

EDMB: Edge Detector with Mamba

Yachuan Li^{1*} Xavier Soria Poma² Yun Bai¹ Qian Xiao¹
Chaozhi Yang¹ Guanlin Li¹ Zongmin Li^{1†}

¹ China University of Petroleum (East China) ² Polytechnic School of Chimborazo (ESPOCH)

liyachuan@s.upc.edu.cn, lizongmin@upc.edu.cn

Abstract

Transformer-based models have made significant progress in edge detection, but their high computational cost is prohibitive. Recently, vision Mamba have shown excellent ability in efficiently capturing long-range dependencies. Drawing inspiration from this, we propose a novel edge detector with Mamba, termed EDMB, to efficiently generate high-quality multi-granularity edges. In EDMB, Mamba is combined with a global-local architecture, therefore it can focus on both global information and fine-grained cues. The fine-grained cues play a crucial role in edge detection, but are usually ignored by ordinary Mamba. We design a novel decoder to construct learnable Gaussian distributions by fusing global features and fine-grained features. And the multi-grained edges are generated by sampling from the distributions. In order to make multi-granularity edges applicable to single-label data, we introduce Evidence Lower Bound loss to supervise the learning of the distributions. On the multi-label dataset BSDS500, our proposed EDMB achieves competitive single-granularity ODS 0.837 and multi-granularity ODS 0.851 without multi-scale test or extra PASCAL-VOC data. Remarkably, EDMB can be extended to single-label datasets such as NYUDv2 and BIPED. The source code is available at <https://github.com/Li-yachuan/EDMB>.

1. Introduction

Edge detection is a fundamental task in computer vision, preserving critical semantic information while mitigating texture noise. Therefore, it is widely used in downstream vision tasks such as image retrieval [50], salient object detection [52], semantic segmentation [45], and depth map

prediction [33].

Edge detection has been studied since the last century [36]. Early image edge detectors [2, 5, 14, 31, 36] are based on local cues, which makes it difficult for these models to understand the high-level semantic information. With the application of deep learning in edge detection [16, 24, 37, 43], the performance of detectors has been improved unprecedented, far surpassing the local feature based methods and even surpassing human beings. Consequently, deep learning based detectors have become the mainstream methods for current edge detection. The alternating stacking of convolutional and pooling blocks enables the deep features to obtain larger receptive fields, which is the key to the success of deep learning-based methods. However, it is found that the effective receptive field is smaller than the theoretical receptive field due to the non-uniform distribution of convolution kernel weights and the inhibition effect of nonlinear activation functions [26].

The self-attention mechanism, as implemented in Transformer, is introduced into edge detection to address the limitations of CNN in capturing sufficient context information [28, 49]. While self-attention effectively models global context and resolves long-range dependence, the high computational cost presents significant challenges, particularly in edge detection, which is critical for real-time applications and resource-constrained edge devices. Therefore, the principal challenge in edge detection lies in effectively capturing semantic context and modeling long-range dependence.

Another concern is multi-granularity edges. People have different edge granularity preferences, so it is reasonable to produce different edge annotation for the same image. The preference is meaningful, and the required edge granularity is vary for different downstream tasks. For example, image segmentation [45] requires coarse-grained edges to refine segmentation boundary, and image retrieval [50] requires fine-grained edges to preserve more semantic information. Therefore, it is necessary to design a multi-granularity edge detector. At present, the research on multi-granularity edge is in the initial stage, and the existing methods [53, 54] rely on multi-label data to define the granularity of edges, which

*This work is supported by National key r&d program (Grant no. 2019YFF0301800), National Natural Science Foundation of China (Grant no. 61379106), Shandong Provincial Natural Science Foundation (Grant nos. ZR2013FM036, ZR2015FM011). China Scholarship Council (CSC) (Grant no. 202406450075).

†Corresponding author

greatly limits the application scope of these methods.

To address the above issues, we propose a novel edge detector with Mamba, termed EDMB, to efficiently generate high-quality multi-granularity edges. First, we run a global Mamba encoder to capture the global context information of the image, then divide the image into non-overlapping windows and capture the fine-grained context information within the window by a fine-grained Mamba encoder. Although Mamba demonstrates greater efficiency compared to the Transformer, its computational cost increases proportionally with feature resolution due to the demands of long-distance modeling. With this in mind, it is advantageous to avoid applying Mamba to high-resolution features to mitigate computational cost. We propose a shared CNN-based high-resolution feature encoder to process high-resolution features, which is composed of several straightforward convolutional and nonlinear modules. In addition, a learnable Gaussian distribution decoder is developed to fuse and decoder these features into learnable Gaussian distributions. And the multi-grained edges are generated by sampling from the distributions. Evidence Lower Bound loss is introduced to supervise the distribution, which can effectively avoid the dependence on multi-label data.

Our contributions are summarized as follows: **1)** We propose a novel Mamba-based edge detector, EDMB, to efficiently generate high-quality edges. To our best knowledge, EDMB is the first Mamba-based edge detector. **2)** We propose a multi-granularity edge generation method based on Evidence Lower Bound loss, so that the generation of multi-granularity edges does not rely on multi-label data. **3)** Extensive experimental results show that the performance of EDMB is highly competitive on both multi-label and single-label datasets such as BSDS [1], NYUDv2 [34] and BIPED [27].

2. Related work

2.1. Edge Detection

Edge detection has been studied since the last century [5, 31]. Early classical methods [2, 10, 14, 36] typically rely on the local cues of images. The lack of global information limits the performance of these detectors. As the pioneer of contemporary edge detection, HED [43] introduces deep supervision mechanism to edge detection and learns multi-scale predictions holistically. Most of the subsequent methods [16, 24, 46] follows HED and focus on feature enhancement and fusion.

In pursuit of higher accuracy, recent studies focus on two aspects. The first is to improve the representational capacity the encoder. In addition to enlarge the scale of CNN [53, 54], an effective method is to introduce Transformer, which have stronger representation ability, into edge detection. EDTER [28] proposes a Transformer combining

global-local features to generate high-quality edges. DiffusionEdge [49] treats edge detection as a denoising process and integrates Swin Transformer with diffusion to generate crisp edges. Similarly, the success of RankED [3] is also inseparable from the powerful representation ability of Transformer. These methods effectively improve the performance of edge detection, but the high computational cost of Transformer limit its application in more downstream tasks.

