

High-throughput feature extraction for measuring attributes of deforming open-cell foams

Steve Petruzza, Attila Gyulassy, Samuel Leventhal, John J. Baglino, Michael Czabaj, Ashley D. Spear, Valerio Pascucci




Fig. 1: Tracking of ligaments and junctions in an aluminum open-cell foam during incremental compressive loading from 0.0 to 2.0 mm. Red boxes highlight the changes in connectivity that we detect when ligaments fracture. The yellow boxes highlight part of the material drifting outside the field of view of during CT scans. Colors are mapped to the ids of ligaments and junctions.

Abstract—Metallic open-cell foams are promising structural materials with applications in multifunctional systems such as biomedical implants, energy absorbers in impact, noise mitigation, and batteries. There is a high demand for means to understand and correlate the design space of material performance metrics to the material structure in terms of attributes such as density, ligament and node properties, void sizes, and alignments. Currently, X-ray Computed Tomography (CT) scans of these materials are segmented either manually or with skeletonization approaches that may not accurately model the variety of shapes present in nodes and ligaments, especially irregularities that arise from manufacturing, image artifacts, or deterioration due to compression. In this paper, we present a new workflow for analysis of open-cell foams that combines a new density measurement to identify nodal structures, and topological approaches to identify ligament structures between them. Additionally, we provide automated measurement of foam properties. We demonstrate stable extraction of features and time-tracking in an image sequence of a foam being compressed. Our approach allows researchers to study larger and more complex foams than could previously be segmented only manually, and enables the high-throughput analysis needed to predict future foam performance.

Index Terms—Topological analysis, foam, features extraction, feature tracking

1 INTRODUCTION

Open-cell metallic foams and lattices are comprised of a network of interconnected ligaments. It is postulated that attributes such as the length of ligaments, connectedness of junctions, or distribution of pore size contribute to the mechanical performance of a foam. However, the relationships among these attributes and specific performance are not well understood. As foams are gaining more interest due to the plethora of manufacturing techniques, and simulation is increasingly able to reproduce performance characteristics, a data-driven approach becomes feasible to understand these relationships. However, high-throughput analysis of metallic foams is needed to populate this high-dimensional data space to analyze.

The state-of-the-art for analysis of metallic foams remains performing segmentation of features by hand, a time-consuming, laborious process that does not scale to the quantity of data sets, data sizes, and morphological complexities that need to be analyzed. Standard

segmentation techniques struggle both to address the variety of structures present in open-cell foams, and to provide robust segmentation as the material is compressed to measure its performance characteristics. Drawing correspondences between junctions and ligaments becomes especially challenging as ligaments deform, self-contact, or fracture during compressive loading. A new approach is needed to address these challenges.

In this paper, we present a high-throughput, end-to-end workflow that provides robust and reliable segmentation of foams into junctions and ligaments. We introduce a new measure, the geodesic density function, that, combined with topological approaches, is able to identify junctions and ligaments consistently across time series images. We employ time tracking to disambiguate and explain the provenance of structures in the later time steps, when the foam has been compressed. Our system involves and guides the user in the analysis of a foam through careful use of visualization tools. Finally, we show that large-scale foams with thousands of ligaments, which were previously deemed too complex and costly to segment "by hand", can now be analyzed with only minutes of user interaction. Our workflow increases the throughput for the analysis of open-cell foams, and marks the first step toward building data-centric models of foam performance.

2 MOTIVATION

Open-cell metallic foams and lattices are a class of structural-material systems that comprise a network of interconnected ligaments, resulting in a hierarchical structure [25]. These low-density, light-weight, load-bearing structures have been used in various multifunctional applications [4, 5, 23]. For example, they are able to serve concurrently as electrodes for energy-storage devices [62], as hosts for newly gener-

• S. Petruzza, A. Gyulassy, S. Leventhal and V. Pascucci are with SCI Institute, University of Utah. E-mail: [spetruzza, jediati, samlev, pascucci]@sci.utah.edu.

• J. J. Baglino, M. Czabaj, A. D. Spear are with Dep. Mechanical Engineering, University of Utah. E-mail: [ashley.spear, m.czabaj, john.baglino]@uuh.edu.

Manuscript received 31 Mar. 2019; accepted 1 Aug. 2019.

Date of publication 16 Aug. 2019; date of current version 20 Oct. 2019.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2019.2934620




Fig. 2: Left: an as-manufactured investment-cast aluminum foam. Right: chamber of Varian BIR 150/130 X-ray CT imaging system with mechanical load frame in place and mechanical load frame (Images reproduced with permission from Matheson et al. [37]).




Fig. 3: The standard model of ligaments and junctions (a) is challenged by configurations that arise from defects (b,c), or from the effects of compression (d).

ated bone and blood vessels in biomedical implants [50], or as impact absorbers and noise insulators for advanced high-speed ground transportation [5].

Open-cell metallic foams and lattices (generically referred to as cellular metals) can be manufactured by a number of different routes, a taxonomy of which is described in Refs. [55] and [32]. The conventional method by which stochastic, open-cell foams are produced is investment casting [55]. More recently, metal-based additive manufacturing has been used to produce lattice structures (e.g., [39, 59, 60]), and tailoring the design of such structures remains an active area of research.

Metal foams and lattices have been manufactured using many types of alloys, including aluminum, titanium, copper, and steel. This work focuses specifically on the design of open-cell aluminum foam, which can serve as a proof-of-concept for other material systems. The complex geometry and resulting mechanical performance of open-cell metal foams and lattices present a high-dimensional design space. For stochastic open-cell foams, manipulating parameters of the investment casting process allows manufacturers to control the pore size (expressed by manufacturers in units of “pores per inch”) and relative density (which is governed by the shape and size of the ligament cross sections). For additively manufactured lattices, the geometry (including lattice spacing and topology) is controlled via the CAD model used to create the part. Discovering the process-structure-property relationships for a high-dimensional design space such as this requires generation of large amounts of data and observations, meriting the implementation and use of the high-throughput feature extraction tool described herein. In this work, we design and deploy a workflow for analyzing grayscale images of metallic foams that were generated from experimental and simulated data.

For experimentally derived data, the images are collected as follows. After manufacturing, each foam sample is mounted in a specialized in situ load frame designed to enable three-dimensional imaging during mechanical crushing. During the mechanical loading, X-ray Computed Tomography (CT) data are collected incrementally [37] as shown in Figure 2. Data collected from the experiment are then post-processed to create a stack of grayscale images representing slices of the foam volume taken at each time step during the mechanical test. Similar image stacks can be generated from numerical simulations. For example, finite-element models of foam volumes can be generated and virtually crushed, and stacks of two-dimensional image slices can be extracted at different time steps, analogously to those extracted from experiments. The foam models can be synthetically generated [2, 47, 57] or created directly from X-ray CT image data of an as-manufactured, uncrushed foam [36].

We postulate that the following features of the foam will impact, to various degrees, the mechanical performance, including the onset and progression of deformation and failure:

- ligaments*: length; orientation; cross section area; perimeter; and shape.
- junctions*: volume; surface area; aspect ratio.
- cells*: size; aspect ratio.
- connectivity*: of junctions and ligaments; and of pores.

However, it is unknown how these specific features interact to govern the mechanical response and whether the features can be rank-ordered by their importance on the performance of the foam. This knowledge will be crucial to enable design optimization of the foams to meet specific performance requirements. The first step in enabling a data-driven design paradigm is collecting and analyzing large amounts of data and relating these properties to measured material performance.

