

# Image Masking for Robust Self-Supervised Monocular Depth Estimation

Hemang Chawla<sup>1</sup>, Kishaan Jeeveswaran<sup>1</sup>, Elahe Arani<sup>1,2\*</sup>, and Bahram Zonooz<sup>1,2\*</sup>

**Abstract**—Self-supervised monocular depth estimation is a salient task for 3D scene understanding. Learned jointly with monocular ego-motion estimation, several methods have been proposed to predict accurate pixel-wise depth without using labeled data. Nevertheless, these methods focus on improving performance under ideal conditions without natural or digital corruptions. The general absence of occlusions is assumed even for object-specific depth estimation. These methods are also vulnerable to adversarial attacks, which is a pertinent concern for their reliable deployment in robots and autonomous driving systems. We propose MIMDepth, a method that adapts masked image modeling (MIM) for self-supervised monocular depth estimation. While MIM has been used to learn generalizable features during pre-training, we show how it could be adapted for direct training of monocular depth estimation. Our experiments show that MIMDepth is more robust to noise, blur, weather conditions, digital artifacts, occlusions, as well as untargeted and targeted adversarial attacks.

## I. INTRODUCTION

Depth estimation is an essential component of vision systems that capture 3D scene structures for applications in mobile robots, self-driving cars, and augmented reality. Although expensive and power-hungry LiDARs offer high-accuracy depth measurements, the ubiquity of low-cost, energy-efficient cameras makes monocular depth estimation techniques a popular alternative. Traditional methods often estimate depth from multiple views of the scene [1]. Instead, deep learning methods have demonstrated depth estimation from a single image. Nevertheless, supervised depth estimation approaches [2], [3] require ground truth labels, making it difficult to scale. On the contrary, self-supervised depth estimation approaches are trained without ground truth labels by using concepts from traditional structure-from-motion and offer the possibility of training on a wide variety of data [4], [5]. However, the deployment requires a focus on the generalizability and robustness of models beyond performance under ideal conditions [6].

Recently, MT-SfMLearner [7] showed that the transformer architecture for self-supervised depth estimation results in higher robustness to image corruptions as well as against adversarial attacks. This is attributed to transformers utilizing the global context of the scene for predictions, unlike convolutional neural networks that have a limited receptive field. However, most research in self-supervised monocular depth estimation focuses primarily on achieving excellent

performance on the independent and identically distributed (i.i.d.) test set. It is assumed that the images are free from noise (e.g. Gaussian) and blur (e.g. due to ego-motion or moving objects in the scene), have clear daylight weather, and are without digital artifacts (e.g. pixelation). Even for the task of object-specific depth estimation [8], it is assumed that the objects are without occlusions. Finally, the robustness of methods against adversarial attacks is not considered, which is a pertinent concern for safety while deploying deep learning models.

Since the performance and robustness of the models are determined by the learned representations, influencing the encoding of features could lead to more robust estimations. We hypothesize that integrating Masked Image Modeling (MIM) into the training of depth estimation would result in learning features that make the model more robust to natural and digital corruptions as well as against adversarial attacks, by modeling the global context in a better way. MIM is a technique that has been used until now for self-supervised representation learning in pre-training of image transformers [9]–[12]. MIM pre-training involves masking a portion of image patches and then using the unmasked portion to predict the masked input [10], [11] or its features [9], [12]. It models long-range dependencies and focuses on the low-frequency details present in the images. However, when pre-trained models are fine-tuned for downstream training, the general features that were learned could possibly be overwritten. Instead, adapting MIM for direct training of a task, such as depth estimation, could lead to richer learned representations that make the model more robust and generalizable.

While both MIM and depth estimation are self-supervised methods, they differ in how they are trained. MIM, used in pre-training of a network, learns by reconstructing the input image generally passed through an autoencoder. Instead, the self-supervised depth estimation network is trained along with a self-supervised ego-motion estimation network, whose output is used to synthesize adjacent images in the training set via the perspective projection transform [13]. Thus, applying MIM to depth estimation requires different considerations than its use for pre-training. Here, we examine the following questions:

- Would integrating MIM to the depth and/or ego-motion estimation networks result in improved robustness of depth estimation?
- MIM has been shown to work well with either block-wise masking [9], or random masking with high mask size [10]. Which masking strategy would work better for depth estimation?
- MIM pre-training uses a relatively high mask ratio and

Authors are with <sup>1</sup>Advanced Research Lab, NavInfo Europe, The Netherlands, and <sup>2</sup>Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands.

