3D Microstructure Reconstruction of X70 Pipeline Steel from 2D EBSD Data Using SliceGAN

Joonas Lahikainen [1] **Abhishek Biswas** [1] **Supriya Nandy** [1] **Napat Vajragupta** [1]

[1] *VTT Technical Research Centre of Finland, Espoo, Finland.* Correspondence to: Napat Vajragupta napat.vajragupta@vtt.fi.

1. Introduction

Polycrystalline metals are composed of many differently oriented grains that collectively create a strong microstructure. Understanding and characterising these microstructures has been central to major advances in materials science, enabling, for example, the design of metals that simultaneously achieve high strength and ductility [1]. As microstructure remains a key determinant of performance in modern engineering alloys, its study continues to be a critical area of research.

Grain structures can be analysed using a range of microscopy methods, with approaches such as light optical microscopy (LOM), scanning electron microscopy (SEM), and transmission electron microscopy (TEM) providing complementary views of the microstructure. In electron microscopes, crystallographic orientation can be mapped using techniques such as electron backscatter diffraction (EBSD) in the SEM or precession electron diffraction (PED) in the TEM, though these measurements are typically restricted to two dimensions. Three-dimensional characterisation is possible but considerably more demanding: FIB-EBSD [2] reconstructs 3D volumes through serial ion-beam milling but can require more than 24 hours for a single volume [3], while TEM-based approaches such as 3D-OMiTEM [4], which builds orientation maps from tilt-series electron diffraction, involve demanding sample preparation. Despite these challenges, the value of three-dimensional data cannot be overstated: it offers a more comprehensive view of the microstructure, including access to true grain shapes, boundary topologies, and the spatial distribution of plastic strain and dislocations [3].

A practical alternative to direct 3D characterisation is to approximate 3D structures from 2D observations. Much of the morphological and crystallographic information is already present in 2D sections, but extrapolating these features into a consistent third dimension remains a fundamental challenge, one that is inherently statistical, since a single 2D slice is compatible with many plausible 3D configurations, even under assumptions of isotropy. A wide range of statistical and computational reconstruction methods have been explored to address this, surveyed comprehensively in [5], but machine-learning approaches offer a complementary next step: they can learn these statistical relationships directly from data, with minimal assumptions.

One prominent method in this direction is SliceGAN [6], a framework based on generative adversarial networks (GANs) that synthesises 3D microstructures from 2D images. SliceGAN has been shown to produce physically meaningful reconstructions across a diverse range of microstructures [7, 8] and has since been extended to handle larger volumes [9] and conditional controls over generated features [10].

In this work, we investigate whether SliceGAN can support 3D reconstructions from routine EBSD scans. Our data consist of a FIB–EBSD volume of X70 pipeline steel, which we aim to reconstruct using only a small number of 2D input slices. The central question is whether a generative model can recover realistic grain morphologies from limited observations while remaining consistent with the underlying crystallography. Addressing SliceGAN's current limitations with complex grain geometries [11] will require deeper integration with orientation data and architectural modifications. To this end, we experiment with a range of orientation representations, including quaternions, axis-angle, continuous representations [12], and inverse pole figure (IPF) maps, as well as quaternion neural networks (QNNs) and modified SliceGAN architectures. The results presented here constitute an initial feasibility study of SliceGAN applied to noisy EBSD maps, and surface the practical challenges that remain.

2. Related Work

A natural approach to addressing SliceGAN's limitations is to combine it with complementary methods. A 2024 study [11] paired SliceGAN with grain packing via DREAM.3D to reconstruct 3D EBSD maps of cold spray-formed aluminium alloys with a bimodal grain structure, exploiting the complementary strengths of each: SliceGAN's capacity to generate diverse morphologies and DREAM.3D's ability to match experimental statistics. Neither method alone produced convincing results, and the hybrid achieved good statistical agreement with the real microstructure. However, validation relied on 2D slice statistics rather than direct 3D comparison, and performance on smaller input datasets remains an open question particularly relevant to our setting.

One limitation of applying SliceGAN directly to orientation data is that quaternions do not live in a Euclidean space, making standard linear convolutions a poor fit for the underlying geometry. QNNs [13] address this by replacing classical operations with their quaternion algebra counterparts, including quaternion-valued convolutions based on the Hamilton product. Beyond the geometric compatibility, bundling features into four-dimensional quaternion

vectors has been shown to encode inter-channel relationships more efficiently than real-valued networks, often achieving comparable or better performance with up to 75% fewer training parameters [14]. The paradigm has been demonstrated in domains such as speech recognition [13] and, more relevant to this work, image generation with GANs [14]. The noisy grain boundaries observed in SliceGAN's quaternion outputs [11] motivate exploring architectures with native quaternion support.

