

000 001 002 003 004 005 DENSEViG: DECOUPLED ENERGY-GUIDED GRAPH 006 STRUCTURE REFINEMENT FOR VISION GNNs 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031

032 ABSTRACT 033

034 Vision Graph Neural Networks (ViG) treats an image as a set of visual patches for
035 graph representation learning and yields promising results across various com-
036 puter vision tasks. However, most existing works primarily focus on static graph
037 construction, ignoring the performance gains and noise reduction benefits of dy-
038 namic structure refinement. Meanwhile, generative models such as Energy-based
039 Models (EBMs) are generally unsuitable for discriminative tasks and struggle with
040 large-scale images. Our goal is to introduce a unified generative-discriminative
041 paradigm for dynamically modeling relationships between visual patches, aim-
042 ing to produce higher-quality representations for improving downstream tasks.
043 Specifically, we propose **D**ecoupled **E**nergy **L**earning (DEL) that defines a joint
044 distribution of sample pairs to approximate the target distribution. It decouples
045 EBMs into energy matching and contrastive learning as a global loss function,
046 which pulls similar pairs closer and pushes dissimilar pairs further apart in the
047 representation space. For implementation, we develop an end-to-end framework,
048 termed **D**ecoupled **E**Nergy learning **g**uided **S**tructure **r**Efinement for improving
049 **ViG** (DenseViG). Structure refinement is deployed within ViG architectures in a
050 plug-and-play manner, dynamically adding or pruning edges based on similarity
051 metrics with a relaxation strategy. Theoretical analyses demonstrate the effective-
052 ness of DenseViG in processing large datasets through graph operations. Empiri-
053 cal evaluations confirm that it outperforms state-of-the-art methods on three major
054 benchmarks, achieving 84.3% Top-1 accuracy on ImageNet-1K, 46.4% mAP on
055 MS COCO, and 50.9% mIoU on ADE20K.

056 INTRODUCTION 057

058 Graph Neural Networks (GNNs) are initially developed to handle graph-structured data due to their
059 strength in modeling relational dependencies and capturing complex topology. More recently, GNNs
060 have been successfully applied to computer vision Tian et al. (2024); Wang et al. (2023); Rahman
061 & Marculescu (2024); Xu et al. (2024); Bhowmik et al. (2023). Among these efforts, Han et al.
062 Han et al. (2022) pioneered Vision GNNs (ViG), which divides images into patches, constructs a
063 graph over visual features, and propagates information via a message-passing scheme. Rather than
064 relying on traditional 2D convolutions or self-attention mechanisms, ViG adopts graph convolution
065 as a novel paradigm for visual representation learning. It achieves competitive performance across
066 various benchmarks, including image classification Wang et al. (2025); Parikh et al. (2025); Li et al.
067 (2023), object detection Yang et al. (2025); Wu et al. (2023), and image segmentation Li et al. (2024);
068 Jiang et al. (2023). However, the main drawback of ViG lies in its largely static graph construction
069 using the conventional k -Nearest-Neighbor (k -NN) algorithm and its variants Han et al. (2022);
070 Wang et al. (2025); Yang et al. (2025); Spadaro et al. (2025); Yao et al. (2024). These methods
071 always suffer from fixed topologies, noise interference and limited high-order semantics. To avoid
072 this, Munir et al. (2024) designed a dynamic axial graph construction that connects patches whose
073 Euclidean distance falls below a threshold set by the distance distribution's mean and standard devi-
074 ation. Han et al. (2023) devised a dynamic hypergraph construction that softly clusters semantically
075 akin patches with Fuzzy C-Means, treating each cluster as a set of hyperedges. Both methods rely
076 on heuristic rules and shallow features without learnable parameters, preventing end-to-end training
077 via gradient updates and restricting the flexibility of ViG.

054
 055 **Motivation** We aim to comprehensively en-
 056 hance the performance of ViG on downstream
 057 tasks. Prior works typically begin by select-
 058 ing expressive generative models, particularly
 059 Energy-based Models (EBMs) LeCun et al.
 060 (2006), known for their strong representational
 061 capacity. To better align EBMs with discrim-
 062 inative tasks such as classification and detec-
 063 tion Bhowmik et al. (2023); Grathwohl et al.
 064 (2019), recent studies by Kim & Ye (2022) and
 065 Wang et al. (2022) have explored several com-
 066 bined approaches using contrastive objectives.
 067 They have demonstrated effective representa-
 068 tions and favorable outcomes on toy datasets
 069 (e.g., MNIST and CIFAR). However, its scal-
 070 ability to large-scale datasets is hindered by
 071 the time-consuming iterations and high mem-
 072 ory consumption associated with EBMs. In-
 073 spired by the success of Zeng et al. (2025) in
 074 applying this investigation to graph-based do-
 075 mains, we utilize the ViG architecture of con-
 076 verting pixel-level images into sparse graphs.
 077 Since reducing the data dimensionality by several folds, it shortens the iteration time and lowers
 078 memory, both of which are beneficial for processing high-resolution imagery. In summary, our
 079 perspective positions graph learning as a critical bridge between EBMs and large-scale visual data.
 080

081 **Contribution** In this paper, we present a Decoupled Energy Learning (DEL) method that addresses
 082 the limitations of single-paradigm modeling. DEL incorporates EBMs with Contrastive Learning
 083 (CL) to define a joint distribution over paired samples for approximating the target data distribution.
 084 Specifically, we decouple its objective as a joint loss including two complementary terms, leverag-
 085 ing the benefits of generative models while retaining strong discriminability for downstream tasks.
 086 During energy matching process, DEL objective maximizes the joint log-likelihood of the similarity
 087 between positive pairs using EBMs, while simultaneously minimizing that of negative pairs using
 088 CL. We theoretically justify that the standard contrastive loss emerges as a special case of the DEL
 089 loss when the generative term is omitted.

090 Building on our motivation, we introduce a novel framework that **Decoupled ENergy** learning
 091 guided Structure rEfinement for improving **ViG**, named DenseViG. This framework performs graph
 092 structure refinement in a modular, plug-and-play fashion, enabling seamless integration into existing
 093 ViG architectures without modifying any components. Within each ViG stage, an intrinsic GNN
 094 encoder embeds the constructed graph to generate node representations. Dynamic refinement en-
 095 hances the raw structure by adding or pruning edges according to learned pairwise representation
 096 similarities with a Gumbel-Sigmoid relaxation strategy. DEL loss provides a global optimization
 097 to facilitate both high-quality representations and optimal structures. DenseViG is evaluated on
 098 three benchmarks for image classification, object detection and semantic segmentation. The major
 099 contributions are threefold as follows:

1. Incorporating EBMs with contrastive learning into DEL approximates the target distribution by pulling similar representations closer and pushing dissimilar ones apart. DEL objective is decoupled into two complementary terms, allowing the unified paradigm to fit more naturally within discriminative tasks while preserving generative capacity.
2. By integrating dynamic refinement into each ViG stage in a plug-and-play manner, we propose a novel end-to-end trainable framework, called DenseViG. Our framework encourages ViG architectures to generate powerful representations and refined structures, consistently improving the performance of downstream vision tasks.
3. Leveraging graph processing, our framework shortens iteration time and reduces memory, enabling training on large-scale images and achieving superior results across three vision benchmarks. Specifically, DenseViG yields 84.3% Top-1 accuracy on ImageNet-1K, 46.4% mAP on MS COCO, and 50.9% mIoU on ADE20K.

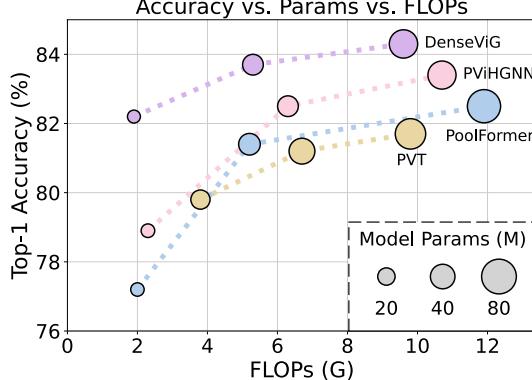


Figure 1: Params, FLOPs, and Top-1 accuracy on ImageNet-1K for representative ViT, ViG, Pool, and HGNN backbones. DenseViG-Ti, S, and M lie on the accuracy and efficiency frontier, reaching 82.2%, 83.7%, and 84.3% Top-1 accuracy. This suggests that decoupled energy-guided graph structure refinement scales to image datasets and points to stronger Vision GNNs.

108 RELATED WORKS
109110 VISUAL REPRESENTATION LEARNING
111

112 Generative models such as EBMs have shown unique potential in computer vision. EBMs represent
113 the data distribution as a scalar-valued energy function Arbel et al. (2020). Recent advances involve
114 deep neural parameterizations of the energy function, widely applied for image reconstruction Pang
115 et al. (2020), object detection Gustafsson et al. (2021), and human motion modeling Zhang et al.
116 (2025). Optimization of EBMs is often considered challenging, and contrastive techniques have
117 proven effective Gutmann & Hyvärinen (2010). Kim & Ye (2022) introduce a hybrid framework that
118 enriches the SimCLR objective with a generative energy-based term, explicitly modeling the joint
119 distribution of positive pairs rather than relying solely on discriminative alignment. Similarly, Wang
120 et al. (2022) propose a unified probabilistic framework that reinterprets adversarial training as the
121 maximum-likelihood estimation of EBMs, clarifying the generative capability of robust classifiers
122 and bridging supervised adversarial training with contrastive learning. These studies motivate a self-
123 supervised framework that leverages both generative expressiveness and discriminative precision for
124 visual representation learning on small- and medium-scale datasets. However, they still suffer from
125 high computational and memory costs on large-scale data.

126 VISION GRAPH NEURAL NETWORKS
127

128 Vision graph neural networks have become a rapidly growing alternative to grid-based architec-
129 tures, allowing flexible interactions between local regions and efficient global context aggregation
130 in images. Among ViG components, graph construction is critical a link of improvement. Static
131 construction methods develop several k -NN variants to reduce time and space complexity. WiGNet
132 Spadaro et al. (2025) mitigates the quadratic overhead of vanilla k -NN through window partition-
133 ing, achieving nearly linear scalability with image size. ClusterViG Parikh et al. (2025) applies
134 the k-means algorithm to cluster image tokens, effectively compressing the search space of k -NN.
135 Dynamic graph approaches, including MobileViG Munir et al. (2023) and GreedyViG Munir et al.
136 (2024), replace expensive fully-connected graphs with sparse, structured patch interactions to ac-
137 celerate computation. These structures dynamically update during the forward passes, improving
138 adaptability across scenarios. More recently, hypergraph methods, such as DVHGNN Li et al.
139 (2025), HgVT Fixelle (2025), and HGFormer Wang et al. (2025), transform patch tokens into hy-
140 pergraph structures, encoding higher-order relationships through dynamic hypergraph convolutions
141 or attention mechanisms, thus capturing richer semantic information. Unlike the aforementioned
142 approaches, we emphasize structure refinement by utilizing a unified paradigm and integrate it into
143 the existing ViG backbones as an end-to-end framework, enhancing downstream tasks.

