Graph-UNet for Materials Image Segmentation Postprint

Authors: Wei Huishan, Han Yuexing, Wang Bing, Chen Qiaochuan

Date: 2022-06-06T00:00:00+00:00

Abstract

Few-shot material image segmentation constitutes one of the challenging research problems in the field of image segmentation. The microstructures of material images typically exhibit characteristics such as diverse shapes, complex textures, and blurred boundaries, which often lead to inaccurate segmentation results. To address the challenge of automatic few-shot material image segmentation, Graph-UNet has been proposed to integrate UNet and graph convolutional neural networks. This method transfers the concepts of multi-dimensional feature fusion and skip connections from convolutional neural networks to graph convolutional neural networks, thereby achieving an effective combination of graph convolution and graph attention. Furthermore, it establishes a universal module for mutual conversion between feature maps and graph structures. Comparative and ablation experiments conducted on material image datasets demonstrate that Graph-UNet achieves superior segmentation results compared to many state-of-the-art methods, accurately identifies various material structures, and advances the exploration of the relationship between material structure and performance.

Full Text

Preamble

Graph-UNet for Material Image Segmentation

Wei Huishan[1], Han Yuexing[1,2]†, Wang Bing[1], Chen Qiaochuan[1]
(1. School of Computer Engineering & Science, Shanghai University, Shanghai 200444, China;
2. Zhejiang Laboratory, Hangzhou 311100, China)

Abstract: Few-shot material image segmentation represents a significant challenge in the field of image segmentation. The microstructures in material images typically exhibit diverse shapes, complex textures, and blurred boundaries, which often lead to inaccurate segmentation results. To address the challenge of automatic segmentation for few-shot material images, we propose Graph-UNet, which integrates UNet and graph convolutional neural networks. Our approach transfers the multi-dimensional feature fusion and skip connection concepts from convolutional neural networks to graph convolutional neural networks, enabling effective combination of graph convolution and graph attention mechanisms. Additionally, we establish a universal module for mutual conversion between feature maps and graph structures. Through comparative and ablation experiments on material image datasets, we demonstrate that Graph-UNet achieves superior segmentation performance compared to many state-of-the-art methods, accurately identifying various material structures and advancing the exploration of structure-property relationships in materials.

Keywords: semantic segmentation; graph convolutional neural network; graph attention; material image; deep learning; few-shot

0 Introduction

With continuous advancements in image segmentation methodologies and increasing interdisciplinary integration, researchers have begun applying sophisticated image segmentation techniques to material image analysis, as evidenced in studies [?, ?]. Accurate segmentation of material images facilitates investigations into the relationships between manufacturing processes, microstructures, and material properties, thereby contributing to the development of novel materials. Image segmentation of material structures serves as the foundation for subsequent image processing tasks, and precise delineation of target regions is crucial for structural analysis and performance evaluation.

Most material images contain intricate textures and microstructures of various shapes. Even within the same material category, structures can exhibit substantial variations in shape and texture. The diversity of material structures makes annotation extremely labor-intensive and time-consuming. Consequently, enhancing model generalization and addressing few-shot image segmentation problems are essential. Distinguishing different material structures in small datasets remains a formidable task.

Several deep learning approaches have been developed for material image segmentation. Azimi et al. [?] proposed MVFCNN, a segmentation algorithm for material microstructures based on Fully Convolutional Networks (FCN). However, this method requires numerous material samples, and the high acquisition cost of most material images limits the application of deep learning in this domain. Jiang et al. [?] introduced a deep learning method that performs pixel-level material segmentation on entire images, combining dilated convolution features with traditional convolution features to remove artifacts caused by dilated convolutions. Xiong et al. [?] developed a rice panicle segmentation algorithm based on Simple Linear Iterative Clustering (SLIC) superpixel generation and convolutional neural network classification. Nevertheless, rice panicles exhibit distinct boundaries and regional features, making the superpixel approach task-specific and limited in generalizability. Decost et al. [?] proposed a deep convolutional neural network capable of extracting distributions of water lime particle size and metamorphic zone widths from micrographs containing multiple microscopic components. However, their approach primarily targets specific application scenarios.

