

GAM-Depth: Self-Supervised Indoor Depth Estimation Leveraging a Gradient-Aware Mask and Semantic Constraints

Anqi Cheng¹, Zhiyuan Yang¹, Haiyue Zhu² and Kezhi Mao^{1*}

Abstract—Self-supervised depth estimation has evolved into an image reconstruction task that minimizes a photometric loss. While recent methods have made strides in indoor depth estimation, they often produce inconsistent depth estimation in textureless areas and unsatisfactory depth discrepancies at object boundaries. To address these issues, in this work, we propose GAM-Depth, developed upon two novel components: gradient-aware mask and semantic constraints. The gradient-aware mask enables adaptive and robust supervision for both key areas and textureless regions by allocating weights based on gradient magnitudes. The incorporation of semantic constraints for indoor self-supervised depth estimation improves depth discrepancies at object boundaries, leveraging a co-optimization network and proxy semantic labels derived from a pretrained segmentation model. Experimental studies on three indoor datasets, including NYUv2, ScanNet, and InteriorNet, show that GAM-Depth outperforms existing methods and achieves state-of-the-art performance, signifying a meaningful step forward in indoor depth estimation. Our code will be available at <https://github.com/AnqiCheng1234/GAM-Depth>.

I. INTRODUCTION

Depth estimation is a pivotal task in the field of robot vision, serving as a crucial component task for several applications, including robot navigation [1], [2], [3], robot manipulation [4], [5], [6], and augmented reality [7]. Although supervised methods [8], [9], [10], [11], [12] have made progress in depth estimation, they demand a large number of labeled RGB-depth pairs during model training, which may not always be viable. To address this limitation, self-supervised learning methods [13], [14], [15], [16], [17], [18] have emerged as popular solutions, as they only demand stereo pairs or monocular sequences. Most of these methods reframe depth estimation as an image reconstruction task, estimating the depth and camera poses simultaneously to perform inverse warping on RGB frames.

Despite efficacy of the self-supervised methods in outdoor autonomous driving scenarios (e.g., KITTI [19]), their performance tends to decline in indoor environments (e.g., NYUv2 [20]). This is partly attributed to the presence of *large textureless regions* in indoor environments, such as walls, ceilings, and floors, whose primary supervision signal (photometric loss) becomes unreliable due to lack of texture




Fig. 1. Depth estimation comparisons. (a) Input, (b) StructDepth [22], (c) Ours, (d) Ground truth, (e) 1. our proposed gradient-aware mask within the blue box region. The gray level of each pixel corresponds to its weight, with closer to white indicating higher weights. 2. The binary mask of keypoints [22] with textureless regions being neglected. 3. Semantic constraints with colors representing proxy semantic labels.

information [21]. To address this issue, notable works such as P²Net [21] and StructDepth [22] restrict the calculation of photometric loss to texture regions represented by keypoints, preventing non-texture areas from "polluting" the photometric loss. Although these attempts produce improved results, we argue that ignoring textureless regions during model training may not be an ideal solution. As illustrated in Fig 1(b), the absence of supervision in textureless regions can result in depth inconsistencies in these areas. Such inconsistent depth prediction for textureless regions such as floors and walls may hinder the robot's obstacle detecting and motion planning processes during indoor navigation [1], [2], [3].

Another issue of the current methods is the inaccurate prediction of *depth discrepancies at object boundaries*, particularly with small and transparent objects. As illustrated in Fig 1(b), blurred boundaries are produced due to inaccurate prediction of depth discrepancies, which could potentially lead to complications in tasks such as robot-arm grasping and robotic manipulation [4], [5], [6]. We believe this problem is due to the limited capability of the feature extractor in the depth backbone, failing to capture object boundary information. To address this issue, we incorporate semantic constraints into feature extractor, aiming to improve the depth discrepancies at object boundaries.

In this work, we propose GAM-Depth, a comprehensive solution that attempts to provide adaptive and robust supervision signals for both textureless regions and key areas meanwhile incorporating semantic constraints into depth boundary predictions. Firstly, we develop a gradient-aware mask that allocates smooth weights to all pixels based on the magnitude of gradient, without having to identify keypoints by thresholding the gradient magnitude. This will balance the

*Corresponding author

¹A. Cheng, Z. Yang and K. Mao are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore (e-mail: {anqi002, yang0674, ekzmao}@ntu.edu.sg).

²H. Zhu is with Singapore Institute of Manufacturing Technology (SIMTech), Agency for Science, Technology and Research (A*STAR), 2 Fusionopolis Way, Singapore, 138634, Singapore.(e-mail: zhu_haiyue@simtech.a-star.edu.sg)

supervision strength in photometric loss for both the texture-rich and texture-less areas. Secondly, drawing inspiration from [23], [24], we incorporate semantic constraints into the depth estimation backbone, aiming for clear depth boundaries. This is accomplished by sharing a feature extractor between a semantic segmentation task and our depth prediction task. Recognizing the crucial role of pose estimation in indoor depth estimation, our model synergizes residual pose optimization [25] to ensure that depth predictions are optimized in the correct direction.