Prior state-of-the-art. The grayscale images acquired from CT scans over the course of a compression sequence, in the past, have been used to generate animations that show the progression of compression. These animations are visually inspected to classify the behaviour of each individual ligament at each step [37]. Ligaments have been classified into one of four categories: those that have demonstrated brittle fracture with little plastic deformation, those that have exhibited only plastic collapse (such as buckling), those that have experienced a significant amount of plastic deformation before eventually fracturing, and those that have remained intact or mostly intact. This manual inspection and classification process is very time consuming and certainly not scalable to large volumes or time series. For instance, even a small example (with 82 ligaments) took domain scientists several days to segment and track only a dozen of ligaments. Hence, there is a need for solutions that can accelerate the segmentation and tracking.

Challenges in segmenting CT foams. One could imagine that the basic problem in creating an embedded graph structure to represent an open-celled foam could be addressed with standard approaches in image processing, such as medial axis transforms, skeletonization, or even topological connectivity structures. However, the neat mental model of “junction” and “ligament” collapses when faced with the reality of time-series of images of a foam being compressed. First, defects in real materials mean there is a continuum of morphologies between something that can be called a ligament and something that can be called a junction, as illustrated in Figure 3. For instance, when a normal ligament becomes too short, the junctions merge to form larger junctions; there is no clear demarcation between a thick ligament and merged junctions. Also, entire faces of the open-cell foam can be filled with material: what is an appropriate skeletal representation of such a shape? A primary contribution of this work is a novel measure of localized density to identify regions that act as junctions. Finally, as the foam is compressed, ligaments break, and previously separated structures now appear to have a ligament between them, precluding techniques that ignore time history and provenance of structures. In this work, we present a set of techniques that reliably and repeatably define and compute a junction and ligament model, and we track those structures throughout the time sequence.

3 RELATED WORK

Skeletonization. For the purpose of analysis and visualization of volumes, a simplified representative structure of shape is important. The exemplary curve-skeleton of the larger object is often interpreted to be a line simplification of the full volume, tracing the object’s center [53]. Various approaches to achieve this simplified skeleton include topological thinning [14], distance field based methods [11, 15], and potential field based methods [19, 63]. Previous distance field approaches to obtain a skeletonization have been based either on penalizing a Dijkstra Shortest Path toward the boundary of the object from a core central line [11, 43], outward flux of gradient vectors from the center of the material [15], or computing the center line through level sets [31]. The interpretability of what defines a representative skeleton has led to numerous approaches of skeletonization as well as attempts at a formal definition [12, 19, 20]. Often skeletonization algorithms approach the problem by starting outside and working their way in, i.e., computing the center line of the full object through erosion [33],

thinning [7, 8, 13, 40, 42], and dilation from the object boundary [34]. A challenge when using these approaches to segment digital images is that the threshold determining the object interface completely determines the connectivity of the skeleton. There is no ability to estimate what components have been omitted. Another limitation of these approaches is that they do not relate the components of a skeleton to the void space outside the manifold of interest, for instance in connecting which portions of a skeleton surround a void or hole.

Of growing interest is the use of geodesics for constructing higher dimensional space summaries as lower dimensional path connections between points of interest or minimum spanning trees that convey significant structures within the higher dimensional shape. Previous approaches for determining these curvilinear structures have relied on computing distances from origin points using the Fast Marching Algorithm while minimizing an energy function or relevant measure [46] or iteratively moving from origin points to sink points through occlusion points of interest [18]. One such implementation used in neurite tracing determines geodesics as those propagated from the origin point that minimize a tubularity measure defined to be a spatial path as well as a curvilinear ‘thickness’ given by intersecting closed balls [6]. Performing image segmentation for holes, ends, and centers through mathematical morphology of geodesics has been of growing interest due to its success and intuitive physical meaning, allowing easy extension of geodesic morphological transformation techniques from two-dimensional image data into three-dimensional volumes or four-dimensional time dependent spaces [38]. Although skeletonization can be performed in geodesic space [9], our primary interest in it is a means of measuring properties from *inside* a topological space, specifically, the manifold pertaining to the metallic material.

Recently, approaches have begun to compute the curved skeleton through topological summaries such as Reeb graphs and merge trees based on persistence [16, 54]. The most successful methods employ the Morse-Smale complex, either directly on the function of interest, or on the signed distance field from a material interface. This approach has been used to construct the filamentary structure of the universe [51], represent bonds between atoms [10], trace lithium diffusion pathways [30], and extract the core structure of a porous material [29]. An advantage of topology-based techniques is they account for the impact of noise through persistence simplification [22]. We use the 1-skeleton of the Morse-Smale complex, as it enables reasoning about the stability of the extraction and provides a means of relating junctions and ligaments to the grains and open faces they surround.

Feature Tracking One central goal of our work is to track the skeletonized ligaments over the compression time series. Overlap between features in subsequent time steps is central to many approaches, and is used to categorize events such as continuation, creation, dissipation, bifurcation, and amalgamation [48, 49]. Such approaches rely on spatial overlap, and they have been extended to topological tracking of contours [3], for varying thresholds with Reeb graphs [21], space-time isosurfaces [56], or building a space-time function to track features defined on level sets [58]. However, compressed foams are imaged relatively sparsely in time, such that feature overlap is not reliable; with these techniques, moving junctions would instead be classified as appearing and disappearing. Basic nearest points have been used to track filamentary structures [29]; however, point correspondences were not used to infer object correspondences.

By providing a robust and concise feature description, skeletonization offers a new mean of feature tracking and matching. Through a hierarchical view of the computed spanning tree object, comparison becomes possible using isomorphic or nearly similar skeleton trees [17, 52]. This approach has also been employed in animation through computation of the centerlines, manipulation, and reconstruction of the original object [24]. Comparing objects based on their shape skeletons has also been used as a similarity measure for feature tracking over time [41] and feature matching following deformation [35, 61]

4 GENERALIZED GEODESIC DENSITY

Often, structures of interest in an image may be determined by *morphology* rather than derived from image intensity. In the case of foams,

ligaments and junctions may be locally indistinguishable in the image intensity (since the metal appears with the same intensity pattern in the images), and it is rather the local arrangement of material that defines a structure. Instead, a practical approach is to compute a field, such as the distance field, that encodes morphology, and then identify intensity features in the derived field. In the case of uniform radius tubes coming together at a junction, the intensity values of the distance field itself are not sufficiently sensitive to identify morphological structures. Instead, we introduce a new measure of localized density to be used to identify junctions. We first present basic geodesic density, which estimates the amount of material connected to and surrounding a point, and then improve the sensitivity of this measure further by generalizing the speed function used in geodesic distance.

4.1 Geodesic Density

Let M be a Riemannian sub-manifold of $\mathbb{R}^3$. The *geodesic distance* between points $a, b \in M$ is defined as

$$d_M(a, b) = \inf_{C_{ab}} \{L(C)\},$$

where

$$L(C) = \int_a^b dl,$$

i.e., the length L of the shortest path among the set of connected paths C_{ab} containing a and b . The geodesic ball of radius r around a point a is

$$B_M(a, r) = \{b \in M \mid d_M(a, b) \leq r\},$$

that is, the set of points connected to a with paths of length r or less.

In our study of metallic foams, we utilized geodesics to characterize properties of points in the material in the context of their neighborhood connected through the material, rather than Euclidean distance. The primary motivation for this is that as the foam is compressed, junctions and ligaments that initially were separated may become abutting. However, the geodesic distances between points in the material are less affected by the deformation.

We introduce the *geodesic density* function ρ' defined on a point in a manifold $a \in M$ and a radius r ,

$$\rho' = V(B_M(a, r)) / V(B_{\mathbb{R}^3}(r)).$$

The geodesic density is simply the ratio of the volume V of a geodesic ball centered at a with radius r to the volume of a Euclidean ball with the same radius in the embedding space $\mathbb{R}^3$. The intuition is that density, $\rho = m/V$, is the amount of mass per unit volume. In our case, the submanifold M represents metallic material, and by using the geodesic distance rather than Euclidean, the density corresponds to a measure local to the points in ligaments and junctions. In other words, adjacent, but structurally separated, ligaments and junctions are not considered when determining the density of a point.

