Explorative Analysis of Dynamic Force Networks in 2D Photoelastic Disks Ensembles

Farhan Rasheed[1], Abrar Naseer[3], Talha Bin Masood[1], Tejas G. Murthy[3], Vijay Natarajan[2], and Ingrid Hotz[1]

***Abstract*—This paper presents an interactive analysis framework for exploring data from photoelastic disk experiments, which serve as a model for two-dimensional granular materials. Granular materials, composed of discrete particles such as sand or gravel, exhibit behaviors resembling fluid or solid states depending on the system configuration. These behaviors arise from interparticle contact forces, which form complex force networks that govern the material's macroscopic behavior. Our framework is specifically designed to analyze such 2D ensembles of dynamic force networks, enabling the identification and characterization of their underlying structures. The framework is built around a topology-based, multiscale data segmentation in terms of force chains and cycles. The analysis methods are structured across three levels: (1) multiscale analysis of individual instances under specific loading conditions, (2) detailed exploration of single experiments encompassing a series of loading and unloading cycles, and (3) comparative analysis across experiments conducted under similar and differing setups. We demonstrate the capabilities of our framework with a case study for each of these levels.**

***Index Terms*—Force network, data abstraction, topological data analysis, explorative framework, multiscale analysis**

I. Introduction

GRANULAR materials, such as sand, grains, or powders, exhibit fascinating and complex behaviors that make them a significant focus of study across physics, engineering, and materials science. What makes these systems particularly intriguing is that their large-scale mechanical behavior can vary significantly depending on their microscale properties: the system setup, including particle arrangement, density, and anisotropy; its microstructure characteristics, also known as *fabric*—including particle shape, size, interparticle contacts, and forces; and the effective applied stress. Granular materials exist mainly in two states: an *unjammed*, fluid-like phase with free particle movement and a *jammed*, solid-like phase with a stable and deformation-resistant configuration. Multiple factors, including *volume fraction*, define the state. This jamming transition is critical in many industrial and natural processes, heavily influencing the durability and safety of constructions, yet predicting when and how it occurs remains a challenging question. This leads to the following main driving research questions when analyzing granular materials: What microscale properties significantly influence the macroscopic mechanical behavior of granular materials, and how do these properties interact to determine macroscopic outcomes? Can jamming behavior in granular materials be accurately predicted based on statistical properties observed at the microscopic scale? What mesoscale aggregation of microscopic properties is most effective for bridging the microscopic and macroscopic scales? This paper is the result of a collaboration with engineers exploring these questions from an experimental perspective, specifically using two-dimensional photoelastic disk experiments (Fig. 1), an essential step in deriving microscopic properties. These experiments use photoelastic disks under varying loading conditions, allowing researchers to visualize force distributions within granular assemblies. When subjected to stress, the photoelastic disks reveal force chains, network-like structures that represent the transmission of forces through the system. These force chains can be mathematically represented as 2D embedded weighted graphs, providing a framework to quantify and analyze the internal dynamics of granular materials. This setup offers an ideal platform for controlled studies of how microscopic arrangements influence large-scale behavior, bridging experimental observation and theoretical modeling. Further details of the experimental setup and methodology are provided in section II.

In this paper, we make the following key contributions: (i) VisPhoD, a novel framework for exploring ensembles of 2D photoelastic granular material experiments across three levels, single force network, single experiment, and a set of experiments; and (ii) a unified and extended formulation of commonly used statistical measures within a topological, multiscale analysis framework. At its core, VisPhoD employs two hierarchical topological segmentations: one that identifies force-connected components and another that captures force cycles enclosing non-participating particles, commonly known as *rattlers*. The proposed framework directly addresses the above-mentioned research questions by linking microscopic, mesoscopic, and macroscopic perspectives within a unified analysis workflow. Through its hierarchical design, VisPhoD enables systematic exploration of how local force chains and cycles (microscale features) aggregate into mesoscale structures and ultimately influence macroscopic phenomena such as jamming. We demonstrate VisPhoD's capabilities through two case studies: the first compares experiments with and without inlays, while the second examines an ensemble of repeated experiments conducted under similar conditions. These case studies illustrate how our approach supports a detailed analysis of the relationships between microscale structure, mesoscale evolution, and macroscopic mechanical response.

This work is a collaboration between two visualization

[1]Department of Science and Technology, Linköping University.
Email: {farhan.rasheed | talha.bin.masood | ingrid.hotz}@liu.se
[2]Department of Computer Science and Automation, Indian Institute of Science.
[3]Department of Civil Engineering, Indian Institute of Science.
Email: {vijayn | tejas | abrarnaseer}@iisc.ac.in

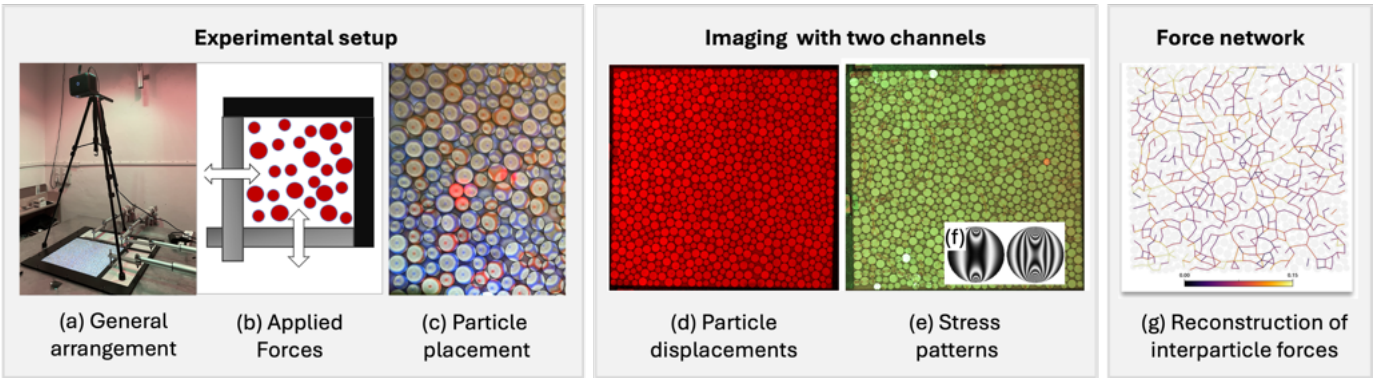

Fig. 1. Photoelastic experiment data generation pipeline. (a) Initial setup of photoelastic disks in a rack, with a camera positioned to capture system evolution. (b) Application of load by inward movement of the rack walls. (c) Projection of the disk layout onto the system to guide accurate disk placement. (d/e) Two channels of images captured by the camera, with (e/f) highlighting stress patterns. (f) Two detailed examples of fringe patterns. (g) Reconstructed force network, with inter-particle forces represented by color.

groups and a granular materials laboratory. The lab conducted the experiments, while the visualization teams developed the analysis framework. Regular discussions ensured seamless integration of experimental and visualization efforts.

II. Domain Background and Goals

Granular material systems are studied through simulations and experiments. Such studies provide direct access to microscopic statistics and interactions at the particle level [1]. This study focuses on analyzing data from experimental approaches, particularly those that use two-dimensional photoelastic disks, see Figure 1.

Photoelastic disk experiments are instrumental in the study of granular systems under controlled conditions. The experimental setup involves arranging photoelastic disks (disks coated with optical polymers) on an air table. This table levitates the particles on a thin layer of pressurized air, creating a near-frictionless base for the particle assembly. To avoid crystallization, disks of different sizes are used. To obtain controlled initial configurations, a predefined initial pattern is projected onto the table, guiding the manual placement of the disks (Figure 1c). During the execution of the experiment, the disks are initially loosely placed on the air table, representing an *unjammed state*. The system is then subjected to quasi-static compression as the rack walls move inward (loading phase), causing the stress levels to increase gradually. This is followed by a gradual relaxation of compression (unloading phase) (Figure 1b). The photoelastic disks respond to these varying stresses by altering their optical properties, which are captured using a high-resolution camera equipped with polarizing filters. In the experiments, a two-channel image is recorded: one channel (red) tracks the disk positions (Figure 1d), the other channel (green) captures stress patterns (Figure 1e). Using the birefringence property of the optical polymers that cover the disks, the contact forces can be calculated, forming a force network (Figure 1g). The experimental setup allows the repetition of experiments under controlled initial conditions and a consistent loading and unloading process. By continuously recording this process, each experiment generates a series of data sets containing information on particle locations, inter-particle connectivity, and inter-particle forces, represented as two distinct networks. More details on photoelastic disk experiments can be found in the paper by Daniels et al. [2]. By analyzing the variations in the system's evolution, researchers seek to identify the microscale properties that play the most significant role in predicting macroscopic outcomes. In our case studies, we explore results from various experimental protocols, which are further detailed in section VII.

Traditionally, similar experiments have focused their analysis on microscale properties, such as particle arrangements, particle connectivity, and force networks. In some applications, an intermediate scale, known as *mesoscale*, is also introduced to bridge the gap between micro- and macroscales. A prominent example, in the context of two-dimensional force networks, is the work by Cambou et al. [3] who proposed subdividing the domain into *meso-domains* defined by loops in the particle connectivity graph. Other approaches have focused on identifying loops in the force network where the values exceed the average force. Such loops enclose particles that do not directly contribute to force propagation within the network. By interpreting the force network as a planar graph embedded in a 2D spatial domain, this concept of loops can be reformulated using the topological notion of cycles in the graph [4]. The advantage of this formulation is that it places a stronger emphasis on the multiscale aspect of the network, extending beyond traditional analysis techniques and integrating these approaches into a unified topological multiscale framework.

