

Gramformer: Learning Crowd Counting via Graph-Modulated Transformer

Hui Lin¹, Zhiheng Ma², Xiaopeng Hong^{3,4*}, Qinnan Shangguan³, Deyu Meng¹

¹School of Mathematics and Statistics, Xi'an Jiaotong University

²Shenzhen Institute of Advanced Technology, Chinese Academy of Science

³Harbin Institute of Technology

⁴Peng Cheng Laboratory

linhuixjtu@gmail.com; zh.ma@siat.ac.cn; hongxiaopeng@iee.org;

sg12qt@gmail.com; dymeng@mail.xjtu.edu.cn

Abstract

Transformer has been popular in recent crowd counting work since it breaks the limited receptive field of traditional CNNs. However, since crowd images always contain a large number of similar patches, the self-attention mechanism in *Transformer* tends to find a homogenized solution where the attention maps of almost all patches are identical. In this paper, we address this problem by proposing Gramformer: a graph-modulated transformer to enhance the network by adjusting the attention and input node features respectively on the basis of two different types of graphs. Firstly, an attention graph is proposed to diverse attention maps to attend to complementary information. The graph is building upon the dissimilarities between patches, modulating the attention in an anti-similarity fashion. Secondly, a feature-based centrality encoding is proposed to discover the centrality positions or importance of nodes. We encode them with a proposed centrality indices scheme to modulate the node features and similarity relationships. Extensive experiments on four challenging crowd counting datasets have validated the competitiveness of the proposed method. Code is available at <https://github.com/LoraLinH/Gramformer>.

Introduction

Crowd counting, which aims to estimate the number of people in crowded scenes, is a core and challenging task in computer vision. It has a wide range of real-world applications, including congestion estimation, traffic monitoring, and other crowd management scenarios, especially after the outbreak of COVID-19. The increasing demand drives the prosperous development of counting methods.

Over the last few decades, Convolution Neural Network (CNN) has been utilized as the mainstream structure to regress the density map and predict the total crowd count (Zhang et al. 2016; Li, Zhang, and Chen 2018; Ma et al. 2019). However, the accuracy of these methods will sometimes be constrained by the fixed-size convolution kernel, where each patch only receives information from its spatial neighbours. This makes traditional CNNs to be difficult to deal with different variations in crowd images. Recently, with the blossom of transformers, some approaches (Liang



Figure 1: Visualization comparison of attention maps in different patches between vanilla attention and the proposed graph-guided attention. Each white cross mark indicates the location of a patch which the attention map is corresponding to. The vanilla attention finds a homogenized solution where the attention maps of most patches are similar (to the final density map), regardless of whether they are background (the first column) or foreground (the last two columns).

et al. 2022; Lin et al. 2022b,a) have leveraged on the self-attention mechanisms to further boost the accuracy. The transformer treats the image as a sequence of patches and each patch receives information from all other ones. With a dot-product operation, the intensity of information received is based on the feature similarity of different nodes.

While the transformer architecture is effective at computing crowd density in an image, it may produce homogenized attention maps as there are numerous regions that appear similar in a crowd image. In our experiments, we are surprised to find that the attention maps of most patches produced by the vanilla transformer are almost identical¹, irrespective of whether they are foreground or background. As depicted in Figure 1, these attention maps are predominantly focused on head positions and *similar* to the crowd density map (*homogenized attention*). It can be also seen

¹The recent study (Wang, Mei, and Yuille 2023) also reveals that the transformer in CLIP (Radford et al. 2021) tends to produce highly similar attention maps across various source points in natural images.

*Corresponding author.

Model	Dataset	ANVar
Vanilla Transformer	ImageNet val	5028.2
Vanilla Transformer	QNRF test	368.6
Our model	QNRF test	4777.6

Table 1: The comparison of the average normalized variance (ANVar) of attention maps generated by an Imagenet pre-trained transformer on natural images and crowded images (resized to a uniform size of 384×384). The ANVar of the attention maps learned by the transformer on crowd images (QNRF) is significantly lower ($<10\%$) than those on natural images (ImageNet). Our Gramformer reduces homogenization and diversifies the attention maps.

that in Table 1, the variance of the attention maps learned by the vanilla transformer on crowd images (QNRF) is below one-tenth of those on natural images (ImageNet). The congruence observed between the homogenized attention maps and the density maps accelerates the production of the target density map by the counting model. However, it comes at the cost of overlooking other details that may be relevant. As a result, this scheme only focuses on the current attention region, ignoring its relationships with other patches, potentially leading to performance loss.

In this paper, we address this problem by proposing a novel solution called *graph-modulated transformer* (*Gramformer* for crowd counting). We introduce two types of graphs, the attention graph and the feature-based centrality encoding graph, to modulate the attention and input node features of the transformer blocks, respectively. Building upon these two representations, we enhance the transformer network for crowd counting from two primary perspectives.

First, we propose a graph-guided attention modulation method to diversify the generated attention maps. We suppose that the key to diverse attention maps lies in attending to complementary information. Based on this understanding, we design an Edge Weight Regression (EWR) network to encode the *dissimilarity between patches* into edge weights. Considering the perspective geometry in crowd images, we further introduce an edge regularization term to limit the dissimilarity encoded by EWR within the same horizontal line to be as little as possible. The graph edges form an attention mask, which is incorporated into the transformer block to modulate the attention maps in an anti-similarity fashion.

Secondly, we propose a graph-guided node feature modulation method that encodes the centrality positions or importance of nodes in the graph. To achieve this, we devise a *Centrality Indices* scheme, which constructs a feature-based neighboring graph and explores the overall similarity of all nodes. Then we encode the graph structural information by the centrality embedding vectors to modulate input node features and thus further enhances the self-attention relations.