144 METHODS
145

146 In this section, we detail the DenseViG framework. First, we formalize the graph notation and
147 review the concepts of EBMs. We then present the core DEL objective, followed by the structure
148 refinement module and the end-to-end optimization pipeline. Finally, we analyze computational
149 complexity and introduce the three DenseViG variants.

152 PRELIMINARIES
153

154 **Notations** Given an image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, we partition it into N patches. Each patch is projected
155 into a D -dimensional feature $\mathbf{x}_i \in \mathbb{R}^D$ via a learnable linear transformation. Consider a graph
156 $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, where $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ is the set of $N = |\mathcal{V}|$ nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set
157 of edges, and $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \in \mathbb{R}^{N \times D}$ contains the set of node features corresponding to
158 each node v_i . The adjacency matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$ indicates edge connectivity, where $\mathbf{A}_{ij} = 1$ if
159 nodes v_i and v_j are connected, and $\mathbf{A}_{ij} = 0$ otherwise.

160 Graph structure learning seeks to jointly generate a refined adjacency matrix \mathbf{A}^* and corresponding
161 representation \mathbf{H}^* for downstream tasks. Notably, this process solely optimizes the raw structure,
while keeping all others unchanged.

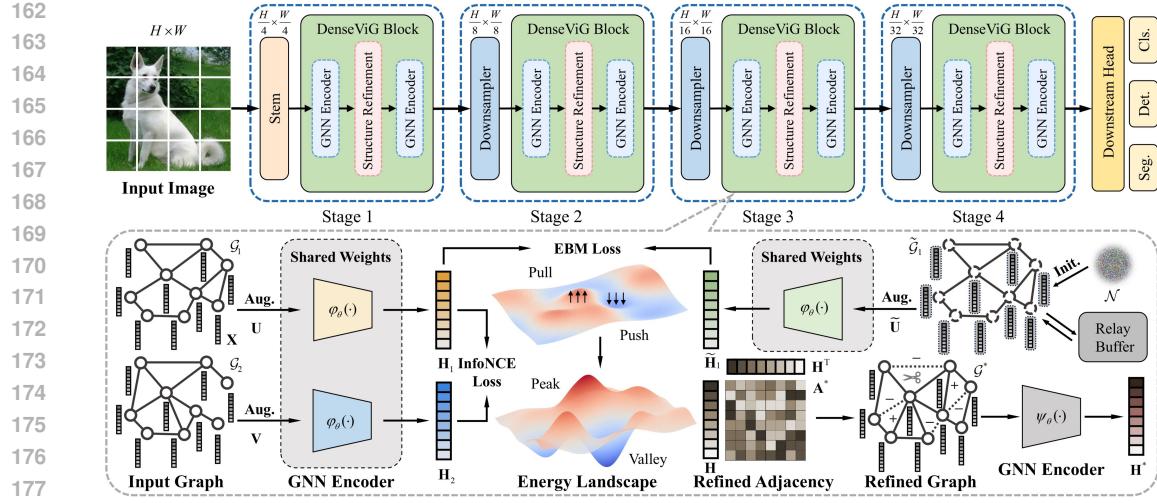


Figure 2: Illustration of the procedure of DenseViG. An image is transformed into node features. Dual graph views feed a shared GNN encoder. DEL improves representations with decoupling EBMs into generative and discriminative aims. Similarities metric and soft relaxation refine structure. Optimal graph boosting downstream task performance in a plug-and-play manner.

Energy-based Models define a probability distribution implicitly by assigning a scalar energy to each data point \mathbf{x} . A parameterized energy function $E_\theta(\mathbf{x}) \in \mathbb{R}$ specifies an unnormalized density using the Boltzmann formulation:

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z(\theta)}, \quad (1)$$

where $Z(\theta) = \int \exp(-E_\theta(\mathbf{x})) d\mathbf{x}$ is the partition function required for normalization.

Since $Z(\theta)$ is usually intractable, training proceeds by calculating the negative log-likelihood of the observed data: $\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [-\log p_\theta(\mathbf{x})]$. A popular optimization method is to minimize Kullback–Leibler (KL) divergence Hinton (2002) via gradient descent: $\mathcal{D}_{\text{KL}}(p_{\text{data}} \| p_\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [E_\theta(\mathbf{x}) + \log Z(\theta)]$, and its objective becomes:

$$\nabla_\theta \mathcal{D}_{\text{KL}}(p_{\text{data}} \| p_\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\nabla_\theta E_\theta(\mathbf{x})] - \mathbb{E}_{\tilde{\mathbf{x}} \sim p_\theta} [\nabla_\theta E_\theta(\tilde{\mathbf{x}})]. \quad (2)$$

where true samples \mathbf{x} map to lower energies, while generated ones $\tilde{\mathbf{x}}$ map to higher energies.

These artificial samples are typically generated by Markov Chain Monte Carlo (MCMC) methods, most often using Langevin dynamics Welling & Teh (2011):

$$\tilde{\mathbf{x}}^{k+1} = \tilde{\mathbf{x}}^k - \lambda \nabla_{\tilde{\mathbf{x}}} E_\theta(\tilde{\mathbf{x}}^k) + \xi^k, \quad (3)$$

where λ is the step size, and ξ^k is i.i.d. standard Gaussian noise $\mathcal{N}(0, 2\lambda)$. With an appropriate step size λ and iterations k , the chain converges to the target distribution p_θ .

DECOPLED ENERGY LEARNING

Let p_d denote the empirical joint data distribution over paired graphs $(\mathcal{G}_1, \mathcal{G}_2)$. To generate two correlated graph views $\mathbf{U}, \mathbf{V} = t_1(\mathbf{X}_1, \mathbf{A}_1), t_2(\mathbf{X}_2, \mathbf{A}_2)$, we independently sample augmentation operators $t_1, t_2 \stackrel{\text{u.a.r.}}{\sim} \mathcal{T}$, where u.a.r. stands for *uniformly at random*. The augmentation set \mathcal{T} consists of random node feature masking and edge perturbation with a fixed ratio Thakoor et al. (2021).

At each ViG stage, a shared-weight GNN encoder φ_θ takes \mathbf{U} and \mathbf{V} as input and returns node representations $\mathbf{H}_1, \mathbf{H}_2 = \varphi_\theta(\mathbf{U}), \varphi_\theta(\mathbf{V}) \in \mathbb{R}^{N \times F}$.

Definition. To approximate p_d , a joint distribution p_θ over the two views (\mathbf{U}, \mathbf{V}) can be defined as an EBM:

$$p_\theta(\mathbf{U}, \mathbf{V}) = \frac{\exp(-E_\theta(\mathbf{U}, \mathbf{V}))}{\int \int \exp(-E_\theta(\mathbf{U}, \mathbf{V})) d\mathbf{U} d\mathbf{V}}, \quad (4)$$

216 where $E_\theta(\mathbf{U}, \mathbf{V}) = \|\mathbf{H}_1 - \mathbf{H}_2\|_2^2 / \tau$ is represented as a metric function that quantifies cross-view
 217 similarity, and τ is a scaling factor.

218 Concretely, $\|\mathbf{H}_1 - \mathbf{H}_2\|_2^2$ indicates the squared ℓ_2 distance between the unit-normalized vectors \mathbf{H}_1
 219 and \mathbf{H}_2 . Our core intuition is to minimize this distance for semantically similar pairs from the same
 220 sample, while maximizing it for dissimilar pairs from distinct samples. Equivalently, p_θ assigns
 221 higher probability mass to matched pairs under p_d and lower mass to mismatched ones.

222 The joint negative log-likelihood of Eq. 4 is given by: $\mathbb{E}_{p_d}[-\log p_\theta(\mathbf{U}, \mathbf{V})]$. According to Eq. 2, its
 223 corresponding gradient can be written as:

$$\nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{U}, \mathbf{V})] = \mathbb{E}_{p_d}[\nabla_\theta E_\theta(\mathbf{U}, \mathbf{V})] - \mathbb{E}_{p_\theta}[\nabla_\theta E_\theta(\mathbf{U}, \mathbf{V})]. \quad (5)$$

227 Eq. 5 requires sampling from $p_\theta(\mathbf{U}, \mathbf{V})$, which is doubly intractable owing to high-dimensional
 228 integrals. Naively applying MCMC is prohibitively expensive for the joint target. To avert this
 229 issue, we employ the product rule to decouple:

$$\nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{U}, \mathbf{V})] = \nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{V}|\mathbf{U})] + \nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{U})], \quad (6)$$

232 where $p_\theta(\mathbf{U})$ denotes the marginal of $p_\theta(\mathbf{U}, \mathbf{V})$.

233 **Theorem.** The marginal distribution $p_\theta(\mathbf{U})$ is also considered as an EBM:

$$p_\theta(\mathbf{U}) = \frac{\exp(-E_\theta(\mathbf{U}))}{\int \exp(-E_\theta(\mathbf{U})) d\mathbf{U}}, \quad (7)$$

238 where $E_\theta(\mathbf{U}) = -\log \int \exp(-E_\theta(\mathbf{U}, \mathbf{V})) d\mathbf{V}$.

239 The negative log-likelihood gradient of Eq. 7 is:

$$\nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{U})] = \mathbb{E}_{p_d}[\nabla_\theta E_\theta(\mathbf{U})] - \mathbb{E}_{p_\theta}[\nabla_\theta E_\theta(\mathbf{U})]. \quad (8)$$

243 Following Eq. 6 and Eq. 8, we obtain the gradient of Decoupled Energy Learning (DEL) objective
 244 $\mathcal{L}_E(\theta)$, incorporating both the discriminative and generative terms:

$$\nabla_\theta \mathcal{L}_E(\theta) = \nabla_\theta \mathbb{E}_{p_d}[-\log p_\theta(\mathbf{V}|\mathbf{U})] + \alpha [\mathbb{E}_{p_d}[\nabla_\theta E_\theta(\mathbf{U})] - \mathbb{E}_{p_\theta}[\nabla_\theta E_\theta(\mathbf{U})]], \quad (9)$$

247 where α is a hyper-parameter to balance the two terms.

248 Eq. 9 allows DEL to indirectly solve the intractable joint distribution. The discriminative term is
 249 straightforward without computing the global gradients w.r.t. $Z(\theta)$, and the generative term only
 250 requires marginal sampling over single view. Negative samples are synthesized via Langevin
 251 dynamics from Gaussian noise initialization, as shown in Eq. 3. Besides, discriminative alignment
 252 alone may eliminate the energy gap between positive and negative samples, resulting in represen-
 253 tation collapse. Generative regularization encourages higher energy for mismatched pairs using
 254 contrastive divergence. This push-pull mechanism accurately aligns the energy landscape with the
 255 underlying data manifold. More proof details are provided in the Appendix A.