III. Related Work

We review key analytical methods from granular materials and visualization, emphasizing those that influenced our research or share related methodologies. Specifically, we examine mesoscale definitions, aggregation, and characterization.

Granular material analysis. A significant body of research has focused on analyzing force and contact networks in granular materials. The role of force networks in granular materials is discussed in the work by Daniels et al. [5]. One frequently discussed concept in this field is the analysis of

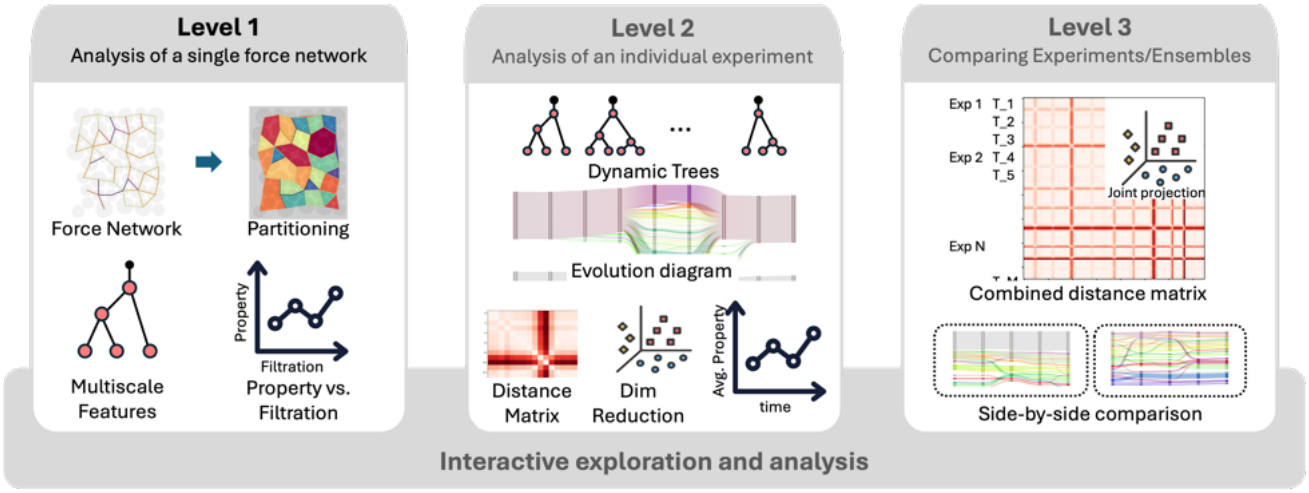

Fig. 2. A set of visualization tasks organized into three interconnected levels: Level 1 analyzes a single force network, Level 2 examines a single experiment, and Level 3 compares multiple experiments. All levels are linked within an interactive exploration framework.

contact loops and their evolution, as explored in works by Smart et al. [6] and [7]. These loops possess the unique property of enclosing particles, referred to as *rattlers*, which do not contribute to load transmission within the network. Building on this concept, Cambou et al. [3] introduced the idea of *meso-domains*, defined as such contact loops, which has been the main inspiration for our multiscale domain segmentation. In addition to studying the cyclic structure, Tordesillas et al. [8], [9] also investigated linear building blocks known as force chains. Unlike cycles, force chains are not as well-defined. They are typically considered as interconnected particles that experience contact forces exceeding a given threshold while satisfying certain straightness criteria. A large variety of properties have been introduced to assess the local organization of grains or particles, either as direct particle properties or as fabric properties aggregated in the meso-domain. These include scalar values such as anisotropy, average coordination number, and tensorial entities [10], [11]. Many of these properties are also available in our framework *VisPhoD* and are described in more detail in subsection V-C.

Topological approaches are well suited for analyzing and comparing data from a multiscale perspective [12], [13]. They are also frequently employed to study graphs and their evolutions, as demonstrated by Hajij et al. [14] or generally to analyze the geometric structure of materials science data [15]. Morse theory-based techniques have been used for segmenting particles and extracting force networks from 3d CT images of granular materials by Pandey et al. [16]. Kondic et al. [17] were pioneers in applying topological methods to the analysis of granular materials, focusing primarily on the zeroth and first Betti numbers, which correspond to force chains (β_0) and cycles (β_1) in the force network. In subsequent work, they proposed using persistence diagrams to derive quantitative measures to describe the evolution of interparticle forces in 2D dense particle systems. Similar to our approach, their multiscale analysis relied on a filtration [18], [19]. Rasheed et al. [4] introduced a method employing dual filtration to hierarchical segmentation of the domain into force-cycle-bounded regions, which is represented as a tree. A review of network-based and algebraic topology-based methods for studying granular materials can be found in the work of Papadopoulos et al. [20]. While all these studies introduce valuable graph analysis and comparison techniques, they do not focus on analyzing particle or fabric properties.

Multiscale feature comparison. Topological descriptors also have great potential for data comparison [13]. The Bottleneck and Wasserstein distances, which are related to optimal transport [21], are among the most established metrics for comparing topological descriptors [22], [23]. Merge trees are often compared using variants of edit distances, which rely on structure-preserving mappings that apply a sequence of simple edit operations [24]. In this paper, we use the deformation-based edit distance from Wetzels et al. [25], which operates directly on the merge tree as a geometric object, unlike branch decomposition methods.

Feature tracking. Observing the evolution of a system during the loading and unloading process is a fundamental task in our work. Since our focus is on topological structures, we will emphasize topological feature-tracking methods in this summary. Many existing tracking methods concentrate on spatial regions that are sparsely distributed in the domain, such as those defined by sub-level or super-level sets. The primary challenge in these approaches is to establish temporal correspondence between contour components, resulting in a temporal tracking graph. A common correspondence criterion is spatial overlap, which, for instance, Sohn. [26] and Widanagamaachchi et al. [27] employed to define tracks for regions represented by specific isovalues. Lukasczyk et al. [28] introduced nested tracking graphs, a method leveraging the hierarchical structure of sub-level and super-level sets, enabling tracking at multiple levels of detail simultaneously. An alternative approach to establishing correspondences uses critical point tracking, as demonstrated by Nilsson et al. [29], [30]. Other work has aimed to improve the layout and clarity of these nested tracking graphs [31], [32]. Saikia and Weinkauf [33] presented a region-tracking algorithm using

information from more than two time steps to establish fea-ture correspondence. In the context of combustion analysis, Schnorr et al. [34] developed a two-step tracking method that tracks space-filling features partitioning the domain using global optimization. In this work, we adopt the method pro-posed by Rasheed et al. [35], which generates nested tracking graphs for hierarchical, space-filling features. It is described in more detail in subsection V-B.

IV. Towards an Analysis Framework – Requirements and Tasks

The overarching motivation for this work was to gain a deeper understanding of the microscale processes involved in loading and unloading, and their influence on the overall evolution of the system. This requires examining both the similarities and differences in the development of the force network under varying initial conditions that represent the microscale. Although the precise approach to achieving this goal was initially open-ended, it was collectively agreed that the system should satisfy several high-level requirements **[R1–R4]** designed to link the system's behavior across multiple scales:

R1 Capture Multiscale Characteristics. Develop a concept for multiscale segmentation that allows for flexible mesoscale definitions and aggregation of fabric properties across different scales.

R2 Summarize Individual Experiments. Provide comprehensive summaries for each experiment to support preliminary data quality assessments and highlight trends.

R3 Monitor System Evolution. Design visualization techniques to track the evolution of aggregated properties during both the loading and unloading phases.

R4 Compare Experiments and Ensembles. Enable comparative analysis of individual experiments and ensembles to identify variations in patterns and overarching behaviors.

Based on these requirements, we derived a set of visualization tasks that collectively address the different aspects of multiscale analysis. To organize them coherently, we grouped the tasks into three interconnected levels, reflecting the increasing complexity of analysis, from examining a single force network, to studying the temporal evolution within an individual experiment, and finally comparing multiple experiments or ensembles. This hierarchical structure mirrors the physical hierarchy of the system, where microscale interactions lead to meso- and macroscale behaviors. It also provides a practical framework for progressive exploration, enabling researchers to move from detailed local inspection to broader comparative studies while maintaining consistent metrics and visualization principles. To support flexible and iterative exploration, we initially adopted a broad set of visualization tasks, which are expected to be refined and streamlined over time as the most effective methods are identified.

Level 1: Analysis of a single force network

L1-Embedding. Visualize the force network within the spatial context, highlighting load-transmission pathways.

L1-Property. Define, configure, or modify the definitions and parameters of network properties.

L1-Multiscale. Navigate across scales to observe how aggregated properties change, for detailed scale analysis.

Level 2: Analysis of an individual experiment

L2-Quality. Evaluate data quality through an experiment summary.

L2-Evolution. Analyze the evolution of aggregated properties during loading and unloading at different scales.

L2-Comparison. Quantitatively compare force networks and derived properties within a single experiment.

Level 3: Comparisons between experiments/ensembles

L3-Trends. Identify outliers and trends using distance matrices and dimensionality reduction techniques.

L3-Juxtaposition. Perform detailed, side-by-side comparisons of selected experiments or loading conditions.

L3-Distance. Quantitatively compare force networks across different experiments.

An overview of the different analysis levels is provided in Fig. 1. Implementing these tasks required the development of several technical concepts, which are introduced in the following section.

V. Technical Analysis Concepts Underlying VisPhoD

This section outlines the core technical concepts of the *VisPhoD* framework, including hierarchical domain segmentation, tracking segment evolution, mesoscale property aggregation, and a quantitative measure for comparing force networks across experiments.