Another research hotspot is edge uncertainty. Different people have different perception of the edge, which leads to the uncertainty of the edge and poses new challenges to edge detection. PEder [12] proposes a collaborative learning framework, which harnesses the disparity in edge perception between two heterogeneous models to refine the labels. RankED [3] introduces RS Loss to edge detection and promotes high confidence edge pixels to have more label certainty. UAED [54] constructs learnable Gaussian distributions to take full use of multiple labels, and BetaNet [22] replaces Bernoulli distribution with Beta distribution and employs recurrent voting strategy to merge multiple labels. MuGE [53] designs an edge granularity network to estimate the edge granularity from multiple labeled edge annotation, which can guide edge detector to generate edges in different granularity. The performance of edge detection can be improved by evaluating the uncertainty of edges. In particular, UAED and MuGE can generate multi-granularity edges, which is beneficial for a wider range of applications of edge detection. While the generation of multi-granularity edges heavily rely on multi-label data, which makes it difficult to apply these methods to the more common single-label data.

2.2. Vision Mamba

Mamba is a kind of selective state-space model (SSM) for modeling long-range dependencies of temporal sequences [13]. Inspired by Vision Transformer, Mamba is applied to serialized patches of images and achieves competitive accuracy on computer vision tasks, such as image classification [55], image segmentation [48] and image restoration [4]. Mamba no longer rely on self-attention, the high computational cost module in Transformer. As a result, Mamba is significantly more efficient than Transformer, and has received continuous attention [40].

Similar to Vision Transformer, early Vision Mamba focuses more on global modeling, which makes it less attentive to local cues. Therefore, LocalMamba [19] introduces a novel scanning method that consists of local scans within different windows to capture local detail information and global context. It searches for scan directions at different network layers to identify and apply the most effective combination of scans. In addition, there are various scanning methods such as Continuous 2D Scanning, Zigzag Scan, and Multi-Path Scan [51], which are mainly aim to achieve the balance between the local cues and the global context of

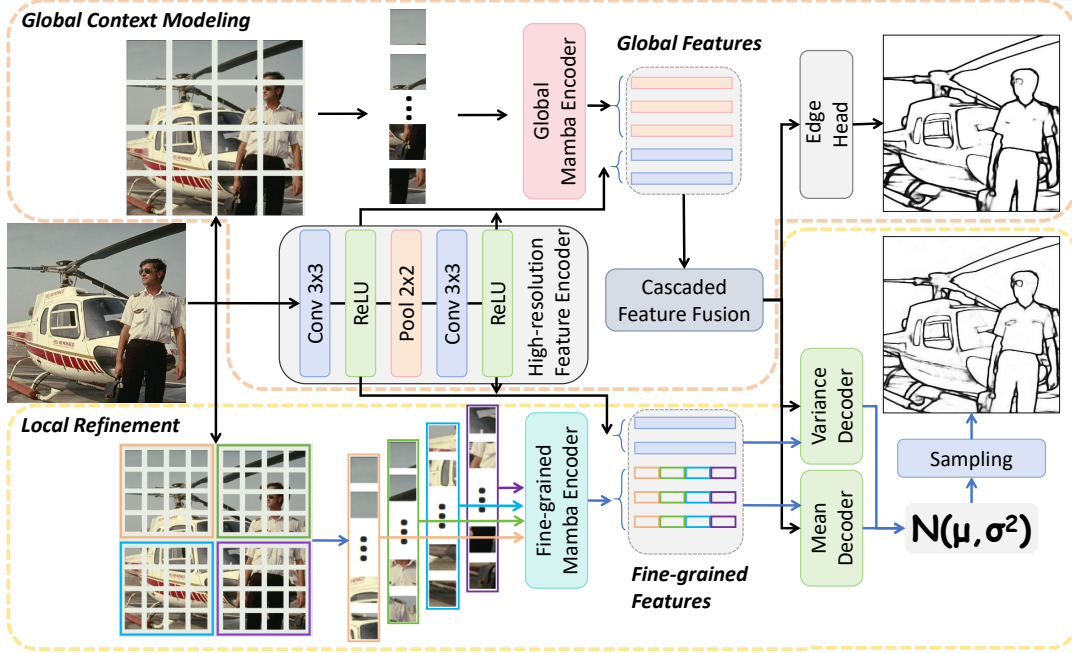


Figure 1. The EDMB’s framework. $N(\mu, \sigma^2)$ means learnable Gaussian distributions, where the means μ and variances σ^2 are predicted by the mean decoder and variance decoder, respectively.

Mamba. To date, modeling Mamba’s long-range and short-range dependencies simultaneously and efficiently remains an open problem.

The computational efficiency of Mamba is between Transformer and CNN. To further improve the efficiency of the model, Mamba is combined with CNN [23, 42], where CNN is used to process high-resolution features. We agree with the idea and further compress the scale of the high-resolution feature encoder, which is only used to maintain the location information of features.

3. Edge Detector with Mamba

3.1. Overview

The overall framework of the proposed EDMB is illustrated in Fig. 1. High-resolution feature encoder E_h , global Mamba encoder E_g and fine-grained Mamba encoder E_f encode image X into features F_h , F_g and F_f , respectively. Multi-scale global features $\{F_h, F_g\}$ and Multi-scale fine-grained features $\{F_h, F_f\}$ are fed into the decoder at the same time. Cascading feature fusion module decode $\{F_h, F_g\}$ into high-resolution global feature F_{hg} , which guides the decoding of $\{F_h, F_f\}$. $\{F_h, F_f\}$ and F_{hg} are combined into a new set, $\{F_h, F_g, F_{hg}\}$, which are input into mean decoder D_m and variance decoder D_v , respectively. Learnable Gaussian distributions are constructed by the output of the two decoders, and the final edge map is sampled from the distributions.