Experiments are conducted on three indoor datasets: NYUv2 [20], ScanNet [26], and InteriorNet [27]. The results demonstrate that GAM-Depth outperforms existing state-of-the-art methods and improves the δ_1 threshold accuracy by a large margin. Fig. 1(b) and (c) show that our approach excels in predicting smooth depth in textureless regions and clear depth boundaries, highlighted in red boxes. The main contributions are as follows:

- We propose a gradient-aware mask in the photometric loss calculation, providing adaptive and robust supervision for both key areas and textureless regions. This gradient-aware mask can be combined with other self-supervised depth estimation methods to improve their performance.
- We incorporate semantic constraints into depth estimation through a co-optimization network and proxy semantic labels of a pretrained lightweight segmentation model. Our approach pioneers the use of semantic constraints in indoor self-supervised depth estimation.
- The proposed model achieves state-of-the-art results on NYUv2 [20], ScanNet [26], and InteriorNet [27].

II. RELATED WORK

A. Self-supervised Monocular Depth Estimation

Although existing supervised monocular depth estimation methods [8], [10], [11], [12], [28], [29] have made significant progress, they demand large quantity of annotated data, which may not be always available. To address this limitation of the supervised methods, [13] introduced the first self-supervised depth estimation method that takes a stereo pair as the input. Godard et al. [14] introduced a left-right consistency loss, ensuring both images in the stereo pair had consistent depth accuracy. The paradigm was then further revolutionized by [15], which demonstrated that promising results could be achieved with just monocular video sequences. As one of the most representative self-supervised depth estimation methods, Monodepth2 [16] generates accurate and sharp depth maps leveraging an encoder-decoder network with skip connections. The main idea is to employ a minimum reprojection loss to handle occluded and out-of-view pixels while dealing with moving objects with automasking. Following this groundbreaking work, several attempts [30], [31] have been made to enhance depth estimation performance for autonomous driving scenarios. [32] presented a novel approach that combines discriminative self-supervised learning techniques for robust feature learning on a large, curated dataset.

Most aforementioned methods were developed in the context of outdoor scenes. For indoor environments, MovingIndoor [33] applied optical flow detection to provide an additional supervisory signal, [34] introduced an auto-rectify network targeting the noise emerging from rotational variations between consecutive frames. P²Net [21] first extracts keypoints by detecting pixels with large gradients, and then limit the photometric loss calculation exclusively to supported regions of the keypoints only. Expanding on this, StructDepth [22] explored indoor structural regularities by a Manhattan normal detection module and a planar region detection module. This ensured the predicted surface normals aligned with main directions while simultaneously generating planar segments. MonoIndoor [25] proposed depth estimation accuracy via a residual pose estimation module, recognizing the challenges of predicting accurate relative poses due to complex indoor rotational motions. Extending this idea further, HI-Net [35] incorporated a homography estimation module and an iterative refinement process. SC-DepthV3 [36] and DistDepth [37] made use of an off-the-shelf depth estimator to generate a single-image depth prior.

B. Semantic Constraints for Depth Estimation

Recent works in depth estimation have witnessed incorporation of semantic constraints for performance enhancement. Hambarde et al. [38] proposed a semantic-to-depth generative adversarial network, while [39] minimized the border consistency between depth and segmentation. Wang et al. [40] leveraged both semantic and instance segmentation to predict category-agnostic and instance-level depth maps. Klinger et al. [41] performed cross-task guidance by incorporating semantic segmentation in a supervised mode. A shared feature extractor is employed to close the gap between the two tasks. Our method is inspired by [23], which jointly predicted depth, relative pose, optical flow, semantic masks, and object labels to perform comprehensive scene understanding, utilizing a proxy semantic network. Unlike [23], which detects dynamic classes to exclude moving objects in autonomous driving scenes, we focus on enhancing depth estimation with semantic constraints in indoor environments.