257 STRUCTURE REFINEMENT

259 Given node representations $\mathbf{H} = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N\}$, we calculate the cosine-similarity matrix $\mathbf{C}_{ij} =$
 260 $\frac{\mathbf{h}_i \mathbf{h}_j^\top}{\|\mathbf{h}_i\|_2 \|\mathbf{h}_j\|_2}$. To obtain a differentiable adjacency matrix, we adopt a Relaxed Bernoulli variable for
 261 each edge via the Gumbel–Sigmoid reparameterization Jang et al. (2016):

$$\mathbf{S} = \text{Sigmoid}\left(\frac{1}{\pi} \left(\log \frac{\mathbf{C}}{1 - \mathbf{C}} + \log \frac{\delta}{1 - \delta}\right)\right), \quad (10)$$

265 where $\delta \sim \text{Uniform}(0, 1)$, and $\pi \in \mathbb{R}^+$ controls the relaxation sharpness. As $\pi \rightarrow 0$, \mathbf{S} converges to
 266 a hard binary decision. For structural validity and numerical stability, we symmetrize and normalize
 267 \mathbf{S} : $\mathbf{A}^* = \mathbf{D}^{-\frac{1}{2}} \left(\frac{\mathbf{S} + \mathbf{S}^\top}{2}\right) \mathbf{D}^{-\frac{1}{2}}$, where \mathbf{D} is the diagonal degree matrix. The refined graph $\mathcal{G}^* =$
 268 $(\mathbf{X}, \mathbf{A}^*)$ is passed to a GNN encoder ψ_θ , architecturally identical to φ_θ , to yield the representation
 269 \mathbf{H}^* for the subsequent processing.

270 END-TO-END OPTIMIZATION
271

272 For the implementation of DEL, we design a tractable MCMC estimator in which the generative
273 term is rewritten as contrastive form, and the discriminative term is expressed as an InfoNCE loss.

274 Given a mini-batch of N paired augmentations $\{\mathbf{U}_i, \mathbf{V}_i\}_{i=1}^N \sim p_d$, each $(\mathbf{U}_i, \mathbf{V}_i)$ forms a positive
275 pair, while $(\mathbf{U}_i, \mathbf{V}_j)$ serve as negative ones for the discriminative term. To balance the two phases
276 of the generative term in KL divergence, we synthesize N views $\{\tilde{\mathbf{U}}_i\}_{i=1}^N \sim p_\theta$ using k Langevin
277 dynamics steps in the negative phase.

278 Since the marginal energy $E_\theta(\mathbf{U}) = -\log \int \exp(-E_\theta(\mathbf{U}, \mathbf{V})) d\mathbf{V}$ is intractable, we approximate
279 it with a log-sum-exp over M complementary views $\{\mathbf{V}_{i,m}\}_{m=1}^M$ from the same batch: $\hat{E}_\theta(\mathbf{U}_i) \approx$
280 $-\log \frac{1}{M} \sum_{m=1}^M \exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_{i,m}))$. Substituting this into Eq. 9 gives the generative loss:

$$282 \quad \mathcal{L}_{\text{Gen}}(\theta) \approx \frac{1}{N} \sum_{i=1}^N (\hat{E}_\theta(\mathbf{U}_i) - \hat{E}_\theta(\tilde{\mathbf{U}}_i)). \quad (11)$$

285 Treating each positive pair $(\mathbf{U}_i, \mathbf{V}_i^+)$ against M' intra-batch negatives $\{\mathbf{V}_{i,j}^-\}_{j=1}^{M'}$ yields the dis-
286 criminative loss:

$$288 \quad \mathcal{L}_{\text{Disc}}(\theta) \approx -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_i^+))}{\exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_i^+)) + \sum_{j=1}^{M'} \exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_{i,j}^-))}. \quad (12)$$

291 In summary, the complete DEL objective is: $\mathcal{L}_E(\theta) = \mathcal{L}_{\text{Disc}}(\theta) + \alpha \mathcal{L}_{\text{Gen}}(\theta) + \beta \mathcal{L}_r(\theta)$, where
292 \mathcal{L}_r is the ℓ_2 regularization that prevents gradient overflow from large energy values. For end-to-
293 end training, we minimize the global loss function: $\mathcal{L}_G = \mathcal{L}_T + \mu \mathcal{L}_E + \gamma \mathcal{L}_R$, where \mathcal{L}_T is the
294 task-specific loss, $\mathcal{L}_R = \sum_{i,j} \mathbb{E}_\delta[\mathbf{S}_{ij}]$ is a sparsity-promoting penalty term, μ and γ are hyper-
295 parameters. The pseudocode of DenseViG is illustrated in Algorithm 1 in the Appendix C.1.

296 ARCHITECTURE OVERVIEW
297

298 As illustrated in Fig. 2, DenseViG is a four-stage architecture that enriches a PyramidViG backbone
299 Han et al. (2022) with DEL-guided structure refinement. We introduce the model’s scaling variants
300 and provide a brief complexity analysis. The complete configuration is in the Appendix C.2.

301 **Scaling Variants** Mirroring the design
302 philosophy of ViG, we instantiate three
303 different configurations: DenseViG-Ti,
304 DenseViG-S, and DenseViG-M. Across
305 all variants, the initial node number is
306 fixed to $N = 3136$, while later stages re-
307 duce it via spatial downsampling. We set 9
308 nearest neighbors in the k -NN graph con-
309 struction, and employ 4 attention heads.

310 Table 1 summarizes the depth, blocks per stage, and dimensionality of each model.

311 **Complexity Analysis** DenseViG introduces additional computation in only two modules: structure
312 refinement computes B cosine-similarity matrices of size $N \times N$, with a cost of $\mathcal{O}(B \cdot N^2 \cdot F)$, where
313 F is the representation dimension. Since ViG already condenses an $H \times W$ image into N nodes, this
314 quadratic term is empirically affordable. DEL performs k Langevin steps on the B graph views in a
315 mini-batch, adding a linear overhead of $\mathcal{O}(k \cdot B \cdot N \cdot F)$. With a small k and F , this cost is negligible
316 relative to structure refinement. Therefore, the total overhead is $\mathcal{O}(B \cdot N^2 \cdot F + k \cdot B \cdot N \cdot F)$, and
317 the overall pipeline retains near-linear scalability w.r.t. image size.

Model	Depth	Blocks	Dim. d_D	Dim. d_F
DenseViG-Ti	12	[2, 2, 6, 2]	192	[48, 96, 240, 384]
DenseViG-S	12	[2, 2, 6, 2]	320	[80, 160, 400, 640]
DenseViG-M	22	[2, 2, 16, 2]	640	[96, 192, 384, 768]

321 EXPERIMENTS
322323 IMAGE CLASSIFICATION
324

325 **Implementation Details** All DenseViG variants are trained on ImageNet-1K Deng et al. (2009),
326 which comprises 1.2M training images, 50K validation images, and 100K test images across 1K

324 Table 2: Image classification performance on ImageNet-1K.  CNN,  ViT,  Pool,  GNN,  CNN-GNN,  HGNN, and  DenseViG (ours). Reporting Params (M), FLOPs (G), and Top-1 & Top-5 Accuracy (%) where available.

Model	Params (M)	FLOPs (G)	Top-1	Top-5	Model	Params (M)	FLOPs (G)	Top-1	Top-5
 ResNet18	11.7	1.8	70.6	89.7	 PoolFormer-S12	12.0	2.0	77.2	93.5
 ResNet50	25.6	4.1	79.8	95.0	 PoolFormer-S36	31.0	5.2	81.4	95.5
 ResNet152	60.2	11.5	81.8	95.9	 PoolFormer-M48	73.0	11.9	82.5	96.0
 ConvNeXt-T	29.0	4.5	82.1	95.9	 MobileViG-Ti	5.2	0.7	75.7	-
 ConvNeXt-S	50.0	8.7	83.1	96.4	 MobileViG-S	7.2	1.0	78.2	-
 ConvNeXt-B	89.0	15.4	83.8	96.7	 MobileViG-M	14.0	1.5	80.6	-
 PVT-Ti	13.2	1.9	75.1	-	 GreedyViG-S	12.0	1.6	81.1	-
 PVT-S	24.5	3.8	79.8	-	 GreedyViG-M	21.9	3.2	82.9	-
 PVT-M	44.2	6.7	81.2	-	 GreedyViG-B	30.9	5.2	83.9	-
 DeiT-Ti	5.7	1.3	72.2	80.1	 PViHGNN-Ti	12.3	2.3	78.9	94.6
 DeiT-S	22.1	4.6	79.8	85.7	 PViHGNN-S	28.5	6.3	82.5	96.3
 DeiT-B	86.4	17.6	81.8	86.7	 PViHGNN-M	52.4	10.7	83.4	96.5
 Swin-T	29.0	4.5	81.3	95.5	 DVHGNN-T	11.1	1.9	79.8	-
 Swin-S	50.0	8.7	83.0	96.2	 DVHGNN-S	30.2	5.2	83.1	-
 Swin-B	88.0	15.4	83.5	96.5	 DVHGNN-M	52.5	10.4	83.8	-
 PyramidViG-Ti	10.7	1.7	78.2	94.2	 HgVT-Ti	7.7	1.8	76.2	93.2
 PyramidViG-S	27.3	4.6	82.1	96.0	 HgVT-S	22.9	5.5	81.2	95.5
 PyramidViG-M	51.7	8.9	83.1	96.4	 DenseViG-Ti	11.1	1.9	82.2	96.1
 WiGNet-Ti	10.8	2.1	78.8	94.6	 DenseViG-S	28.5	5.3	83.7	96.7
 WiGNet-S	27.4	5.7	82.0	95.9	 DenseViG-M	53.3	9.6	84.3	96.9
 WiGNet-M	49.7	11.2	83.0	96.3					

347
348 classes. Training is conducted for 310 epochs on 224×224 crops with the AdamW optimizer (momentum $\beta_1=0.9$, $\beta_2=0.99$, weight decay=0.05). The learning rate is initialized at 0.001 and annealed by a cosine scheduler. We apply RandAugment, CutMix, RandomResizedCrop, and label smoothing for data augmentation, and track an Exponential Moving Average (EMA) of model parameters. All experiments are implemented in PyTorch and run on 8 RTX3090 GPUs with a batch size of 64. The hyperparameters are specified in the Appendix C.3. to ensure the reproducibility.

349
350
351
352
353
354 **Experimental Results** Table 2 compares DenseViG with leading image classification methods, including ResNet He et al. (2016), ConvNeXt Liu et al. (2022d), PVT Wang et al. (2021b), DeiT Touvron et al. (2021), Swin Liu et al. (2022c), PoolFormer Yu et al. (2022), PyramidViG Han et al. (2022), WiGNet Spadaro et al. (2025), MobileViG Munir et al. (2023), GreedyViG Munir et al. (2024), PViHGNN Han et al. (2023), DVHGNN Li et al. (2025), and HgVT Fixelle (2025). Notably, DenseViG-S boosts Top-1 accuracy by 1.7% over the strongest GNN baseline WiGNet-S and by 0.6% over the best HGNN baseline DVHGNN-S. These margins confirm that DenseViG captures higher-order visual information interactions more effectively than prior approaches, without appreciably increasing model size or computational cost.

