

# SwinMTL: A Shared Architecture for Simultaneous Depth Estimation and Semantic Segmentation from Monocular Camera Images

Pardis Taghavi<sup>1</sup>, Reza Langari<sup>1</sup> and Gaurav Pandey<sup>2</sup>

**Abstract**—This research paper presents an innovative multi-task learning framework that allows concurrent depth estimation and semantic segmentation using a single camera. The proposed approach is based on a shared encoder-decoder architecture, which integrates various techniques to improve the accuracy of the depth estimation and semantic segmentation task without compromising computational efficiency. Additionally, the paper incorporates an adversarial training component, employing a Wasserstein GAN framework with a critic network, to refine model’s predictions. The framework is thoroughly evaluated on two datasets - the outdoor Cityscapes dataset and the indoor NYU Depth V2 dataset - and it outperforms existing state-of-the-art methods in both segmentation and depth estimation tasks. We also conducted ablation studies to analyze the contributions of different components, including pre-training strategies, the inclusion of critics, the use of logarithmic depth scaling, and advanced image augmentations, to provide a better understanding of the proposed framework. The accompanying source code is accessible at <https://github.com/PardisTaghavi/SwinMTL>.

## I. INTRODUCTION

Recent advancements in computer vision have enabled various applications ranging from self-driving cars and robot navigation to immersive augmented reality experiences. The key to these advancements lies in scene understanding, where algorithms for semantic segmentation and depth estimation play pivotal roles. Leveraging these algorithms allows for 3D scene reconstruction from camera images alone, reducing the need for expensive sensors such as Lidars, and increasing scalability in real-world scenarios. An additional benefit of obtaining detailed scene understanding from low-cost cameras is that it reduces the need for costly high-definition maps, which may not always be reliable in dynamic environments.

While individual tasks like depth estimation and semantic segmentation have attracted considerable attention and attained impressive accuracy through the evolution of neural network architectures, the concurrent execution of multiple networks on embedded systems poses challenges due to the significant memory footprint and real-time performance constraints. In response, multi-task learning (MTL) has emerged as a solution, seeking to enhance generalization by leveraging domain-specific information within the training signals of interconnected tasks [1].

<sup>1</sup>Pardis Taghavi and Reza Langari are with the Department of Mechanical Engineering, Texas A&M University, College Station, TX 77840, USA {ptgh, rlangari}@tamu.edu

<sup>2</sup>Gaurav Pandey is with the Department of Engineering Technology & Industrial Distribution, Texas A&M University, College Station, TX 77840, USA gpandey@tamu.edu




Fig. 1: Unified Vision: This figure provides a general overview of the seamless integration designed for joint semantic segmentation and depth estimation, ultimately contributing to the creation of a 3D scene.

This paper presents a novel multi-task learning approach utilizing a shared encoder and decoder for concurrent dense semantic segmentation and depth prediction tasks using a single camera. We also integrate adversarial training [2] into the training framework that involves training our model alongside a critic network. This results into an iterative refinement process where the critic guides the model towards better predictions for both tasks. The seamless integration of the encoder and decoder in our model significantly reduces the memory footprint of the inference engine, and the adversarial training improves accuracy. Overall our approach yields improved results compared to state-of-the-art methods. Our method is validated on the challenging Cityscapes [3] outdoor dataset and the NYU Depth V2 [4] indoor dataset demonstrating superior performance in comparison to existing approaches for each task.

## II. RELATED WORK

### A. Semantic Segmentation

The pioneering work in deep learning for semantic segmentation was the introduction of the Fully Convolutional Networks (FCN) architecture by Long et al. [5]. Despite its effectiveness, FCNs struggle with capturing fine image details due to multiple pooling operations and neglected pixel relationships [6]. Subsequent works, such as U-net [7], addressed these challenges by introducing an encoder-decoder architecture with skip connections to enhance feature propagation. DeepLab [8] incorporated Atrous Spatial Pyramid Pooling (ASPP) to capture both local and global contextual information for precise segmentation. The recent introduction of attention mechanisms and the development of

Vision Transformers (ViTs) [9] has surpassed state-of-the-art convolution-based methods in semantic segmentation [10].