4.2 Identifying ligaments/junctions with geodesic density

We model ligaments as cylindrical ‘tubes’ and junctions as intersections of at least three cylindrical segments. Intuitively, the density ρ' at junctions is higher than along ligaments, as more material surrounds points in junctions. Consider a Euclidean ball in $\mathbb{R}^3$ with radius R intersected by a tubular segment (subset of M) with radius r , as illustrated in Figure 4. The volume of the tube with spherical caps is determined by the two radii, r and R , as is the volume of an overall tube. We list the equations relating various structures with respect to a ball of radius R .

We wish to use ρ' to distinguish between a tube, intersecting tubes, and disks. Given a distance field d from the boundary of M , in Figure 4, the center of each has a value r in d (or $h = r$ in the case of the disk). These points are indistinguishable in the distance field, however, they have different values in ρ' , as is shown in the plots in Figure 5. The larger the radius R of the ‘search ball’, the greater the difference between the values of the point in a tube, intersection, and disk. Figure 6 illustrates that the larger the search radius, the easier it is to define an iso-value that is localized to a junction.




Fig. 4: Sphere (dark blue) transected (light blue) by tube, a junction, and a disk.

4.3 Generalized Geodesic Density

As the material collapses on itself during compression, large merged areas become difficult to separate, even with the geodesic density. By modulating the local speed, it is possible to amplify variation in the image intensity that could correspond to the hairline interface between merged regions, improving the sensitivity of the geodesic distance. The junctions in Figure 8 need this improved sensitivity to be separable. A generalized geodesic distance simply takes into account varying speeds along paths between a and b , and instead is defined as the time to traverse the *fastest* path between them. Let $s : M \rightarrow \mathbb{R}$ be a function that defines the speed at any point in the domain. Then the generalized geodesic distance between points is defined as

$$d_M^*(a, b, s) = \inf_{C_{ab}} \{L^*(C, s)\},$$

where

$$L^*(C, s) = \int_a^b \frac{1}{s(l)} dl,$$

i.e., the time L^* of the fastest path among the set of connected paths C_{ab} containing a and b given the speed function s . The generalized geodesic ball of time-radius t_r around a point a then becomes

$$B_M^*(a, t_r, s) = \{b \in M \mid d_M^*(a, b, s) \leq t_r\}.$$

We introduce the *generalized geodesic density* function ρ^* defined on a point in a manifold $a \in M$ and a radius r with speed function s as,

$$\rho^* = V(B_M^*(a, t_r, s)) / V(B_{\mathbb{R}^3}^*(t_r, 1)),$$

that is, the ratio of the volume reachable from a in time t_r given speed function s , to the volume reachable in Euclidean space with constant speed of 1.

Picking a speed function The generalized geodesic density ρ^* is more expressive, capable of adapting to the underlying intensity values in the image. We use the speed function to slow down propagation to avoid crossing cracks that appear as faintly diminished intensities in the material, to avoid gathering material from the other side in the density estimate. Therefore, our speed function is defined as $s : M \rightarrow [0, 1]$, where $s(v) = 0$ if v is a voxel that is unambiguously in the background, and $s(v) = 1$ if v is a voxel that is unambiguously interior to the material. We allow the user to set two thresholds on the image intensity, a low threshold for background, and a high threshold for interior of material. The speed function parameterizes (and clamps if necessary) all intensity values to the range $[0, 1]$ with linear interpolation. Section 5.2.1 discusses more how a user selects the speed function.

5 A WORKFLOW FOR SEGMENTING METALLIC FOAMS

Input and Desired Result. In the metallic foam segmentation task, we are given a time series of images $I_0, \dots, I_n$, and an estimated material density value ρ_{mat} . The first desired result is, for each image I_i , to identify each existing ligament and report its descriptive properties, such as length, curvature, cross section parameters, and orientation. The second




Fig. 5: The ratio of the volume of a junction to the volume of a tube is plotted in red as a function of the ratio of the radius of the tube r to the search radius R . The dotted line represents the portion where the closed-form equation $V_{junction}$ is not valid, and an approximation is used. The ratio of the volume of a disk-like feature with thickness a to a tube with radius $a/2$ is plotted in green as a function of a/R . The tube, junction, and disk all have identical values in the distance function, but they can be separated by the geodesic density function ρ^* . As the search radius R increases, the ratios r/R and a/R decrease, and the curves increase. For instance, if $R = 2r$, then the intersection points at junctions are expected to have 1.28x the value in ρ^* as points along the center of a ligament of the same radius.




Fig. 6: A T-junction of tubes with radius r is rasterized onto a volumetric grid. As the search radius R increases (indicated with a yellow sphere, left to right), the contours of the geodesic density function ρ^* become more localized in the vicinity of the junction. The colormap on each slice is rescaled.

desired result is time-tracking each ligament, i.e., for each ligament ℓ_j in image I_i , to identify its representative ℓ_k in I_{i+1} . Furthermore, the time tracking should detect when a ligament breaks, that is, when ℓ_j in I_i has no corresponding representative in I_{i+1} .

Strategy of Approach. Ligaments, on their own, are extremely challenging to extract robustly in later time steps of a compression sequence, because 1) ligaments can break, and 2) apparent ligaments can form where the previously disconnected portions of the material become adjacent. Compounding the challenge is that segmentation of CT scans is sensitive to noise in the data. Our overall strategy is to generate new features, junctions, extracted from a generalized geodesic density field in a manner that is more robust, and then define ligaments as portions of a topological skeleton connecting junctions. Junctions are then tracked through time, allowing inference of ligament tracking and detecting fracture points. Finally, as no new ligaments can form in the material, skeletonization artifacts that form during the compression sequence can be detected as newly connected junctions, and discarded.

5.1 Overview of Entire Workflow

The end-to-end workflow consists of four main stages: 1) creating junctions, 2) creating proto-ligaments, 3) time tracking and resolving ligaments, and 4) metrology. The first two stages of the workflow involve hands-on exploration of the data to determine thresholds for the feature extraction. Figure 7 provides an overview of the workflow, with numbered steps referred to in the following text. The green blocks in this diagram indicate user intervention to select parameter settings. The first, and most labor-intensive, part of the overall workflow is creating junctions. Once junctions have been identified, a skeleton is created of the material using a distance field and the Morse-Smale complex. The skeleton is combined with the junctions to create proto-




Fig. 7: Workflow for extracting and tracking features from foam image data over time, with numbers ordering the steps. Red numbers indicate steps where manual interaction is required and black indicates automated computation. The user first estimates (1-3) thresholds to correctly segment the junctions using the geodesic density metric (4-5). The junctions segmented can guide the extraction of the Morse-Smale complex (MSC) representation of the signed distance field of the input data (6) to generate a 1-skeleton containing a superset of the ligaments (7). Finally, time tracking of junctions and ligaments (9-11) reports foam attributes over time.

ligament, structures that *might* be ligaments, but that need further verification. Next, time-tracking of the junctions is performed, and the corresponding time-tracking of the ligaments is inferred. The ligament tracking information allows classification and discarding of proto-ligaments that have no precursor in the prior time step. Finally, several metrics of interest are computed for each ligament over the time series.

5.2 Workflow for Creating Junctions

Given the time series of images $I_0, \dots, I_n$, our goal is to generate geodesic density fields $\rho^*(I_0), \dots, \rho^*(I_n)$ that have the property that junctions are identifiable as connected regions above a threshold h . Formally, the junctions are the super-level set components given a threshold in the geodesic density field. Furthermore, ideally, for each $i, j \in 0, \dots, n$ a 1-to-1 map can be created between connected components of $\rho^*(I_i)$ and $\rho^*(I_j)$. As demonstrated in section 4, given an appropriate search radius and speed function, junctions can be distinguished with a threshold from ligaments in an idealized model. However, given the noise and morphological variability in real-world data, the radius selection, speed function specification, and threshold selection become a trial-and-error approach that relies on visual validation.

