

Squishicalization: Exploring Elastic Volume Physicalization

Daniel Pahr, Michal Piovarči, Hsiang-Yun Wu, and Renata G. Raidou

Abstract—We introduce *Squishicalization*, a pipeline for generating physicalizations of volumetric data that encode scalar information through their physical characteristics—specifically, by varying their “squishiness” or local elasticity. Data physicalization research is increasingly exploring multisensory information encoding, with a particular focus on enhancing direct interactivity. With *Squishicalization*, we leverage the tactile dimension of physicalization as a means of direct interactivity. Inspired by conventional volume rendering, we adapt the concept of transfer functions to encode scalar values from volumetric data into local elasticity levels. In this way, volumetric scalar data are transformed into sculptures, where the elasticity represents physical properties such as the material’s density distribution within the volume. In our pipeline, scalar values guide the weighted sampling of the scalar field. The sampled data is then processed through Voronoi tessellation to create a sponge-like structure, which can be printed with consumer-grade 3D printers and readily available filament. To validate our pipeline, we conduct a computational and mechanical evaluation, as well as a two-stage perceptual study of the capabilities of our generated squishicalizations. To further investigate potential application scenarios, we interview experts across several domains. Finally, we summarize actionable insights and future avenues for the application of our *Squishicalization*. All supplemental materials are available at https://osf.io/35gnv/?view_only=605e5085061f40439a98545f0c447cf3.

Index Terms—Data Physicalization, Digital Fabrication, Volume Data, Tactile Information Encoding

I. INTRODUCTION

DATA physicalization involves encoding information into physical properties of artifacts [16]. Although static physicalizations have long been used for data exploration [8], recent approaches have proposed interactive data physicalizations [3]. As data physicalization goes beyond visually inspecting physical representations and towards multisensory experiences [13], [14], the integration of haptic feedback could hold significant promises for tangible data exploration. Tactile feedback offers a means for enhancing interactivity and introduces new opportunities for engagement, allowing users to interact directly with their data [6]. In data physicalization, tactile feedback is typically realized by making use of the surface properties of objects. Yet, when considering volumetric data, existing research on their physical representation has primarily focused on developing interactive and engaging sculptures that purely use the visual channel [30], [33], [35].

Inspired by volume rendering, our work seeks to bridge a gap in the literature by introducing *Squishicalization*, a

method for translating scalar information in volumetric data into tactile properties of physical artifacts. Similar to how volume rendering techniques map scalar data values to visual properties through transfer functions (TFs), our pipeline extends this concept to density transfer functions (DTFs). These map scalar values of volumetric data onto local density, which subsequently determine the elasticity or “squishiness” of the resulting physical artifacts. This results in squishicalizations, i.e., physicalizations that represent scalar distributions within the data by leveraging physical properties related to elasticity. By adjusting the DTFs, designers control the tactile properties of the physical artifacts to reflect the density distribution of materials within the volumetric dataset. To achieve this, we employ the concept of Voronoi tessellated microstructures [27] from digital fabrication within our squishicalizations, which are later fabricated with elastic materials using 3D printing. While we show examples created with fused filament fabrication (FFF), any 3D printing method that processes elastic material is equally suited.

Unlike previous works, our approach introduces direct interactivity by allowing users to engage with the data through tactile exploration of the material’s elasticity, offering a novel way to explore volumetric data beyond surface structures. With *Squishicalization*, we aim to answer the following **research questions**: (Q1) How can we use consumer-grade 3D printing to create tangible artifacts that map volumetric data onto physical properties, specifically elasticity? (Q2) How can we control the elasticity of our tangible artifacts to ensure an effective correlation to the scalar values of the volumetric data? (Q3) Which practical applications could benefit from elastic volume physicalization?

We design an interactive elastic volume physicalization pipeline, *Squishicalization*, to create tangible sculptures from scalar fields in volumetric data. The individual steps of the pipeline can be likened to those of the volume rendering pipeline. We subsequently evaluate our pipeline’s computational, mechanical, and perceptual capabilities. From this evaluation, we distill recommendations for the design of our squishicalizations. We also conducted interviews with expert users from three fields—extended reality (XR), materials science, and medical education—to gather insights into the applicability of our proposed *Squishicalization*.

II. RELATED WORK

Various levels of interactivity are seen in data physicalization—from full-body experiences to surface-level tactile interactions. Incorporating data directly into

Daniel Pahr (contact author, e-mail: daniel.pahr@gmail.com) and Renata G. Raidou are with TU Wien, Austria. Michal Piovarči is with ETH Zurich, Switzerland. Hsiang-Yun Wu is with St. Pölten University of Applied Sciences, Austria and TU Wien, Austria.

material characteristics like elasticity remains largely unexplored, even though fabrication pipelines that leverage direct interactivity (also in the form of elasticity) exist. However, encoding volumetric data in the internal physical properties of objects has not yet been explored and we anticipate that microstructure-based fabrication approaches could offer a promising method in this direction.

Data Physicalization and Direct Interactivity. Data physicalizations bring data into the physical space by encoding information into the material properties of artifacts [18]. The value of physical representations has been underscored multiple times. For instance, Jansen et al. [17] show that in the context of 3D bar charts, a physical version can outperform its screen-based counterpart, indicating that the audience's visual perception can be enhanced by physical manifestation. Building upon this work, Taher et al. [41] propose physical bar charts with added interactivity allowing users to manipulate the display by indirect interaction.

Interaction in physicalization can be designed to directly stimulate the user's senses [3]. This is known in the literature as *direct interactivity*. Karyda et al. [19] use this interactivity to emotionally engage an audience with personal data. Djavaherpour et al. [8] demonstrate the use of the surface properties of 3D-printed objects as a form of information encoding. Ion et al. [15] describe a metamaterial-based approach to create textures that morph between discrete states to provide continuous information. Multisensory data representations are also common [13] with a prominent example being *Move&Find* [14]. This work integrates the user into the physical representation generating a full-body experience, where the data is encoded as physical feedback to the users' actions.

Physicalization Toolkits and Pipelines. While there are numerous frameworks to create data visualizations for different types of data, the large design space for data physicalization does not allow the easy development of a general framework. As an exception, *MakerVis* [40] supports the design of physical versions of well-known visualizations from data selection to fabrication. Conversely, in the context of volume data physicalization, several fabrication workflows and pipelines for the computer-aided generation of physical representations have been presented, targeting mostly laypeople edutainment. Schindler et al.'s [34], [35] workflow creates 3D (nested) papercrafts from digital meshes of anatomical or biological data. However, the underlying volumetric data have to be transformed into geometry meshes in a preprocessing step. Raidou et al.'s [33] *Slice&Dice* demonstrates a pipeline for the computer-aided generation of 3D sculptures from volumetric data. They use regular printers to process transparent foils and octree-based spatial arrangements to create a 3D effect. Pahr et al. [30] use similar materials and configurations for segmented medical volumetric data to allow the 2D/3D exploration of anatomical structures.

3D Printed Physicalization. On top of solutions with conventional printing, such as those discussed in the previous subsection, 3D-printed solutions are also prevalent. Djavaherpour et al. [9] summarize the different possibilities for creating data physicalizations using digital fabrication methods. They identify additive fabrication as a suitable method for non-

expert use to create detailed representations of limited size. Because of its accessibility, cost-effectiveness, and ease of use [11], several physicalization concepts have been fabricated using fused filament fabrication (FFF). For instance, Stusak et al. [38] investigate different methods to create exercise feedback using 3D-printed artifacts. Similarly, Khot et al. [21] propose 3D printed physicalization concepts for increasing the understanding of physical activity. Using *Fantibles* [20], the social media usage of sports fans has also been embodied in compact artifacts. In these cases, FFF printing is employed to create physicalizations for immediate feedback, supported by the relatively fast printing times of the method.

Additionally, FFF printers often offer simple modifications to create more complex prints. Ang et al. [1] use multi-nozzle FFF printing to create three-dimensional representations of blood flow in the human heart. Bae et al. [2] use multi-material printing to embed conductive material into their network physicalization to detect physical interaction on a microcontroller. Instead of using electronics, *AirTouch* [42] uses pneumatic sensing in elastic objects to detect user interaction. Also using pneumatics, Kim et al. [22] present a pipeline to create inflatable elastic 3D-printed shape-changing objects. Torres et al. [43] introduce a fabrication toolkit, *HapticPrint*, to design haptically different interactive objects. Among other things, they use infill patterns to create different levels of elastic compliance.

3D Printed Microstructures. Fabrication devices typically offer only a limited selection of materials [37]. To enhance the design space, *metamaterials* are often used [4]. These leverage the high spatial resolution of the fabrication hardware to generate tiny microstructures. Bickel et al. [5] propose to control the deformation behavior of 3D printed objects by using microstructures. In this way, it is possible to manifest bulk behavior similar to materials not available for printing. For instance, *Built-to-last* [25] uses a honeycomb tessellation of 3D objects based on internal force distribution to maximize the strength-to-weight ratio in solid objects. Also, Martinez et al. [27] adapt Voronoi diagrams to the context of fabrication by replacing the distance function. They propose a special family of distance functions that guarantee fabrication constraints. Unfortunately, achieving control over the deposition at the micro-scale requires a tight coupling of the microstructure generation with the fabrication hardware. To this end, specific techniques were devised for inkjet printers [47], powder sintering machines [36], photosensitive resin printers [26], and fused filament fabrication [44].