3.2. Review Vision Mamba

The Mamba encoders of EDMB follow the Vision Mamba (ViM) from [55], which is applicable in 2-D sequence. A 2-D image $X \in \mathbb{R}^{H \times W \times 3}$ is first transformed into flattened 2-D patches $X_p \in \mathbb{R}^{C \times (P^2 \times 3)}$, where $(H \times W)$ is the image size, P is the patch size. Next, the X_p is linearly projected to latent embedding $E_l \in \mathbb{R}^{C \times N}$ and add position embeddings $E_{pos} \in \mathbb{R}^{C \times N}$. Different from ViT [11], which directly computes the correlation between all patches through costly Self-Attention, the core of ViM is advanced state space model (SSM). In SSM, the correlation between patches is established through the state matrix $A \in \mathbb{R}^{N \times N}$ and projection parameters $B \in \mathbb{R}^{C \times 1}$, $C \in \mathbb{R}^{C \times 1}$. The process can be represented as

$$\begin{aligned} h'(t) &= Ah(t) + Bx(t), \\ y(t) &= Ch(t) \end{aligned} \quad (1)$$

where $h(t) \in \mathbb{R}^N$ represents intermediate implicit states. To transform the continuous parameters A , B to discrete parameters \bar{A} , \bar{B} , transformation method zero-order hold is used:

$$\begin{aligned} \bar{A} &= \exp(\Delta A), \\ \bar{B} &= (\Delta A)^{-1}(\exp(\Delta A) - I) \cdot \Delta B. \end{aligned} \quad (2)$$

The output y of SSM is compute through $y = x * (CB, CAB, \dots, C\bar{A}^{M-1}\bar{B})$, where M is the length of the input sequence x . Vanilla SSM is unidirectional correlation,

which is suitable for temporal sequence. ViM constructs bidirectional SSM, $SSM_f(\cdot)$ and $SSM_b(\cdot)$, by reversing the input sequence to capture the context information between image patches. The final result can be roughly expressed as:

$$\mathbf{T}_l = \text{Lin}(SSM_f(\mathbf{T}_{l-1}) + SSM_b(\mathbf{T}_{l-1})) + \mathbf{T}_{l-1}, \quad (3)$$

where T is the token sequence. l is the l -th layer of the Vim encoder. $\text{Lin}(\cdot)$ is a linear projection.

3.3. Mamba-based Encoder

The Mamba-based encoder is shown in Fig. 1, where features are divided into global features, fine-grained features, and high-resolution features. Those features are encoded by global Mamba encoder E_g , fine-grained Mamba encoder E_f and high-resolution feature encoder E_h , respectively. Following ViM, E_g performs $4 \times$ downsampling during image patching to avoid Mamba’s high computational cost on high-resolution features. And only the first three stages of the ViM encoder are adopted, avoiding meaningless computations on low-resolution features. Specifically, the input image is split into a sequence of 4×4 patches, and then the sequence is linearly projected to latent embedding. The output \mathbf{F}_g of E_g can be compute by

$$\mathbf{F}_g = \{\mathbf{F}_g^1, \mathbf{F}_g^2, \mathbf{F}_g^3\} = E_g(\mathbf{X}), \quad (4)$$

where \mathbf{F}_g^i is the output of stage i in E_g .

Inspired by EDTER [28], we propose fine-grained Mamba encoder E_f to capture fine-grained cues which is ignored by E_g . Although E_f has the same structure as E_g , the data processing is different. We decompose image $X \in \mathbb{R}^{H \times W \times 3}$ into a sequence $\{X_1, X_2, X_3, X_4\}$ through a $\frac{H}{2} \times \frac{W}{2}$ non-overlapping sliding window and encode the sequence with the fine-grained Mamba encoder. We randomly drop the gradient of X_i to reduce the computational cost during training. While the whole sequence is encoded to obtain the full fine-grained image features F_f . The output \mathbf{F}_f of E_f can be compute by

$$\mathbf{F}_f = \{\mathbf{F}_f^1, \mathbf{F}_f^2, \mathbf{F}_f^3\} = E_f(\{X_1, X_2, X_3, X_4\}). \quad (5)$$

Compared to CNN, Mamba is more expensive to compute on high-resolution features, therefore, the data is downsampled before input to Mamba-based encoders. In order to make up for the location error caused by downsampling with a smaller cost, we propose high-resolution feature encoder E_h , which is a lightweight encoder based on CNN. E_h contains two vanilla Convolution blocks, and each vanilla convolution block consists of a 3×3 convolution layer and a ReLU activation. To ensure the accuracy of the location information, the feature resolution in the first vanilla Convolution block is consistent with the input image. A down-sampling module is inserted between the two

convolution blocks to generate two-scale location features. Consistent with VGG16 [35], the number of channels doubles when the feature resolution is halved. Since accurate location information rather than rich semantic information is mainly consider, the features in E_h is compressed to 16 channels, which is much smaller than existing models such as VGG16 [35] (64 channels) and ResNet50 [17] (64 channels). Compared to VGG16 and ResNet50, E_h not only has fewer modules, but also has fewer channels. Therefore, even though the features keep high resolution, the computational cost remains low. The output \mathbf{F}_h of E_h can be computed by

$$\mathbf{F}_h = \{\mathbf{F}_h^1, \mathbf{F}_h^2\} = E_h(\mathbf{X}). \quad (6)$$

3.4. Learnable Gaussian Distributions Decoder

Following previous works [53, 54], the key to generate multi-granularity edges is to construct learnable Gaussian distributions, and adjust the edge granularity by changing the variance. Therefore, We propose a learnable Gaussian distributions (LGD) decoder that utilizes multi-scale features to predict the means and variances of learnable Gaussian distributions, as shown in fig. 2.

The input of the decoder is multi-scale global features $\{\mathbf{F}_g, \mathbf{F}_h\}$ and multi-scale fine-grained features $\{\mathbf{F}_f, \mathbf{F}_h\}$. A Cascaded Feature Fusion (CFF) module fuses $\{\mathbf{F}_g, \mathbf{F}_h\}$ to obtain high-resolution global features \mathbf{F}_{hg} . The other two CFF modules utilize $\{\mathbf{F}_f, \mathbf{F}_h\}$ to generate fine-grained mean features \mathbf{F}_{fm} and fine-grained variance features \mathbf{F}_{fv} , respectively. The structure of CFF module is shown in the left part of fig. 2. The upsampled low-resolution features are concatenated with the adjacent higher-resolution features in the channel dimension, the result are fused by MB-Conv [18]. The output of the module is the features with highest resolution. CFF module is widely applied [30] due to its balance in effectiveness and efficiency. Global features contain richer global context for distinguishing edges and textures more effectively from the semantic level. Fine-grained features are better at capturing short-range local cues and are able to perceive thinner edges. Therefore, \mathbf{F}_{hg} is combined with \mathbf{F}_{fm} and \mathbf{F}_{fv} to generate the mean and variance respectively in the mean decoder D_m and variance decoder D_v .

Following EDTER [28], a spatial feature transform block is hired in D_m to fully modulate the global features and fine-grained features. A mean head consisting of several simple Conv-BN-ReLU modules maps the results to single-channel means of distributions. D_v is basically the same as D_m in structure. The difference is that the variance head contains an additional SoftPlus function to process the output to ensure the non-negativity of the variances of distributions.

