

3D Visualization and Quantitative Evaluation of Graph Neural Network–Predicted Force Chains in Granular Materials


A THESIS

SUBMITTED FOR THE DEGREE OF
Master's of Technology

in Computer Science and Automation

by

Deepak Solanki



भारतीय विज्ञान संस्थान

Department of Computer Science and Automation
Indian Institute of Science
Bangalore – 560012

June 2026

© Deepak Solanki
June 2026
All rights reserved

Dedicated To
my family and friends

Declaration of Originality

I, Deepak Solanki, with SR No. 24796, hereby declare that the material presented in the thesis titled

3D Visualization and Quantitative Evaluation of Graph Neural Network–Predicted Force Chains in Granular Materials

represents original work carried out by me in the Department of CSA at the Indian Institute of Science during the years 2024-26.

With my signature, I certify that:

- I have not manipulated any of the data or results.
- I have not committed any plagiarism of intellectual property.
- I have clearly indicated and referenced the contributions of others.
- I have explicitly acknowledged all collaborative research and discussions.
- I have understood that any false claim will result in severe disciplinary action.
- I have understood that the work may be screened for any form of academic misconduct.

29th June, 2026

Date



Student Signature

In my capacity as supervisor of the above-mentioned work, I certify that the above statements are true to the best of my knowledge, and that I have carried out due diligence to ensure the originality of the report.

Prof. Vijay Natrajan

Advisor Name



Advisor Signature

Acknowledgements

I would like to express my sincere gratitude to my advisor, Prof. Vijay Natrajan, for his constant guidance, encouragement, and patience throughout this work.

I am also thankful to my examiner, Prof. Chiranjib Bhattacharya, for his valuable feedback, thoughtful suggestions, and encouragement, which helped improve this work.

I would also like to thank my colleagues and friends in the VGL Lab for creating a supportive and motivating environment.

Finally, I am deeply grateful to my parents, family, and friends for their endless love, understanding, patience, and encouragement.

Abstract

Force chains govern stress transmission in granular materials, yet direct measurement of inter-particle forces remains challenging. In this work, a Graph Neural Network (GNN) is implemented to predict the Normalized Particle Maximum Normal Contact Force (NPMNCF) at the particle level using data from discrete element method simulations. Granular assemblies are represented as graphs with particles and contacts modeled as nodes and edges, and an encode-process-decode architecture with iterative message passing is employed. The model is evaluated through training convergence, particle-scale correlation, and comparison of predicted and ground-truth force distributions. A persistent error analysis is further conducted to identify spatial regions where the model consistently underperforms, revealing that prediction failures concentrate at specimen boundaries due to incomplete graph neighborhoods and force chain termination gradients. Results show that the GNN captures particle-scale force heterogeneity and enables reliable identification of force-chain particles, providing a foundation for future contact-level force prediction and force-chain analysis

Contents

Acknowledgements	i
Abstract	ii
List of Figures	iv
List of Tables	v
1 Introduction	2
2 Literature Review	4
3 Methodology	6
3.1 DEM Data Description	6
3.2 Graph Representation of Granular Assemblies	7
3.3 Node and Edge Features	7
3.4 Prediction Target: NPMNCF	8
3.5 GNN Architecture Overview	8
3.6 Encoder Implementation	9
3.7 Processor: Message Passing with Residual Updates	9
3.8 Decoder	10
3.9 Training Configuration	10
3.10 Force-Chain Identification	11
4 Results	13
4.1 Training Convergence and Learning Behavior	13
4.2 Probability Density Distribution of NPMNCFs	15
4.3 NPMNCF Prediction During Uniaxial Compression	18
4.4 Spatial Evolution of Force Chains	19
4.5 Spatial Analysis of Persistent Prediction Errors	21
4.6 Factors Contributing to Boundary Error	23
5 Conclusion and Future Work	25
References	26

List of Figures

3.1	Encode–Process–Decode architecture of the proposed GNN model. . .	8
4.1	Training and testing loss curves for L1 and L2 loss functions.	13
4.2	Correlation vs epoch.	14
4.3	Probability density functions (PDFs) of normalized particle maximum normal contact forces (NPMNCFs) obtained from DEM and GNN predictions at different axial strain levels.	15
4.4	Prediction of NPMNCFs of a granular assembly at different axial strains under uniaxial compression.	18
4.5	Force Transmission Evolution	20
4.6	Spatial distribution of persistent absolute prediction error for Test Assembly 59. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.	22
4.7	Spatial distribution of persistent absolute prediction error for Test Assembly 60. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.	22
4.8	Spatial distribution of persistent absolute prediction error for Test Assembly 61. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.	23