A. Multiscale Domain Segmentation

We introduce two distinct segmentations of the domain based on different perspectives: *cycle-bounded regions* as a basis for property aggregation and *connected components* approximating force chains. Both segmentations can be formulated within a topological framework. Therefore, each force network, represented as a straight-line weighted planar graph $G = (V, E)$, is interpreted as a one-dimensional simplicial complex F. The vertices V correspond to 0-simplices (disk centers), while the weighted edges correspond to 1-simplices (inter-disk forces). Some vertices in V are isolated and correspond to non-participating disks, called *rattlers*. In topological terms, the zeroth (H_0) and first (H_1) homology groups of F represent connected components and cycles, respectively. These homology groups form the basis for our two multiscale segmentations, computed through respective filtrations with respect to weights assigned to all simplices. Vertex (0-simplex) weights are set to the largest force value in the system increased by a small positive number. Edges (1-simplices) inherit their weights from the corresponding force values.

Segmentation via Cycle-bounded Regions: Representative volume elements (RVE), which can be understood as the smallest volume that effectively captures macroscopic behavior [36], are a key concept in the analysis of granular materials. However, defining such RVEs remains an open problem. Cambou et al. [3] introduced the concept of *meso-domains*, characterized by *loops* in the particle connectivity graph that segments the domain. This approach is based on the observation that the particles enclosed by these loops, termed *rattlers*, do not contribute to the propagation of forces within the material. Building on this idea, we propose a multiscale decomposition of the domain into regions bounded by force loops and constrained by variable force thresholds. In the context of force networks, these loops correspond to cycles in a weighted graph restricted to edges that exceed the respective threshold. This problem, well-known in topol-ogy, can be effectively addressed using persistent homology by applying filtration over edge weights [37]. Rather than computing cycles directly in the complex F, we adopt a dual filtration approach, utilizing Alexander's duality theorem [38]. A key advantage of this approach is that it provides a compact representation of homology generators that remain consistent across filtration levels, enabling hierarchical domain segmentation. Additionally, dual filtration naturally yields a tree representation of the segmentation, offering a structured view of the multiscale topology. More details on the method can be found in [4].

Segmentation via Connected Components: A complementary view to the force cycles is given by connected components, related to force chains, and used to quantify the force network [39]. Each component represents a connected set of particles that interact by a force. The multiscale description of these particles is obtained by varying the force threshold, generating a filtration starting with the maximum force threshold. As for the cycle-bounded regions, this segmentation can be represented as a tree.

B. Multiscale Tracking of Cycle Bounded Regions

Each photoelastic disk experiment produces a sequence of force networks, capturing the evolution of forces as loading and unloading are gradually applied. To analyze structural changes and identify consistent trends within the force network, we track the segmentations defined by cycle-bounded regions (subsection V-A) across multiple pre-selected hierarchy levels, resulting in a multiscale nested tracking graph. Typically, these hierarchy levels include the base level and the median force level at each time-step. To establish correspondences between segments at a given hierarchy level in consecutive time steps, we employ a spatial overlap-based method proposed in [35]. To efficiently compute overlaps, we exploit the hierarchical structure of the segmentation represented in the trees. In practice, this means starting with overlap computations at the leaf nodes, which correspond to the base level, and propagating the values upward to the selected hierarchy level. Formally, let T_t and T_{t+1} denote trees in consecutive time steps with leaf nodes N for T_t and M for T_{t+1} respectively. Each leaf represents one segment at the lowest filtration level. To quantify the similarity of the nodes $n \in N$ and $m \in M$, we employ the overlap coefficient:

$$O(n, m) = \frac{|n \cap m|}{\min(|n|, |m|)}$$

where $|.|$ denotes the spatial area covered by the corresponding segment. A correspondence between nodes is established when their overlap coefficient exceeds a user-adjustable threshold τ. In all our experiments we used $\tau = 0.3$ as a good compromise between discarding spurious connections while keeping significant correspondences. Computing correspondences between all leaf nodes across all consecutive time steps builds the basis for the correspondences at higher levels. For any arbitrary hierarchy level, the overlap coefficient between consecutive time steps is computed recursively by following the parent nodes up from the leaves of the tree and aggregating the node overlap. By computing the correspondences across all consecutive time points for each selected hierarchy level, we construct a tracking graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ represents the segments at all hierarchy levels and $\mathcal{E}$ represents the correspondences between segments across time.

C. Fabric Property Aggregation

One of the primary motivations for segmenting through cycle-bounded regions is to aggregate properties relevant to force network analysis (subsection V-A). These properties are accumulated on different scales to observe their variation effectively. The aggregation process across scales can be formalized as:

$$\mathcal{A}_{i \in S_f}(P(i))$$

where S_f are segments at the filtration level f, $\mathcal{A}$ is an aggregation operator (e.g., Sum, $Mean$, Max, Min) according to the requirement. and $P(i)$ is a property of a segment i in the domain.

VisPhoD enables the aggregation of properties commonly found in the literature while also allowing for their modification and the definition of new properties. A selection of scalar and tensorial properties is described below: *Scalar properties* include disk-based attributes such as the disk coordination number, contact-based properties such as contact forces, as well as segment-based attributes such as the number of disks enclosed by cycles (i.e. rattlers) or the perimeter of the segment. *Tensorial properties* encompass various fabric tensors and derived scalar attributes, such as anisotropy (absolute

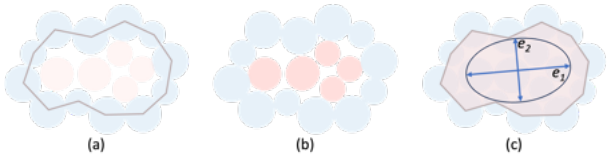

Fig. 3. Property aggregation over cycles. (a) iteration over the boundary, (b) iteration over the interior, (c) fabric tensor represented as an ellipse highlighting its eigenvectors.

difference in eigenvalues) and elongation (orientation of the major eigenvector). Fabric tensors $\mathfrak{F}$ are defined as:

$$\mathfrak{F} = \frac{1}{N}\sum_{i=1}^{N} \omega_i \vec{n}_i \otimes \vec{n}_i$$

where $\vec{n}_i$ is the normalized direction of contact i. The sum either iterates over all contacts in a cycle or only the cycle boundary, Figure 3. For connected components, it iterates over all contacts within components. N represents the total number of contacts considered. ω_i is a weight assigned to the contact i that typically represents the contact force. The symbol $\otimes$ denotes the outer product. Other *global properties* of interest include the total number of cycles in the domain or the number of n-cycles, where n denotes the cycle length.

D. Quantitative Comparative Measures of Segmentations

Tracking segmentation at specific hierarchy levels (subsection V-B) is well-suited for visualizing the evolution of the force network (subsection VI-B). However, this approach does not fully leverage the complete hierarchical segmentation represented by segmentation trees. To address this limitation, a complementary distance measure between segmentation trees quantifies similarities and differences between force networks. It enables comparison and clustering across experiments within an ensemble. Among the various tree comparison measures available, we adopt the edit-distance-based Path Mapping Distance Metric [25], a state-of-the-art method known for its stability and computational efficiency. This metric is derived from the edit distance, which quantifies the minimum cost required to transform one tree into another through operations such as insertion, deletion, and relabeling.

VI. Design of VisPhoD

VisPhoD is a visual analysis system built upon the hierarchical segmentations and aggregated properties discussed in the previous section. Its design is guided by the requirements and tasks outlined in section IV. The system is structured into three interconnected levels, enabling seamless exploratory analysis. These levels are described in detail in the following sections, with an overview provided in Figure 2. The optimal entry point for exploration depends on the specific scientific question and the available data.

A. Level 1 Visual Components – Single-instance Analysis

The first level includes visual components for an in-depth analysis of an individual force network.

Direct Force Network Rendering: A visualization option enabling direct inspection of the disk arrangement and force network. In this view, disks are shown in light gray, while contact forces are depicted as edges connecting the disks, with colors indicating their force values (Figure 4a). This representation facilitates the task *L1-Embedding*.

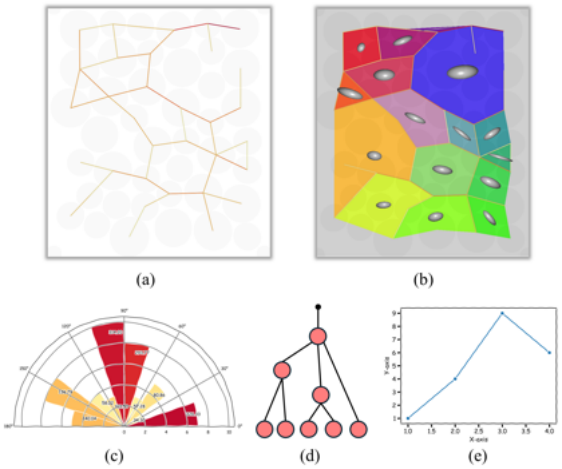

Fig. 4. VisPhoD Level 2 - Single Experiment Instance. (a) Force network visualization. (b) Partitioning of the domain, overlaid with glyphs representing fabric tensor. (c) Polar plots illustrating cycle or edge orientations, colors encode statistical properties of the cycles. (d) Hierarchical tree representation of cycles across filtration levels. (e) Line plot displaying key properties across filtration levels, including the number of cycles, number of n-cycles, average anisotropy, and other relevant metrics.

Segmentation Rendering: This view shows the domain segmentation at a selected hierarchy level, integrating it with the disks and force network to provide contextual insights. Each segmented region is colored with a uniquely assigned random color or a 2D colormap (Figure 9), where the color corresponds to the segment's centroid position (Figure 4b). This representation supports the task *L1-Embedding*. Users can explore segment-specific properties by hovering over a region. Additionally, segment orientations and shapes, derived from their fabric tensors, are visualized using ellipsoid glyphs (Figure 4b). Interactive exploration allows users to dynamically adjust the filtration value and observe how the segmentation evolves, thus supporting the task *L1-Multiscale*.