III. METHODOLOGY OF *Squishicalization*

We present *Squishicalization*, a pipeline for the computer-aided generation of tangible physicalizations that map scalar information from volumetric data onto their internal elasticity. The goal is to directly stimulate the user's senses by employing tangibility as a means of direct interactivity [3].

We refer to **elasticity** as a material's ability to deform by an applied force and subsequent to its original shape and size. Here, more elastic means that an object or a region in an object deforms more than a less elastic object when the




Fig. 1. The steps of our *Squishicalization* pipeline. (a) Scalar volumetric data is read from a file or created by a mathematical function. (b) A subsampling (gray: random sample, red: selected seeds) based on the scalar distribution yields seed points to be used for the (c) Voronoi tessellation of the volume. (d) The regions corresponding to the tessellation borders are extracted using iso-contouring.

same force is applied. **Stiffness** refers to an object’s resistance to (elastic) deformation. Our volume physicalization pipeline comprises four steps (Figure 1) and draws upon foundational concepts of the volume rendering pipeline [28], which creates visual representations of volumetric data using properties, such as color and opacity, within well-defined TFs [23]. Comparably, our approach aims to generate physical objects from volumetric data that physically encode scalar information onto their physical characteristics—specifically, their elasticity.

We adapt the TF concept to reflect the scalar distribution within the volume onto the distribution of elasticity inside the physicalization. Instead of manipulating color and opacity to enhance structure visibility, we manipulate the “squishiness” based on the scalar values to enhance tangibility. This allows us to create physical objects where the distribution of elasticity maps the distribution of scalar values in the data.

Physicalizing **volumetric data**, as those depicted in Figure 1a, is challenging due to the complexity of representing their inherent 3D scalar fields. Slicing and filtering [30], [33] have been commonly employed in literature to enhance structure readability. We propose to map the scalar values of a volumetric data set to local elasticity. In a 3D medical imaging dataset, for example, we can represent denser tissues like bones with lower elasticity (i.e., denser means stiffer). This is done by our DTFs, which are comparable to the concept of TFs in conventional volume rendering, and influence sampling densities. Here, a **sampling step** generates a number of seeds per region, reflecting locally the density of the underlying structures (Figure 1b). Following up on the medical imaging example above, bones will be sampled with a higher number of seed points than soft tissues. We opt for weighted sampling, to mitigate clustering of the samples [46]. Subsequently, using **Voronoi tessellation** we divide the volume data into labeled regions using the seed points resulting from the sampling step as the region centers (Figure 1c). The Voronoi regions will produce a sponge-like structure, such as the one shown in Figure 1d. At this point, dense areas are tessellated into a high number of small regions, whereas sparse areas result in fewer regions of larger volumes [27]. In the final step, we produce the squishicalizations by **FFF 3D printing**, chosen for its accessibility and cost efficiency [11]. FFF printing allows the use of elastic materials, which, combined with

sponge-like microstructures, can produce soft artifacts. In the resulting squishicalization, denser microstructures can be physically compressed to a lesser degree. This means that our soft artifacts directly encode the scalar values of volumetric data in the form of local elasticity. We elaborate on the steps of our pipeline in the remainder of this section and we denote with colored blocks algorithmic the substeps within each step. Several printed examples and their respective elasticities are demonstrated in Figures 4 and 5.

A. Data Acquisition and Basic Processing

We consider a volumetric dataset $V : \mathbf{R}^3 \rightarrow \mathbf{R}$ consisting of scalar values in a regularly spaced three-dimensional grid.

Such data can be obtained in various ways, e.g., through imaging modalities such as computed tomography (CT) or magnetic resonance imaging (MRI). In these cases, the scalar values correspond to measurements of physical phenomena, such as the absorption of X-rays in materials or the response of molecules to magnetic polarization. These scans vary in voxel size, from around 1mm in full or partial body scans to much smaller in tissue scans, where individual cells can be viewed. Although prevalent in medicine, **3D image data** are also used in other fields, e.g., in materials science. Alternatively, scalar volumetric data can be synthetically generated by **mathematical functions** in the form $f(x, y, z) = s$, which can be sampled over discrete points in a given interval. The result can be visualized in the 3D space, where each voxel encodes the function value at a given point. The resulting volumetric data are commonly viewed on screens as axial slices, using volume rendering techniques [28], or by extracting one or multiple isosurfaces [24].

A typical step in volume rendering is **thresholding**. It can be used to segment volumetric data by eliminating unwanted scalar values within user-determined ranges (e.g., under a certain value). This way, artifacts from data acquisition can be removed or the background of an image can be ignored, reducing the amount of data to be processed. By extracting an **isosurface** from the scalar field, we can additionally create a 3D mesh of a structure. Such meshes can be combined in the printing

process to add solid parts, achieving the upper threshold of stiffness for a given material, bypassing the need to create a microstructure.

B. Weighted Volume Sampling

Our approach aims to transform our volumetric data into regions of different elasticity, depending on their scalar values. Previous work by Lu et al. [25] deals with the fabrication of solid models based on a force distribution, while Martinez et al. [27] investigate the generation of elastic models based on given point distributions. In contrast, our method introduces an interactive approach for generating point distributions directly from diverse volume data, based on the following simple concept: a continuous range of scalar values (e.g., each of the four colored regions in Figure 1a) should correspond to a specific stiffness within the final physical object. Therefore, we adapt the approach of Martinez et al. to create point distributions based on the underlying scalar values of the volumetric data. We design our approach **independently of prior segmentation**, and directly rely on the scalar values.

The tessellation is discussed in the upcoming section, but its basis requires a set of points, which can be produced through sampling that follows these three requirements (Figure 1b): First, **regions with similar scalar values should have similar**

on scalar values does not guarantee this. For instance, Lu et al. [25] use an error diffusion approach to create their sample distribution, which also does not guarantee regularly spaced sampling points. Random sampling can introduce undesirably clustered points in regions where the scalar field distribution is stable, introducing tangible artifacts instead of reflecting the scalar field accurately.

Second, **point density of a region should relate to its scalar values**. Ideally, the relationship between scalars and sampling density should be customizable to allow flexibility in manufacturing. Here, we call back to the application of TFs in volume rendering, and we employ a similar concept to relate scalar values to elasticity through sampling density (density TF or DTF, further described below). Third, we need to **limit computational complexity** as much as possible. Tessellation is a computationally expensive process, whose performance depends on the selected number of sample points. In our case, this is further complicated by the dimensionality of our data.

A method that fulfills all of these requirements is **weighted Poisson-disk sampling** [7]. In this method, a minimum distance between two sample points is defined to create an equally spaced point set. This distance can be a function of a local weight, and in our case, this relates to the scalar values of the data. However, the performance of Poisson-disk sampling algorithms is impacted by the dimensionality of the sampling space, as well as the amount of candidate points. Yuksel [46] proposes a sample elimination approach to create sets with Poisson-disk characteristics from a given point set. This method enables us to create sample sets with equal spacing of points, while the minimal distance between points can be a function of the local scalars. A random sample of points of a sufficient size inside the original volume can serve as the input for the sample reduction. The size of the desired sample is also a parameter of our algorithm.

To obtain a set of sample points we use a sample reduction technique that **iteratively eliminates points in a given set until a desired sample size is reached**. We follow Yuksel's approach [46] for the sample reduction, using the recommended parameter values $\alpha = 8$, $\beta = 0.65$, and $\gamma = 1.5$. A higher α increases the influence of point distance on the weight function, while β and γ values are used to lessen the influence of very close points compared to far-away ones. The approach also allows for the use of a weight function that considers local scalar values in addition to point distances.

For the customization of the weighted sampling, we propose the use of **density transfer functions (DTFs)**, an analogous concept to color and/or opacity TFs in volume rendering. We define a DTF as a function $f(s) = d_{max}(s)$, where s is a scalar value in a volume dataset and d_{max} is the maximum radius for the given scalar value. Conventional TFs make use of the concept of **piecewise linear TFs** [23]. As we do not know the precise relation between sample density and elasticity and users may prefer to emphasize certain structures over others (Figure 2), piecewise linear functions are also suited for DTFs.




Fig. 2. Weighted sampling of a volumetric dataset for different DTFs. The volume rendering of a head CT dataset is overlaid with gray points (sampling input) and red points (sampling output). The chosen (red) sample points obtained with each of the custom DTFs are more clustered around cerebral structures (soft structure) and less around the skull (hard structure).

elasticity. Thus, their underlying tessellations or cells should be sized approximately equally to ensure similar elastic properties [27]. Limiting the number of points used to represent a given volume, or adjusting the sampling probability based

We define a **DTF** by supplying a set of control points $\mathcal{C} = C_i$, consisting of value pairs $C_i = (s_i, d_i)$, where d_i is the maximum radius parameter for a given scalar value s_i . The maximum radius for a scalar value at a given s is calculated as the linear interpolation between d_k and d_{k+1} , where $s_k < s < s_{k+1}$.