Contact: hemang.chawla@navinfo.eu

\*Contributed equally.

Code: <https://github.com/NeurAI-Lab/MIMDepth>

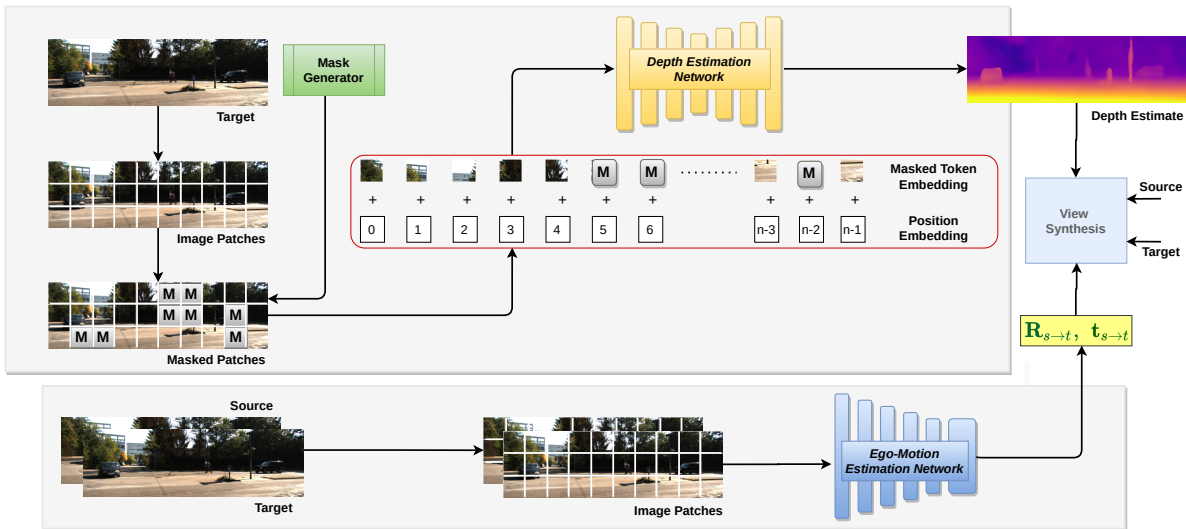


Fig. 1: An overview of Masked Image Modeling for Depth Estimation. The method learns to predict depth from the masked as well as unmasked patches with a better understanding of the global context.

mask size [9]–[11] due to more information redundancy in images than in sentences. Would the high mask ratio or high mask size used for MIM pre-training be suitable for integrating it into depth estimation?

- MIM has been shown to result in better features for downstream tasks when its loss is applied only to masked regions [10]. Would similarly applying the loss on only the masked regions result in a more robust depth estimation?

With our proposed method MIMDepth, we demonstrate that applying blockwise masking with a relatively lower mask ratio (than MIM pre-training) only to the depth estimation network, with a loss on the complete image, results in improved robustness to natural and digital corruptions, occlusions, as well as untargeted and targeted adversarial attacks. It is additionally found to improve the performance of the ego-motion estimation network, while maintaining a competitive performance of the depth estimation network on the i.i.d. test sets.

## II. RELATED WORKS

**Self-supervised Monocular Depth Estimation** One of the challenging tasks of interest in 3D scene understanding is monocular depth estimation. Although self-supervised depth estimation was introduced for stereo pairs [14], it was soon extended to a monocular setup [13]. Monocular self-supervised approaches to depth estimation have the advantage of not requiring any labels and can learn from a wide variety of data from multiple sources. Over the years, improvements have been made to deal with challenges due to occlusions [15], dynamic objects [16]–[18], and scale-consistency issues [4], [19] and more. While most methods generally used 2D convolutional architectures, a 3D convolutional architecture was proposed to estimate depth from symmetrical packing and unpacking blocks that preserve depth details [5]. Recently, MTSfMLearner [7] has

shown that transformers can also be used for depth and pose estimation resulting in comparable performance, but improved robustness to natural corruptions and adversarial attacks due to their global receptive fields. Although other methods that use transformers have also been proposed [20], [21], they do not consider the robustness of their proposed approaches. We show that integrating mask image modeling trains networks to identify long-range dependencies and could further improve the robustness of depth estimation.

