRVis Research Center, Austria

Abstract

During the last decades, anatomy has become an interesting topic in education—even for laymen or schoolchildren. As medical imaging techniques become increasingly sophisticated, virtual anatomical education applications have emerged. Still, anatomical models are often preferred, as they facilitate 3D localization of anatomical structures. Recently, data physicalizations (i.e., physical visualizations) have proven to be effective and engaging—sometimes, even more than their virtual counterparts. So far, medical data physicalizations involve mainly 3D printing, which is still expensive and cumbersome. We investigate alternative forms of physicalizations, which use readily available technologies (home printers) and inexpensive materials (paper or semi-transparent films) to generate crafts for anatomical edutainment. To the best of our knowledge, this is the first computer-generated crafting approach within an anatomical edutainment context. Our approach follows a cost-effective, simple, and easy-to-employ workflow, resulting in assemblable data sculptures (i.e., semi-transparent sliceforms). It primarily supports volumetric data (such as CT or MRI), but mesh data can also be imported. An octree slices the imported volume and an optimization step simplifies the slice configuration, proposing the optimal order for easy assembly. A packing algorithm places the resulting slices with their labels, annotations, and assembly instructions on a paper or transparent film of user-selected size, to be printed, assembled into a sliceform, and explored. We conducted two user studies to assess our approach, demonstrating that it is an initial positive step towards the successful creation of interactive and engaging anatomical physicalizations.

CCS Concepts

• **Human-centered computing** → **Visualization application domains**; • **Applied computing** → **Life and medical sciences**;

1. Introduction

Patient education and involvement in treatment have become standard practice in many places around the world [GRF*10]. At the same time, there is increasing public interest in learning anatomy and physiology. For laymen without medical training or for a schoolchild, understanding the available information and recognizing structures within medical images, such as CT or MRI, would be overwhelming. Hence, anatomical illustrations (2D or 3D, in books or on-screen, static or interactive) are often employed. Sometimes anatomical models are preferred, as they facilitate 3D localization of structures within the human anatomy.

Hand-drawn anatomical illustrations appeared as early as 1522 [Cho52]. Figure 1(a) presents a well-known illustration of the human muscular system by A. Vesalius, drawn in 1543. People have been looking for alternatives to these 2D illustrations, for example by producing anatomical wax sculptures, such as the Anatomical Venus, shown in Figure 1(b). These 3D sculptures provided a reusable and durable anatomical education and training alternative to dissections. Later, the invention of medical imaging revolutionized the domain of anatomical education [PB14]. Nowadays, as medical imaging techniques become increasingly sophisticated, many anatomical education applications have emerged [BCK*11, SHK*16, HDK17, PS18]. These provide

the immediate advantage of bringing on-screen personalized 2D or 3D representations of patients, in a time- and cost-efficient manner.

Naturally, the use of physical representations of anatomical data has decreased. However, in recent years, the domain of data physicalization demonstrated that physical visualizations can be effective [Moe08, ZM08, JDI*15]—sometimes, even more than their virtual counterparts [JDF13]. This is due to some of their unique properties, which are not encountered in traditional, screen-based visualizations, such as the sense of scale, or full engagement of the user's perception and cognition. Medical applications within the domain of data physicalization revolve nowadays around the topic of 3D printing. Still, 3D printing is expensive and cumbersome for private use [RMVTK*10]. At the same time, common 3D prints do not provide a lot of interactivity, due to their inherently static nature. Medical data physicalizations that do not involve 3D printing, but are made of affordable and accessible materials, are still limited [SCS17, SB17, Fra08].

An additional consideration relates to the anatomical education of schoolchildren, where engagement is often sought by activity, such as playing, building, and crafting. Depending on the age, papercrafts, such as origamis, pop-ups, or sliceforms, are often used for educational purposes [EE98]. An example of a paper-based sliceform of a sphere is shown in Figure 1(c). Computer-assisted re-




Figure 1: (a) Illustration of the human muscular system, drawn in 1543 by A. Vesalius [Ves43]. (b) An Anatomical Venus sculpture, produced at La Specola in c. 1785. (c) A sliceform of a sphere.

construction of papercrafts has been investigated in the past, mainly for recreational purposes [LNLRL13, RJLYL14, PDRK18]. Crafting approaches have not been yet investigated for anatomical education, and they do not go beyond the use of other materials, except for paper (e.g., transparent).

The *contribution* of our work is the investigation of alternative forms of medical data physicalizations, which use readily available technologies and inexpensive materials. We make use of the advantages of data sculptures (in the shape of semi-transparent sliceform crafts) within a cost-effective, simple, and easy-to-employ workflow. Our *goal* is to support the creation of interactive and engaging anatomical physicalizations, to be used in laymen or children's education. The main *components* of our approach include:

- The support of both meshes and volumetric medical data.
- An octree formulation and optimization for sliceforms to support volume data, based on a user-defined transfer function.
- A packing algorithm to relatively preserve an optimal assembly order and the volume structure while minimizing the used space.

2. Related Work

Anatomical Education—Preim and Saalfeld [PS18] presented a survey on virtual anatomy education systems, focusing mainly on systems for medical students. A vast range of techniques is reviewed, spanning from surface and volume visualization [PHP*01, LSHL16, VGHN08] to animations [BPC*14], and from anatomical labeling [BG05] to virtual and augmented reality [SSPOJ16, PPS19, JPC*16, MWDE*16]. The most influential examples include the VOXEL-MAN [PHP*01], the open anatomy browser [HDK17], and the Online Anatomical Human [SHK*16]. Applications for the general public include the ZygoteBody [BCK*11] and the BioDigitalHuman [QSO*12]. Preim and Saalfeld discuss *constructivism* [HRL10], according to which active learning supports knowledge construction with less cognitive load, and *embodiment* [JVJB17], according to which learning can benefit from the involvement of motorics and physical interaction. This has been our starting point in the investigation of the suitability of data physicalization approaches for anatomical *edutainment* (i.e., education and entertainment).

(Medical) Data Physicalization—Data physicalization can be summarized as the process of mapping data to objects and their properties of the physical world [Moe08, JDI*15]. The generated

physical models aid or replace digital representations and allow data exploration with other senses, beyond the optical channel. The physicalization of medical data has revolved, so far, mainly around 3D printing [KNSV15, RMVTK*10, ASS*19]. More complex models are even possible, due to the recent advancements of 3D printing, such as multicolor and polymaterial printing. A recent example includes the work of Ang et al. [ASS*19], who developed a tangible physicalization of cardiac blood flow, to explore 4D MRI data in a slice-based physical model that complements traditional on-screen visualizations. Despite its high educational value, 3D printing remains time-consuming, while it is not affordable. It also involves a complex workflow, which starts with the selection of an anatomical area, the creation of its 3D geometry, the optimization of the anatomical geometry for printing, and the selection of adequate printing technology and suitable materials. This makes it non-trivial for the general population, while the manipulation of material properties to mimic tissues is still not possible.

Medical data physicalizations that do not involve 3D printing are limited. Except for the aforementioned wax models [MMŽ10], other materials, such as wood, ivory, cardboard, and fabric, have been employed [Olr00]. Notable are the papier-mâché anatomical models of Auzoux [Auz20], which could also be taken apart for exploration. Two more recent examples involve the creation of rearrangeable wooden models for the physical exploration of MRI brain scan cross-sections [Fra08] and the usage of volvelles, which are interactive wheel charts of concentric, rotating disks [SB17]. This physicalization mimics the on-screen fine-tuning of transfer functions to display meaningful volumetric data to non-knowledgeable users. Using paper, the production of volvelles is cost-effective and can be supported by home printers.

Sliceforms and Papercrafts—The topic of automatic creation of sliceforms and papercrafts has been tackled before, but not within the context of seeing through volumetric data or of an anatomical education application. For instance, there are automatic algorithms for generating stable, foldable, and physically plausible sliceforms, for designing and producing paper pop-ups from 3D meshes [RJLL13, RJLYL14, XZM*18]. Alternatively, origami architecture papercrafts [LLL*13], iris papercrafts [IIM16] or animated pop-ups that show motions of articulated characters [RLL15] are available. Such approaches have been investigated for large-scale models [SCS16, CS17], as well as for entire 3D scenes through developable surfaces [PDRK18], and the automatic generation of paper architecture [LSH*10]. None of these approaches looks into generating physical models from volumetric data (only from mesh data), while they only employ paper.