In addition, we hire an Edge Head, a Mean Head and a Variance Head to directly deal with high-resolution global

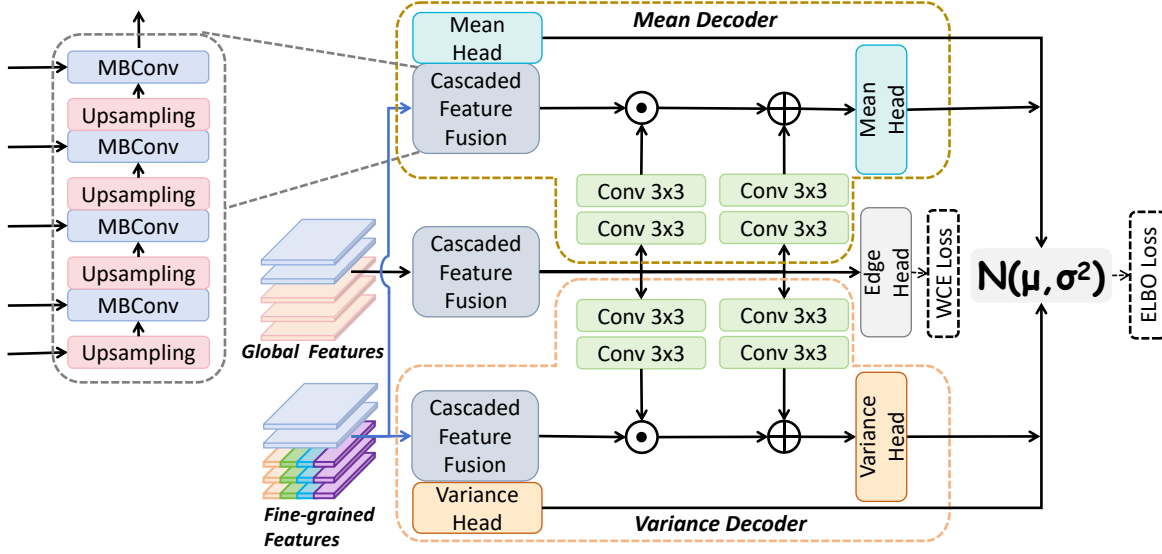


Figure 2. The detailed architecture of the LGD decoder. *MBCConv* is a kind of efficient convolution proposed by MobileNetv2 [18].

features \mathbf{F}_{hg} , fine-grained mean features \mathbf{F}_{fm} and fine-grained variance features \mathbf{F}_{fv} , which are used to aid training and will be discarded during inference. The outputs of the three auxiliary head are termed p_a , μ_a and σ_a^2 . The learnable Gaussian distributions can be expressed as

$$\begin{aligned} \mathcal{N}(\mu, \sigma^2) &= \mathcal{N}(D_{LGD}(\{\mathbf{F}_{\text{g}}, \mathbf{F}_{\text{h}}\}, \{\mathbf{F}_{\text{f}}, \mathbf{F}_{\text{h}}\})) \\ &= \mathcal{N}(D_m(\mathbf{F}_{\text{hg}}, \mathbf{F}_{\text{fm}}), D_v(\mathbf{F}_{\text{hg}}, \mathbf{F}_{\text{fv}})) \end{aligned} \quad (7)$$

3.5. Loss Function

To force the outputs obey Gaussian distributions, we introduce ELBO (Evidence Lower Bound) loss [21] to supervise the learning of EDMB. ELBO loss consists a Kullback-Leibler (KL) divergence, which measures the similarity between the predicted distribution and the prior distribution, and a Weighted Cross Entropy (WCE) loss, which measures the distance between the ground truth y and a sampling p of the distributions.

For image X , KL loss can be formulated as

$$\begin{aligned} L_{kl}(N(\mu, \sigma^2), N(0, 1)) &= \int \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{\xi} \left(\log \frac{e^{\xi}/\sqrt{2\pi\sigma_i^2}}{e^{-x^2/2}/\sqrt{2\pi}} \right) dx \\ &= \int \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{\xi} \log \left\{ \frac{1}{\sqrt{\sigma_i^2}} \exp \left\{ \frac{1}{2} (x^2 + 2\xi) \right\} \right\} dx \\ &= \frac{1}{2} \int \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{\xi} (-\log \sigma_i^2 + x^2 + 2\xi) dx \\ &= \frac{1}{2} \sum_{i=1}^N (\mu_i^2 + \sigma_i^2 - \log \sigma_i^2 - 1), \end{aligned} \quad (8)$$

where $\xi = -(x - \mu_i)^2/2\sigma_i^2$ and N is the total number of pixels in image X . WCE loss [53, 54] can be formulated as

$$L_{wce}(p, y) = - \sum_{i=1}^N \alpha \log(p_i) + \beta \log(1 - p_i), \quad (9)$$

where p denotes a sampling p of the learnable distributions $N(\mu, \sigma^2)$ and y is the ground truth. $\alpha = |Y_+|/|Y|$ and $\beta = \lambda \cdot |Y_-|/|Y|$. $|Y_+|$ and $|Y_-|$ are used to represent the number of edge and non-edge, separately. λ controls the weight of positive over negative samples. $|Y| = |Y_+| + |Y_-|$. N is the total number of pixels in image X . Therefore, ELBO loss is defined as

$$\begin{aligned} L_{elbo}(p, y, N(\mu, \sigma^2), N(0, 1)) \\ = L_{wce}(p, y) + \varphi L_{kl}(N(\mu, \sigma^2), N(0, 1)), \end{aligned} \quad (10)$$

where φ is a hyperparameter to balance the reliability and randomness of the output.

In practice, we easily encounter the predicament that global features and fine-grained features do not converge at the start of training. We propose two auxiliary losses to supervise the learning of global features and fine-grained features. As shown in Fig 2, WCE loss $L_{wce}(p_a, y)$ is used for the high-resolution global features \mathbf{F}_{hg} . ELBO loss $L_{elbo}(p_a, y, N(\mu_a, \sigma_a^2), N(0, 1))$ is used for the fine-grained mean features \mathbf{F}_{fm} and the fine-grained variance features \mathbf{F}_{fv} .

