Granular Matter manuscript No.
(will be inserted by the editor)

Morse Theory-based Segmentation and Fabric Quantification of Granular Materials

Karran Pandey · Talha Bin Masood · Saurabh Singh · Ingrid Hotz · Vijay Natarajan · Tejas G. Murthy

Received: date / Accepted: date

Abstract This article presents a robust Morse theory-based framework for segmenting 3D x-ray computed tomography image (CT) and computing the fabric, relative arrangement of particles, of granular ensembles. The framework includes an algorithm for computing the segmentation, a data structure for storing the segmentation and representing both individual particles and the connectivity network, and visualizations of topological descriptors of the CT image that enable interactive exploration. The Morse theory-based framework produces superior quality segmentation of a granular ensemble as compared to prior approaches based on the watershed transform. The accuracy of the connectivity network also improves. Further, the framework supports the efficient computation of various distribution statistics on the segmentation and the connectivity network. Such a comprehensive characterization and quantification of the fabric of granular ensembles is the first step towards a multiple length scale understanding of the behavior.

Keywords Granular ensemble · fabric · segmentation · particle connectivity network · Morse theory · Morse-Smale complex · topological persistence.

Karran Pandey · Vijay Natarajan
Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India.
E-mail: vijayn@iisc.ac.in

Talha Bin Masood · Ingrid Hotz
Department of Science and Technology (ITN), Linköping University, Norrköping, Sweden.
E-mail: ingrid.hotz@liu.se

Saurabh Singh · Tejas G. Murthy
Department of Civil Engineering, Indian Institute of Science, Bangalore, India.
E-mail: tejas@iisc.ac.in

1 Introduction

Understanding of the mechanical behavior of geomaterials from the micro-scale (*i.e.,* inter-particle interactions) to the macro-scale (or, continuum scale) has significantly benefited engineering of physical infrastructure. Such an understanding straddling multiple length scales has established beyond reasonable doubt that the ensemble mechanical response is an integration of the arrangement and interaction of individual particles in geomaterials such as sand, clay, structured sand, and structured clay [46, 34, 7]. Numerical techniques such as discrete element methods and imaging experiments using x-ray computed tomography (CT) have enormously contributed to this understanding of geomaterial behavior. Prior to use of x-ray CT, fabric and microstructure studies were limited to only two-dimensional packing [33, 34]. Access to CT has enabled extraction and quantification of the three-dimensional fabric and its evolution under different stress paths [11, 49, 20, 21, 3, 16, 1]. Within the overall workflow of quantification of fabric of soils, the method used for image segmentation crucially determines the quality of the results. In this paper, we present a novel approach for segmentation based on Morse theory, which is a branch of topology that studies the relationship between the connectivity of spaces and scalar functions defined on them. Crucial to the goal of quantification of fabric, this Morse theory-based approach simultaneously supports a data structure that efficiently stores both the segmentation and the connectivity network that represents the inter-particle arrangement.

1.1 Previous work

Fabric Quantification. Initial quantification of the fabric of sands using CT techniques has been reported by multiple authors [2, 3, 16, 24, 26]. More recently, the fabric of weakly cohered glass-ballotini has also been investigated [44, 43]. The basic premise of this CT-based analysis of fabric is the collection of images through the cross-section of a laboratory or natural sample, and concomitant image segmentation and analysis. Segmentation refers to the process of labeling or identification of different phases or objects present in the image. In this study, images of weakly cemented granular materials captured through x-ray CT are analyzed. The small amount of epoxy present mediates the contact between particles. Air voids present in this ensemble, are also segmented.

Once the particles and contacts are identified, the inter-particle arrangement is studied by computing the coordination number distribution, or more commonly by computing a fabric tensor. The *fabric tensor* is a measure of the orientation distribution of the particles [32]. Characteristics of the fabric tensor and its formulation have been enumerated in recent work [43].

Watershed-based segmentation. In previous work, including the above-mentioned studies, variants of the watershed algorithm [30, 37] have been typically used to segment the particles and to identify the contacts (or epoxy bonds). These watershed-based algorithms require as input a scalar field that contains, in its topology, exactly one catchment for each particle *i.e.,* the presence of exactly one field minimum / maximum corresponding to each particle. Further, [48] show the contact orientation obtained through conventional and random walker based watershed implementation is either error-prone or computationally expensive.

The watershed algorithm identifies and labels each catchment (particle) and its ridge line (epoxy bonds). The segmentation method begins by first binarizing the volume image corresponding to a particle-void gray level cutoff. Following this, the binarized particles are tested to ensure that they do not contain inherent flaws such as regions of low attenuation intensity. Next, a distance transform is applied on the complementary image to obtain a scalar field such that each particle contains exactly one maximum. Due to the irregular shape of the particles, the distance transform generates one dominant peak and several small peaks. These small peaks are removed by applying the h-maxima transform, which suppresses all the maxima (peaks) that are lower than the specified tolerance [45]. An appropriate threshold for the h-maxima transform is chosen manually by an expert user. This transform provides exactly one maximum for each particle. Next, the field is inverted such that each particle contains exactly one minimum. The labeled particles are computed by applying the watershed algorithm to the inverted modified distance field. This algorithm progressively fills the basins defined by each minimum from below by uniformly raising the "water level". Water from two adjacent basins meet along their common boundary, the separating ridge, that passes through the contact region. Figure 1 illustrates the different steps in the watershed-based approach and contrasts it with our proposed Morse theory-based approach thereby highlighting the benefits of the latter.

An important limitation of the watershed technique is that it fails for non-convex shaped particles, particularly those that contain a neck, and results in over-segmentation. A carefully chosen threshold for the h-maxima transform addresses this problem to some extent. However, a single threshold uniformly applied to the entire scalar field fails to capture all possible cases and the method often results in over-segmentation. Also, the quality of the result naturally depends on the manual choice of threshold for the h-maxima. Other extensions and variants have been proposed [50, 26] to handle the over-segmentation problem. These methods either rely on a single absolute user-specified threshold or adaptive local application of watershed, which is computationally inefficient and expensive.

Topology-based segmentation. Topological frameworks like discrete Morse theory [17] and persistent homology [13] enable effective tools for the study and analysis of scientific data due to their ability to robustly extract shape and structure even in the presence of noise [19, 18, 6]. The physical properties of porous and granular materials are closely related to the geometry and topology of the material structure. So, topological methods have been found to be useful in the study of diverse phenomena within this domain, ranging from the crystallization process in bead packings [38] to permeability in sandstone [22]. Persistent homology-based pipelines have been used for summarizing essential structural properties [38, 25] and Morse theory has been used to extract robust geometric representations of the material structure [19, 36, 10]. Robins et al. [36] and Freidrichs et al. [10] first described the benefits of discrete Morse theory for the segmentation and skeletonization of 3D images, including an application to the computation of the pore network of a silica sphere pack from 3D CT scan data. The focus of their work is on computing the Morse-Smale (MS) complex and hence the other steps of the CT scan analysis pipeline such as the boundary surface computation and geometry extraction are not discussed. Further, the use of a manually selected simplification threshold results in the lack of an intuitive

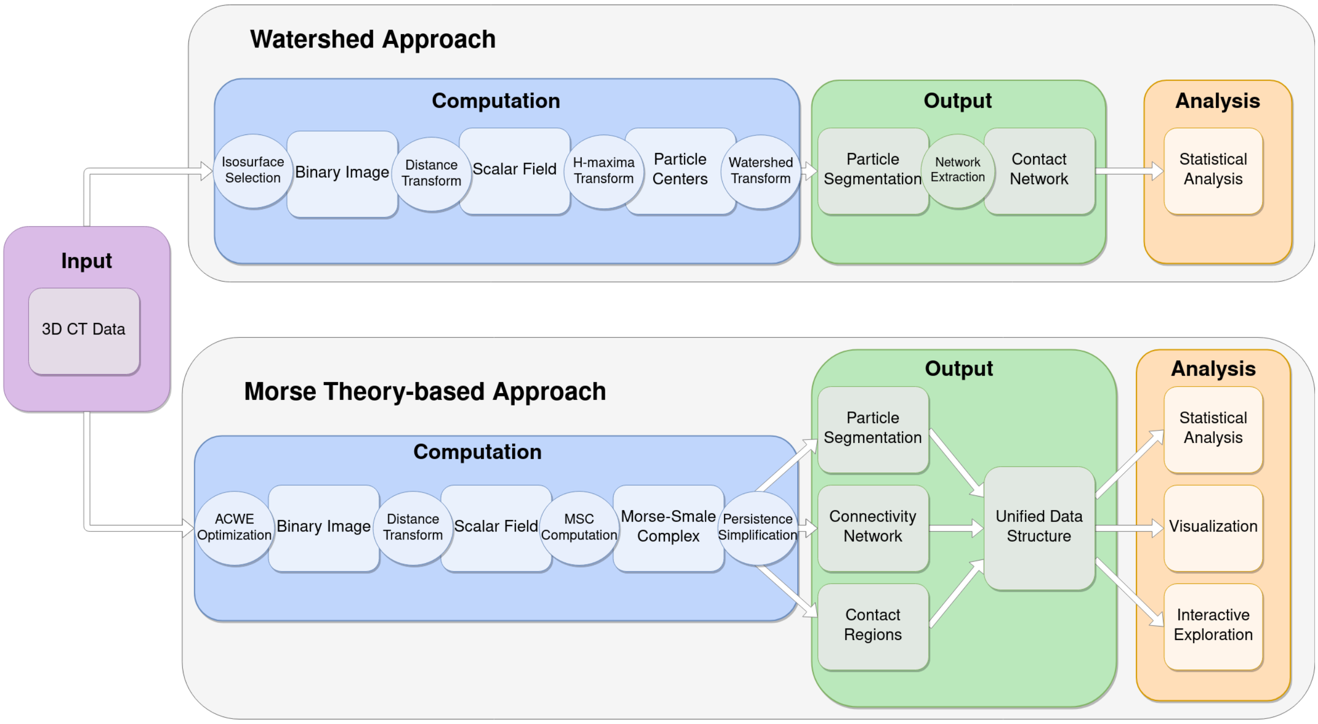

Fig. 1: Visual comparison of the watershed-based and the proposed Morse theory-based approaches to process and analyze 3D CT data. The comparison helps identify key differences between the two method pipelines in terms of computation, output, and supported analysis. The unified data structure retrieved from the MS complex supports various analysis tasks.

choice or guide for selecting noise simplification thresholds. Gyulassy et al. [19] also adapted Morse theory in their pipeline to extract the core filament structure of porous materials. Their curved skeleton-based representation is computed by carefully simplifying a distance field in order to identify and extract features of interest from the MS complex. They are able to identify consistent and stable core structures of porous solids, thus extracting meaningful results about the material properties. Their method design exemplifies the benefits of adapting Morse theoretic ideas to fit the context of an application focused pipeline.

Our method falls into the category of Morse theory-based approaches that aim to study material packing structures, but it is distinct from previous work on three essential points. First, our proposed method packages the benefits of discrete Morse theory as a part of a comprehensive and coherent pipeline for analysis of CT images with the objective of computing a robust geometric representation of the granular material structure. Second, the pipeline is designed to support a smooth and convenient end-user experience with intuitively tunable parameters that generalizes well to diverse use-cases. Third, our focus is on the study of the material phase of the packing structure of granular materials as opposed to the study of pore spaces in earlier work.