List of Tables

4.1	Persistent error and spatial statistics.	21
-----	--	--------------------

Force chains govern stress transmission in granular materials, yet direct measurement of inter-particle forces remains challenging. In this work, a Graph Neural Network (GNN) is implemented to predict the Normalized Particle Maximum Normal Contact Force (NPMNCF) at the particle level using data from discrete element method simulations. Granular assemblies are represented as graphs with particles and contacts modeled as nodes and edges, and an encode–process–decode architecture with iterative message passing is employed. The model is evaluated through training convergence, particle-scale correlation, and comparison of predicted and ground-truth force distributions. A persistent error analysis is further conducted to identify spatial regions where the model consistently underperforms, revealing that prediction failures concentrate at specimen boundaries due to incomplete graph neighborhoods and force chain termination gradients. Results show that the GNN captures particle-scale force heterogeneity and enables reliable identification of force-chain particles, providing a foundation for future contact-level force prediction and force-chain analysis

Chapter 1

Introduction

Granular materials exhibit complex mechanical behavior governed by highly heterogeneous transmission of inter-particle contact forces. Rather than being uniformly distributed, contact forces organize into localized chain-like structures, commonly referred to as force chains, which form the primary load-bearing skeleton of granular assemblies. Radjai et al. [8] first demonstrated that force transmission in granular packings is bimodal in nature, consisting of strong and weak contact networks that play distinct mechanical roles. The strong-force network, composed of contacts carrying forces higher than the system average, dominates stress transmission and controls macroscopic responses such as stiffness, strength, and strain localization.

Understanding the formation and evolution of force chains is therefore central to the micromechanical interpretation of granular materials. Numerical studies using the discrete element method (DEM) have provided valuable insights into three-dimensional force-chain structure and anisotropy; however, such approaches rely on simulated data and are not directly accessible in physical experiments. Experimental investigation of force chains has traditionally been limited to two-dimensional photoelastic systems, which cannot be readily extended to realistic three-dimensional granular packings.

Recent advances in X-ray micro-computed tomography (CT) have enabled non-destructive, three-dimensional measurement of particle kinematics and inter-particle contact evolution in granular materials under mechanical loading. Using CT-based experiments, Cheng and Wang [1] provided detailed grain-scale observations of contact network evolution during shearing, offering new opportunities to study granular micromechanics in three dimensions. Despite this progress, CT imaging does not directly measure inter-particle contact forces, making experimental identification of three-dimensional force chains a long-standing challenge.

To overcome this limitation, machine learning approaches based on Graph Neural Networks (GNNs) have recently been introduced to infer inter-particle force measures from particle kinematics and contact information. In particular, Cheng et al. [3] proposed a GNN-based framework for predicting particle-scale force descriptors and identifying force-chain particles in three-dimensional granular materials. While this

approach represents a significant methodological advance, its practical implementation and validation across different granular assemblies remain areas of active investigation.

Motivated by this context, the present work focuses on the implementation and evaluation of a GNN-based framework for particle-scale force prediction in granular materials using discrete element method data. Emphasis is placed on assessing model convergence, prediction accuracy, and the ability of the learned force measures to distinguish particles associated with strong force networks. This work focuses on establishing and validating a computational framework for particle-scale force prediction, which provides a basis for future force-chain studies..

Chapter 2

Literature Review

The study of force chains has been an important topic in understanding the mechanical behavior of granular materials. Previous research has shown that stress transmission in granular assemblies is highly non-uniform and occurs through localized chains of strong inter-particle contacts. Radjai et al. [8] demonstrated that contact forces in granular packings exhibit a bimodal distribution, consisting of strong and weak contact networks. The strong contact network forms the main load-bearing structure of the material and is commonly used to identify force chains.

To describe the internal structure and anisotropy of granular materials, several researchers have introduced tensor-based measures. The fabric tensor was proposed to quantify the directional distribution of contacts and particle arrangements, and its relationship with macroscopic mechanical behavior has been widely studied [10, 7]. Rothenburg and Bathurst [9] further showed that induced fabric anisotropy plays a key role in stress-strain behavior. These studies established that fabric and stress tensors provide an effective way to characterize the orientation and evolution of force chains under loading.

Later studies extended force-chain analysis to three-dimensional systems. Voivret et al. [12] investigated multiscale force networks in granular materials and showed that force transmission exhibits strong anisotropy across different scales. Tordesillas et al. [11] studied the evolution and buckling of force chains during deformation, highlighting the dynamic nature of load-bearing structures. Guo and Zhao [4] analyzed shear-induced anisotropy using fabric and stress tensors, further linking force-chain evolution to macroscopic response.