We propose a graph-modulated transformer architecture for crowd counting to modulate the attention and input node features of the transformer blocks, respectively. The contributions in detail can be summarized as follows:

- We propose a graph-guided attention modulation method

to diversify the generated attention maps in an anti-similarity manner.

- We propose an edge weight regression network with an edge regularization training term, to determine and regularize the edge weights for the *attention graph*.
- We propose a graph-guided node feature modulation method to capture centrality information and modulate the input feature node.
- Our approach consistently achieves promising counting performance on popular benchmarks.

Related Works

Graph Transformer

Recently, with the blossom of Transformer (Vaswani et al. 2017), researchers gradually focus on the generalization and combination between Transformers and GNNs (Han et al. 2022). GraphTransformer (Dwivedi and Bresson 2020) introduces a generalization of transformer networks to homogeneous graphs. It incorporates the attention mechanism into graph structure and proposes to use Laplacian eigenvector as positional encoding. GraphTrans (Wu et al. 2021) applies a permutation-invariant Transformer module after a standard GNN module. SAN (Kreuzer et al. 2021) learns positional encodings by proposing the Laplace spectrum of a given graph. LSPE (Dwivedi et al. 2021) presents a random-walk diffusion-based positional encoding scheme to initialize the positional representations of the nodes. It enables GNNs to learn both structural and positional representations. Graphi (Mialon et al. 2021) encodes structural and positional information by positive definite kernels and local sub-structures such as paths of short length. GPS (Rampásek et al. 2022) focuses and categorizes three main ingredients in unified transformer and graph neural networks: positional/structural encoding, local message-passing mechanism, and global attention mechanism. Graphormer (Ying et al. 2021) proposes several structural encodings to incorporate structural information reflected on nodes and the relation between node pairs. The proposed graph-guided transformer differs from existing approaches by modulating both the attention mechanism and node features respectively based on different types of constructed graph structures.

Crowd Counting

Early crowd counting methods adopt state-of-the-art detectors to perform the task, i.e., Faster R-CNN (Ren et al. 2015), YOLO (Redmon et al. 2016), RetinaNet (Lin et al. 2017), etc. However, the detection-based methods (Wu and Nevatia 2005; Idrees et al. 2013) appear to perform poorly in images with heavy occlusion. To improve this, density-based methods are introduced and gain promising counting accuracy.

CNN is popularly used to generate predicted density maps since it possesses translation equivariance and is effective in extracting local details (Liu et al. 2022), which has driven rapid developments in crowd countings (Zhang et al. 2016; Li, Zhang, and Chen 2018; Ma et al. 2019). Further improvements such as multi-scale mechanisms (Zeng et al. 2017; Ma et al. 2020; Sindagi and Patel 2017; Cao et al. 2018; Liu

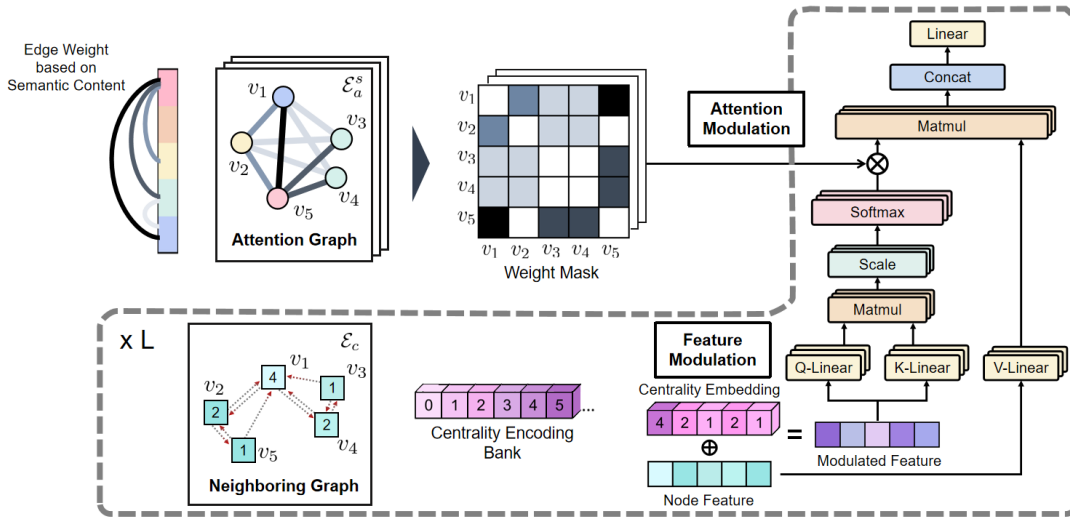


Figure 2: The framework of Gramformer, which contains two main parts: an attention graph to modulate the attention mechanism by its edge weight, and a feature-based centrality encoding graph to encode the centrality or importance of each node. In the attention graph, different colors represent different semantic values predicted by EWR, and the color difference corresponds to the strength of connecting edges. In centrality encoding, each node is assigned a centrality index, which is linked to its in-degree. The centrality index is used to find the corresponding embedding from a learnable bank to modulate node features.

et al. 2019), perspective estimation (Yang et al. 2020b; Yan et al. 2019; Yang et al. 2020a), auxiliary task (Liu et al. 2020; Zhao et al. 2019; Yang et al. 2020c), density refinement (Liu et al. 2018), optimal transport (Ma et al. 2021; Lin et al. 2021; Wang et al. 2020a; Wan, Liu, and Chan 2021) also show their effectiveness based on CNN.