1.2 Contributions

This paper presents a new Morse theory-based framework for computing, storing, and exploring the fabric in geomaterials. We apply this framework on x-ray CT images of a mono-dispersed steel sphere packing for benchmarking and on an epoxy bonded sand packing to demonstrate its utility and advantages over existing approaches. The framework incorporates a new method for segmentation, introduces a new connectivity network that supports further statistical analysis, and a visualization tool that supports interactive exploration of the granular ensemble, see Figure 1. Key contributions of this paper include

- A robust algorithm for computing a segmentation of a 3D x-ray CT image. The segmentation is represented as the Morse complex of the scalar field corresponding to the CT image. Further, the algorithm is amenable to parallel computation.
- An effective method for simplification and noise removal. A topological persistence [15, 13] based

method removes both geometric and topological noise resulting in the desired segmentation. The method enables both a flexibility in the choice of integrated geometric and topological criterion for simplification and an intuitive choice of threshold for the criterion.
- An improved boundary surface computation method that constructs the necessary input for the segmentation algorithm. An active-contour [8] based method computes an optimal boundary surface of the particles in the packing. The surface is locally adapted to the best separating scalar value and is thus free of the assumption of a single global threshold.
- A unified data structure for storing both the segmentation and contacts. The data structure representing the Morse complex supports fast queries. Further, it enables efficient noise removal while guaranteeing consistency in the topological relationship between the segments.
- Construction of a geometrically meaningful connectivity network between contacting particles in the packing. The arcs of the network adapt to the geometry of the particles and are guaranteed to lie within the packing as they connect particle centers and contacts, thus accurately representing the packing structure even for complex non-convex particle geometries.
- An effective method for visualizing the segmentation and the connectivity network. The segmentation, contacts, and connectivity network are visualized via high quality rendering of the cells of the Morse complex.

Experimental results demonstrate that the segmentation results are superior to watershed-based methods, both in terms of segmentation quality and identification of contacts. The proposed method is suitable for both simple and complex CT granular datasets with non-convex morphology, and has the potential to generalize to broad range of particle sizes including those with particle breakage.

2 Challenges

In this section, we describe the main challenges faced in the process of analyzing 3D CT images using existing techniques. These challenges serve as the underlying motivation for the design of our approach.

Characteristics of reconstructed CT data. Reconstructed 3D CT images of granular ensembles can be very challenging to process because of the reconstruction artifacts, noise, and irregular geometry of the particles. The x-ray CT scans suffer from noise and artifacts, the incompatibility between the reconstructed values and their attenuation densities, due to the inconsistencies in x-ray source (polychromatic source), object (cylindrical shape of the object), detector, and reconstruction algorithm [23]. Figure 2 shows the reconstructed volume and a representative slice of the epoxy-bonded sand specimen. The background intensity and the reconstructed values of sand particles vary across the specimen due to artifacts and noise. The reconstructed image is affected by beam hardening — the average reconstructed values of air, epoxy, and sand are significantly higher near the periphery than the central region. The partial volume effect results in dark streaks near the periphery. We also observe occasional low reconstructed valued regions in sand due to flaws within the particles that manifest into intra-particle voids after thresholding. Further, the epoxy and pore air have similar attenuation coefficients that, together with the background noise, makes the segmentation of the epoxy phase challenging. In addition to reconstruction-based artifacts, the actual geometry of the particles in the packing too can often be highly non-convex, irregular, and show great variation in shape, size, and volume. This combination of artifacts and diversity in particle geometry necessitates specifically designed segmentation techniques that are not only inherently robust but make no assumptions about the particle geometry and distribution of noise.

Manual selection of boundary surface. A watershed-based segmentation method begins by computing a surface that bounds the particles. Otsu's method is the popular choice of boundary surface extraction method. It often reports a surface that fails to accurately capture the geometry of the packing. One approach to address this shortcoming is to manually explore the histogram of values in the CT scan and fine tune the choice of isosurface that represents the boundary surface. This manual approach, in addition to being tedious, may still fail to correctly capture the packing geometry. For instance, we highlight one such situation in Figure 3, where a volume rendering of the CT scan with a well tuned color and opacity map is used to illustrate the difference between the actual particle geometry and what is identified by a carefully selected isosurface. Further, we also observe that isosurfaces corresponding to different isovalues capture the correct packing geometry within different regions of the CT data. This indicates that a single uniform choice of isosurface is not sufficient to model the boundary surface geometry. An incorrect boundary surface has the potential to grossly affect downstream analysis, making it difficult to compute accurate segmentation and contact information.

Simplification and noise removal. The tendency of watershed-based segmentation to over-segment non-

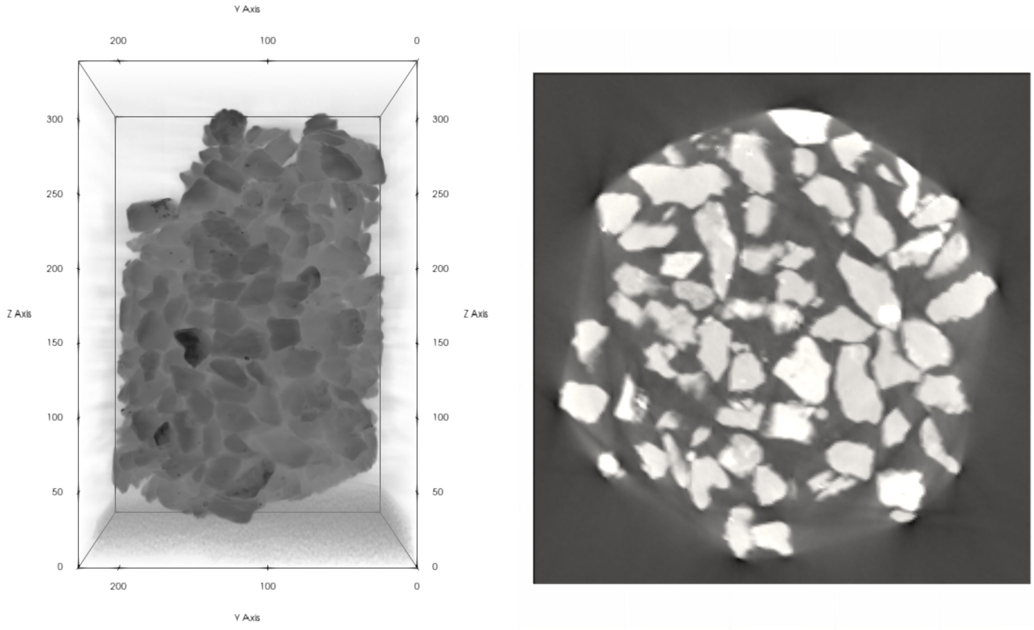

Fig. 2: Left: Direct volume rendering of CT scan of a sand sample using a grayscale color map and opacity which increases linearly with the CT intensity values. Note how the particle boundaries are not clear and there is noise in the data. Moreover, the sand particles have varying intensity, and some imaging artifacts appear as horizontal bands of varying intensity values. Right: A cross section of the CT volume highlights similar issues. Additionally, imaging artifacts at the periphery of the sample is apparent.

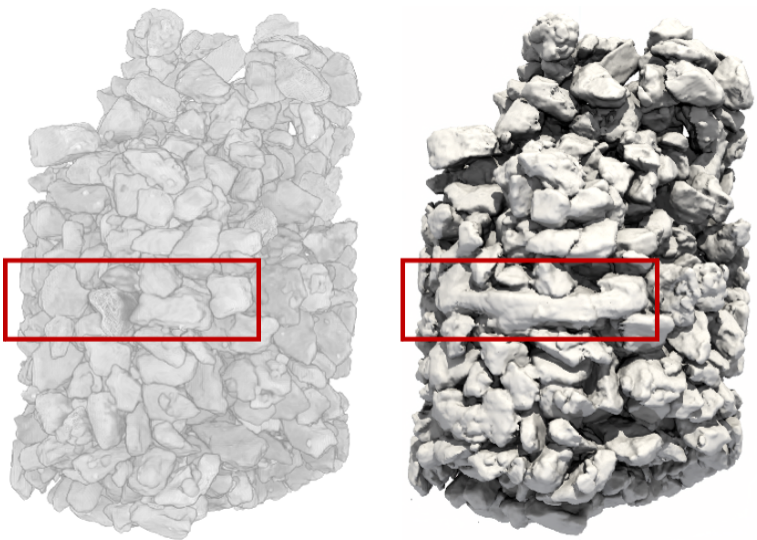

Fig. 3: Left: Direct volume rendering of the CT data using a fine-tuned color map. Right: Manually selected isosurface guided by a histogram. Notice how the isosurface does not agree well with the CT data, specially in the region enclosed by the red box.

convex geometry necessitates the employment of region-merging heuristics such as the h-maxima transform. The h-maxima transform requires the user to select the desired depth value up to which local maxima are to be suppressed. A good manual choice is essential for achieving good results. However, this choice of parameter is tedious and unintuitive. It is often difficult to choose a parameter value with a high level of confidence because evaluating the effect of this choice requires the computation of the entire segmentation followed by statistics

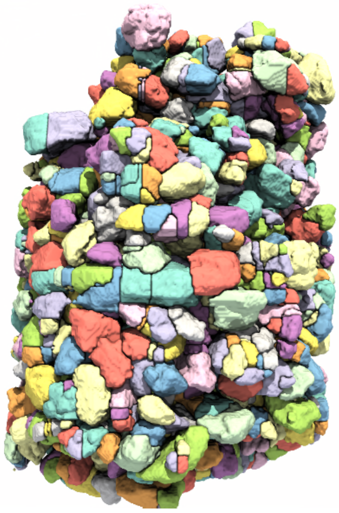

(a) h-maxima threshold=0.0

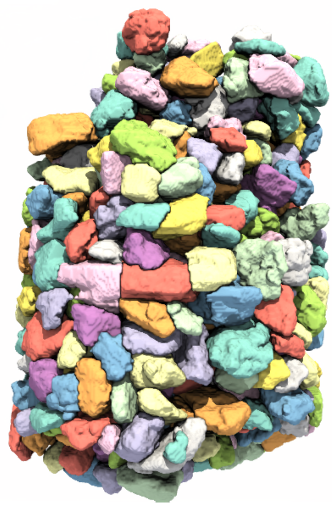

(b) h-maxima threshold=1.0

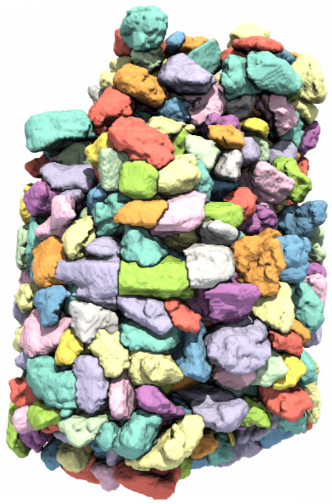

(c) h-maxima threshold=2.0

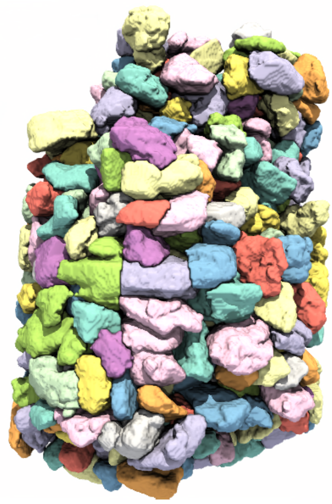

(d) h-maxima threshold=3.0

Fig. 4: In the watershed-based approach to segmentation, it is challenging to identify a "correct" h-maxima threshold. Note how the chosen threshold affects the final segmentation. Often, this threshold is chosen using a trial-and-error process. In some cases, the threshold identification is guided by prior knowledge of the sample, say the expected particle size or expected number of particles in the sample.

on the segmented particles. Further, small changes in h-maxima value may result in large variations in the segmentation, which implies that approximate solutions may not be satisfactory. In Figure 4, we illustrate the watershed segmentation of a granular ensemble at different h-maxima thresholds. Note how relatively similar h-maxima parameter values result in different segmentation results.