3. Requirements and Conceptual Choices

The basic concept behind *Slice and Dice* is to provide a workflow for the cost-effective, simple, and easy generation of engaging physicalizations, which can be used in anatomical edutainment. Our main motivation is that there is a lot of evidence that physical models can be effective in communicating data and processes [Moe08, ZM08, JDI*15]. They have also been found to be more engaging than their virtual counterparts [JDF13]. At the same time, we want to take advantage of the so-called “IKEA effect” [NMA12], according to which there is a tendency that people




Figure 2: The main steps of the *Slice and Dice* workflow.

value things more if they have created them themselves. To the best of our knowledge, such an approach does not exist in anatomical edutainment. Our main *requirements* were:

- (R1) The user has no (or little) knowledge of anatomy, and/or (physical) visualization.
- (R2) The workflow should accommodate anatomical meshes, and medical volumes.
- (R3) The user interaction with the workflow should be limited.
- (R4) The outcome of the workflow should be an easy-to-assemble and engaging physical model.
- (R5) The assembly of the physical model should be intuitive, time- and cost-efficient.
- (R6) The physical model should require easy-to-find and inexpensive materials, and common technology (e.g., home printers).

Based on these requirements, our *conceptual choices* were:

- (C1) **Supported data and functionality:** We consider the usage of both anatomical meshes and medical volumes. However, we do not focus on providing functionality for fine-tuning the visual properties of the employed surface or volume renderings, respectively. We limit ourselves to providing a number of anatomical mesh models, and a number of medical image volumes, with preset color and opacity (or transfer function) assignments.
- (C2) **Target physicalization:** We base our anatomical physicalizations on the generation of sliceforms. As we discussed in Section 2, sliceforms have been used before as a means of edutainment, but not as a means of edutainment. They are relatively popular and common among the general population, such as in greeting cards, 3D puzzles, and decorations. Sliceforms come also as a natural choice, given the usage of slice-based views for imaging data in medical practice.
- (C3) **Required materials and technologies:** For the creation of our anatomical physical models, we employ transparent films (acetate sheets), which are compatible with common printers, inexpensive, and easy-to-acquire at any stationery shop. Normal paper could also be used, but transparent films ensure visibility on the inner structures. An unintentional effect of the usage of transparent films is that the assembled sliceforms can have a pseudo-3D appearance.

4. The *Slice and Dice* Workflow

Figure 2 depicts the workflow of our *Slice and Dice* strategy. The workflow has been designed to primarily support volumetric data, but mesh data can also be employed. Meshes are converted into a volume (Figure 2(a)), before undergoing the remainder of the workflow. The volume is initially processed and rendered on-screen, employing preset transfer functions (Figure 2(b)). Basic interactions,

such as zooming, rotating, and panning, are possible. Subsequently, the data is partitioned into slices, with the use of an octree that takes into account the visibility of the different structures, as conveyed by the preset transfer function of the volume rendering (Figure 2(c)). The slice configuration is then simplified, to minimize the number of separate planes that are used for the final sliceform assembly—thus, minimizing the assembly effort and time (Figure 2(d)). The slices are characterized based on their configuration within the octree structure. These might be up-down slices (i.e., intersections), or cut-throughs (i.e., slices that are passing through one or more slices), or have a stopper (i.e., slices that “hang” from others). Then, the assembly order is optimized to ensure that the sliceform can be assembled. Finally, a packing algorithm is used to place the resulting slices, together with eventual annotations about the slice characterization and order onto the selected medium (Figure 2(e)). A home printer can be used, and after cutting, the slices are assembled into a sliceform to be explored (Figure 2(f)). All steps are discussed in detail, in the upcoming sections.

4.1. Preprocessing and Rendering

The first step is to convert the data into a format that is usable by our workflow. If mesh data are imported, these are converted into volumetric data, where each mesh is mapped to one intensity value. Therefore, if a complex anatomical model of several meshes is loaded, each is treated as a distinct component, and is mapped to one distinct intensity value. For the mesh data, we employ a traditional surface rendering for aesthetic reasons, and the converted volumetric data are used for computations only, from the next step on. Instead, for the volumetric data, we employ a GPU-based volume rendering. Basic interaction (zoom, pan, rotate) is supported.

The visual properties of this initial rendering are predetermined, and cannot be altered by the user. There are three main reasons for this choice. First, allowing the users to interact and change the transfer functions might be difficult, as they most probably have no knowledge of rendering, or of the underlying (medical) data. Second, volumetric medical imaging data comes in many different formats and standards, and designing a transfer function that can accommodate all possible imaging options is impossible [LKG*16]. Third, there is a vast literature on transfer function design [LKG*16], and strategies for visibility in rendering [VG05, BGKG06, HBH03], which cannot be all included. Hence, we consider the tuning of transfer functions out of scope for this work, and we provide a preset of volumetric data with predetermined transfer functions. For the mesh data, we assign opacities based on the location, where inner structures have higher opacity, and colors are assigned in a semi-random way so that each structure is discernible from the rest. However, in the future, it would be




Figure 3: An example of an octree-based spatial partitioning. (a) The transfer function of the volume. (b) The input volume. (c) Four axial slices (top view). (d) Two-level octree partitioning. (e) Unification of small adjacent slices.

interesting to investigate optimization strategies for the employed transfer functions in the physicalizations, as well as the effect of their superimposition (and order) within the physical model.

4.2. Slice Retrieval

The next step in our workflow is to generate representative slices from the available volumetric data. Lattices are used to describe strips of materials crossed and fastened together. Mathematically, lattices have interesting properties that can be used to investigate and formalize the structure of objects, in two or three dimensions. For this reason, lattices are used in sliceform papercrafts, to better represent the features of a target object. In contrast to classical sliceform designs, which try to maximally preserve the features of the outer surfaces, in this paper, we aim to find a lattice partition that better preserves the inner structures of our target volume. In computer graphics, an octree is a space partitioning technique for fast searching neighbors, using a tree data structure. By recursively subdividing the input 3D space into eight octants, an octree allows to hierarchically manage sparse and coarse information, in the 3D space. In this paper, we are inspired by the concept of an octree, and introduce an *octree formulation for sliceform papercrafts of volumetric data*. After the octree-based spatial partitioning, we optimize the slice configuration by *unifying adjacent slices* to form bigger ones. This is done to facilitate, and speed up the assembly.

Octree-Based Volume Partitioning—Classical octrees are built for searching in a fast way which objects are nearby in a 3D space; for example, points of a mesh that are nearby in a 3D scene [Bru92]. The algorithm begins with a root node, and subdivides iteratively each tree node into eight children (octants)—if and only if, dense information is stored within the node region. In our approach, we replace this conditional density statement by the total number of quantized voxel types within the corresponding tree node. This replacement is done, because each quantized value in the transfer function reflects a distinct structure in a medical dataset, which we consider an interesting sub-region that requires more slices. The “importance” of the quantized voxels can be retrieved from the transfer function (Figure 3(a)). To determine which structures are more “important”, the scalar value of a voxel within the volume data is multiplied by its visibility factor (e.g., its opacity), as given by the transfer function. For example, if the transfer function of a CT volume dictates that all intensity values of 1000 have 0.5 opacity, the scalar value 1000 will be scaled by 0.5. Therefore, values that are more opaque are “prioritized” in our approach, and transparent values (i.e., hidden structures) are not considered. Fig-

ure 3(b,c) shows an octree partitioning example, where three quantized values are included in the volume. Gray color is considered as the background (opacity = 0) and is ignored in our calculation. Red and yellow voxels indicate different structures in the dataset (opacity = 0.5 and 1, respectively). In Figure 3(d), our octree algorithm subdivides the entire volume (Level 1), and recursively subdivides the sub-volumes until no different features exist, or at the finest voxel resolution (Level 2, in this example). We also introduce a parameter L that allows users to specify the highest tree level, in order to control the complexity of the final papercraft model. However, with an octree, we may generate several small slices adjacent to each other, which complicate the assembly process. Next, we explain how we integrate those slices into bigger, for easier assembly.

Slice Generation and Unification—In our papercraft design, we take two perpendicular slices, starting from the center of each tree node; for instance, one coronal (yz plane) and one sagittal (xz plane) slice (Figure 3(d)). This gives our papercraft flexibility to fold the 3D model into 2D, for better portability. We exclude the bounding planes of each tree node, to avoid having slices on the same 2D plane—especially, if the tree nodes belong to different tree levels. With this choice, when two slices share a boundary edge, they might be still distinct from each other. Figure 3(e) shows such an example. If we consider the coronal slice x_3 in Figure 3, the upper half is subdivided twice, while the lower half is subdivided once. With an octree, we will receive two separate small slices: two subdivided half-red and half-yellow slices. For these adjacent slices, the assembly will be tedious. We, therefore, combine these slices into one, by computing their minimum and maximum coordinates to synthesize their unification. Note that this is done by first searching the connected slices on the same 2D plane, and then performing the unification process, to find the appropriate number of rectangular slices that covers the entire extent on the same 2D plane.