### B. Depth Estimation

Eigen et al [11] pioneered the use of a learning-based approach, introducing a multi-scale convolutional neural network (CNN) coupled with a scale-invariant loss function for monocular depth estimation. Subsequently, several CNN-based methods have tackled depth estimation, treating it as a continuous regression problem based on a single RGB image [12], [13], [14]. Some methodologies, exemplified by [15] and [16] focused on discretizing continuous depth into intervals, treating depth prediction as a per-pixel classification problem. Another notable approach is found in the work of Bhat et al. [17], which combines regression and classification models, formulating depth estimation as a per-pixel classification-regression model. More recently, transformer architectures have made their way into depth estimation networks, aiming to capture global features and enhance prediction accuracy. One notable example is the work of Xie et al. [18], where transformers were incorporated into depth estimation networks, showcasing their ability to reveal intricate global patterns and improved overall performance.

### C. Multi Task Learning

Multi-task learning (MTL) is a machine learning paradigm where a single model is trained on multiple related tasks [19]. MTL has gained attention for its capacity to enhance performance on all tasks compared to independent single-task training. This is achieved by leveraging shared information and representations across the tasks, allowing the model to learn more efficiently. Methods such as [20], [21], focus on designing the encoder architecture to learn features that are relevant to multiple tasks. This approach can be effective for tasks with significant shared visual features. However, it may not be optimal for capturing both shared characteristics and task-specific variation features [1]. Other works, like [22], propose architectures that predict intermediate auxiliary tasks before using them as additional input for the final task predictions. This approach encourages the model to learn intermediate representations useful for multiple tasks. Beyond architecture design, some studies focused on the optimization process or proposed novel loss functions. For instance, [23] weighs multiple loss functions based on task uncertainty. [24] introduced the gradient normalization in deep multitask models to address issues with imbalanced gradients in MTL settings. Additionally, [25] proposed dynamic task prioritization during training to focus on more challenging tasks. Dynamic prioritization is achieved by adaptively adjusting the weight of each task's loss function based on its perceived difficulty, which is inversely proportional to the model's performance on that task.

works such as [22], [26] rely on a joint CNN model to fuse information from predicted auxiliary tasks, enhancing accuracy but at the cost of increased model complexity. Hoyer et al. [27] employ self-supervised depth estimation and

transfer learned knowledge to semi-supervised semantic segmentation. This method addresses the challenge of requiring large labeled training data, but it depends on the quality of self-supervised depth estimations and struggles with domain discrepancies. The approach proposed by Liu et al. [28] tackles the limitations of prior works by employing a single shared network with task-specific attention modules. This combination enables the network to learn shared features across different tasks while enabling each task to focus on its specific feature set through the attention modules. While attention modules can automatically select features, they may not always capture the exact features most relevant to the task. To address this, Ye et al [29] introduced explicit task prompts, which are learnable tokens capturing task-specific information. Nevertheless, the inclusion of learnable task prompts introduces significant computational overhead to the network.

The work published by [22], [27], [28] is closest to ours as they apply MTL to joint estimation of depth and semantic segmentation. Similar to Liu et al [28], we employ an end-to-end training approach for joint depth estimation and semantic segmentation. However, unlike their use of task-specific attention modules, our method directly shares features across tasks, further simplifying the model. Additionally, the incorporation of an adversarial training framework maintains high accuracy. This design choice is crafted to overcome limitations observed in related works, where the introduction of extra network components resulted in increased complexity. Our primary objective is to showcase enhanced accuracy through the seamless optimization of weights for both depth regression and pixel classification. By sidestepping the challenges associated with auxiliary tasks such as task-specific attention modules and complexities, our approach presents a more practical and effective solution.