Figure 1b schematically depicts the weighted sampling concept for four regions with different scalar values, Figure 2 shows a concrete example of the sampling process for two different DTFs.

C. Polyhedral Voronoi Structure Generation

After obtaining the set of sample points $\mathcal{C}$, we create a tessellation from the original volumetric dataset (Figure 1c). From this, we extract the desired sponge-like microstructure. For the tessellation, we pose the following requirements: First, our goal is to create **printable sculptures** that can be produced using FFF, which allows the printing of elastic objects using specific materials. Such 3D printers create objects by adding layer-by-layer to the printing area. Yet, large overhanging areas have to be supported by creating support structures underneath, which would influence the mechanical properties of the print. Also, the microstructures we aim to create do not allow the extraction of support structures after print. We, therefore, require that printing the microstructures does not require support. Second, squishicalizations should exhibit **isotropic elasticity**. This means that the tessellation should not influence the mechanical properties in a non-uniform way (e.g., along a specific axis as in the case of Martinez et al. [27]). Otherwise, this would result in inconsistent behavior when exploring artifacts encoding scalar values into local densities, where the direction of compression could significantly impact tactility. Third, **shapes** inherent to the volumetric data should be preserved as much as possible. This could mean that the shape of the anatomical area depicted in medical images can be reproduced with printing. Additionally, limiting the volume to “visible” areas (e.g., by suppressing the background with a threshold, as discussed in Section III-A) can increase the speed of the operation performed significantly, since “invisible” areas do not have to be tessellated.

We use a **Voronoi tessellation**-based approach to create elastic objects with customizable density. In a **Voronoi tessellation** of a volume, given a set of sample points, each voxel contains the label of its assigned sample point. A **distance function** determines the proximity of all samples to a single voxel. The labels are chosen by finding the closest sample point.

Martinez et al.’s [27] approach uses **polyhedral distance functions**, based on parametrizable cones. They argue that, when using a Euclidean distance function, the borders of the resulting areas could have arbitrary angles, causing overhangs that cannot be printed without support structures. In contrast, their customizable cone-shaped distance function can limit border angles resulting in the tessellation to exclude steep




Fig. 3. (a) Labeled regions and (b) extracted border mesh, created by Voronoi tessellation with a bi-truncated honeycomb distance function. Sample point set taken from Figure 2, bottom right.

drops. We fulfill our first requirement, by using a similar polyhedral distance function for the tessellation.

A **polyhedral distance function** is defined as follows: To calculate the distance between points P_1 and P_2 , we construct a convex polyhedron that contains P_1 , defined by a set of planes $\mathcal{B}$ traversing the vertices of the polygon, defining its faces. We first calculate the intersections with every plane on a connecting line between P_1 and P_2 . The closest intersection point I_{min} to P_1 indicates where the connecting line intersects with the polygon. The polyhedral distance between P_1 and P_2 is the ratio between the Euclidean distance between the points and the distance between P_1 and I_{min} so that $d_{\mathcal{B}}(P_1, P_2) = \frac{|P_2 - P_1|}{|I_{min} - P_1|}$. This corresponds to a **scaling factor** applied to the given polygon around P_1 so that it just touches P_2 , without rotating the polygon. In the case where the line between two candidate seed points is parallel to a face of the polygon, i.e., the two computed distances would be equal, we opt for the seed point with the lower index.

From our second requirement, we do not desire customizable anisotropy, i.e., we want to avoid adjusting directional variations in the properties of the model. Thus, we use the bitruncated cubic honeycomb polyhedron. This way, the **distance function guarantees isometry** along every axis. As a honeycomb shape, this polyhedron can be arranged to create a tessellation of 3D space using only translation. It also contains only steep angles and horizontal planes, small enough to bridge small gaps to reduce the need for support—in accordance with the first requirement. To fulfill the third requirement, we limit the seed points for the tessellation by defining a background scalar value, which will be suppressed and excluded from both subsampling and tessellation. For example, in Figure 1 the white region around the volume is not considered. This limits performance impacts and produces a volume that **preserves the shape** of a region of interest. Figure 1c schematically depicts this process, while we show a labeled tessellation performed with this method in Figure 3a. Figure 9 shows examples with different regular point densities.

D. Meshing

To create a printable mesh, we need to obtain the region borders from our tessellation result (Figure 1d). Martinez et al. [27] investigate a layerwise path extraction. These methods create a large performance bottleneck. As the Voronoi regions in the mesh are already labeled with integer values, we opt for a labeled iso-surface extraction approach [10] instead:

To extract the region border we use an implementation of Frisken’s surface nets algorithm [10] in PyVista [39]. It produces a single mesh of region borders for a labeled volume. Because we also obtain a label for the background voxel value in the first step, the object’s outer shape is preserved. In addition to the extraction of borders, this approach also produces smoother edges and more equally spaced vertices, removing undesirable staircase artifacts and simplifying the generated geometry.

Figure 1d schematically depicts the mesh extraction process, while we also show a 3D result in Figure 3b.

E. Slicing and Printing

In the last step, we process the resulting mesh using readily available slicing software. After preparation, the squishicalization can be printed. Because a multitude of 3D printers and a wide range of materials exist, we require this step to be **independent of printing hardware** and as **customizeable** as possible. For accessibility, we also require the **materials to be compatible** with consumer-grade 3D printers. Slicing software serves to transform triangle meshes into toolpaths. These are created by cutting a mesh along the vertical axis and analyzing the geometry on the cutting plane. Many aspects of this process are customizable, such as the selection of printer, wall thickness, and parameters of support structures. Such software is usually designed for closed, non-manifold meshes, as a slicer needs to determine the inside and outside of a mesh to create infill patterns that are used to limit filament usage while preserving the mechanical stability of a 3D-printed object. However, multi-label segmentation is not designed to produce such well-behaved meshes. *UltiMaker Cura* provides a **surface mode** that processes a mesh independently of composition and extracts any borders in a mesh as a single wall to print. *UltiMaker Cura* supports toolpath creation for a multitude of contemporary 3D printers. The use of slicer software also allows us to **combine different meshes** in a single print. This way, solid objects can be suspended in soft material, increasing the range of elasticity we can achieve. We use **thermoplastic polyurethane (TPU)** as it is commercially available in a wide range of colors, from different manufacturers, and compatible with regular FFF 3D printers. An example of the head dataset is demonstrated in Figure 4.

F. Implementation

We implemented our pipeline in *Python*, with slicing handled externally by *UltiMaker Cura*. The weighted Poisson-disk sample reduction algorithm proposed by Yuksel [46] was extended to be compatible with the concept of DTFs using




Fig. 4. A demonstration of the result (right) of the *Squishicalization* pipeline for a volumetric dataset of a human head (left), acquired by Magnetic Resonance Imaging (MRI). Local densities in the squishicalization correspond to scalar values within the volumetric dataset.

NumPy. We parallelize the calculation on the GPU using *Numba* to create *CUDA* kernels for the polyhedral voronoi tessellation approach proposed by Martinez et al. [27]. An implementation of Friskens’ [10] surface nets algorithm in *PyVista* serves for surface extraction and mesh smoothing. Our implementation is available at *OSF*.

IV. Squishicalization RESULTS

We now show some results of our physicalization pipeline and test the capabilities of our approach in a computational, a mechanical, and a perceptual evaluation. For more details on the tasks, procedure, results of the study, and prop descriptions, we refer the reader to our supplemental materials.

A. Performance

We tested different parameter combinations and evaluated the impact of individual settings on the performance of *Squishicalization*, using a total of four datasets (Table I). The **generic** dataset consists of $300 \times 300 \times 300$ voxels with randomized scalar values. The **foot** dataset is a CT scan of a human foot and the **head** dataset is an MRI of a woman—both taken from the Open SciVis Datasets. The **fibers** dataset

TABLE I
PERFORMANCE FOR DIFFERENT DATASETS. *Head** MARKS THE USE OF A SCALAR THRESHOLD, REDUCING THE IMPACT OF THE TESSELLATION PERFORMANCE. n_r INDICATES THE NUMBER OF RANDOM SAMPLES USED IN THE SAMPLE ELIMINATION AND n_s INDICATES THE NUMBER OF SEEDS, t INDICATES MEASURED TIME ($m : ss$).