4.3. Papercraft Realization

The quality of a papercraft does not only rely on the feature representation of the input data. It also depends on the feasibility of its realization, based on slice intersections and order of assembly. To support the realization of the sliceform assembly, we first *categorize all pairwise slice intersections*, as hinge types within the papercraft. Then, we *determine a reasonable hinge order*, to guide users through easier assembly.

Hinge Characterization and Construction—To realize the sliceform papercraft, we first need to find a way to stitch (i.e., intersect) two slices together. As introduced by Le-

Nguyen et al. [LNLRL13], each line segment in the intersection of two slices is called a *hinge*. Two types of hinges are commonly used in a sliceform model: *up-down intersections* and *cut-through hinges*. For *up-down* hinges, we can cut the upper hinge segment of the coronal slices and the lower hinge segment of the sagittal slices, and stitch the two slices by moving them vertically along the hinges (Figure 4(a)). The cut-away upper or lower hinge segment is called a *slot*. For *cut-through* hinges, the cut-away hinge segment is located on the bigger slice along the *z*-axis, and no cutting will be performed on the smaller one (Figure 4(b)). Its corresponding slot is called a *none* slot, since no cut will be performed, and it will be inserted through the *cut-through* slot.

Assembly Order Optimization—Our assembly order concept stems from common practices in 3D anatomy puzzle games, where small, inner structures need to be assembled, before large, outer ones. However, the feasibility of a sliceform depends on the condition that there are no interlocks in the assembly order, to facilitate the reconstruction. The idea of an outwards assembly is less rigorous (up to a certain degree) since it primarily influences user experience in assembly. The feasibility of reconstruction, though, needs to be formulated by hard constraints; otherwise, the physical model cannot be assembled. These observations can be summarized as:

- (O1) Order Inequality:** The order of each hinge is distinct, since we assume that hinges are stitched sequentially.
- (O2) Cut-Through Hinge Dependency:** Given a slice with a *none* slot between two *up-down* hinges, the *cut-through* hinge should be assembled earlier than the *up-down* ones.
- (O3) Backbone Hinge:** The coronal and sagittal slices of the root node in the octree form the backbone hinge. This is the longest hinge with the largest slices. The backbone hinge is expected to be stitched first, to maintain the papercraft stability.
- (O4) Radial Dependency:** The assembly of hinges on a slice should be composed starting from the center.

These statements led us to consider the assembly order problem as a constrained permutation problem, where each element in the set should be distinct, and certain partial orders in the set should be preserved. We formulate this problem as an integer programming problem, where **(O1)**–**(O3)** are hard constraints and **(O4)** is a soft constraint. The reason for choosing integer programming rather than another optimization technique is that it allows us to separate hard and soft constraints. It is also capable of finding the optimal solution, instead of a feasible solution, if the scale of the problem is relatively small. Our input in the formulation are the indices of the hinges $H = \{1, 2, \dots, n\}$, and the output will be a sequential order of

hinge indices. Mathematically, the **(O1) Order Inequality** is equal to $x_i < x_j$ or $x_i > x_j$, and can be formulated as:

$$\forall i, j \in H \text{ and } i \neq j, \quad \begin{aligned} x_i + 1 &\leq x_j + M \cdot \alpha_{ij} \\ x_j + 1 &\leq x_i + M \cdot (1 - \alpha_{ij}), \end{aligned} \quad (1)$$

where x_i and x_j correspond to integer hinge indices. α_{ij} is a binary variable for validating one of the above equations in the solution space. It is applied together with a large constant M . M in our setting should be at least larger than the total number of hinges. The **(O2) Cut-Through Hinge Dependency** is formulated as:

$$\forall i, j, k \in C \text{ and } i \neq j \neq k, \quad \begin{aligned} x_j + 1 &\leq x_i \\ x_j + 1 &\leq x_k, \end{aligned} \quad (2)$$

where C collects all triples (i, j, k) where i, k are *up-down* hinges, and j is a *cut-through* hinge, and hinge j is adjacent to hinges i and k on a slice. The incorporation of **(O3) Backbone Hinges** is intuitive and can be formulated as $x_0 = 0$, since we force to place the backbone hinge x_0 , in the first place. As for the soft constraint **(O4) Radial Dependency**, we formulate it as the objective function in our optimization:

$$objective_{O4} = \sum_{x_i \in H} w_{distance}(i) \cdot x_i, \quad (3)$$

where $w_{distance}(i)$ denotes the Euclidean distance of a hinge to the center of the volume. Once we retrieve the order of hinges, we use this order to guide the users, to assemble the papercraft. Two instructions can be provided. One is the order of hinges, and the other is the order of slices, which follows the order of hinges.

4.4. Paper Placement

Now that we have the optimal ordering for the assembly of the sliceform, we need to put the distinct slices onto a piece of transparent film of a user-selected format (A3 or A4). For this, we need to take into account three considerations: First, slices that are “nearby” in the 3D volume space (and, subsequently, in the octree) need to be “nearby” on the printable configuration **(P1)**. This is to support the intuitive assembly of the sliceform. Second, the optimized order of assembly needs to be respected **(P2)**. This is to support the easy assembly of the sliceform, given that there might also be cut-through slices (with or without a stopper), which are more complex to put together. Third, we have to make the best possible use of the given materials (i.e., optimal use of the size of the transparent film) **(P3)**. This is to guarantee that the object is big enough to be explored, but also fits within the available sizes.

To this end, we construct for each slice a 4D vector $\vec{v} = (x, y, z, d)$ that contains the normalized coordinates in the 3D space of the center (x, y, z) of the slice **(P1)** and the order d **(P2)** of the assembly. Normalization ensures that all variables are weighted the same within the vector, which is important if our approach is based on Euclidean distances. For example, if we have five slices ($d \in [0, 4]$) within an \mathbb{R}^3 space $[-100, 100]^3$, without normalization, the coordinates (x, y, z) and the order d will be given unequal weights depending on their variance. Principal Component Analysis (PCA) [Pea01] is done in order to find if any variables of the vector \vec{v} might be correlated, to get rid of redundancy and to identify the two components that maximize variance. Then, a k -means




Figure 4: Types of hinges: (a) an up-down (intersection) hinge, and cut-through hinges (b) without, and (c) with a stopper.

clustering [Llo82] partitions the slices into k clusters. Each slice belongs to the cluster with the nearest mean. For calculating the optimal number of clusters, we employ the elbow method [Tho53]. The clustering gives an indication of slices that are located “nearby” in the 3D space (P1), and in order of assembly (P2).

The users can select the paper format, and whether to use one or more sheets. To make optimal use of the available printing material (P3), we apply the following packing approach. First, a kd -tree [Ben75] is applied to partition the available paper space based on the results of the previous clustering, while taking into account the sizes of each of the slices. If a cluster contains more (or bigger) slices, its respective paper partition will be bigger to fit all slices. The sizes of the partitions are constrained so that all partitions fit within the available paper format. This step results in a number of partitions equal to the number of clusters and each of these partitions contains the slices of the respective cluster, which are uniformly scaled to preserve their proportions. The next step involves taking a rectangle bin packing algorithm [Jyl10] to determine how the slices should be packed into the partitions, as well as their optimal sizes, to cover as much as possible space in the available paper format. Rotations of slices within the bins are possible. We use a Maximal Rectangles algorithm with the best short side fit [Jyl10]. The algorithm stores the slices into the partition that they belong in such a way that the length of the shortest leftover side of the partition is minimized. Thus, $\min(w_f - w, h_f - h)$ is the smallest, where w_f and h_f are the partition dimensions, and w and h are the slices dimensions. Once the packing is completed, the partitions are assembled into the selected size of the paper and the number of sheets.