5.2.1 Generating generalized geodesic density

To generate a generalized geodesic density field that can reliably extract junctions as connected components given a threshold, a user first examines the *last* and then the *first* image in a time series. We have observed that in practice, the last time step represents the most complex scenario, where most of the ligaments are collapsed, and the material is compressed in a very limited portion of the initial volume. The last image gives an overview of what level of merging geometry happens, and how/if morphological structures are represented in the compressed state. The last image is used to identify the material and background values to be used to parameterize the speed function (from section 4.3), or if an additional processing step must be undertaken. The first image provides a pristine view of the ligaments, and is used both to select a search radius, and also visually verify that the chosen super-level set threshold creates a component at every junction.

Defining a speed function. During our thresholds selection, we want to understand which values in the volume under examination represent background and which instead represent material. The goal of using a generalized geodesic density function, as opposed to a regular geodesic density function, is to account for situations where a single threshold is not sufficient to determine which voxels are material and which are background. By controlling the speed, faint features separating morphological structures can be amplified, or in extreme cases where separate structures are not separated by image intensity, to fill holes by working on a distance field instead.




Fig. 8: When structures touch in the last time step, a user determines if they are separable by a simple intensity thresholding. In the CT scan on the left, the structures are separable, and the blue represents user-selected background, and the green represents user-selected material. The CT scan on the right is not able to separate merged structures, and therefore a 3d distance field is computed from the material interface, and a user selects background/material in this image.

The input images are loaded in either ImageJ [44] or Paraview [1], depending on a user's preference. A user navigates via slices to identify locations where the foam is pressed on itself. Certain CT scans have the sensitivity to represent hairline gaps between the structures, in which case the user picks a threshold (Figure 7, step 1) that represents “definite background” and another threshold that completely separates ambiguous regions (Figure 7, step 2) and can be interpreted as “definite material”. In other CT examples, and examples arising from voxelizing finite element simulations, there is often no intensity value that separates merged structures. In this case, we compute a distance field to the material interface threshold (see section 5.3), and select thresholds on the distance field, such that the morphological structures are separated, while retaining as much of the material as possible as interior. The thresholds selected for the last time step are reused for the entire image sequence.

Picking a Search Radius R . The relative ratios between the density value measured along tubes and at intersections increase with R , up to a limit of 2:3. Therefore, as R increases, the selectivity of a threshold of ρ^* increases as well; it is easier to pick a value of ρ^* that includes a super-level-set component at each junction, while omitting unwanted points along ligaments. Two factors provide an upper limit to how large R should be relative to r . First, the influence of a junction (or other thick material body) will extend at least R from the ends of a ligament. With R




Fig. 9: Selection of junction threshold. The goal is to find a threshold that both separates compressing junctions at the last timestep and does not remove any existing junction at the first time step. On the left, the threshold is too low, over connecting junctions, and on the right the threshold is too high, removing junctions.

too large, ligaments can erroneously count the material of the junctions at either end, making junctions harder to separate from ligaments. For the materials in this study, however, ligament lengths were much larger than the spacing between junctions, meaning that this concern did not impact the selection of R . The second limiting factor for the maximum size of R is performance. For each voxel inside the material, $\frac{4\pi}{3}R^3$ potential voxels must be visited during the computation of the geodesic density function. Although evaluating ρ^* is embarrassingly parallel, large R values still prove prohibitively costly. In practice, we have found that a ratio of $\frac{r}{R} = 0.6$ consistently produced density fields from which it was simple to extract junctions using a threshold. For the materials studied, this corresponded to a search radius $R < 30$ voxels.

To select the search radius (Figure 7, step 3), a user loads the first time step, and similar to selecting the speed function thresholds, navigates through slices of the data. Using the visualization system's measuring tools, the user estimates the radius r of the largest tube to be considered a ligament. The value of R is recommended to be $1.6r$.

Selecting Junction Threshold. The generalized geodesic density function ρ^* is computed (Figure 7, step 4) given the user-selected parameters for the first and last time steps and written as raw files. A user loads these two files in Paraview as well as the two original images, and then uses the contour filter to investigate the density threshold to select in the context of a semitransparent material interface surface (Figure 7, step 5). Starting with the last time step, the lowest threshold value is picked that produces separated connected components for each junction. This same threshold is applied to the first time step where with the less cluttered configuration, a user is better able to visually verify that each junction has been extracted. Figure 9 illustrates this process. Checking that the number of connected components is the same for the first and last time steps is a strong indicator of stability of the junction extraction, however, not strictly necessary. For instance, in certain datasets from simulated foams, material is moved out of the field of view of the simulation domain, leading to fewer junctions in later time steps. Whereas in the vast majority of cases users reported that the original density function ρ^* produced a satisfactory set of junctions, in rare cases the thresholds used to produce ρ^* needed to be adjusted. For instance, in the case where no threshold exists that separates the morphological structures in the last time step, both the background threshold and the material threshold must be increased, and the density field regenerated. In the case that junctions are overly connected in the first time step, the radius R needs to increase and a new ρ^* must be computed. No case was encountered where R was too large.

Process Entire Time Series. We reuse the four thresholds identified (R , background, material, and junction) and process the entire time series. If the background/material threshold selection required it, a distance field is computed for each image, and that field is used in place of the image itself for the rest of the workflow. The density field ρ^*




Fig. 10: Using Morse-Smale complex (MSC) representation to compute a 1-skeleton from the signed distance field

is generated for each image. The super-level set threshold identified during visual inspection is used to create an index volume from ρ^* , where each voxel is labeled either as background if it has value below the threshold, or with the identifier of its junction. A disjoint-set Union-Find data structure is used in a pass through the index volume to identify the connected components in the volume. Additionally, we generate boundary surfaces of each connected component using VTK [45] to enable rapid loading and visual verification in Paraview.

5.3 Workflow for Creating Proto-Ligaments

In a foam, the ligaments are portions of material connecting two junctions. For the analysis of ligaments, the desired objective is to extract a line in the medial axis of ligaments connecting junctions. Our strategy is to extract a skeleton representation of the foam and then identify the subset of its lines that form the ligaments. Our approach to skeletonization is based on computing the Morse-Smale complex (MSC) of a distance field, simplifying it to remove spurious connections, and clipping the 1-skeleton using junctions. First, a threshold is selected to identify the interface between material and background. Next, a signed distance field is computed using this interface. The MSC is computed and simplified, and the ridge-like arcs entirely within the material are saved. The arcs are then clipped by the components of the junction to create proto-ligaments. These are piecewise linear paths connecting two different junctions. Determining whether a proto-ligament corresponds to an actual ligament or to an artifact is left to the time tracking stage of the overall workflow.

Estimating Material Interface Threshold. The surface of the material is an isosurface of the image intensity. Picking the right intensity value to separate material from background is a first step to skeleton generation (Figure 7, step 6). The material scientists know, *a priori*, the relative density of the material for most samples. In this case, they supply a number that is the percentage of the foam volume that is material vs. the total enclosed volume of the sample, for the uncompressed (time = 0) foam. Foam images are available either as filling an axis-aligned rectangular domain, or as a cylinder inside the rectangular domain. In the former case, we use a threshold that simply includes the correct percentage of voxels. In the latter case, we found that manual intervention was needed to either set the threshold, or pick the subdomain on which the automated threshold could be computed. For both simulated and imaged data, the image acquisition parameters are kept constant so that the same threshold separates material from background in the entire time series. However, we found that for certain time series CT images, the average image intensity fluctuated between time steps. In this case, a new threshold was computed for each subsequent time step by including the same number of high-intensity voxels as the prior step. In particular, given a threshold $TR1$ at time step $T1$, we compute the number of voxels in an iso-volume at $TR1$ (i.e., number of voxels with value lower or equal to $TR1$) and find the corresponding threshold $TR2$ at time step $T2$ that will produce the same number of voxels in an iso-volume at $TR2$. This approach uses the fact that no matter is created or lost during the time series.