Network extraction. The natural step after computing the segmentation and a set of contacts between particles is to construct a connectivity network. To achieve this, the particles associated with each contact are identified by searching within the neighborhood (26 pixel connectivity) of each contact. A simple network, where connectivity of particles is represented by straight lines between particle centres that share a contact, is used to represent the packing. This is an abstract representation and does not represent the geometry of the contact or connectivity.

3 Background

This section reviews the necessary mathematical background for the topological methods described in this paper. We introduce terminology and key concepts on Morse functions, MS complex, and topological simplifica-

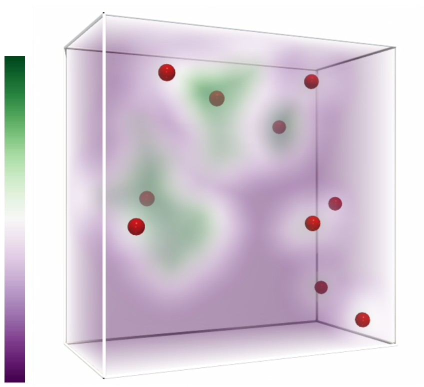

(a)

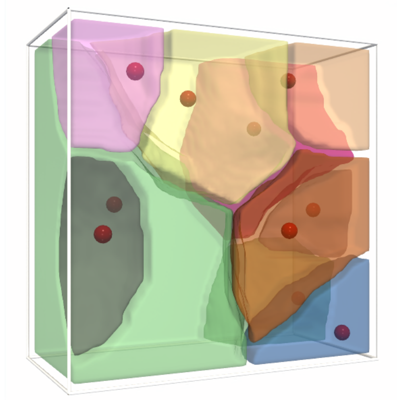

(b)

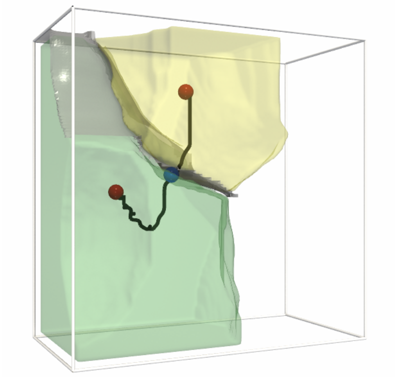

(c)

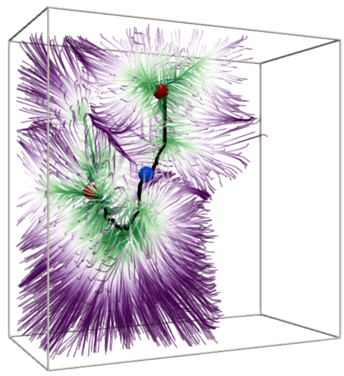

(d)

Fig. 5: Morse complex and its cells. (a) A synthetic scalar field shown using pseudo colors. Points identified as *maxima* are shown as red spheres. (b) Morse complex partitions the domain of the scalar field based on *descending manifolds* of maxima. (c) A pair of adjacent maxima are highlighted, their descending manifolds are shown in yellow and green. The *2-saddle* at the interface of these two maxima is shown as a blue sphere. The descending manifold of this 2-saddle is the grey surface at the interface between the descending manifolds of the two maxima. The bold black line is the *ascending manifold* of the 2-saddle which also connects the two maxima. (d) The integral lines within the descending manifolds of the two maxima. Notice how these lines converge towards the two maxima and form a separation surface around the 2-saddle. This structure and partition is robustly extracted using the MS complex.

tion, which are required for understanding the definition of our proposed connectivity network, the segmentation, and particle extraction methods described in the next section.

3.1 Morse function and Morse complex

Morse theory studies the relationship between the topological properties of a space and the critical points of a smooth real valued function defined on the space. We introduce the necessary terms from Morse theory and refer the reader to books on this topic for further details [28, 13].

Consider a smooth (twice differentiable) scalar function $f : \mathbb{R}^3 \to \mathbb{R}$. A point $p_c \in \mathbb{R}^3$ is called a critical point of f if the gradient of f at p_c is zero,

$$\nabla f_{p_c} = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z} \right)_{p_c} = \mathbf{0}. \tag{1}$$

The critical point p_c is called non-degenerate if its Hessian (matrix of second partial derivatives at p_c) is non-singular. The function f is called a *Morse function* if all its critical points are non-degenerate and have distinct function values. Any smooth function f can be infinitesimally perturbed into a Morse function. The number of negative eigenvalues of the Hessian corresponds to the *Morse index* of the critical point. Critical points of a three-dimensional function f can be of four types: minima (index-0), 1-saddle (index-1), 2-saddle (index-2), and maxima (index-3). An *integral line* is a curve in $\mathbb{R}^3$

whose tangent vector at a point on the curve is parallel to the gradient of f at that point. The integral line passing through a point p is the solution to

$$\frac{\partial}{\partial t} L(t) = \nabla f(L(t)), \forall t \in \mathbb{R} \quad (2)$$

with initial value $L(0) = p$. The function f monotonically increases or decreases along an integral line. The end points of an integral line corresponding to the limits as t approaches $-\infty$ and ∞, its origin and destination respectively, are critical points of f. The set of all integral lines originating at a critical point p_c together with p_c is called the *ascending manifold* of p_c. Similarly, p_c together with the set of all integral lines whose destination is p_c is called the *descending manifold* of p_c. Figure 5 shows a simple 3D scalar field, its critical points, and descending manifolds.

The descending manifolds (similarly, ascending manifolds) of all critical points of f partition the domain of f. This partition is called the *Morse complex* and consists of cells of dimensions 0, 1, 2, and 3. The descending manifold of a critical point with index i has dimension i. Therefore, the descending manifold of a maximum is a three dimensional volume (3-manifold), the descending manifold of a 2-saddle is a two dimensional sheet (2-manifold), the descending manifold of a 1-saddle is a one dimensional arc (1-manifold) and the descending manifold of a minimum is the minimum itself. Conversely, the ascending manifold of a critical point with index i has dimension $3 - i$. We propose the use of the Morse complex to represent the individual particles and the connectivity network between the particles.

The overlay of the ascending and descending manifolds results in a finer partition of the domain of f called the *Morse-Smale complex* (MS complex). Each cell of the MS complex is a set of integral lines that share a common origin and destination. The MS complex is a well studied topological structure and there are several efficient algorithms developed for computing hierarchical representations of the complex [40, 42]. More important in the context of this paper, the Morse complex can be extracted directly as a subset of the MS complex.

3.2 MS complex computation

CT scan data is available as a 3D density image, which can be modeled as a continuous scalar function defined on the vertices of a 3D cube grid and interpolated within each cell. The cube grid represents the domain of the scalar function. Ideas from Morse theory can be transferred to piecewise continuous functions to define critical points [15] via local comparison of function values. Further, a simulated perturbation ensures that neighboring vertices do not have equal function values and hence the gradient is zero only at critical points. This perturbation helps classify the critical points. Earlier algorithms employed numerical approaches to compute the MS complex and were affected by the errors in the computation of interpolants and derivatives. In contrast, we use combinatorial methods for MS complex computation. We use a parallel algorithm based on *discrete Morse theory*, which introduces a discrete analog of gradients as a directed edge between grid cells and computes integral lines as paths in a directed graph [40]. By avoiding derivative computations, this method ensures robustness of the results. In addition, the method provides theoretical guarantees for the correctness of adjacency relationship between cells of the output complex.

The 3D cube grid represents the domain as a collection of cells, a *cubical complex*. Voxels of the input are vertices (0-cells) of the grid; edges, faces, and cubes are the 1-, 2-, and 3-cells. The parallel algorithm of Shivashankar and Natarajan [40] computes discrete representations of the ascending and descending manifolds of the critical points of the function and supports operations for topological simplification as described in the following section. The descending manifold of a maximum is a collection of connected 3-cells in the grid, a volumetric region. The descending manifold of a 2-saddle is a collection of connected 2-cells resulting in a surface. The descending manifold of a 1-saddle is an arc consisting of a sequence of edges (1-cells) and the descending manifold of a minimum is a solitary 0-cell, namely the minimum. The ascending manifolds are computed similarly and represented as a collection of 3-cells (minimum), 2-cells (1-saddle), 1-cells (2-saddle), and 0-cell (maximum).

3.3 Topological simplification

A pair of critical points of f, that is connected via a single arc in the MS complex, may be removed via a *critical point pair cancellation*, resulting in a simpler MS complex [15, 14]. Morse theory guarantees the existence of a simpler function f' corresponding to the simpler MS complex. Pairs of critical points represent *topological features*. For example, the 2-saddle-maximum pairs in Figure 5c represents the two volumetric regions (yellow and green). We employ the critical point pair cancellation for topological simplification, iteratively removing small features that result from high frequency noise or sampling artifacts in the CT scan. Two crucial advantages of this approach are that (a) the algorithm is combinatorial and does not suffer from errors due to numerical computation, and (b) it supports controlled simplification where the data outside of a local neighborhood is not affected.

The cancellation changes both the combinatorial structure of the MS complex as well as the geometry of the ascending and descending manifolds of the surviving critical points. In this paper, we focus on 2-saddle-maximum cancellation because we are interested in computing the individual particle segments. Let p_2q_3 represent a pair of critical points where the subscript indicates the index of the critical point. First, p_2, q_3, and all arcs incident on them are deleted. Next, arcs incident on q_3 are routed to the surviving maximum that is connected to p_2 in the MS complex. Finally, the geometry of the descending manifold of q_3 is merged into the descending manifold of the surviving maximum that was connected to p_2.

The order of the cancellations plays an important role in determining the structure and geometry of the resulting simplified MS complex. Intuitively, given the aim of the simplification, the order should depend on the importance or size of a feature represented by a pair of critical points. While this importance can be defined in different ways, a commonly used and effective definition is the difference in function value between the pair of critical points. We iteratively apply critical point pair cancellation ordered by the difference in function value between the pair of adjacent critical points in the MS complex.

3.4 Persistence diagrams and persistence curves

An important practical consideration is the choice of the amount of simplification to be performed. The simplification threshold is defined as the maximum difference in function value between the canceled critical point pairs. This difference in function value is also referred to as the *persistence* of a critical point pair. Visual representations called persistence diagrams and persistence curves are often used to aid the choice of a threshold that separates noise from features of interest [13, 9]. A *persistence diagram* is a 2D scatter plot of the function values of the canceled critical point pairs. A *persistence curve* is a graph plot of the number of surviving critical points on the y-axis against an increasing simplification threshold on the x-axis. A dense collection of points very close to the diagonal in the persistence diagram is visually representative of a set of noisy features. A simplification threshold to remove such a collection of critical point pairs can be computed by locating the knee of the persistence curve, which corresponds to a sharp change in the slope of the curve.