Segmentation Tree: A tree visualization provides an overview of the entire segmentation hierarchy (Figure 4d). The leaf nodes represent cycles at zero filtration level, whereas the inner nodes indicate the filtration values at which two cycles merge. This visualization summarizes the evolution of cycles across different filtration levels and helps guide the selection of relevant hierarchy levels. A red horizontal line, positioned at the current filtration level, serves as a reference and can be dragged up or down to explore the segmentation at different scales. This representation provides direct support for the task *L1-Multiscale*.

Polar Plots: A directional histogram plot, commonly used in granular material analysis, visualizes the directional statistics of segments by grouping occurrences into directional bins (Figure 4c). Users can choose between two plots: one directly summarizing the contact directions and another based on fabric directions derived from the eigenvectors of the corresponding fabric tensor. Optional weightings can be applied to emphasize specific directions. One option highlights directions associated with strong force values (resp. eigenvalues). For eigenvector-based plots, anisotropy is an alternative weighting measure, as the eigenvector directions of isotropic tensors

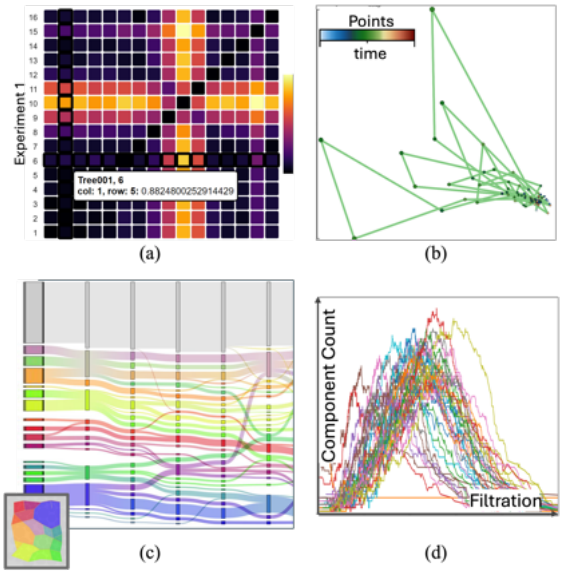

Fig. 5. VisPhoD Level 2: (a) Heatmap of the distance matrix, brighter colors indicate greater distances. (b) Multi-Dimensional Scaling (MDS) projection of the distance matrix, with lines connecting adjacent time points. (c) Evolution diagram illustrating cycle development over time. Columns represent 1d projection of the segmentations (for initial time step also showing the 2D segmentation). Links indicate spatial overlap between segments in consecutive time steps, colors encode the spatial location of each segment's center. (d) Line plot comparing aggregated parameters over time.

are less informative and stable. The coloring of the bins can optionally be used to display other measures averaged over the respective bin. This representation supports the task *L1-Property*.

Statistics/Property Graphs: Aggregated properties can be analyzed as a function of filtration levels, providing insights into their evolution across scales (Figure 4e). The plot supports various aggregation operations, including sum, mean, and median, with adjustable weights and modes available via drop-down menus. At the base level, properties capture local characteristics. As segments grow, these properties transition to regional attributes and, at the highest level, they represent averages over the entire domain. This visual representation supports tasks *L1-Property* and *L1-Multiscale*.

B. Level 2 Visual Components – Experiment Analysis

The second level summarizes and presents statistical insights for a single experiment.

Evolution Diagram: The evolution diagram provides a comprehensive summary of a single experiment by visualizing the nested tracking graph (subsection V-B) using a Sankey diagram (Figure 5c). It effectively captures the evolution of cycles under varying loading conditions, emphasizing merge and split events over time. In the evolution diagram, the nodes represent cycles, while the links indicate their relationships over time, mapped to the horizontal axis. The size of the nodes corresponds to the size of the cycle, and the thickness of the link reflects their degree of overlap. The node positions along the vertical axis are arranged using a layout that minimizes edge crossings. Each vertical axis represents an instance of the experiment as a bar, which effectively serves as a one-dimensional projection of the segmentation. Segment sizes are accurately displayed, color indicating the location of the centroid using a two-dimensional color map (Figure 9a). Alternatively, colors can be used to highlight hierarchical information. Users can interactively select the desired filtration level to explore the corresponding flow diagram. Since the representation provides an overview of the entire experiment, it is well suited for evaluating the quality of the data. This representation supports the task *L2-Evolution* and *L2-Quality*.

Distance Matrix and Projection Plots: Complementary to the evolution graph, we provide a matrix view that visualizes the distance between individual instances based on the segmentation tree, see Sec V-D. Although these measures reduce the complexity of force networks to a numerical representation, they effectively highlight time steps where significant changes occur in the network. The distance matrix is presented as an interactive heat map (Figure 5a), allowing users to explore detailed visualizations of selected pair of instances. Additionally, an alternative perspective on the experiment can be obtained through dimensionality reduction techniques. VisPhoD supports both MDS and T-SNE projections to facilitate this analysis Figure 5b. This representation supports tasks *L2-Evolution* and *L2-Comparison*.

Property Comparison and Evolution Graph: All aggregated property graphs from level 1 are available for analyzing an entire experiment. The property graphs for all instances can be plotted within a single image (Figure 5d) or alternatively, against time or instance ID to observe their evolution. For aggregated properties, this requires selecting a specific scale or hierarchical level. This representation supports tasks *L2-Evolution* and *L2-comparison*.

C. Level 3 Visual Components – Ensemble Analysis

The third level visualizes multiple experiments, highlighting similarities, differences, and patterns to guide selection for a deeper analysis.

Distance Matrix and Projection Plots: In a matrix view, heat maps visualize all instances in multiple experiments (Figure 6, middle). Filtering supports focusing on instances with, e.g., high packing fraction. Hovering over the distance matrix displays a tooltip with details about the time point and distance value. The distance matrix is interconnected with dimensionality reduction plots (Figure 6, top), using MDS and t-SNE. In these plots, each experiment is depicted as a line colored by experiment ID, with nodes encoding time or packing fraction through color and size by default. Users can customize these mappings to represent other properties, such as the number of n-cycles or total anisotropy. These visualizations support the tasks *L3-Distance* and *L3-Trends*.

Side-by-side Comparison: Users can select specific experiments or instances for side-by-side comparison (Figure 6 bottom), supporting the task *L3-Juxtaposition*. Furthermore, the *L2-Comparison* task is inherently supported within these views, allowing users to examine the tree of a single experiment and explore its heat map and projection plots.

Table View and Filtering: VisPhoD includes a table view (Figure 6, middle) that provides access to the metadata of ensemble members. This includes labels indicating whether

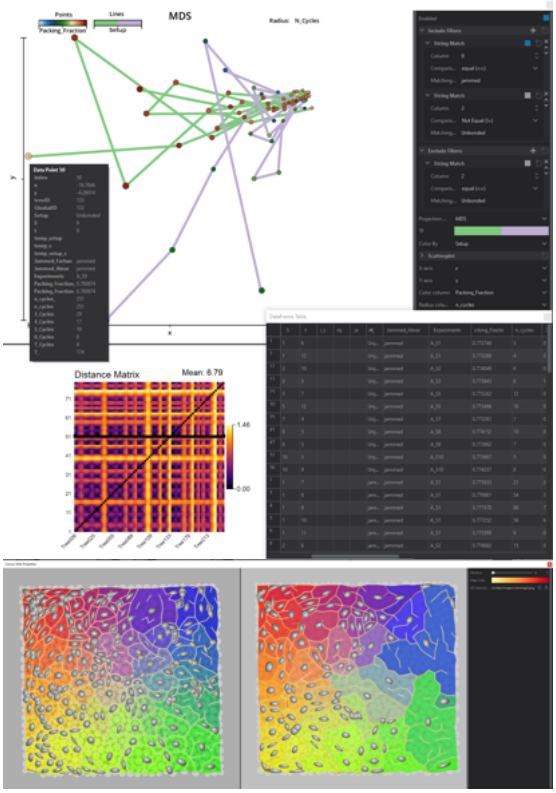

Fig. 6. VisPhoD Level 3: Top row: Joint distance matrix projection for two datasets using MDS. Lines, colored by dataset, connect adjacent time points, colored according to their packing fraction. Hovering over points reveals details of the corresponding time step. The interaction panel on the right allows for adjustments to the representation. Middle row: Joint heat map representing the distance matrix. The interaction panel supports filtering. Bottom row: Side-by-side comparison of the segmentations for two selected time points.

a system is jammed or unjammed, its packing fraction, the corresponding setup and experiment, and its time point. Users can apply logical filters to refine the dataset, such as excluding unjammed time points or selecting specific packing fraction ranges. The table view is seamlessly integrated with the heat map and projection plots, ensuring that only the filtered instances are displayed. Additionally, hovering over a point in the projection plots reveals its metadata for quick reference.

D. Across Level Interaction

VisPhoD supports interaction across levels if enabled. Selecting points (max 2) in the projection plots in Level 3 displays the corresponding visualization in the Level 1 view side-by-side. Similarly, the possibility to select experiments in Level 3 shows the evolution diagrams in Level 2 view side by side. Selecting a column in the evolution diagram opens the respective visualization in Level 1.