**Masked Image Modeling** Masked image modeling is a method for self-supervised representation learning through images corrupted by masking. This was developed following masked language modeling (MLM) [22]. These methods are based on replacing a portion of the tokenized input sequence with learnable mask tokens and learning to predict the missing context using only the visible context. iGPT [23] operates on clustered pixel tokens and predicts unknown pixels directly. ViT [24] explores masked patch prediction by predicting the mean color. On the contrary, BEIT [9] operates on image patches but uses an additional discrete Variational AutoEncoder (dVAE) tokenizer [25] to learn to predict discrete tokens corresponding to masked portions. BEIT uses special blockwise masking that mitigates the wastage of modeling capabilities on short-range dependencies and high-frequency details similar to BERT. Instead, SiMMIM [10] and MAE [11] show that even random masking with a higher mask ratio or mask size can similarly perform well for self-supervised pretraining from image data. However, MIM has not been explored for directly training a relevant task of interest rather than as a pre-training method. We demonstrate how it can be adapted to improve the robustness of self-supervised monocular depth estimation.

## III. METHOD

We propose MIMDepth, a method for masked self-supervised monocular depth estimation (see Figure 1) that in-

tegrates masked image modeling (MIM) into self-supervised monocular depth estimation. MIMDepth trains the depth estimation network by masking a portion of the input, while the network learns to predict complete pixelwise depth from the partial observations, including for the masked portions. Hereafter, we elaborate upon the components of the method.

### A. Self-supervised Monocular Depth Estimation

We use MT-SfMLearner [7] as the baseline monocular depth estimation method, which we briefly summarize here. Similar to other self-supervised methods, it trains depth and ego-motion estimation networks simultaneously. The training input to the networks is a set of consecutive image triplets  $\{I_{-1}, I_0, I_1\}$  from a video. The depth network is trained to output pixel-wise depth for the target image  $I_0$ , while the ego-motion estimation network is trained to output relative translation and rotation between the image pairs  $\{I_{-1}, I_0\}$  and  $\{I_0, I_1\}$ . Note that for the ego-motion estimation network, the input pair of images is concatenated along the channel dimension. Both the depth and ego-motion estimation networks utilize a transformer-based encoder DeiT-Base [26] which processes the tokenized inputs at constant resolution.

The predicted depth and ego-motion are linked together via the perspective projection transform for a given set of camera intrinsics  $K$ , through which the source images  $I_{-1}$  and  $I_1$  are warped to synthesize the target image  $I_0$ . Both the depth and the ego-motion estimation networks are trained simultaneously using a common loss  $\mathcal{L}_{depth}$  between  $\hat{I}_0$  and  $I_0$ , which is a weighted sum of two losses. First is the appearance-based photometric loss between the original and the synthesized target image,  $\mathcal{L}_{photo}$ . The photometric loss is itself composed of a structural similarity ( $\mathcal{L}_{SSIM}$ ) loss [27] and an  $\ell_1$  distance loss between the images. Second is a smoothness loss on the depth output,  $\mathcal{L}_{smooth}$ , for regularization in low-texture portions of the scene. Therefore,

$$\mathcal{L}_{depth} = \lambda_1 \cdot \mathcal{L}_{SSIM} + \lambda_2 \cdot \ell_1 + \lambda_3 \cdot \mathcal{L}_{smooth}, \quad (1)$$

where  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are the hyperparameter weights for the losses.

### B. Integrating Masked Image Modeling

Our method MIMDepth, integrates MIM with self-supervised monocular depth estimation for masked self-supervised depth estimation. While MIM in pre-training applies masking for the network that reconstructs the image, we could mask the input to the depth or the ego-motion estimation networks. The input to the depth estimation network is a single image, which is converted into non-overlapping patches with token and position embeddings. For integrating MIM with the depth network, we use a mask generator to sample the patch indexes to be masked. The method for mask generation is described in Algorithm 1. A block with a random aspect ratio is repeatedly selected to be masked in the image until the required mask ratio is achieved. The indexes to be masked are replaced with learnable masked token embeddings, which are initialized as Gaussian noise.

