CrackGNN: Revealing a Crack Severity Manifold

Shantanu Motiani[1,2], Maarten Stol[2], and Arun Mukundan[1]

[1] University of Amsterdam, Faculty of Science
[2] BrainCreators, Amsterdam

1 Introduction and Methodology

The structural integrity and serviceability of transportation infrastructure, particularly roadway networks, are paramount for economic prosperity and public safety. Pavement cracking is a primary form of distress that, if left unaddressed, can propagate rapidly, leading to significant structural failures like potholes and rutting, compromising pavement quality, service life, and user safety, while drastically increasing maintenance costs [1]. Pavement Management Systems face critical bottlenecks from traditional labor-intensive manual inspections with high error rates [1][8]; and while supervised computer vision approaches (such as CNN-based segmentation models) effectively detect and outline cracks [13][2], they fall short in quantifying crack *severity* without large, high-quality labeled datasets [3][10]. These black-box methods struggle with domain-shift and the lack of transparency required for high-stakes infrastructure decisions. Furthermore, current industry practice treats severity as a series of discrete classes, a paradigm that overlooks the continuous nature of pavement degradation [4].

To address these challenges, we introduce CrackGNN, an unsupervised framework that reframes crack severity assessment as a representation learning problem. We extract seven interpretable features from crack skeletons (junction count, box dimension, tortuosity, orientation, propagation potential, branching factor, continuity) to encode complexity and morphology relevant to severity [9]. Using these as nodes (with features), we construct inter-crack dependency graphs by estimating pairwise mutual information (MI) with the Kraskov-Stögbauer-Grassberger method [6], sparsifying by retaining top-50% MI edges for noise reduction. This graph is fed to our GNN network which fuses an autoencoder and a two-layer GCN [5], producing embeddings that reflect both features and relationships. A learnable fusion parameter combines these, and embeddings are clustered with silhouette-guided adjustment.

In our loss function, we encourage correlation across embedding dimensions. This approach can be considered counter-intuitive, as a common goal in machine learning is to learn disentangled representations where each embedding dimension corresponds to an independent feature [7]. The naive expectation would be that such de-correlated dimensions would offer a clearer and more modular understanding of the data and avoid "representation collapse" [11]. However, by encouraging correlation, we allow the model to identify and amplify the most critical underlying factor in the data. Instead of distributing the concept of severity across multiple, independent dimensions, our method enables the model to consolidate this information into a primary axis of variation, creating a powerful and unified representation of the feature we aim to measure

The novelty of this work is the unsupervised discovery of an intrinsic severity manifold (Section 2). This discovery offers both immediate practical value and fundamental theoretical insights. For the industry, it provides a principled foundation for crack severity assessment that eliminates subjective bias and the need for vast labelled datasets, while maintaining interpretability. For researchers, it opens new directions in infrastructure informatics, unsupervised learning, and confirms the manifold hypothesis in a novel, real-world application. By bypassing the data bottlenecks of supervised methods, CrackGNN demonstrates that unsupervised learning can uncover hidden engineering data structures, offering a scientifically rigorous and practically valuable solution for modern Pavement Management Systems.

2 The Crack Severity Manifold

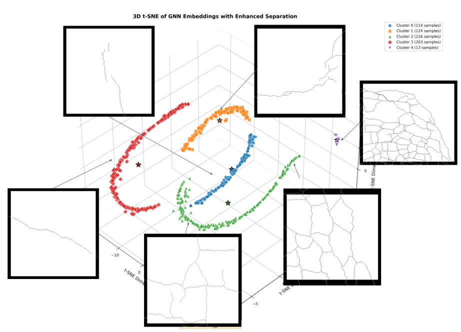

Fig. 1. t-SNE visualization of CrackGNN's embeddings (severity manifold) s

Visualizing CrackGNN embeddings through t-SNE (Figure 1) reveals organization into a spiral-shaped manifold where severity increases monotonically, revealing that pavement crack severity can be understood as a single, dominant dimension despite apparent morphological complexity. Simple linear cracks occupy one end while complex alligator networks occupy the opposite. High-severity clusters exhibit elevated junction counts, fractal dimensions, and tortuosity—signatures expected for advanced alligator cracking. Low-severity clusters display minimal junction activity and linear morphologies consistent with early-stage cracking. This provides evidence that the manifold hypothesis [12] could govern pavement crack analysis, this means that the ostensibly intricate, multi-feature representations can in fact be unfurled into a single, interpretable severity axis. Rather than a pathological collapse, maximizing latent correlation in CrackGNN reveals the intrinsic simplicity of the task: *severity is the overwhelming source of variation in the feature set.*

Acknowledgments

Author: Shantanu Motiani `shantanumotiani@gmail.com`
University Supervisor: Arun Mukundan (Uniersity of Amsterdam)
`a.mukundan@uva.nl`
External Supervisor: Maarten Stol (BrainCreators)
`maarten.stol@braincreators.com`.
The authors would like to thank the team at Inspech (inspech.com) for their support and collaboration.