4 Methodology

In this section, we describe our method for the analysis of 3D CT scans of material packing. The method first constructs a geometric and a topological structure – a segmentation of the packing into its constituent particles, and a topologically accurate contact-based connectivity network between the particles. The structures are constructed in two steps. First, we extract the material's boundary surface and compute its associated distance field. Second, we use the shape and connectivity information as captured by the gradient of the distance field to inform the segmentation and network extraction. Figure 1 presents an overview of the method while comparing it with the watershed-based approach and Figure 6 illustrates the individual steps of the method using a synthetic 3D dataset.

We begin with a description of the dataset followed by the individual steps of the method.

4.1 Dataset and preprocessing

We analyze a 3D CT scan of a packing of sand particles coated with epoxy at the contacts. The packing has a contact bound structure *i.e.,* the cementation exists only at the contacts in a skeleton of sand particles [44]. The scan has a dimension of 888 × 912 × 1360 voxels with a resolution of 12.5 µm per voxel. For the data to fit in memory, we first downsample the scan by a factor of 4 across each dimension. The downsampled image is computed using a simple mean operation across 4 × 4 × 4 sized cubes. We then run our analysis algorithms on the downsampled data of size 222 × 228 × 340. Further, the epoxy bonds are not segmented separately due to the weak contrast in the reconstructed values of epoxy and pore air. However, this does not affect the segmentation accuracy because the packing has a contact bound structure and the epoxy bonds form at the particle contact.

4.2 Boundary surface extraction

The material boundary surface is extracted using an automatic bi-modal threshold computation followed by an iterative local refinement of the corresponding isosurface.

We use Otsu's method [35] to compute the threshold for the initial boundary isosurface. Otsu's method takes as input the 3D image and essentially searches for an optimal threshold that divides the image into two classes - foreground and background. Optimality,

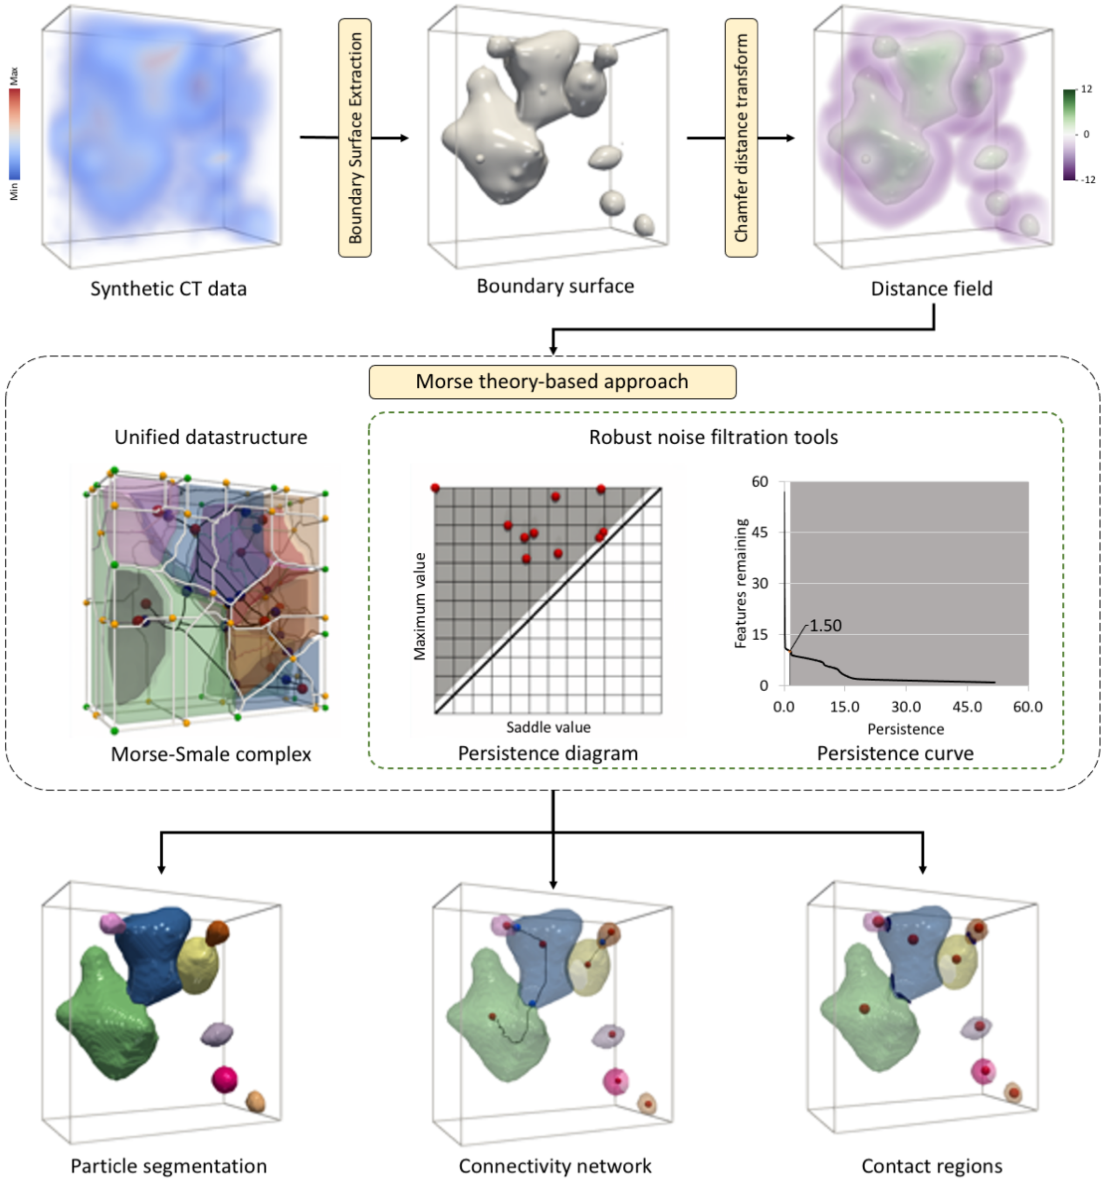

Fig. 6: Method pipeline. First, compute a bounding surface based on an automatically determined bi-modal threshold and a local active contour-based refinement. Next, compute a signed distance field from this surface and the MS complex for this scalar field. Simplify the MS complex via iterative cancellation of pairs of adjacent critical points in the MS complex. The difference in function value between the pair of critical points, called its persistence, is used to determine the order of cancellation. The threshold for simplification is determined by analyzing the persistence diagram. The segmentation is computed as the descending manifolds of the maxima in the simplified complex. Apply geometry driven pruning to compute the connectivity network.

in this search, is defined as a minimum intra-class variance. While the Otsu threshold-informed isosurface is a good starting point, attempting to pick a single scalar threshold that minimizes intra-class variance globally for the two classes may result in a surface that is locally non-optimal. Specifically, a global threshold restricts the search to isosurfaces of the function. The boundary surface is not necessarily modeled well by an isosurface, especially in regions where a locally better surface geometry is ignored due to the restriction and assumption of a single globally optimal isovalue. We therefore employ a local refinement procedure using an active contour model to construct a better boundary surface. Active contour models apply local updates on an input surface to optimize a given objective function. In particular, we adapt the Chan-Vese method [8], also known as the

Active Contours Without Edges (ACWE) method. The ACWE method, as opposed to using a gradient based objective function to extract a boundary surface, tries to minimize intra-class variance both inside and outside the surface within an iteration. The primary advantage of the ACWE method is its ability to robustly handle images with noisy boundaries without relying on numerical and potentially unstable gradient computations. We adapt the ACWE model to our setting by ignoring the smoothness and rigidity terms in the standard ACWE objective function to allow for sharp features in the boundary surface. As a result, our model effectively minimizes the following objective:

$$F(C, c_1, c_2) = \lambda_1 \int_{Inside(C)} |f(x,y,z) - c_1|^2 \, dxdydz + \lambda_2 \int_{Outside(C)} |f(x,y,z) - c_2|^2 \, dxdydz,$$

where C is the surface in the current iteration, f is the image intensity scalar function defined at each voxel, c_1 is the average intensity of voxels inside C, c_2 is the average intensity of voxels outside C, λ_1 is the normalizing constant for initial number of points and range of intensity inside C, and λ_2 is the normalizing constant for initial number of points and range of intensity outside C.

Our active contour framework successfully gives us high quality particle boundary surfaces across packings of a diverse set of materials and geometries.

4.3 Segmentation and connectivity network computation

Next, we segment and compute the connectivity network of the individual particles in the ensemble based on the geometry of the extracted boundary surface. A popular two-step approach for segmenting connected regions that touch each other first computes a distance transform from a bounding surface followed by a watershed transform [37], described in section 1.1. In watershed transform, the central point of the contact region between two catchment basins is a 2-saddle of the distance field. The 2-saddles are therefore natural representatives of the regions of contact between two particles in the packing. Watershed-based methods are however prone to over-segmentation of non-convex regions. Such regions may contain multiple local maxima of the distance field. So, it is necessary to design region merging protocols that fix the over-segmentation errors. In our work, we achieve robustness and computational efficiency in the two-step pipeline by using a fast Chamfer distance transform [27] and the topological framework of the MS complex. We use the MS complex for (a) characterizing and computing the catchment basins and regions of contact, and (b) to fix over-segmentation using an elegant combinatorial topological simplification procedure.

Initial catchment basin computation. The descending 3-manifolds of maxima of the distance field correspond to the catchment basins and aid in computing the segmentation. We retrieve the set of all descending manifolds of maxima from the Morse complex of the distance field. The set of these descending manifolds, further restricted to lie within the material boundary surface, is a segmentation of the packing into its individual particles. This step may produce over-segmented regions due to imaging noise or small variations in the distance field may introduce multiple local maxima within some non-convex regions.

Topological simplification. Next, we apply topological simplification to merge such regions and fix the over-segmentation. As described earlier, the topological simplification proceeds through an iterative cancellation procedure, canceling critical point pairs with persistence below a given simplification threshold. In our method, we restrict such cancellations to 2-saddle-maxima pairs. Intuitively, the cancellation operation essentially merges regions represented by low persistent maxima connected with the higher persistence regions. Persistence of a 2-saddle-maxima pair, can be loosely defined as the difference in function values, thus representing the difference between the radius of the contact (represented by the distance value of the 2-saddle) and the radius of the particle (represented by the distance value of the maxima). High persistent pairs are therefore particles where the contacts are much closer to the boundary surface than the central points in the particles. Conversely, low persistent pairs are particles where the width of the contact is comparable to the width of the entire particle thus implying that the particle may possibly be over-segmented.

We select a persistence threshold for simplification by identifying a knee in the persistence curve, a plot of decreasing persistence values over all critical point pairs. The presence of a clear knee in the curve indicates a clear separation of features from noise. Following noise removal via topological simplification, the quality of the segmentation can be further improved based on additional geometric criteria. For instance, we observed over-segmentation in particles with high-width contacts relative to particle size. We fix it by using the ratio of the distance field values at the 2-saddle and maximum. In particular, particles with the aforementioned ratio greater than 0.75 were often found to be over-segmented and were thus merged with their neighbor particle, as defined by the 2-saddle.