To address end-to-end material image segmentation, this paper proposes a universal network model, Graph-UNet, which combines UNet and graph convolutional neural networks, as illustrated in [Figure 1: see original paper]. We utilize deep convolutional networks to model graphs and apply graph convolution methods to solve semantic segmentation tasks for material images. The model constructs a graph structure from feature maps processed by convolutional neural networks, employs a Multi-dimensional Graph Attention Module (MGAM) to exchange features between neighboring nodes, and then encapsulates the updated graph structure back into feature maps for subsequent convolutional operations. This approach makes two primary contributions: (a) we propose a novel graph attention module that integrates graph convolution and graph attention for multi-dimensional feature fusion; (b) we develop a graph model for material images that encapsulates images into graph structures for participation in graph convolution operations.

The remainder of this paper is organized as follows: Section 1 introduces the network architecture and specific details. Section 2 presents experimental results and discussions on material datasets. Section 3 concludes the paper and outlines future research directions.

1 Graph Model for Semantic Segmentation

Graph Convolutional Neural Networks (GCNNs) are designed to address learning problems on graph-structured data. Recently, researchers have begun exploring how graph convolution can be applied to structured data. Lu et al. [?] first combined graph convolution with fully convolutional networks for semantic segmentation, achieving a 1.34% performance improvement. Zhang et al. [?] proposed the Dual Graph Convolutional Network (DGCNet), integrating graph convolution for irregular data into a dual attention model for semantic segmentation tasks. Ma et al. [?] introduced a novel Attention Graph Convolutional Network (AGCN) for superpixel segmentation of remote sensing images, utilizing graph convolution combined with graph attention for image segmentation.

Currently, few graph convolution-based network models target material image segmentation. Drawing inspiration from previous research, the graph structure in this work features two key characteristics. First, graph structures enable more flexible skip connections. Unlike convolutional networks that establish an n×n local region around a pixel, graph convolution allows custom-defined neighborhoods for each node. Second, the message propagation mechanism in graph convolution facilitates information exchange between nodes, and as graph convolution layers are stacked, each node' s receptive field expands, enabling comprehensive capture of broader contextual features.

To leverage the advantages of attention mechanisms in image segmentation tasks and enhance network focus on important features, we incorporate a multi-dimensional feature fusion-based graph attention module into UNet [?] to address few-shot material image segmentation.

We selected UNet as the backbone network because it was originally developed for medical images, which share significant similarities with material images, such as monotonic colors and predominant use of grayscale. Notable UNet variants include UNet++ [?] and Attention-UNet [?]. UNet++ enhances feature supplementation through additional skip connections, while Attention-UNet improves network performance by incorporating attention mechanisms in the decoder. Inspired by the application of graph convolution and attention mechanisms in image segmentation, we propose the graph attention module.

In our network, the input image first passes through five convolutional layers, with pooling layers added after the first four. Following the fifth convolutional layer, we convert the extracted feature maps into graph structures and employ the multi-dimensional feature fusion-based graph attention module for node feature propagation and learning. Subsequently, the module converts the graph structure back into feature maps for the following four upsampling layers, progressively enlarging the feature map dimensions. After each upsampling operation, the current feature map is concatenated with the encoder feature map of the same size. Finally, the network outputs a prediction map through a Sigmoid function to obtain the segmentation result.

1.1 Graph Structure

A graph structure is an irregular data structure that can be represented as a triplet $G(N, E, U)$. N denotes the node set, which is an $|N| \times S$ matrix where $|N|$ is the number of nodes and S is the dimension of node feature vectors. E represents the edge set, and U denotes graph features.

We construct the graph using a directed approach, assigning directionality to each edge. To ensure no discriminative difference in connectivity between nodes after graph construction, we initialize each node's total degree to a specific value using a four-neighborhood scheme. The model is trained through supervised learning.

1.2.1 Graph Encoder

The input to the graph attention module is a graph structure. To feed feature maps into the graph attention module, we create a graph encoder that converts feature maps into graph structures.

In our model, node features are initialized by the encoder portion of UNet. In the graph model, connections between nodes are represented using an adjacency matrix, with each node connected to its nearest neighbors. We establish edge connections in a four-neighborhood pattern based on the spatial relationships between pixel points. These connections indicate that features can propagate through edges within the graph attention module.