The Vision Transformer (ViT) (Dosovitskiy et al. 2020) has demonstrated its outstanding performances in a variety of vision tasks (Carion et al. 2020; Zheng et al. 2021; Chen et al. 2021). Lately, some works adopt transformer to boost the counting performance. TransCrowd (Liang et al. 2022) uses transformer to reformulate the weakly-supervised crowd counting problem from the perspective of sequence-to-count. MAN (Lin et al. 2022b) incorporates global attention, learnable local attention, and instance attention into a counting model. CLTR (Liang, Xu, and Bai 2022) introduces an end-to-end transformer framework targeting on the crowd localization task. DACount (Lin et al. 2022a) proposes an agency-guided semi-supervised approach and uses a foreground transformer to refine crowd information.

Recently, HyGnn (Luo et al. 2020) uses a Hybrid GNN to formulate information from different scales and domains. RRP (Chen et al. 2020) employs a vanilla GCN to propagate information between different regions. STGN (Wu et al. 2022) and CoCo-GCN (Zhai et al. 2022) target on video crowd counting and multi-view crowd counting respectively from graph perspective. Our method is distinct from them in two ways. First, we construct two different types of graphs based on the dissimilarity between each two patches and the overall similarity of all nodes. Second, these graphs will be leveraged to modulate the attention mechanisms with an anti-similarity fashion, and node features with a centrality indices scheme, respectively.

The Proposed Method

Given a crowd image I , we denote the feature map extracted by the backbone as $F \in \mathbb{R}^{C \times W \times H}$, where C , W , and H are the channel, width, and height, respectively. These features can be represented as a set of nodes, denoted as $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$, where² $N = W \times H$. We then build two graphs with two different ways of edge construction, i.e., an attention graph with edges \mathcal{E}_a and a feature-based centrality encoding graph with edges \mathcal{E}_c . Guided by these two graphs, we improve the transformer by modulating the attention mechanism using the edge weights produced by the attention graph and biasing the input node features based on the in-degree of the node in the neighboring graph, which is termed by the graph-modulated transformer (Gramformer).

Structure of Graph-Modulated Transformer

The structure of Gramformer is shown in Figure 2. Gramformer’s main features include an attention modulation scheme and a node feature modulation scheme. In details, given a single attention head $s \in \{0, 1, \dots, S\}$ and corresponding learnable matrices $W_Q^s, W_K^s \in \mathbb{R}^{C \times \frac{C}{S}}$, we define the graph-modulated multi-head attention as:

$$\mathcal{R}^s(v_i^l, v_j^l) = E_{ij}^s \cdot \mathcal{S}\left(\frac{(\hat{v}_i^l W_Q^s)(\hat{v}_j^l W_K^s)^T}{\sqrt{C}}\right). \quad (1)$$

\mathcal{S} is the softmax function and l is the transformer layer. \hat{v}_i^l is the modulated node features, which is defined by

$$\hat{v}_i^l = v_i^l + \mathcal{P}(v_i^l). \quad (2)$$

²Without being ambiguous, the symbol v is also used to denote node features for simplicity.

E^s is a matrix of modulation factors, which is produced by the attention graph. $\mathcal{P}(v_i^l)$ is a biased vector, which is from a learnable embedding bank corresponding to the centrality index produced by the neighboring graph. Note that E^s is set in the first layer of the transformer and kept constant for the remaining layers, while $\mathcal{P}(v_i^l)$ is re-computed in every layers. In the following subsections, we will elaborate on how to calculate E^s and $\mathcal{P}(v_i^l)$ using two tailored graphs.

Attention Graph

The key to diverse attention maps lies in attending to complementary information. Based on this understanding, we propose an Edge Weight Regression (EWR) network to construct edges of the attention graph while encoding the dissimilarity between patches into edge weights.

Specifically, EWR estimates the semantic content value contained in each node and then determines the edges based on the degree of difference between the semantic values. For example, if the semantic content of nodes v_A and v_B have a large gap, a strong edge will be formed between v_A and v_B to encourage the model to explore the relationship between distinctive elements. On the other hand, if the node v_C has a similar information content to the node v_D , the model will create a weaker or no connection between v_C and v_D to reduce excessive attention in the subsequent attention module. We also extend the EWR to multiple heads to construct heterogeneous graphs with different edge connections to encode different kinds of semantics, which is consistent with multi-head attention. For head s , the weight and intensity of edge $E_{ij}^s \in \mathcal{E}_a$ between nodes v_i and v_j can be obtained by

$$E_{ij}^s = |\mathcal{F}_E^s(v_i^0) - \mathcal{F}_E^s(v_j^0)|. \quad (3)$$

\mathcal{F}_E^s denotes a head-specific EWR module, consisting of a feed-forward network (FFN) followed by an activating layer. We use a sigmoid function to normalize the output within a range of values from 0 to 1. In particular, 0 prohibits an edge construction. We directly apply the edges to modulate the attention mechanism based on their weights and intensities, i.e., $E = \mathcal{E}_a$. EWR is in an anti-similarity fashion, diversifying the attention by appropriately reducing the intensity between weakly connected nodes. In this way, the attention is guided to be diversified, which can effectively discover the missing context and residual details. In our model, the attention graph will be generated by EWR after the backbone and remain constant throughout subsequent transformer layers.

We further introduce an Edge Regularization term to encourage the dissimilarity encoded by EWR within the same horizontal line to be as little as possible. The idea is motivated by the existence of the perspective geometry of pin-hole cameras in crowd images, which usually generates perspective or scale variations only along the vertical axis (Shi et al. 2019). And within the same horizontal axis, these semantics remain consistent. Based on this observation, we propose a regularization to minimize the variance of the estimated semantic content value among nodes within the same horizontal axis. The regularization can be written as

$$\mathcal{Q} = \sum_h \sum_{\substack{i \\ s.t. y_i=h}} (\mathcal{F}_E(v_i^0) - \bar{\mathcal{F}}_h), \quad (4)$$

where y_i is the horizontal coordinate of each node and $\bar{\mathcal{F}}_h$ is the mean semantic value of the h_{th} row.