Contact identification. A 2-saddle of the distance field with a positive value represents a connection between two contacting particles because the descending manifolds of the maxima associated with the particles meet at the 2-saddle. The ascending 1-manifold of the 2-saddle is therefore a steepest ascent arc connecting the 2-saddle to the maxima associated with the contacting particles. The set of all ascending manifolds of 2-saddles is thus a geometrically meaningful connectivity network between the contacting particles in the packing. However, retrieving this entire set from the MS complex results in a number of spurious and clustered 2-saddles and multiple arcs connecting a pair of maxima. Since multiple contacts may exist between a pair of particles, it is important to distinguish between the primary 2-saddles that represent the central point within a contact region and spurious 2-saddles that are alternate representations of a given contact region. We achieve this by computing all explicit contact regions between two connected particles and picking the saddle with the highest function value from each such extracted region.

For each pair of particles, the contact regions are computed by clipping the descending 2-manifold of the highest 2-saddle. The descending 2-manifold is clipped by the boundary surface *i.e.,* all 2-cells containing vertices with distance value less than 0 are removed. The connected components of the remaining 2-cells represent the contact regions. While extracting the connected components, we mark the highest 2-saddle within each component and retain its arc in the connectivity network. We remove all other spurious arcs and therefore extract a robust contact-based connectivity network for the granular ensemble packing.

5 Implementation

In this section, we describe implementation details of the proposed method and the practical experience of executing the pipeline shown in Figure 6. We highlight the unified data structure that supports efficient and robust processing, and the principled simplification that results in good quality segmentation. From a user's perspective, executing the pipeline consists of three main steps:

Boundary extraction. The first step processes the input CT image, computes and stores the boundary surface that separates material from the surrounding volume. The surface is stored as a distance field, sampled at the same resolution as the input CT image. The user provides the CT scan as input and the software runs the active contour optimization, automatically computing the required parameters, and returns the distance field.

MS complex computation and simplification. The second step processes the distance field to construct the MS complex, computes the persistence curve, and displays it to the user. Guided by the curve, the user selects an appropriate persistence threshold. An appropriate choice is the knee in the curve as indicated in Figure 6. The iterative cancellation based simplification directed by a measure of importance such as difference in scalar function values simplifies all gradient based manifolds consistently. The persistence curve is a succinct and abstract visual representations of the distribution of extracted particle segments corresponding to a selected simplification measure. It provides a data-driven process to the user for selecting a simplification threshold. The MS complex is simplified based on the selected threshold to separate scan noise and concave protrusions from potentially important features.

The resulting simplified MS complex is stored. All ascending and descending manifolds are stored in a topologically consistent fashion. Specifically, three types of manifolds are computed – the 3D descending manifolds of maxima that represent the segmentation into individual particles, the 2D descending manifolds of 2-saddles that are processed to compute the contact regions between particles, and the 1D ascending manifolds of 2-saddles that represent the connectivity network. The computed manifolds are stored in a unified data structure that supports fast queries by separating the storage of the connectivity of the MS complex from its geometry. The nodes and arcs (0-cells and 1-cells) of the MS complex are stored as a graph and the geometry is extracted on demand. The data structure also enables flexible and efficient processing of the manifolds while guaranteeing that the topological relationships between the manifolds are maintained consistently. Specifically, it supports global simplification of the MS complex to remove noise and the local clustering of 2-saddles in the connectivity network using the contact regions.

The algorithms for computing the MS complex are amenable to parallel computation and execute on multi-core CPUs and GPUs resulting in small run times [40]. Note that seemingly time-intensive tasks such as computing the persistence curve are executed efficiently in our pipeline due to fast combinatorial algorithms that require only the connectivity of the MS complex for these steps.

Geometry extraction. The third step focuses on extracting the desired geometric and topological structures from the simplified MS complex. The user can optionally extract and store the segmentation, contact information, and the connectivity network. The segmentation consists of the segmented volume together with associated particle centres. The contact information consists of contact regions between a pair of particles and the associated contact points. The connectivity network connects the particle centers and contact points. This comprehensive collection of structures can then conveniently be used to facilitate user-desired downstream analysis.

In addition to the intuitive parameter-free pipeline, the code is easy to install and execute. The core algorithms used in the pipeline, including the boundary surface and discrete MS complex computation are available from existing software packages [47, 39, 29, 41]. We further plan to release our source code in the public domain for use by the community.

6 Benchmarking and Experimental Results

In this section, we present results and experiments to validate our approach and illustrate the key benefits of our method compared to standard watershed-based pipelines for the analysis of granular materials. The section is divided into three main parts. We first benchmark our method by computing physically expected statistics for a packing of mono-dispersed steel beads. After validating our approach on a well-studied dataset, we demonstrate the utility of the method for the analysis of complex packing structures by processing a cemented sand packing. We compute and visualize the segmentation and connectivity network for the packing and highlight observed statistical trends. Finally, we compare the quality of computed results and ease of computation with a recent watershed-based pipeline for the analysis of cemented granular materials [44]. All computational experiments were performed on a workstation with an Intel Xeon E5-1660 v4 @ 3.8 GHz processor with 8 cores, 128 GB main memory, and an Nvidia Quadro P2000 graphics card with 6 GB RAM. As described earlier, the MS complex is computed using a GPU parallel algorithm [41]

6.1 Benchmarking: mono-dispersed steel sphere packing

We use the tomography scan of a packing of mono-dispersed steel beads to benchmark our algorithm. We compare physically expected statistics like the average coordination number and diameter (in mm) in the packing to perform the benchmarking. The original scan has a dimension of $681 \times 681 \times 1004$. In order to segment this large volume, the analysis is carried out by splitting the volume into five blocks (1-201, 202-402, 403-603, 604-804, and 805-1004 along the height) and four intermediate blocks (131-271, 332-472, 533-673, 734-874 along the height). These blocks were chosen such that each particle is completely captured in at least one of the blocks. The center of each particle was remapped to the original volume and all duplicate entries and partially split particles were removed. Table 1 presents the average diameter and coordination number for both methods. The average diameters (0.98 mm and 1.0 mm) are consistent within the resolution ($16\mu m$) of X-ray CT. Theoretically, the average coordination number of a random packing of mono-dispersed spherical particles with infinite friction is known to lie in the interval [4,6], as calculated using Maxwell counting [4]. A deviation from the perfect spherical shape increases the average coordination number [12]. The average coordination number in our computation (6.04 and 6.08) is close to the theoretical range and within the experimentally obtained range [5].

Table 1: Comparing statistics for mono-dispersed steel beads.

	Diameter (mm)	Coordination number
MS Complex - High res. ($681 \times 681 \times 1004$)	0.98	6.04
Watershed - High res. ($681 \times 681 \times 1004$)	1.00	6.08

We also study the effect of data resolution by computing the statistics for downsampled versions of the data. We create a medium resolution ($341 \times 341 \times 502$) and a low resolution ($171 \times 171 \times 251$) version by downsampling by a factor of 2 and 4, respectively. Figure 7 shows the segmentation and connectivity network for medium resolution. The average diameters were found to be 0.96 mm and 0.91 mm for medium and low resolution, respectively. The difference from the actual diameter increases with decreasing resolution (error in the measurement), namely 32 μm (medium) and 64 μm (low). The average coordination numbers were 6.03 and 6.06 for medium and low resolution volumes. These results are also consistent with the watershed results at the original resolution. This indicates that our method is robust to downsampling for the computation of average statistics.

6.2 Visualizations of extracted geometric and topological structures

In this section, we demonstrate our proposed method's potential for analyzing complex packing structures by

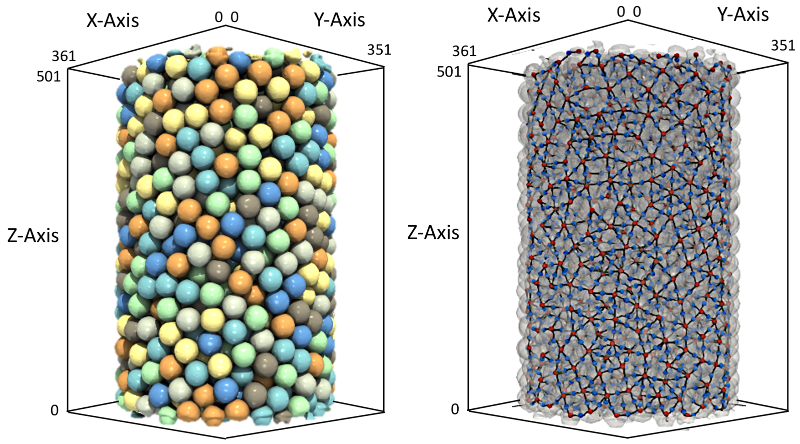

Fig. 7: The segmentation (left) and connectivity network (right) for the mono-dispersed steel bead packing.

processing a cemented sand packing with varied particle geometries. From the computed MS complex, we first extract the following geometric and topological structures:

- Segmentation of the packing structure into its constituent individual particles.
- Topological network of contact-based connections between particles, embedded in the domain using geometrically meaningful arcs.
- Geometric regions of contact between connected particles.

Access to these structures facilitates both visual exploration and statistical analysis of the material packing structure. In particular, the computed segmentation and connectivity network is used to generate high-quality visualizations of the individual particles in the packing (Figure 8) and their contact structure (Figure 9).

The MS complex data structure and persistence-based visualizations of the extracted segments (see Figure 11) further allow for the convenient selection and visualization of specific segments and contacts of interest. The selection may be directed by a variety of geometric and topological criteria such as the coordination number, volume, and relative width of the particle compared to the contact region. Figure 8 (right) and Figure 9 (right) highlight visualizations of such selected particles and Figure 10 illustrates the visualization of a selected contact and its extracted properties.

6.3 Statistical Analysis

In addition to localized exploration and analysis, our framework supports robust and efficient computation of detailed statistical metrics to study global trends in the packing structure. Figure 12 presents a comparison between the statistics. We highlight the similarities and differences between the histograms of the coordination number (with and without multiplicity), volume, and the orientation of contacts for individual particles in the packing. We report statistics computed using watershed-based segmentation computed using both manual and active contour based boundary surfaces. The histogram of coordination number with the inclusion of multiple contacts is significantly different from the watershed-based method. The histogram computed by excluding multiple contacts show similar trends, for example the most frequent coordination number is observed in the range of 4 to 6. The distribution of volume of particles is similar but there is a small difference in the actual counts. Both methods predict that most contacts are oriented along the z-direction (direction of gravity); this observation is consistent with previous studies on fabric of granular materials and cemented granular materials [31, 44]. The differences in counts of histograms is due to over-segmentation of the particles by the watershed-based method (discussed in section 6.4).

We also present a spherical histogram representation of the sand packing obtained from this analysis, see Figure 13. Details on creating these spherical histograms are provided elsewhere [43]. Further details on the quantification of the fabric of this cemented sand ensemble is presented in the supplementary material.