1.2.2 Graph Attention Module

We propose a Multi-dimensional Graph Attention Module (MGAM), as illustrated in [Figure 2: see original paper].

Graph Convolutional Networks (GCN) represent one deep learning approach for processing graph structures. Its implementation formula [?] is as follows:

$$Z^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}X^{(l)}W^{(l)})$$

where X is the graph structure converted from UNet encoder output, D and A are the degree and adjacency matrices of X, $\tilde{}$ indicates that self-connections are added to each node with corresponding degree matrix $\tilde{D}$ and adjacency matrix $\tilde{A}$, $W^{(l)}$ is the weight matrix for the l-th graph convolution layer, σ is the ReLU activation function, and $Z^{(l+1)}$ is the output of the $(l+1)$-th graph convolution operation.

Graph Attention Networks (GAT) represent an attention mechanism applied to graph structures. Its implementation formula [?] is as follows:

$$h_i' = \sigma\left(\frac{1}{K}\sum_{k=1}^{K}\sum_{j\in\mathcal{N}i}\alpha_{ij}^k W^k h_j\right)$$

where K denotes the number of attention heads, $\mathcal{N}_i$ is the set of neighbor nodes of node i, α_{ij}^k is the attention coefficient between node i and its j-th neighbor in the k-th attention head, W^k is the weight matrix for the k-th attention head, h_j is the feature vector of node j, σ is the activation function, and h_i' is the aggregated output feature of node i after the multi-head graph attention layer. The specific formula for α_{ij}^k is:

$$\alpha_{ij}^k = \frac{\exp(\mathrm{LeakyReLU}(a^T[W^k h_i||W^k h_j]))}{\sum_{r\in\mathcal{N}i}\exp(\mathrm{LeakyReLU}(a^T[W^k h_i||W^k h_r]))}$$

where LeakyReLU() is a specific activation function, a and W are weight matrices, and h_i, h_j, and h_r are the feature vectors of node i, node j, and the r-th neighbor of node i, respectively.

Both GCN and GAT aggregate neighbor node features to the central node: GCN utilizes the Laplacian matrix, while GAT employs attention coefficients. However, GCN cannot assign different weights to each neighbor, treating all neighbors equally during convolution without considering node importance. In contrast, GAT effectively incorporates correlations between node features into the model and uses multi-head attention to stabilize the learning process, as shown in [Figure 3: see original paper].

Generally, GCN for node classification or segmentation employs one to two layers; stacking too many GCN layers leads to over-smoothing. However, in convolutional neural network models, deeper stacks of convolutional layers typically demonstrate better performance and effectiveness. Deeper models provide stronger non-linear expressive capabilities, enabling learning of more complex transformations to fit more sophisticated feature inputs, as exemplified by residual networks. To incorporate this CNN philosophy, we deepen the graph convolution model by introducing a hyperparameter a that controls the proportion of node features flowing into the next graph convolution layer. This approach simulates skip connections in residual networks but only partially mitigates over-smoothing. Performance still degrades when too many layers are stacked.

We utilize two GCN layers to simultaneously propagate node features and perform dimensionality reduction, followed by two GCN layers for feature dimensionality expansion while summing the dimensionality-reduced node features at each step. Initially, the input graph structure contains $H \times W$ nodes, each with feature dimension C. After the first graph convolution layer, node feature dimensions are reduced to $C/2$. The graph with updated node features is then fed into the second graph convolution layer, reducing dimensions to $C/4$. Subsequently, we perform dimensionality expansion by passing through two additional graph convolution layers. The dimensionality-reduced and expanded features correspond one-to-one, with the reduced-dimension node features connected to a GAT layer and then summed with the expanded graph convolution output. Graph structures with dimensions C, $C/2$, and $C/4$ correspond to high-resolution, medium-resolution, and low-resolution feature maps in the image domain. By learning coarse features from low-resolution graph structures and fine details from high-resolution graph structures, fusing different resolution graph structures enables accurate learning and representation of few-shot material images. After the graph attention layer, we set a hyperparameter a to control the proportion of node features flowing to the next graph convolution layer, preventing node features from converging to nearly global features after excessive graph convolution layers. The specific fusion implementation for the feature expansion portion is shown in Equation (4):

$$Z_{add}^{(l+1)} = Z^{(l+1)} + a \cdot H_{add}$$

where $Z_{add}^{(l+1)}$ is the summation result of the graph attention layer output and the graph convolution layer expansion output, $Z^{(l+1)}$ is the output feature of the $(l+1)$-th graph convolution layer, a is the hyperparameter, and H_{add} is the output of the graph attention layer.