The total loss of EDMB is

$$\begin{aligned} L &= L_{elbo}(p, y, N(\mu, \sigma^2), N(0, 1)) \\ &\quad + \alpha_1 L_{wce}(p_a, y) \\ &\quad + \alpha_2 L_{elbo}(p_a, y, N(\mu_a, \sigma_a^2), N(0, 1)), \end{aligned} \quad (11)$$

where α_1 and α_2 are hyperparameters to balance multiple loss functions.

3.6. Training Strategy

In theory, EDMB can be trained end-to-end. While we follow the practice of EDTER [28] and adopt multi-stage training. As shown in Fig 1, the global Mamba encoder is employed at first to catch the global context, and the loss function in this stage is

$$L_g = L_{wce}(p_a, y). \quad (12)$$

Then, the global information can guide the learning of the fine-grained Mamba encoder and generate learnable Gaussian distributions (LGD). The loss function in this stage is

$$L_f = L_{elbo}(p, y, N(\mu, \sigma^2), N(0, 1)) + \alpha_2 L_{elbo}(p_a, y, N(\mu_a, \sigma_a^2), N(0, 1)). \quad (13)$$

Compared with end-to-end training, multi-stage training has three main advantages: **1)** In the local refinement stage, gradients are not applied to the global Mamba encoder or the high-resolution feature encoder, in order to minimize computational overhead and storage demands. **2)** The gradient of the model is more stable and makes model training easier. **3)** Compared to Eq. 11, optimizing Eq. 12 and Eq. 13, respectively, avoids the dependence on the hyperparameter α_1 .

Different edges can be obtained by sampling independently multiple times from the learnable Gaussian distributions, but they are in the same granularity. To generate multi-granularity edges, we follow the practice of MuGE [53] and control the granularity of the edges by adjusting the weight of the learned variance. The final edge map can be formulated as

$$p_\gamma = \mu + \gamma\sigma^2, \quad (14)$$

where γ is a hyperparameter that controls the granularity.

4. Experiments

4.1. Datasets

We evaluate the proposed EDMB on three widely used datasets BSDS500, NYUDv2, and BIPED. **BSDS500** [1] is a popular natural scenes dataset for edge detection, consisting of 200 training images, 100 validation images, and 200 test images. Each image is annotated by 4 to 9 experts. In accordance with previous studies [24, 43], the model is trained on both the training and validation sets, and evaluated on the test set. **NYUDv2** [34] is a indoor scene semantic segmentation dataset. The edge ground truth is generated from the segmentation maps. NYUDv2 comprises 1,449 groups of meticulously annotated RGB and depth images, with each group containing a single annotated edge.

For our experiments, we utilized 795 images for training the model and evaluated its performance on the remaining images. **BIPED** [38] contains 250 carefully annotated high-resolution Barcelona Street View images. There are 200 images for training, and the rest 50 images for testing. All images are carefully annotated at single-pixel width by experts in the computer vision. The data augmentation strategy during training is shown in Table 1.

Datasets	Augmentation strategies	Resize
BSDS	F (4×), R (25×)	481 × 321
NYUDv2	F (2×), S (3×), R (4×)	400 × 400
BIPED	F (2×), S (3×), R (16×), G (3×)	400 × 400

Table 1. Augmentation strategies adopted on three edge detection benchmarks. **F**: flipping, **S**: scaling, **R**: rotation, **C**: cropping, **G**: gamma correction.

4.2. Implementation Details

The proposed EDMB is implemented using PyTorch library. All parameters are updated by Adam optimizer. The Mamba based encoders are pre-trained on ImageNet, and the remaining modules are randomly initialized. The batch-size is 4 for the global context modeling stage. And the batchsize of the local refinement stage is reduced to 3 due to the limitation of the memory. Weight decay is set to 5e-4. For the pretrained module, the learning rate is 1e-5, and the rest is set to 1e-4. Following previous works [24, 27], the hyperparameter λ is set to 1.1 for both BSDS500 and BIPED, and 1.3 for NYUDv2. The hyperparameter α_2 in Eq. 13 is set to 0.4 for balancing the multiple loss functions. A standard Non-Maximum Suppression (NMS) is performed to produce the final edge maps before the quantitative evaluation. $F_measure = \frac{2 \times P \times R}{P + R}$ is utilized as the quality evaluation standard of the generated edge map, where P represents the accuracy rate and R represents the recall rate. The edges are binarized by an optimal threshold value for each image (OIS) or the same optimal threshold for the whole dataset (ODS). For more experimental details, please refer to previous works [24, 27].

In addition to accuracy measurements, key metrics are included to evaluate the model compactness and inference efficiency. They are measured by model size (Param) and Giga Floating-Point Operations (GFLOPs). GFLOPs is a hardware independent metric of computational requirements that is more fair compared to Frames Per Second (FPS).

4.3. Comparison with State-of-the-arts

We compare our method with classic edge detectors (Canny [2], SCG [29], OEF [15], etc.), CNN-based detectors (RCF [24], BDCN [16], UAED [54], etc.) and the Transformer-based detectors (EDTER [28], DiffusionEdge

Methods	ODS	OIS	Param	GFLOPs
Canny _{TPAMI'86} [2]	0.611	0.676	-	-
gPb-UCM _{TPAMI'10} [1]	0.729	0.755	-	-
SCG _{NeurIPS'12} [29]	0.739	0.758	-	-
SE _{TPAMI'14} [9]	0.743	0.764	-	-
OEF _{CVPR'15} [15]	0.746	0.770	-	-
DeepContour _{CVPR'15} [32]	0.757	0.776	-	-
Deep Boundary _{ICLR'15} [20]	0.789	0.811	-	-
CEDN _{CVPR'16} [47]	0.788	0.804	-	-
RDS _{CVPR'16} [25]	0.792	0.810	-	-
AMH-Net _{NeurIPS'17} [44]	0.798	0.829	-	-
HED _{TPAMI'17} [43]	0.788	0.808	14.6	57.5
CED _{CVPR'17} [41]	0.803	0.820	21.8	-
LPCB _{ECCV'18} [8]	0.800	0.816	-	-
RCF _{TPAMI'19} [24]	0.798	0.815	14.8	75.3
DSDC _{ACMMM'20} [6]	0.802	0.817	-	-
LDC _{ACMMM'21} [7]	0.799	0.816	-	-
BDCN _{TPAMI'22} [16]	0.806	0.826	16.3	103.4
FCL _{NN'22} [46]	0.807	0.822	16.5	134.4
EDTER _{CVPR'22} [28]	0.824	0.841	468.8	802.3
PiDiNet _{TPAMI'23} [39]	0.789	0.803	0.71	11.6
UAED _{CVPR'23} [54]	0.829	0.847	72.5	72.8
RankED _{CVPR'24} [3]	0.824	0.840	121.0	181.0
MuGE _{CVPR'24} [53]	0.831	0.847	72.5	73
DiffusionEdge _{AAAI'24} [49]	0.834	0.848	225	494
EDMB (Our)	0.837	0.851	71.0	266.3
UAED*	0.841	0.847	72.5	-
MuGE*	0.850	0.856	72.5	-
EDMB* (Our)	0.851	0.854	71.0	-

Table 2. Quantitative results on the BSDS500 dataset. All edges are generated without multi-scale testing or additional PASCAL-VOC data. * means generating multi-granularity edges. The best two results are denoted as **red** and **blue** respectively, and the same for other tables.