6.4 Watershed comparisons

In this section, we present a detailed comparison between our method and a recently proposed watershed-based pipeline for the analysis of cemented granular materials [43]. Using the cemented sand packing as a case study, the key benefits of our method are highlighted below in a comparison divided across the fundamental steps

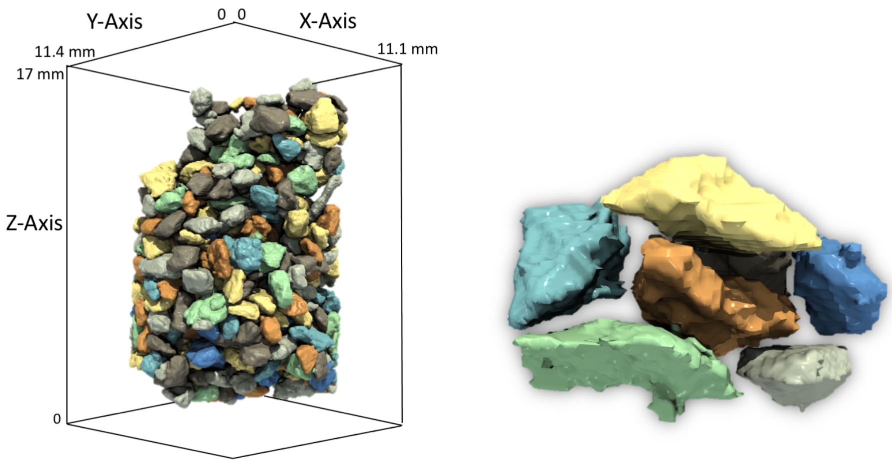

Fig. 8: Global (left) and local (right) visualization of the segmentation of the packing into its constituent individual particles. Particles are colored by their segmentation identifiers.

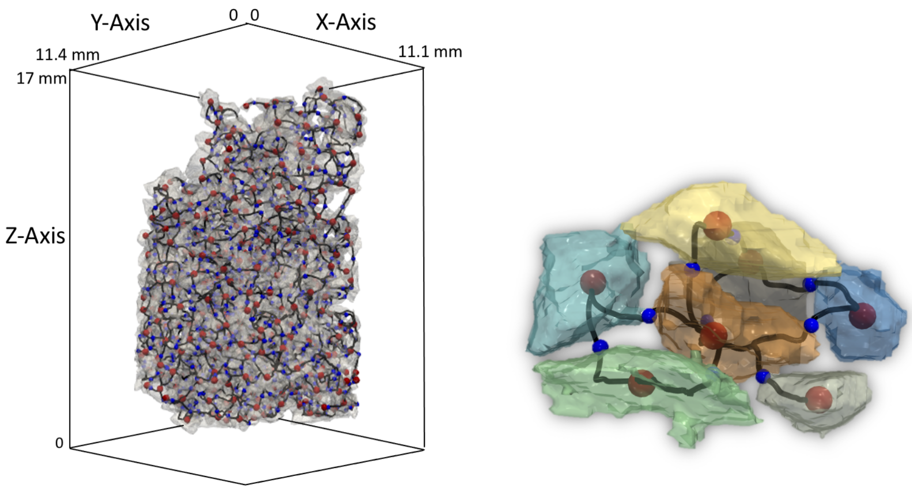

Fig. 9: Global (left) and local (right) visualization of the geometric embedding of the extracted connectivity network of the packing. In the local view (right), bead centres are represented as red spheres and contacts as blue spheres.

of boundary extraction, segmentation, and connectivity network extraction.

Boundary surface extraction. The histogram of scalar values is often used to select a global scalar threshold that separates the scan into regions containing material and the surrounding medium. A good separating surface consists of a collection of surface pieces that are defined by different scalar thresholds within different regions of the scan. As a result, a boundary surface that is defined by a single threshold contains noticeable artifacts that range from incorrectly merged particles to highly distorted particle geometries.

In comparison, our method's active contour-based optimization computes an optimal boundary surface (Figure 14) that locally adapts to the best separating value, and is thus free of the assumption of a single global threshold. Figure 14 highlights artifacts in a boundary surface that is computed using a global threshold and compares it against the improved surface computed using our active contour-based approach. The higher quality of the boundary surface results in an improvement in the quality of downstream analysis.

Segmentation. The difference in quality of the extracted boundary, coupled with the watershed transform's tendency to over-segment in the presence of noise and non-convex geometry, causes a tangible quality difference between the segmentation results of our Morse theory-based approach and watershed-based methods. As discussed in Section 2, post-processing heuristics such as the h-maxima transform merge over-segmented regions based on a given threshold. However, a precise and optimal choice of threshold for the transform is

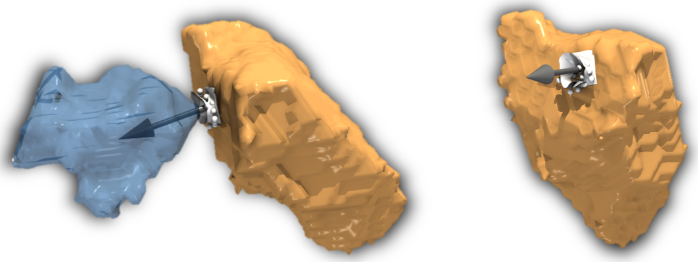

Fig. 10: Extracted region and orientation of contact between two sand particles colored in blue and brown. The translucent gray surface represents a clip of the extracted boundary surface of the particles near the contact region. This surface is the descending 2-manifold of the 2-saddle that defines the contact. Constituent points of the contact region are represented as white spheres, with the arrow displaying a fitted normal vector to the contact.

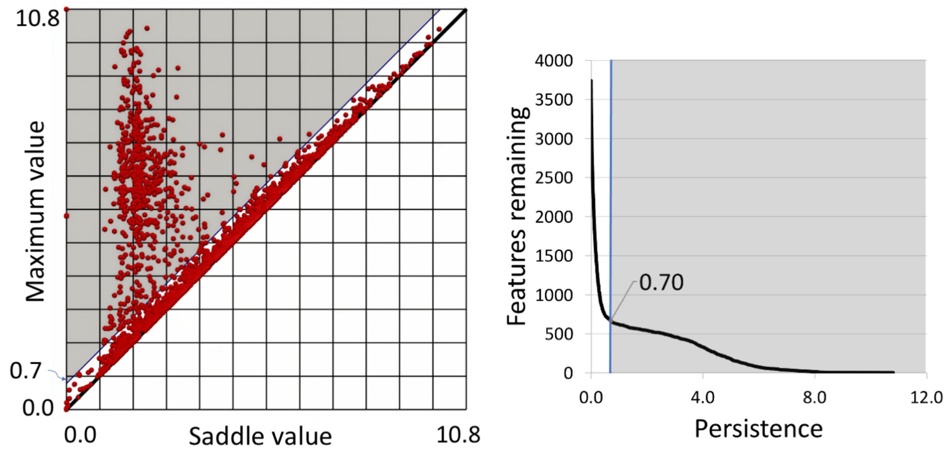

Fig. 11: Persistence diagram (left) and curve (right) for the cemented sand packing. The knee in the persistence curve (0.7) informs the choice of a noise threshold that separates features from noise. We can see how the selected threshold removes the cluster of low-persistence noise (white region) from the remaining features (grey region) in the persistence diagram.

elusive, and thus this merging process can potentially both struggle to correct over-segmentation or result in under-segmentation.

Our persistence-based framework alleviates this inconvenience by providing an intuitive and statistically optimal choice of threshold for simplification. In particular, we use the knee of the persistence curve (see Figure 11) to identify a precise and well-motivated choice of threshold. The persistence curve is a plot of the number of segments for increasing values of simplification threshold. It can be computed efficiently due to the unified data structure, which supports simple and efficient updates. Computing a similar curve for an h-maxima transform requires the repeated application of the transform for a large number of values while counting the segments, and still arriving at an approximation of the persistence curve. Figure 15 shows the watershed segmentation using a manually selected h-maxima threshold, which is identified via a visual inspection process followed by an expert, and highlights the improvement in segmentation achieved using the automatically computed persistence curve-based threshold. The segmentation quality is better both in terms of the geometry of individual particles and in terms of fewer cases of under- and over-segmentation.

The quality of segmentation also implicitly influences the accurate identification of contacts in the packing. While geometric inaccuracy in a segmentation due to a non-optimal boundary surface can affect the identified geometric position of the contacts, issues such as under-segmentation and over-segmentation can result in missing or spurious additional contacts. For instance, true contacts may be missed due to under-segmentation, and over-segmentation of a single particle leads to spurious contacts between incorrectly identified segments. Figure 16 compares the set of contacts identified by the