Dataset	Size	Sampling			Tessel.	Meshing	Total
		n_r	n_s	t			
Generic	$300 \times 300 \times 300$	1k	0.3k	<1s	00:53	<1s	0:54
		5k	1.5k	0:01	4:10	0:01	4:12
		10k	3k	0:04	8:27	0:02	8:33
Foot	$256 \times 256 \times 256$	10k	3k	0:12	1:32	<1s	1:45
Fibers	$350 \times 350 \times 400$	10k	0.3k	0:13	1:36	<1s	1:51
Head	$256 \times 256 \times 70$	60k	20k	17:09	15:09	<1s	32:20
Head*	$256 \times 256 \times 70$	60k	20k	17:11	2:18	0:01	19:31




Fig. 5. Resulting squishicalizations for different volumetric datasets printed on a Prusa i3 Mk3s using an opaque TPU filament. (a) Foot CT dataset from the Open SciVis Datasets. (b) Head MRI dataset from the Open SciVis Datasets. (c) Industrial CT dataset provided by a materials science expert.

was provided by a domain expert in the later course of an interview, described in Section V. We limit the parameter space exploration of the performance to the parts of the pipeline we implemented (Sections III-A– III-D), as slicing (Section III-E) is performed in external software. The tests were performed on a system with an AMD Ryzen 9 3900X 12-Core Processor with 3800 Mhz, an NVIDIA GeForce GTX 1080, and 64GB RAM. We show the results in Table I. Besides the FFF 3D printing bottleneck, our benchmarks show that the tessellation is the biggest time sink ranging between 0 : 53 and 8 : 27 minutes. The sampling becomes an intense factor for big sample sizes. While processing the head dataset, we created a sample of 20k from 60k random points to account for the size of the final print. Sampling took 17 : 09 minutes, while tessellation of the whole volume took 15 : 09. Limiting the tessellation to voxels with values greater than a predefined threshold reduced the time for tessellation to 2 : 18. Mesh creation is always negligible (approx. one second).

All prints were done on a Prusa i3 Mk3s, with a wall thickness of 0.2mm. For solid parts, we used 100% infill density. Regarding **printing times** and **material use**, the foot dataset on a scale of $10cm \times 10cm \times 10cm$ took about one day and 5 hours, using 78g of filament. The foot was printed combined with a mesh created from an isosurface extracted the scalar values representing solid bones. In the DTF, only soft tissue is considered in the sampling process and bones are initially excluded; and then, reinserted as an isosurface. The outcome is shown in Figure 5a. The fibers dataset took about 19 hours to print on a scale of $5cm \times 5cm \times 6cm$, and used 44g of filament. Here, we also used an isosurface of the fiber material in the print and a dense sampling of the material in which the fibers are embedded. This allows the fibers to move more freely while still maintaining a coherent structure. The outcome is shown in Figure 5c. Finally, the head dataset, at a scale of $23cm \times 20cm \times 9m$, took two days and 7 hours to print and consumed 194g of filament. The DTF was designed to emphasize a tactile difference between gray and white matter, while harder tissues were sampled more sparsely. We did not include solid parts in the print. To save time and

material, and to facilitate an inside view of our microstructure, we clipped the dataset along the coronal axis and excluded all scalar values under 7% of the maximum, to limit noise and performance impact. The *squishicalization* of the head dataset is shown in Figures 4 and 5b.

Upon visual inspection, our test prints show that vertical walls are well pronounced. As the structures approach more flat angles, staircase artifacts consisting of disconnected strings become visible. This can be explained by the slicing technique we employed to process the meshes. *UltiMaker*’s “Surface Mode” produces a single wall for every intersection of the mesh with the horizontal plane. When the distance between two sequential layers becomes too large, this leads to layers becoming disconnected at places. We further discuss the impact of these artifacts in later sections.

B. Perceptual Study

We conducted a study with 18 participants to investigate the abilities of our squishicalizations to stimulate tactile perception. The objective of our study was to investigate whether it is feasible for humans to recognize and interpret complex 3D structures by directly interacting with their respective squishicalizations. We are particularly interested in properties such as partitions, gradients, and the presence of heterogeneous regions within 3D synthetic stimuli.

Participants. We recruited our participants from university staff and students without offering compensation. Our participants (male: 16, female: 2) were between 22 and 36 years old (Mean=29, Std=4.05). Two hold a Ph.D., 14 a M.Sc., one a B.Sc., and one has a high school education.

Tasks. Our tasks were designed based on three objectives: *recognizing partitions* within the volume, i.e., part A is less elastic/denser than part B, *elasticity gradients*, i.e., directional changes of elasticity, and *position and size* of heterogeneous regions in objects with variable density. In the **partitions** task, participants had to decide how a cube consisting of two equally sized regions differing in stiffness was partitioned. We handed participants one cube at a time, each cube split along a different plane. The **gradients** task had participants determining

the direction of a stiffness gradient running through a cube. Again, we handed them one cube at a time, each created with a stiffness gradient along a different diagonal. Finally, in the **regions** task, participants had to determine the diameter and position of a spherical region with a higher stiffness than the surrounding volume.

Datasets. 3D scalar data with varying partitions, gradients, and heterogeneities can be easily expressed by 3D arrays, simplifying the creation of the datasets for the study objects. Hence, we created three cubes with 50mm side lengths for each task. An additional top and bottom thin layer was added to the cubes to visually obscure the inner microstructures. All samples were printed with 0.3mm wall thickness using an opaque TPU filament. For the **partitions** task, we created synthetic cubes with two regions of distinct densities. These were parted in the middle along either one of the coordinate axes (P1), one of the face diagonals (P2, see Figure 6a), or the body diagonal of the cube (P3). In the **gradients** tasks, we used synthetic cubes with a linear density gradient. We created three datasets, with increasing values along one of the coordinate axes (G1), one of the face diagonals (G2), and the body diagonal of the cube (G3, see Figure 6b). Lastly, for the **regions** tasks, we created synthetic cubes of homogeneous density containing an additional spherical region with a different density than the rest of the cube. We used three different diameters for the regions, 20% (R1), 33% (R2, see Figure 6c), and 50% (R3) and placed the regions in randomized positions within the cube.

Procedure. Each participant performed all three tasks with three cubes per task presented in a balanced, pseudorandom order. Every cube was presented to the study participants an equal amount of times in all three steps in each task by using a Latin square scheme. Nine of the participants completed the tasks in order (partitions, gradients, regions), and nine completed them in reverse order to combat the influence of learning effects. We measured the **time** participants spent on each experiment. Additionally, participants were asked to rate perceived task **difficulty** and the **confidence** in their answer on a scale of 1–10, ranging from low to high. For the **partitions** and the **gradients** task, we gave the participants **single-choice questions** with 7 possible answers, meaning a 14.23% chance of randomly guessing the answer. The possible answers in the partitions task were three axis-parallel partitions, three partitions along the face diagonal, and the body diagonal of




Fig. 6. Three different datasets were used in our perceptual study. Different colors represent different scalar values. (a) Dataset for the partition task P2 (face diagonal partition), (b) Dataset for the gradients task G3 (gradient along body diagonal), (c) Dataset for the regions task R2 (medium-sized region)

TABLE II
QUANTITATIVE STUDY RESULTS FOR THE PARTITIONS AND GRADIENTS TASK. TIMES IN *m : ss*, ACCURACY FOR POSITIONS AND GRADIENTS IN THE PERCENTAGE OF TOTAL RIGHT ANSWERS, REGION POSITIONING ERROR IN PERCENT OF CUBE SIDE LENGTH, AND REGION SIZE ERROR SHOWN IN RELATION TO CUBE SIZE (BOTH INCL. $\pm$ STANDARD DEVIATIONS AND THE TIMES, DIFFICULTY, AND CONFIDENCE ARE INDICATED FOR BOTH POSITIONING AND SIZE ESTIMATIONS TOGETHER).

Dataset	Accuracy	Time		Difficulty		Confidence	
		avg	std	avg	std	avg	std
P1	50%	1:19	0:42	4.83	2.75	6.61	2.45
P2	17%	1:31	1:03	4.72	2.61	6.83	2.12
P3	22%	1:45	1:04	5.56	2.36	5.94	1.92
P1-3	29.63%	1:32	0:57	5.04	2.55	6.46	2.17
G1	39%	1:47	0:49	6.06	2.58	5.78	2.18
G2	11%	1:51	1:03	5.94	1.86	5.50	1.58
G3	22%	1:55	1:03	5.50	1.58	5.67	2.00
G1-3	24.07%	1:51	0:58	5.98	2.11	5.58	1.87
R1 pos.	38% $\pm$ 18%						
R1 size	4% $\pm$ 10%	2:22	1:13	6.28	2.63	5.33	2.35
R2 pos.	15% $\pm$ 12%						
R2 size	-4% $\pm$ 17%	1:45	0:29	4.22	2.24	7.17	2.09
R3 pos.	16% $\pm$ 9%						
R3 size	-2% $\pm$ 22%	2:25	1:04	5.56	2.09	6.28	1.99
R1-3		2:11	1:00	5.35	2.44	6.26	2.24

the cube. In the gradients task, the possible answers were the three gradients along the coordinate axes, three along the face diagonals, and one along the body diagonal. In the regions task, we asked the participants to indicate the **position** of the center and the **diameter** of the denser region in the cube, measured proportional to the cube side length. We let our participants freely interact with our study props. We also asked participants for **qualitative feedback** after each task and after they completed all three tasks.