Distance field computation. A distance transform is a field where each voxel in a volume represents the shortest distance from that voxel to an object. In our case, the distance transform computed (Figure 7, step 7) for each voxel is the signed Euclidean distance to the surface of the material, with positive values inside the material and




Fig. 11: The 1-skeleton generated by the Morse-Smale complex places a maximum (red spheres) in each junction (pink blobs) (a), with arcs (yellow tubes) connecting them. This graph structure is used to derive the connectivity of junctions. The locations of the 1-skeleton within disk-like regions is not stable. In (b) the skeleton in consecutive time steps (red for gray time step, blue for purple) shifts from one side of the disk to the other, highlighting the need for stable junction extraction.

negative outside. We use a distance field to ensure that subsequent skeleton computation places the arcs in the middle of each ligament. **MS complex computation/simplification.** The Morse-Smale complex has been well established as a means of extracting the 1-skeleton of a distance field to the interface surface of a porous solid [29], (Figure 7, step 8). Key aspects that make it attractive are the ability to simplify the representation with respect to noise, and to explore the reconstruction as thresholds are varied interactively [26]. We use a parallel discrete gradient algorithm [28] and MSC computation in the open source library MSCEER [27]. The MSC of the signed distance field, without any simplification, contains a super set of the features we are interested in (see Figure 10a).

This introduces a new challenge, that is, to find the right amount of persistence simplification to produce a skeleton that is consistent with the input foam structure. In other words, we want our 1-skeleton to include all and only the arcs connecting maxima that correspond to junctions in the foam. In all the examples we have seen, each junction contained a maximum in the MSC (see Figure 11a). Although the exact location of that maximum within a junction was unstable, for instance, moving between sides of disk-like structures between time steps (see Figure 11b), each junction had a high persistence maximum. Intuitively, the accumulation of material at junctions tends to increase the distance value, which also corresponds to a local increase in the geodesic density function ρ^* . To find the right persistence to simplify the complex, removing noise while keeping the main ridge-like structures, we use the number of junctions found in the previous step. Specifically, the number of junctions extracted using the geodesic density is the target the number of maxima that we want to have in our 1-skeleton. Given this input, we repeatedly apply cancellation operations on the MSC [26] until only the desired number of maxima remain. The 1-skeleton is formed by the critical points and the 2-saddle-maximum arcs (ridge-like structure) that remain after simplification. However, we note that there is no guarantee that each junction has exactly one maximum. An exact matching is not needed; instead, we use the 1-skeleton as a graph embedded in the domain that illuminates the connections between junctions and their pathways.

Proto-ligaments: connecting junctions with the 1-skeleton. The ridge-like arcs of the 1-skeleton cross between junctions. In the interior of junctions, the exact location of the maximum, and the arcs themselves, are unstable, as shown in Figure 11b. Therefore, we remove the portion of the 1-skeleton that is interior to a junction (Figure 7, step 9). We define a *proto-ligament* as any remaining path connecting two different junctions. We use the index volumes created during the connected component sweep of the junctions workflow to query whether a point on an arc is inside a junction. We use the terminology proto-ligament, because the existence of a path between junctions does not necessarily indicate the existence of a ligament; in later time steps, newly created paths can be found between junctions due to the material self-contacting. These are not to be considered in the final analysis.

5.4 Time Tracking Junctions, Resolving Ligaments

Our feature tracking system does a forward trace of junctions and ligaments, using the resolved features of time step $i - 1$ to resolve




Fig. 12: We show the result of closest point matching (a) between junctions in consecutive time steps. For visual verification, we also produce closest point matching between tracked ligaments (b). The ligaments and their matches are colored according to feature id.

features at time step i . It is a two-stage strategy, first creating a mapping between junctions, and then using that mapping to infer a mapping of ligaments. Proto-ligaments in time step i are accepted as true ligaments if a precursor is identified from time step $i - 1$.

Mapping Junctions. The relatively large time steps taken between CT images restrict the kind of feature tracking that is possible. For instance, overlap tracking fails because the junctions often move a distance several times their own size. Instead, we employ a greedy matching approach, where the largest, closest junctions are matched first, followed by smaller, farther junctions (Figure 7, step 10).

Let J_i, J_{i-1} be the sets of junctions in time steps i , and $i - 1$, respectively. We wish to build a set of pairs $\sigma = \langle \gamma_j, \gamma_k \rangle$ uniquely mapping between junctions of time steps i and $i - 1$. Each point p of each junction $\gamma_j \in J_i$ is inserted into a spatial acceleration structure, in our case a kd-tree K_j . Similarly, each point of each junction $\gamma_k \in J_{i-1}$ is inserted into K_{i-1} . Let $\text{ClosestJ}(p, K)$ return the identifier of the junction that has the closest point to p in K . We build a $\|J_i\| \times \|J_{i-1}\|$ matrix M , where for $\gamma_j \in J_i$ and $\gamma_k \in J_{i-1}$, the $m_{j,k}$ element of M is

$$m_{j,k} = \|p \in \gamma_j \mid \text{ClosestJ}(p, K_{i-1}) = k\| + \|q \in \gamma_k \mid \text{ClosestJ}(q, K_i) = j\|$$

Intuitively, the matrix M keeps track of how many times the closest point q in the next time step to point p in the current time step occurred for each possible pairing of junctions. Given M , the pairing P is computed by finding the highest valued element $m_{j,k}$, pairing junctions γ_j and γ_k into pair $\sigma_{jk} = \langle \gamma_j, \gamma_k \rangle$, adding σ_{jk} to P , and then removing row j and column k from M . This greedy approach is repeated until M is empty. We illustrate the results of the closest point queries from every point in time $i - 1$ to time i in Figure 12a. In practice, M is very sparse and is stored as a sparse graph, and the temporal matching is nearly linear in the number of junctions.

Resolving, and Tracking Ligaments. Ligaments are resolved and tracked inductively (Figure 7, step 11). We accept every proto-ligament in the first time step $t = 0$ as a ligament. Let P be the set of pairings produced by the junction tracking between time step i and $i - 1$. Each proto-ligament ℓ_i has endpoints connecting two junctions in the same time step. We accept proto-ligament ℓ in time i as a ligament, if the junctions at its endpoints, γ_k, γ_l , are mapped to the junctions time $i - 1$ that are also connected by an accepted ligament. Intuitively, a proto-ligament becomes a ligament if, mapping its junctions back in time, the mapped junctions have a ligament in the previous time step. Therefore, only connections initially present in the uncompressed foam can be carried forward by the time tracking, a desirable property when the crushing can cause new connections to appear for material that is merely adjacent. As a visual aide, and for debugging, we also provide users with a visualization drawing a line segment between each vertex along a ligament and the closest vertex on the tracked, matching ligament, seen in Figure 12b.