The Rationality of the Edge Regularization. Due to the imaging process of perspective geometry, crowd images often exhibit similar scales along the horizontal axis, resulting in similar density values in the density map. As the vanilla transformer tends to produce attention maps highly similar to the predicted density map, as observed in the introduction, the attention map will also exhibit similar values along the horizontal axis, leading to excessive focus between different locations on the same horizontal axis. Overly homogeneous attention maps are detrimental to counting performance. Therefore, to diversify the attention maps, the proposed edge regularization penalty suppresses cases where similar attention values are produced along the same horizontal axis and breaks the downside of excessive attention.

Feature-based Centrality Encoding

We then go on to modulate the node features by looking at the overall structure and discovering the role or importance of each node. To achieve this, we devise a Centrality Indices scheme and a feature-based neighboring graph, which explores the overall similarity of all nodes and modulate the node features by encoding the centrality representation \mathcal{P} into the node features.

For each node v_i , we select its similarity-based top q -NN, denoted as $\mathcal{N}(v_i)$ where q is the percentage. We then construct directed graph edges $E'_{ij} \in \mathcal{E}_c$ as follows:

$$E'_{ij} = \begin{cases} 1, & \text{if } v_j \in \mathcal{N}(v_i) \\ 0, & \text{otherwise} \end{cases}. \quad (5)$$

Here we adopt the L2 norm Euclidean distance to calculate the similarity. Then, we gather all the neighbour sets of the entire graph as $\mathcal{N} = \{\mathcal{N}(v_1), \mathcal{N}(v_2), \dots, \mathcal{N}(v_N)\}$. To identify the centrality, we count the number of occurrences of each node in \mathcal{N} . Nodes with high occurrences (in-degree in the graph) represent central positions in the feature space and we consider the occurrence as the centrality index. Suppose that $\{0, 1, \dots, m\}$ are the set of centrality indices, where m is the upper bound. All occurrences will be normalized to ensure that the maximum value does not exceed the bound.

We then introduce a learnable bank consisting of a series of centrality embedding vectors, denoted as $\{p_1, p_2, \dots, p_m\}$, $p \in \mathbb{R}^C$. Each node is encoded using the corresponding representation obtained through embedding matching. Specifically, we match the node v_j^l with its centrality index i from the bank, denoted as $\mathcal{P}(v_j^l) = p_i$, where i corresponds to the occurrences of node v_j^l within \mathcal{N} . In this way, we modulate features by encoding with distinctive vectors to incorporate the feature structural information into representations.

Further Processing

Based on the attention map, the features will be further processed by:

$$\tilde{v}_i^l = LN(v_i^l + \text{Concat}_s \left[\sum_{j=1}^N \mathcal{R}^s(\hat{v}_i^l, \hat{v}_j^l) \cdot (v_j^l W_v^s) \right] W_o), \quad (6)$$

where $W_v^s \in \mathbf{R}^{C \times \frac{C}{S}}$ is the projection weight matrix and \mathcal{R}^s is the graph-modulated attention map obtained by Equation 1. The layer normalization (LN) function is used to balance the scales of the values. All information from S graphs will be concatenated and finally integrated by the weight matrix $W^o \in \mathbf{R}^{C \times C}$ to update the node features. With an additional FFN (feedforward network) \mathcal{K} to fuse the features, the updated nodes can be obtained by:

$$\mathcal{V}^{l+1} = LN(\tilde{\mathcal{V}}^l + \mathcal{K}(\tilde{\mathcal{V}}^l)). \quad (7)$$

To be general, we denote the initial nodes extracted by the backbone as \mathcal{V}^0 , and each transformer layer updates the node features to $\mathcal{V}^1, \dots, \mathcal{V}^L$, where L is the number of layers.

It is important to note that while the attention graph \mathcal{E}_a remains constant throughout the L layers, the neighboring and similarity relationships will change in each layer after the update of features. This means that the feature-based neighboring graph is dynamic and the centrality indices in each layer will adjust to represent these changes. It allows the model more effectively and timely to capture the dynamic centrality feature positions of nodes.

Finally, we transmit graph nodes \mathcal{V}^L into the density regression head to obtain the predicted density map. For counting supervision, we utilize the Instance Attention Loss (IAL) to mitigate the negative impact of annotation noise (Lin et al. 2022b). IAL dynamically focuses on the most important annotations and ignores supervisions with high uncertainty. Thus the final loss is expressed by

$$\mathcal{L} = \mathcal{L}_{IA} + \lambda \mathcal{Q}, \quad (8)$$

where λ is the weight of the regularization.

Experimental Results

Datasets

We evaluate our crowd counting method and compare it with other state-of-the-art methods on four largest crowd counting benchmarks. They are widely used in recent papers and are described as follows.

- ShanghaiTech A (Zhang et al. 2016) includes 482 images with 244,167 annotated points. 300 images are divided for the training set and the remaining 182 images are for the test set. The numbers of people vary from 33 to 3,139.
- UCF-QNRF (Idrees et al. 2018) contains 1,535 crowded images with about 1.25 million annotated points. It has a wide range of people count and images with an average resolution of $2,013 \times 2,902$. The training and test sets include 1,201 and 334 images, respectively.
- NWPU-CROWD (Wang et al. 2020b) contains 5,109 images and 2.13 million annotated instances with point and box labels. 3,109 images are used in the training set; 500 images are in the validation set; and the remaining 1,500 images are in the test set. Images in NWPU-CROWD are in largely various density and illumination scenes. It is the largest dataset in which test labels are not released.
- JHU-CROWD++ (Sindagi, Yasarla, and Patel 2020) has a more complex context with 4,372 images and 1.51 million annotated points. 2,272 images are chosen for the

training set; 500 images are for the validation set; and the rest 1,600 images for the test set. The dataset contains diverse scenarios and is collected under different environmental conditions of weather and illumination.