## III. METHOD

### A. Model Architecture

In our proposed model, we adopt the vision transformer introduced by Liu et al. [30] as our backbone. The inherent hierarchical structure and attention mechanisms of transformers make them well-suited for capturing complex spatial relationships, which are crucial for tasks such as depth estimation and semantic segmentation. The Swin Transformer[30], in particular, demonstrates superior performance owing to its additional shifted windows approach in the backbone. In our method, we design a multi-task learning framework to simultaneously address both depth estimation and semantic segmentation task. Notably, we deviate from the prevalent trend of employing intricate architectures for multi-task scenarios, opting for a consistent encoder and decoder for both tasks. This departure allows us to maintain shared components, simplifying the overall architecture and rendering it suitable for embedded systems. Additionally, we incorporate a generative adversarial training approach that enhances the overall accuracy. While previous studies leveraging GANs typically utilize binary classification discriminators in similar settings, we opt for the Wasserstein GAN (WGAN) [2] to measure

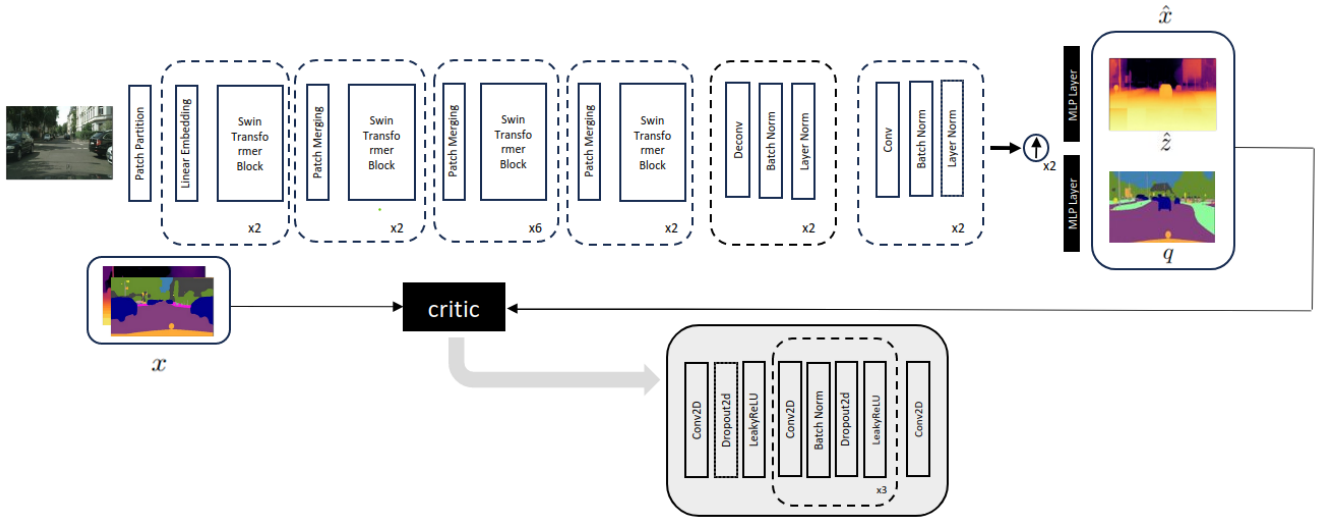


Fig. 2: Network Architecture: Detailed network architecture tailored for the precise prediction of both depth map ( $\hat{z}$ ) and segmentation map ( $q$ ). Notably, the network employs a critic discriminator during training, enhancing its ability to discern and refine the joint output ( $\hat{x}$ ) representing depth and segmentation information.

the similarity distance between predictions and ground truth distributions. Our approach showcases superior performance on both tasks when compared to training individual networks for each task with the same computational cost.

To enable individualized predictions for each task, we introduce two task-specific Multi-Layer Perceptron (MLP) layers at the end of the decoder. One MLP layer is dedicated to depth prediction, while the other is tailored for semantic segmentation. This streamlined and shared architecture contributes to the model’s memory efficiency and ensures high-quality task-specific predictions. The shared components allow for seamless integration of information, promoting generalization across tasks, while the task-specific MLP layers enable the model to fine-tune its predictions for depth estimation and semantic segmentation independently. Next we describe the loss functions used to train the proposed network.

**Depth Estimation Loss:** We utilize the pixel-wise scale-invariant loss proposed by [11] for depth estimation:

$$L_{depth}(\hat{z}, z_{gt}) = \frac{1}{N} \sum_{i,j} (e_{i,j}^2) - \frac{1}{N^2} (\sum_{i,j} e_{i,j})^2 \quad (1)$$

where

$$e_{i,j} = \log \hat{z}^{i,j} - \log z_{gt}^{i,j} \quad (2)$$

Here  $z^{i,j}$  is the depth at pixel  $(i, j)$ . The loss is designed to be invariant to global scale changes in the predicted depth map. The Logarithmic term (2) reduces the impact of large errors, making it robust to outliers and more accurate as compared to vanilla  $L1$  or  $L2$  loss.