Experimental investigation of granular materials at the particle scale has been significantly advanced by the use of X-ray micro-computed tomography (CT). Using CT-based experiments, Cheng and Wang [1] studied inter-particle contact evolution in sheared granular materials and provided detailed grain-scale observations of contact and fabric changes. While CT enables accurate measurement of particle kinematics and contact networks in three dimensions, it does not allow direct measurement of inter-particle contact forces. As a result, experimental identification of force chains in three-dimensional granular materials remains challenging.

To overcome this limitation, machine learning approaches have been introduced to estimate contact forces from experimentally measurable quantities. Cheng and Wang [2] proposed a machine learning model to estimate contact forces in granular materials under uniaxial compression using particle-scale features. Building on this work, Cheng et al. [3] developed a Graph Neural Network (GNN)-based strategy to experimentally estimate force chains using CT-measured particle kinematics and contact evolution. Their approach models granular assemblies as graphs composed of particles and contacts and predicts force-related quantities that can be used to construct strong-force networks. The evolution of force chains is analyzed using force statistics as well as fabric and stress tensor measures.

In addition to the work by Cheng and co-workers, other studies have explored the use of deep learning and graph-based methods for predicting force networks in granular materials. Mandal et al. [6] showed that GNNs can robustly predict force-chain formation by explicitly considering contact network connectivity. Li et al. [5] demonstrated that GNN-based models can predict contact force networks across different packing conditions. Although these studies show the effectiveness of GNNs for force prediction, force-chain analysis is often limited to qualitative visualization or a small number of global statistical metrics.

From the literature, it can be observed that while significant progress has been made in predicting force chains using machine learning, comprehensive three-dimensional visualization and quantitative evaluation of force-chain geometry, connectivity, and evolution remain limited. This motivates the present work, which focuses on improving the interpretation of GNN-predicted force chains through enhanced 3D visualization and physically meaningful evaluation metrics.

Chapter 3

Methodology

This work implements a Graph Neural Network (GNN) framework to predict particle-scale force measures in granular materials, following the machine learning-based strategy proposed by Cheng et al. [3]. The methodology is based on an encode-process-decode architecture and is designed to predict the Normalized particle Mean Normal Contact Force (NPMNCF) for each particle. All mathematical symbols and model parameters used in the implementation are explicitly defined in this section.

3.1 DEM Data Description

The data used in this study are obtained from three-dimensional Discrete Element Method (DEM) simulations of granular assemblies subjected to uniaxial compression. DEM simulations provide full particle-scale information, including particle geometry, kinematics, contact networks, and contact forces, which are not directly accessible in experimental CT-based measurements.

Each granular assembly consists of spherical particles simulated under controlled boundary conditions. The simulations are performed incrementally in strain, and particle-level information is recorded at each strain level. For every strain increment, the DEM outputs include particle positions, particle radii, particle displacements, coordination numbers, inter-particle contact lists, and normal contact forces.

From these DEM outputs, particle kinematics and contact evolution information are extracted and used as inputs to the GNN model. Particle radius and coordination number are used to construct node features, while contact status and relative displacements are used to construct edge features. The corresponding NPMNCF values computed from DEM simulations serve as supervised learning targets.

3.2 Graph Representation of Granular Assemblies

A granular assembly at a given strain level is represented as a graph

$$G = (V, E),$$

where V denotes the set of particles and E denotes the set of inter-particle contacts. Each particle is represented as a node $i \in V$, and each contact between particles i and j is represented as an edge $(i, j) \in E$.

Let $N = |V|$ denote the total number of particles and $M = |E|$ denote the total number of contacts. To enforce physical symmetry of contact forces, each contact is represented using two directed edges $(i \rightarrow j)$ and $(j \rightarrow i)$, resulting in $2M$ directed edges. This bidirectional formulation enables symmetric information exchange between contacting particles and improves physical consistency.

3.3 Node and Edge Features

Each particle (node) i is associated with a feature vector

$$\mathbf{x}_i \in \mathbb{R}^2,$$

defined as

$$\mathbf{x}_i = [r_i, CN_i],$$

where r_i is the particle radius normalized by the maximum particle radius in the assembly, and CN_i is the coordination number, defined as the number of contacts associated with particle i . These features represent basic geometric and topological properties influencing force transmission.

Each directed contact (edge) $(i \rightarrow j)$ is associated with a feature vector

$$\mathbf{e}_{ij} \in \mathbb{R}^3,$$

defined as

$$\mathbf{e}_{ij} = [s_{ij}, u_{ij}^n, u_{ij}^t],$$

where s_{ij} denotes the contact status (new, persistent, or lost between consecutive strain increments), u_{ij}^n denotes the relative displacement along the contact normal direction, and u_{ij}^t denotes the relative displacement perpendicular to the contact normal. These kinematic features encode the local deformation mechanisms governing

normal contact force evolution.