Implement Details

Network Structure: VGG-19 (Pretrained on ImageNet) (Simonyan and Zisserman 2014) is adopted as the backbone network to extract features which will be formulated into graphs. The Edge Weight Regression (EWR) adopts two 3×3 convolution layers with a ReLU activation function in the middle and a sigmoid function at the end. The density regression head consists of an upsampling layer and three convolution layers with two ReLUs in middle. The kernel sizes of the first two layers are 3×3 and that of the last is 1×1 .

Training Details: We set the training batch size as 1 and crop images with a size of 512×512 . As some images in ShanghaiTech A contain smaller resolution, the crop size for this dataset changes to 256×256 . Random scaling of $[0.75, 1.25]$ and horizontal flipping are also adopted to augment each training image. We use Adam algorithm (Kingma and Ba 2014) with a learning rate 10^{-5} to optimize the parameters. We set the percentage of nearest neighbors q as 30%, and the maximum in-degree bound m as 18. The number of transformer layers L is 2 and the loss weight λ is 0.1.

Comparison with State-of-the-art Methods

We evaluate our model on the above four datasets and list recent state-of-the-art methods for comparison. The quantitative results, measured by MAE (Mean Absolute Error), MSE (Mean Square Error) and NAE (Normalized Absolute Error), are listed in Tables 2 and 3.

Our method achieves great counting accuracy on all four benchmark datasets. Generally, by comparing Gramformer against other state-of-the-art methods, Gramformer outperforms or at least is on par with these competitors on all datasets. Specifically, on dense datasets NWPU and UCF-QNRF, our method achieves the best MAE and MSE. On JHU++, Gramformer performs best on one of the evaluation metrics. And on the sparse dataset, ShanghaiTech A, Gramformer ranks second and shows a competitive performance.

Furthermore, we observe that Gramformer outperforms other approaches by large margins and is on par with the recent methods MAN and CLTR, which are built based on transformer. We believe the characteristic of the attention module in the transformer works effectively on crowd counting task. And especially, it will help gain great accuracy improvements on dense datasets, i.e., UCF-QNRF, JHU++ and NWPU. Therefore, our graph-guided improvement in the application of attention to the counting task is an inspiring direction. We also present a visualization comparison in the appendix.

Key Issues and Discussion

Ablation Studies: We perform the ablation experiments on UCF-QNRF to study the effect of each proposed term in

Dataset Method	ShanghaiTech A		UCF-QNRF		JHU++	
	MAE	MSE	MAE	MSE	MAE	MSE
MCNN (Zhang et al. 2016) (CVPR 16)	110.2	173.2	277	426	188.9	483.4
CA-Net (Liu, Salzmann, and Fua 2019) (CVPR 19)	61.3	100.0	107.0	183.0	100.1	314.0
CG-DRCN-CC (Sindagi, Yasarla, and Patel 2020) (PAMI 20)	60.2	94.0	95.5	164.3	71.0	278.6
DPN-IPSM (Ma et al. 2020) (ACMMM 20)	58.1	91.7	84.7	147.2	-	-
BL (Ma et al. 2019) (ICCV 19)	62.8	101.8	88.7	154.8	75.0	299.9
DM-Count (Wang et al. 2020a) (NIPS 20)	59.7	95.7	85.6	148.3	-	-
UOT (Ma et al. 2021) (AAAI 21)	58.1	95.9	83.3	142.3	60.5	252.7
GL (Wan, Liu, and Chan 2021) (CVPR 21)	61.3	95.4	84.3	147.5	59.9	259.5
S3 (Lin et al. 2021) (IJCAI 21)	57.0	96.0	80.6	139.8	59.4	244.0
P2PNet (Song et al. 2021) (ICCV 21)	52.7	85.1	85.3	154.5	-	-
ChfL (Shu et al. 2022) (CVPR 22)	57.5	94.3	80.3	137.6	57.0	235.7
STEERER (Han et al. 2023)	55.6	87.3	76.7	135.1	55.4	221.4
MAN (Lin et al. 2021) (CVPR 22)	56.8	90.3	77.3	<u>131.5</u>	53.4	209.9
CLTR (Liang, Xu, and Bai 2022) (ECCV 22)	56.9	95.2	85.8	141.3	59.5	240.6
Gramformer	<u>54.7</u>	<u>87.1</u>	76.7	129.5	53.1	<u>228.1</u>

Table 2: Comparisons with the state of the arts on ShanghaiTech A, UCF-QNRF and JHU-Crowd++. The best performance is shown in bold and the second best is shown in underlined. Note that the results of STEERER are based on the VGG19, which shares the same backbone as our method.