E. A Note on Implementation

VisPhoD is implemented using Inviwo, a prototyping framework for scientific visualization [40]. Inviwo is primarily written in C++, but also supports wrappers for Python and JavaScript. The segmentation tree layout, evolution diagram, and heatmap are computed using D3 and JavaScript. Different views can be displayed in separate rendering windows with

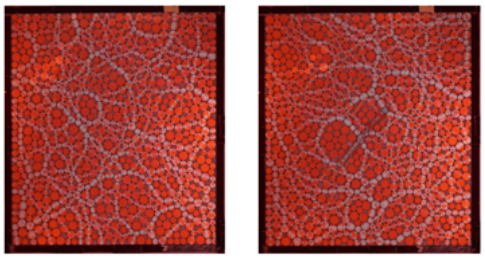

(a) No Inclusion (A2) (b) With Inclusion (A1)

Fig. 7. Dataset A: An instance of dataset A (a) without inclusion and (b) with inclusion having disks with relatively dark color.

dedicated property widgets or combined based on hierarchical levels.

The proposed system is designed as an interactive framework that allows users to explore data and adjust parameters in real time. Most computations are performed on demand, providing immediate visual feedback and supporting flexible analysis. Only two components benefit from precomputation: the evolution diagram and the pairwise path mapping distance used to compare merge trees across time points and experiments. For the evolution diagram, the most computationally intensive step is constructing the overlap matrix $C(t_i, t_{i+1})$. For a dataset consisting of 39 steps computing each overlap matrix of size 1650×1650 takes on average 4.39 seconds, resulting in a total of approximately four minutes to generate the full evolution diagram.

The path mapping distance is also potentially expensive, given that the algorithm has a runtime complexity that is superquadratic in the size of the trees. However, in practice the targeted two-dimensional disc datasets yield moderately sized merge trees. Consequently, even for the complete set of experiments, total computation time typically remains below twenty minutes. For example, computing all pairwise distances among three large trees (with 673, 737, and 539 nodes, respectively) requires 87 seconds.

VII. Case Study

We demonstrate VisPhoD through two use cases utilizing photoelastic disk datasets with different experimental protocols. The first use case compares two experiments, one with an inlay and one without, with similar but not identical load-unload cycles. Although the force values in this study may not be entirely reliable, notable trends can still be observed. The second use case explores two sets of experiments: one with uniform and one with varying initial fabric.

A. Effect of Inclusions in Granular Materials

A common approach to enhancing the structural properties of granular materials involves embedding inclusions or inlays within them. Inclusions have material properties distinct from the surrounding granular medium, and provide a cost-effective strategy for reinforcing the material. This case study investigates the impact of such inclusions in photoelastic disk experiments.

The dataset comprises results from two experimental configurations: *A1*, in which an inclusion is embedded within

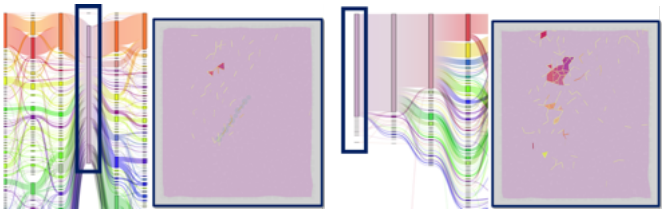

(a) (b)

Fig. 8. A closeup of two selected time-steps from the evolution diagram and the corresponding rendering of the cycle-based segmentation view of (a) *A1* at time point 12 and (b) *A2* at time point 1.

the granular system, and *A2*, where no inclusion is present (Figure 7). The inclusion is formed by gluing multiple photoelastic disks into a chain, while the remaining particles are free to move within the granular material. Both configurations contain 820 disks subjected to load-unload cycles. Biaxial stress is applied quasi-statically by moving the lower and right boundaries inward or outward. The goal is to observe variations in the distribution and propagation of stresses caused by the inclusion.

We begin with a Level 2 analysis (*L2-Evolution*) to inspect the evolution diagrams for A1 and A2 across multiple filtration thresholds. These diagrams provide a comprehensive overview of both datasets. A cropped version at a median filtration level, using the default colormap, is shown in Figure 9(left). As a first step, we assess the overall correctness of the data. In doing so, we identify a few time steps where only a few or no cycles are present, which piques our interest. To verify the correctness of these time steps (*L2-Quality*), we examine two specific cases in their contextual view using Level 1 tasks (*L1-Embedding*). The rendering view in Figure 8 displays the corresponding force networks and the respective time steps in the evolution diagram, revealing that no force paths connect the two walls and no major cycles exist, which is consistent with the evolution diagram. This observation suggests a potential artifact introduced during the data-generation pipeline.

Next, we analyze the evolution, merging, and splitting of cycle-bounded regions throughout the experiments. A striking observation in experiment *A1* (with inclusion) is the presence of a large and persistent cycle that remains throughout the entire experiment. This cycle appears in the orange-red region and is visible as an orange band in the diagram (Figure 9, bottom,left). In experiment *A2* (without inclusion), cycles in the orange-red region also tend to be larger than in other areas, though they are less pronounced and less stable compared to *A1*. Since large cycles correspond to regions filled with non-participating, loose particles (rattlers), this suggests that this area is less affected by the external load applied to the opposing walls (right and bottom). Its stability may result from being shielded by the inclusion, as the region lies approximately behind it. To validate this observation, we designed a customized colormap that divides the domain into six distinct regions, each assigned a specific color, with the region behind the inclusion highlighted in purple (Figure 9, right). The resulting evolution diagram confirms that the purple nodes, representing the region behind the inclusion, remain more consistent in *A1* compared to *A2*. This effect becomes even more pronounced at a zero-threshold filtration level (*Supplementary Figure 1*).

While evolution diagrams offer a quick overview of regional development, they lack quantitative detail. Statistical plots help bridge this gap by providing complementary insights (L1-Property, L2-Comparison). To quantify differences, we examine the number of rattlers in the regional divisions from the colormap Figure 9(b), revealing significant variations between experiments. Notably, in *A1*, the number of rattlers in the purple region (most shielded by the inlay) is significantly lower in *A2* compared to *A1*, with a p-value of 0.005 from the *Mann-Whitney U test* (Figure 10(a)). Conversely, in other regions, *A2* has fewer rattlers than *A1*, with the most pronounced difference in the blue regions. These findings align with trends observed in the evolution diagram.

Beyond the evolution of segmentation over time, we analyze characteristic properties across filtration levels, such as the total cycle count and the n-cycle count, to provide additional perspectives on the data. Specifically, we focus on graphs that display the total number of cycles and the number of 3-cycles (Figure 10(b,c)), where each line represents a time point within the respective experiment. Both cycle count plots reveal a higher degree of clustering in the experiment *A1* compared to *A2*. Furthermore, in *A1*, the number of cycles decreases more rapidly with increasing filtration levels. Furthermore, the maximum absolute number of cycles and 3-cycles is higher in *A2* than in *A1*. A higher number of 3-cycles that do not contain any particles corresponds to fewer rattlers, which is consistent with the findings in Figure 10(a).

In addition, we used the distance matrix and projection using t-SNE to provide a further overview of the data. As expected, there are no clear patterns visible in both plots Figure 11 as the tree distance measure only considers the hierarchical structure of the segmentation and is ignorant to the geometric embedding. This is also demonstrated in the two segmentation plots that belong to the different experiments Figure 11 botom, which are very close in the t-SNE projection (highlighted by the red arrow). The two segmentations are indeed similar, except for the area close to the inclusion.

Summary: In this case study, the domain scientists particularly appreciated the evolution diagram's concise overview of the entire experiment, as it provides a novel perspective on the data, including the entire loading and unloading sequence. The ability to query the diagram for more detailed information was also well received. However, while the observations from this case study were considered interesting, the domain scientists noted that the results from these experiments should be taken with care, as they were part of a preliminary study with force-value extraction that might not have been entirely reliable.

B. Investigating Effect of Initial Fabric on Jamming

The key question in this case study is whether and, if so, how the fabric (defined by the initial spatial arrangement of particles, their geometric properties, and their local coordination with neighbors) influences the system's overall behavior.

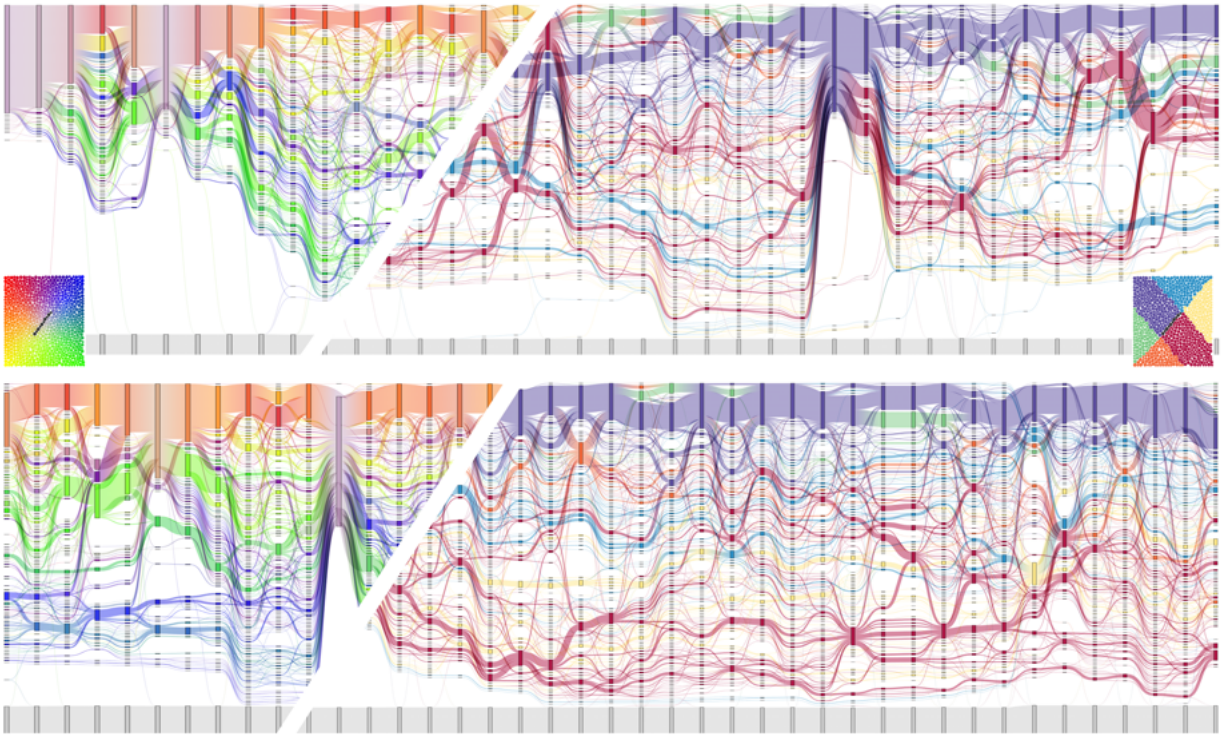

Fig. 9. Evolution diagrams at the median level for Dataset *A1* (bottom) and *A2* (top). The left side of each diagram uses the default color map to provide spatial context, where segment colors are assigned based on their centroids. The right side features a dataset-specific color map that emphasizes the area behind the link (highlighted in black).