A related concern is the choice of orientation format. Although quaternions are widely used for 3D orientations, they are discontinuous in the mathematical sense. The authors of [12] prove that no 3D rotation representation in four or fewer dimensions can be continuous, making quaternions and Euler angles fundamentally difficult for neural networks to learn. They propose alternative 5D and 6D formats and demonstrate empirically that these outperform discontinuous ones on several learning tasks. We include these as part of our investigation into more principled orientation encoding for microstructure generation.

3. Materials And Methods

Building on these approaches, we use an $80 \times 400 \times 105$ voxel FIB-EBSD volume of X70 pipeline steel obtained from VTT, shown in Fig. 1(a). Our goal is to reconstruct a comparable EBSD volume using only a subset of this data. For anisotropic materials such as X70 steel, at least three orthogonal EBSD slices are typically required to capture the characteristic features of each plane, and in favourable cases these slices can be sufficient on their own. The number of slices required ultimately depends on their representativeness of the full volume and the level of measurement noise.

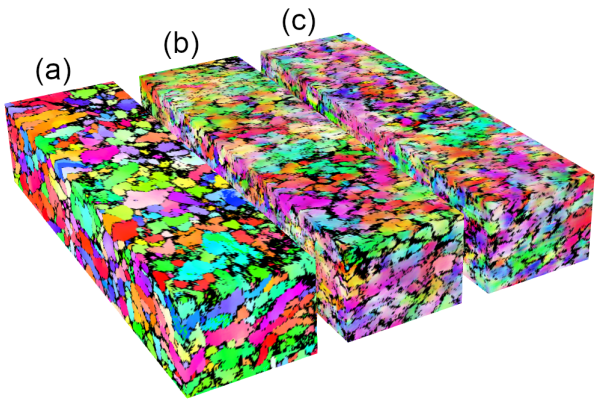

Fig. 1: Inverse pole figure (IPF-Z) views of (a) the raw FIB–EBSD volume of X70 pipeline steel; (b) a SliceGAN reconstruction from six IPF-Z slices per spatial direction; and (c) a SliceGAN reconstruction using the same six slices with additional IPF-Y input. Black represents missing sensor data.

We currently focus on three orientation representations: quaternions, the continuous 6D representation from [12], and IPF maps encoded as RGB images. SliceGAN also supports discretised label inputs, which can produce visually cleaner outputs [6, 11] but at the cost of orientation continuity, making them unsuitable for practical use. For quaternions, we use the original SliceGAN architecture with standard modules replaced by their quaternion counterparts [13]. The 6D and IPF representations use the original Wasserstein GAN with support for both gradient penalty (GP) and spectral norm regularisation. Microstructure-informed loss terms will be considered once satisfactory visual quality has been achieved.

We assume cubic crystal symmetry throughout. An important caveat of IPF maps is that, except in the triclinic case, the mapping from orientation to colour is non-unique [15]: it is not possible to recover the full orientation from IPF-X/Y/Z colours alone. Despite this, we observe that SliceGAN produces its best reconstructions when trained on IPF maps (see Fig. 1). To exploit this, we plan to encode orientations reduced to the cubic fundamental zone using axis-angle representation: the rotation axis mapped to a colour using the triclinic IPF colour scheme, and the rotation angle retained as a fourth channel. This would give a representation that is both visually structured and invertible.

4. Experiments

QGAN. Our initial QNN-based results exhibit limited reconstruction fidelity. The generator fails to produce convincing grain boundaries, allowing the discriminator to become too strong too quickly and leaving the generator without a useful learning signal. Switching to spectral norm regularisation in place of GP leads to mode collapse, resulting in overly simple geometries. In both cases, output quality deteriorates as training proceeds, and further work on model architecture and hyperparameters is needed.

Vanilla SliceGAN. The IPF representation produces the cleanest outputs so far. The 6D continuous representation, despite its theoretical advantages for orientation learning, suffers from similar training instability as the QNN models, suggesting that pure orientation data may be too challenging for the generators in their current forms. Combining quaternion channels with IPF maps does not yield any improvement either. As visible in Fig. 1(b) and (c), the outputs are blurry, a known limitation of GANs, which makes grain boundary identification more difficult. This may be partially alleviated by increasing the patch size, which is currently constrained to 64×64 by the size of the dataset.