III. METHOD

A. Method Overview

In this work, we develop GAM-Depth, comprising two new components: gradient-aware mask and semantic constraints. The gradient-aware mask M_{gra} enhances depth estimation consistency by providing robust supervision signals for both textureless regions and texture-rich areas (introduced in Sec. III-B) while the semantic constraints L_{seg} improve boundary depth discrepancies by merging low-level details with high-level scene comprehension (introduced in Sec. III-C). For textureless regions with large proportions in indoor environments, we assign small and smooth weights instead of excluding them from supervision [21], [22]. For texture-rich regions with abrupt gradient changes, the gradient-aware mask is also able to assign smooth weights adaptively instead of a uniform allocation [21], [22]. For object boundaries,

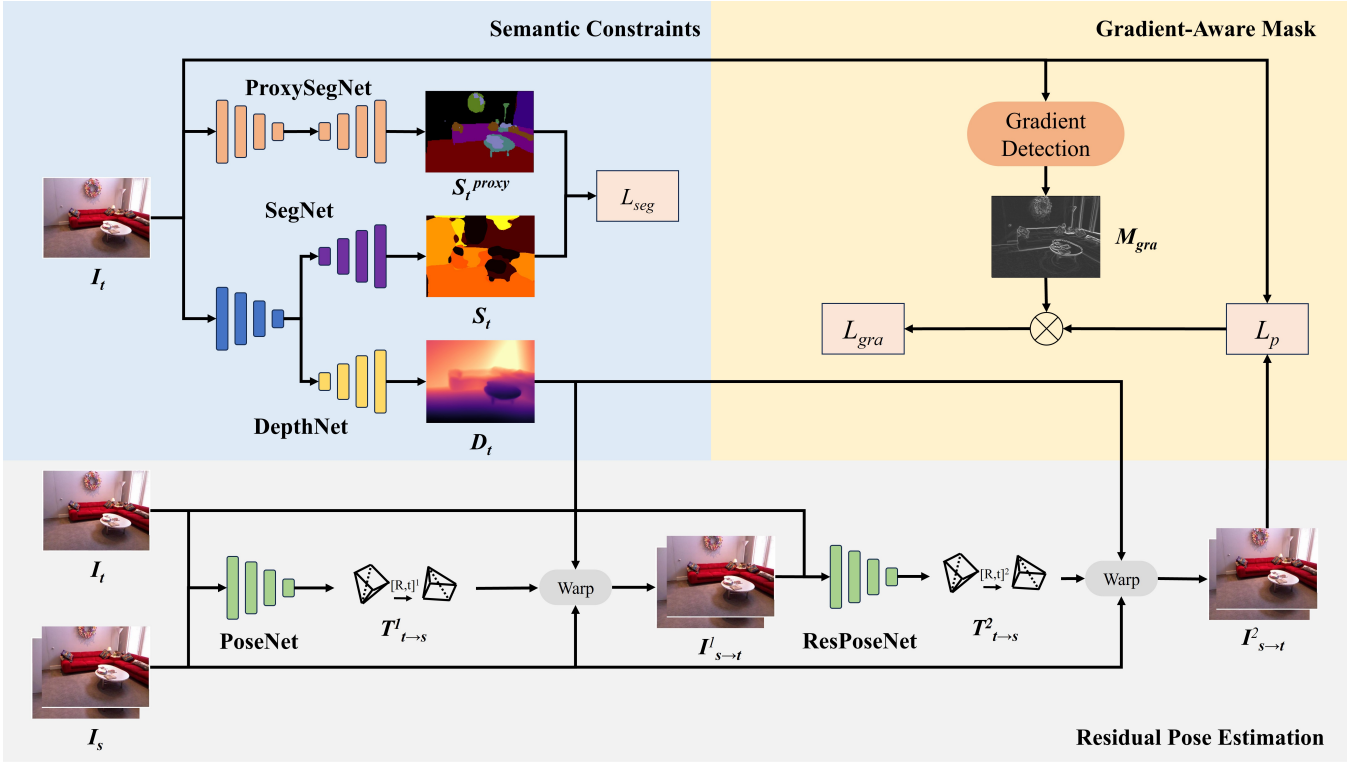


Fig. 2. Overall training pipeline of GAM-Depth. Given a target frame I_t , DepthNet estimates its depth map D_t and SegNet predicts its semantic labels S_t . DepthNet shares its encoder with SegNet. ProxySegNet predicts proxy semantic labels S_t^{proxy} to supervise the training of GAM-Depth through segmentation loss L_{seg} . The reference frame I_s is warped to I_t 's view through PoseNet and ResPoseNet proposed by MonoIndoor [25]. Our gradient-aware mask M_{gra} is generated by a gradient detection method. Finally, a gradient-aware photometric loss L_{gra} is calculated as the multiplication of the photometric loss L_p between I_t and the warped frame $I_{s \rightarrow t}$ and M_{gra} .

under the guidance of the semantic constraints, clearer depth discrepancies are produced.