---

### Algorithm 1: Blockwise Masking for Depth Estimation Network

---

```

input : unmasked tokens  $x_u$ 
          mask size  $m_s$ 
          mask ratio  $m_r$ 
          image shape  $(h \times w)$ 
          block aspect ratio  $a$ 
output: masked tokens  $x_m$ 
1  $n \leftarrow (h \times w) / m_s^2$  // num maskable patches
2  $M \leftarrow []$  // initialize mask
3 do
4    $s \leftarrow \text{random}(m_s, m_r \times n - |M|)$  // select block size
5    $r \leftarrow \text{random}(a, a^{-1}), a < 1$  // select block aspect ratio
6    $x \leftarrow \text{random}(0, w - \sqrt{s/r}); y \leftarrow \text{random}(0, h - \sqrt{s \cdot r})$ 
7    $M_{block} \leftarrow \{(i, j) : i \in [y, y + \sqrt{s \cdot r}], j \in [x, x + \sqrt{s/r}]\}$ 
8    $M \leftarrow M \cup M_{block}$  // add block to mask
9 while  $|M| < m_r \times n$ 
10  $M \leftarrow M.\text{repeat}([h/m_s], [w/m_s])$  // repeat mask to match the transformer patch size
11  $x_m \leftarrow x_u \times (1 - M) + \mathcal{N}(0, \delta) \times M$  // generate masked tokens
12 return  $x_m$ 

```

---

We use a blockwise masking strategy similar to BEIT [9], but we do not use any dVAE to generate targets. We use a mask ratio of 25% which is higher than that used in MLM but less than that used in MIM pre-training. This is due to the lower information redundancy in the depth of objects than in their texture. As an example, a gray color car would have similar texture across the pixels, but its depth would vary depending upon how it is oriented with respect to the camera. The network is trained to directly predict the depth for the complete image from the partially masked input. However, we do not apply masking to the ego-motion network, as it does not improve depth estimation (see Section IV).

SimMIM [10] shows that it is better to learn to predict over the masked patches rather than the complete image in order to pre-train representations that assist in downstream tasks. This is because training the network on the complete image would waste learning capability on the reconstruction quality of unmasked portions. However, with MIMDepth, the objective is to directly learn pixel-wise dense prediction including the unmasked portions. Hence, we apply the loss  $\mathcal{L}_{depth}$  to the complete target image.

## IV. RESULTS

Here, we perform a comparative analysis between the proposed method MIMDepth and existing self-supervised monocular depth estimation methods in terms of their performance on i.i.d. test set, as well as robustness to natural and digital corruptions, occluded views of the scene, and against adversarial attacks. Thereafter, through an ablation study, we resolve the questions posed in Section I regarding the mask sampling strategy, size and ratio, as well as the training target for integrating MIM into depth estimation. Finally, we show how integrating MIM into self-supervised monocular depth estimation also results in improved ego-motion estimation on the i.i.d. test set.

### A. Settings

We train the networks on a TeslaV100 GPU for 20 epochs with AdamW optimizer [28] at a resolution of  $640 \times 192$  with batchsize 12. The depth and ego-motion encoders are initialized with ImageNet [29] pre-trained weights. The learning rate is set at  $1e^{-5}$  and decays by a factor of 10 after 15 epochs. Unless otherwise stated, the networks are trained on the Eigen-Zhou split [13] of the KITTI dataset [30] with 39,810 training, 4424 validation, and 697 test images, respectively. We set the mask size  $m_s$ , mask ratio  $m_r$ , aspect ratio  $a$  for Algorithm 1 to 16, 25%, and 0.3, respectively.

### B. Robustness

While multiple works focus on the performance of depth estimation models, deployment in the real world necessitates models that not only perform well in the ideal conditions, but are also robust to various corruptions as well as adversarial attacks.

To evaluate the robustness, we study the impact of natural corruptions and adversarial attacks on depth performance against other methods, namely Monodepth2 [15], Packnet [5], and MT-SfMLearner [7] which use 2D convolutions, 3D convolutions, and transformer-based architectures, respectively. We report the robustness of MIMDepth through the average RMSE for models trained with three different seeds.

1) *Robustness to Natural Corruptions*: The camera used in an autonomous driving vehicle or a driver assistance system can be subject to different corruptions. Particularly, these corruptions can be due to noise (Gaussian, shot, impulse), or blur (defocus, glass, motion, zoom), different weather conditions (snow, frost, fog, brightness) and even digital aberrations (contrast, elastic, pixelation, JPEG). Figure 2 compares the different methods on the four types of corruptions, as well as on clean uncorrupted data. Note that the proposed MIMDepth outperforms other methods on clean, noisy, blur, and weather corruptions. However, PackNet-SfM is better on digital corruption due to its better performance on pixelated corruption, attributed to its architecture’s property of preserving depth details through 3D convolutions. Nevertheless, MIMDepth outperforms it on contrast, elastic, and JPEG digital corruptions. MIMDepth’s robustness can be attributed to the higher level of long-range semantic context that is learned by introducing MIM. The impact of a few corruptions on depth estimates is visualized in Figure 3.