[49], etc.) on BSDS500. Quantitative results are shown in Table 2. By training on the train-val set of BSDS500, our method achieves the F-measure ODS of 0.837 with single-granularity testing, which is superior to all the state-of-the-art edge detectors. Compared with the second best method DiffusionEdge, both ODS and OIS of EDMB are improved by 0.3%. More importantly, the model size of EDMB is only 1/3 of that of DiffusionEdge, and the computational cost is reduced to 1/2 of that of DiffusionEdge. It reflects the advantages of Mamba-based models over Transformer-based methods in edge detection.

Following MuGE [53], we evaluate the multi-granularity edge of EDMB and compare it with UAED and MuGE, which is denoted as EDMB*, UAED* and MuGE*. By tuning the hyperparameter γ in Eq. 14, we can obtain multi-granularity edges. EDMB* generates $M = 11$ samples with $\gamma = \frac{n}{2} - 5, n \in \{0 \sim 10\}$. The sampling

Methods	NYUDv2		BIPED	
	ODS	OIS	ODS	OIS
AMH-Net _{NeurIPS'17}	0.744	0.758	-	-
HED _{TPAMI'17}	0.722	0.737	0.829	0.847
RCF _{TPAMI'19}	0.745	0.759	0.843	0.859
BDCN _{TPAMI'22}	0.748	0.762	0.839	0.854
PiDiNet _{ICCV'21}	0.733	0.747	0.868	0.876
EDTER _{CVPR'22}	0.774	0.789	0.893	0.898
DexiNed _{PR'23}	-	-	0.895	0.900
DiffusionEdge _{AAAI'24}	0.761	0.766	0.899	0.901
RankED _{CVPR'24}	0.780	0.793	-	-
EDMB (Our)	0.774	0.787	0.903	0.907
EDMB* (Our)	0.783	0.789	0.906	0.908

Table 3. Quantitative comparisons on NYUDv2 and BIPED. All results are computed with a single scale input. * means generating multi-granularity edges.

E_g	E_f	E_h	LGD	MG	ODS	OIS
✓					0.821	0.830
✓	✓				0.836	0.846
✓	✓	✓			0.839	0.848
✓	✓	✓	✓		0.837	0.851
✓	✓	✓	✓	✓	0.851	0.854

Table 4. Ablation study of each component on BSDS500.

number M of EDMB* is the same as MuGE* and UAED* for fair comparison. EDMB* achieves the best ODS of 0.851, surpassing UAED* and MuGE*. Although compared with MuGE*, EDMB* has no significant advantage in accuracy but higher computational cost. However, EDMB* gets rid of the dependence on multi-label data, which makes EDMB* can be applied to a wider range of single-label data. Moreover, the computational cost of multi-granularity edges is proportional to the granularity M , which is much higher than that of single-granularity edges. Therefore, the computational cost is usually not considered when evaluating multi-granularity edges.

The experimental results on the single-label datasets NYUDv2 and BIPED are recorded in Table 3. The ODS of EDMB* on NYUDv2 is 0.783, which can match the SOTA method RankED and significantly better than other methods. On BIPED, EDMB* achieves 0.924 ODS, which is the new SOTA method and 0.7% higher than the second best DiffusionEdge. The improvement in accuracy is astounding.

Qualitative results are presented in Fig. 3. Overall, EDMB* gets crisper edges. The granularity of EDMB* is similar to that of MuGE*, and is significantly better than that of UAED*.

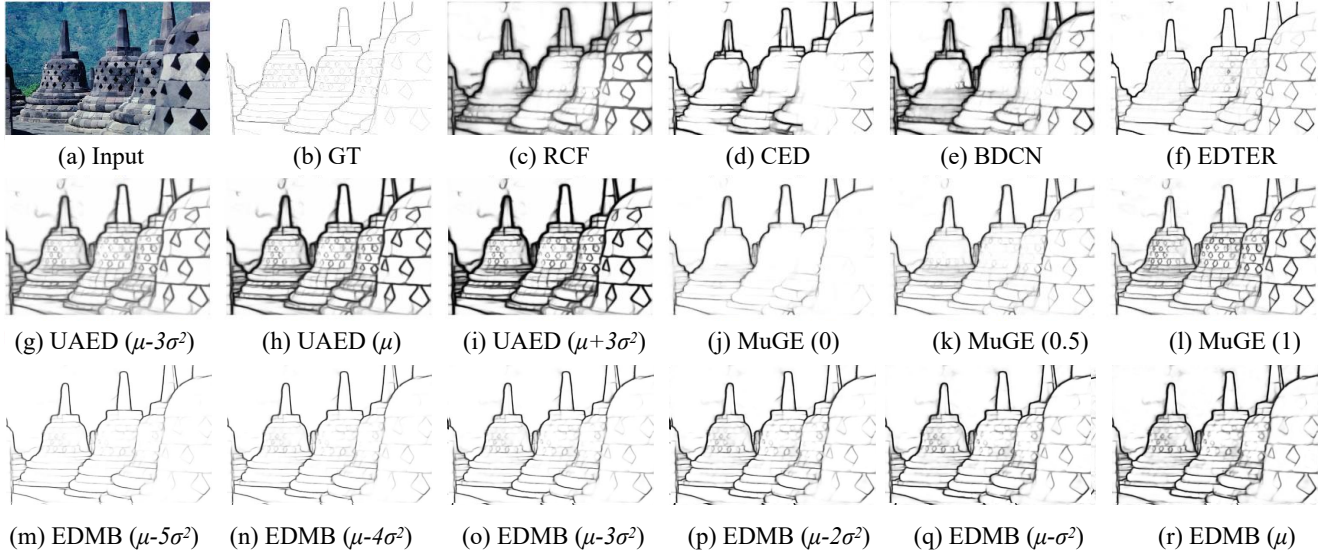


Figure 3. Qualitative comparisons on challenging samples in the BSDS500 test set. MuGE produces diverse results with edge granularity of 0, 0.5, and 1, respectively. EDMB and UAED samples from the learned distribution with $\mu + \gamma\sigma^2$. Fig. (c)~(l) come from [53].