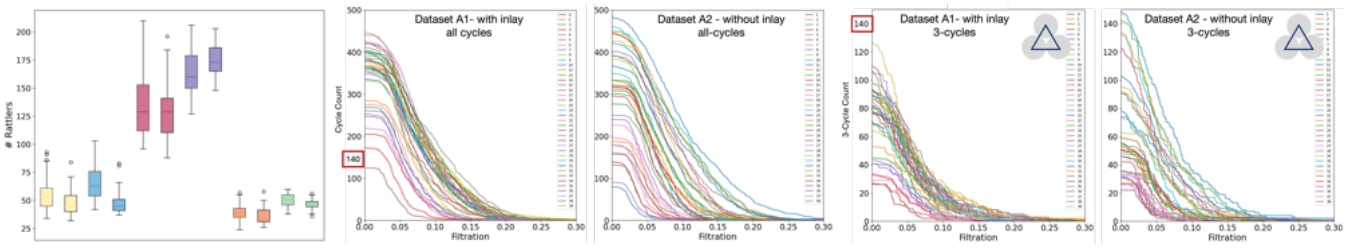

(a) (b) (c)

Fig. 10. Statistical plots give the quantitative view of the experiments. (a) Box plot showing the statistical summary of the rattlers at a median level over the entire experiment. The distribution of rattlers in the purple region in A2 is significantly less than the corresponding region in A1 with p-value 0.005 (they are different 0.01) by *Mann-Whitney U test*. The color corresponds to the regions in Figure 9(b) the boxes with dashed outlines belong to experiment *A2* with solid outlines to *A1*. The line plots show the development of the number of 3-cycles across filtration for (b) *A1* and (c) *A2*. Each line shows a time point in the experiment.

Two sets of carefully controlled photoelastic disk experiments have been conducted to investigate this systematically. In the first set, dataset *B*, the experiment was repeated 10 times with identical initial fabric and similar load-unload cycles. In the second set, dataset *C*, the experiment was repeated with a randomized initial fabric while maintaining consistent loading and unloading conditions. Consistent behavior is expected for the experiments in dataset *B*, while a more variable and less predictable pattern is anticipated in dataset *C*. Since these experiments require high precision and effort, some trials were completed before reaching the unloading phase. Figure 12 presents the dataset generated from the experiments, where time points are classified as jammed (red) or unjammed (blue) based on a packing fraction threshold. Time-steps where no data is available are marked in white. This threshold is determined by computing the average contact force per particle, with the jamming packing fraction identified at the point where this value exhibits a steep increase. This case study focuses mainly on visualization tasks to compare trends within ensembles.

Initial Overview Comparison: To investigate structural similarities and differences, we begin our analysis with t-SNE and MDS dimensionality reduction plots. These methods generate joint projections of all instances across experiments in both datasets, as shown in Figure 13. The projections are derived from the pairwise distance between the cycle-bounded segmentation trees represented in a distance matrix. Each point in these projections represents a force network at a specific time step within an experiment. The point radius corresponds to the number of cycles at filtration level 0. Consecutive time

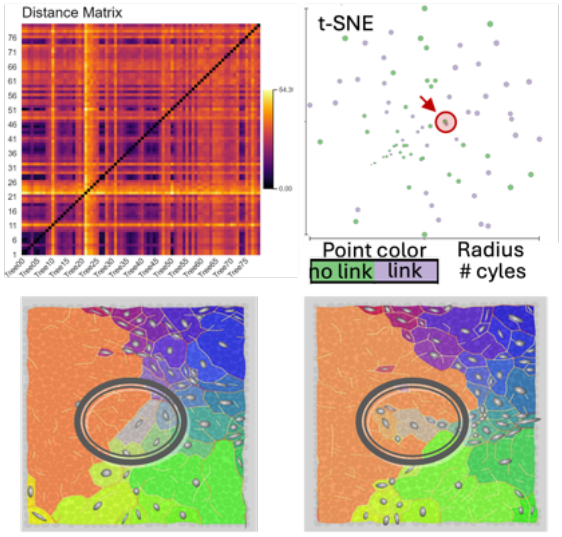

Fig. 11. Top: Distance matrix and t-SNE projection for both experiments *A1* and *A2*. Side-by-side visualization of segmentation for two instances selected in the t-SNE plot.

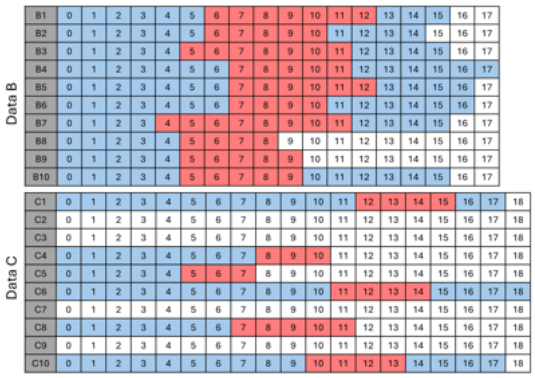

Fig. 12. Overview of the data from experiments *B* & *C*: Each row represents an individual experiment. Instances are labeled as jammed (blue) or unjammed (red), and white indicates missing data.

steps within an experiment are connected by lines, color-coded by dataset (green for *C*, purple for *B*). To differentiate packing fractions, we use a categorical color map with three levels: low, medium, and high.

A direct observation from these plots is that dataset *B* generally exhibits higher packing fractions than dataset *C*, as indicated by the color of their respective points. The two projections (MES and t-SNE) offer distinct perspectives on the data. It is evident that the MDS plot more accurately represents the distances between different states. However, the t-SNE plot provides greater detail on states with low packing fractions. In the t-SNE plot (Figure 13), time points with higher packing fractions and more cycles cluster on the left, forming two distinct groups, one for each dataset, highlighted in the oval. In contrast, these same time points appear to be more dispersed in the MDS plot. The opposite trend is observed for time points with low packing fractions and fewer cycles. In the MDS plot, they form a tight cluster, whereas in the t-SNE plot, they are more widely spread. Notably, a distinct group of points in the bottom right, highlighted in the gray box, caught our attention.

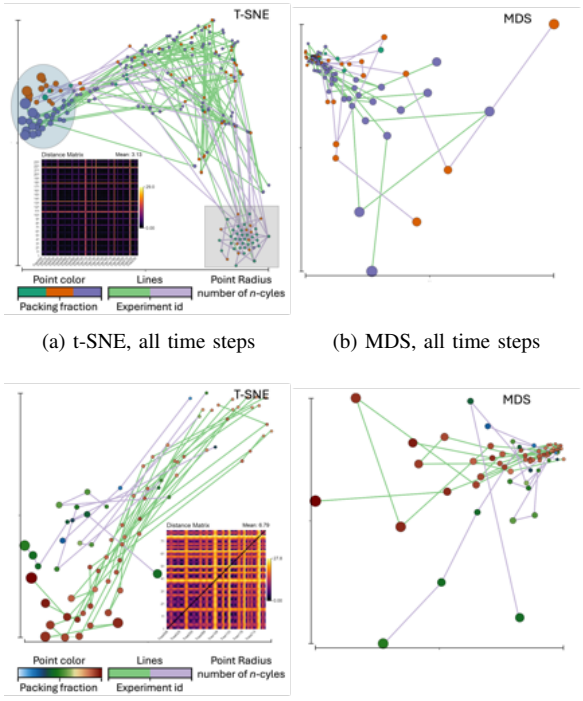

(a) t-SNE, all time steps (b) MDS, all time steps

(c) t-SNE, jammed time steps (d) MDS, jammed time steps

Fig. 13. Combined dimensionality reduction plots for data sets B&C. Dataset B is colored by green line while dataset C is color by purple.

This group consists mainly of early time points in dataset *C*, which we hypothesize correspond to configurations with negligible force networks and may be an artifact of t-SNE.

To test this assumption, we selected one experiment from each dataset and analyzed their force networks in detail, as shown in *Supplementary Figure 2*. The evolution diagram reveals a clear difference in behavior at the start of the experiments. In dataset *B*, small contact cycles (low filtration level) appear from the beginning, represented by the short segments highlighted in the figure's inset. In contrast, dataset *C* initially lacks any contact cycles. These early time steps in dataset *C* are mapped by t-SNE to the highlighted group in Figure 13. A similar trend is observed in the connected force components ("force chains" at the median force value filtration level), shown in the bottom row of *Supplementary Figure 2*. Each connected component is assigned a random color, while gray discs represent isolated discs at this filtration level. In dataset *B*, connected components are present from the beginning, whereas in dataset *C*, they are initially absent.