2) *Occlusion Corruptions*: Monocular depth estimation training requires that the target image is reconstructed using adjacent images. However, due to static as well as dynamic objects, certain regions visible in the target image are occluded in the source image (and vice versa) and cannot be reconstructed. Occlusions are also an important consideration in object-specific depth estimation [8]. However, the role of occlusions has not been studied directly during the test time. Here, we evaluate the impact of occluded context in estimating the depth of objects in the scene.

Consequently, we simulate occlusions through modified test sets in which portions of the scene are masked out. Pixels in the occluded region are replaced by the average pixel value

of the complete image. In particular, we generate occlusions in the images using blockwise and random masking at the 25% mask ratio.

The aim is to evaluate the performance of the depth networks on the unoccluded regions of the scene, as well as to study the direct impact of masked depth estimation in predicting the depth on the complete image. This can also be thought to correspond to scenarios in which only a portion of the target object is visible for object-specific depth estimation. Figure 2 compares the different methods on the occluded test sets. Occlusions are found to impact depth estimates not only in the masked regions but also in the unmasked regions. We observe that MIMDepth is more robust to occlusions than other methods. By using the global context in a better way, MIMDepth is able to have the best predictions on the unmasked regions and is also able to infer plausible depth for the masked regions. The impact of occlusions on depth estimates for an example image is visualized in Figure 3.

### C. Adversarial Attacks

Adversarial attacks are used to fool networks into making an incorrect prediction through imperceptible changes in input images. These also help to measure the generalizability of the network as the data distribution is altered. We consider two types of adversarial attacks. For the untargeted attack, we generate adversarial examples following [31] at attack strengths of  $\epsilon = \{1.0, 2.0, 4.0, 8.0, 16.0\}$ . The gradients are computed with respect to the training loss  $\ell_{depth}$  (see Section III) and the perturbations are accumulated over  $\min(\epsilon + 4, \lceil 1.25 \cdot \epsilon \rceil)$  iterations. For the targeted attack, we generate adversarial examples that are intended to fool networks into predicting flipped depth estimates [32]. This attack uses the gradients on the RMSE values corresponding to when the target depths are horizontal or vertical flips of the original prediction. Similar to the PGD attack above, we generate adversarial examples with attack strengths of  $\epsilon = \{1.0, 2.0, 4.0\}$ .

The results of the untargeted adversarial attack are shown in Figure 2. The results of horizontal and vertical flipped adversarial attacks on depth estimation are shown in Figure 4. We note that MIMDepth is the most robust to untargeted and targeted attacks. This can be attributed to the better global context that is learned by adapting MIM to depth estimation.

### D. Performance

It is noted that the accuracy of the model is often at odds with the robustness [33]. Hence, we also compare MIMDepth with other state-of-the-art methods for its performance on the i.i.d. test set in Figure 5. For a fair comparison, we consider methods that predict depth only from single images, do not use additional supervised semantic information and train only on KITTI. MIMDepth is able to achieve performance comparable to that of other state-of-the-art methods on both error (RMSE) and accuracy ( $\delta < 1.25$ ) metrics [34]. Overall, this shows that integrating MIM into monocular depth estimation improves the robustness of depth estimation

	Corruptions					Oclusions (Complete)		Oclusions (Unmasked)		Untargeted Attack				
	clean	noise	blur	weather	digital	blockwise	random	blockwise	random	1	2	4	8	16
Monodepth2	4.855	9.129	8.176	7.375	6.414	9.290	6.190	5.101	5.880	5.736	6.269	6.742	7.192	7.882
Packnet-SfM	4.788	8.682	6.037	6.041	5.373	8.238	6.475	5.062	6.243	4.904	5.086	5.436	5.947	6.936
MT-SfMLearner	4.783	6.075	6.260	5.884	5.493	6.722	5.172	4.697	5.037	5.475	5.983	6.486	6.971	7.464
MIMDepth	4.760	5.721	5.963	5.786	5.415	5.570	4.983	4.550	4.911	4.875	5.002	5.155	5.457	5.821

Fig. 2: Robustness to natural corruptions, oclusions, and untargeted adversarial attack. Values denote RMSE (m).