Dataset Method	MAE	Overall MSE	NAE	Avg.	Scene Level (MAE)
					$S_0 \sim S_4$
MCNN (Zhang et al. 2016)	232.5	714.6	1.063	1171.9	356.0 / 72.1 / 103.5 / 509.5 / 4818.2
SANet (Cao et al. 2018)	190.6	491.4	0.991	716.3	432.0 / 65.0 / 104.2 / 385.1 / 2595.4
CSRNet (Li, Zhang, and Chen 2018)	121.3	387.8	0.604	522.7	176.0 / 35.8 / 59.8 / 285.8 / 2055.8
BL (Ma et al. 2019)	105.4	454.2	0.203	750.5	66.5 / 8.7 / 41.2 / 249.9 / 3386.4
SFCN+ (Wang et al. 2020b)	105.7	424.1	0.254	712.7	54.2 / 14.8 / 44.4 / 249.6 / 3200.5
UOT (Ma et al. 2021)	87.8	387.5	0.185	566.5	80.7 / 7.9 / 36.3 / 212.0 / 2495.4
DMCount (Wang et al. 2020a)	88.4	388.6	<u>0.169</u>	498.0	146.7 / <u>7.6</u> / 31.2 / 228.7 / 2075.8
GL (Wan, Liu, and Chan 2021)	79.3	346.1	0.180	508.5	92.4 / 8.2 / 35.4 / 179.2 / 2228.3
S3 (Lin et al. 2021)	81.7	349.8	0.222	466.5	78.2 / 10.5 / 35.3 / 206.2 / <u>2002.4</u>
ChfL (Shu et al. 2022)	76.8	343.0	0.170	470.1	56.7 / 8.4 / <u>32.1</u> / 195.1 / 2058.0
MAN (Lin et al. 2022b)	<u>76.5</u>	<u>323.0</u>	0.170	<u>464.6</u>	43.3 / 8.5 / 35.3 / 190.9 / 2044.9
Gramformer	72.5	316.4	0.160	441.9	26.9 / 7.4 / 33.4 / <u>186.8</u> / 1955.2

Table 3: Comparisons on NWPU-CROWD, which is further divided into five refined subsets $S_0 \sim S_4$ according to crowd size.

Components	MAE	MSE
Baseline (VGG+IAL)	84.7	150.9
B + EWR	79.3	135.4
B + Centrality Indices	79.8	138.1
B + EWR + CI	77.9	130.2
B + EWR + Edge Regularization	78.1	132.8
B + EWR + CI + ER	76.7	129.5

Table 4: Ablation study on UCF-QNRF. After combining all proposed components, the model achieves its best accuracy.

Gramformer and provide quantitative results in Table 4. We start with the baseline of VGG+IAL, which details can be referred to MAN (Lin et al. 2022b). We first test the contribution of graph-modulated attention by EWR. The performance from baseline is improved by 5.4 and 15.5, for MAE and MSE, respectively. Then, we incorporate the cen-

tral indices scheme into our model, the counting accuracy in terms of MAE is further improved by 3.9 without EWR and 1.4 with EWR. And finally, we adopt Edge Regularization to insert the prior knowledge of perspective geometry in crowd images. The model then achieves its best with 76.7 and 129.5 for MAE and MSE.

m	MAE	MSE
36	80.2	140.1
24	78.4	136.9
18	76.7	129.5
12	80.2	138.7

Table 5: The influence of the number of centrality embedding vectors. Experiments are conducted on UCF-QNRF.

Graph-guided Attention and Graph Transformer: We study the effectiveness of the proposed graph-guided attention. We choose Graphormer (Ying et al. 2021), one of the most popular models among graph transformers, as our comparison model. We extract its attention module to replace our proposed graph-guided attention. Specifically, its edge connections are constructed according to the nearest neighbor similarities. The feature of each edge is obtained by concatenating the features of two endpoint nodes and smoothing with an MLP (Multilayer Perceptron). The feature of its edge is then encoded into a single value and added to the attention map, which is R in Equation 1. More detailed information about the model can be found in the appendix. The results of this baseline on UCF-QNRF is 80.9 for MAE and 142.5 for MSE, with a worse of 4.2 for MAE and 13.0 for MSE respectively. The comparison shows the efficacy of building edges by diversifying the attention maps for more abundant information.

The Influence of Centrality Indices: We hold experiments on UCF-QNRF to study the influence of the number of centrality embedding vectors. m is the upper bound of in-degree ranking values, determining the number of distinguished embedding vectors. The results are shown in Table 5. When $m = 76.7$, our model performs best. When m is reduced or elevated, the accuracy drops accordingly. We posit that an ill-suited value of m , resulting in either a limited number of encodings or an excessive set of embedding choices, has the potential to undermine the performance.

The Influence of The Percentage of Nearest Neighbors q : q determines how many nodes are selected as neighbors of a node when calculating its in-degree value. A lower q allows each node to select only the most similar ones, which reduces the number of nodes with high occurrences and corresponding centrality indices. And when q grows larger, the neighbor selection process becomes more relaxed. At this point, nodes with high in-degree values will increase, influencing the learning ability of graph structure via embedding vectors. We hold experiments to study these influences. The results on UCF-QNRF are shown in Table 6. Gramformer achieves its best when $q = 0.3$. And the accuracy drops when q gradually increases or decreases. It justifies the importance of the nearest neighbor selection process in the updating of graph structure and subsequent count predictions. An appropriate number of selected neighbors will aid in improving the counting accuracy.

We also study the combined influence of m and q . Nine models are trained with different settings: a cross-product of $m \in \{12, 18, 24\}$ and $q \in \{0.2, 0.3, 0.4\}$, to show a clear comparison. Table 7 reports MAE of counting results on UCF-QNRF. The counting accuracy of the model is affected by different settings. And Gramformer performs best under the setting in the center with $m = 18$ and $q = 0.3$.

Running Cost Evaluation: In Table 8, we compare the model size, the floating point operations (FLOPs) computed on one 384×384 input and the inference time for 100 images between Gramformer and other models including of ViT-B (Dosovitskiy et al. 2020), Bayesian (Ma et al. 2019),

q	MAE	MSE
0.06	81.6	140.1
0.08	78.3	136.0
0.1	78.8	135.4
0.2	77.9	132.3
0.3	76.7	129.5
0.4	78.2	133.1

Table 6: The influence of the nearest neighbor selection. When calculating the in-degree value, the decision of each node to connect to too many or too few (q) neighbors will reduce the counting accuracy.