The overall training pipeline of GAM-Depth is shown in Fig. 2. In our self-supervised training architecture, monocular sequences serve as the exclusive inputs. DepthNet estimates the depth map D_t of the target frame I_t , while the SegNet predicts its semantic labels S_t simultaneously. These two networks share an encoder for a comprehensive understanding of the scene. ProxySegNet, which is an off-the-shelf lightweight semantic segmentation model [42], provides proxy semantic labels S_t^{proxy} to guide the training of SegNet. PoseNet and ResPoseNet are utilized to generate accurate relative poses $T_{t \rightarrow s}^1$ between the target frame I_t and the reference frame I_s . With D_t and $T_{t \rightarrow s}^1$ being predicted, following [15], the warping process can be described as:

$$p_s \sim \mathbf{K} \mathbf{T}_{t \rightarrow s}^1 \mathbf{D}_t(p_t) \mathbf{K}^{-1} p_t \quad (1)$$

where p_t and p_s are pixels of I_t and I_s , and \mathbf{K} denotes the camera intrinsic matrix. The warped frame $I_{s \rightarrow t}$ can be reconstructed following [15]:

$$I_{s \rightarrow t} = I_t \langle \text{proj}(\mathbf{D}_t, \mathbf{T}_{t \rightarrow s}^1, \mathbf{K}) \rangle \quad (2)$$

where $\langle \cdot \rangle$ is a locally sub-differentiable bilinear sampling operator [43]. Our gradient-aware mask M_{gra} is generated through a smoothing function on the gradient magnitude, which contributes to the photometric loss L_p calculation.

B. Gradient-Aware Mask

In self-supervised training, the absence of ground truth depth information necessitates the reliance on the photometric loss L_p [14], which is a weighted combination of the L1 and SSIM loss [44]:

$$L_p(I_t, I_{s \rightarrow t}) = (1 - \alpha) L_1(I_t, I_{s \rightarrow t}) + \alpha \text{SSIM}(I_t, I_{s \rightarrow t}) \quad (3)$$

where I_t , I_s , and $I_{s \rightarrow t}$ are the target frame, reference frame, and warped frame, respectively. α denotes the weight of SSIM loss [44].

The minimum reprojection loss L_{min} , as proposed by Monodepth2 [16], proves to be effective in self-supervised learning for its ability to mask out occluded and out-of-view pixels:

$$L_{min} = \sum_{p_t} \min_s (L_p(I_t, I_{s \rightarrow t})) \quad (4)$$

where s is the index of reference frames.

However, L_{min} often fails in indoor environments with the presence of large textureless regions. To address this issue, P²Net [21] and StructDepth [22] first extract keypoints [45] based on pixel gradient magnitude thresholding, and then warp to calculate L_{min} within the local support regions of these keypoints. Despite the success of these approaches, we identify two limitations: a lack of robustness in keypoints selection and an absence of supervision in textureless regions.




Fig. 3. Different weights assignment between our gradient-aware mask M_{gra} and keypoints-only photometric loss. (a) M_{gra} allocates different weights to various pixels based on their gradient magnitudes m , providing more robust supervision for both textureless regions and key areas. (b) P²Net [21] and StructDepth [22] allocate weight of 1 to keypoints and 0 to non-keypoints.

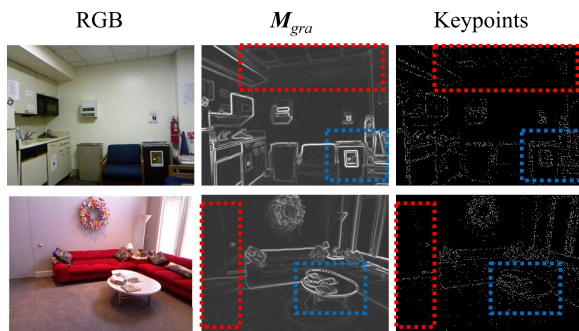


Fig. 4. Comparisons between our gradient-aware mask M_{gra} and keypoints detected by DSO [45]. The gray level of each pixel corresponds to its weight, with closer to white indicating higher weights. Our M_{gra} provides adaptive supervision at both textureless (indicated by red boxes) and texture-rich regions (indicated by blue boxes). Weights of textureless regions are completely excluded in P²Net [21] and StructDepth [22].

To address these limitations, we propose a gradient-aware mask M_{gra} and a gradient-aware photometric loss L_{gra} :

$$M_{gra} = \beta + \frac{1 - \beta}{1 + e^{-\gamma_1 \cdot m + \gamma_2}} \quad (5)$$

$$L_{gra} = \sum_{p_t} M_{gra}(p_t) \min_s (L_p(\mathbf{I}_t, \mathbf{I}_{s \rightarrow t})) \quad (6)$$

where m denotes pixel gradient magnitude, and β , γ_1 , and γ_2 are hyperparameters.