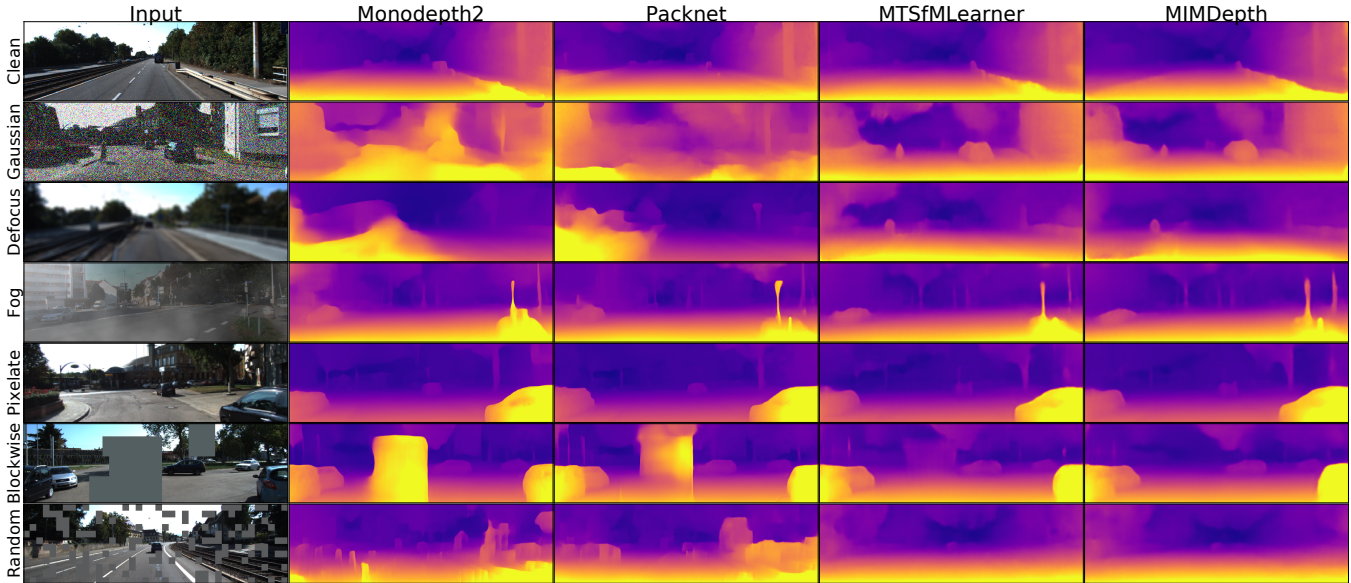


Fig. 3: Examples of depth estimation with different corruptions of the scene.

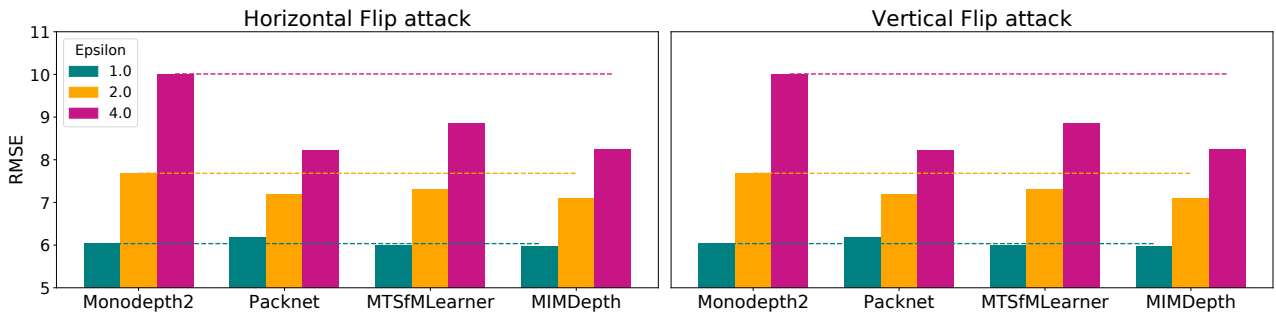


Fig. 4: Targeted horizontal and vertical flip adversarial attacks.

while maintaining comparable performance on the i.i.d. test set.

*Model Size and Computational Efficiency:* The expansion in terms of depth model size is limited to the mask token with dimension  $d=768$  parameters, which is insignificant compared to the baseline depth model size (183 M). Thus, the memory overhead is marginal. Moreover, since image masking is not used during prediction, the inference speed remains the same.

#### E. Ablation Studies

Here, we resolve the questions posed in Section I about the integration of MIM into self-supervised monocular depth

estimation through ablation over different configurations. We use the RMSE on clean input and robustness to occlusion corruptions (for complete image) as the metrics for comparison.