363 364 OBJECT DETECTION

365
366 **Implementation Details** Object detection is conducted on MS COCO Lin et al. (2014), containing 118K training, 5K validation, and 20K test images. RetinaNet Lin et al. (2017) and Mask R-CNN He et al. (2017) are employed as detection heads and implemented in the MMDetection Chen et al. (2019). All models are pre-trained on ImageNet-1K and fine-tuned for 12 epochs at the resolution of 1280×800 with an AdamW optimizer, an initial learning rate of 2×10^{-4} , and an effective batch size of 16. We additionally apply data augmentations including random horizontal Flipping and Multi-scale jittering. Other settings are the same as the image classification setup. Bounding box regression uses a weighted sum of ℓ_1 and GIoU losses.

367
368
369
370
371
372 **Experimental Results** As depicted in Table 3, DenseViG outperforms the state-of-the-art approaches. Specifically, under the RetinaNet head, DenseViG-S achieves 44.6% mAP that surpasses PyramidViG by 2.8%, ViHGNN-S by 2.4%, and PViHGNN-S by 1.3%, respectively. Under the Mask R-CNN head, DenseViG-S achieves 46.4% mAP that surpasses PyramidViG by 3.8%, ViHGNN-S by 3.3%, and PViHGNN-S by 1.6%, respectively. More results are in the Appendix D.

378 Table 3: Object detection performance on MS Table 4: Semantic segmentation performance on
 379 COCO. $1\times$ denotes the standard training sched- ADE20k. Head method includes input configura-
 380 ule of 12 epochs. Indicating Params (M), tion: -1 last backbone layer only, -4 last four lay-
 381 FLOPs (G), and mAP (%) where available. 382
 382 FLOPs (G), and mAP (%) where available.

Backbone	RetinaNet $1\times$				
	Params	FLOPs	AP _{box}	AP ₅₀	AP ₇₅
ResNet-50	38M	239G	36.3	55.3	38.6
PVT-S	34M	227G	40.4	61.3	44.2
Swin-T	39M	245G	41.5	62.1	44.2
PyramidViG-S	36M	240G	41.8	63.1	44.7
PViHGNN-S	38M	244G	42.2	63.8	45.1
DVHGNN-S	38M	242G	43.3	64.3	46.3
DenseViG-S	40M	247G	44.6	66.1	47.8

Backbone	Mask R-CNN $1\times$				
	Params	FLOPs	AP _{box}	AP ₅₀	AP ₇₅
ResNet-50	44M	260G	38.0	58.6	41.4
PVT-S	44M	245G	40.4	62.9	43.8
Swin-T	48M	264G	42.2	64.6	46.2
PyramidViG-S	46M	259G	42.6	65.2	46.0
PViHGNN-S	48M	262G	43.1	66.0	46.5
DVHGNN-S	49M	261G	44.8	66.8	49.0
DenseViG-S	51M	266G	46.4	68.7	50.6

Backbone	Params	Head	Params	mIoU	Acc.
ResNet-18	11.5M	PPM-1	12.9M	33.8	76.1
ResNet-50	25.6M	SFPN	1.6M	36.7	-
ResNet-50	25.6M	PPM-1	23.2M	41.3	79.7
ResNet-101	44.5M	PPM-1	23.2M	42.2	80.6
ConvNeXt-T	29.0M	UperNet	29.8M	46.1	-
PVT-S	24.5M	SFPN	1.6M	39.8	-
PoolFormer-S12	12.0M	SFPN	1.6M	37.2	-
Swin-T	28.3M	SFPN	1.6M	41.5	-
Swin-T	28.3M	UperNet	29.8M	46.1	-
GreedyViG-S	12.0M	SFPN	1.6M	43.2	-
GreedyViG-B	30.9M	SFPN	1.6M	47.4	-
HgVT-S	22.9M	MLP-4	235K	28.5	71.8
HgVT-S	22.9M	PPM-4	15.5M	36.0	75.7
DenseViG-Ti	11.1M	MLP-4	235K	33.0	76.8
DenseViG-Ti	11.1M	SFPN	1.6M	42.7	79.4
DenseViG-S	28.5M	PPM-4	15.5M	48.8	81.7
DenseViG-S	28.5M	UperNet	29.8M	50.9	82.5

402 SEMANTIC SEGMENTATION

404 **Implementation Details** We validate semantic segmentation on ADE20K Zhou et al. (2017), which
 405 provides 20K training, 2K validation, and 3K test images covering 150 categories. The backbone
 406 is initialized with ImageNet-1K weights and fine-tuned for 16 epochs with 4 segmentation heads:
 407 MLP, PPM Zhao et al. (2017), Semantic FPN Kirillov et al. (2019), and UPerNet Xiao et al. (2018).
 408 All models are implemented in the MMSegmentation Contributors (2020). Training uses 512×512
 409 crops with a batch size of 16, and the same data augmentation pipeline, GPU configuration, opti-
 410 mizer settings, and learning rate schedule as in our object detection setup. Supervision is provided
 411 by the standard pixel-wise cross-entropy loss.

412 **Experimental Results** Table 4 presents the superiority of DenseViG with other famous backbones.
 413 Specifically, using lightweight linear heads such as MLP and PPM, DenseViG-Ti achieves competi-
 414 tive results with faster inference speed. With SFPN and UPerNet head, DenseViG-S achieves 48.8%
 415 and 50.9% mIoU, outperforming the suboptimal GreedyViG-B by 1.4% and 3.5%, respectively.

417 ABLATION STUDIES

418 **Effectiveness of DEL Objective** Table 5 investigates the influence of
 419 generative and discriminative terms in DEL on classification perfor-
 420 mance. Using only the generative term improves Top-1 accuracy by
 421 0.5-3.1% over the baseline, suggesting EBM_s enhance representation
 422 quality. Replacing EBM_s with InfoNCE independently achieves accu-
 423 precacies of 81.7% and 83.0% on ViG-Ti
 424 and ViG-S, demonstrating the effec-
 425 tiveness of the discriminative term.
 426 Involving EBM_s with SimCLR yields
 427 a 0.6% gain, while pairing EBM_s
 428

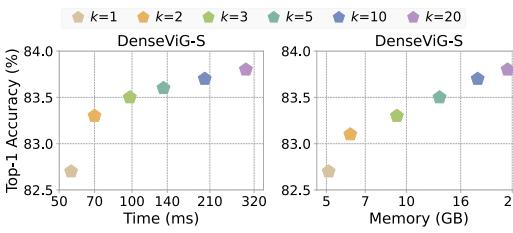
429 Table 5: Performance ablations of generative and discriminative terms in DEL on ImageNet-1K using the DenseViG-
 430 Ti/S backbone. Denoting used (✓) or not used (✗).

Method	Gen. term		Disc. term		Top-1
	EBMs	InfoNCE	SimCLR		
PyramidViG-Ti	✗	✗	✗		78.2
DenseViG-Ti	✓	✗	✗		81.3
DenseViG-Ti	✗	✓	✗		81.7
DenseViG-Ti	✓	✗	✓		82.4
DenseViG-Ti	✓	✓	✗		82.2
PyramidViG-S	✗	✗	✗		82.1
DenseViG-S	✓	✗	✗		82.6
DenseViG-S	✗	✓	✗		83.0
DenseViG-S	✓	✗	✓		83.6
DenseViG-S	✓	✓	✗		83.7

432 with InfoNCE obtains the optimal balance of 82.2% and 83.7%. These findings reinforce our view
 433 that standard contrastive loss represents a special case of DEL.
 434

435 **Effectiveness of Structure Refinement** Table 6 presents the impact of
 436 varying refinement strategy. Static k -NN and the axial graph score 82.1%
 437 and 81.1%. With $\pi=0.1$, adding
 438 edges only lifts accuracy to 83.3%, while pruning edges only reaches
 439 82.9%, proving that each step alone
 440 reduces noise or restores missing
 441 links. Applying both adding and
 442 pruning at every stage pushes accu-
 443 racy to 83.7%, surpassing the half-
 444 stage schedule and the milder relax-
 445 ation. Full-stage refinement with a softer relaxation therefore yields the best result.
 446

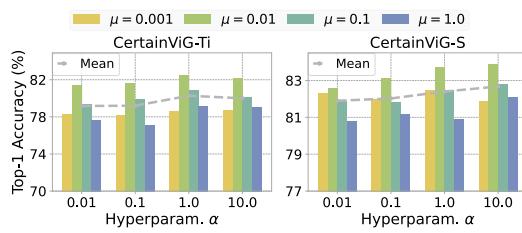
447 **Analysis of Iterative Steps** Fig. 3 shows that raising k from 1 to 5 lifts Top 1 accuracy by 1.0%,
 448 adding 79 ms and 8.22 GB. Increasing to 10 or 20 yields little extra accuracy but much higher cost.
 449 We choose $k=5$, which provides near peak accuracy with acceptable time and memory, confirming
 450 the scalability of our EBM sampler on current environments.
 451



452 Figure 3: Impact of k Langevin steps in nega-
 453 tive sampling for DenseViG-S on ImageNet-1K.
 454 Evaluated on 8 RTX 3090 GPUs (batch size 64).
 455

456 Table 6: Impact of various structure refinement methods on
 457 ImageNet-1K. ‘Half’ means placing at every second stage.
 458 The temperature π controls relaxation sharpness.

Model	Refinement	Relaxation	Stage	Top-1
PyramidViG-S	Static k -NN	Hard	All	82.1
GreedyViG-S	Axial graph	Hard	All	81.1
DenseViG-S	Add-only	$\pi=0.1$	All	83.3
DenseViG-S	Prune-only	$\pi=0.1$	All	82.9
DenseViG-S	Add-prune	$\pi=0.1$	Half	83.1
DenseViG-S	Add-prune	$\pi=1.0$	All	83.5
DenseViG-S	Add-prune	$\pi=0.1$	All	83.7



459 Figure 4: Hyper-parameter α and μ analysis of
 460 DenseViG-Ti and DenseViG-S on ImageNet-1K.
 461 ‘Mean’ reports the average of the four columns.
 462

463 **Setting of Hyper-parameters** Fig. 4
 464 indicates that raising α from 0.01 to
 465 10.0 significantly improves Top 1 accuracy of DenseViG-Ti and S. Perfor-
 466 mance peaks at $\alpha=10.0$, yet generated
 467 samples reduce diversity, which means
 468 an excessive discriminative weight
 469 weakens the EBM’s generative power.
 470 We therefore set $\alpha=1.0$ for a trade-off
 471 outcome. Increasing μ from 0.001 to 0.1 brings rapid gains and confirms the benefit of the Dense
 472 framework. Above 0.1 the improvement levels off and training slows, so we fix $\mu=0.1$ for all models.
 473

474 **Plug-and-play Verification** Table 7 demonstrates a 1.5% Top-1 gain on WiGNet-S and 1.7% on
 475 GreedyViG-S. Added parameters and FLOPs are minimal, confirming that Dense framework can be
 476 seamlessly plugged into existing backbones and boosts performance with minimal extra cost.
 477

478 CONCLUSION

480 This paper proposes DenseViG, a novel end-to-end framework that decouples EBMs into generative
 481 and discriminative aims to enable differentiable graph structure refinement for enhancing ViG. At
 482 each stage, DenseViG dynamically adds or prunes edges between visual patches, yielding cleaner
 483 and higher-quality representations while suppressing noise. The resulting sparse graphs significantly
 484 decrease the computational and memory costs of EBMs on high-resolution images. Extensive exper-
 485 iments on three vision benchmarks show that DenseViG sets new state-of-the-art results, validating
 the effectiveness of decoupled energy-guided graph structure refinement.