As shown in Fig. 3, M_{gra} adaptively and robustly allocates different weights to various pixels, contingent on their gradient magnitudes m . Specifically, pixels with larger gradient magnitude, often indicating keypoints marked by significant RGB shift, receive higher weights during photometric loss calculation. Conversely, pixels with smaller gradient magnitudes, indicating textureless regions, are assigned lower weights, rather than being entirely ignored in P²Net [21] and StructDepth [22]. Gradient magnitudes m are calculated using a 3x3 Sobel operator. Fig. 4 demonstrates the difference between our gradient-aware mask M_{gra} and keypoints detected by DSO [45]. The gray level of each pixel corresponds to its weight, with closer to white indicating higher weights. As shown in red boxes, in contrast to simply neglecting textureless regions, M_{gra} provides supervision for walls, ceilings, and floors.

Remarkably, M_{gra} not only elevates our baseline’s performance but also enhances the depth estimation accuracy of other indoor self-supervised methods [36] when combined with them. Details are presented in Sec. IV-D.2.

C. Semantic Constraints

Different from some established methods [41], [23] for outdoor scenes which utilize semantic and instance segmentation to detect dynamic objects, our approach pioneers the integration of semantic segmentation with depth estimation in an indoor self-supervised framework. Specifically, this is implemented through a shared encoder between DepthNet and SegNet, augmented by a cross-entropy segmentation loss L_{seg} in the final loss function.

We employ the Light-weight RefineNet [42], a lightweight and pretrained semantic segmentation model, as our Proxy-SegNet. To follow the principles of self-supervised training, neither depth nor semantic segmentation ground-truth information is required during the training process. The ProxySegNet produces proxy semantic labels S_t^{proxy} , guiding SegNet’s predicted semantic labels S_t and refining depth estimation by capturing both low-level and high-level features through a shared encoder.

The cross-entropy segmentation loss L_{seg} can be described as:

$$L_{seg} = - \sum_c \sum_{p_t} (S_t(p_t) \log(S_t^{proxy}(c, p_t))) \quad (7)$$

where c denotes class labels.

The shared encoder empowers the model with holistic scene comprehension by capturing shared features. The decoder of SegNet has the same architecture as DepthNet’s decoder except for the last Softmax layer. Their weights are learned independently. To focus on depth training, L_{seg} is reduced to 0 at the middle stage of training.

D. Final Training Loss

We combine the L_{gra} and L_{seg} with loss proposed in the baseline:

$$L_{total} = L_{gra} + \lambda_1 L_{seg} + \lambda_2 L_{smooth} + \lambda_3 L_{norm} + \lambda_4 L_{planar} \quad (8)$$

where L_{smooth} , L_{norm} and L_{planar} are the smoothness loss, the Manhattan normal loss, and the co-planar loss proposed by StructDepth [22]. λ_1 , λ_2 , λ_3 , and λ_4 are the hyperparameter weights for corresponding losses.

IV. EXPERIMENTS

In this section, we test the performance of GAM-Depth on NYUv2 [20] and its generalization ability on ScanNet [26] and InteriorNet [27]. We also conduct ablation studies to validate the effectiveness of each individual component.

A. Datasets

Following P²Net [21] and StructDepth [22], we evaluate our model using three publicly available indoor scene datasets: NYUv2, ScanNet, and InteriorNet. Our training is based on NYUv2.

TABLE I

QUALITATIVE RESULT ON NYUV2 [20]. THE BEST RESULTS IN EACH CATEGORY ARE IN **BOLD**. \downarrow INDICATES THE LOWER THE BETTER WHILE \uparrow INDICATES THE HIGHER THE BETTER. \checkmark DENOTES SUPERVISED LEARNING, Δ DENOTES ADDITIONAL SUPERVISION FROM PRETRAINED DEPTH ESTIMATION MODEL, AND \times DENOTES SELF-SUPERVISED LEARNING

Methods	Supervision	AbsRel \downarrow	RMS \downarrow	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
Eigen et al. [8]	\checkmark	0.158	0.641	0.769	0.950	0.988
Laina et al. [10]	\checkmark	0.127	0.573	0.811	0.953	0.988
DORN [11]	\checkmark	0.115	0.509	0.828	0.965	0.992
BTS [12]	\checkmark	0.110	0.392	0.885	0.987	0.994
P3Depth [29]	\checkmark	0.104	0.356	0.898	0.981	0.996
NewCRFs [28]	\checkmark	0.095	0.334	0.922	0.992	0.998
DistDepth [37]	Δ	0.130	0.517	0.832	0.963	0.990
SC-Depthv3 [36]	Δ	0.123	0.486	0.848	0.963	0.991
Monodepth2 [16]	\times	0.161	0.600	0.771	0.948	0.987
P ² Net [21]	\times	0.150	0.561	0.796	0.948	0.986
StructDepth [22]	\times	0.142	0.540	0.813	0.954	0.988
Bian et al. [34]	\times	0.138	0.532	0.820	0.956	0.989
MonoIndoor [25]	\times	0.134	0.526	0.823	0.958	0.989
HI-Net [35]	\times	0.130	0.522	0.828	0.959	0.989
Ours(GAM-Depth)	\times	0.131	0.507	0.836	0.960	0.990




Fig. 5. Qualitative results on NYUv2 [20]. RGB images, P²Net [21], StructDepth [22], our predictions, and ground truth depth maps are presented for comparison. GAM-Depth obtains more accurate results at object boundaries and textureless regions, as indicated by the red circles.