Study Results. For the partitions and gradients task, we report the accuracy of the answers of our participants for each dataset. For the regions task, we report the mean deviation of our participants' prediction of the center and diameter of the spherical region from the ground truth in percent. We also show the results of the regions task in Figure 7. For all tasks, we report the average time taken, perceived difficulty, and confidence. The results are shown in detail in Table II.

Analyzing the **quantitative** results shows that some participants had trouble recognizing partitions (17%) and gradients (22%) along face diagonals. They were only slightly more accurate in recognizing partitions and gradients along the body diagonal (both 22%). Half of the participants correctly




Fig. 7. Study results for the regions task. The red sphere in each picture represents the ground truth, while each grey sphere represents the answer of one participant for each of the three cases ((a) R1, (b) R2, (c) R3).

indicated the partition along the coordinate axis (50%) and slightly fewer participants were able to identify the axis parallel gradient (39%). On average, participants rated the difficulty of the gradients task higher than the partitions task (5.98 vs 5.04), whereas indicated confidence was lower (5.58 vs 6.46). Figure 7a shows that participants were not able to accurately determine the position of the smallest region (R1) in the regions task. However, with an increased size of the region in R2 and R3 (Figures 7b, 7c), the predictions get increasingly accurate. On average, participants slightly underestimated the size of the region in R2 (−13%) and R3 (−4%).

The **qualitative** feedback from the participants gives additional insights into why especially the partitions and the gradients task were perceived so difficult and had such low accuracies. Most participants (67%) pointed out at some point during the study that they were irritated by the lower elasticity of the cubes along the z-axis. A number of participants (28%) also remarked that the surface of the cubes was uneven or that the sides, bottom, and top of the cubes were very stiff. Some participants (22%) also directly referenced that the region in R1 was too small or too hard to find. One additional participant (6%) suspected there was no region of higher density in R1 and gave an invalid answer. Four participants (22%) directly mentioned not being used to tactile perception before the experiment. A few other participants (11%) also mentioned that they felt fatigued during the experiment, but did not quit the experiment. Finally, the majority of participants (67%) found the experience “cool”, “fun”, “pleasant”, or “interesting”.

C. Mechanical Evaluation

The next step in our evaluation is to assess the mechanical performance of the squishicalizations through uniaxial loading tests [26], [32]. These tests aim to understand how the samples respond to compression forces and measure their deformation behavior under such conditions (Figure 8). As the testing geometry, we opt for a metallic cylinder with an 80mm diameter to ensure consistent force application over the whole sample. The geometry is lowered towards the sample at a constant velocity of 0.2mm/s. The sample is compressed until 25mm of displacement or until reaching 100N of force — whichever happens first. We test and measure each sample along all primary axes to understand how the sample deforms in different directions. This aims to provide insights into the mechanical behavior of the samples, which can inform their suitability for various applications or provide feedback for further design improvements. We show the force–displacement perpendicular to the printing direction observed during the loading tests in Figure 9.

Comparing force–displacement curves to evaluate elasticity is highly challenging. To enable quantitative analysis of the measured data we employ the perceptual model for elasticity proposed by Piovarci et al. [32]. The model proposes a one-dimensional space of elasticity ordering measured samples

from softest to stiffest. Furthermore, the perceptual space can be anchored by selecting a reference object. This allows us to express the difference in perceived stiffness between the reference and other objects as so-called Just Noticeable Differences (JNDs). When two samples are 1 or more JND apart they are easy to perceptually distinguish by human observers. Below 1 JND the samples get progressively harder and are generally considered indistinguishable at 0.5 JND. Applying the model to our samples anchored on the softest material reveals that the span between the softest and hardest material covers a total of 287 JNDs. The difference in elasticity in the two directions perpendicular to the printing direction is challenging to distinguish. However, the difference in stiffness along the printing axis reveals a significant increase in stiffness on average by 21 JNDs. This change is caused by both the different behavior along the printing axis and the outer wall of the cube that introduced significant buckling into the samples (see inset). Lastly, to visualize the range of achievable material properties with our system we compare them against the measurements of human tissues gathered by Guimaraes et al. [12], as depicted in Figure 10. We observe that the manufactured samples span a wide range of mechanical properties ranging from almost mucus-like to cartilage-like tissues. We discuss the implications of our mechanical analysis in Section VI.

D. Follow-Up Study

With the results of our perceptual and mechanical studies in hand, we set out to further support the feasibility of our method. This time we focused on different aspects of tactile perception in addition to improving the fidelity of the props according to our prior findings.

Participants. We again recruited 18 participants, 10 of which had participated in the previous study. The participants were between 24 and 42 years old (Mean=30.2, Std=5.29). Three hold a Ph.D., 13 a M.Sc., and two a B.Sc. 12 of the participants identify as male and 6 as female.

Tasks. The tasks were designed around *sorting objects according to their stiffness* to determine the ability to recognize different levels of elasticity, *recognizing compositions of different complexity* to determine if one can perceive how many different parts of varying elasticity an object is made up of, and *recognizing elasticity variations in objects* to determine the perceptibility of elasticity variations in nested objects. In the **sorting** task, participants had to sort three spheres based on their stiffness. All three spheres were handed out at once and completion time was measured. For the **counting** task, participants were presented with one homogeneous sphere, one consisting of two nested regions, and one consisting of three nested regions (Figure 11). Again, we handed the participants all three spheres at once and asked them to assign each sphere to a composition (1, 2, or 3 regions), with each composition occurring exactly once. In the **assignment** task, we handed the participants spheres, each consisting of three distinct regions, and let them determine their composition. Each region within the sphere had a different level of stiffness (“soft”, “hard”, and “medium”; similar to Figure 11c). The task was divided into two phases, each employing three spheres. In the first




Fig. 8. Assessing the mechanical performance of a squishicalization through uniaxial loading.




Fig. 9. Force-displacement measurements for the manufactured squishicalizations. The samples were fabricated with different seed point counts and a fixed wall width of 0.3mm . For comparison, we manufactured one sample in 0.2 and 0.4mm wall thickness (in the last column).




Fig. 10. Perceived stiffness range covered by real human tissues superimposed with the range of our squishicalizations (in blue). The perceptual space is anchored with mucus as the reference. One unit of perceived stiffness corresponds to one Just Noticeable Difference (JND).

phase, we provided the spheres' orientations to assess whether this information aids in determining the sphere composition. In the second phase, participants had to identify the alignment without any orientation cues. Lastly, we handed the participants one of our three prototype squishicalizations at a time and encouraged them to **exploration** freely, sharing their observations.

Datasets. In our first study, the outer layer of the cubes, especially on the top and bottom was perceived as confusing by many participants and the wall thickness used to create the props led to a noticeable difference in elasticity along the print axis. To combat this, we designed spherical props with a 75mm diameter, which do not have vertical or horizontal surfaces that could influence perception like cubes do. Additionally, we opted to use a wall thickness of 0.2mm as opposed to 0.3mm to reduce the anisotropic effects along the print axis. We used three different datasets, composed of one single homogeneous

region (Figure 11a), two nested regions (Figure 11b), or three nested regions (Figure 11c) with varying stiffness. We created multiple physical representations of each dataset by varying the compositions in terms of elasticity. For the homogeneous dataset, we created three versions with increasing stiffness. For the two regions, we created both permutations of the dataset: soft outside and hard inside, and vice versa. For the three regions, we created representations of all six permutations of combinations of three different levels of stiffness.

Procedure. Participants completed the tasks in the order sorting – counting – assignment (unlabelled) – assignment (labeled) – free exploration. We again measured **time**, perceived **difficulty**, and **confidence**, analogously to the first study. Again, we let our participants freely interact with our study props. The data for each task was collected using a questionnaire. For the **sorting** and **counting** tasks, participants were able to compare three spheres at once and had to assign each sphere to exactly one of three possible answers. The props in the **counting** task were assigned pseudo-randomly for each participant, alternating different combinations of stiffnesses in the two-region and three-region props. For the **assignment** tasks, the participants were handed one sphere at a time according to a balanced, pseudo-randomized scheme. The first three spheres were labeled using a red dot on an indicated position, and then the task was repeated with three different spheres without labels. Finally, we handed the participants the **foot**, **head**, and **fibers** prototypes and briefly explained how they were created. For each prototype, the participants shared their thoughts and findings during the interaction. Lastly, we asked people who participated in the previous study about how their experiences differed in comparison. After each task, we asked the participants for qualitative feedback.