3.4 Prediction Target: NPMNCF

The supervised learning target is the Normalized particle Mean Normal Contact Force (NPMNCF) associated with each particle. For a particle i with coordination number CN_i , the mean normal contact force $MNCF_i$ is computed as

$$MNCF_i = \frac{1}{CN_i} \sum_{j \in \mathcal{N}(i)} F_{ij}^n,$$

where $\mathcal{N}(i)$ denotes the set of neighboring particles in contact with particle i , and F_{ij}^n denotes the normal contact force between particles i and j .

The particle value of $MNCF_i$ over the loading history is then normalized by the mean normal contact force of the entire assembly to obtain the NPMNCF value. This particle-level force descriptor is robust and sufficient for identifying force-chain particles.

3.5 GNN Architecture Overview

The GNN architecture consists of three components: an encoder, a processor, and a decoder. All node and edge embeddings are represented in a latent space of dimension $d = 128$, where d denotes the hidden feature dimension. This value is selected to provide sufficient capacity to represent the heterogeneous and long-tailed force distributions observed in granular materials.




Figure 3.1: Encode-Process-Decode architecture of the proposed GNN model.

3.6 Encoder Implementation

The encoder maps raw node and edge features into the latent space. For each directed edge ($i \rightarrow j$), the input to the edge encoder is constructed by concatenating the feature vectors of the receiver particle i , the sender particle j , and the edge feature:

$$\mathbf{z}_{ij}^{(0)} = [\mathbf{x}_i \parallel \mathbf{x}_j \parallel \mathbf{e}_{ij}],$$

where $\parallel$ denotes vector concatenation. The resulting vector $\mathbf{z}_{ij}^{(0)} \in \mathbb{R}^7$ is passed through a multilayer perceptron (MLP) to produce an encoded edge embedding

$$\mathbf{h}_{ij}^{(0)} \in \mathbb{R}^d.$$

For each node i , the encoded edge embeddings connected to the node are aggregated using mean aggregation:

$$\bar{\mathbf{h}}_i^{(0)} = \frac{1}{CN_i} \sum_{j \in \mathcal{N}(i)} \mathbf{h}_{ji}^{(0)}.$$

The node encoder input is then defined as

$$\mathbf{u}_i^{(0)} = [\mathbf{x}_i \parallel \bar{\mathbf{h}}_i^{(0)}],$$

which is passed through a second MLP to obtain the initial node embedding

$$\mathbf{v}_i^{(0)} \in \mathbb{R}^d.$$

3.7 Processor: Message Passing with Residual Updates

The processor performs $K = 7$ message-passing iterations to propagate force-related information across the contact network. At iteration k , edge and node embeddings are updated sequentially.

For each directed edge ($i \rightarrow j$), the edge update is computed as

$$\Delta \mathbf{h}_{ij}^{(k)} = \text{MLP}_e([\mathbf{v}_i^{(k-1)} \parallel \mathbf{v}_j^{(k-1)} \parallel \mathbf{h}_{ij}^{(k-1)}]),$$

followed by a residual update

$$\mathbf{h}_{ij}^{(k)} = \mathbf{h}_{ij}^{(k-1)} + \Delta \mathbf{h}_{ij}^{(k)}.$$

For each node i , the updated edge embeddings are aggregated as

$$\mathbf{m}_i^{(k)} = \frac{1}{CN_i} \sum_{j \in \mathcal{N}(i)} \mathbf{h}_{ji}^{(k)}.$$

The node update is computed as

$$\Delta \mathbf{v}_i^{(k)} = \text{MLP}_v([\mathbf{v}_i^{(k-1)} \parallel \mathbf{m}_i^{(k)}]),$$

followed by a residual update

$$\mathbf{v}_i^{(k)} = \mathbf{v}_i^{(k-1)} + \Delta \mathbf{v}_i^{(k)}.$$

Residual connections enable progressive refinement of embeddings and improve gradient flow during training.

3.8 Decoder

After $K = 7$ message-passing iterations, the final node embedding

$$\mathbf{v}_i^{(K)}$$

is passed through a decoder MLP:

$$\hat{y}_i = \text{MLP}_{\text{dec}}(\mathbf{v}_i^{(K)}),$$

where $\hat{y}_i$ denotes the predicted NPMNCF value for particle i . No activation function is applied at the output layer, allowing prediction over the full force range.