References

1. Ashraf, A., Sophian, A., Bawono, A.A.: Crack detection, classification, and segmentation on road pavement material using multi-scale feature aggregation and transformer-based attention mechanisms. Construction Materials **4**(4), 655–675 (Oct 2024). https://doi.org/https://doi.org/10.3390/constrmater4040036, https://www.mdpi.com/2673-7108/4/4/36
2. Baduge, S.K., Thilakarathna, S., Perera, J.S., Ruwanpathirana, G.P., Doyle, L., Duckett, M., Lee, J., Saenda, J., Mendis, P.: Assessment of crack severity of asphalt pavements using deep learning algorithms and geospatial system. Construction and Building Materials **401**, 132684–132684 (Aug 2023). https://doi.org/https://doi.org/10.1016/j.conbuildmat.2023.132684, https://www.sciencedirect.com/science/article/pii/S0950061823024005
3. Estilong, J., Palaoag, T.: Literature review on road damage detection and severity recognition: Leveraging computer vision. Journal of Information Systems Engineering and Management **10**(5s), 498–512 (Jan 2025). https://doi.org/https://doi.org/10.52783/jisem.v10i5s.670, https://www.jisem-journal.com/index.php/journal/article/view/670
4. Huang, S., Chen, H., Yan, L., Zou, X., Li, B., Bi, Y.: A review of the progress in machine vision-based crack detection and identification technology for asphalt pavements. Digital Transportation and Safety **4**(1), 65–79 (Jan 2025). https://doi.org/https://doi.org/10.48130/dts-0025-0006
5. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv:1609.02907 [cs, stat] (Feb 2017), https://arxiv.org/abs/1609.02907
6. Kraskov, A., Stögbauer, H., Grassberger, P.: Estimating mutual information. Physical Review E **69**(6) (Jun 2004). https://doi.org/https://doi.org/10.1103/physreve.69.066138, https://journals.aps.org/pre/abstract/10.1103/PhysRevE.69.066138
7. Liu, Y., Wenxuan, T., Zhou, S., Zhu, E.: Deep graph clustering via dual correlation reduction (Dec 2021)
8. Lu, S.H., Lin, J.D., Chiou, Y.S.: Analysis and verification results of manual inspection of pavement condition index. Applied Sciences **15**(10), 5600 (May 2025). https://doi.org/https://doi.org/10.3390/app15105600, https://www.mdpi.com/2076-3417/15/10/5600
9. Qian, L., Toda, H., Uesugi, K., Kobayashi, M., Kobayashi, T.: Direct observation and image-based simulation of three-dimensional tortuous crack evolution inside opaque materials. Physical Review Letters **100**(11) (Mar 2008). https://doi.org/https://doi.org/10.1103/physrevlett.100.115505

10. Qureshi, W.S., Hassan, S.I., McKeever, S., Power, D., Mulry, B., Feighan, K., O'Sullivan, D.: An exploration of recent intelligent image analysis techniques for visual pavement surface condition assessment. Sensors **22**(22), 9019–9019 (Nov 2022). https://doi.org/https://doi.org/10.3390/s22229019
11. Sun, J., Chen, R., Li, J., Ding, Y., Wu, C., Liu, Z., Yan, J.: Understanding and mitigating dimensional collapse of graph contrastive learning: A non-maximum removal approach. Neural Networks **181**, 106652 (Jan 2025). https://doi.org/https://doi.org/10.1016/j.neunet.2024.106652
12. Whiteley, N., Gray, A., Rubin-Delanchy, P.: Statistical exploration of the manifold hypothesis (2022), https://arxiv.org/abs/2208.11665
13. Yuan, Q., Shi, Y., Li, M.: A review of computer vision-based crack detection methods in civil infrastructure: Progress and challenges. Remote Sensing **16**(16), 2910–2910 (Aug 2024). https://doi.org/https://doi.org/10.3390/rs16162910, https://www.mdpi.com/2072-4292/16/16/2910