Training behaviour. Both the GP and Wasserstein losses drop rapidly in early epochs before rising steadily for the remainder of training, with the generator and discriminator losses diverging considerably. Output quality does not recover and continues to worsen, suggesting that neither model has yet achieved a stable training regime on this dataset.

References

[1] J. Shen, B. Chen, X. Ye, H. Imai, J. Umeda, and K. Kondoh. The formation of bimodal multilayered grain structure and its effect on the mechanical properties of powder metallurgy pure titanium. *Materials & Design*, 116:99–108, 2017.

[2] Z. Liu, S. Bai, S. Burke, J. N. Burrow, R. Geurts, C.-J. Huang, C. Jiao, H.-B. Lee, Y. S. Meng, L. Novák, B. Winiarski, J. Wang, K. Wu, and M. Zhang. Fib-sem: Emerging multimodal/multiscale characterization techniques for advanced battery development. *Chemical Reviews*, 125(11):5228–5281, 2025.

[3] Oxford Instruments NanoAnalysis. 3d electron backscatter diffraction (3d ebsd). `https://www.ebsd.com/ebsd-techniques/3d-ebsd`, 2026. Accessed: 2026-01-13.

[4] H. H. Liu, S. Schmidt, H. F. Poulsen, A. Godfrey, Z. Q. Liu, J. A. Sharon, and X. Huang. Three-dimensional orientation mapping in the transmission electron microscope. *Science*, 332(6031):833–834, 2011.

[5] Ramin Bostanabad, Yichi Zhang, Xiaolin Li, Tucker Kearney, L. Catherine Brinson, Daniel W. Apley, Wing Kam Liu, and Wei Chen. Computational microstructure characterization and reconstruction: Review of the state-of-the-art techniques. *Progress in Materials Science*, 95:1–41, 2018.

[6] Steve Kench and Samuel J. Cooper. Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion. *Nature Machine Intelligence*, 3(4):299–305, April 2021.

[7] Y. Zhou, T. Zhao, W. Zhang, F. Teng, and X. Nie. Digital core reconstruction of tight carbonate rocks based on slicegan. *Artificial Intelligence in Geosciences*, 6(1):100116, 2025.

[8] Y. Ghazal, M. Awadalla, D. Barradas, A. Ayyad, A. Nasr, and S. Ghani. Optimizing gan training for 3d seismic microstructure generation. In *Eighth EAGE High Performance Computing Workshop*, pages 1–3. European Association of Geoscientists & Engineers, 2024.

[9] Keiya Sugiura, Toshio Ogawa, Yoshitaka Adachi, Fei Sun, Asuka Suzuki, Akinori Yamanaka, Nobuo Nakada, Takuya Ishimoto, Takayoshi Nakano, and Yuichiro Koizumi. Big-Volume SliceGAN for Improving a Synthetic 3D Microstructure Image of Additive-Manufactured TYPE 316L Steel. *Journal of Imaging*, 9(5), 2023.

[10] Hyungjin Chung and Jong Chul Ye. Reusability report: Feature disentanglement in generating a three-dimensional structure from a two-dimensional slice with sliceGAN. *Nature Machine Intelligence*, 3(10):861–863, October 2021.

[11] B. Murgas, J. Stickel, and S. Ghosh. Generative adversarial network (gan) enabled statistically equivalent virtual microstructures (sevm) for modeling cold spray formed bimodal polycrystals. *npj Computational Materials*, 10(1):32, 2024.

[12] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[13] Titouan Parcollet, Mirco Ravanelli, Mohamed Morchid, Georges Linarès, Chiheb Trabelsi, Renato De Mori, and Yoshua Bengio. Quaternion recurrent neural networks. In *International Conference on Learning Representations*, 2019.

[14] Eleonora Grassucci, Edoardo Cicero, and Danilo Comminiello. Quaternion generative adversarial networks. In Roozbeh Razavi-Far, Ariel Ruiz-Garcia, Vasile Palade, and Juergen Schmidhuber, editors, *Generative Adversarial Learning: Architectures and Applications*, volume 221 of *Intelligent Systems Reference Library*, pages 57–86. Springer Cham, 2022.

[15] G. Nolze and R. Hielscher. Orientations – perfectly colored. *Journal of Applied Crystallography*, 49(5):1786–1802, October 2016.