4.5. Papercraft Stability

There are several factors that influence the stability of the final papercraft. Some essential factors include (S1) the weight balance of the octree partition, (S2) unexpected easy-to-open hinges, and (S3) the materials used for the papercraft. In rare cases, if the features of a volume are only located at one of the top corners, which requires a large number of slices for appropriate representation, the weight balance of the papercraft can be skewed. This did not happen in our experiments, since we normalized the volume coordinates, after reading the data. Nevertheless, one can sum up the gravitational torques of slices to test if the current model is weight-balanced, or add additional “empty” slices at the bottom, to compensate for this issue. Since our algorithm is based on octree partitioning, we may encounter a case, where a *none* slot is located exactly at the boundary of the slice. For this, we introduce a stopper, as shown in Figure 4(c). In our experiments, we tested our papercraft using copy papers, cardstock papers, and acetate sheets (transparent films). All paper types worked, but the slots should have a minimum width (e.g., 1 mm) to facilitate the sliceform assembly.

4.6. Print and Assembly

For the printing, the user can choose between A3, A4, and a user-determined size, and the partitions can be printed on one or multiple papers, depending on how big the physical model should be. The assembly printout contains labels and annotations for the slices, denoting the order in which the slices should be assembled, slots

for the up-down and cut-through hinges, as well as stoppers. We include also a rendered 3D model of the data, annotated with the octree cuts, and the order of slices, to facilitate the papercraft assembly. For the cutting, a cutting plotter can be used, but simple scissors, or a utility knife, are sufficient, and more accessible.

5. Results with Mesh Datasets

Nested Spheres Dataset—We created a synthetic dataset consisting of four sphere meshes, nested in the 3D spatial configuration, shown in Figure 5(a). The octree computation with $L = 3$ yields twelve slices, the assembly order of which, and the single A3 paper placement is shown in Figure 5(b). The resulting sliceform is shown in Figure 5(c). This example, despite being simple in terms of depicted structures, is quite complex, as it has six cut-throughs for the deeper-nested structures. The configuration of the nested spheres is identifiable, and the inner red sphere is visible.

Head Dataset—We obtained from the BodyParts3D database [MFT*09] 87 meshes comprising the anatomy of the head, and including the skull bones and brain components, as shown in Figure 5(d). The octree computation with $L = 2$ yields six slices, and with $L = 3$ yields fourteen slices. The order of the highest level, and its multi-A3 paper placement is shown in Figure 5(e). The resulting sliceforms are shown in Figure 5(f), for both octree levels. This example, despite consisting of many meshes, is quite easy to assemble. The $L = 2$ sliceform is easier, as it consists of only intersections—thus, it is straightforward to assemble, even for smaller sized slices. The $L = 3$ sliceform includes also two complex cut-throughs, which had to cross seven slices each. Therefore, it is advised to use bigger paper formats. The configuration of the different structures in the human head can be clearly seen, although in the $L = 2$ sliceform, the structures might be easier to identify.

Heart Dataset—We obtained from the BodyParts3D database [MFT*09] 125 meshes comprising the anatomy of the heart, as shown in Figure 5(g). The octree computation with $L = 2$ yields six slices, and with $L = 3$ yields fourteen slices. The order of the lowest level, and its single A3 paper placement is shown in Figure 5(h). The resulting sliceform is shown in Figure 5(i), for $L = 2$. This example consists only of intersecting slices, therefore, it is quite easy to assemble, for both levels. For $L = 3$, it is advised to use bigger paper formats or multiple pages. The configuration of the different structures in the human heart is clearly represented, and the heart chambers are identifiable.

Stuffed Bunny Dataset—We created a synthetic dataset consisting of five meshes (Stanford bunny, and four hearts) in the 3D spatial configuration shown in Figure 6. This example includes meshes, more topologically complex than the nested spheres. In Figure 6, we showcase the effect of adjusting the parameter L from 2 to 4, in the octree partitioning. Increasing the value of L creates more slices in the denser sub-region, to reveal higher details of the structures.

6. Results with Volumetric Datasets

Aneurysm Dataset—Our first volumetric example comes from a CT dataset of the lower torso of a patient with an aneurysm,

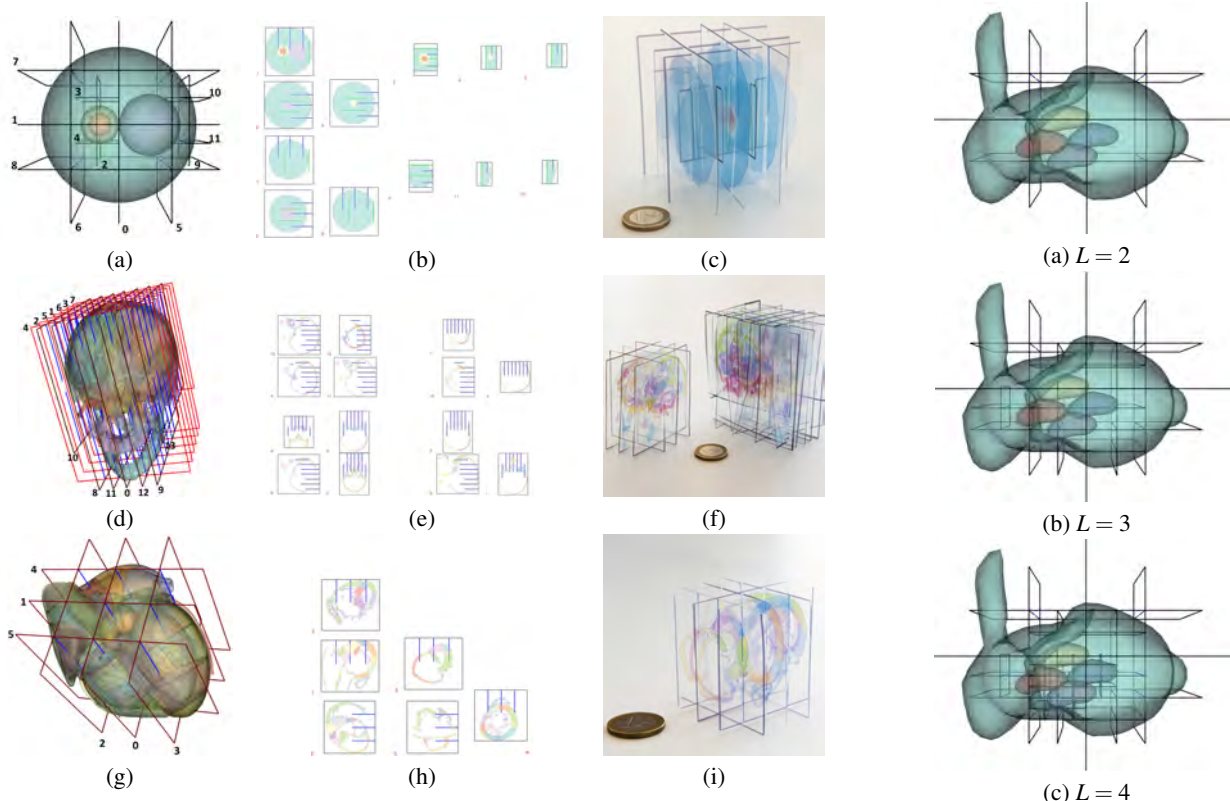


Figure 5: Results for three mesh datasets. In (a)–(c), we show a synthetic dataset, with a configuration of four nested spheres. In (d)–(f), we present a head mesh dataset, and in (g)–(i) a heart mesh dataset. The head and heart meshes are from the BodyParts3D database [MFT*09].

Figure 6: Results for different levels (L) of octree partitioning, for an artificial configuration of five meshes (top view).

as shown in Figure 7(a). The spatial resolution of the dataset is $256 \times 257 \times 119$ voxels, with a spacing of $0.74 \times 0.74 \times 1.5 \text{ mm}^3$. The octree computation with $L = 2$ yields six slices, which correspond to coronal and sagittal planes of the dataset. The order of assembly, and its single A3 paper placement are shown in Figure 7(b). The resulting sliceform is shown in Figure 7(c). This example consists only of intersecting slices, therefore, it is quite easy to assemble, both in A4 and in A3 size. The pathology (aneurysm) of the patient is clearly visible, and the sliceform gives a pseudo-3D appearance on the data, while at the same time it shows them in a slice-based manner, similar to clinical practice. With an optimized transfer function, the view on the inner structure of the patient's body would also improve. Figure 8 shows this dataset with a different transfer function and at various levels ($L = 2$ to $L = 4$) of partitioning. As the octree level increases, additional slices provide a better representation of fine structures within the dataset, such as more details on the blood vessels, and on smaller bones.