Study Results. For the **sorting**, **counting**, and **assignment** tasks, we report the percentage of totally correct, partially correct, and completely wrong sorting results. For the **assignment** task, we also list these, but calculated per subtask (labeled, unlabeled). We also report time taken, perceived difficulty, and confidence per task. The detailed results are shown in Table III. For the **exploration** task, we code individual insights and list the number of occurrences. We also




Fig. 11. Datasets used for the follow-up study, upper half removed to illustrate composition. Different colors represent different scalar values. (a) Single homogeneous region (sorting and counting task), (b) Two nested regions (counting task), (c) Three nested regions (counting and assignment task).

provide a summary of the received qualitative feedback. The **quantitative** results confirm that all participants managed to **sort** homogeneous squishicalizations correctly, in order of stiffness. In the **counting** task, most participants (56%) were able to identify the number of regions with varying stiffness within the spheres. Five participants confused the sphere with two regions with the sphere with three regions, while two confused the homogeneous sphere with the sphere with two regions. The influence of the labeling of the orientation in the **assignment** task was noticeable, with 78% totally correct assignments in the labeled condition and 43% in the unlabeled condition. In the labeled condition, the innermost region of the props was identified correctly in 77.78% of cases, while the outer and middle regions were identified correctly more often (87.03% and 90.47% times, respectively). For the unlabeled condition, we observe that participants identified the innermost region correctly in 74.07% of the cases, while the outer and middle regions were only identified in 46.3% of the cases.

Our collected **qualitative** data sheds further light on different aspects of the follow-up study. In the free **exploration** task, we counted individual insights per prop. For the foot prototype, 15 participants reported being able to feel the bones inside the structure. 12 participants remarked positively on the realistic feel in terms of stiffness compared to a real foot. We also collected 7 comments on two factors that were perceived as unrealistic: the missing nails and the unnatural surface texture. While the head prototype was supposed to illustrate the position and structure of the brain, only two participants reported feeling a difference between white and grey matter. However, 10 participants noticed a large cavity behind the ear, which is not seen from the outside, and 8 participants found it interesting to explore cavities, such as the nasal cavity and the larynx. In the fibers dataset, 13 participants remarked upon the increased stiffness along the direction of the fibers. Additionally, 4 participants found a difference in elasticity along two axes in the direction parallel to the fibers. 5 participants commented on the visible aspects of the composition of the material's microstructure, such as fiber alignment.

The **sorting** task was considered "easy" or "trivial" by seven participants, but three comments indicate that it is harder to tell apart the medium sphere from the hard than the soft one. For the **counting** task, 8 participants indicated difficulty in identifying a difference between two and three regions. The feedback on the **assignment (labeled)** task shows that four participants found it difficult to identify the inner region and three had trouble identifying region borders. In contrast, for the **assignment (unlabelled)**, five participants found the outer layers harder to determine—with five comments explicitly referring to the missing labeling. Interestingly, three participants found the inner region easier to determine. In **comparison to the first study**, 8 of the 10 returning participants confirmed that the z-axis anisotropy was less noticeable with the new props. In all of the qualitative data, only one statement refers to this effect being noticeable.

TABLE III

QUANTITATIVE STUDY RESULTS FOR THE FOLLOW-UP STUDY. TIMES IN *m:ss*. ACCURACY IS REPORTED AS THE NUMBER OF RESULTS WITH CORRECT SORTINGS (ALL), THE NUMBER OF RESULTS WITH PARTIALLY CORRECT SORTINGS (PART), AND THE NUMBER OF COMPLETELY WRONG SORTING (NONE). FOR THE SORTING (S) AND COUNTING (C) TASKS, THE TOTAL NUMBER OF TRIALS IS 18 (ONE PER PARTICIPANT); FOR THE ASSIGNMENT TASKS, LABELED (A*) AND UNLABELED (A), THE TOTAL NUMBER OF TRIALS IS 54 (THREE PER PARTICIPANT).

Task	Accuracy			Time		Difficulty		Confidence	
	All	Part	None	avg	std	avg	std	avg	std
S	18	0	0	0:14	0:06	1.28	0.57	9.39	1.54
C	10	7	1	01:19	1:05	5.89	1.88	5.61	2.20
A*	42	12	0	02:24	1:03	5.72	2.16	5.67	1.91
A	23	21	10	02:56	1:12	6.78	2.26	4.72	2.22

V. EXPERT INTERVIEWS

To determine potential future usage scenarios for *Squishicalization*, we reached out to domain experts in the fields of extended reality (XR), materials science, and medical education. Every session took 30–60 minutes. The focus was put specifically on a potential application of our proposed workflow as opposed to the general use of data physicalization. We first presented the concept to the experts and then held semi-structured interviews targeting the following questions: *What kind of data do you think squishicalizations are useful for? For what kind of users are Squishicalizations suited? Which tasks could be performed using squishicalizations? Where do you see weaknesses in squishicalizations in their current form?* In the interview summaries below, we refer to the domain experts by pseudonyms based on their profession and using their actual pronouns.

Extended Reality. We interviewed on-site a senior XR researcher from our faculty, Dr. X. We brought a prototype version of the foot squishicalization (Figure 5a) to explain the concept and let Dr. X freely interact with it. The discussion focused on the simulation of an immersive and realistic environment. As such, the data used in this hypothetical usage scenario would be largely synthetic. The potential users indicated by Dr. X are professionals in training, as well as XR researchers and consumers. Dr. X identified potential applications for our squishicalizations in **mobile haptics** for VR applications. Here, users enter a VR environment while tactile feedback for specific objects in the virtual space is provided using real-life props. The props are handled by a robot that simulates the behavior of the virtual object. Dr. X added that even though squishicalizations cannot replicate the weight of represented objects accurately, a robot could simulate this by adding force. Furthermore, he indicated possible use in **surgical simulations**. Since squishicalizations can be replicated relatively easily and cheaply using 3D printing, they could be used as destructible props. To provide accurate feedback for such applications, it would be necessary to replicate tactile tissue properties accurately, while a head-mounted display would provide realistic visuals. Lastly, Dr. X proposed that, augmented with pressure sensors, our squishicalizations could serve as **provenance** tools to analyze user interactions. In combination with VR applications, they could provide insights into how people interact with objects in a virtual space.

Materials Science. We subsequently reached out to a visualization expert in the field of materials science, M. This interview was held virtually as the expert is situated in a different city, but the props had been mailed to him in advance. M described the data used in the hypothetical usage scenario as industrial CT scans ranging up to $2000 \times 2000 \times 2000$ voxels. Our method would be suitable for foam-like structures or fiber-reinforced composites. The user group of the applications discussed with M includes material scientists and researchers. As a first possible application, M mentioned **immersive analytics**. Current approaches in materials science visualization use VR in an attempt to enhance 3D spatial perception. M mentioned that physicalization could represent material properties with tactile feedback. Other application scenarios could include the **analysis of fabrication characteristics** of different materials, such as pores and voids, or **defect analysis** and the **temporal analysis of material properties**.

After our initial interview, M provided us with a CT scan from a fiber-reinforced material (the squishicalization of which is shown in Figure 5c) and recommended a **follow-up interview** with a fiber-reinforced composite material expert, F. Subsequently, we generated, printed, and shipped the specimen to the two experts for a second virtual interview with both. F commented on the accurate representation of the isotropy in the fiber material. The way the sample cannot be compressed along the fiber direction resembles how such materials may behave in reality, despite the individual fibers being noticeably more elastic than in reality. F found that the compression behavior of the sample in the horizontal plane showed some anisotropy, which could provide further insights into the fiber microstructure. F also mentioned other kinds of **compound materials** that could be interesting to analyze using our approach, such as composite laminates and multi-material components. Both experts indicated limitations for applications in their domain. According to M, some datasets have very small anomalies, which can be undetected or lost during our sampling process. F also mentioned that a more accurate parametrization of the process, for example, based on additional local mechanical properties caused by the microstructure, would help create a more realistic behavior, while larger samples could represent the characteristics more faithfully. To increase realism, they also discussed the possibility of multi-material printing.

Medical Education. We interviewed also Dr. W, a University Professor and specialist in anatomy and anatomical education. He extended the invitation to two medical doctors in training, Dr. K and Dr. A. The interview was held in person in their office and we brought prototypes of the foot dataset (Figure 5a) for them to examine. Dr. W mentioned tangible representation of **3D histology** as the first possible application. He described the idea of creating enlarged squishicalizations of cellular structures. Using microscopy data, our method could create a “*tangible cell*”, which would be focused on K12 students rather than doctors in training. Additionally, Dr. W inquired about the ability of our method to **display large cavities** in human bodies, in a way that their proportion and position are understandable. He also proposed creating waterproof models of these cavities to **illustrate blood flow** through cardiac

volume. As doctors in training, Dr. K and Dr. A identified a potential for creating **tangible representations of rare conditions** for anatomical education. Doctors use palpation to detect anomalies in tissue using their sense of touch. Some of these conditions do not appear regularly in cadavers and are exemplified through prepared specimens. Dr. A explained that the tactile feel of prepared specimens differs from live examples, while Dr. K added that such tactile stimuli can not be taught in books. Dr. W pointed out the importance of accurate tissue reproduction for the method to be usable for professionals in training. He stated an accurate mapping from tissue type to elasticity would be needed. Additionally, Dr. K pointed out that the surface properties of the presented squishicalizations were unrealistic.