**Segmentation Loss:** We adopt the standard cross-entropy loss for segmentation, effectively guiding the model towards accurate pixel-level classification:

$$L_{seg}(p, q) = - \sum_i p_i \log(q_i) \quad (3)$$

In equation (3),  $p$  represents the ground truth label of each pixel and  $q$  is the predicted probability distribution obtained from the log-softmax function. For a given class  $i$ ,  $q_i$  is calculated using the log-softmax function:

$$q_i = \log \left( \frac{e^{x_i}}{\sum_j e^{x_j}} \right) \quad (4)$$

In equation (4),  $x_i$  represents the logit corresponding to class  $i$ , and  $\sum_j e^{x_j}$  represents the sum of logits for all classes. This formulation ensures a robust and normalized representation of predicted probabilities for accurate segmentation loss computation.

**Adversarial Loss:** We integrate the loss function proposed in [31] for the generative adversarial networks, incorporating a gradient penalty (GP) regularization term to enforce Lipschitz continuity on the critic ( $c$ ), promoting stability and smoothness in training. This technique prevents the critic’s gradients from becoming excessively large, which can lead to training instability and convergence issues.

$$L_{GAN}(\hat{x}, x) = \mathbb{E}(c(x)) - \mathbb{E}(c(\hat{x})) + \lambda * GP \quad (5)$$

$$GP = \mathbb{E}[(\|\nabla_{\tilde{x}} c(\tilde{x})\|_2 - 1)^2] \quad (6)$$

$$\tilde{x} = \epsilon x + (1 - \epsilon) \hat{x} \quad (7)$$

In equation (5),  $L_{GAN}(\hat{x}, x)$  represents the adversarial network loss, where  $\hat{x}$  is the generated sample and  $x$  is the real sample (Fig. 2). The gradient penalty term GP (6) enforces smoothness by penalizing deviations from Lipschitz continuity. The parameter  $\lambda$  controls the strength of the regularization. Equation (7) introduces  $\tilde{x}$ , a linear interpolation between real and generated samples with  $\epsilon$  as a weighting

coefficient, facilitating gradient computation for the penalty term.

**Total Loss:** Total loss is a balanced combination of depth estimation, segmentation, and adversarial network losses, with  $\alpha$  and  $\beta$  controlling the influence of each loss function:

$$L_{\text{tot}} = \alpha L_{\text{depth}} + (1 - \alpha)L_{\text{seg}} + \beta L_{\text{GAN}} \quad (8)$$

Adversarial learning is a pivotal component of our approach. Inspired by Generative Adversarial Networks (GANs) [32], our approach establishes a dynamic interplay between two competing networks: the generator ( $g$ ), represented by SwinMTL, and the discriminator (referred to as a critic in the context of [2]). This interplay engages these two networks in a min-max competition, as shown in (9), where predictions are refined towards ground truth labels, enhancing overall accuracy.

$$\min_g \max_c \mathbb{E}_{x \sim p_{\text{data}}} [c(x)] - \mathbb{E}_{x \sim p_g} [c(g(x))] \quad (9)$$

In implementing adversarial training, we introduce a critic network during the training process to refine the output of the generator. Unlike previous approaches [33] that employ discriminators through binary classification, we leverage the Wasserstein GAN framework with Earth Mover’s Distance (EMD) [2], [31]. EMD quantifies dissimilarity between two distributions, namely predictions( $\hat{x}$ ) vs ground truth labels( $x$ ), addressing issues such as mode collapse and preventing the generator from outperforming the discriminator in the early epochs.

The critic in this approach receives as input the concatenated depth and segmentation maps of both ground truths and predictions. It maximizes the value function defined in (9), aiming to differentiate between the ground truth and the predicted data by outputting a probability ranging from zero to one. Conversely, the generator endeavors to minimize the value function to reduce the EDM distance between these datasets. This adjustment allows the critic to thoroughly assess the alignment between generated and authentic data across various dimensions. The detailed architecture of the critic is delineated in Fig. 2.