Foot Dataset—Our second volumetric example comes from a CT dataset of a human foot, as shown in Figure 7(d). The spatial resolution of the dataset is $256 \times 256 \times 256$ voxels, with a spacing of $1 \times 1 \times 1 \text{ mm}^3$. The octree computation with $L = 3$ yields fourteen slices, which correspond to axial and sagittal planes of the dataset. The order of assembly, and its single A3 paper placement are shown in Figure 7(e). The resulting sliceform is shown in Figure 7(f). This example consists of ten intersecting slices, and four cut-throughs

without stoppers, and is more complex than the previous one. The sliceform gives a pseudo-3D appearance to the foot, similar to a hologram. The bones and other small structures are visible, and can be explored, but probably an optimized transfer function would be required for a better view on the volume.

Shoulder Dataset—The last volumetric example comes from a CT dataset of a patient with a broken shoulder, as shown in Figure 7(g). The spatial resolution of the dataset is $131 \times 94 \times 166$ voxels, with a spacing of $1 \times 1 \times 1 \text{ mm}^3$. The octree computation with $L = 3$ yields fourteen slices, which correspond to axial and coronal planes of the dataset. The order of assembly, and its multi-A3 paper placement are shown in Figure 7(h). The resulting sliceform is shown in Figure 7(i). This example consists only of intersecting slices, which are quite close to each other, therefore, it is quite difficult to assemble the sliceform, and also to discern clearly the bone structures.

7. Evaluating the Feasibility of Sliceforms

We conducted an evaluation with 10 participants, to determine the feasibility and time-efficiency of our proposed approach. Our group of participants was gender-balanced (five males, five females), and they were between 28 and 37 years old. At this initial stage, we decided not to include children, as our evaluation would require a different design than for adults. Two of our participants have normal vision, three were wearing lenses, and five were wear-

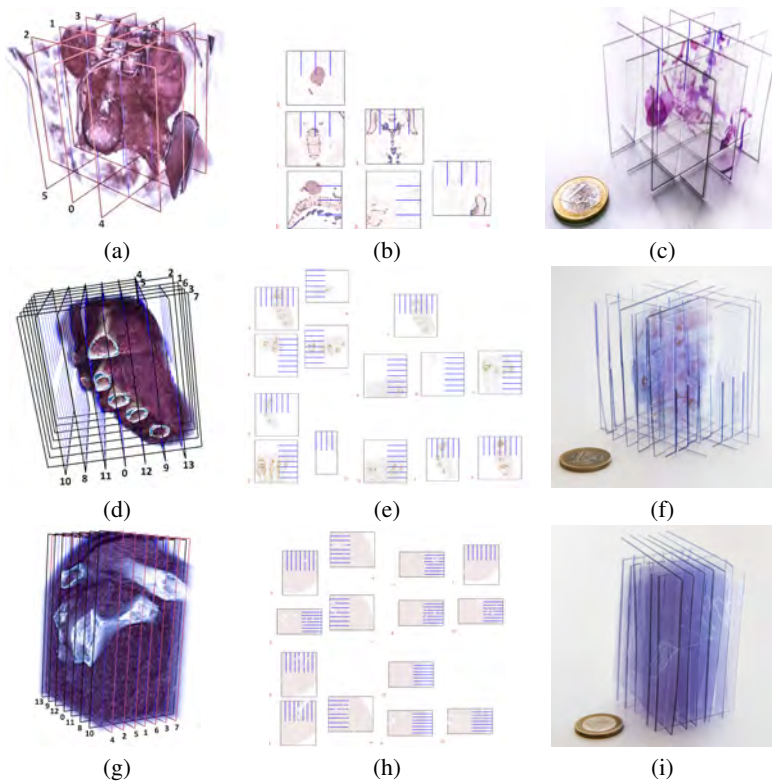


Figure 7: Results for three CT volumes. In (a)–(c), we show a dataset of an aneurysm, in (d)–(f) of a foot, and in (g)–(i) of a shoulder.

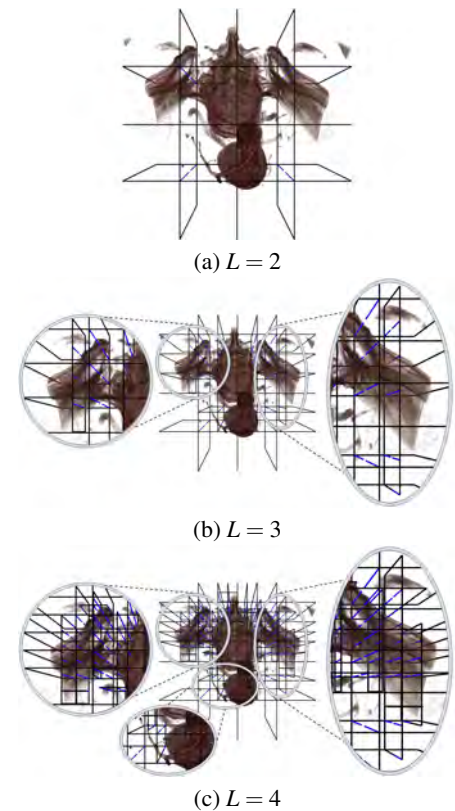


Figure 8: Results for different levels (L) of octree partitioning, for the aneurysm dataset (top view).

ing glasses. Six participants are from the visualization community (one in BioVis, one in BioMedVis, one in MedVis, and the rest in InfoVis or Visual Analytics), one is a software engineer, one is a software developer, one is an account manager, and one is a teacher. One of the participants is quite keen on paper crafting, and puzzles.

We designed our evaluation based on guidelines from the seven scenarios by Lam et al. [LBI*11]. We were interested in evaluating three aspects: (i) the assembly performance of the users, (ii) the communicative value of our proposed physicalizations, and (iii) the experience of the users, while constructing the physicalizations. To this end, we conducted a controlled experiment, where we asked our evaluation participants to assemble a set of sliceforms, while we measured the time for completion (*User Performance*). At the same time, we observed their behavior and interaction with the sliceforms, during and after assembly (*Communicative Value*). After the controlled experiment, we asked them open questions about their experience with the sliceform assembly, to obtain feedback (*User Experience*).

We prepared three cases, where slices were printed, and cut with the use of a cutting plotter, for precision. A *pop-out cutting* setting, which performs perforation cuts at slice boundaries and hinges, was selected to simulate the pre-cut sheets commonly used in papercraft books, to avoid that our users need to cut themselves, and to ensure the precision of the cuts. The participants had only to take out the pre-cut slices from within the paper frame, and to assemble them.

This time was not considered part of the experiment. Initially, we gave a short introduction and instructions on how the participants should interpret the annotations, and the labels on the printouts. We did not disclose to the participants the intended purpose of the sliceforms, but we showed them a printout of the 3D virtual counterpart of the dataset that they were building, and explained to them what the dataset is, and what it shows. The participants could not interact with the 3D representation, as we wanted them to interact only with the physical object, and to determine whether they actually needed the interactive 3D representation. We asked our participants to construct the first two (aneurysm and spheres datasets, shown in the previous sections), while the third one (head mesh dataset, shown in Section 5) was optional, as it has more slices, is more complex and would definitely require more time, and more experience. We did not randomize our datasets, as they were intended to be shown in order of increasing difficulty. We did not train the participants beforehand, to test the feasibility of our designed papercraft. Similarly, in a real-life scenario, they would start from an easy sliceform, and move on to harder cases. We measured the time for the completion of the assembly and upon completion, we stopped the time, inspected the sliceform, and informed them whether they had assembled it correctly, or not. If the assembly was wrong, we asked them to redo it, and measured the time again.

User Performance (UP)—For assessing UP, we measured the assembly completion time. All participants constructed the aneurysm and the spheres dataset successfully. Only six participants volun-

Table 1: Results of the user performance (UP) assessment, for evaluating the feasibility of the sliceforms.