486 ETHICS STATEMENT
487488 All authors have read and agree to follow the ICLR Code of Ethics. This work aims to advance
489 representation learning for the public good. We use only public data and do not process personal
490 data. We respect licenses and uphold fairness, privacy, security, and research integrity.
491492 REPRODUCIBILITY STATEMENT
493494 We will release the code after acceptance. Our experiments use public data and standard settings.
495 We will share scripts, configuration files, and clear steps so others can reproduce the results. All
496 resources will have transparent sources.
497498 REFERENCES
499500 Michael Arbel, Liang Zhou, and Arthur Gretton. Generalized energy based models. *arXiv preprint*
501 *arXiv:2003.05033*, 2020.
502503 Aritra Bhowmik, Yu Wang, Nora Baka, Martin R Oswald, and Cees GM Snoek. Detecting objects
504 with context-likelihood graphs and graph refinement. In *Proceedings of the IEEE/CVF Interna-*
505 *tional Conference on Computer Vision*, pp. 6524–6533, 2023.506 Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen
507 Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark.
508 *arXiv preprint arXiv:1906.07155*, 2019.
509510 MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox
511 and benchmark. <https://github.com/open-mmlab/mmsegmentation>, 2020.512 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
513 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
514 pp. 248–255. Ieee, 2009.515 Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances*
516 *in neural information processing systems*, 32, 2019.
517518 Joshua Fixelle. Hypergraph vision transformers: Images are more than nodes, more than edges. In
519 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 9751–9761, 2025.
520521 Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi,
522 and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like
523 one. *arXiv preprint arXiv:1912.03263*, 2019.524 Fredrik K Gustafsson, Martin Danelljan, and Thomas B Schon. Accurate 3d object detection using
525 energy-based models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
526 *Pattern Recognition*, pp. 2855–2864, 2021.527 Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle
528 for unnormalized statistical models. In *Proceedings of the thirteenth international conference on*
529 *artificial intelligence and statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings,
530 2010.
531532 Kai Han, Yunhe Wang, Jianyuan Guo, Yehui Tang, and Enhua Wu. Vision gnn: An image is worth
533 a graph of nodes. *Advances in neural information processing systems*, 35:8291–8303, 2022.534 Yan Han, Peihao Wang, Souvik Kundu, Ying Ding, and Zhangyang Wang. Vision hgnn: An image
535 is more than a graph of nodes. In *Proceedings of the IEEE/CVF International Conference on*
536 *Computer Vision*, pp. 19878–19888, 2023.
537538 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
539 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
770–778, 2016.

540 Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the*
 541 *IEEE international conference on computer vision*, pp. 2961–2969, 2017.
 542

543 Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural*
 544 *computation*, 14(8):1771–1800, 2002.

545 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv*
 546 *preprint arXiv:1611.01144*, 2016.
 547

548 Juntao Jiang, Xiyu Chen, Guanzhong Tian, and Yong Liu. Vig-unet: vision graph neural networks
 549 for medical image segmentation. In *2023 IEEE 20th International Symposium on Biomedical*
 550 *Imaging (ISBI)*, pp. 1–5. IEEE, 2023.

551 Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph structure
 552 learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD international*
 553 *conference on knowledge discovery & data mining*, pp. 66–74, 2020.

554 Beomsu Kim and Jong Chul Ye. Energy-based contrastive learning of visual representations. *Ad-*
 555 *vances in Neural Information Processing Systems*, 35:4358–4369, 2022.

556 Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid net-
 557 *works*. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 558 pp. 6399–6408, 2019.

559 Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, Fujie Huang, et al. A tutorial on energy-
 560 based learning. *Predicting structured data*, 1(0), 2006.

561 Caoshuo Li, Tanzhe Li, Xiaobin Hu, Donghao Luo, and Taisong Jin. Dvhggn: Multi-scale dilated
 562 vision hggn for efficient vision recognition. In *Proceedings of the Computer Vision and Pattern*
 563 *Recognition Conference*, pp. 20158–20168, 2025.

564 Kuan Li, Yang Liu, Xiang Ao, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. Reliable
 565 representations make a stronger defender: Unsupervised structure refinement for robust gnn. In
 566 *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pp.
 567 925–935, 2022.

568 Tanzhe Li, Wei Lin, Xiawu Zheng, and Taisong Jin. Glvig: Global and local vision gnn may be
 569 what you need for vision. In *Chinese Conference on Pattern Recognition and Computer Vision*
 570 (*PRCV*), pp. 368–382. Springer, 2023.

571 Xinhong Li, Geng Chen, Yuanfeng Wu, Junqing Yang, Tao Zhou, Yi Zhou, and Wentao Zhu. Med-
 572 segvig: Medical image segmentation with a vision graph neural network. In *2024 IEEE Interna-*
 573 *tional Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 3408–3411. IEEE, 2024.

574 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 575 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European*
 576 *conference on computer vision*, pp. 740–755. Springer, 2014.

577 Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense
 578 object detection. In *Proceedings of the IEEE international conference on computer vision*, pp.
 579 2980–2988, 2017.

580 Nian Liu, Xiao Wang, Lingfei Wu, Yu Chen, Xiaojie Guo, and Chuan Shi. Compact graph structure
 581 learning via mutual information compression. In *Proceedings of the ACM web conference 2022*,
 582 pp. 1601–1610, 2022a.

583 Yixin Liu, Yu Zheng, Daokun Zhang, Hongxu Chen, Hao Peng, and Shirui Pan. Towards unsu-
 584 *pervised deep graph structure learning*. In *Proceedings of the ACM Web Conference 2022*, pp.
 585 1392–1403, 2022b.

586 Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng
 587 Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of*
 588 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12009–12019, 2022c.

594 Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie.
 595 A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and*
 596 *pattern recognition*, pp. 11976–11986, 2022d.

597

598 Mustafa Munir, William Avery, and Radu Marculescu. Mobilevig: Graph-based sparse attention for
 599 mobile vision applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 600 *Pattern Recognition*, pp. 2211–2219, 2023.

601 Mustafa Munir, William Avery, Md Mostafijur Rahman, and Radu Marculescu. Greedyvig: Dy-
 602 namic axial graph construction for efficient vision gnns. In *Proceedings of the IEEE/CVF Con-*
 603 *ference on Computer Vision and Pattern Recognition*, pp. 6118–6127, 2024.

604

605 Bo Pang, Tian Han, Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. Learning latent space
 606 energy-based prior model. *Advances in Neural Information Processing Systems*, 33:21994–
 607 22008, 2020.

608 Dhruv Parikh, Jacob Fein-Ashley, Tian Ye, Rajgopal Kannan, and Viktor Prasanna. Clustervig:
 609 Efficient globally aware vision gnns via image partitioning. *arXiv preprint arXiv:2501.10640*,
 610 2025.

611 Tianhao Peng, Xuhong Li, Haitao Yuan, Yuchen Li, and Haoyi Xiong. Sola-gcl: Subgraph-oriented
 612 learnable augmentation method for graph contrastive learning. In *Proceedings of the AAAI Con-*
 613 *ference on Artificial Intelligence*, volume 39, pp. 19875–19883, 2025.

614

615 Md Mostafijur Rahman and Radu Marculescu. G-cascade: Efficient cascaded graph convolutional
 616 decoding for 2d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference*
 617 *on Applications of Computer Vision*, pp. 7728–7737, 2024.

618 Gabriele Spadaro, Marco Grangetto, Attilio Fiandrotti, Enzo Tartaglione, and Jhony H Giraldo.
 619 Wignet: Windowed vision graph neural network. In *2025 IEEE/CVF Winter Conference on Ap-*
 620 *plications of Computer Vision (WACV)*, pp. 859–868. IEEE, 2025.

621

622 Shantanu Thakoor, Corentin Tallec, Mohammad Gheshlaghi Azar, Mehdi Azabou, Eva L Dyer,
 623 Remi Munos, Petar Veličković, and Michal Valko. Large-scale representation learning on graphs
 624 via bootstrapping. *arXiv preprint arXiv:2102.06514*, 2021.

625 Yuchuan Tian, Hanting Chen, Chao Xu, and Yunhe Wang. Image processing gnn: Breaking rigidity
 626 in super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
 627 *recognition*, pp. 24108–24117, 2024.

628

629 Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and
 630 Hervé Jégou. Training data-efficient image transformers & distillation through attention. In
 631 *International conference on machine learning*, pp. 10347–10357. PMLR, 2021.

632

633 Hao Wang, Shuo Zhang, and Biao Leng. Hgformer: Topology-aware vision transformer with hy-
 634 *pergraph learning. IEEE Transactions on Multimedia*, 2025.

635

636 Ruijia Wang, Shuai Mou, Xiao Wang, Wanpeng Xiao, Qi Ju, Chuan Shi, and Xing Xie. Graph
 637 structure estimation neural networks. In *Proceedings of the web conference 2021*, pp. 342–353,
 2021a.

638

639 Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo,
 640 and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without
 641 convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.
 568–578, 2021b.

642

643 Yifei Wang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. A unified contrastive energy-
 644 based model for understanding the generative ability of adversarial training. *arXiv preprint*
 645 *arXiv:2203.13455*, 2022.

646

647 Yuan Wang, Kun Yu, Chen Chen, Xiyuan Hu, and Silong Peng. Dynamic graph learning with
 content-guided spatial-frequency relation reasoning for deepfake detection. In *Proceedings of the*
IEEE/CVF conference on computer vision and pattern recognition, pp. 7278–7287, 2023.

648 Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In
 649 *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688,
 650 2011.

651

652 JiaFu Wu, Jian Li, Jiangning Zhang, Boshen Zhang, Mingmin Chi, Yabiao Wang, and Chengjie
 653 Wang. PvG: Progressive vision graph for vision recognition. In *Proceedings of the 31st ACM*
 654 *international conference on multimedia*, pp. 2477–2486, 2023.

655

656 Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. Nodeformer: A scalable graph
 657 structure learning transformer for node classification. *Advances in Neural Information Processing*
 658 *Systems*, 35:27387–27401, 2022.

659

660 Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for
 661 scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*, pp.
 662 418–434, 2018.