Centrality Bank Size	Nearest Neighbors		
	0.2	0.3	0.4
12	78.8	80.2	81.5
18	77.9	76.7	78.2
24	81.4	78.4	80.6

Table 7: Performances of Gramformer under different settings of centrality embedding and nearest neighbors on UCF-QNRF (measured as MAE).

Model	Model Size (M)	GFLOPs	Inference time
ViT-B	86.0	55.4	21.3
Bayesian	21.5	56.9	10.3
VGG19+Trans	29.9	58.0	10.8
MAN	30.9	58.2	11.3
Gramformer	29.0	60.9	12.6

Table 8: Running Cost Evaluation.

the combination of VGG19 and Transformer (Vaswani et al. 2017) and MAN (Lin et al. 2022b). All experiments are conducted with a single RTX 3080 GPU. Gramformer has a slight increase in FLOPs and inference time, for the computation of nearest neighbours in forward predictions. However, our model is smaller than VGG19+Trans and MAN, and significantly smaller than ViT-B. This means that the proposed components are lightweight.

Conclusion

This paper aims to enhance the ability of transformers to modulate the attention mechanism and input node features respectively on the basis of two different types of graphs. We contribute to diversifying the attention map to attend to more complementary information by proposing a graph-guided attention modulation. We also encode the centrality or importance of nodes by designing a centrality indices scheme to adjust the input node features. The proposed Gramformer achieves high counting accuracy on popular crowd counting datasets. Improving the transformer network by graph modulation is an inspiring direction, and we will apply it to a wider range of vision tasks.

Acknowledgements

This work is funded by the National Natural Science Foundation of China (62076195, 62206271, 62376070, 12226004 and 61721002), the Fundamental Research Funds for the Central Universities (AUGA5710011522), the Guangdong Basic and Applied Basic Research Foundation (2020B1515130004), and the Shenzhen Key Technical Projects under Grant (JSGG20220831105801004, CJGJZD2022051714160501, and JCYJ20220818101406014).

References

- Cao, X.; Wang, Z.; Zhao, Y.; and Su, F. 2018. Scale aggregation network for accurate and efficient crowd counting. In *ECCV*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*.
- Chen, X.; Bin, Y.; Gao, C.; Sang, N.; and Tang, H. 2020. Relevant region prediction for crowd counting. *Neurocomputing*.
- Chen, X.; Yan, B.; Zhu, J.; Wang, D.; Yang, X.; and Lu, H. 2021. Transformer tracking. In *CVPR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Dwivedi, V. P.; and Bresson, X. 2020. A generalization of transformer networks to graphs. *arXiv preprint*.
- Dwivedi, V. P.; Luu, A. T.; Laurent, T.; Bengio, Y.; and Bresson, X. 2021. Graph Neural Networks with Learnable Structural and Positional Representations. In *ICLR*.
- Han, K.; Wang, Y.; Guo, J.; Tang, Y.; and Wu, E. 2022. Vision GNN: An Image is Worth Graph of Nodes. In *NIPS*.
- Han, T.; Bai, L.; Liu, L.; and Ouyang, W. 2023. Steerer: Resolving scale variations for counting and localization via selective inheritance learning. In *ICCV*, 21848–21859.
- Idrees, H.; Saleemi, I.; Seibert, C.; and Shah, M. 2013. Multi-source multi-scale counting in extremely dense crowd images. In *ICCV*.
- Idrees, H.; Tayyab, M.; Athrey, K.; Zhang, D.; Al-Maadeed, S.; Rajpoot, N.; and Shah, M. 2018. Composition loss for counting, density map estimation and localization in dense crowds. In *ECCV*.
- Kingma, D. P.; and Ba, J. L. 2014. ADAM: AMETHOD FOR STOCHASTIC OPTIMIZATION. *arXiv preprint*.
- Kreuzer, D.; Beaini, D.; Hamilton, W.; Létourneau, V.; and Tossou, P. 2021. Rethinking graph transformers with spectral attention. *NIPS*.
- Li, Y.; Zhang, X.; and Chen, D. 2018. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In *CVPR*.
- Liang, D.; Chen, X.; Xu, W.; Zhou, Y.; and Bai, X. 2022. Transcrowd: weakly-supervised crowd counting with transformers. *Science China Information Sciences*.
- Liang, D.; Xu, W.; and Bai, X. 2022. An end-to-end transformer model for crowd localization. In *ECCV*.
- Lin, H.; Hong, X.; Ma, Z.; Wei, X.; Qiu, Y.; Wang, Y.; and Gong, Y. 2021. Direct Measure Matching for Crowd Counting. In *IJCAI*.
- Lin, H.; Ma, Z.; Hong, X.; Wang, Y.; and Su, Z. 2022a. Semi-supervised Crowd Counting via Density Agency. In *ACM MM*.
- Lin, H.; Ma, Z.; Ji, R.; Wang, Y.; and Hong, X. 2022b. Boosting Crowd Counting via Multifaceted Attention. In *CVPR*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *ICCV*.
- Liu, L.; Qiu, Z.; Li, G.; Liu, S.; Ouyang, W.; and Lin, L. 2019. Crowd counting with deep structured scale integration network. In *ICCV*, 1774–1783.
- Liu, L.; Wang, H.; Li, G.; Ouyang, W.; and Lin, L. 2018. Crowd Counting using Deep Recurrent Spatial-Aware Network. In *IJCAI*.
- Liu, W.; Salzmann, M.; and Fua, P. 2019. Context-aware crowd counting. In *CVPR*.
- Liu, Y.; Liu, L.; Wang, P.; Zhang, P.; and Lei, Y. 2020. Semi-supervised crowd counting via self-training on surrogate tasks. In *ECCV*.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *CVPR*.
- Luo, A.; Yang, F.; Li, X.; Nie, D.; Jiao, Z.; Zhou, S.; and Cheng, H. 2020. Hybrid graph neural networks for crowd counting. In *AAAI*.
- Ma, Z.; Wei, X.; Hong, X.; and Gong, Y. 2019. Bayesian loss for crowd count estimation with point supervision. In *ICCV*.
- Ma, Z.; Wei, X.; Hong, X.; and Gong, Y. 2020. Learning scales from points: A scale-aware probabilistic model for crowd counting. In *ACMMM*.
- Ma, Z.; Wei, X.; Hong, X.; Lin, H.; Qiu, Y.; and Gong, Y. 2021. Learning to count via unbalanced optimal transport. In *AAAI*.
- Mialon, G.; Chen, D.; Selsosse, M.; and Mairal, J. 2021. Graphit: Encoding graph structure in transformers. *arXiv preprint*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*. PMLR.
- Rampásek, L.; Galkin, M.; Dwivedi, V. P.; Luu, A. T.; Wolf, G.; and Beaini, D. 2022. Recipe for a general, powerful, scalable graph transformer. *arXiv preprint*.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *CVPR*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*.
- Shi, M.; Yang, Z.; Xu, C.; and Chen, Q. 2019. Revisiting perspective information for efficient crowd counting. In *CVPR*.
- Shu, W.; Wan, J.; Tan, K. C.; Kwong, S.; and Chan, A. B. 2022. Crowd counting in the frequency domain. In *CVPR*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint*.
- Sindagi, V.; Yasarla, R.; and Patel, V. M. 2020. Jhu-crowd++: Large-scale crowd counting dataset and a benchmark method. *PAMI*.
- Sindagi, V. A.; and Patel, V. M. 2017. Generating high-quality crowd density maps using contextual pyramid cnns. In *ICCV*.
- Song, Q.; Wang, C.; Jiang, Z.; Wang, Y.; Tai, Y.; Wang, C.; Li, J.; Huang, F.; and Wu, Y. 2021. Rethinking Counting and Localization in Crowds: A Purely Point-Based Framework. In *ICCV*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*.
- Wan, J.; Liu, Z.; and Chan, A. B. 2021. A Generalized Loss Function for Crowd Counting and Localization. In *CVPR*.