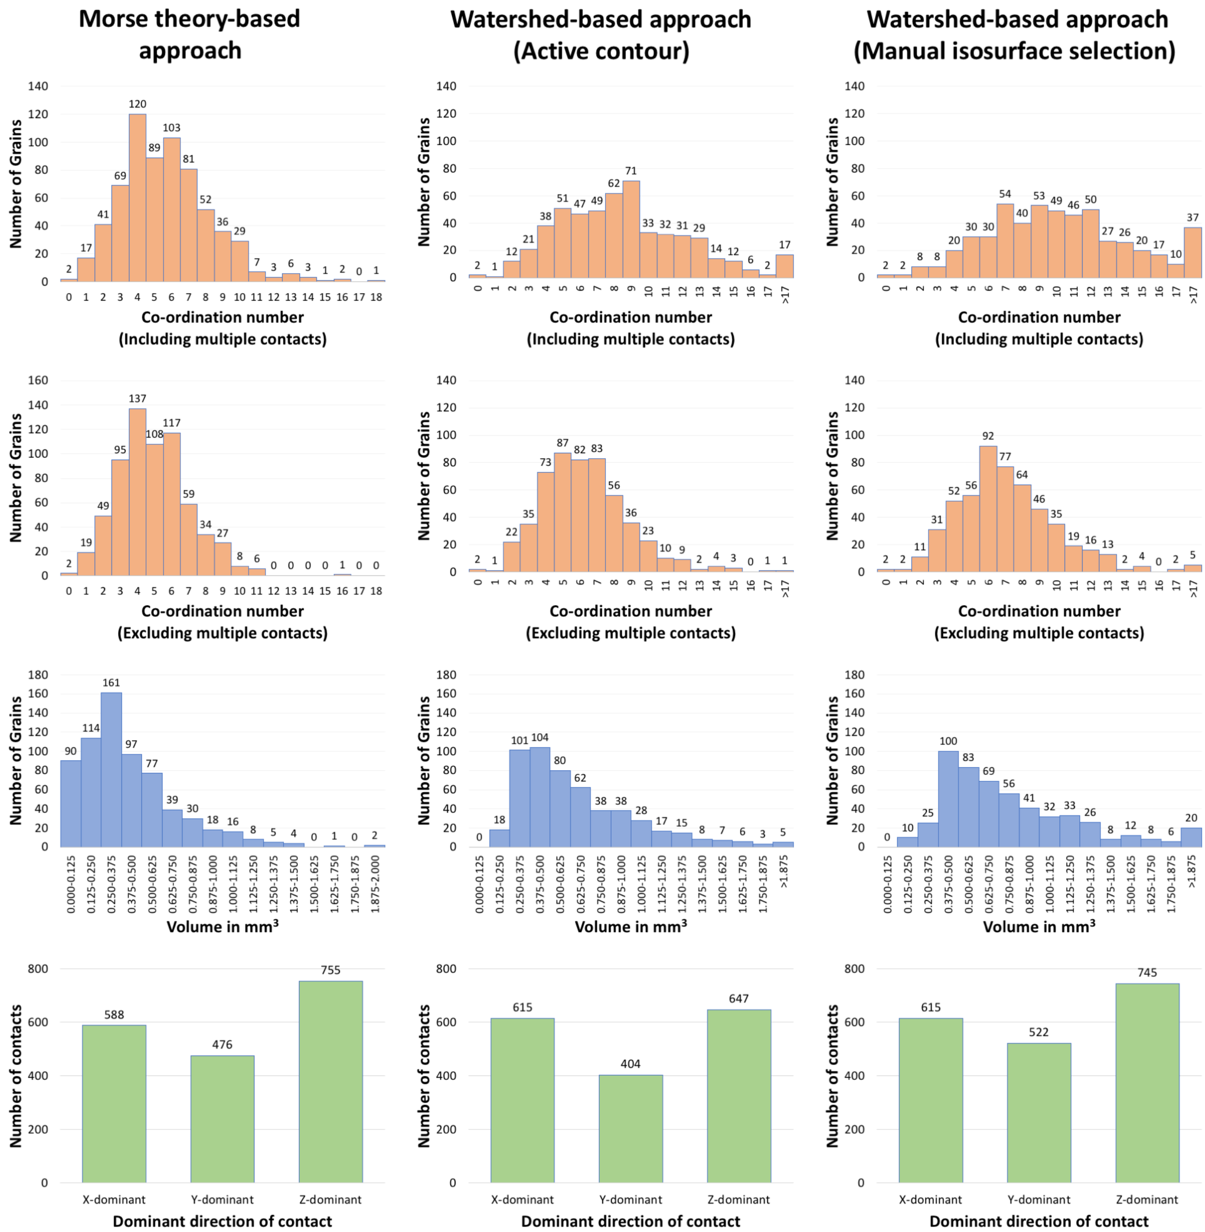

Fig. 12: Statistical analysis of sand granular ensemble computed using the Morse theory-based and watershed-based methods. Histogram of coordination numbers is similar when counting contacts with multiplicity, else the histograms are similar. The distribution of volume of particles is similar but the actual counts are different.

two methods by plotting them on a manually tuned volume rendering of the CT scan. The volume rendering helps us visualize key differences in the contacts identified by the two methods. The highlighted circles bring to attention particularly problematic regions, where either essential contacts are missed or are incorrectly placed or spurious additional contacts are identified by the watershed-based approach. Further, the large disparity in the number of identified contacts highlights the magnitude of difference in results between the two approaches.

An additional benefit of our method is that the criterion used for simplification can be suitably modified to accommodate different packing structures with variation in particle size and geometry. The topological simplification step naturally supports alternative criteria for iterative region-merging. The persistence diagram can be used as an intuitive visual tool to explore the distribution of the packing based on the selected criteria,

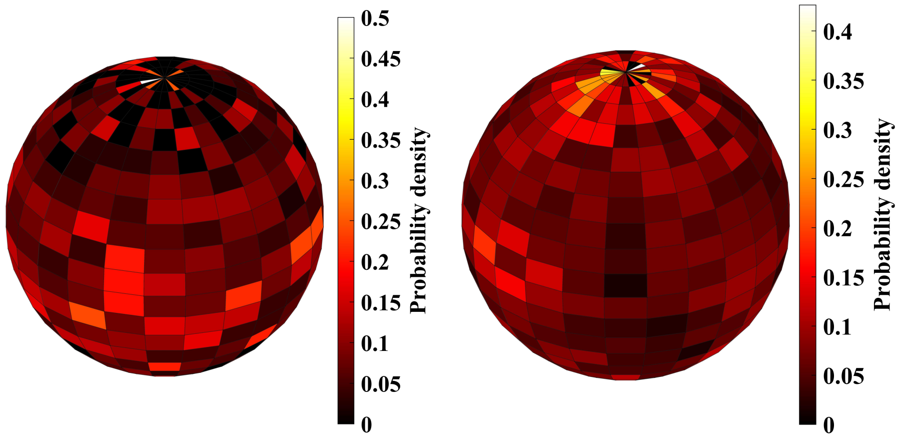

Fig. 13: (left) Normalized particle stereographic histogram and (right) normalized contact spherical histogram for sand packing.

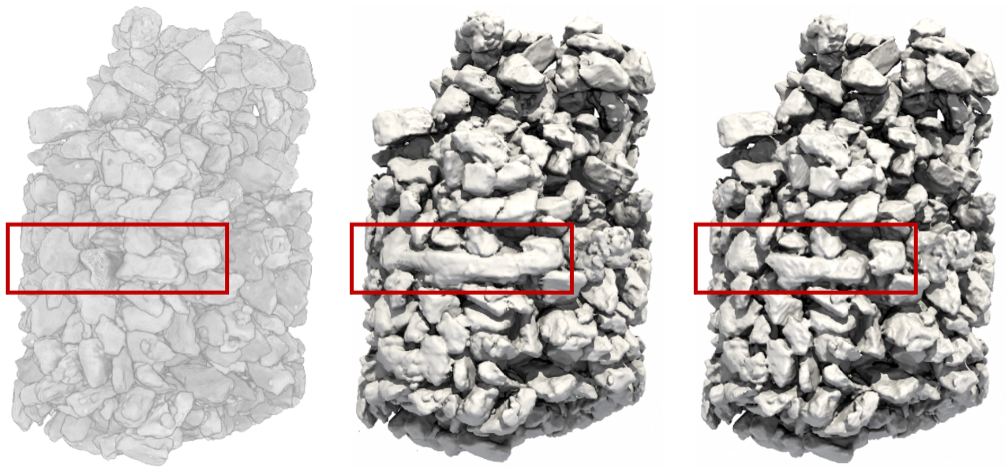

Fig. 14: High quality boundary surface extraction. A side-by-side visualization of a volume rendering of the CT scan (left) along with the boundary surfaces extracted using the isosurface method (middle) and active contour optimization (right). The inset highlights a region containing multiple artifacts in the isosurface boundary. The geometry of an elongated particle is not captured correctly, it merges with other particles in the packing. The volume rendering is generated using a manually tuned color and opacity map that delineates the individual particles.

thus providing a unified framework for the intuitive exploration of diverse choices of criterion and associated thresholds.

Connectivity network extraction. In order to represent the contact structure, existing methods use a network of straight lines connecting particle centers and contacts. However, in the case of non-convex particle geometries, such a network fails to account for the geometry of the packing [44]. For example, arcs of the network that are represented as line segments may lie outside the individual particles.

Our method instead uses the gradient of the distance field to compute a geometry-aware connectivity network that resides within the packing structure, see Figure 9. The combination of the discrete Morse theory-based gradient and the unified persistence-based simplification allows us to efficiently compute a numerically robust connectivity network that accurately captures the global packing structure.

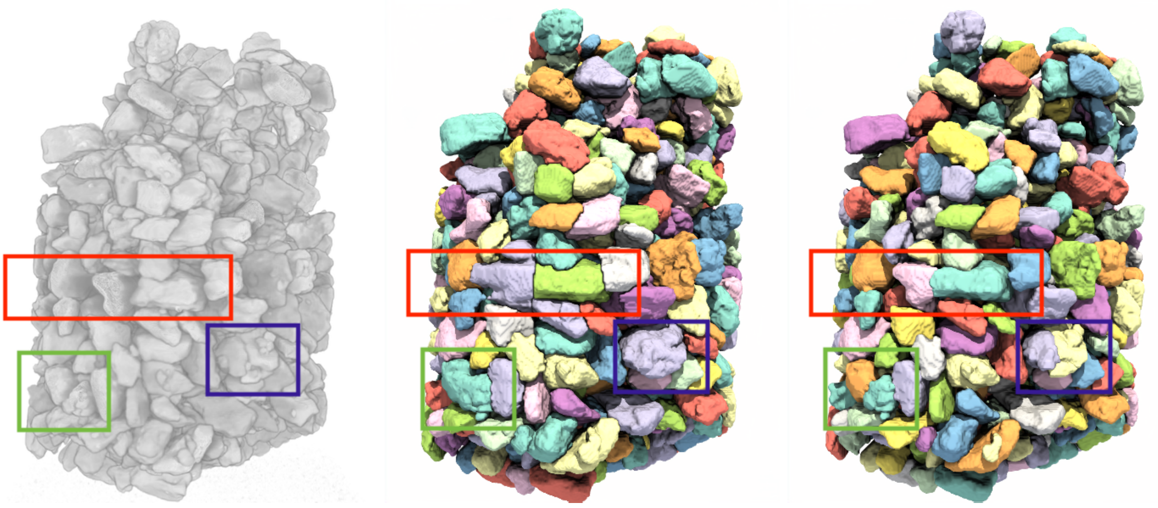

(a) Direct volume rendering (b) Watershed-based approach with h-maxima threshold = 2.0 (c) Morse theory-based approach

Fig. 15: Improved segmentation. Comparison of particle segmentation obtained using the watershed and our Morse theory-based approach. Red box: Geometry of the segments computed using the Morse complex agree better with the original CT data thanks to the improved quality of the boundary surface. Green and Blue boxes: Multiple cases of under segmentation in the watershed approach. In both cases, two particles are incorrectly identified as a single particle whereas the Morse complex identifies them correctly. The volume rendering uses a manually tuned color and opacity map to delineate individual particles.

7 Conclusions

In summary, we outline a robust and efficient approach to compute the segmentation and connectivity network of granular material packings from x-ray CT images. Through the combination of active-contour optimization and Morse theory, our approach functions as a unified and fully-automated framework to compute, query, and visualize the geometric and topological properties of the packing structure. The automated nature of the algorithm allows for convenient large-scale computation and the persistence-based simplification approach enables the intuitive exploration of variation in particle geometries. Salient features of our method, such as the locally-optimal boundary extraction, an efficient and flexible framework for noise removal and geometry-aware connectivity network, together ensure that the approach generalizes to diverse packing structures with variation in particle size, geometry, and material density.

Acknowledgements This work is supported by an Indo-Swedish joint network project (DST/ INT/ SWD/VR/P-02/2019) and Swedish Research Council (VR) grant 2018-07085 and the VR grant 2019-05487. KP and VN are partially supported by a Swarnajayanti Fellowship from the Department of Science and Technology, India (DST/SJF/ETA-02/2015-16) and a Mindtree Chair research grant.