663

664 Yankun Xu, Junzhe Wang, Yun-Hsuan Chen, Jie Yang, Wenjie Ming, Shuang Wang, and Mohamad
 665 Sawan. Vsvig: Real-time video-based seizure detection via skeleton-based spatiotemporal vig. In
 666 *European Conference on Computer Vision*, pp. 228–245. Springer, 2024.

667

668 Kunshan Yang, Lin Zuo, Mengmeng Jing, Xianlong Tian, Kunbin He, and Yongqi Ding. Flexible
 669 vig: Learning the self-saliency for flexible object recognition. *IEEE Transactions on Circuits and*
 670 *Systems for Video Technology*, 2025.

671

672 Ruijie Yao, Sheng Jin, Lumin Xu, Wang Zeng, Wentao Liu, Chen Qian, Ping Luo, and Ji Wu.
 673 Gkgnet: Group k-nearest neighbor based graph convolutional network for multi-label image
 674 recognition. In *European Conference on Computer Vision*, pp. 91–107. Springer, 2024.

675

676 Donghan Yu, Ruohong Zhang, Zhengbao Jiang, Yuexin Wu, and Yiming Yang. Graph-revised con-
 677 volutional network. In *Joint European conference on machine learning and knowledge discovery*
 678 *in databases*, pp. 378–393. Springer, 2020.

679

680 Junliang Yu, Xin Xia, Tong Chen, Lizhen Cui, Nguyen Quoc Viet Hung, and Hongzhi Yin. Xsimgcl:
 681 Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on*
 682 *Knowledge and Data Engineering*, 36(2):913–926, 2023a.

683

684 Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and
 685 Shuicheng Yan. Metaformer is actually what you need for vision. In *Proceedings of the IEEE/CVF*
 686 *conference on computer vision and pattern recognition*, pp. 10819–10829, 2022.

687

688 Yue Yu, Xiao Wang, Mengmei Zhang, Nian Liu, and Chuan Shi. Provable training for graph con-
 689 trastive learning. *Advances in Neural Information Processing Systems*, 36:50327–50345, 2023b.

690

691 Xianlin Zeng, Yufeng Wang, Yuqi Sun, Guodong Guo, Wenrui Ding, and Baochang Zhang. Graph
 692 structure refinement with energy-based contrastive learning. In *Proceedings of the AAAI Confer-
 693 ence on Artificial Intelligence*, volume 39, pp. 22326–22335, 2025.

694

695 Jianrong Zhang, Hehe Fan, and Yi Yang. Energymogen: Compositional human motion generation
 696 with energy-based diffusion model in latent space. In *Proceedings of the Computer Vision and*
 697 *Pattern Recognition Conference*, pp. 17592–17602, 2025.

698

699 Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing
 700 network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
 701 2881–2890, 2017.

702

703 Jianan Zhao, Qianlong Wen, Mingxuan Ju, Chuxu Zhang, and Yanfang Ye. Self-supervised graph
 704 structure refinement for graph neural networks. In *Proceedings of the sixteenth ACM international*
 705 *conference on web search and data mining*, pp. 159–167, 2023.

706

707 Tong Zhao, Yozen Liu, Leonardo Neves, Oliver Woodford, Meng Jiang, and Neil Shah. Data aug-
 708 mentation for graph neural networks. In *Proceedings of the aaai conference on artificial intelli-
 709 gence*, volume 35, pp. 11015–11023, 2021.

702 Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene
703 parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and*
704 *pattern recognition*, pp. 633–641, 2017.

705 Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Deep graph contrastive
706 representation learning. *arXiv preprint arXiv:2006.04131*, 2020.

708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

APPENDIX

A MISSING PROOFS AND DERIVATIONS

A.1 REFORMULATION OF THE MARGINAL EBM OBJECTIVE AS A GENERATIVE LOSS

Starting from Eq. 1 and the marginal probability $p_\theta(\mathbf{U})$, we have

$$\begin{aligned}
 \nabla_\theta \log p_\theta(\mathbf{U}) &= \nabla_\theta \log \frac{\exp(-E_\theta(\mathbf{U}))}{Z_\theta} \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) - \nabla_\theta \log Z_\theta \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) - \nabla_\theta \log \sum_{\mathbf{U}} \exp(-E_\theta(\mathbf{U})) \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) - \frac{\sum_{\mathbf{U}} \nabla_\theta \exp(-E_\theta(\mathbf{U}))}{\sum_{\mathbf{U}'} \exp(-E_\theta(\mathbf{U}'))} \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) + \sum_{\mathbf{U}} \frac{\exp(-E_\theta(\mathbf{U}))}{\sum_{\mathbf{U}'} \exp(-E_\theta(\mathbf{U}'))} \nabla_\theta E_\theta(\mathbf{U}) \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) + \sum_{\mathbf{U}} p_\theta(\mathbf{U}) \nabla_\theta E_\theta(\mathbf{U}) \\
 &= -\nabla_\theta E_\theta(\mathbf{U}) + \mathbb{E}_{\tilde{\mathbf{U}} \sim p_\theta} [\nabla_\theta E_\theta(\tilde{\mathbf{U}})],
 \end{aligned}$$

where we use $\nabla_\theta \log Z_\theta = -\mathbb{E}_{\tilde{\mathbf{U}} \sim p_\theta} [\nabla_\theta E_\theta(\tilde{\mathbf{U}})]$. Therefore, the gradient of the negative log-likelihood \mathcal{L}_θ under the data distribution p_d is

$$\nabla_\theta \mathcal{L}_\theta = -\nabla_\theta \mathbb{E}_{\mathbf{U} \sim p_d} [\log p_\theta(\mathbf{U})] = \mathbb{E}_{\mathbf{U} \sim p_d} [\nabla_\theta E_\theta(\mathbf{U})] - \mathbb{E}_{\tilde{\mathbf{U}} \sim p_\theta} [\nabla_\theta E_\theta(\tilde{\mathbf{U}})].$$

Integrating both sides shows that \mathcal{L}_θ is equivalent (up to a θ -independent constant) to an energy matching objective, corresponding to the generative loss $\mathcal{L}_{\text{Gen}}(\theta)$:

$$\mathcal{L}_{\text{Gen}}(\theta) = \mathbb{E}_{\mathbf{U} \sim p_d} [E_\theta(\mathbf{U})] - \mathbb{E}_{\tilde{\mathbf{U}} \sim p_\theta} [E_\theta(\tilde{\mathbf{U}})] + \text{const} \approx \frac{1}{N} \sum_{i=1}^N E_\theta(\mathbf{U}_i) - \frac{1}{M} \sum_{j=1}^M E_\theta(\tilde{\mathbf{U}}_j).$$

Finally, minimizing $\mathcal{L}_{\text{Gen}}(\theta)$ is equivalent to minimizing the KL divergence $\mathcal{D}_{\text{KL}}(p_d \| p_\theta)$, since

$$\begin{aligned}
 \nabla_\theta \mathcal{D}_{\text{KL}}(p_d \| p_\theta) &= \nabla_\theta \sum_{\mathbf{U}} p_d(\mathbf{U}) \log \frac{p_d(\mathbf{U})}{p_\theta(\mathbf{U})} \\
 &= \nabla_\theta \sum_{\mathbf{U}} p_d(\mathbf{U}) \log p_d(\mathbf{U}) - \nabla_\theta \sum_{\mathbf{U}} p_d(\mathbf{U}) \log p_\theta(\mathbf{U}) \\
 &= \nabla_\theta \mathbb{E}_{\mathbf{U} \sim p_d} [\log p_d(\mathbf{U})] - \nabla_\theta \mathbb{E}_{\mathbf{U} \sim p_d} [\log p_\theta(\mathbf{U})] \\
 &= 0 + \nabla_\theta \mathcal{L}_{\text{Gen}}(\theta) \\
 &= \nabla_\theta \mathcal{L}_{\text{Gen}}(\theta),
 \end{aligned}$$

where $\mathbb{E}_{\mathbf{U} \sim p_d} [\log p_d(\mathbf{U})]$ is θ -independent.

A.2 REGULARIZATION FOR THE EBM OBJECTIVE

Following Du & Mordatch (2019), we add an ℓ_2 penalty on the magnitudes of the energy values for both real and model samples when optimizing $\mathcal{L}_E(\theta)$. While the relative energy differences between true and negative samples shape the energy landscape, the absolute scale of energies can drift during training, causing large values and numerical instability. The penalty anchors the scale and improves gradient behavior in both positive and negative phases. Concretely, real samples $\mathbf{U}_i \sim p_d(\mathbf{U})$ and negatives $\tilde{\mathbf{U}}_j \sim p_\theta(\tilde{\mathbf{U}})$ in the energy function, we use

$$\mathcal{L}_r(\theta) = \frac{1}{N} \sum_{i=1}^N E_\theta(\mathbf{U}_i)^2 + \frac{1}{M} \sum_{j=1}^M E_\theta(\tilde{\mathbf{U}}_j)^2,$$

810 where N and M are the mini-batch sizes for real and negative samples, respectively. In the overall
 811 objective, this term is weighted by β to prevent gradient overflow and stabilize Langevin dynamics.
 812

813 A.3 DERIVATION OF THE CONDITIONAL EBM OBJECTIVE AS A DISCRIMINATIVE LOSS 814

815 The discriminative term of \mathcal{L}_E can be defined as the expected negative log-likelihood of the target
 816 \mathbf{V} given context \mathbf{U} :

$$818 \quad \mathcal{L}_{\text{Disc}}(\theta) = \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim p_d} [-\log p_\theta(\mathbf{V} | \mathbf{U})]. \\ 819$$

820 Since the conditional probability $p_\theta(\mathbf{V} | \mathbf{U}) = p_\theta(\mathbf{U}, \mathbf{V}) / p_\theta(\mathbf{U})$, we can rewrite this as
 821

$$823 \quad \mathcal{L}_{\text{Disc}}(\theta) = \mathbb{E}_{p_d} [-\log \frac{p_\theta(\mathbf{U}, \mathbf{V})}{p_\theta(\mathbf{U})}]. \\ 824$$

825 Consider the conditional EBM over paired views (\mathbf{U}, \mathbf{V}) :
 826

$$827 \quad p_\theta(\mathbf{V} | \mathbf{U}) = \frac{\exp(-E_\theta(\mathbf{U}, \mathbf{V}))}{Z_\theta(\mathbf{U})}, \quad Z_\theta(\mathbf{U}) = \int \exp(-E_\theta(\mathbf{U}, \mathbf{V}')) d\mathbf{V}', \\ 828$$

829 we have

$$832 \quad p_\theta(\mathbf{V} | \mathbf{U}) = \frac{\exp(-E_\theta(\mathbf{U}, \mathbf{V}))}{\int \exp(-E_\theta(\mathbf{U}, \mathbf{V}')) d\mathbf{V}'}. \\ 833$$