5.5 Metrology

The quantities of interest to domain experts pertain mainly to the ligaments. The ligaments are available as a sequence of 3d points, and the identifiers of the ligaments at each end point. Additionally, the original image, the distance field, and the density fields are sampled at each point along a ligament. Understanding what metric impacts

material performance relies on collecting a broad spectrum of information about the ligaments. The geometric realization of a ligament originates from the 1-skeleton of the MSC, which was computed using a discrete gradient approach. As a result, initial line segments on a ligament are aligned with axis orientation. We perform constrained iterative smoothing to remove these effects and obtain a smooth line for each ligament. For each ligament, for each time step we compute:

Length: The accumulated Euclidean distance between adjacent points on ligament

Curvature: For computing curvature $\mathcal{C}$, we consider a normal vector $\hat{n}_{p_i}$ to each point p_i along ligament ℓ of length b . Adjacent points, p_i to p_{i+1} , are considered to curve proportionally to the cosine similarity of these normal vectors. The curvature of the ligament is then the sum of all cosine similarities between adjacent normal vectors along the ligament.

$$\mathcal{C}(\ell) = \sum_{i=0}^b \cosim(\hat{n}_{p_i}, \hat{n}_{p_{i+1}})$$

Orientation: Similarly, the orientation of a ligament is defined to be the cosine similarity of the initial values of the 1-skeleton of a ligament to the $\hat{x}$, $\hat{y}$ and $\hat{z}$ axis vectors.

Cross section area and perimeter: We collect cross section areas and perimeters along several slices normal to the ligament along its path. As is done with curvature, to compute the cross section of a ligament we obtain the normal vector from a voxel on the 1-skeleton to its neighbor. With this normal vector and the perpendicular vector tangential to the ligaments 1-skeleton, we compute and take a cross sectional slice. Since multiple ligaments can be very near, potentially almost touching, it is necessary to identify which ligament was initially sliced. For this reason, the ligament of interest is marked, and once the slice is obtained we compute all connected components. The connected component with the marked voxel is then the cross sectional slice.

6 EVALUATION

The workflow of the entire process (depicted in Figure 7) was evaluated using both simulated and imaged data. In particular, we considered a small simulated dataset of size 300x300x300 voxels and four large experimental datasets with XY resolution 1024x1024 and different number of slices over the Z direction (see Table 1 for details). The evaluation was performed by four users who experimented with the process and stages of the workflow, and results were verified and validated at each stage.

User Evaluation. For a given foam data, the users performed all the interactive steps described in the previous section, and then visualized and evaluated the results. Of the four domain scientists involved in testing our system, two were considered experts (professors), as they collaborated during the development of the techniques, and two were undergraduates in materials science, previously involved only in sample generation and imaging. Users were guided only by a flow chart similar to the figures in Section 5 for running the tools, and a slide deck with screen-shots for guiding the user through Paraview's interface. All software tools are linked with scripts, and produce state files for visual inspection in Paraview.

User Feedback. During testing of an early prototype, users needed several passes through the workflow to ensure the thresholds selected separated junctions of the compressed foam. After reordering the workflow to *first* look at the *most compressed* time step, and defaulting to use the generalized geodesic density, users reported that only occasionally would they have to complete a second pass of the entire workflow. For real CT image data, all users reported the ability to create satisfactory results with only a single pass for threshold generation. The greatest challenge was posed by simulated foams, where occasionally users required two or three iterations of the junction extraction, as the last time steps have heavily merged geometry. To ensure that the merged junctions could be separated required users to switch to computing the geodesic density on the distance field instead.

Performance. In Table 2, the computation times of the different analysis steps are reported for one simulated and one experimental dataset. The analysis were performed on a 12-core Intel(R) Xeon(R) Gold 6136




Fig. 13: Curvature of ligaments over time and details of detected ligament fracture. From the plot it is possible to detect deforming ligaments as curves increasing over time. Separation occurs if the curve terminates before the last time step. The images above show the last three time steps before separation of the 4 highlighted ligaments. The images on the last row show full separation, however fracture might occur at any time step.

Table 1: Number of junctions and ligaments extracted from four experimental foams (see corresponding foams in Figure 14). All datasets have images of resolution 1024x1024 and number of slices as listed.

Foam	N. slices	N. junctions	N. ligaments	Rel. density
a	759	1858	3651	5.34 %
b	759	983	2304	12.04 %
c	1043	1525	3130	3.95 %
d	1001	1063	1985	3.90 %

CPU @ 3.00GHz. The time for the users to select the various thresholds was under 5 minutes. In Figure 14 we report images of the features extracted in the four experimental foams, and we include a summary of the number of junctions and ligament extracted from each of them in Table 1. Users verified the results via visual inspection.

Visual Validation. Given that no ground-truth segmentation exists for the data being analyzed, all the metrics and features extracted were visually inspected and validated by the users at each step of the process. In particular, Paraview was used to visualize isocontours of the junctions in combination with tube visualization of the ligaments in the context of the material interface surface (as in Figure 14). Furthermore, coloring the ligaments by metrics (as in Figure 15) helped the user in validating the values of each. For example, in Figure 15, ligaments are either colored by inside/outside a junction (left), ligament length (middle), and orientation respect to the Z axis (right).

Finally, since ligament and junctions are tracked over time, we were able to observe the evolution of their metrics in order to detect changes in connectivity. In particular, a change in the connectivity of junctions and ligaments, or a drastic variation of cross section area over time can be indicators of onset of ligament failure (e.g., fracture). For example, in Figure 1 we highlight in the red boxes connectivity changes that have been reported by the tracking analysis. In particular, those changes




Fig. 14: Junctions (pink) and ligaments (blue tubes) extracted from the CT scans of four foams traditionally manufactured (transparent gray isosurface). These results were obtained without supervision by nonexpert users.

Analysis	Synthetic Data	Experimental data
Geodesic Density	22.159s	58m 17.670s
Distance Field	32.182s	27m 9.592s
Gradient	31.947s	19m 24.507s
MSC graph extract	0.684s	21.342s
Tracking junctions	7.923s	7m 23.488s
Tracking ligaments	0.798s	3.183s

Table 2: Computation time of the analysis were performed using a synthetic dataset of size 300x300x300 and an experimental dataset of size 1043x1024x1024.




Fig. 15: Some of the metrics extracted for the ligaments. On the left, a binary field indicates when a point along the ligament is inside a junction or not, in the middle, the color of the tubes indicates the length of each ligament; and on the right, the color indicates the orientation respect to the Z axis of the ligament.

correspond to ligaments that fracture at a certain time step during the compression. In the yellow box, instead, is part of the material drifting out of the field of view of the CT scan.

Finally, to demonstrate that our process can help classify the behavior of the ligaments, we observed the curvature value of each ligament tracked over time (see Figure 13). From these results, it is possible to quickly identify ligaments that are collapsing or fracturing during the compression (i.e., increasing curvature value over time, followed by disappearance). In Figure 13, we report four examples of ligaments with increasing curvature that are fracturing at different time/compression steps. In one particular case (i.e., second column from left), the ligament feature extracted (in red) is still contacting the two junctions even after the ligament separated (third image from top) because the ligament is still touching the junction and producing two connected maxima in the MSC. However, in the next time step, the connectivity is lost and the ligament is correctly classified as fully fractured.

6.1 Limitations

Both experimental and simulated data present certain challenges that our workflow is not currently addressing. For example, given the limited field of view of the CT scans, parts of the material can drift in and out of the boundaries from one time step to the other (e.g., see yellow box in Figure 1). This can cause junctions and ligaments to appear or disappear from the boundaries. In our process, we currently ignore the features at the boundaries, also because for large datasets the number of features at the boundaries represents a very small percentage of the total. A limitation of our approach is the assumption that junctions can be identified independently for each time step. Although not yet encountered in the images processed so far, junctions that are completely molded to each other are unlikely to be separable, and will cause a

failure in their segmentation and tracking. Currently, our use of the generalized density does not leverage local material anisotropy, which could be added as a penalty term to the local speed function, for improved sensitivity to local morphology. Another limitation of this work is the reliance on the existence of maxima within each junction; one could imagine geometries, for instance a perfectly smooth T-junction, where no maximum is detected inside a junction defined by ρ^* . In this case, the persistence simplification might remove potential ligaments from the MSC 1-skeleton. Finally, computing the geodesic density function is expensive, practically limiting the search radius used, and discouraging experimentation with the speed function thresholds. We will look to accelerate this computation in future work. Finally, as foams become even larger and more complex, the visual validation parts of the workflow may become onerous and error-prone.