VI. TAKEAWAYS AND FUTURE DIRECTIONS

Data physicalization research leverages the tactile properties of physical sculptures. As opposed to other recent works that use surface properties of 3D printed objects to express characteristics in data [8], [29], [31] or other fabrication solutions with [2], [42] or without sensing capabilities [40], with *Squishicalization*, **we express the characteristics of data in physicalizations beyond the surface**. We encode scalar values as local densities in printable artifacts to add a tactile dimension to the generated data physicalizations. Moreover, our *Squishicalization* pipeline is **designed for use with any consumer-grade 3D printing that can make use of elastic materials**. All tools and materials are available and affordable for home use, and no manual labor is needed for assembly.

Performance. Compared to other microstructure-based approaches [27], [36], *Squishicalization* is **applicable to volumetric data without prior segmentation**. Similar to volume rendering, we employ piecewise linear DTFs to map local density to scalar intervals. In terms of performance, Martinez et al. [27] name extraction of region borders as the biggest bottleneck in their Voronoi tessellation-based approach. In *Squishicalization*, we eliminate the performance impact of the tessellation by using labeled volume isosurface extraction [10]. Parallelization of the region extraction part on the GPU and restricting tessellation to regions of interest increases our performance, as we showed in Section IV-A. Compared to Lu et al. [25], who demonstrate their approach with just 100 points, our method scales well with a much higher number of seed points (see Table I).

Parameter Space Exploration. *Squishicalization* has a vast parameter space to adjust for material and fabrication-related variations. In Section IV-C, we document the results of our mechanical tests on samples with different densities and variations in wall thickness. The axial load tests show that **our method can reproduce a wide range of elasticity**. We also show that the samples have near isotropic behavior in the horizontal plane. As highlighted later by some experts in our interviews (Section V), a more accurate mapping between scalar values and elasticity is needed to accurately reproduce desired behavior. Currently, we only use voxel distance mapping scalar values to construct our DTFs. **A first step to more accurate representations could be to incorporate**

voxel spacing and point densities into DTFs. Determining the precise relationship between sample density and elasticity for a given wall thickness and material is essential for this. Our mechanical and perceptual evaluations point towards a nonlinear relationship between these factors. Combined with traditional dart-throwing Poisson-disk sampling approaches [7], precision can increase at the expense of control over the currently user-selected amount of samples.

Suitability for Selected Applications. Our results show that our prototypes could stimulate the tactile sense of our participants to determine the size and positions of medium and large-sized heterogeneous regions, as well as nested components. We also show that the elastic properties of human tissue can be emulated realistically through our mechanical evaluation. This indicates that our technique could be, for instance, useful as a **training tool for palpitation** of rare conditions for medical professionals and demonstration of self-palpitation of irregularities for laypeople. Also, our interviews with material scientists revealed that our method could be useful to **analyze and illustrate material characteristics** of compound materials. Extending the method to provide sensing capabilities, similar to Bae et al. [2] and Tejada et al. [42], could further support the analysis of users' **haptic interactions**. Tangible encoding of **uncertainty** in scalar data could be encoded in stiffness, e.g., with harder structures indicating higher certainty, and softer structures the opposite.

Oppositely, as study participants struggled to determine the directions of gradients, we have indications that *squishicalizations* might **not be suitable** for identifying the structure of more complex mathematical formulations, such as probability distributions. Our technique produces rather rough surfaces, which could confuse sightless people in **accessibility scenarios**, which could be improved in the future with an additional sensory channel. Finally, several participants stated that using tactile perception is **unusual**. During our study, participants were sometimes concerned they could destroy the props, or felt inhibited in applying too much force. This indicates that "hidden" direct interactivity, such as exploring the sub-surface properties produced by *Squishicalization*, needs to be encouraged to be fully leveraged.

Refinements. The initial perceptual study and the mechanical evaluation used samples with a wall thickness of 0.3mm. This way we could produce a large number of samples in reduced time and with moderate material use. However, the axial load tests indicate that this wall strength increases the stiffness of the samples in the printing direction. The low vertical resolution of these prints leads to disconnected staircase artifacts between the layers, resulting from the slicing method. Combined with the additional walls, bottom, and top cover of the cubes used in our user study, we see a large impact on tactile perception, reflected by the qualitative comments of our participants and also their low performance. In our follow-up study, we reduced the wall thickness of the props and used a spherical shape without stiff outer surfaces, guided by the mechanical evaluation results. **Using wall thicknesses around 0.2mm yielded more accurate results and led to less noticeable anisotropy.** This only requires an increase of sample points to reach comparable levels of stiffness. Staircase

artifacts could be reduced by increasing the vertical resolution of the prints, or by refining the slicing to allow for horizontal walls. Also, conical slicing [45], as opposed to vertical slicing, could be used to create more complex squishicalizations, even with overhanging structures.

ACKNOWLEDGMENTS

The authors acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme.

REFERENCES

- [1] K. D. Ang, F. F. Samavati, S. Sabokrohiyeh, J. Garcia, and M. S. Elbaz. Physicalizing cardiac blood flow data via 3D printing. *Computers & Graphics*, 85:42–54, Dec. 2019. doi: 10.1016/j.cag.2019.09.004
- [2] S. S. Bae, T. Fujiwara, A. Ynnerman, E. Y.-L. Do, M. L. Rivera, and D. A. Szafrir. A Computational Design Pipeline to Fabricate Sensing Network Physicalizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):913–923, Jan. 2024. doi: 10.1109/TVCG.2023.3327198
- [3] S. S. Bae, C. Zheng, M. E. West, E. Y.-L. Do, S. Huron, and D. A. Szafrir. Making Data Tangible: A Cross-disciplinary Design Space for Data Physicalization. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–18. New Orleans LA USA, Apr. 2022. doi: 10.1145/3491102.3501939
- [4] E. Barchiesi, M. Spagnuolo, and L. Placidi. Mechanical metamaterials: a state of the art. *Mathematics and Mechanics of Solids*, 24(1):212–234, Jan. 2019. doi: 10.1177/1081286517735695
- [5] B. Bickel, M. Bäcker, M. A. Otaduy, H. R. Lee, H. Pfister, M. Gross, and W. Matusik. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.*, 29(4):63:1–63:10, July 2010. doi: 10.1145/1778765.1778800
- [6] S. Brewster, F. Chohan, and L. Brown. Tactile feedback for mobile interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pp. 159–162. New York, NY, USA, Apr. 2007. doi: 10.1145/1240624.1240649
- [7] R. Bridson. Fast Poisson disk sampling in arbitrary dimensions. *SIGGRAPH sketches*, 10(1):1, 2007.
- [8] H. Djavaherpour, L. Moorman, F. Samavati, and Y. Jansen. First Insights Into INTUIT: An Interactive Tactile Physicalization for User Interpretation of RADAR Technology. *IEEE Computer Graphics and Applications*, 43(5):91–98, Sept. 2023. doi: 10.1109/MCG.2023.3286228
- [9] H. Djavaherpour, F. Samavati, A. Mahdavi-Amiri, F. Yazdanbakhsh, S. Huron, R. Levy, Y. Jansen, and L. Oehlberg. Data to Physicalization: A Survey of the Physical Rendering Process. *Computer Graphics Forum*, 40, Feb. 2021.
- [10] S. F. Frisken. SurfaceNets for Multi-Label Segmentations with Preservation of Sharp Boundaries. *The Journal of computer graphics techniques*, 11(1):34–54, 2022.
- [11] J. Garcia, Z. Yang, R. Mongrain, R. L. Leask, and K. Lachapelle. 3D printing materials and their use in medical education: a review of current technology and trends for the future. *BMJ Simulation & Technology Enhanced Learning*, 4(1):27–40, Dec. 2017. doi: 10.1136/bmjstel-2017-000234
- [12] C. F. Guimarães, L. Gasperini, A. P. Marques, and R. L. Reis. The stiffness of living tissues and its implications for tissue engineering. *Nature Reviews Materials*, 5(5), May 2020. doi: 10.1038/s41578-019-0169-1
- [13] T. Hogan and E. Hornecker. Towards a Design Space for Multisensory Data Representation. *Interacting with Computers*, p. iwc:iww015v1, May 2016. doi: 10.1093/iwc/iww015
- [14] J. Hurtienne, F. Maas, A. Carolus, D. Reinhardt, C. Baur, and C. Wienrich. Move and Find: The Value of Kinaesthetic Experience in a Casual Data Representation. *IEEE Computer Graphics and Applications*, 40(6):61–75, Nov. 2020. doi: 10.1109/MCG.2020.3025385
- [15] A. Ion, R. Kovacs, O. S. Schneider, P. Lopes, and P. Baudisch. Metamaterial Textures. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–12. Montreal QC Canada, Apr. 2018. doi: 10.1145/3173574.3173910
- [16] Y. Jansen and P. Dragicevic. An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405, 2013. doi: 10.1109/TVCG.2013.134