Method	Label	Granularity (γ)	ODS	OIS
UAED	Random	$\pm \frac{n}{2}, n \in \{0, 2 \sim 6\}$	0.841	0.847
EDMB	Random	$\pm \frac{n}{2}, n \in \{0, 2 \sim 6\}$	0.849	0.855
EDMB	Mixed	$\pm \frac{n}{2}, n \in \{0, 2 \sim 6\}$	0.847	0.853
EDMB	Mixed	$\frac{n}{2} - 5, n \in \{0 \sim 10\}$	0.845	0.852
EDMB	Random	$\frac{n}{2} - 5, n \in \{0 \sim 10\}$	0.851	0.854

Table 5. Effects of multi-label handling methods and granularity choices on BSDS500 dataset. Multi-granularity is obtained by sampling $\mu + \gamma\sigma^2$ from LGD.

4.4. Ablation Study

We perform ablation experiments on BSDS500, and the results are recorded in Table 4. E_g , E_f and E_h mean the global Mamba encoder, the fine-grained Mamba encoder and the high-resolution feature encoder. The contribution of each module can be visualized by the improvement of accuracy. Although LGD does not directly contribute to the accuracy, the construction of LGD is a prerequisite for multi-granularity (MG) edge and cannot be omitted.

We further explore the influence of multi-label selection and the granularity setting on the experimental results. The results are recorded in Table 5 and Figure 4. EDMB no longer relies on multi-label data, there are two label processing methods for multi-label data, mixed label of RCF [24] or randomly selected label of UAED [54]. Granularity setting means the setting of hyperparameter γ in Eq. 14. By observing the results of different granularities in Figure 4, it is found that when $\gamma \in [-5, 0]$, the granularity of EDMB changes most obviously, and the accuracy of multi-

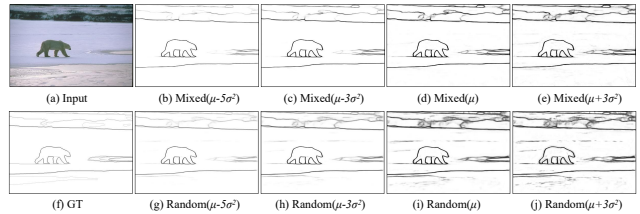


Figure 4. Results of different multi-label handling methods. The first row shows the results for the mixed label, and the second row shows the results for the randomly selected label

granularity edge in Table 5 is also the highest. When $\gamma > 0$, the granularity of the edges hardly changes anymore. The mixed label performs worse than randomly selected label, not only less edge granularity change as shown in Figure 4, but also lower multi-granularity accuracy.

5. Conclusion

EDMB leverages Mamba to efficiently capture long-range dependencies and fine-grained cues of images, which surpasses Transformer-based detectors in terms of accuracy and efficiency. In addition, we introduce ELBO loss to supervise learnable Gaussian distributions that are able to generate multi-granularity edges without relying on multi-labeled data. EDMB is the first edge detector that can generate multi-granularity edges on single-label datasets and achieves competitive results. **Limitation.** The computational cost of EDMB is higher than CNN. And there is a lack of discussion on the application of multi-granularity edges to downstream tasks.

References

- [1] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010. 2, 6, 7
- [2] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, 8(6):679–698, 1986. 1, 2, 6, 7
- [3] Bedretin Cetinkaya, Sinan Kalkan, and Emre Akbas. Ranked: Addressing imbalance and uncertainty in edge detection using ranking-based losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3239–3249, 2024. 2, 7
- [4] Chi-Sheng Chen, Guan-Ying Chen, Dong Zhou, Di Jiang, and Dai-Shi Chen. Res-vmamba: Fine-grained food category visual classification using selective state space models with deep residual learning. *arXiv preprint arXiv:2402.15761*, 2024. 2
- [5] Larry S Davis. A survey of edge detection techniques. *Computer graphics and image processing*, 4(3):248–270, 1975. 1, 2
- [6] Ruoxi Deng and Shengjun Liu. Deep structural contour detection. In *Proceedings of the 28th ACM international conference on multimedia*, pages 304–312, 2020. 7
- [7] Ruoxi Deng, Shengjun Liu, Jinxin Wang, Huibing Wang, Hanli Zhao, and Xiaoqin Zhang. Learning to decode contextual information for efficient contour detection. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4435–4443, 2021. 7
- [8] Ruoxi Deng, Chunhua Shen, Shengjun Liu, Huibing Wang, and Xinru Liu. Learning to predict crisp boundaries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 562–578, 2018. 7
- [9] Piotr Dollár and C Lawrence Zitnick. Fast edge detection using structured forests. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1558–1570, 2014. 7
- [10] Piotr Dollár and C. Lawrence Zitnick. Fast edge detection using structured forests. *arXiv: Computer Vision and Pattern Recognition*, Jun 2014. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [12] Yuanbin Fu and Xiaojie Guo. Practical edge detection via robust collaborative learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2526–2534, 2023. 2
- [13] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 2
- [14] Sam Hallman and Charless C. Fowlkes. Oriented edge forests for boundary detection. *Cornell University - arXiv*, Dec 2014. 1, 2
- [15] Sam Hallman and Charless C Fowlkes. Oriented edge forests for boundary detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1732–1740, 2015. 6, 7
- [16] Jianzhong He, Shiliang Zhang, Ming Yang, Yanhu Shan, and Tiejun Huang. Bdcn: Bi-directional cascade network for perceptual edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):100–113, 2022. 1, 2, 6, 7
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [18] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 4, 5
- [19] Tao Huang, Xiaohuan Pei, Shan You, Fei Wang, Chen Qian, and Chang Xu. Localmamba: Visual state space model with windowed selective scan. *arXiv preprint arXiv:2403.09338*, 2024. 2
- [20] Kokkinos Iasonas. Pushing the boundaries of boundary detection using deep learning. *arXiv preprint arXiv:1511.07386*, 2015. 7
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5
- [22] Mingchun Li, Dali Chen, and Shixin Liu. Beta network for boundary detection under nondeterministic labels. *Knowledge-Based Systems*, 266:110389, 2023. 2
- [23] Weibin Liao, Yinghao Zhu, Xinyuan Wang, Cehngwei Pan, Yasha Wang, and Liantao Ma. Lightm-unet: Mamba assists in lightweight unet for medical image segmentation. *arXiv preprint arXiv:2403.05246*, 2024. 3
- [24] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Jia-Wang Bian, Le Zhang, Xiang Bai, and Jinhui Tang. Richer convolutional features for edge detection. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 41(08):1939–1946, 2019. 1, 2, 6, 7, 8
- [25] Yu Liu and Michael S Lew. Learning relaxed deep supervision for better edge detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 231–240, 2016. 7
- [26] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. *Advances in neural information processing systems*, 29, 2016. 1
- [27] Xavier Soria Poma, Edgar Riba, and Angel Sappa. Dense extreme inception network: Towards a robust cnn model for edge detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1923–1932, 2020. 2, 6
- [28] Mengyang Pu, Yaping Huang, Yuming Liu, Qingji Guan, and Haibin Ling. Edter: Edge detection with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1402–1412, 2022. 1, 2, 4, 6, 7
- [29] Xiaofeng Ren and Liefeng Bo. Discriminatively trained sparse code gradients for contour detection. In *Proceedings*