1) *NYUv2*: The NYUv2 dataset comprises 582 indoor scenes captured by a Microsoft Kinect camera. Following the split of [8] and [22], 283 scenes (approximately 230K images) are used for training. The training set is sampled at intervals of 10 frames, using ± 10 and ± 20 frames as reference frames. After sampling, the total number of training samples is reduced to 21465. Following the official split, the validation set contains 654 dense labeled RGB-depth pairs.

2) *ScanNet*: ScanNet comprises approximately 2.5 million RGBD videos captured across 1513 scenes using a depth camera attached to an iPad.

3) *InteriorNet*: InteriorNet is a synthetic dataset containing millions of video sequences representing various interior design layouts, furniture, and object models.

B. Implementation Details

The model is implemented using PyTorch [46]. To allow a fair comparison with previous models, GAM-Depth adopts the implementation of StructDepth [22], training the network for 50 epochs based on P²Net [21]. We adopt Adam [47] as the optimizer. Due to limited computing resources, the batch size is decreased from 32 to 16, which impacts the results negatively. For fair comparisons with existing

TABLE II

SCANNET [26] RESULTS WITH THE TRAINED MODEL ON NYUV2 [20]

Methods	AbsRel ↓	RMS ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [16]	0.191	0.451	0.693	0.926	0.983
P ² Net [21]	0.175	0.420	0.740	0.932	0.982
StructDepth [22]	0.165	0.400	0.754	0.939	0.985
GAM-Depth	0.159	0.390	0.770	0.939	0.985

TABLE III

INTERIORNET [27] RESULTS WITH THE TRAINED MODEL ON NYUV2 [20]

Methods	AbsRel ↓	RMS ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [16]	0.368	0.817	0.586	0.815	0.898
P ² Net [21]	0.346	0.737	0.642	0.833	0.902
StructDepth [22]	0.330	0.715	0.660	0.840	0.905
GAM-Depth	0.313	0.693	0.676	0.846	0.907

indoor self-supervised methods, the backbone model adopts ResNet18 [48] and pose networks follow MonoIndoor [25]. Following [16], [21] and [22], we normalize the depth to 10m and use the median scaling strategy to avoid the scale ambiguity of self-supervised monocular depth estimation. Hyperparameter α for photometric loss calculation is set to 0.85 as [14], [16]. Our gradient-aware photometric loss L_{gra} weighting factors β is set to 0.1 to balance the weights between textureless and texture-rich regions while γ_1 and γ_2 are set to 0.1 and 40, respectively. Hyperparameters λ_1 , λ_2 , λ_3 , and λ_4 for the final training loss are set to 0.001, 0.1, 0.05, 0.1, respectively.

C. Depth Estimation Results

1) *Quantitative Results*: According to the quantitative results on NYUV2 [20] shown in Tab. I, GAM-Depth achieves the best result and outperforms existing self-supervised indoor depth estimation methods by a large margin, especially on δ_1 accuracy. Even if GAM-Depth is trained in the self-supervised mode, it outperforms some supervised and pretrained model-supervised methods. Tab. II and Tab. III show the depth estimation results on ScanNet [26] and InteriorNet [27] with the model trained on NYUV2. Although they have not been used for training, the results show that GAM-Depth has better generalization ability than existing methods.

2) *Qualitative Results*: Fig. 5 showcases the qualitative results of NYUV2. The depth predictions of GAM-Depth, StructDepth [22], and P²Net [21] are processed with the same median scaling method. Results are better viewed by coloring. The red circles show that GAM-Depth produces accurate object boundaries and continuous depth in textureless regions.