The graph attention module constitutes a crucial component of the network model, representing a combination of graph convolution and graph attention. It primarily serves to deepen the network, enhance its ability to fit complex input images, alleviate excessive focus on local features, prevent neglect of global feature information, and achieve superior material image segmentation performance. Graph convolution preserves the spatial location information of feature map pixels without node disappearance, while the graph attention layer focuses on important nodes, enhancing feature flow for significant nodes and suppressing flow for irrelevant ones.

1.2.3 Graph Decoder

The output of the graph attention module remains a graph structure. To enable continued processing in convolutional neural networks, we create a graph decoder that converts the graph structure back into feature maps.

2.1 Experimental Environment and Evaluation Metrics

In our experiments, we used a computer with an i7-8750H CPU@2.20GHz processor and Python 3.8 as the development environment. The network was implemented using PyTorch. We standardized all images to 512×512 pixels as network input and converted label images to binary format for training. We employed binary cross-entropy loss as the network' s loss function and trained the model using the Adam optimizer with momentum. The momentum and decay coefficients were set to 0.9 and 0.0005, respectively, to mitigate overfitting. We trained for 100 epochs with a batch size of 1 and an initial learning rate of 0.001, adjusting the learning rate to approach optimal network parameters. Each material structure dataset contained two classes: target material structure and background. We trained on the training dataset, minimizing the loss between input and label images, and saved the optimal network parameters across 100 iterations for testing. Network performance was evaluated using Mean Intersection over Union (MIoU).

We trained the network on the training dataset, saved the optimal parameters, and then performed predictions on the test dataset, calculating MIoU between test predictions and corresponding annotations as the performance metric. Finally, we converted network outputs into visualized binary annotation results.

Our experimental dataset comprises four material structure categories: Spheroidite, Wood, Pearlite1, and Pearlite2. Spheroidite, Pearlite1, and Pearlite2 are open-source material datasets from literature [**?**]. Due to the complexity of experiments and high annotation costs, each category contains only 4 to 6 images. The dataset distribution is shown in .

We primarily use MIoU as the evaluation metric for semantic segmentation networks, calculated as follows:

$$\text{MIoU} = \frac{1}{k}\sum_{i=1}^{k}\frac{p_{ii}}{\sum_{j=1}^{k} p_{ij} + \sum_{j=1}^{k} p_{ji} - p_{ii}}$$

where p_{ii} represents the number of pixels with true value i predicted as i, p_{ij} represents pixels with true value i predicted as j, p_{ji} represents pixels with true value j predicted as i, and k is the number of classes. MIoU is calculated on a per-class basis, computing IoU for each class, summing them, and averaging to obtain a global evaluation metric.

2.2 Comparative Experiments with Existing Methods

To further demonstrate the effectiveness of our proposed model for material image segmentation, we compared it against current state-of-the-art methods on the same dataset. We evaluated traditional segmentation methods including KMeans [?] and Watershed [?], as well as deep learning approaches FCN [?], UNet [?], UNet++ [?], CENet [?], R2U-Net [?], Attention-UNet [?], and UNet3+ [?]. Experimental results are shown in [Figure 4: see original paper] and .

KMeans is overly sensitive to noise in complex texture images, causing most textural details to be segmented out. Consequently, KMeans cannot obtain complete material structures, leading to over-segmentation of phases. Watershed produces inconsistent segmentation results; as shown in figures 5, 6, and 7 of the Watershed results in [Figure 4: see original paper], it barely segments target structures and even produces large black areas as erroneous results. For grayscale images with low contrast, Watershed struggles to achieve satisfactory segmentation. Both KMeans and Watershed lack generality for image segmentation tasks.