#### IV. EXPERIMENTS

In this section, we present the details of our experimental setup, showcasing results on the Cityscapes and NYU Depth V2 datasets. Additionally, we conduct ablation studies to gain insights into the individual contributions of key components in our approach.

##### A. Experimental Setup

Cityscapes is an outdoor dataset capturing urban street scenes, Cityscapes provides high-quality pixel-level annotations for 30 classes grouped into 8 categories (19 classes for training). The dataset comprises approximately 5000 finely annotated images, distributed into 2975, 500, and




Fig. 3: The first row shows a stereo pair from the Cityscapes dataset. The second row illustrates the dense disparity map (bottom right) and the dense depth map (bottom left), both generated using the advanced Cascaded Recurrent Stereo Matching Network (CREStereo) [34]

1525 images for training, validation, and testing, respectively. Spanning 50 different cities in Germany, the dataset’s resolution is 2048x1024. The Cityscapes dataset does not have dense depth annotation but includes stereo images. We leverage these stereo images to generate the dense disparity map using the Cascaded Recurrent Stereo Matching Network (CREStereo) [34] and, accordingly, the dense depth map (Fig. 3) required for training our network. NYU Depth V2 dataset consists of 1449 RGB-D images of indoor scenes, each paired with dense labeled depth images and semantic labels across 40 categories. Each image is of size 640x480, distributed into 795 training and 654 testing images.

Inspired by [35] we integrated a modified version of MixUp augmentation into our training pipeline to increase the robustness and generalization capabilities of our model. we randomly select a fixed number ( $N$ ) of square patches from each input image. These patches vary in size, ranging from  $1/8$  to  $1/2$  of the image height ( $h$ ). The patches are then shuffled, disrupting the original spatial relationships within the image. The shuffled patches are displaced within images of each batch to create new augmented images in each epoch. This process effectively blends visual information from multiple images, generating a greater diversity of training samples. MixUp Augmentation proves particularly beneficial in the NYU Depth V2 dataset, mitigating challenges arising from its limited number of images.

To underscore the importance of the near-range region for depth estimation in Cityscapes dataset, we transform the depth from linear space to logarithmic space [42] for the critic network. As illustrated in Fig. 4, a significant concentration of pixels falls within the 0-10 meter range, constituting over 80% of pixels in each image. In order to effectively capture this distribution and assign greater significance to this critical zone, we employ the logarithmic depth parameterization proposed in [42]. This transformation aligns model learning with the natural distribution of depth values, resulting in substantial improvements in depth estimation accuracy for the Cityscapes dataset. Here,  $d_{\text{metric}}$  denotes the metric depth, while  $d_{\text{min}}$  and  $d_{\text{max}}$  represent the minimum

TABLE I: Quantitative comparison of single-task depth prediction methods on Cityscapes test set (starred methods are self-supervised)

model	Resolution	AbsRel ↓	SqRel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Pilzer et al.[36]*	512x256	0.438	5.713	5.745	0.400	0.711	0.877	0.940
Struct2Depth[37]*	416x128	0.145	1.737	7.280	0.205	0.813	0.942	0.976
Li et al.[38]*	416x128	0.119	1.290	6.980	0.190	0.846	0.952	0.982
DynamicDepth[39]*	416x128	0.103	1.000	5.867	0.157	0.895	0.974	0.991
Lee et al.[40]*	832x256	0.116	1.213	6.695	0.186	0.852	0.951	0.982
InstaDM[41]*	832x256	0.111	1.158	6.437	0.182	0.868	0.961	0.983
<b>SwinMTL</b>	<b>512x1024</b>	<b>0.089</b>	<b>1.051</b>	<b>5.481</b>	<b>0.139</b>	<b>0.921</b>	<b>0.976</b>	<b>0.990</b>




Fig. 4: Histogram of depth values of an instance in the Cityscapes dataset, visually highlighting the dominance of values within the 0 – 10 meter range. The distribution’s peak and long tail motivate the logarithmic transformation.

and maximum depth values respectively.