D. Ablation Studies

1) *Effectiveness of M_{gra} and L_{seg}* : Tab. IV shows the improvements of our gradient-aware mask M_{gra} and se-

TABLE IV

ABLATION STUDIES ON OUR GRADIENT-AWARE MASK M_{gra} AND SEMANTIC CONSTRAINTS L_{seg} . OURS(KEY) AND OURS(AVG) REPRESENT REPLACING M_{gra} WITH ONLY CONSIDERING KEYPOINTS' PHOTOMETRIC LOSSES OR AVERAGING ALL PIXELS' LOSSES

Methods	M_{gra}	L_{seg}	AbsRel ↓	RMS ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Ours	✗	✗	0.136	0.526	0.824	0.958	0.989
Ours	✓	✗	0.134	0.520	0.828	0.959	0.989
Ours(key)	✗	✓	0.133	0.516	0.830	0.959	0.989
Ours(avg)	✗	✓	0.132	0.510	0.834	0.960	0.990
Ours(full)	✓	✓	0.131	0.507	0.836	0.960	0.990

TABLE V

ABLATION STUDIES ON THE GENERALIZATION ABILITY OF OUR GRADIENT-AWARE MASK M_{gra}

Methods	AbsRel ↓	RMS ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
StructDepth [22]	0.142	0.540	0.813	0.954	0.988
StructDepth [22] + M_{gra}	0.150	0.567	0.796	0.949	0.986
Ours(baseline)	0.136	0.526	0.824	0.958	0.989
Ours(baseline) + M_{gra}	0.134	0.520	0.828	0.959	0.989
SC-DepthV3 [36]	0.123	0.486	0.848	0.963	0.991
SC-DepthV3 [36] + M_{gra}	0.120	0.471	0.854	0.963	0.991

semantic constraints L_{seg} . Excluding either M_{gra} or semantic constraints L_{seg} shows degradation. Our baseline achieves better performance compared to [22] through residual pose estimations [25]. Replacing our gradient-aware mask M_{gra} with either only considering keypoints or averaging all pixels' losses shows deficiencies.

2) *Generalization ability of M_{gra}* : In this part of the experiments, we apply our M_{gra} to a couple of photometric loss-based self-supervised depth estimation methods, including StructDepth [22], and SC-DepthV3 [36]. As shown in Tab. V, the use of M_{gra} elevates the performance of our baseline and SC-DepthV3, but degrades the performance of StructDepth. The degradation in StructDepth may be due to its limited supervision [22] whereas the improvement in SC-DepthV3 [36] is because of depth guidance from an off-the-shelf depth estimation model, which serves as a strong supervision signal.

V. CONCLUSIONS

In this paper, we have developed GAM-Depth for self-supervised indoor depth estimation, built upon our proposed gradient-aware mask and incorporation of semantic constraints. The gradient-aware mask provides robust supervision for both textureless and texture-rich regions by adaptive weight allocation in the photometric loss calculation. The semantic constraints enable a co-optimization training pipeline that integrates low-level details and high-level scene comprehension. Results have proved that our model can produce smooth depth in textureless regions and sharp depth at boundaries.