- of the 25th International Conference on Neural Information Processing Systems-Volume 1, pages 584–592, 2012. 6, 7
- [30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 4
- [31] Mohsen Sharifi, Mahmood Fathy, and Maryam Tayefeh Mahmoudi. A classified and comparative study of edge detection algorithms. In *Proceedings. International conference on information technology: Coding and computing*, pages 117–120. IEEE, 2002. 1, 2
- [32] Wei Shen, Xinggang Wang, Yan Wang, Xiang Bai, and Zhi-jiang Zhang. Deepcontour: A deep convolutional feature learned by positive-sharing loss for contour detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3982–3991, 2015. 7
- [33] Zhu Shengjie, Brazil Garrick, and Liu Xiaoming. The edge of depth: Explicit constraints between segmentation and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13116–13125, 2020. 1
- [34] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pages 746–760. Springer, 2012. 2, 6
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 4
- [36] Irwin Sobel, Gary Feldman, et al. A 3x3 isotropic gradient operator for image processing. *a talk at the Stanford Artificial Project in*, pages 271–272, 1968. 1, 2
- [37] Xavier Soria, Gonzalo Pomboza-Junez, and Angel Domingo Sappa. Ldc: Lightweight dense cnn for edge detection. *IEEE Access*, 10:68281–68290, 2022. 1
- [38] Xavier Soria, Angel Sappa, Patricio Humanante, and Arash Akbarinia. Dense extreme inception network for edge detection. *Pattern Recognition*, 139:109461, 2023. 6
- [39] Zhuo Su, Jiehua Zhang, Longguang Wang, Hua Zhang, Zhen Liu, Matti Pietikäinen, and Li Liu. Lightweight pixel difference networks for efficient visual representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 7
- [40] Xiao Wang, Shiao Wang, Yuhe Ding, Yuehang Li, Wentao Wu, Yao Rong, Weizhe Kong, Ju Huang, Shihao Li, Haoxiang Yang, et al. State space model for new-generation network alternative to transformers: A survey. *arXiv preprint arXiv:2404.09516*, 2024. 2
- [41] Yupei Wang, Xin Zhao, and Kaiqi Huang. Deep crisp boundaries. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3892–3900, 2017. 7
- [42] Renkai Wu, Yinghao Liu, Pengchen Liang, and Qing Chang. Ultralight vm-unet: Parallel vision mamba significantly reduces parameters for skin lesion segmentation. *arXiv preprint arXiv:2403.20035*, 2024. 3
- [43] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. *International Journal of Computer Vision*, 125(1):3–18, 2017. 1, 2, 6, 7
- [44] Dan Xu, Wanli Ouyang, Xavier Alameda-Pineda, Elisa Ricci, Xiaogang Wang, and Nicu Sebe. Learning deep structured multi-scale features using attention-gated crfs for contour prediction. *Advances in neural information processing systems*, 30, 2017. 7
- [45] Jiacong Xu, Zixiang Xiong, and Shankar P Bhattacharyya. Pidnet: A real-time semantic segmentation network inspired from pid controller. *arXiv preprint arXiv:2206.02066*, 2022. 1
- [46] Wenjie Xuan, Shaoli Huang, Juhua Liu, and Bo Du. Fcl-net: Towards accurate edge detection via fine-scale corrective learning. *Neural Networks*, 145:248–259, 2022. 2, 7
- [47] Jimei Yang, Brian Price, Scott Cohen, Honglak Lee, and Ming-Hsuan Yang. Object contour detection with a fully convolutional encoder-decoder network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 193–202, 2016. 7
- [48] Yuhuan Yang, Chaofan Ma, Jiangchao Yao, Zhun Zhong, Ya Zhang, and Yanfeng Wang. Remember: Referring image segmentation with mamba twister. *arXiv preprint arXiv:2403.17839*, 2024. 2
- [49] Yunfan Ye, Kai Xu, Yuhang Huang, Renjiao Yi, and Zhiping Cai. Diffusionedge: Diffusion probabilistic model for crisp edge detection. *arXiv preprint arXiv:2401.02032*, 2024. 1, 2, 7
- [50] Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M Hospedales, and Chen-Change Loy. Sketch me that shoe. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 799–807, 2016. 1
- [51] Hanwei Zhang, Ying Zhu, Dan Wang, Lijun Zhang, Tianxiang Chen, Ziyang Wang, and Zi Ye. A survey on visual mamba. *Applied Sciences*, 14(13):5683, 2024. 2
- [52] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnet: Edge guidance network for salient object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8779–8788, 2019. 1
- [53] Caixia Zhou, Yaping Huang, Mengyang Pu, Qingji Guan, Ruoxi Deng, and Haibin Ling. Muge: Multiple granularity edge detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25952–25962, 2024. 1, 2, 4, 5, 6, 7, 8
- [54] Caixia Zhou, Yaping Huang, Mengyang Pu, Qingji Guan, Li Huang, and Haibin Ling. The treasure beneath multiple annotations: An uncertainty-aware edge detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15507–15517, 2023. 1, 2, 4, 5, 6, 7, 8
- [55] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*, 2024. 2, 3