$$d_{\log} = \frac{\log(d_{metric}/d_{min})}{\log(d_{max}/d_{min})} \quad (10)$$

SwinMTL offers flexibility in choosing different backbones for its model; however, our reported results are based on the SwinV2-B model, pretrained with Masked Image Modeling (MIM) [18]. We employ the AdamW optimizer [43] with a base learning rate of  $5e-4$  and weight decay of  $5e-2$ . The learning rate follows a polynomial strategy with a decay factor of 0.9. Additionally, we incorporate random adjustments in brightness, contrast, gamma correction, hue saturation value, and horizontal flip, aligning with established methodologies in prior research. For direct comparison with other approaches, we set the maximum depth range to 80m for the Cityscapes dataset and 10m for the NYU Depth V2 dataset. We implemented our method in PyTorch 2.0.1 (Python 3.8) with CUDA 11.8, utilizing two NVIDIA GeForce RTX 3090 GPUs.

### B. Results and Comparisons

We conduct a thorough performance comparison of SwinMTL against state-of-the-art models on the Cityscapes and NYU Depth V2 datasets for dense prediction tasks. The

TABLE II: Quantitative Performance Comparison of Our MTL Framework versus SOTA Methods on Cityscapes Test Set on The Full Resolution (1024x2048)

model	RMSE ↓	mIoU ↑
Hybridnet[44]	12.090	66.61
MGMNet[45]	8.300	55.70
MTL[1]	6.797	70.43
PAD-Net[22]	6.777	70.23
3-ways[27]	6.528	75.00
[46]	6.649	74.95
<b>SwinMTL</b>	<b>6.352</b>	<b>76.41</b>

TABLE III: Quantitative Performance Comparison of SwinMTL Framework versus SOTA Methods on NYU Test Set

model	RMSE ↓	mIoU ↑
Cross-Stitch[47]	0.6290	36.34
MTL[1]	0.6380	39.44
PAD-Net[22]	0.6270	36.61
3-ways[27]	0.622	39.47
[46]	0.6040	38.93
MTAN[28]	0.5906	40.01
MTI-Net[48]	0.5365	45.97
ATRC[49]	0.5363	46.33
InvPT[50]	0.5183	53.56
TaskPrompter[29]	<b>0.5152</b>	55.30
<b>SwinMTL</b>	0.5179	<b>58.14</b>

results for the Cityscapes dataset are compared with other depth estimation methods in Table I and with other MTL networks in Table II. Table III showcases the performance of NYU Depth V2 dataset compared to other pioneering MTL networks. We ensure a fair comparison by using the best-reported numbers from prior works. As demonstrated by the tables for both Cityscapes and the NYU Depth V2 dataset, SwinMTL consistently outperforms or shows comparable performance with enhanced computational efficiency. For additional insights, Fig. 5 presents qualitative results on the Cityscapes dataset.

### C. Ablation Studies

To efficiently assess the individual contributions of various components within our proposed framework, we conducted ablation studies on lower resolution images (256x128) from the Cityscape dataset for accelerated computation. These findings provide valuable insights into the effectiveness of each module while maintaining computational feasibility.

TABLE IV: Ablation Studies on the Cityscapes dataset

Ablation	Model	AbsRel↓	mIoU↑
Single-Task vs. Multi-Task Network Performance	only Depth	0.110	-
	only Seg SwinMTL	- <b>0.104</b>	42.82 <b>50.35</b>
Impact of Different Pretrained Networks	w/o pretrained	0.617	17.24
	w normal pretrained	0.432	25.73
	w MIM pretrained	<b>0.104</b>	<b>50.35</b>
Impact of Critic Integration	w/o Critic	0.234	50.15
	w one Critic	<b>0.104</b>	<b>50.35</b>
	w two Critics	0.105	50.30
Linear vs. Logarithmic	Linear Space	0.152	-
	Log Space	<b>0.104</b>	-

**Comparison of Single and Multi-task Networks:** We conducted a comparative analysis of single-task and multi-task networks using a Swin-Base backbone. As summarized in Table IV, the multi-task network achieved superior evaluation metrics for both depth estimation and semantic segmentation tasks, despite a marginal increase of only 3% in parameter count. This evidence showcases the potential of multi-task learning to enhance model performance without compromising efficiency.