7 CONCLUSIONS

There is a high demand for means to understand and correlate the design space of material properties to the material performance metrics with respect to attributes of their features for open cell metallic foams. Prior to our work, the fine grain characterization of their features was accomplished via manual segmentation or with skeletonization approaches that may not accurately model the variety of shapes present in nodes and ligaments. In this work we introduce a high-throughput end-to-end workflow to segment foams into junctions and ligaments and track their behavior and properties over time. The task of tracking junctions and ligaments for each time-series of images is reduced to finding five thresholds: three to define the geodesic speed and radius, one for the material interface, and one for the junction threshold. Experiments performed by users on both simulated and imaged foams demonstrated the large benefits of this workflow, generating sets of thousands of ligaments and junctions with minimal user intervention. Furthermore, metrics collected from this process helped classify the behavior of ligaments over time, such as bending or breaking. The high throughput of our workflow is an essential step toward populating databases to be used to understand the parameters that regulate foam performance. In the future we plan to more fully automate the threshold selection (e.g., for cylindrical foams) and implement ρ^* and other computational stages of the workflow using accelerators (e.g., GPUs). Furthermore, we will investigate utilizing neighborhoods in the junction/ligament graph to improve time tracking, and make it more robust to large-scale dislocations. Finally, we will investigate using existing junctions to postulate a segmentation for time steps where junctions are lost, indicating a catastrophic merging that is not separable using traditional approaches.

ACKNOWLEDGMENTS

This work was supported in part by NSF:CGV Award: 1314896, NSF:IIP Award: 1602127, NSF:ACI Award: 1649923, DOE/SciDAC DESC0007446, PSAAP CCMSC DE-NA0002375, NSF:OAC Award: 1842042 and NSF CMMI-1629660. Additional support comes from Intel Graphics and Visualization Institutes of XeLLENCE program.

REFERENCES

- [1] J. Ahrens, B. Geveci, and C. Law. Paraview: An end-user tool for large data visualization. *The visualization handbook*, 717, 2005.
- [2] A. August, J. Ettrich, M. Röhl, S. Schmid, M. Berghoff, M. Selzer, and B. Nestler. Prediction of heat conduction in open-cell foams via the diffuse interface representation of the phase-field method. *International Journal of Heat and Mass Transfer*, 84:800–808, 2015.
- [3] C. Bajaj, A. Shamir, and B.-S. Sohn. Progressive tracking of isosurfaces in time-varying scalar fields, 2002.
- [4] J. Banhart. Manufacture, characterisation and application of cellular metals and metal foams. *Progress in materials science*, 46(6):559–632, 2001.
- [5] J. Banhart. Aluminum foams: On the road to real applications. *Mrs Bulletin*, 28(04):290–295, 2003.
- [6] F. Benmansour, E. Turetken, and P. Fua. Tubular geodesics using oriented flux: An itk implementation. *The Insight Journal*, 6:27–29, 2013.
- [7] G. Bertrand. A parallel thinning algorithm for medial surfaces. *Pattern Recogn. Lett.*, 16(9):979–986, Sept. 1995. doi: 10.1016/0167-8655(95)00034-E
- [8] G. Bertrand and Z. Aktouf. Three-dimensional thinning algorithm using subfields. In *Vision Geometry III*, vol. 2356, 1995. doi: 10.1117/12.198601
- [9] S. Beucher. Digital skeletons in euclidean and geodesic spaces. *Signal Processing*, 38(1):127–141, 1994.
- [10] H. Bhatia, A. G. Gyulassy, V. Lordi, J. E. Pask, V. Pascucci, and P. Bremer. Topoms: Comprehensive topological exploration for molecular and condensed-matter systems. *Journal of Computational Chemistry*, 0(0), 2018. doi: 10.1002/jcc.25181
- [11] I. Bitter, A. E. Kaufman, and M. Sato. Penalized-distance volumetric skeleton algorithm. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):195–206, 2001.
- [12] H. Blum. A transformation for extracting new descriptors of shape. In W. W. Dunn, ed., *Models for the Perception of Speech and Visual Form*, pp. 362–381. MIT Press, 1967.
- [13] G. Borgefors, I. Nyström, and G. Sanniti di Baja. Surface skeletonization of volume objects. In *International Workshop on Structural and Syntactic Pattern Recognition*, pp. 251–259. Springer, 08 1996. doi: 10.1007/3-540-61577-6_26
- [14] G. Borgefors, I. Nyström, and G. Sanniti di Baja. Computing skeletons in three dimensions. *Pattern Recognition*, 32:1225–1236, 07 1999. doi: 10.1016/S0031-3203(98)00082-X
- [15] S. Bouix, K. Siddiqi, and A. Tannenbaum. Flux driven fly throughs. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, vol. 1, pp. I–I. IEEE, 2003.
- [16] P.-T. Bremer, G. Weber, V. Pascucci, M. Day, and J. Bell. Analyzing and tracking burning structures in lean premixed hydrogen flames. *IEEE Transactions on Visualization and Computer Graphics*, 16(2):248–260, 2010.
- [17] A. Brennecke and T. Isenberg. 3d shape matching using skeleton graphs. In *SimVis*, pp. 299–310, 01 2004.
- [18] R. Cárdenes, S. K. Warfield, E. Macias, and J. Ruiz-Alzola. Occlusion points propagation geodesic distance transformation. In *Proceedings 2003 International Conference on Image Processing (Cat. No. 03CH37429)*, vol. 1, pp. I–361. IEEE, 2003.
- [19] N. D. Cornea, D. Silver, X. Yuan, and R. Balasubramanian. Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer*, 21(11):945–955, 2005.
- [20] T. K. Dey and J. Sun. Defining and Computing Curve-skeletons with Medial Geodesic Function. In A. Sheffer and K. Polthier, eds., *Symposium on Geometry Processing*. The Eurographics Association, 2006. doi: 10.2312/SGP/SGP06/143-152
- [21] H. Edelsbrunner, J. Harer, A. Mascarenhas, and V. Pascucci. Time-varying Reeb graphs for continuous space-time data. In *20th Symp. on Computational Geometry*, pp. 366–372. ACM, ACM Press, New York, NY, USA, 2004.
- [22] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science, FOCS '00*, pp. 454–. IEEE Computer Society, Washington, DC, USA, 2000.
- [23] A. G. Evans, J. Hutchinson, and M. Ashby. Multifunctionality of cellular metal systems. *Progress in Materials Science*, 43(3):171–221, 1998.
- [24] N. Gagvani and D. Silver. Parameter controlled skeletonization of three dimensional objects. *Dept. Elect. and Comput. Eng., Rutgers Univ., Piscataway, NJ, Tech. Rep. CAIP-TR-216*, 1997.
- [25] L. J. Gibson and M. F. Ashby. *Cellular solids: structure and properties*. Cambridge university press, 1999.
- [26] A. Gyulassy, V. Pascucci, P. Bremer, and B. Hamann. A topological approach to simplification of three-dimensional scalar functions. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):474–484, July 2006. doi: 10.1109/TVCG.2006.57
- [27] A. Gyulassy. MSCEER: Morse-Smale Complex Extraction, Exploration, Reasoning. <https://github.com/sci-visus/MSCEER>, 2018.
- [28] A. Gyulassy, T. Bremer, and V. Pascucci. Shared-memory parallel computation of morse-smale complexes with improved accuracy. *IEEE Transactions on Visualization and Computer Graphics*, 25:1183–1192, 2019.
- [29] A. Gyulassy, M. Duchaineau, V. Natarajan, V. Pascucci, E. Bringa, A. Higginbotham, and B. Hamann. Topologically clean distance fields. *IEEE Transactions on Computer Graphics and Visualization*, 13(6):1432–1439, 2007.
- [30] A. Gyulassy, A. Knoll, K. C. Lau, B. Wang, P. Bremer, M. E. Papka, L. A. Curtiss, and V. Pascucci. Interstitial and interlayer ion diffusion geometry extraction in graphitic nanosphere battery materials. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):916–925, Jan 2016. doi: 10.1109/TVCG.2015.2467432
- [31] M. S. Hassouna and A. A. Farag. Robust centerline extraction framework using level sets. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1, pp. 458–465. IEEE, 2005.
- [32] C. Körner and R. F. Singer. Processing of metal foams—challenges and opportunities. *Advanced Engineering Materials*, 2(4):159–165, 2000.
- [33] C. Lantuéjoul and S. Beucher. On the use of the geodesic metric in image analysis. *Journal of Microscopy*, 121(1):39–49, 1981.
- [34] C. Lantuéjoul and F. Maisonneuve. Geodesic methods in quantitative image analysis. *Pattern recognition*, 17(2):177–187, 1984.
- [35] A. S. Lowet, C. Firestone, and B. J. Scholl. Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception, & Psychophysics*, 80(5):1278–1289, 2018.
- [36] E. Maire, A. Fazekas, L. Salvo, R. Dendievel, S. Youssef, P. Cloetens, and J. M. Letang. X-ray tomography applied to the characterization of cellular materials. related finite element modeling problems. *Composites science and technology*, 63(16):2431–2443, 2003.
- [37] K. E. Matheson, K. K. Cross, M. M. Nowell, and A. D. Spear. A multi-scale comparison of stochastic open-cell aluminum foam produced via conventional and additive-manufacturing routes. *Materials Science and Engineering: A*, 707:181–192, 2017.
- [38] F. Meyer. Mathematical morphology: from two dimensions to three dimensions. *Journal of microscopy*, 165(1):5–28, 1992.
- [39] L. E. Murr, S. M. Gaytan, D. A. Ramirez, E. Martinez, J. Hernandez, K. N. Amato, P. W. Shindo, F. R. Medina, and R. B. Wicker. Metal fabrication by additive manufacturing using laser and electron beam melting technologies. *Journal of Materials Science & Technology*, 28(1):1–14, 2012.
- [40] K. Palágyi, E. Sorantin, E. Balogh, A. Kuba, C. Halmai, B. Erdöhelyi, and K. Hausegger. A sequential 3d thinning algorithm and its medical applications. In *IPMI*, 2001.
- [41] F. Reinders, F. H. Post, and H. J. Spoelder. Visualization of time-dependent data with feature tracking and event detection. *The Visual Computer*, 17(1):55–71, 2001.
- [42] P. K. Saha and B. B. Chaudhuri. Detection of 3-d simple points for topology preserving transformations with application to thinning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):1028–1032, Oct 1994. doi: 10.1109/34.329007
- [43] M. Sato, I. Bitter, M. A. Bender, A. E. Kaufman, and M. Nakajima. Teasar: tree-structure extraction algorithm for accurate and robust skeletons. In *Proceedings the Eighth Pacific Conference on Computer Graphics and Applications*, pp. 281–449, Oct 2000. doi: 10.1109/PCCGA.2000.883951
- [44] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri. Nih image to imagej: 25 years of image analysis. *Nature methods*, 9(7):671, 2012.
- [45] W. J. Schroeder, B. Lorensen, and K. Martin. *The visualization toolkit: an object-oriented approach to 3D graphics*. Kitware, 2004.
- [46] J. A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, vol. 3. Cambridge university press, 1999.
- [47] P. Siegkas, V. Tagarielli, and N. Petrinic. Modelling stochastic foam geometries for fe simulations using 3d voronoi cells. *Procedia Materials Science*, 4:221–226, 2014.