Due to limited samples, deep learning-based networks cannot obtain rich training data, resulting in insufficient feature learning for material structures. FCN results exhibit varying degrees of noise. As shown in figure 3 of the FCN results in [Figure 4: see original paper], numerous small but noticeable noise points appear near boundaries, indicating inadequate feature learning in boundary regions. In figure 7 of the FCN results, blocky material structures show obvious color discrepancies, with centers tending toward black and edges appearing white. For such structures, FCN fails to segment blocks as unified regions, creating many non-existent boundaries. Clearly, FCN does not effectively incorporate contextual information, leading to segmentation of many small regions.

UNet, UNet++, and UNet3+ produce good segmentation results, but UNet3+ exhibits slightly inferior visualization quality with some noise, suggesting that excessive skip connections are not universally beneficial for material datasets. CENet demonstrates stable evaluation metrics and segmentation maps but still contains small segmentation errors, as shown in figures 2 and 7 of the CENet results. R2U-Net exhibits local segmentation errors, while Attention-UNet produces inconsistent results. Detailed statistics in show that our method achieves superior MIoU performance, generally exceeding existing methods by 2-5 percentage points.

Our network model demonstrates stable performance across multiple datasets. Most noise points in material structures are effectively filtered, large connected regions are accurately segmented, and narrow boundaries between material structures are segmented better than most methods. Even with low color contrast, the method exhibits strong learning capability, achieving a good balance between classification accuracy and localization precision.

2.3 Validation Experiments on Graph Attention Module Effectiveness

We investigated network performance with different backbones to validate the effectiveness of the graph attention module. We compared models with and without the graph attention module using ResNet34-UNet, ResNet34-FCN, and UNet++ as backbones, finding that the graph attention module positively impacts network performance, as shown in .

To demonstrate the positive effect of our proposed graph attention module on material image segmentation, we compared it with classical attention methods: SE-Net [?], CBAM [?], DANet [?], CCNet [?], and PSANet [?], as shown in . We added each of these five attention methods as separate modules to UNet and compared them with our graph attention module across four datasets. Experiments confirm that our proposed module has a positive impact on material image segmentation.

2.4 Ablation Experiments

In this work, we primarily propose a novel multi-dimensional feature fusion-based graph attention module. To understand the impact of graph convolution, graph attention, and hyperparameter a on the network model, we conducted detailed ablation experiments with five configurations: (1) UNet with one graph convolution layer, (2) UNet with one graph attention layer, (3) UNet with a multi-dimensional graph convolution module, (4) UNet with a multi-dimensional graph attention module without hyperparameter a, and (5) our complete network model. We compared the segmentation results of these five configurations on material datasets, with different graph convolution module structures shown in [Figure 5: see original paper].

Ablation experiment results are presented in , with numbers reported as percentage segmentation accuracy. The table reveals that our proposed multi-dimensional feature fusion-based graph attention module positively influences model performance, with each component contributing to segmentation improvements to varying degrees.

Visualization results in [Figure 6: see original paper] show that each network model produces relatively accurate segmentation results for material structures, making differences difficult to discern visually. However, upon examining boundary segmentation, particularly in figures 5 and 6 of each experimental group, our proposed module demonstrates positive effects on boundary identification. This confirms that the multi-dimensional feature fusion-based graph attention module enhances segmentation of fine details.

3 Conclusion

In this work, we extracted a semantic segmentation method for few-shot material structure images. By incorporating a multi-dimensional feature fusion-based graph attention module into UNet, our approach obtains more feature information under limited sample conditions to achieve automatic segmentation of material images. The method has certain limitations requiring improvement. For instance, it is relatively computationally expensive. Additionally, boundary details with indistinct material structures remain challenging to segment effectively.

For future research, we will further strengthen our method by leveraging the advantages of graph convolution, employing deeper graph convolutional neural networks to enhance model fitting capabilities, and applying the approach to more novel material structures.

References

[1] He Dong, Tang Ting. Image segmentation of carbon fiber wood composite based on fuzzy C-means clustering algorithm [J]. Synthetic Materials Aging and Application, 2020 (4): 78-80.

[2] Zhang Lixin, Che Shijie, Xu Zhengguang, et al. Image segmentation method of superalloy microstructure based on residual network [J]. Science Technology and Engineering, 2020, 20 (1): 246-251.

[3] Azimi S, Mücklich F, Fritz M. Advanced steel microstructural classification by deep learning methods [J]. ArXiv: 1706. 06480, 2018.