MIMDepth applies blockwise masking only on the depth network. We have shown that masking the depth network improves the robustness of the model. Here, we study the utility of applying MIM also on the ego-motion network. We also compare the use of blockwise and random mask sampling strategies for training the networks. From Table I, we observe that blockwise masking helps the depth estimation network to reduce the mean RMSE over clean and corrupted

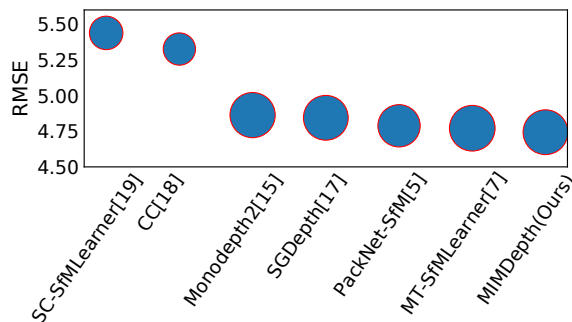


Fig. 5: Performance of state-of-the-art self-supervised monocular depth estimation methods. The size of the bubble correlates with the accuracy ( $\delta < 1.25$ ) for each method (bigger is better).

input. Random masking does not perform as well. This is because block-wise masking is better at inhibiting short-range connections to prevent the network from inpainting depth. Furthermore, block-wise masking is more reflective of how specific objects may be occluded rather than random occlusions across the scene. We also note that masking the ego-motion network along with the depth network does not improve depth estimation performance on clean data or robustness to occlusions. This could be attributed to the ego-motion network learning the translation and rotation from the shared context of the image pair, which is reduced by masking and is ineffective in improving the dense depth predictions.

MIMDepth uses a mask ratio and size of 25% and 16 respectively. We also compare with the use of a larger mask size or mask ratio in Table II. We observe that mask size 16 and mask ratio of 25% have the lowest mean RMSE over clean and corrupted input. Note that this value is lower than the masking ratio of 40% used by BEIT [9] or the mask size of 32 used by SimMIM [10] for pre-training. This can be attributed to the lower information redundancy for depth than that for image reconstruction (e.g. a vehicle with low texture variation and color consistency could have varying relative depth across its body, depending upon its orientation).

MIMDepth applies the training loss to the complete target image. We also compare with applying the loss only on the masked region as used by SimMIM [10]. Table II shows that MIMDepth works better when the loss  $\ell_{depth}$  is applied to the entire image. This is due to the training model’s focus on learning to predict depth for the complete image, as opposed

Depth	Ego-Motion	Clean	Blockwise	Random
B	-	4.760	5.570	4.983
R	-	4.777	6.592	5.071
B	B	4.788	5.614	5.009
B	R	4.789	5.579	5.010

TABLE I: Comparing impact of masking strategies for depth and ego-motion networks. B refers to blockwise masking and R refers to random masking. Values denote RMSE (m).

to the pre-training model’s focus on learning representations that only assist downstream tasks at the cost of poor image reconstruction.

### F. Impact on Ego-motion estimation Performance

We additionally evaluate the impact of integrating MIM with depth estimation on the ego-motion estimation predictions. We train MIMDepth on the KITTI odometry split and compare with the baseline MT-SfMLearner in Table III. Note that ego-motion predictions on the i.i.d. test set are improved by integrating MIM into the depth network. This can be understood by the use of a common loss for depth and ego-motion networks that is reduced by learning a better context for depth estimation as well as by improving the ego-motion prediction.

## V. CONCLUSION

We propose a method for adapting masked image modeling (MIM), used until now for self-supervised pre-training of models, into direct training of monocular depth estimation. Our method, MIMDepth, shows improved robustness against noise, blur, weather conditions, digital artifacts, and untargeted and targeted adversarial attacks on depth estimation. While model robustness is often at odds with accuracy, our method is able to maintain comparable depth estimation performance, while also showing improvements in ego-motion estimation. Through ablation studies, we reason about the differences in applying MIM to pre-training and adapting it for depth estimation. We find that blockwise masking (as opposed to random masking) of only the depth estimation network with a relatively lower mask size (than that used for pre-training), and a loss on the complete image (instead of only the masked regions, as used for pre-training) results in more robust depth estimates. We contend that this work could inspire the community to focus on the robustness required in the deployment of depth estimation networks. In the future, such methods could be extended to transformer-based depth estimation with multi-frame input (e.g. [20] or overlapping image patches (e.g. [21]). Additionally, MIM could also be adapted for direct training of more downstream tasks.

Size	Ratio	Loss	Clean	Blockwise	Random
16	25%	Complete	4.760	5.570	4.983
16	40%	Complete	4.872	5.594	5.079
32	25%	Complete	4.793	5.612	5.016
16	25%	Masked	4.794	5.652	5.031

TABLE II: Comparing the impact of mask ratio, mask size, and training target. Values denote RMSE (m).