- [17] Y. Jansen, P. Dragicevic, and J. D. Fekete. Evaluating the efficiency of physical visualizations. *Conference on Human Factors in Computing Systems - Proceedings*, pp. 2593–2602, 2013. doi: 10.1145/2470654.2481359
- [18] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and Challenges for Data Physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pp. 3227–3236. New York, NY, USA, Apr. 2015. doi: 10.1145/2702123.2702180
- [19] M. Karyda, D. Wilde, and M. G. Kjærsgaard. Narrative Physicalization: Supporting Interactive Engagement With Personal Data. *IEEE Computer Graphics and Applications*, 41(1):74–86, Jan. 2021. doi: 10.1109/MCG.2020.3025078
- [20] R. A. Khot, J. Andres, J. Lai, J. Von Kaenel, and F. Mueller. Fantibles: Capturing cricket fan's story in 3D. *DIS 2016 - Proceedings of the 2016 ACM Conference on Designing Interactive Systems: Fuse*, pp. 883–894, June 2016. doi: 10.1145/2901790.2901886
- [21] R. A. Khot, L. Hjorth, and F. F. Mueller. Understanding physical activity through 3D printed material artifacts. *Conference on Human Factors in Computing Systems - Proceedings*, pp. 3835–3844, 2014. doi: 10.1145/2556288.2557144
- [22] H. Kim, A. Everitt, C. Tejada, M. Zhong, and D. Ashbrook. MorphesPlug: A Toolkit for Prototyping Shape-Changing Interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*, pp. 1–13. New York, NY, USA, 2021. doi: 10.1145/3411764.3445786
- [23] P. Ljung, J. Krüger, E. Groller, M. Hadwiger, C. D. Hansen, and A. Ynnerman. State of the Art in Transfer Functions for Direct Volume Rendering. *Computer Graphics Forum*, 35(3):669–691, 2016. doi: 10.1111/cgf.12934
- [24] W. E. Lorensen and H. E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. In *Seminal graphics*, pp. 347–353. New York, NY, USA, July 1998. doi: 10.1145/280811.281026
- [25] L. Lu, A. Sharf, H. Zhao, Y. Wei, Q. Fan, X. Chen, Y. Savoye, C. Tu, D. Cohen-Or, and B. Chen. Build-to-last: strength to weight 3D printed objects. *ACM Trans. Graph.*, 33(4):97:1–97:10, July 2014. doi: 10.1145/2601097.2601168
- [26] J. Martínez, J. Dumas, and S. Lefebvre. Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics*, 35(4):1–12, July 2016. doi: 10.1145/2897824.2925922
- [27] J. Martínez, S. Hornus, H. Song, and S. Lefebvre. Polyhedral voronoi diagrams for additive manufacturing. *ACM Transactions on Graphics*, 37(4):1–15, Aug. 2018. doi: 10.1145/3197517.3201343
- [28] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.
- [29] M. J. McGuffin, R. Servera, and M. Forest. Path Tracing in 2D, 3D, and Physicalized Networks. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–14, 2023. doi: 10.1109/TVCG.2023.3238989
- [30] D. Pahr, H.-Y. Wu, and R. G. Raidou. Vologram: An Educational Holographic Sculpture for Volumetric Medical Data Physicalization. In *VCBM 2021: 11th Eurographics Workshop on Visual Computing for Biology and Medicine, Paris, France, 22-24 September 2021*, pp. 19–23, 2021. doi: 10.2312/vcbm.20211341
- [31] L. J. Perovich, B. Rogowitz, V. Crabb, J. Vogelsang, S. Hartleben, and D. Offenhuber. The tactile dimension: a method for physicalizing touch behaviors. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23*, pp. 1–15. New York, NY, USA, Apr. 2023. doi: 10.1145/3544548.3581137
- [32] M. Piovarči, D. I. W. Levin, J. Rebello, D. Chen, R. Đuriković, H. Pfister, W. Matusik, and P. Didyk. An interaction-aware, perceptual model for non-linear elastic objects. *ACM Transactions on Graphics*, 35(4):1–13, July 2016. doi: 10.1145/2897824.2925885
- [33] R. G. Raidou, M. E. Gröller, and H. Y. Wu. Slice and Dice: A Physicalization Workflow for Anatomical Edutainment. *Computer Graphics Forum*, 39(7):623–634, Oct. 2020. doi: 10.1111/cgf.14173
- [34] M. Schindler, T. Korpitsch, R. G. Raidou, and H.-Y. Wu. Nested Papercrafts for Anatomical and Biological Edutainment. *Computer Graphics Forum*, 41(3):541–553, June 2022. doi: 10.1111/cgf.14561
- [35] M. Schindler, H. Y. Wu, and R. G. Raidou. The Anatomical Edutainer. In *Proceedings - 2020 IEEE Visualization Conference, VIS 2020*, pp. 1–5, Oct. 2020. doi: 10.1109/VIS47514.2020.00007
- [36] C. Schumacher, B. Bickel, J. Rys, S. Marschner, C. Daraio, and M. Gross. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics*, 34(4):136:1–136:13, July 2015. doi: 10.1145/2766926
- [37] P. Sitthi-Amorn, J. E. Ramos, Y. Wangy, J. Kwan, J. Lan, W. Wang, and W. Matusik. MultiFab: a machine vision assisted platform for multi-material 3D printing. *ACM Transactions on Graphics*, 34(4):1–11, July 2015. doi: 10.1145/2766962
- [38] S. Stusak, A. Tabard, F. Sauka, R. A. Khot, and A. Butz. Activity sculptures: Exploring the impact of physical visualizations on running activity. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2201–2210, Dec. 2014. doi: 10.1109/TVCG.2014.2352953
- [39] C. Sullivan and A. Kaszynski. PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *Journal of Open Source Software*, 4(37):1450, May 2019. doi: 10.21105/joss.01450
- [40] S. Swaminathan, C. Shi, Y. Jansen, P. Dragicevic, L. Oehlberg, and J.-D. Fekete. Creating physical visualizations with makervis. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems, CHI EA '14*, pp. 543–546. New York, NY, USA, Apr. 2014. doi: 10.1145/2559206.2574788
- [41] F. Taher, Y. Jansen, J. Woodruff, J. Hardy, K. Hornbæk, and J. Alexander. Investigating the Use of a Dynamic Physical Bar Chart for Data Exploration and Presentation. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):451–460, Jan. 2017. doi: 10.1109/TVCG.2016.2598498
- [42] C. E. Tejada, R. Ramakers, S. Boring, and D. Ashbrook. AirTouch: 3D-printed Touch-Sensitive Objects Using Pneumatic Sensing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, pp. 1–10. New York, NY, USA, Apr. 2020. doi: 10.1145/3313831.3376136
- [43] C. Torres, T. Campbell, N. Kumar, and E. Paulos. HapticPrint: Designing Feel Aesthetics for Digital Fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology, UIST '15*, pp. 583–591. New York, NY, USA, Nov. 2015. doi: 10.1145/2807442.2807492
- [44] T. Tricard, V. Tavernier, C. Zanni, J. Martínez, P.-A. Hugron, F. Neyret, and S. Lefebvre. Freely orientable microstructures for designing deformable 3D prints. *ACM Transactions on Graphics*, 39(6):1–16, 2020.
- [45] M. Wüthrich, M. Gubser, W. J. Elspass, and C. Jaeger. A Novel Slicing Strategy to Print Overhangs without Support Material. *Applied Sciences*, 11(18), Jan. 2021. doi: 10.3390/app11188760
- [46] C. Yuksel. Sample Elimination for Generating Poisson Disk Sample Sets. *Computer Graphics Forum*, 34(2):25–32, 2015. doi: 10.1111/cgf.12538
- [47] B. Zhu, M. Skouras, D. Chen, and W. Matusik. Two-Scale Topology Optimization with Microstructures. *ACM Transactions on Graphics*, 36(4):1, Aug. 2017. doi: 10.1145/3072959.3095815

BIOGRAPHY SECTION



Daniel Pahr is a PhD Student at TU Wien Austria. He completed his Master's degree in Medical Informatics at TU Wien in 2021, focusing on the physicalization of medical data. His current research revolves around investigating the interactive capabilities of data physicalizations.



Michal Piovarči is a postdoctoral researcher at ETH Zürich. He received his doctoral degree at USI Lugano under the supervision of prof. Piotr Didyk. His thesis entitled Perception-Aware Fabrication was awarded the Eurographics 2021 PhD award. Currently, Michal continues his work at the intersection of computational fabrication, perception, and machine learning.



Hsiang-Yun Wu is a senior scientist and lecturer at St. Pölten University of Applied Sciences, Austria, and TU Wien, Austria. She received her Ph.D. degree in 2013 from The University of Tokyo, Japan. Her principal research interests cover computational approaches for network analytics and visualization, data physicalization, and human-centered techniques. More information about Dr. Wu can be found at <http://yun-vis.net/>.



Renata G. Raidou (<http://www.renataraidou.com>) is Assistant Professor at TU Wien, Austria. She is the holder of the 1st Dirk Bartz Prize for Visual Computing in Medicine 2018, the EuroVis Best PhD Award in 2018, and the EuroVis Young Researcher Award in 2022. She works at the interface between Visual Analytics, Image Processing, and Machine Learning, with a strong focus on medical applications.