- Wang, B.; Liu, H.; Samaras, D.; and Nguyen, M. H. 2020a. Distribution Matching for Crowd Counting. *NIPS*.
- Wang, F.; Mei, J.; and Yuille, A. 2023. SCLIP: Rethinking Self-Attention for Dense Vision-Language Inference. *arXiv preprint*.
- Wang, Q.; Gao, J.; Lin, W.; and Li, X. 2020b. NWPUCrowd: A large-scale benchmark for crowd counting and localization. *PAMI*.
- Wu, B.; and Nevatia, R. 2005. Detection of multiple, partially occluded humans in a single image by bayesian combination of edgelet part detectors. In *ICCV*.
- Wu, Z.; Jain, P.; Wright, M.; Mirhoseini, A.; Gonzalez, J. E.; and Stoica, I. 2021. Representing long-range context for graph neural networks with global attention. *NIPS*.
- Wu, Z.; Zhang, X.; Tian, G.; Wang, Y.; and Huang, Q. 2022. Spatial-Temporal Graph Network for Video Crowd Counting. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Yan, Z.; Yuan, Y.; Zuo, W.; Tan, X.; Wang, Y.; Wen, S.; and Ding, E. 2019. Perspective-guided convolution networks for crowd counting. In *ICCV*.
- Yang, Y.; Li, G.; Du, D.; Huang, Q.; and Sebe, N. 2020a. Embedding perspective analysis into multi-column convolutional neural network for crowd counting. *IEEE Transactions on Image Processing*.
- Yang, Y.; Li, G.; Wu, Z.; Su, L.; Huang, Q.; and Sebe, N. 2020b. Reverse perspective network for perspective-aware object counting. In *CVPR*.
- Yang, Y.; Li, G.; Wu, Z.; Su, L.; Huang, Q.; and Sebe, N. 2020c. Weakly-supervised crowd counting learns from sorting rather than locations. In *ECCV*.
- Ying, C.; Cai, T.; Luo, S.; Zheng, S.; Ke, G.; He, D.; Shen, Y.; and Liu, T.-Y. 2021. Do transformers really perform badly for graph representation? *NIPS*.
- Zeng, L.; Xu, X.; Cai, B.; Qiu, S.; and Zhang, T. 2017. Multi-scale convolutional neural networks for crowd counting. In *ICIP*.
- Zhai, Q.; Yang, F.; Li, X.; Xie, G.-S.; Cheng, H.; and Liu, Z. 2022. Co-Communication Graph Convolutional Network for Multi-View Crowd Counting. *IEEE Transactions on Multimedia*.
- Zhang, Y.; Zhou, D.; Chen, S.; Gao, S.; and Ma, Y. 2016. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*.
- Zhao, M.; Zhang, J.; Zhang, C.; and Zhang, W. 2019. Leveraging heterogeneous auxiliary tasks to assist crowd counting. In *CVPR*.
- Zheng, S.; Lu, J.; Zhao, H.; Zhu, X.; Luo, Z.; Wang, Y.; Fu, Y.; Feng, J.; Xiang, T.; Torr, P. H.; et al. 2021. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*.