Another notable feature, visible primarily in the MDS plot (Figure 13), is that most experiments form a cycle in the projection, starting and ending in the top-left cluster, which represents states with low packing fractions. However, experiments that were interrupted and lack unloading data do not exhibit these cycles. The states at which these experiments were disrupted appear to be the most distinct, suggesting that their unique characteristics may have contributed to the disruption.

Comparison Overview of Jammed States: In the second phase of our analysis, we focus on time points identified as

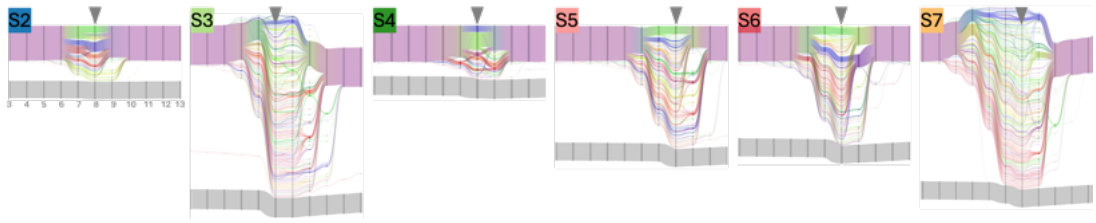

(a) Evolution diagrams, time step 3 to 12 - low filtration level. Triangles highlight the time points selected for the segmentation.

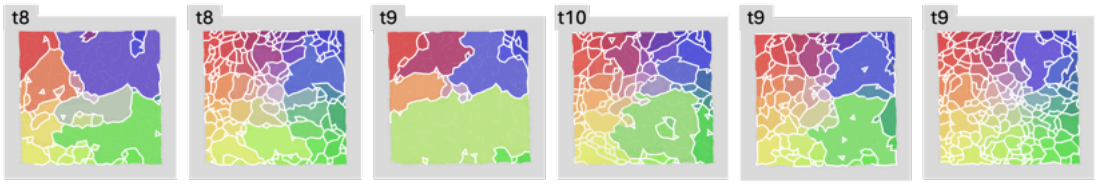

(b) Segmentation for the time steps selected from the MDS plot

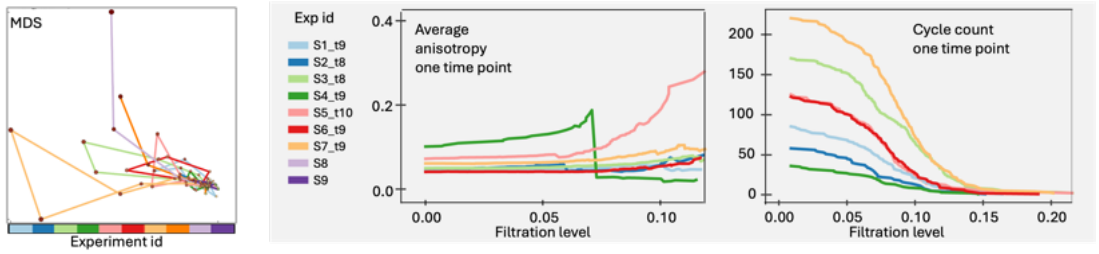

(c) Left: MDS projection, where each line represents a separate experiment. The time point farthest from the cluster in each experiment has been selected for further analysis. Right: Cycle count and anisotropy as functions of filtration level for the selected time steps.

Fig. 14. Comparison of experiments in dataset B. Despite a uniform initial fabric across all experiments, the evolution of the load-unload cycle varies significantly, though some common patterns emerge.

jammed, as these phases are of particular interest to domain scientists. The corresponding dimensionality reduction plots (Figure 13, bottom row) present a condensed version of the full data projection. In the MDS plot, the closed-cycle structure remains visible, suggesting a symmetry between the jamming and unjamming phases. The t-SNE plot provides a clearer distinction between the two experimental setups, especially in configurations with the highest packing fraction.

Since assessing variance differences between datasets is challenging in projection plots, we analyze their distance matrices directly. To quantify these differences, we compute the arithmetic mean of all pairwise distances within each experiment's matrix by summing all values and dividing by the total number of elements. The results indicate that dataset *B* has a lower mean distance than dataset *C*. However, this finding should be interpreted with caution, as several parameters differ between datasets, limiting their direct comparability. Further investigation is needed to draw definitive conclusions.

Side-by-side Comparison of Experiments in Dataset B: In the final phase of the case study, we focus on experiments with a uniform initial fabric (dataset *B*). As before, we begin with an overview of all experiments using the evolution diagrams (*Supplementary Figure 2*, top row). The significant differences between the experiments were a major surprise to domain scientists. Although they were aware of variations in the development of jamming (Figure 12), the extent of these differences was unexpected.

The evolution diagrams clearly highlight the progression of jamming; refer to Figure 14. Jammed areas are more strongly interconnected and efficiently transmit forces. In terms of cycle size, this corresponds to many small cycles (segments containing few rattlers). Unlike a classification based solely on packing fraction, the evolution diagrams reveal that jamming does not evolve simultaneously across the entire domain but progresses gradually, an observation consistent across all experiments. For example, in experiment S6, the first region to fragment into many small cycles is the red area, followed by the orange-yellow region in the next time step. In contrast, large parts of the blue and green areas never fully develop jamming. This pattern is largely consistent across all experiments, though they differ in how completely jamming progresses. Among them, experiments S3 and S7 achieve the most complete jamming.

The *Supplementary Figure 2* presents detailed segmentations at the lowest filtration level for the time steps with the highest cycle count. These time steps correspond to the points in the MDS projection that are farthest from the main cluster. For these points, we also plot the evolution of cycle count across filtration levels. Interestingly, experiments S5 and S6 show nearly identical progressions, also having similar evolution diagrams. However, their average anisotropy differs significantly at higher filtration levels. A closer examination of their force networks reveals that S5's high anisotropy is driven by a single persistent cycle with extremely high force

values that can be considered as an outlier.

Summary: For datasets *B* and *C*, VisPhoD provided, for the first time, a comprehensive overview of all experiments simultaneously, with the evolution diagram enabling detailed examination of the jamming process. The unexpected variation in dataset *B* highlights the strong stochastic nature of granular material evolution and underscores the value of visualiza-tions for understanding complex data. Comparisons between datasets reveal differences in projection plots and distance variance, suggesting potential underlying distinctions, though the limited number of experiments, especially in dataset *C*, precludes definitive conclusions.

C. Summary of User Feedback

Domain scientists emphasized that the interactive workflow and multi-level structure of VisPhoD greatly improved data exploration and validation. In the inclusion study, they particularly valued the evolution diagram for providing an intuitive overview and for supporting the identification of artifacts and persistent structures. In the jamming study, users highlighted the ability to compare multiple experiments simultaneously, clearly showing the strong stochastic variability of the system even under identical initial conditions. Overall, users reported that the system enhanced the interpretability and verification of the data in collaborative analysis. As a future goal, it was mentioned that such a system, with minor adaptations, could also be applied to simulation data for comparing experimental and computational results.

VIII. Conclusion

In this work, we introduced VisPhoD, a framework for exploring data from photoelastic disk experiments, developed in close collaboration with domain scientists. A key challenge in designing the framework was the absence of clearly defined tasks and open-ended research questions. Rather than enforcing rigid analytical pathways, our goal was to create an interactive tool that encourages reasoning and discovery through direct engagement with the data. The resulting topological framework unifies common analysis tasks within an interactive environment for studying force networks in granular materials. VisPhoD effectively bridges the scales central to granular physics by linking local force chains and cycles at the microscale with their mesoscale aggregation and macroscopic manifestation in jamming behavior. This multi-level perspective enables intuitive comparison of experimental conditions and reveals how subtle variations in microstructure can lead to significant differences in macroscopic outcomes.

Beyond the specific application presented, the hierarchical workflow, spanning analysis of a single force network, a single experiment, and ensembles of experiments, provides a generalizable methodological principle for organizing complex data analysis across scales. The use of Sankey diagrams with color-based linkages to the original data, while preserving segment size and connectivity, allows effective visualization of temporal relationships as one-dimensional projections of higher-dimensional data. Although Sankey diagrams have been applied elsewhere, our implementation introduces a novel interpretative dimension that enhances interpretability and facilitates exploration of evolving spatial patterns.

VisPhoD has been well received by collaborators and establishes a robust foundation for future research. With minor adaptations, it can be applied to simulation data, enabling direct comparison between experimental and computational results. Our case studies demonstrated the framework's potential for multiscale analysis but also revealed areas for improvement, particularly in refining aggregated properties to improve robustness against outliers in the force network. Due to its modular set up, VisPhoD is well positioned to support more comprehensive, quantitative investigations of jamming and related phenomena as more and larger datasets become available.

Acknowledgements

This collaboration was initiated as a result of the Indo-Swedish joint network project funded by the Swedish Research Council (VR, Sweden) and the Department of Science and Technology (DST, India) with grant numbers: DST/INT/SWD/VR/P-02/2019 and VR 2018-07085. Additionally, this work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP), the Swedish e-Science Research Center (SeRC), the ELLIIT environment for strategic research in Sweden, and the Swedish Research Council (VR) grants 2019-05487 and 2023-04806. The authors acknowledge Peter Steneteg and Martin Falk for their support with Inviwo.

References

[1] S. J. Antony, "Power of photo-stress analysis in unravelling the mechanics of granular materials and its applications in interdisciplinary research," *Optics and Lasers in Engineering*, vol. 183, no. 108512, 2024.