	Gender	Age	Vision	Job	Dataset 1: Aneurysm CT			Dataset 2: Nested Spheres			Dataset 3: Head Mesh Data		
					Attempt 1	Attempt 2	Total attempts	Attempt 1	Attempt 2	Total attempts	Attempt 1	Attempt 2	Total attempts
A	F	33	lenses	Account manager	00:59	01:35	2	12:55	-	1	34:00	-	1
B	M	28	lenses	BioVis	01:41	-	1	21:25	-	1	-	-	-
C	F	28	normal	Software Developer	02:40	04:44	2	15:30	29:07	2	-	-	-
D	M	33	glasses	BioMedVis	01:16	02:31	2	18:59	-	1	-	-	-
E	M	36	lenses	Visual Analytics	01:15	-	1	15:00	-	1	-	-	-
F	M	33	normal	Software Engineer	01:48	03:18	2	08:17	-	1	23:33	-	1
G	M	29	glasses	Visual Analytics	01:29	-	1	14:34	-	1	25:15	-	1
H	F	33	glasses	MedVis	01:24	-	1	08:32	-	1	24:09	-	1
I	F	37	glasses	InfoVis	02:35	-	1	08:27	-	1	26:38	-	1
J	F	30	glasses	Teacher	02:14	-	1	10:51	-	1	29:29	-	1
					Time (mm:ss)		Attempts (N)	Time (mm:ss)		Attempts (N)	Time (mm:ss)		Attempts (N)
Average (μ)					02:17		1.40	14:45		1.10	27:11		1
Standard deviation (σ)					01:05		0.52	06:47		0.32	03:57		0

teered to do the third experiment, as they were motivated by the previous experiments. The completion times are summarized in Table 1, as well as the number of attempts. The average time ($mm:ss$) for the completion of the sliceforms was 02:17 ($\sigma=01:05$) for the aneurysm dataset, 14:45 ($\sigma=06:47$) for the spheres dataset and 27:11 ($\sigma=03:57$) for the head mesh dataset. For the completion of the aneurysm sliceform, six people finished it with one attempt, and the others needed a second attempt. For the spheres dataset, only one needed a second attempt, and for the head mesh dataset, all six participants completed it with one attempt. This could be due to learning, from the previous simpler cases. Although the fastest participant has a MedVis background, we did not observe any statistically significant difference in performance times between participants with a visualization (or Bio/MedVis) background, and the people that had no prior knowledge. Also, we did not observe any difference between men and women. All participants were able to recognize that they had committed an error in the assembly, and we only confirmed that, so that they could start the second attempt. We cannot quantify whether the second attempt needed more or less time, as participants committed errors at different steps of the assembly, and some needed to restart from scratch, while others from just half-way through.

Communication Value (CV)—For assessing CV, we observed how the participants were assembling, interacting, and exploring the sliceform, during the previously described controlled experiment. Some participants were using the order of assembly and others not, since we did not force them to follow the instructions, but kept this as a recommendation. The fastest participant used strictly the order for all datasets. The second fastest, though, used the order only for the last case. He did not use the order in the other two cases, because he thought that starting from the corners was easier than from the center, but in the interview, he commented that he should have used the order. The third, fourth, and fifth fastest participants strictly used the order, in all cases. Two of the participants never used the order, and they had the most difficulties and the longest completion times. We also noticed during this phase that most of the participants were really motivated to finish the sliceforms. Even if they found the assembly difficult, and were frustrated, none of them gave up. In some cases, if the sliceforms were too small, and if the slices were too close together, they had difficulties with the assembly, and were slower, but started improvising and using objects to push the slices into the appropriate position. They all always wanted to see the 3D model with the slice place-

ment, at all times. Most considered the aneurysm sliceform very easy, even those, who got it wrong. This is probably due to having no prior experience with sliceforms, or the underlying data. The spheres sliceform was considered tougher, but most of the people got it correct with the first attempt. Most of the participants, after building the sliceform, started looking at it from all possible sides, to see if they built it correctly, and see how it looked like. All were able to tell if they made a mistake by themselves. Most showed satisfaction with their successful assembly of the sculpture, and some wanted to keep it.

User Experience (UE)—After the experiment, we asked each participant a set of open questions. When asking whether sliceforms are helpful in explaining and communicating how the data looks in reality, and if they could extract useful information from building up, interacting, and observing the sliceforms, a participant (account manager) answered that “it gives a more 3D view on the data”, and another participant (visualization researcher) that “It helps you build a mental image of the slice configuration. It looks like you are interpolating the space mentally and build a 3D view in your head and imagine how the entire space looks like.” When asking about interaction and assembly experience, many of the participants commented that they wanted to finish it, even if it was difficult, and/or frustrating. This shows that sliceforms are actually engaging. About interaction after the assembly, one participant stated that “during the assembly, I was focused only on the building and did not look at the data, but at the end, I realized that I had gotten a totally new view on the data.” With regard to missing features, points for improvements, and current limitations, the participants mentioned that sometimes the cuts are difficult to access, and that a stiffer material could be easier to handle. Many participants noticed that the complexity increases fast with an increasing number of slices, and with more complex configurations, in case of cut-throughs. One important point for improvement is that the labels for the order should actually be printed on the slices, maybe at a corner, instead next to the slices, on the frame. One participant would have liked “a step-by-step IKEA-like instruction”. Two participants would have liked to have some visual indicators showing the three orientations (both visualization researchers). All participants thought that the assembly process was understandable, “but it needs practice and patience.” With regard to the material, the participants said that using transparent films is “even more useful”, because “when you see the intersections, it looks more spatial”, and because “with paper, you cannot see through it. You need the transparent film to make a

Table 2: Results of the educational value assessment.

	Gender	Age	Vision	Job	Correct Answers	Incorrect Answers	Doubting
A	F	33	lenses	Journalist	9	0	1
B	M	35	glasses	Programmer	8	0	2
C	M	39	normal	Business analyst	9	0	1
D	M	31	glasses	Programmer	6	2	2
E	F	26	normal	French teacher	5	3	2
F	M	29	normal	Photographer	8	0	2
G	F	30	lenses	Nurse	10	0	0
Average (μ)					7.86	0.57	1.43
Standard deviation (σ)					1.77	0.98	0.79

hologram out of it.” When asked how they would use sliceforms, people with kids commented that they would use it with them, “as a recreational game”, or “as an educational toy”. Others commented that it is “good to show the intrinsic structure of medical data”, “a great example for teaching 3D rendering, to explain it better and easier” (visualization researcher), “[good] for education or for explaining concepts, like an octree” (software engineer).

8. Evaluating the Educational Value of Sliceforms

We conducted an additional evaluation with 7 participants, to determine the educational value of our approach for the general public. Our group of participants included four males, and three females, between 26 and 39 years old. Similarly to the feasibility evaluation above, we did not include children. Three of our participants have normal vision, two were wearing lenses, and two were wearing glasses. The educational and occupational background of our participants is broad, and only one person has good knowledge of anatomy—being a nurse.

We conducted a controlled experiment with three cases (aneurysm, head, and heart sliceforms), where we asked our evaluation participants to identify several anatomical and/or pathological structures in pre-assembled sliceforms. These structures were shown to them using traditional, static, labeled anatomical illustrations, retrieved online, and printed prior to the evaluation. The purpose of this experiment was to evaluate whether anatomical structures can be identified and recognized in the sliceforms. During the experiment, we randomized the order of datasets and structures. We did not measure the time that was needed to identify a structure, only whether it was successfully pointed out (identified, not identified, and doubting: correct, but unsure answer). These measurements are summarized in Table 2. Most of the people could identify the anatomical structures correctly, or with some doubts. The nurse participant identified everything correctly. The missed structures were either the aneurysm (a pathology that is not realistically depicted in an illustration), or structures that are deeper in the volume, such as the corpus callosum in the brain, or structures that are more difficult to tell apart from others, such as the right vs. left atrium in the heart. In general, the orientation (left–right) was an issue for most participants. We did not observe any difference between men and women.

After the controlled experiment, we asked them open questions about their learning experience with the sliceforms. The first question was whether there is something that they could see better or easier, either in the sliceform, or in the illustration. Most of the

participants commented that they appreciate the three-dimensional view of the organs, as well as the ability to look at the structures from different viewpoints. One participant commented that it is “easier to see the shape and the size of things”. One person commented that “the colors confuse me a little bit in model”, and another one that “sometimes, I cannot see very deep into the model, but maybe I can take out a few slides”. The nurse commented that the sliceform would be “a nice toy for kids”, and that she cannot see something that she did not already know. Regarding the effect of the physical model on learning anatomy compared to a simple anatomical illustration, five participants commented that it is more fun, and more interesting, being interactive. Two commented that it was easy, but they would have appreciated labels (“If I would have the names of the structures on the model, I could learn a lot—and faster”). One participant commented that “if it would take me a long time to make [the sliceform], maybe I would still prefer to google [an illustration]”. When we asked whether the tool is helpful in explaining and communicating how anatomy looks in reality, all participants (except for the nurse) answered positively. The nurse commented that she might use it for explaining anatomy to a patient, while another participant commented that “it would be cool to know what these structures do”, in conjunction with an app. A participant commented: “It is a very simple tool to explain where things are. For example, I broke my metatarsal and could not tell much from the X-rays. With this, I would have understood”.