3.9 Training Configuration

The model is trained in a supervised manner using DEM-generated data. Training is performed with a batch size of one graph due to the large graph size. Model parameters are optimized using the Adam optimizer with a learning rate of 10^{-3} .mean

squared error loss is used:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2,$$

where y_i is the ground-truth NPMNCF value and w_i assigns higher weight to particles associated with force chains.

3.10 Force-Chain Identification

Force chains are identified by applying a threshold to the predicted NPMNCF values. Particles with predicted values exceeding the threshold are classified as force-chain particles. Connected components of these particles within the contact network are interpreted as force chains, following the strong-force network concept introduced in [3].

Summary of Work Completed

This project focused on implementing and evaluating a graph-based learning approach for particle-scale force prediction in granular materials. The main tasks completed are summarized below:

- Studied and understood the Graph Neural Network framework proposed by Cheng et al., with emphasis on graph construction, input features, and the encode–process–decode architecture.
- Selected a uniaxial compression case from discrete element method (DEM) simulations to serve as a representative loading condition for force transmission analysis.
- Processed DEM output data to construct particle–contact graphs, including particle features, contact connectivity, and relative displacement-based edge features.
- Implemented the GNN model independently based on the description in the literature and trained it to predict the Normalized particle Mean Normal Contact Force (NPMNCF) at the particle level.
- Evaluated the model performance using training convergence, correlation with DEM ground-truth forces, and comparison of predicted and actual force distributions.


Overall, this work establishes a working and validated implementation of a GNN-based force prediction model, which can be extended in future studies to more detailed force-chain analyses.

Chapter 4

Results

This section presents the results of the implemented Graph Neural Network (GNN) model for predicting the Normalized particle Mean Normal Contact Force (NPMNCF) in granular materials. The evaluation focuses on training convergence, prediction accuracy, and reproduction of particle-scale force distributions. All results correspond to the bidirectional GNN model evaluated at epoch 100.

4.1 Training Convergence and Learning Behavior



(a) Training and testing L1 loss versus epoch.

(b) Training and testing L2 loss versus epoch.

Figure 4.1: Training and testing loss curves for L1 and L2 loss functions.

Figure 4.1a shows how the MAE (L1 loss) changes during training and testing. In the beginning, both losses decrease very quickly, in initial phase model is able to learn the relationship between particle motion and contact force effectively. After around 70–80 epochs, the curves become almost stable, showing that the model has

converged properly. The training and testing curves are very close to each other throughout the process, which indicates that the model generalizes well on unseen data. Small fluctuations in the testing curve are expected because the DEM data contains different particle arrangements and continuously changing contact networks. The final MAE value of nearly 0.29 indicates that the predicted normalized contact forces are reasonably close to the actual DEM results.




Figure 4.2: Correlation vs epoch.

Figure 4.1b presents the evolution of the MSE (L2 loss) during training and testing. Similar to the MAE behavior, the MSE loss decreases steadily with increasing epochs and eventually reaches a stable plateau. This confirms that the optimization process is stable and that the proposed GNN model successfully captures the force transmission behavior in granular materials. The close agreement between training and testing losses further demonstrates stable learning and reliable prediction capability.

4.2 Probability Density Distribution of NPMNCFs




Figure 4.3: Probability density functions (PDFs) of normalized particle maximum normal contact forces (NPMNCFs) obtained from DEM and GNN predictions at different axial strain levels.

Figure 5 presents a systematic comparison of normalized particle maximum normal contact force (NPMNCF) distributions between discrete element method (DEM) simulations and graph neural network (GNN) predictions across seven axial strain levels (0.5% to 8.0%). The horizontal axis spans NPMNCF values from 0 to approximately 8, where each value represents the maximum normal contact force experienced by a particle normalized by the assembly's characteristic force scale. This normalization procedure enables direct comparison across different strain states and assemblies, with unity representing the mean normalized force.

The linear-scale plots (5a and 5b) reveal the bulk distribution characteristics

through the probability density function (PDF). Both DEM and GNN distributions exhibit a prominent peak centered at $\text{NPMNCF} \approx 1.0$, indicating that the majority of particles experience forces near the normalized mean. The DEM distribution reaches a maximum PDF value of approximately 2.8, while the GNN prediction achieves approximately 3.9, suggesting the neural network model predicts slightly sharper concentration around the mean. This indicates a mild underprediction in the intermediate force regime ($0.5 \leq \text{NPMNCF} \leq 2.0$), though the magnitude of this deviation is quantitatively minor and does not substantially compromise model fidelity.

A salient observation across the linear plots is the near-universal collapse of all seven strain-level distributions. This strain-independent behavior of normalized force statistics—spanning a compression range from 0.5% to 8.0%—is physically significant. It demonstrates that the normalized force distribution follows a characteristic universal form, indicating that granular materials maintain consistent load-sharing mechanisms and stress transmission patterns across varying deformation states. This universality suggests that the fundamental principles governing force distribution in granular systems operate independently of the instantaneous strain condition within the examined range.