834 Hence, the discriminative loss can be written as
 835

$$837 \quad \mathcal{L}_{\text{Disc}}(\theta) = \mathbb{E}_{p_d} [-\log p_\theta(\mathbf{V} | \mathbf{U})] = \mathbb{E}_{p_d} [-\log \frac{\exp(-E_\theta(\mathbf{U}, \mathbf{V}))}{\int \exp(-E_\theta(\mathbf{U}, \mathbf{V}')) d\mathbf{V}'}]. \\ 838$$

840 In practice, the expectation over p_d is estimated by the empirical average over N data pairs
 841 $\{\mathbf{U}_i, \mathbf{V}_i\}_{i=1}^N$. Instead of approximating the intractable normalizer, we optimize the InfoNCE
 842 surrogate by normalizing over one positive and several negatives. Concretely, for each positive pair
 843 $(\mathbf{U}_i, \mathbf{V}_i^+)$, we select M' negative samples $\{\mathbf{V}_{i,j}^-\}_{j=1}^{M'}$:

$$845 \quad \mathcal{L}_{\text{Disc}}(\theta) \approx -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_i^+))}{\exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_i^+)) + \sum_{j=1}^{M'} \exp(-E_\theta(\mathbf{U}_i, \mathbf{V}_{i,j}^-))}, \\ 846$$

847 where \mathbf{V}_i^+ denotes the positive response for \mathbf{U}_i , and $\mathbf{V}_{i,j}^-$ are the negative samples.
 848

850 B ADDITIONAL RELATED WORKS 851

852 B.1 GRAPH STRUCTURE LEARNING

853 Graph structure learning attempts to infer and optimize explicit or implicit relationships between
 854 nodes directly from data. It has delivered excellent results on graph-related tasks such as node classi-
 855 fication, link prediction, and graph generation. Direct optimization methods including NodeFormer
 856 Wu et al. (2022), STABLE Li et al. (2022), and ProGNN Jin et al. (2020) treat the adjacency matrix
 857 as a learnable and adjust edge weights through task-driven objectives and regularization, thereby
 858 transforming graph construction from a fixed pre-processing step into a trainable module. Metric
 859 learning approaches including SUBLIME Liu et al. (2022b), GRCN Yu et al. (2020), and CoGSL
 860 Liu et al. (2022a) employ a GNN encoder to embed node features into a latent space, compute
 861 pairwise similarities, and normalize them to define connectivity. The resulting structures capture
 862 task-relevant relations while mitigating noise and structural bias. Probability estimation techniques
 863 including GEN Wang et al. (2021a), GAuG-O Zhao et al. (2021), and SGSR Zhao et al. (2023)

864 model links as probabilistic variables, estimate their existence likelihood via neural networks or statistical methods, and then refine the graph by selecting high-probability edges. To date, the dynamic graph structure learning methods have not been explored in ViG architectures. We therefore present the evaluation of their transferability to the vision domain.
 865
 866
 867
 868

869 **B.2 GRAPH CONTRASTIVE LEARNING**
 870

871 Graph contrastive learning uses self supervised objectives to align views of the same node, subgraph, or graph and to separate unrelated samples, producing robust and discriminative embeddings. Zhu
 872 et al. (2020) introduced node level contrastive learning with GRACE. It builds two corrupted graph
 873 views by removing edges and masking features, then maximizes agreement of node embeddings.
 874 The objective follows InfoNCE and can be viewed as maximizing mutual information. Its mix of
 875 structural and attribute changes gives diverse contexts and has links to InfoMax and triplet losses. Yu
 876 et al. (2023b) showed that contrastive training treats nodes unevenly. Hard nodes keep high InfoNCE
 877 loss under many augmentations. They define node compactness and use it as a regularizer in a novel
 878 provable training model. SOLA-GCL Peng et al. (2025) learns how to augment at the subgraph level.
 879 It detects dense subgraphs and applies learned changes such as node dropping, feature masking, and
 880 edge edits. This preserves key subgraph meaning. XSimGCL Yu et al. (2023a) found that most
 881 gains come from the contrastive loss, not from graph changes. It replaces topology edits with simple
 882 noise on embeddings and still excels. BGRL Thakoor et al. (2021) removes negatives and learns
 883 by prediction with an online and a target network. It scales well with simple changes. Collectively,
 884 these studies reveal that principled objectives and adaptive structure learning are key to advancing
 885 graph contrastive methods. Our DenseViG also contributes to the field of graph contrastive learning
 886 by decoupling EBMs into contrastive learning and generative modeling.
 887

888 **C IMPLEMENTATION DETAILS**
 889

890 **C.1 TRAINING ALGORITHM**
 891

892 **Algorithm 1:** DenseViG training algorithm

893 **Input:** Training image \mathcal{I} , graph \mathcal{G} , GNN encoders $\phi_\theta(\cdot)$ and $\psi_\theta(\cdot)$, augmentations t_1, t_2 ,
 894 temperature τ , relaxation π , weights $\alpha, \beta, \mu, \gamma$, batch size N , number of batches B ,
 895 Langevin step size λ and iterations K , and training epochs P
 896 **Output:** Trained DenseViG model (updated parameters θ)
 897 Convert image \mathcal{I} into graph \mathcal{G} , extract patch features \mathbf{X} , and construct the raw adjacency \mathbf{A} via
 898 k -NN algorithm;
 899 Randomly initialize $\phi_\theta(\cdot)$ and $\psi_\theta(\cdot)$ with hyper-parameters $\tau, \lambda, \pi, \alpha, \beta, \mu$, and γ ;
 900 **for** $p = 1, 2, \dots, P$ **do**
 901 **for** $b = 1, 2, \dots, B$ **do**
 902 Sample a mini-batch of N graphs from the training examples;
 903 Generate two correlated graph views $\mathbf{U} = t_1(\mathcal{G}_1)$ and $\mathbf{V} = t_2(\mathcal{G}_2)$;
 904 Compute corresponding node representations $\mathbf{H}_1 = \varphi_\theta(\mathbf{U})$ and $\mathbf{H}_2 = \varphi_\theta(\mathbf{V})$;
 905 Calculate the discriminative InfoNCE loss $\mathcal{L}_{\text{Disc}}(\theta)$ with Eq. 12;
 906 Initialize $\tilde{\mathbf{U}}^0 \sim \mathcal{N}(0, 2\lambda)$;
 907 **for** $k = 1, 2, \dots, K$ **do**
 908 Sample $\tilde{\mathbf{U}}^k$ from relay buffer;
 909 Update $\tilde{\mathbf{U}}^{k+1}$ from $\tilde{\mathbf{U}}^k$ with Eq. 3;
 910 **end**
 911 Calculate the generative EBM loss $\mathcal{L}_{\text{Gen}}(\theta)$ with Eq. 11;
 912 Calculate the complete DEL objective $\mathcal{L}_E(\theta) = \mathcal{L}_{\text{Disc}}(\theta) + \alpha \mathcal{L}_{\text{Gen}}(\theta)$;
 913 Build cosine similarity matrix \mathbf{C} and draw relaxed adjacency \mathbf{S} with Eq. 10;
 914 Symmetrize and normalize \mathbf{S} to form \mathbf{A}^* and compute $\mathbf{H}^* = \psi_\theta(\mathbf{X}, \mathbf{A}^*)$;
 915 Compute the task-specific loss \mathcal{L}_T and form the global objective \mathcal{L}_G ;
 916 Update model parameters θ by gradient descent to minimize $\nabla_\theta \mathcal{L}_G$;
 917 **end**
end

Algorithm 1 outlines training for DenseViG. Each iteration samples a mini-batch of images, converts them to graphs with patch features and a k -NN adjacency, and produces two augmented views. A shared GNN encoder $\phi_\theta(\cdot)$ yields node representations. We calculate the contrastive energy learning objective with a discriminative InfoNCE term in Eq. 12 and a generative term in Eq. 11. The generative negatives are synthesized from Gaussian noise via K steps of Langevin dynamics from Eq. 3. We then refine the graph structure. From the cosine similarity matrix we draw a relaxed adjacency with the Gumbel Sigmoid rule in Eq. 10, symmetrize and normalize it to obtain the refined adjacency, and compute the corresponding representation with the GNN encoder $\psi_\theta(\cdot)$ for the task head. The global loss contains the task objective, the weighted DEL terms, and a sparsity regularizer on the relaxed edges. Model parameters are updated by gradient descent.

C.2 MODEL CONFIGURATION

We scale the DenseViG to five sizes Ti, S, M, B, and L under a four stage pyramid ViG, as shown in Table 8. A DenseViG block applies a Graph encoder followed by an FFN. The encoder builds a k -NN graph in feature space with size 9 and performs multi head updates with 4 heads. Linear projections are applied before and after the graph convolution. The encoder expands channels and then projects back. We use GELU activations, LayerNorm, residual connections, and DropPath. The FFN has two linear layers with an expansion ratio $d_E = 4$.

Table 8: Configurations of DenseViG series. All variants use k -NN size 9, 4 heads, and input resolution 224×224 . Stage 1 applies a convolutional stem to produce 56×56 tokens ($H/4 \times W/4$). Stages 2–4 downsample with stride-2 projections to 28×28 , 14×14 , and 7×7 tokens ($H/8$, $H/16$, $H/32$), widening channels at each stage. d_D denotes the feature dimension, d_F the representation dimension, and d_E the FFN expansion dimension ratio. ‘Ti’ denotes tiny, ‘S’ small, ‘M’ medium, ‘B’ base, and ‘L’ large.

Stage	Tokens	Layer specification	DenseViG variants							
			Ti	S	M	B	L			
1	$\frac{H}{4} \times \frac{W}{4}$	Input	Dimension d_D	192	320	640	768	1024		
			Patch Size	Stem, 56×56 , stride 4						
		DenseViG Block	Dimension d_F	48	80	96	128	192		
			Dimension d_E	4						
		Blocks	2	2	2	2	2	2		
2	$\frac{H}{8} \times \frac{W}{8}$	Input	Patch Size	Downsample, 28×28 , stride 2						
			Dimension d_F	96	160	192	256	384		
		DenseViG Block	Dimension d_E	4						
			Blocks	2	2	2	2	2		
		Input	Patch Size	Downsample, 14×14 , stride 2						
3	$\frac{H}{16} \times \frac{W}{16}$	DenseViG Block	Dimension d_F	240	400	384	512	768		
			Dimension d_E	4						
			Blocks	6	6	16	18	24		
		Input	Patch Size	Downsample, 7×7 , stride 2						
			Dimension d_F	384	640	768	1024	1536		
4	$\frac{H}{32} \times \frac{W}{32}$	DenseViG Block	Dimension d_E	4						
			Blocks	2	2	2	2	2		
Params (M)				11.1	28.5	53.3	98.5	120.3		
FLOPs (G) @ 224×224				1.9	5.3	9.6	18.4	24.9		

All variants have four stages with output resolutions $H/4 \times W/4$, $H/8 \times W/8$, $H/16 \times W/16$, and $H/32 \times W/32$. For input 224×224 , a three layer 3×3 convolutional stem reduces the input to quarter resolution and produces 56×56 tokens, thus $N=3136$. Later stages reduce the token grid by spatial downsampling. After each stage, a stride 2 3×3 projection reduces the number of tokens by a factor of 4 and increases the representation dimension d_F . The block schedule is 2 in Stage 1, 2 in Stage 2, model dependent counts in Stage 3, and 2 in Stage 4, with exact values in Table 8. We

972 use a linear DropPath schedule with end rates 0.1, 0.2, 0.3, 0.5, and 0.6 for T_i through L . Absolute
 973 positional encoding is used, with additional relative positional encoding within the pyramid.
 974

976 C.3 EXPERIMENT SETTING

978 We adopt a unified training pipeline for DenseViG on classification, detection, and segmentation.
 979 Hyper-parameters follow the main ablations. Optimizer is AdamW with $\beta_1=0.9$, $\beta_2=0.99$, weight
 980 decay 0.05. We fix $\tau=0.1$, $\alpha=1.0$, $\beta=10^{-5}$, $\mu=0.1$, $\gamma=10^{-4}$, $k=5$ Langevin iterations and $\lambda=0.05$
 981 step size for negative sampling, relaxation $\pi=0.1$ for structure refinement. We maintain an EMA of
 982 weights with decay 0.999. We apply stochastic depth with drop path rate 0.1. Table 9 summarizes
 983 the training hyper-parameters for each task. In the following, we detail any task-specific setup
 984 differences, while other settings remain consistent across tasks.
 985

986
 987 Table 9: Unified training hyper-parameter settings for DenseViG across ImageNet-1K classification,
 988 MS COCO detection, and ADE20K segmentation. Others not listed are shared across the three tasks.