9. Discussion

Concerning our requirements, we evaluated the sliceforms with people without visualization and/or anatomical knowledge, and we found that they could all complete their sliceforms, and saw a value in them (R1). Our workflow is able to tackle both meshes and volumes (R2), as shown in the examples of Sections 5 and 6. We did not test the interaction of the user with the workflow, but it requires only three clicks (loading the data, slicing the volume, packing), which could also be reduced to one (loading the data) (R3). The outcome of the workflow is an engaging physical model, as resulting from the UE part of our evaluation, although it is not always easy to assemble (R4). This, however, relates more to the personal aptitude of each person. Some evaluation participants found the assembly easier than others, but most of them learned quickly how to assemble the sliceforms, and became faster. The assembly of the physical model was intuitive, when people were using the order of the assembly and the instructions, as we discussed in the UP and CV part of the evaluation. All participants could complete their sliceforms in reasonable times, and the cost of one sliceform is less than €1, while printing requires a home printer (R5, R6). An initial evaluation was conducted to assess sliceform in a learning context, but a more thorough (and potentially comparative) study with laymen has to be performed in the future, where the sliceforms can be tested against traditional methods of anatomical education, such as anatomical illustrations from textbooks, or on-screen 3D illustrations, and against anatomical education settings in VR [PS18].

Regarding the material, using transparent films provides a more insightful pseudo-3D view on the data than normal paper. With more slices, the view might become overloaded, especially if many different structures with different color assignments are present. The volume datasets look aesthetically more pleasing, but this

could be due to the fact that the mesh datasets consist of many different structures, with different visual properties. In this work, the selection and fine-tuning of the employed transfer function was considered to be out of scope. However, the selection of visual properties, the order of slices, and the superimposition of slices within the physical model plays a very big role in the visual outcome of the sliceforms. In the future, it should be more thoroughly investigated. This implies that transfer functions are not only a challenging research topic for on-screen visualizations [LKG*16]—but also, for physical visualizations. Additional constraints on the octree level-setting could prioritize user-specified structures. The size of the sliceforms, and the slice distances also influence the complexity of assembly and visibility. Small structures, such as blood vessels, might not be so visible in small or less dense constructs. The background also affects visibility, and a white background should be preferred. In the assembly order optimization, more sophisticated constraints could control the complexity of the papercraft, in addition to the computation of permutation performance. For the paper placement, we are now supporting only rectangular shapes, but sliceforms could “follow” the geometry of the depicted data, for a more organic appearance. Moreover, in the current approach, we allow users to define the camera position manually. An automatic approach, such as mutual information [OTWO17], could be introduced, to find the best viewpoint, before slicing.

10. Conclusions and Future Work

We present the first computer-generated crafting approach for anatomical edutainment, which uses accessible technologies (common printers) and inexpensive materials (paper or semi-transparent films) to generate cost-effective assemblable data sculptures (semi-transparent sliceforms), from anatomical meshes or volumetric medical imaging data. We introduce an octree formulation, and optimization for sliceform papercrafts, to support volume data based on a given transfer function, and we employ a packing that preserves the optimal assembly order and the structure of the volume, while minimizing the used resources. Among directions for future work, we prioritize the integration of slices that “follow” the geometry of the anatomical structures, where slices are not necessarily rectangular or equidistant, as well as the ability to fine-tune the visual properties of the data in a more flexible way. Visual indicators could also be integrated to make the assembly easier and more intuitive, while an additional evaluation to assess further the educational purpose of the sliceforms (not only the engagement and feasibility) would be required—also, compared to virtual applications. All in all, *Slice and Dice* is an initial positive step towards interactive and engaging physicalizations for anatomical edutainment.

References

- [ASS*19] ANG K. D., SAMAVATI F. F., SABOKROHIYEH S., GARCIA J., ELBAZ M. S.: Physicalizing cardiac blood flow data via 3D printing. *Computers & Graphics* 85 (2019), 42–54. 2
- [Auz20] AUZOUX L. T. J.: Papier-mâché anatomical models, 1820. URL: <http://dataphys.org/list/papier-mache-anatomical-models/>. 2
- [BCK*11] BLUME A., CHUN W., KOGAN D., KOKKEVIS V., WEBER N., PETERSON R. W., ZEIGER R.: Google body: 3D human anatomy in the browser. In *ACM SIGGRAPH 2011 Talks* (2011), p. 19. 1, 2
- [Ben75] BENTLEY J. L.: Multidimensional binary search trees used for associative searching. *Communications of the ACM* 18, 9 (1975), 509–517. 6
- [BG05] BRUCKNER S., GRÖLLER M. E.: Volumeshop: An interactive system for direct volume illustration. In *IEEE Visualization* (2005), 671–678. 2
- [BGKG06] BRUCKNER S., GRIMM S., KANITSAR A., GRÖLLER M. E.: Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization & Computer Graphics* 12, 6 (2006), 1559–1569. 3
- [BPC*14] BAUER A., PACLET F., CAHOUE V., DICKO A. H., PALOMBI O., FAURE F., TROCCAZ J.: Interactive visualization of muscle activity during limb movements: Towards enhanced anatomy learning. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2014), pp. 191–198. 2
- [Bru92] BRUNET P.: *3-D Structures for the Encoding of Geometry and Internal Properties*. 1992, pp. 159–188. 4
- [Cho52] CHOULANT L.: *Geschichte und Bibliographie der anatomischen Abbildung nach ihrer Beziehung auf anatomische Wissenschaft und bildende Kunst*. R. Weigel, 1852. 1
- [CS17] CHEN L., SASS L.: Generative computer-aided design: multi-modality large-scale direct physical production. *Computer-Aided Design and Applications* 14, 1 (2017), 83–94. 2
- [EE98] EISENBERG M., EISENBERG A.: Shop class for the next millennium: Education through computer-enriched handicrafts. *Journal of Interactive Media in Education* 1998, 2 (1998). 1
- [Fra08] FRASER N.: Rearrangeable wooden model of brain scan, 2008. URL: <http://dataphys.org/list/rearrangeable-wooden-model-of-3d-mri-data/>. 1, 2
- [GRF*10] GRUMAN J., ROVNER M. H., FRENCH M. E., JEFFRESS D., SOFAER S., SHALLER D., PRAGER D. J.: From patient education to patient engagement: implications for the field of patient education. *Patient education and counseling* 78, 3 (2010), 350–356. 1
- [HBH03] HADWIGER M., BERGER C., HAUSER H.: High-quality two-level volume rendering of segmented data sets on consumer graphics hardware. In *IEEE Visualization* (2003), pp. 301–308. 3
- [HDK17] HALLE M., DEMEUSY V., KIKINIS R.: The open anatomy browser: a collaborative web-based viewer for interoperable anatomy atlases. *Frontiers in neuroinformatics* 11 (2017), 22. 1, 2
- [HRL10] HUANG H.-M., RAUCH U., LIAW S.-S.: Investigating learners’ attitudes toward virtual reality learning environments: Based on a constructivist approach. *Computers & Education* 55, 3 (2010), 1171–1182. 2
- [IIM16] IGARASHI Y., IGARASHI T., MITANI J.: Computational design of iris folding patterns. *Computational Visual Media* 2, 4 (2016), 321–327. 2
- [JDF13] JANSEN Y., DRAGICEVIC P., FEKETE J.-D.: Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), pp. 2593–2602. 1, 2
- [JDI*15] JANSEN Y., DRAGICEVIC P., ISENBERG P., ALEXANDER J., KARNIK A., KILDAL J., SUBRAMANIAN S., HORNBAEK K.: Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), pp. 3227–3236. 1, 2
- [JPC*16] JOHN N. W., PHILLIPS N. I., CENYDD L. A., POP S. R., COOPE D., KAMALY-ASL I., DE SOUZA C., WATT S. J.: The use of stereoscopy in a neurosurgery training virtual environment. *Presence: Teleoperators and Virtual Environments* 24, 4 (2016), 289–298. 2
- [JVJB17] JANG S., VITALE J. M., JYUNG R. W., BLACK J. B.: Direct manipulation is better than passive viewing for learning anatomy in a three-dimensional virtual reality environment. *Computers & Education* 106 (2017), 150–165. 2