The semi-logarithmic representation (5c and 5d) reveals force-chain structure invisible on linear scales. By compressing the vertical axis logarithmically, rare but mechanically critical events become statistically visible as an approximately exponential tail extending from $\text{NPMNCF} \approx 1$ toward higher values. Particles with $\text{NPMNCF} > 3.0$ represent force-chain carriers, occurring at frequencies of 10^{-3} to 10^{-4} statistically rare but mechanically essential for stress transmission. The exponential decay pattern indicates that high-force particles obey a consistent statistical law rather than appearing as random anomalies.

The GNN successfully reproduces this tail structure, with the slope and general behavior closely matching DEM predictions. Quantitative differences emerge at extreme force values ($\text{NPMNCF} > 3.0$), where the GNN predicts slightly lower probability densities. This modest underprediction of extreme events reflects the inherent difficulty in learning rare-event statistics from finite training datasets and does not substantially compromise the model’s predictive utility.

The effectiveness of this four-panel comparison derives from the complementary information provided by both scales. Linear representation reveals that approximately 80% of particles cluster between $\text{NPMNCF} = 0.2$ and 3.0, emphasizing typical bulk behavior. Semi-log representation illuminates the mechanically critical force-chain population following exponential statistics. The GNN demonstrates simultaneous

success at both perspectives: accurately reproducing bulk force distributions while capturing force-chain physics through exponential tail prediction. This dual capability—correctly predicting both typical and extreme events—indicates that the model has internalized fundamental principles of stress heterogeneity in granular systems rather than memorizing strain-specific patterns, suggesting robust generalization capability across different mechanical states.

4.3 NPMNCF Prediction During Uniaxial Compression




Figure 4.4: Prediction of NPMNCFs of a granular assembly at different axial strains under uniaxial compression.

Figure 4.4 The model performed remarkably well across all sixteen strain levels tested, maintaining correlation coefficients between 0.79 and 0.85. This consistency across different compression stages suggests the network has learned meaningful features of granular force transmission that transfer well to unseen test assemblies.

During early compression (0.5-2.5% strain), predictions were most accurate with correlations hovering around 0.84-0.85. The scatter plots show points clustering

tightly around the ideal prediction line, indicating the model could reliably estimate forces when the granular packing was still relatively loose and disordered. At this stage, forces correlate fairly directly with local particle contacts and displacements, making the task more straightforward for a network trained on kinematic data.

However, we observed a distinct performance dip in the intermediate loading range (3.0-6.5% strain), where correlations dropped to 0.796 at 6.5% strain. This isn't surprising given what's physically happening the granular system undergoes major reorganization during this phase. Particles rotate significantly, the contact network reshuffles, and force chains begin forming. The scatter plots widen noticeably here, particularly in the upper range (high force particles), indicating genuine difficulty predicting forces for load-bearing particles. The model seems to systematically underpredict high forces during this transition, clustering points slightly below the diagonal line.

Interestingly, performance recovered at higher strains (7.0-8.0%), with correlations rising back to 0.82-0.84. We attribute this to the stabilization of force chain geometry at large deformations. Once the system reaches mature force chains aligned with the loading direction, forces become more predictable again from kinematic inputs alone. The network essentially trades the transient complexity of reorganization for the stability of established load paths.

One observation worth highlighting: predictions for low-force particles remained consistently good across all strain levels, with minimal scatter regardless of compression stage. Since most particles in the system carry relatively low contact forces, this accuracy on the bulk population partly explains why overall correlations stay above 0.79 even during the problematic intermediate regime.

4.4 Spatial Evolution of Force Chains

Figure 7 reveals one of the most visually striking aspects of granular mechanics: how force chains—the load-bearing networks within a compressed material—evolve as stress increases. The visualization compares actual contact force distributions from DEM simulations against GNN predictions across five strain levels (1.5%, 3.0%, 4.5%, 6.0%, 7.5%), presented at two perpendicular planes cutting through the assembly's center.




Figure 4.5: Force Transmission Evolution

At the lowest strain level (1.5%), the granular system appears almost uniformly red, with scattered yellow hotspots distributed seemingly at random throughout both the XZ and YZ planes. This randomness is not accidental—it reflects the granular material’s initial state where particles are still rearranging and contacts are forming relatively uniformly. The DEM and GNN predictions align closely at this stage, both showing the same chaotic force distribution pattern. The material hasn’t yet "decided" where to place its load-bearing structures.