Conflict of interest

The authors declare that they have no conflict of interest

References

1. Andò, E.: Experimental investigation of microstructural changes in deforming granular media using x-ray tomography. Theses, Université de Grenoble (2013)
2. Andò, E., Hall, S.A., Viggiani, G., Desrues, J., Bésuelle, P.: Experimental micromechanics: grain-scale observation of sand deformation. Géotechnique Letters **2**(3), 107–112 (2012). DOI 10.1680/geolett.12.00027
3. Andò, E., Hall, S.A., Viggiani, G., Desrues, J., Bésuelle, P.: Grain-scale experimental investigation of localised deformation in sand: a discrete particle tracking approach. Acta Geotechnica **7**(1), 1–13 (2012). DOI 10.1007/s11440-011-0151-6
4. Baule, A., Morone, F., Herrmann, H.J., Makse, H.A.: Edwards statistical mechanics for jammed granular matter. Reviews of modern physics **90**(1), 015006 (2018)
5. Bernal, J., Mason, J.: Packing of spheres: co-ordination of randomly packed spheres. Nature **188**(4754), 910–911 (1960)

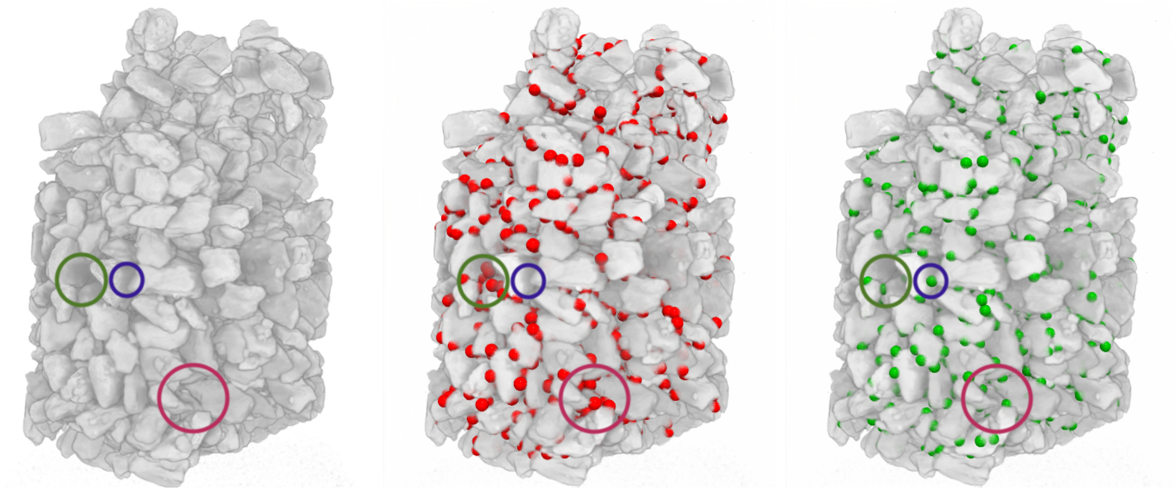

(a) Direct volume rendering (b) Watershed-based approach with h-maxima threshold = 2.0 (c) Morse theory-based approach

Fig. 16: Improved identification of contacts. Comparison of particle contacts obtained using the watershed and our Morse theory-based approach. A total of 2746 contact points are identified using watershed approach while 1842 points are identified by our method. Blue circle: The watershed-based method fails to identify the contact point. Red and green circles: The watershed-based approach incorrectly identifies a large number of contact points. The contacts identified by our approach appear to be reasonable when compared to the volume rendering.

6. Bhatia, H., Gyulassy, A.G., Lordi, V., Pask, J.E., Pascucci, V., Bremer, P.T.: Topoms: Comprehensive topological exploration for molecular and condensed-matter systems. Journal of computational chemistry **39**(16), 936–952 (2018)
7. Burland, J.B.: On the compressibility and shear strength of natural clays. Géotechnique **40**(3), 329–378 (1990). DOI 10.1680/geot.1990.40.3.329
8. Chan, T., Vese, L.: An active contour model without edges. In: International Conference on Scale-Space Theories in Computer Vision, pp. 141–151. Springer (1999)
9. Chazal, F., Michel, B.: An introduction to topological data analysis: fundamental and practical aspects for data scientists. arXiv preprint arXiv:1710.04019 (2017)
10. Delgado-Friedrichs, O., Robins, V., Sheppard, A.: Skeletonization and partitioning of digital images using discrete morse theory. IEEE transactions on pattern analysis and machine intelligence **37**(3), 654–666 (2014)
11. Desrues, J., Chambon, R., Mokni, M., Mazerolle, F.: Void ratio evolution inside shear bands in triaxial sand specimens studied by computed tomography. Géotechnique **46**(3), 529–546 (1996). DOI 10.1680/geot.1996.46.3.529
12. Donev, A., Connelly, R., Stillinger, F.H., Torquato, S.: Underconstrained jammed packings of nonspherical hard particles: Ellipses and ellipsoids. Physical Review E **75**(5), 051304 (2007)
13. Edelsbrunner, H., Harer, J.: Computational Topology: An Introduction. American Mathematical Soc. (2010)
14. Edelsbrunner, H., Harer, J., Natarajan, V., Pascucci., V.: Morse-Smale complexes for piecewise linear 3-manifolds. In: Proc. 19th Ann. Symposium on Computational Geometry, pp. 361–370 (2003)
15. Edelsbrunner, H., Harer, J., Zomorodian, A.: Hierarchical Morse complexes for piecewise linear 2-manifolds. In: Proceedings of the seventeenth annual symposium on Computational geometry, pp. 70–79 (2001)
16. Fonseca, J., O'Sullivan, C., Coop, M.R., Lee, P.D.: Non-invasive characterization of particle morphology of natural sands. Soils and Foundations **52**(4), 712 – 722 (2012). DOI https://doi.org/10.1016/j.sandf.2012.07.011
17. Forman, R.: A user's guide to discrete morse theory. Sém. Lothar. Combin **48**, 35pp (2002)
18. Günther, D., Boto, R.A., Contreras-Garcia, J., Piquemal, J.P., Tierny, J.: Characterizing molecular interactions in chemical systems. IEEE transactions on visualization and computer graphics **20**(12), 2476–2485 (2014)
19. Gyulassy, A., Duchaineau, M., Natarajan, V., Pascucci, V., Bringa, E., Higginbotham, A., Hamann, B.: Topologically clean distance fields. IEEE Transactions on Visualization and Computer Graphics

13(6), 1432–1439 (2007)
20. Hall, S.A., Bornert, M., Desrues, J., Pannier, Y., Lenoir, N., Viggiani, G., Bésuelle, P.: Discrete and continuum analysis of localised deformation in sand using x-ray micro-CT and volumetric digital image correlation. Géotechnique **60**(5), 315–322 (2010). DOI 10.1680/geot.2010.60.5.315
21. Hasan, A., Alshibli, K.: Three dimensional fabric evolution of sheared sand. Granular Matter **14**(4), 469–482 (2012). DOI https://doi.org/10.1007/s10035-012-0353-0
22. Herring, A., Robins, V., Sheppard, A.: Topological persistence for relating microstructure and capillary fluid trapping in sandstones. Water Resources Research **55**(1), 555–573 (2019)
23. Hsieh, J.: Computed tomography: principles, design, artifacts, and recent advances, vol. 114. SPIE press (2003)
24. Imseeh, W.H., Druckrey, A.M., Alshibli, K.A.: 3d experimental quantification of fabric and fabric evolution of sheared granular materials using synchrotron micro-computed tomography. Granular Matter **20**(2) (2018). DOI 10.1007/s10035-018-0798-x
25. Kader, M., Brown, A., Hazell, P., Robins, V., Escobedo, J., Saadatfar, M.: Geometrical and topological evolution of a closed-cell aluminium foam subject to drop-weight impact: An x-ray tomography study. International Journal of Impact Engineering **139**, 103510 (2020)
26. Kong, D., Fonseca, J.: Quantification of the morphology of shelly carbonate sands using 3d images. Géotechnique **68**(3), 249–261 (2018). DOI 10.1680/jgeot.16.P.278
27. Krissian, K., Westin, C.F.: Fast and accurate redistancing for level set methods. Computer Aided Systems Theory (EUROCAST'03) pp. 48–51 (2003)
28. Matsumoto, Y.: An Introduction to Morse Theory. Amer. Math. Soc. (2002). Translated from Japanese by K. Hudson and M. Saito
29. McCormick, M.M., Liu, X., Ibanez, L., Jomier, J., Marion, C.: Itk: enabling reproducible research and open science. Frontiers in neuroinformatics **8**, 13 (2014)
30. Meyer, F., Beucher, S.: Morphological segmentation. Journal of Visual Communication and Image Representation **1**(1), 21–46 (1990)
31. Oda, M.: Initial fabrics and their relations to mechanical properties of granular material. Soils and foundations **12**(1), 17–36 (1972)
32. Oda, M.: Fabric tensor for discontinuous geological materials. Soils and Foundations **22**(4), 96–108 (1982). DOI 10.3208/sandf1972.22.4_96
33. Oda, M., Konishi, J., Nemat-Nasser, S.: Experimental micromechanical evaluation of strength of granular materials: Effects of particle rolling. Mechanics of Materials **1**(4), 269–283 (1982). DOI 10.1016/0167-6636(82)90027-8
34. Oda, M., Nemat-Nasser, S., Konishi, J.: Stress-induced anisotropy in granular masses. Soils and Foundations **25**(3), 85–97 (1985). DOI 10.3208/sandf1972.25.3_85
35. Otsu, N.: A threshold selection method from gray-level histograms. IEEE transactions on systems, man, and cybernetics **9**(1), 62–66 (1979)
36. Robins, V., Wood, P.J., Sheppard, A.P.: Theory and algorithms for constructing discrete morse complexes from grayscale digital images. IEEE Transactions on pattern analysis and machine intelligence **33**(8), 1646–1658 (2011)
37. Roerdink, J.B., Meijster, A.: The watershed transform: Definitions, algorithms and parallelization strategies. Fundamenta informaticae **41**(1, 2), 187–228 (2000)
38. Saadatfar, M., Takeuchi, H., Robins, V., Francois, N., Hiraoka, Y.: Pore configuration landscape of granular crystallization. Nature communications **8**(1), 1–11 (2017)
39. Schroeder, W., Martin, K., Lorensen, B.: The visualization toolkit, 4th edn. kitware. New York (2006)
40. Shivashankar, N., Natarajan, V.: Parallel computation of 3D Morse-Smale complexes. Computer Graphics Forum **31**(3), 965–974 (2012)
41. Shivashankar, N., Natarajan, V.: Efficient software for programmable visual analysis using Morse-Smale complexes. In: Topological Methods in Data Analysis and Visualization, pp. 317–331 (2017). DOI 10.1007/978-3-319-44684-4_19
42. Shivashankar, N., Pranav, P., Natarajan, V., van de Weygaert, R., Bos, E.P., Rieder, S.: Felix: A topology based framework for visual exploration of cosmic filaments. IEEE Transactions on Visualization and Computer Graphics **22**(6), 1745–1759 (2015)
43. Singh, S.: Weakly cemented granular materials: study at multiple length scales. Ph.D. thesis, Indian Institute of Science, Bengaluru, Department of Civil Engineering (2020)
44. Singh, S., Miers, J., Saldana, C., Murthy, T.: Quantification of fabric in cemented granular materials. Computers and Geotechnics **125**, 103644 (2020). DOI https://doi.org/10.1016/j.compgeo.2020.103644
45. Soille, P.: Morphological image analysis: principles and applications. Springer Science & Business Media (1999)

46. Sowers, G.B., Sowers, G.F.: Introductory soil mechanics and foundations. LWW (1951). DOI 10.1097/00010694-195111000-00014
47. Van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T.: scikit-image: image processing in python. PeerJ **2**, e453 (2014)
48. Wiebicke, M., Andò, E., Šmilauer, V., Herle, I., Viggiani, G.: A benchmark strategy for the experimental measurement of contact fabric. Granular Matter **21**(3), 1–13 (2019). DOI https://doi.org/10.1007/s10035-019-0902-x
49. Yang, Z.X., Li, X.S., Yang, J.: Quantifying and modelling fabric anisotropy of granular soils. Géotechnique **58**(4), 237–248 (2008). DOI 10.1680/geot.2008.58.4.237
50. Zheng, J., Hryciw, R.D.: Segmentation of contacting soil particles in images by modified watershed analysis. Computers and Geotechnics **73**, 142 – 152 (2016). DOI https://doi.org/10.1016/j.compgeo.2015.11.025