- [Jyl10] JYLÄNKI J.: A thousand ways to pack the bin—a practical approach to two-dimensional rectangle bin packing. **6**
- [KNSV15] K N H., SAMUEL B., VETTUKATTIL J.: Hybrid 3D printing: A game-changer in personalized cardiac medicine? *Expert Review of Cardiovascular Therapy* 13 (2015). **2**
- [LBI*11] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization & Computer Graphics* 18, 9 (2011), 1520–1536. **8**
- [LKG*16] LJUNG P., KRÜGER J., GRÖLLER E., HADWIGER M., HANSEN C. D., YNNERMAN A.: State of the art in transfer functions for direct volume rendering. In *Computer Graphics Forum* (2016), vol. 35, pp. 669–691. **3, 11**
- [LLL*13] LE S. N., LEOW S.-J., LE-NGUYEN T.-V., RUIZ C., LOW K.-L.: Surface and contour-preserving origamic architecture paper pop-ups. *IEEE Transactions on Visualization & Computer Graphics* 20, 2 (2013), 276–288. **2**
- [Llo82] LLOYD S.: Least squares quantization in pcm. *IEEE Transactions on Information Theory* 28, 2 (1982), 129–137. **6**
- [LNLRL13] LE-NGUYEN T.-V., LOW K.-L., RUIZ C., LE S. N.: Automatic paper sliceform design from 3D solid models. *IEEE Transactions on Visualization & Computer Graphics* 19, 11 (2013), 1795–1807. **2, 5**
- [LSH*10] LI X.-Y., SHEN C.-H., HUANG S.-S., JU T., HU S.-M.: PopUp: automatic paper architectures from 3d models. *ACM Transactions on Graphics* 29, 4 (2010), 111. **2**
- [LSHL16] LICHTENBERG N., SMIT N., HANSEN C., LAWONN K.: Sline: Seamless line illustration for interactive biomedical visualization. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2016), pp. 133–142. **2**
- [MFT*09] MITSUHASHI N., FUJIEDA K., TAMURA T., KAWAMOTO S., TAKAGI T., OKUBO K.: Bodyparts3d: 3D structure database for anatomical concepts. *Nucleic acids research* 37 (2009), D782–D785. **6, 7**
- [MMŽ10] MARKOVIĆ D., MARKOVIĆ ŽIVKOVIĆ B.: Development of anatomical models-chronology. *Acta Medica Medianae* 49, 2 (2010). **2**
- [Moe08] MOERE A. V.: Beyond the tyranny of the pixel: Exploring the physicality of information visualization. In *2008 12th International Conference Information Visualisation* (2008), pp. 469–474. **1, 2**
- [MWDE*16] MESSIER E., WILCOX J., DAWSON-ELLI A., DIAZ G., LINTE C. A.: An interactive 3D virtual anatomy puzzle for learning and simulation - initial demonstration and evaluation. *Studies in health technology and informatics* 220 (2016), 233–40. **2**
- [NMA12] NORTON M. I., MOCHON D., ARIELY D.: The IKEA effect: When labor leads to love. *Journal of consumer psychology* 22, 3 (2012), 453–460. **2**
- [Olr00] OLR Y.: Wax, wooden, ivory, cardboard, bronze, fabric, plaster, rubber and plastic anatomical models: Praiseworthy precursors of plastinated specimens. *Journal of the International Society for Plastination* 15, 1 (2000), 30–5. **2**
- [OTWO17] OHTAKA Y., TAKAHASHI S., WU H.-Y., OHTA N.: Using mutual information for exploring optimal light source placements. In *Smart Graphics* (2017), pp. 155–166. **11**
- [PB14] PREIM B., BOTHA C. P.: *Visual computing for medicine: theory, algorithms, and applications*. MK Publisher, Elsevier, 2014. **1**
- [PDRK18] PACZKOWSKI P., DORSEY J., RUSHMEIER H., KIM M. H.: PaperCraft3D: paper-based 3D modeling and scene fabrication. *IEEE Transactions on Visualization & Computer Graphics* 25, 4 (2018), 1717–1731. **2**
- [Pea01] PEARSON K.: On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2, 11 (1901), 559–572. **5**
- [PHP*01] POMMERT A., HÖHNE K. H., PFLESSER B., RICHTER E., RIEMER M., SCHIEMANN T., SCHUBERT R., SCHUMACHER U., TIEDE U.: Creating a high-resolution spatial/symbolic model of the inner organs based on the visible human. *Medical Image Analysis* 5, 3 (2001), 221–228. **2**
- [PPS19] POHLANDT D., PREIM B., SAALFELD P.: Supporting anatomy education with a 3D puzzle in a vr environment-results from a pilot study. In *Proceedings of Mensch und Computer 2019*. 2019, pp. 91–102. **2**
- [PS18] PREIM B., SAALFELD P.: A survey of virtual human anatomy education systems. *Computers & Graphics* 71 (2018), 132–153. **1, 2, 10**
- [QSO*12] QUALTER J., SCULLI F., OLIKER A., NAPIER Z., LEE S., GARCIA J., FRENKEL S., HARNIK V., TRIOLA M.: The biodigital human: a web-based 3D platform for medical visualization and education. *Studies in health technology and informatics* 173 (2012), 359–361.
- [RJLL13] RUIZ JR C. R., LE S. N., LOW K.-L.: Generating multi-style paper pop-up designs using 3D primitive fitting. In *SIGGRAPH Asia 2013 Technical Briefs*. 2013, pp. 1–4. **2**
- [RJLYL14] RUIZ JR C. R., LE S. N., YU J., LOW K.-L.: Multi-style paper pop-up designs from 3D models. In *Computer Graphics Forum* (2014), vol. 33, pp. 487–496. **2**
- [RL15] RUIZ C., LE S. N., LOW K.-L.: Generating animated paper pop-ups from the motion of articulated characters. *The Visual Computer* 31, 6-8 (2015), 925–935. **2**
- [RMVTK*10] RENGIER F., MEHNDIRATTA A., VON TENGG-KOBLIGK H., ZECHMANN C., UNTERHINNINGHOFEN R., KAUCZOR H.-U., GIESEL F.: 3d printing based on imaging data: Review of medical applications. *International Journal of Computer Assisted Radiology and Surgery* 5 (2010), 335–341. **1, 2**
- [SB17] STOPPEL S., BRUCKNER S.: Vol2velle: Printable interactive volume visualization. *IEEE Transactions on Visualization & Computer Graphics* 23, 01 (2017), 861–870. **1, 2**
- [SCS16] SASS L., CHEN L., SUNG W. K.: Embodied prototyping: exploration of a design-fabrication framework for large-scale model manufacturing. *Computer-Aided Design and Applications* 13, 1 (2016), 124–137. **2**
- [SCS17] STEINMAN D., COPPIN P., STEINMAN D.: Data physicalization for improving communication and memorization. In *16^o Encontro Internacional de Arte e Tecnologia 16th International Meeting of Art and Technology* (2017), p. 512. **1**
- [SHK*16] SMIT N., HOFSTED E. C.-W., KRAIMA A., JANSMA D., DERUITER M., EISEMANN E., VILANOVA A.: The online anatomical human: Web-based anatomy education. In *Proceedings of the 37th Annual Conference of the European Association for Computer Graphics: Education Papers* (2016), pp. 37–40. **2**
- [SSPOJ16] SAALFELD P., STOJNIC A., PREIM B., OELTZE-JAFRA S.: Semi-immersive 3D sketching of vascular structures for medical education. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2016), pp. 123–132. **2**
- [Tho53] THORNDIKE R. L.: Who belongs in the family? *Psychometrika* 18, 4 (1953), 267–276. **6**
- [Ves43] VESALIUS A.: De humani corporis fabrica libri septem, 1543. **2**
- [VG05] VIOLA I., GRÖLLER M. E.: Smart visibility in visualization. In *Computational aesthetics* (2005), pp. 209–216. **3**
- [VGHN08] VÁZQUEZ P.-P., GÖTZELMANN T., HARTMANN K., NÜRNBERGER A.: An interactive 3D framework for anatomical education. *International journal of computer assisted radiology and surgery* 3, 6 (2008), 511–524. **2**
- [XZM*18] XIAO N., ZHU Z., MARTIN R. R., XU K., LU J.-M., HU S.-M.: Computational design of transforming pop-up books. *ACM Transactions on Graphics* 37, 1 (2018), 1–14. **2**
- [ZM08] ZHAO J., MOERE A. V.: Embodiment in data sculpture: a model of the physical visualization of information. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts* (2008), pp. 343–350. **1, 2**