As strain increases to 3.0%, the visualization still shows a fairly distributed pattern, though certain regions begin to develop slightly higher force concentrations. The material is still in a transitional phase, gradually organizing its internal structure but without yet forming the dramatic force chains that appear at higher strains. The GNN continues to match the DEM results well, capturing this intermediate organizational state accurately.

The transformation becomes pronounced at 4.5% and especially at 6.0% strain. Here, the imagery shifts dramatically. Yellow and orange lines begin appearing—these

are force chains, the load-bearing structures. They manifest as linear features oriented predominantly in the vertical direction, aligned with the applied compression axis. This makes physical sense: particles aligned vertically can efficiently transmit the vertical load down through the assembly. The network becomes less democratic and more organized, with clear "strong" pathways (yellow) interspersed with "weak" regions (red). The GNN predictions capture this transition remarkably well, showing the same chain structures in the same locations as the actual DEM simulations.

At 7.5% strain, the pattern crystallizes further. Prominent vertical load-bearing chains dominate both planes, creating a distinctive lattice-like pattern where most of the stress flows through these organized pathways. The weak network (red regions) provides lateral support, preventing chain buckling and collapse. This hierarchical force distribution is what allows granular materials to support large vertical stresses without continuous failure.

4.5 Spatial Analysis of Persistent Prediction Errors

Aggregate metrics such as MAE and Pearson correlation measure overall model accuracy but do not reveal *where* predictions fail. To address this, the absolute error for each particle was averaged across all 81 compression timesteps, yielding a **persistent absolute error** that identifies particles the GNN consistently mispredicts throughout the loading history. This field was computed for three test assemblies (59, 60, 61) and visualized in ParaView with reduced opacity so that only high-error hotspots remain visible .

Table 4.1: Persistent error and spatial statistics.

	Assembly 59	Assembly 60	Assembly 61
Number of particles	9,019	9,127	9,076
Mean error	0.291	0.293	0.300
Median error	0.276	0.277	0.285
95th percentile	0.528	0.540	0.547
99th percentile	0.780	0.788	0.781
Maximum error	1.415	1.504	1.535
Top 5% particles near boundary	66.7%	72.9%	67.0%
Coordination number ≤ 3	34.6%	36.4%	35.8%
NPMNCF ($\sigma/\text{maximum}$)	0.76 / 7.1	0.77 / 8.0	0.76 / 5.7
Contacts ($t = 1 \rightarrow 80$)	13k $\rightarrow$ 22k	12k $\rightarrow$ 22k	13k $\rightarrow$ 22k




Figure 4.6: Spatial distribution of persistent absolute prediction error for Test Assembly 59. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.




Figure 4.7: Spatial distribution of persistent absolute prediction error for Test Assembly 60. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.




Figure 4.8: Spatial distribution of persistent absolute prediction error for Test Assembly 61. Error values are averaged across all 81 compression timesteps. High-error regions are predominantly concentrated near specimen boundaries and force-chain endpoints.

4.6 Factors Contributing to Boundary Error

Six interrelated mechanisms explain why the GNN persistently fails at boundary regions:

(1) Incomplete neighborhoods. Interior particles have 4–7 neighbors providing rich context for message passing. Boundary particles have contacts on only one side, often with $CN \leq 3$ (34–36% of observations). With one or two incoming messages, noise is amplified rather than averaged out, yielding unreliable latent representations.

(2) Force chain termination gradients. Force chains transmit compressive load through the assembly and terminate abruptly at free surfaces, creating sharp force gradients. The NPMNCF distribution is heavy-tailed ($\sigma \approx 0.76$, max up to $8\times$ mean), and MSE loss causes the model to regress extreme values toward the mean. Since force chain endpoints coincide with boundary particles, this systematic bias directly produces the observed hotspots.

(3) Truncated receptive field. Seven message-passing steps give each particle a receptive field spanning 7 hops. For boundary particles, much of this field extends beyond the specimen into empty space, creating an asymmetric cone-shaped information region. The node features (normalized radius and CN) lack explicit boundary encoding—a low CN could indicate either a boundary particle or an early-compression

interior particle—so the model cannot distinguish these structurally different cases.

(4) Mean-aggregation bottleneck. Node updates aggregate edge messages via mean pooling: $\bar{\mathbf{e}}_i = |\mathcal{N}(i)|^{-1} \sum_j \mathbf{e}_{ij}$. At a boundary particle with two contacts—one strong chain and one weak contact—the mean washes out the strong signal. Interior particles with the same chain plus several moderate contacts produce a more representative average. Few samples from a heavy-tailed distribution yield high-variance estimates.