990 Task	991 Classification (ImageNet-1K)	992 Detection (MS COCO)	993 Segmentation (ADE20K)
994 Optimizer	995 AdamW ($\beta_1=0.9$, $\beta_2=0.99$)	996 AdamW (same as left)	997 AdamW (same as left)
998 Weight decay	999 5×10^{-2}	1000 5×10^{-2}	1001 5×10^{-2}
1002 Layer-wise LR decay	1003 —	1004 0.75 per layer (fine-tune)	1005 0.75 per layer (fine-tune)
1006 Initial learning rate	1007 1×10^{-3}	1008 2×10^{-4}	1009 2×10^{-4}
1010 Training epochs	1011 310	1012 12 (1× schedule)	1013 16 (fine-tune)
1014 Batch size	1015 64	1016 16	1017 16
1018 Input resolution	1019 224×224	1020 1280×800	1021 512×512
1022 LR schedule	1023 Cosine annealing	1024 Step decay (8, 11 ep)	1025 Step decay (proportional)
1026 Data augmentation	1027 RandAug, CutMix, 1028 RandomResizedCrop	1029 RandomFlip, 1030 Multi-scale jittering	1031 RandomFlip, 1032 RandomResizedCrop
1033 Label smoothing	1034 0.1	1035 —	1036 —
1037 EMA decay rate	1038 0.999	1039 0.999	1040 0.999
1041 Drop-path rate	1042 0.1	1043 0.1	1044 0.1
1045 \mathcal{L}_E temperature τ	1046 0.1	1047 0.1	1048 0.1
1049 \mathcal{L}_E balance α	1050 1.0	1051 1.0	1052 1.0
1053 \mathcal{L}_E balance β	1054 1×10^{-5}	1055 1×10^{-5}	1056 1×10^{-5}
1057 \mathcal{L}_G weight μ	1058 0.1	1059 0.1	1060 0.1
1061 \mathcal{L}_G sparsity γ	1062 1×10^{-4}	1063 1×10^{-4}	1064 1×10^{-4}
1066 Langevin iterations k	1067 5	1068 5	1069 5
1071 Langevin step size λ	1072 0.05	1073 0.05	1074 0.05
1077 Relaxation temp. π	1078 0.1	1079 0.1	1080 0.1

1009
 1010
 1011 **ImageNet 1K classification** We train from scratch for 310 epochs on 224×224 images. Initial
 1012 learning rate is 1×10^{-3} with a cosine schedule for total batch size 64 on 8 RTX 3090 GPUs.
 1013 Augmentations include RandAugment, CutMix, RandomResizedCrop, and label smoothing 0.1.
 1014 The backbone uses ViG initialization without pretraining.

1015 **MS COCO detection** We fine tune ImageNet pretrained DenseViG with MMDetection, using RetinaNet
 1016 and Mask R-CNN heads while keeping the backbone and Dense module fixed. Training runs
 1017 for 12 epochs at 1280×800 . AdamW and regularization match classification. We use layer wise
 1018 learning rate decay of 0.75 per ViG layer from output to input. Initial learning rate is 2×10^{-4}
 1019 with step drops at epochs 8 and 11, batch size 16. Augmentations are RandomFlip and Multi-scale
 1020 jittering. EMA is enabled. Losses are sigmoid focal for classification and a weighted sum of ℓ_1 and
 1021 GIoU for boxes.

1022 **ADE20K segmentation** We fine tune ImageNet pretrained DenseViG with MMSegmentation for 16
 1023 epochs on 512×512 random crops. Optimizer, EMA, and layer wise decay match detection. Initial
 1024 learning rate is 2×10^{-4} with a proportional multi step schedule, batch size 16. Augmentations are
 1025 RandomFlip and RandomResizedCrop. We attach linear MLP, PPM, Semantic FPN, and UPerNet
 1026 heads. The backbone is fine tuned with the pixel wise cross entropy loss.

1026 D MORE EXPERIMENT RESULTS

1028 D.1 OBJECT DETECTION AND INSTANCE SEGMENTATION

1030 We evaluate object detection and instance segmentation on MS COCO using RetinaNet and Mask
 1031 R-CNN, as stated in Table 10 and Table 11. All models use ImageNet 1K pretraining and follow the
 1032 MMDetection recipe. Both tables report parameters, FLOPs, and COCO metrics, with RetinaNet
 1033 giving AP^{box} averaged over IoU 0.50 to 0.95, AP^{box} at IoU 0.50 and 0.75, and AP^{box} at small,
 1034 medium, and large sizes, and Mask R-CNN giving both AP^{box} and AP^{mask} with the same breakdown.

1035 Under the same protocol, DenseViG-S improves both tasks across heads and schedules. With Reti-
 1036 naNet 1x it reaches 44.6 AP^{box} , exceeding DVHGNN-S by 1.3 AP^{box} , with consistent gains at AP^{box}_{50}
 1037 and AP^{box}_{75} and the largest gain on large objects with plus 2.3 AP^{box}_L . The 3x with MS setting lifts
 1038 this to 47.0 AP^{box} . With Mask R-CNN 1x, DenseViG-S attains 46.4 AP^{box} and 42.7 AP^{mask} , im-
 1039 proving over DVHGNN-S by 1.6 AP^{box} and 2.5 AP^{mask} . The 3x with MS schedule reaches 50.2
 1040 AP^{box} and 45.2 AP^{mask} . These gains come with similar parameters and FLOPs, which shows that
 1041 DenseViG improves the detection and segmentation performance with high efficiency and scales
 1042 well with stronger heads and longer training.

1044 Table 10: Object detection on MS COCO. All backbones are ImageNet-1K pretrained. “1x” =
 1045 12 epochs; “3x” = 36 epochs; “MS” = multi-scale training. AP : mean over IoU 0.50:0.95 with
 1046 step 0.05 and over categories. AP^{box} : bounding boxes. $AP^{box}_{50/75}$: AP at IoU 0.50 and 0.75.
 1047 $AP^{box}_{S/M/L}$: COCO area; S if area $< 32^2$, M if $32^2 \leq$ area $< 96^2$, L if area $\geq 96^2$, in px^2 .

1049 Backbone	Params	FLOPs	1050 RetinaNet 1x						1051 RetinaNet 3x + MS					
			AP^{box}	AP^{box}_{50}	AP^{box}_{75}	AP^{box}_S	AP^{box}_M	AP^{box}_L	AP^{box}	AP^{box}_{50}	AP^{box}_{75}	AP^{box}_S	AP^{box}_M	AP^{box}_L
1051 ResNet-50	38M	239G	36.3	55.3	38.6	19.3	40.0	48.8	39.0	58.4	41.8	22.4	42.2	51.6
1052 PVT-S	34M	227G	40.4	61.3	43.0	25.0	42.9	55.7	42.2	62.7	45.0	26.2	45.2	57.2
1053 Swin-T	39M	245G	41.5	62.1	44.2	25.1	44.9	55.5	43.9	64.8	47.1	28.4	47.2	57.8
1054 PyramidViG-S	36M	240G	41.8	63.1	44.7	28.5	45.4	53.4	44.2	65.8	47.6	31.8	47.7	55.7
1055 PViHGNN-S	38M	244G	42.2	63.8	45.1	29.3	45.9	55.7	44.6	66.5	48.0	32.6	48.2	58.0
1056 DVHGNN-S	38M	242G	43.3	64.3	46.3	28.3	47.9	54.6	45.7	67.0	49.2	31.6	50.2	56.9
1057 DenseViG-S	40M	247G	44.6	66.1	47.8	29.7	48.5	56.9	47.0	68.8	50.7	33.0	50.8	59.2

1059 Table 11: Object detection and instance segmentation on MS COCO. All backbones are ImageNet-
 1060 1K pretrained. “1x” = 12 epochs; “3x” = 36 epochs; “MS” = multi-scale training. AP^{mask} : instance
 1061 masks. $AP^{mask}_{50/75}$: AP at IoU 0.50 and 0.75.

1063 Backbone	Params	FLOPs	1064 Mask R-CNN 1x						1065 Mask R-CNN 3x + MS					
			AP^{box}	AP^{box}_{50}	AP^{box}_{75}	AP^{mask}	AP^{mask}_{50}	AP^{mask}_{75}	AP^{box}	AP^{box}_{50}	AP^{box}_{75}	AP^{mask}	AP^{mask}_{50}	AP^{mask}_{75}
1066 ResNet-50	44M	260G	38.0	58.6	41.4	34.4	55.1	36.7	41.0	61.7	44.9	37.1	58.4	40.1
1067 PVT-S	44M	245G	40.4	62.9	43.8	37.8	60.1	40.3	43.0	65.3	46.9	39.9	62.5	42.8
1068 Swin-T	48M	264G	42.2	64.6	46.2	39.1	61.6	42.0	46.0	68.1	50.3	41.6	65.1	44.9
1069 PyramidViG-S	46M	259G	42.6	65.2	46.0	39.4	62.4	41.6	46.4	68.8	50.0	41.9	65.9	44.4
1070 PViHGNN-S	48M	262G	43.1	66.0	46.5	39.6	63.0	42.3	46.9	69.6	50.5	42.1	66.5	45.1
1071 DVHGNN-S	49M	261G	44.8	66.8	49.0	40.2	63.5	43.1	48.6	70.4	53.0	42.7	67.0	45.9
1072 DenseViG-S	51M	266G	46.4	68.7	50.6	42.7	65.8	45.4	50.2	72.3	54.6	45.2	69.3	48.2