[4] Jiang X, Du J, Sun B, et al. Deep dilated convolutional network for material recognition [C]. 2018 Eighth International Conference on Image Processing Theory, Tools and Applications (IPTA), 2018. http://dx.doi.org/10.1109/IPTA.2018.8608160.

[5] Xiong X, Duan L, Liu L, et al. Panicle-SEG: A robust image segmentation method for rice panicles in the field based on deep learning and superpixel optimization [J]. Plant Methods, 2017, 13 (1): 104-118.

[6] Decost B L, Lei B, Francis T, et al. High throughput quantitative metallography for complex microstructures using deep learning: A case study in ultrahigh carbon steel [J]. ArXiv: 1805. 08693, 2018.

[7] Lu Y, Chen Y, Zhao D, et al. Graph-FCN for image semantic segmentation [J].

Advances in Neural Networks, 2019. https://doi.org/10.1007/978-3-030-22796-8_{11}.

[8] Huang Z, Wang X, Huang L, et al. CCNet: Criss-cross attention for semantic segmentation [C]. International Conference on Computer Vision, 2019. http://dx.doi.org/10.1109/ICCV.2019.00069.

[9] Ma F, Gao F, Sun J, et al. Attention graph convolution network for image segmentation in big SAR imagery data [J]. Remote Sensing, 2019, 11 (21): 2586-2606.

[10] Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation [J]. ArXiv: 1505. 04597, 2015.

[11] Zhou Z, Siddiquee M, Tajbakhsh N, et al. UNet+: A nested U-Net architecture for medical image segmentation [J]. ArXiv: 1807. 10165, 2018.

[12] Oktay O, Schlemper J, Folgoc L L, et al. Attention U-Net: Learning where to look for the pancreas [J]. ArXiv: 1804. 03999, 2018.

[13] Kip F, M. Welling. Semi-supervised classification with graph convolutional networks. ArXiv: 1710. 10903, 2016.

[14] Velikovi P, Cucurull G, Casanova A, et al. Graph attention networks [J]. ArXiv: 1710. 10903, 2017.

[15] Decost B L, Hecht M, Francis T, et al. UHCSDB: UltraHigh carbon steel micrograph database Integrating Materials&Manufacturing Innovation, 2017, 6 (8): 1-9.

[16] Macqueen J. Some methods for classification and analysis of multiVariate observations [C]. In 5-th Berkeley Symposium on Mathematical Statistics and Probability, 1965, 66 (1): 281-297.

[17] Bieniek A, Moga A. An efficient watershed algorithm based on connected components [J]. Pattern Recognition, 2000, 33 (6): 907-916.

[18] Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2015, 39 (4): 640–651.

[19] Gu Z, Cheng J, Fu H, et al. CE-Net: Context encoder network for 2D medical image segmentation [J]. IEEE Transactions on Medical Imaging, 2019, 38 (10): 2281-2292.

[20] Alom M Z, Hasan M, Yakopcic C, et al. Recurrent residual convolutional neural network based on U-Net (R2U-Net) for medical image segmentation [J]. ArXiv: 1802. 06955, 2018.

[21] Huang H, Lin L, Tong R, et al. UNet 3+: A full-scale connected UNet for medical image segmentation [J]. ArXiv: 2004. 08790, 2020.

[22] Jie H, Li S, Gang S, et al. Squeeze-and-excitation networks [J]. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2017, 42 (8): 2011-2023.

[23] Woo S, Park J, Lee J Y, et al. CBAM: Convolutional block attention module [J]. ArXiv: 1807. 06521, 2018.

[24] Fu J, Liu J, Tian H, et al. Dual attention network for scene segmentation [C]. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020. http://dx.doi.org/10.1109/CVPR.2019.00169.

[25] Huang Z, Wang X, Huang L, et al. CCNet: Criss-cross attention for semantic segmentation [C]. International Conference on Computer Vision, 2019. http://dx.doi.org/10.1109/ICCV.2019.00069.

[26] Liu H, Liu F, Fan X, et al. Polarized self-attention: towards high-quality pixel-wise regression [J]. ArXiv: 2107. 00782, 2021.

Note: Figure translations are in progress. See original paper for figures.