Method	09		10	
	$t_{err}(\%)$	$r_{err}(\text{°}/100\text{m})$	$t_{err}(\%)$	$r_{err}(\text{°}/100\text{m})$
MT-SfMLearner [7]	9.574	2.570	13.851	5.568
MIMDepth (Ours)	6.236	2.065	8.178	3.865

TABLE III: Comparison of the performance of ego-motion estimation against the baseline in the KITTI odometry split.

## REFERENCES

- [1] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.
- [2] J. Lee, M. Han, D. Ko, and I. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation. arxiv 2019," *arXiv preprint arXiv:1907.10326*, 1907.
- [3] W. Yin, Y. Liu, C. Shen, and Y. Yan, "Enforcing geometric constraints of virtual normal for depth prediction," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 5684–5693.
- [4] H. Chawla, A. Varma, E. Arani, and B. Zonooz, "Multimodal scale consistency and awareness for monocular self-supervised depth estimation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [5] 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.
- [6] H. Chawla, A. Varma, E. Arani, and B. Zonooz, "Adversarial attacks on monocular pose estimation," in *2022 IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*. IEEE (in press), 2022.
- [7] A. Varma, H. Chawla, B. Zonooz, and E. Arani, "Transformers in self-supervised monocular depth estimation with unknown camera intrinsics," in *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 4: VISAPP, INSTICC*. SciTePress, 2022, pp. 758–769.
- [8] S. Lee, C. Lee, H. Kim, and H. J. Kim, "Realtime object-aware monocular depth estimation in onboard systems," *International Journal of Control, Automation and Systems*, vol. 19, no. 9, pp. 3179–3189, 2021.
- [9] H. Bao, L. Dong, S. Piao, and F. Wei, "Beit: Bert pre-training of image transformers," in *International Conference on Learning Representations*, 2021.
- [10] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu, "Simmim: A simple framework for masked image modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9653–9663.
- [11] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16000–16009.
- [12] C. Wei, H. Fan, S. Xie, C.-Y. Wu, A. Yuille, and C. Feichtenhofer, "Masked feature prediction for self-supervised visual pre-training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14668–14678.
- [13] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," 2017.
- [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] 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.
- [16] A. Gordon, H. Li, R. Jonschkowski, and A. Angelova, "Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras," 2019.
- [17] M. Klingner, J.-A. Termöhlen, J. Mikolajczyk, and T. Fingscheidt, "Self-Supervised Monocular Depth Estimation: Solving the Dynamic Object Problem by Semantic Guidance," in *ECCV*, 2020.
- [18] A. Ranjan, V. Jampani, L. Balles, K. Kim, D. Sun, J. Wulff, and M. J. Black, "Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2019, pp. 12240–12249.
- [19] J. Bian, Z. Li, N. Wang, H. Zhan, C. Shen, M.-M. Cheng, and I. Reid, "Unsupervised scale-consistent depth and ego-motion learning from monocular video," in *Advances in Neural Information Processing Systems*, 2019, pp. 35–45.
- [20] V. Guizilini, R. Ambrus, D. Chen, S. Zakharov, and A. Gaidon, "Multi-frame self-supervised depth with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 160–170.
- [21] J. Yang, L. An, A. Dixit, J. Koo, and S. I. Park, "Depth estimation with simplified transformer," in *CVPR Workshops*, 2022.
- [22] J. D. M.-W. C. Kenton and L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [23] M. Chen, A. Radford, R. Child, J. Wu, H. Jun, D. Luan, and I. Sutskever, "Generative pretraining from pixels," in *International conference on machine learning*. PMLR, 2020, pp. 1691–1703.
- [24] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," *ICLR*, 2021.
- [25] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8821–8831.
- [26] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International Conference on Machine Learning*. PMLR, 2021, pp. 10347–10357.
- [27] 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.
- [28] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [29] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [30] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [31] A. Kurakin, I. Goodfellow, S. Bengio, *et al.*, "Adversarial examples in the physical world," 2016.
- [32] A. Wong, S. Cicek, and S. Soatto, "Targeted adversarial perturbations for monocular depth prediction," in *Advances in neural information processing systems*, 2020.
- [33] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, "Robustness may be at odds with accuracy," in *International Conference on Machine Learning*, 2019.
- [34] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *Advances in neural information processing systems*, 2014, pp. 2366–2374.