[2] K. E. Daniels, J. E. Kollmer, and J. G. Puckett, "Photoelastic force measurements in granular materials," *Review of Scientific Instruments*, vol. 88, no. 5, 2017.

[3] B. Cambou, H. Magoariec, and N.-S. Nguyen, *Granular Materials at Meso-scale - Towards a Change of Scale Approach*. Elsevier, 2016.

[4] F. Rasheed, T. B. Masood, T. G. Murthy, V. Natarajan, and I. Hotz, "Multi-scale visual analysis of cycle characteristics in spatially-embedded graphs," *Visual Informatics*, vol. 7, no. 3, pp. 49–58, 2023.

[5] K. E. Daniels, "The role of force networks in granular materials," in *EPJ Web of Conferences*, vol. 140. EDP Sciences, 2017, p. 01006.

[6] A. G. Smart and J. M. Ottino, "Evolving loop structure in gradually tilted two-dimensional granular packings," *Physical Review E*, vol. 77, no. 4, p. 041307, 2008.

[7] R. Arévalo, L. A. Pugnaloni, I. Zuriguel, and D. Maza, "Contact network topology in tapped granular media," *Physical Review E*, vol. 87, no. 2, p. 022203, 2013.

[8] A. Tordesillas, D. M. Walker, G. Froyland, J. Zhang, and R. P. Behringer, "Transition dynamics and magic-number-like behavior of frictional granular clusters," *Physical Review E*, vol. 86, no. 1, p. 011306, 2012.

[9] A. Tordesillas, C. Steer, and D. M. Walker, "Force chain and contact cycle evolution in a dense granular material under shallow penetration," *Nonlinear Processes in Geophysics*, vol. 21, no. 2, pp. 505–519, 2014.

[10] M. R. Kuhn, W. Sun, and Q. Wang, "Stress-induced anisotropy in granular materials: fabric, stiffness, and permeability," *Acta Geotechnica*, 2015.

[11] R. Wan, P. Guo, and M. Al-Mamlun, "Behaviour of granular materials in relation to teir fabric dependencies," *Soils and foundations*, vol. 45, no. 2, pp. 77–86, 2005.

[12] C. Heine, H. Leitte, M. Hlawitschka, F. Iuricich, L. De Floriani, G. Scheuermann, H. Hagen, and C. Garth, "A survey of topology-based methods in visualization," in *Computer Graphics Forum*, vol. 35, no. 3. The Eurographs Association & John Wiley & Sons, Ltd. Chichester, UK, 2016, pp. 643–667.

[13] L. Yan, T. B. Masood, R. Sridharamurthy, F. Rasheed, V. Natarajan, I. Hotz, and B. Wang, "Scalar field comparison with topological descriptors: Properties and applications for scientific visualization," *Computer Graphics Forum*, vol. 40, no. 3, pp. 599–633, 2021.
[14] M. Hajij, B. Wang, C. Scheidegger, and P. Rosen, "Visual detection of structural changes in time-varying graphs using persistent homology," *IEEE Pacific Visualization Symposium (PacificVis)*, 2018.
[15] I. Obayashi, T. Nakamura, and Y. Hiraoka, "Persistent homology analysis for materials research and persistent homology software: Homcloud," *J. of the physical society of Japan*, vol. 91, no. 9, 2022.
[16] K. Pandey, T. B. Masood, S. Singh, I. Hotz, V. Natarajan, and T. Murthy, "Morse theory-based segmentation and fabric quantification of granular materials," *Granular Matter*, vol. 24, no. 1, 2022.
[17] L. Kondic, A. Goullet, C. O'Hern, M. Kramar, K. Mischaikow, and R. Behringer, "Topology of force networks in compressed granular media," *Europhysics Letters*, vol. 97, no. 5, p. 54001, 2012.
[18] M. Kramár, A. Goullet, L. Kondic, and K. Mischaikow, "Persistence of force networks in compressed granular media," *Physical Review E*, vol. 87, no. 4, p. 042207, 2013.
[19] ——, "Evolution of force networks in dense particulate media," *Physical Review E*, vol. 90, no. 5, p. 052203, 2014.
[20] L. Papadopoulos, M. A. Porter, K. E. Daniels, and D. S. Bassett, "Review: Network analysis of particles and grains," *Journal of Complex Networks*, vol. 6, pp. 485–565, 2018.
[21] C. Villani, *Topics in Optimal Transportation*. American Mathematical Society, 2003, no. 58.
[22] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, "Stability of persistence diagrams," *Discrete & Computational Geometry*, vol. 37, pp. 103–120, 2007.
[23] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko, "Lipschitz functions have l_p-stable persistence," *Foundations of Computational Mathematics*, vol. 10, no. 2, pp. 127–139, 2010.
[24] R. Sridharamurthy, T. B. Masood, A. Kamakshidasan, and V. Natarajan, "Edit distance between merge trees," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 26, no. 3, pp. 1518–1531, 2020.
[25] F. Wetzels and C. Garth, "A deformation-based edit distance for merge trees," *2022 Topological Data Analysis and Visualization (TopoInVis)*, pp. 29–38, 2022.
[26] B.-S. Sohn and C. Bajaj, "Time-varying contour topology," *IEEE Transaction on Visualization and Computer Graphics*, vol. 12, no. 1, pp. 14–25, 2006.
[27] W. Widanagamaachchi, C. Christensen, V. Pascucci, and P.-T. Bremer, "Interactive exploration of large-scale time-varying data using dynamic tracking graphs," in *IEEE Symposium on Large Data Analysis and Visualization (LDAV'12)*, 2012, pp. 9–17.
[28] J. Lukasczyk, G. Weber, R. Maciejewski, C. Garth, and H. Leitte, "Nested Tracking Graphs," *Computer Graphics Forum*, vol. 36, no. 3, pp. 12–22, 2017.
[29] E. Nilsson, J. Lukasczyk, T. B. Masood, C. Garth, and I. Hotz, "Probabilistic gradient-based extrema tracking," in *2023 Topological Data Analysis and Visualization (TopoInVis)*. IEEE, 2023, pp. 72–81.
[30] F. Rasheed, D. J"onsson, E. Nilsson, T. B. Masood, and I. Hotz, "Subject-specific brain activity analysis in fmri data using merge trees," in *Topological Methods in Visualization, IEEE workshop*, 2022.
[31] W. Köpp and T. Weinkauf, "Temporal treemaps: Static visualization of evolving trees," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 534–543, 2019.
[32] A. Dobler and M. Nöllenburg, "Improving temporal treemaps by minimizing crossings," *Computer Graphics Forum*, vol. 43, no. 3, p. e15087, 2024.
[33] H. Saikia and T. Weinkauf, "Global feature tracking and similarity estimation in time-dependent scalar fields," *Computer Graphics Forum*, vol. 36, no. 3, pp. 1–11, 2017.
[34] A. Schnorr, D. N. Helmrich, D. Denker, T. W. Kuhlen, and B. Hentschel, "Feature tracking by two-step optimization," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 26, no. 6, 2020.
[35] F. Rasheed, A. Naseer, E. Nilsson, T. B. Masood, and I. Hotz, "Multi-scale cycle tracking in dynamic planar graphs," *2024 IEEE Topological Data Analysis and Visualization (TopoInVis)*, pp. 44–54, 2024.
[36] Z. Mróz, *Comprehensive Structural Integrity*. Elsevier Science, 2003, vol. Volume 2: Fundamental Theories and Mechanisms of Failure, ch. Representative material elements, pp. 1–46.
[37] H. Edelsbrunner and J. Harer, "Persistent homology – a survey," in *Surveys on Discrete and Computational Geometry: Twenty Years Later*, J. E. Goodman, J. Pach, and R. Pollack, Eds. AMS Bookstore, 2008, vol. 458, pp. 257–282.
[38] A. Hatcher, *Algebraic Topology*. Cambridge University Press, 2002.
[39] M. Kramár, A. Goullet, L. Kondic, and K. Mischaikow, "Evolution of force networks in dense particulate media," *Physical Review E*, vol. 90, no. 5, p. 052203, 2014.
[40] D. Jönsson, P. Steneteg, E. Sundén, R. Englund, S. Kottravel, M. Falk, A. Ynnerman, I. Hotz, and T. Ropinski, "Inviwo - a visualization system with usage abstraction levels," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 26, no. 11, pp. 32–3254, 2020.

IX. Biography Section

Farhan Rasheed is a PhD student at Linköping University, specializing in scientific visualization. His research focuses on topological methods in scientific visualization. He received his Master of Science in Scientific Computing from Heidelberg University in Germany.

Abrar Naseer is a PhD student at the Indian Institute of Science, focusing on experimental research in granular materials.

Talha Bin Masood received his PhD in scientific visualization from the Indian Institute of Science and is currently an assistant professor in the Scientific Visualization Group at Linköping University, Sweden. His research focuses on the application of geometric and topological methods in natural sciences and engineering fields.

Tejas G. Murthy is a professor in the Department of Civil Engineering at the Indian Institute of Science. His research focuses on understanding granular materials. He received his PhD from Purdue University.

Vijay Natarajan is a Professor in the Department of Computer Science and Automation at the Indian Institute of Science. He received his PhD in Computer Science from Duke University. His current research interests include scientific visualization, computational topology, and geometry.

Ingrid Hotz is currently a Professor in Scientific Visualization at the Linköping University in Sweden. She received her Ph.D. degree from the Computer Science Department at the University of Kaiserslautern, Germany. Her research interests lie in data analysis and scientific visualization, ranging from basic research questions to effective solutions to visualization problems in applications.