- [48] D. Silver and X. Wang. Tracking and visualizing turbulent 3d features. *IEEE Trans. Vis. Comp. Graph.*, 3(2):129–141, apr-jun 1997. doi: 10.1109/2945.597796
- [49] D. Silver and X. Wang. Tracking scalar features in unstructured data sets. In *Visualization '98. Proceedings*, pp. 79–86, Oct 1998. doi: 10.1109/VISUAL.1998.745288
- [50] R. Singh, P. Lee, R. Dashwood, and T. Lindley. Titanium foams for biomedical applications: a review. *Materials Technology*, 25(3-4):127–136, 2010.
- [51] T. Sousbie. The persistent cosmic web and its filamentary structure - I. theory and implementation. *Monthly Notices of the Royal Astronomical Society*, 414(1):350–383, 6 2011.
- [52] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *2003 Shape Modeling International.*, pp. 130–139. IEEE, 2003.
- [53] S. Svensson, I. Nyström, and G. Sanniti di Baja. Curve skeletonization of surface-like objects in 3d images guided voxel classification. *Pattern Recognition Letters*, 23:1419–1426, 10 2002. doi: 10.1016/S0167-8655(02)00102-2
- [54] D. Ushizima, D. Morozov, G. H. Weber, A. G. C. Bianchi, J. A. Sethian, and E. W. Bethel. Augmented topological descriptors of pore networks for material science. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2041–2050, Dec 2012. doi: 10.1109/TVCG.2012.200
- [55] H. N. Wadley. Cellular metals manufacturing. *Advanced engineering materials*, 4(10):726–733, 2002.
- [56] G. Weber, P.-T. Bremer, M. Day, J. Bell, and V. Pascucci. Feature tracking using reeb graphs. In *Topological Methods in Data Analysis and Visualization*, Mathematics and Visualization, pp. 241–253. Springer, 2011. doi: 10.1007/978-3-642-15014-2_20
- [57] G. D. Wehinger, H. Heitmann, and M. Kraume. An artificial structure modeler for 3d cfd simulations of catalytic foams. *Chemical Engineering Journal*, 284:543–556, 2016.
- [58] W. Widanagamaachchi, J. Chen, P. Klacansky, V. Pascucci, H. Kolla, A. Bhagatwala, and P. Bremer. Tracking features in embedded surfaces: Understanding extinction in turbulent combustion. In *2015 IEEE 5th Symposium on Large Data Analysis and Visualization (LDAV)*, pp. 9–16, Oct 2015. doi: 10.1109/LDAV.2015.7348066
- [59] C. Yan, L. Hao, A. Hussein, S. L. Bubb, P. Young, and D. Raymont. Evaluation of light-weight alsi10mg periodic cellular lattice structures fabricated via direct metal laser sintering. *Journal of Materials Processing Technology*, 214(4):856–864, 2014.
- [60] C. Yan, L. Hao, A. Hussein, P. Young, and D. Raymont. Advanced lightweight 316l stainless steel cellular lattice structures fabricated via selective laser melting. *Materials & Design*, 55:533–541, 2014.
- [61] C. Yang, O. Tiebe, K. Shirahama, and M. Grzegorzec. Object matching with hierarchical skeletons. *Pattern Recognition*, 55:183–197, 2016.
- [62] C.-Y. Zhao, W. Lu, and Y. Tian. Heat transfer enhancement for thermal energy storage using metal foams embedded within phase change materials (pcms). *Solar Energy*, 84(8):1402–1412, 2010.
- [63] Y. Zhou and A. W. Toga. Efficient skeletonization of volumetric objects. *IEEE Transactions on visualization and computer graphics*, 5(3):196–209, 1999.