REFERENCES

- [1] Y. Pan, P. Xiao, Y. He, Z. Shao, and Z. Li, "MULLS: Versatile LiDAR SLAM via multi-metric linear least square," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 633–11 640.
- [2] S. Kareer, N. Yokoyama, D. Batra, S. Ha, and J. Truong, "ViNL: Visual navigation and locomotion over obstacles," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2018–2024.
- [3] G. Georgakis, B. Bucher, A. Arapin, K. Schmeckpeper, N. Matni, and K. Daniilidis, "Uncertainty-driven planner for exploration and navigation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 295–11 302.
- [4] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396.
- [5] H. Zhu, Y. Li, F. Bai, W. Chen, X. Li, J. Ma, C. S. Teo, P. Y. Tao, and W. Lin, "Grasping detection network with uncertainty estimation for confidence-driven semi-supervised domain adaptation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9608–9613.
- [6] K. Takahashi, W. Ko, A. Ummadisingu, and S.-i. Maeda, "Uncertainty-aware self-supervised target-mass grasping of granular foods," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2620–2626.
- [7] H. Liu, Y. Zhang, W. Si, X. Xie, Y. Zhu, and S.-C. Zhu, "Interactive robot knowledge patching using augmented reality," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1947–1954.
- [8] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in neural information processing systems*, vol. 27, 2014.
- [9] F. Liu, C. Shen, and G. Lin, "Deep convolutional neural fields for depth estimation from a single image," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 5162–5170.
- [10] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, "Deeper depth prediction with fully convolutional residual networks," in *2016 Fourth international conference on 3D vision (3DV)*. IEEE, 2016, pp. 239–248.
- [11] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, "Deep ordinal regression network for monocular depth estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2002–2011.
- [12] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation," *arXiv preprint arXiv:1907.10326*, 2019.
- [13] R. Garg, V. K. Bg, G. Carneiro, and I. Reid, "Unsupervised CNN for single view depth estimation: Geometry to the rescue," in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*. Springer, 2016, pp. 740–756.
- [14] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 270–279.
- [15] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1851–1858.
- [16] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow, "Digging into self-supervised monocular depth estimation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3828–3838.
- [17] Y. Almalioglu, M. R. U. Saputra, P. P. De Gusmao, A. Markham, and N. Trigoni, "GANVO: Unsupervised deep monocular visual odometry and depth estimation with generative adversarial networks," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 5474–5480.
- [18] K. Wang, Y. Chen, H. Guo, L. Wen, and S. Shen, "Geometric pre-training for monocular depth estimation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4782–4788.
- [19] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361.
- [20] N. Silberman, D. Hoiem, P. Kohli, and R. 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*. Springer, 2012, pp. 746–760.
- [21] Z. Yu, L. Jin, and S. Gao, "P2Net: Patch-match and plane-regularization for unsupervised indoor depth estimation," in *European Conference on Computer Vision*. Springer, 2020, pp. 206–222.
- [22] B. Li, Y. Huang, Z. Liu, D. Zou, and W. Yu, "StructDepth: Leveraging the structural regularities for self-supervised indoor depth estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 663–12 673.
- [23] F. Tosi, F. Aleotti, P. Z. Ramirez, M. Poggi, S. Salti, L. D. Stefano, and S. Mattoccia, "Distilled semantics for comprehensive scene understanding from videos," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 4654–4665.
- [24] D. Seichter, M. Köhler, B. Lewandowski, T. Wengefeld, and H.-M. Gross, "Efficient rgb-d semantic segmentation for indoor scene analysis," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 13 525–13 531.
- [25] P. Ji, R. Li, B. Bhanu, and Y. Xu, "Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 787–12 796.
- [26] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3D reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.
- [27] W. Li, S. Saedi, J. McCormac, R. Clark, D. Tzoumanikas, Q. Ye, Y. Huang, R. Tang, and S. Leutenegger, "Interiornet: Mega-scale multi-sensor photo-realistic indoor scenes dataset," *arXiv preprint arXiv:1809.00716*, 2018.
- [28] W. Yuan, X. Gu, Z. Dai, S. Zhu, and P. Tan, "Neural window fully-connected crfs for monocular depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3916–3925.
- [29] V. Patil, C. Sakaridis, A. Liniger, and L. Van Gool, "P3depth: Monocular depth estimation with a piecewise planarity prior," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1610–1621.
- [30] V. Guizilini, R. Ambrus, S. Pillai, A. Raventos, and A. Gaidon, "3d packing for self-supervised monocular depth estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2485–2494.
- [31] J. Watson, O. Mac Aodha, V. Prisacariu, G. Brostow, and M. Firman, "The temporal opportunist: Self-supervised multi-frame monocular depth," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1164–1174.
- [32] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [33] J. Zhou, Y. Wang, K. Qin, and W. Zeng, "Moving indoor: Unsupervised video depth learning in challenging environments," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8618–8627.
- [34] J.-W. Bian, H. Zhan, N. Wang, T.-J. Chin, C. Shen, and I. Reid, "Auto-rectify network for unsupervised indoor depth estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 9802–9813, 2021.
- [35] G. Wu, K. Li, L. Wang, R. Hu, Y. Guo, and Z. Chen, "HI-Net: Boosting self-supervised indoor depth estimation via pose optimization," *IEEE Robotics and Automation Letters*, vol. 8, no. 1, pp. 224–231, 2022.
- [36] L. Sun, J.-W. Bian, H. Zhan, W. Yin, I. Reid, and C. Shen, "SC-DepthV3: Robust self-supervised monocular depth estimation for dynamic scenes," *arXiv preprint arXiv:2211.03660*, 2022.
- [37] C.-Y. Wu, J. Wang, M. Hall, U. Neumann, and S. Su, "Toward practical monocular indoor depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3814–3824.

- [38] P. Hambarde, A. Dudhane, P. W. Patil, S. Murala, and A. Dhall, "Depth estimation from single image and semantic prior," in *2020 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2020, pp. 1441–1445.
- [39] S. Zhu, G. Brazil, and X. Liu, "The edge of depth: Explicit constraints between segmentation and depth," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 13 116–13 125.
- [40] L. Wang, J. Zhang, O. Wang, Z. Lin, and H. Lu, "Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 541–550.
- [41] M. Klingner, J.-A. Termöhlen, J. Mikolajczyk, and T. Fingscheidt, "Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. Springer, 2020, pp. 582–600.
- [42] V. Nekrasov, C. Shen, and I. Reid, "Light-weight refinenet for real-time semantic segmentation," *arXiv preprint arXiv:1810.03272*, 2018.
- [43] M. Jaderberg, K. Simonyan, A. Zisserman *et al.*, "Spatial transformer networks," *Advances in neural information processing systems*, vol. 28, 2015.
- [44] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [45] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [46] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [47] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [48] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.