**Performance with Different Pretrained Networks:** Understanding the influence of pretraining variations on downstream task performance is crucial for model optimization. In this ablation study (Table IV), we evaluate our model’s performance with diverse pretrained networks: a standard Swin-V2 model pretrained on ImageNet-22K and a Swin-V2 pretrained with masked image modeling (MIM) [18]. By comparing results, we can see the effectiveness of MIM pretraining on the accuracy of depth and segmentation predictions which is also aligned with results shown in [18].

**Critic Addition Performance:** The integration of a critic into the network architecture led to performance gains, as evidenced in Table IV. Notably, incorporating a single critic during training, without increasing testing parameters, produced the most significant improvements in depth estimation and semantic segmentation accuracy. This addition highlights the potential for critics to refine model learning without sacrificing computational overhead.

**Integrating Logarithmic Depth Scaled for Enhanced Accuracy:** Fig.4 reveals a major concentration of depth values within the 0-10 meter range, accounting for over 80% of the pixels in our dataset. To effectively model this distribution and prioritize precision in this critical zone, we adopt the logarithmic depth parameterization proposed in [42]. This transformation aligns model learning with the natural distribution of depth values, leading to substantial improvements in depth estimation accuracy, as evidenced in Table IV.

**Computation Footprint:** Lastly, table V showcases a comparison against state-of-the-art methods, revealing our model’s ability to achieve high accuracy with significantly fewer parameters and lower computational costs.

TABLE V: The parameter count comparison across different methods, emphasizing our model’s lean architecture.

Models	FLOPs	N. of Parameters
TaskPrompter[29]	416G	373M
ATRC w/ ViT-B[49]	260G	119M
PAD-Net w/ ViT-B[22]	166G	109M
SwinMTL w/ Swin-B	<b>65G</b>	<b>87.38M</b>

## V. CONCLUSIONS

In conclusion, this paper proposed a streamlined yet effective multi-task learning framework for joint depth estimation and semantic segmentation. We demonstrate the effectiveness of our method on the Cityscapes and NYU Depth V2 datasets, consistently outperforming leading-edge methods in both depth estimation and semantic segmentation. The contributions of our work encompass the introduction of a multi-task learning design that considers computational efficiency along with accuracy, the integration of a critic network to refine dense predictions. Furthermore, our work studies the impact of various components, including pre-training strategies, the incorporation of critics, the utilization of logarithmic depth scaling, and the inclusion of MixUp augmentation through extensive ablation studies. Future work could explore further optimizations on joint weight training procedures to deepen our understanding of the intricate relationship between depth and segmentation. Additionally, extending the application of our framework to diverse datasets and varied environmental conditions would contribute to its robust generalization capabilities.

## REFERENCES

- [1] S. Vandenhende, S. Georgoulis, W. Van Gansbeke, M. Proesmans, D. Dai, and L. Van Gool, “Multi-task learning for dense prediction tasks: A survey,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 7, pp. 3614–3633, 2021.
- [2] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *International conference on machine learning*. PMLR, 2017, pp. 214–223.
- [3] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [4] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, “Indoor segmentation and support inference from rgb-d images,” in *Computer Vision – ECCV 2012*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 746–760.
- [5] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [6] B. Chen, X. Qi, Y. Zhou, G. Yang, Y. Zheng, and B. Xiao, “Image splicing localization using residual image and residual-based fully convolutional network,” *Journal of Visual Communication and Image Representation*, vol. 73, p. 102967, 2020.
- [7] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [8] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017.