(5) Training data imbalance. Boundary particles constitute $\sim 15\text{--}20\%$ of the specimen. Uniform MSE loss weights every particle equally, so gradients are dominated by the $\sim 80\%$ interior majority. Over 100 epochs, the model is optimized primarily for interior particles, under-optimizing the difficult boundary minority.

(6) Evolving contact topology. The contact network grows by 68–79% during compression, producing a structurally different graph at each timestep. The model processes timesteps independently without temporal memory, so particles in topologically unstable regions—where contacts are rapidly created or broken—yield inconsistent predictions that accumulate into high persistent error.

Summary

The persistent error analysis demonstrates that GNN prediction failures are systematically concentrated at specimen boundaries and force chain endpoints, driven by incomplete graph neighborhoods, extreme target value gradients, truncated receptive fields, aggregation bottlenecks, training imbalance, and evolving topology. Potential improvements include explicit boundary encoding in node features, attention-based aggregation, weighted loss functions prioritizing boundary particles, and recurrent architectures for temporal context.

Chapter 5

Conclusion and Future Work

The present study investigates force prediction in granular materials using a graph neural network trained on DEM simulations of spherical glass bead packings under uniaxial compression. Although the model demonstrates strong predictive performance for the considered dataset, its applicability to other granular systems remains an open question. A natural extension of this work is to evaluate the proposed framework on different types of granular materials, including particles with varying size distributions, shapes, frictional properties, and material characteristics. Such investigations would help determine the generalizability and robustness of graph-based force prediction models across a wider range of granular assemblies.

Another promising direction is the incorporation of physical constraints directly into the learning process through Physics-Informed Graph Neural Networks (PI-GNNs). The current model learns force relationships purely from data and does not explicitly enforce governing physical principles. By embedding constraints such as force balance, contact mechanics laws, stress equilibrium, and energy conservation into the training objective, the model could produce predictions that are not only accurate but also physically consistent. Such an approach would provide deeper insight into the underlying mechanics governing force transmission during uniaxial compression.

Future work may also explore the integration of temporal information into the graph-learning framework. Since the contact network evolves continuously during compression, incorporating temporal graph neural networks or recurrent architectures could enable the model to capture the dynamic evolution of force chains and contact topology more effectively than independent timestep predictions.

Finally, combining physics-informed learning with detailed contact-level force prediction would enable a comprehensive study of force-chain evolution, stress transmission, and fabric anisotropy in granular materials. Such developments would move graph-based models beyond predictive tools toward interpretable frameworks capable of revealing the fundamental physics governing granular deformation and force propagation.

References

- [1] Z. Cheng and J. Wang. Experimental investigation of inter-particle contact evolution of sheared granular materials using x-ray micro-tomography. Soils and Foundations, 58(6):1492–1510, 2018.
- [2] Z. Cheng and J. Wang. Estimation of contact forces of granular materials under uniaxial compression based on a machine learning model. Granular Matter, 24(1):1–15, 2022.
- [3] Z. Cheng, J. Wang, and W. Xiong. A machine learning-based strategy for experimentally estimating force chains of granular materials using x-ray micro-tomography. Géotechnique, 74(12):1291–1303, 2024.
- [4] N. Guo and J. Zhao. The signature of shear-induced anisotropy in granular media. Computers and Geotechnics, 47:1–15, 2013.
- [5] Y. Li, N. Brodu, and J. A. Dijksman. Predicting contact force networks in granular materials using graph neural networks. The Journal of Chemical Physics, 159(4):044902, 2023.
- [6] S. Mandal, D. M. Sussman, and D. S. Bassett. Robust prediction of force chains in granular materials using graph neural networks. Soft Matter, 18(14):2644–2655, 2022.
- [7] M. Oda. Initial fabrics and their relations to mechanical properties of granular material. Soils and Foundations, 12(1):17–36, 1972.
- [8] F. Radjai, D. E. Wolf, M. Jean, and J.-J. Moreau. Bimodal character of stress transmission in granular packings. Physical review letters, 80(1):61, 1998.
- [9] L. Rothenburg and R. J. Bathurst. Analytical study of induced anisotropy in idealized granular materials. Géotechnique, 39(4):601–614, 1989.
- [10] M. Satake. Fabric tensor in granular materials. Proceedings of the IUTAM Conference on Deformation and Failure of Granular Materials, pages 63–68, 1982.
- [11] A. Tordesillas, D. M. Walker, and Q. Lin. Force chains and contact cycle evolution in granular materials. Physical Review E, 81(1):011302, 2010.

-
- [12] C. Voivret, F. Radjai, J.-Y. Delenne, and M. S. El Youssoufi. Multiscale force networks in granular media. Physical Review Letters, 102(17):178001, 2009.