Fig. 5: Exploring SwinMTL's Qualitative Results on the Cityscapes Dataset

- [9] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [10] H. Thisanke, C. Deshan, K. Chamith, S. Seneviratne, R. Vidanaarachchi, and D. Herath, "Semantic segmentation using vision transformers: A survey," *Engineering Applications of Artificial Intelligence*, vol. 126, p. 106669, 2023.
- [11] 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.
- [12] B. Li, C. Shen, Y. Dai, A. van den Hengel, and M. He, "Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1119–1127.
- [13] 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.
- [14] D. Kim, S. Lee, J. Lee, and J. Kim, "Leveraging contextual information for monocular depth estimation," *IEEE Access*, vol. 8, pp. 147808–147817, 2020.
- [15] 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.
- [16] R. Diaz and A. Marathe, "Soft labels for ordinal regression," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4738–4747.
- [17] S. F. Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4009–4018.
- [18] Z. Xie, Z. Geng, J. Hu, Z. Zhang, H. Hu, and Y. Cao, "Revealing the dark secrets of masked image modeling," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 14475–14485.
- [19] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, pp. 41–75, 1997.
- [20] O. Sener and V. Koltun, "Multi-task learning as multi-objective optimization," *Advances in neural information processing systems*, vol. 31, 2018.
- [21] M. Teichmann, M. Weber, M. Zoellner, R. Cipolla, and R. Urtasun, "Multinet: Real-time joint semantic reasoning for autonomous driving," in *2018 IEEE intelligent vehicles symposium (IV)*. IEEE, 2018, pp. 1013–1020.
- [22] D. Xu, W. Ouyang, X. Wang, and N. Sebe, "Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 675–684.
- [23] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7482–7491.
- [24] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, "Grad-norm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *International conference on machine learning*. PMLR, 2018, pp. 794–803.
- [25] M. Guo, A. Haque, D.-A. Huang, S. Yeung, and L. Fei-Fei, "Dynamic task prioritization for multitask learning," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 270–287.
- [26] A. Mousavian, H. Pirsiavash, and J. Košecák, "Joint semantic segmentation and depth estimation with deep convolutional networks," in *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016, pp. 611–619.
- [27] L. Hoyer, D. Dai, Y. Chen, A. Koring, S. Saha, and L. Van Gool, "Three ways to improve semantic segmentation with self-supervised depth estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 11130–11140.
- [28] S. Liu, E. Johns, and A. J. Davison, "End-to-end multi-task learning with attention," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 1871–1880.
- [29] H. Ye and D. Xu, "Taskprompter: Spatial-channel multi-task prompting for dense scene understanding," in *The Eleventh*

- International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=CwPopPJda>
- [30] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [31] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," *Advances in neural information processing systems*, vol. 30, 2017.
- [32] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [33] C. Zhang, Y. Tang, C. Zhao, Q. Sun, Z. Ye, and J. Kurths, "Multitask gans for semantic segmentation and depth completion with cycle consistency," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 12, pp. 5404–5415, 2021.
- [34] J. Li, P. Wang, P. Xiong, T. Cai, Z. Yan, L. Yang, J. Liu, H. Fan, and S. Liu, "Practical stereo matching via cascaded recurrent network with adaptive correlation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 16 263–16 272.
- [35] S. Li, Z. Wang, Z. Liu, D. Wu, and S. Z. Li, "Openmixup: Open mixup toolbox and benchmark for visual representation learning," *arXiv preprint arXiv:2209.04851*, 2022.
- [36] A. Pilzer, D. Xu, M. Puscas, E. Ricci, and N. Sebe, "Unsupervised adversarial depth estimation using cycled generative networks," in *2018 international conference on 3D vision (3DV)*. IEEE, 2018, pp. 587–595.
- [37] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova, "Unsupervised monocular depth and ego-motion learning with structure and semantics," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.
- [38] H. Li, A. Gordon, H. Zhao, V. Casser, and A. Angelova, "Unsupervised monocular depth learning in dynamic scenes," in *Conference on Robot Learning*. PMLR, 2021, pp. 1908–1917.
- [39] Z. Feng, L. Yang, L. Jing, H. Wang, Y. Tian, and B. Li, "Disentangling object motion and occlusion for unsupervised multi-frame monocular depth," in *European Conference on Computer Vision*. Springer, 2022, pp. 228–244.
- [40] S. Lee, F. Rameau, F. Pan, and I. S. Kweon, "Attentive and contrastive learning for joint depth and motion field estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4862–4871.
- [41] S. Lee, S. Im, S. Lin, and I. S. Kweon, "Learning monocular depth in dynamic scenes via instance-aware projection consistency," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 3, 2021, pp. 1863–1872.
- [42] S. Saxena, J. Hur, C. Herrmann, D. Sun, and D. J. Fleet, "Zero-shot metric depth with a field-of-view conditioned diffusion model," *arXiv preprint arXiv:2312.13252*, 2023.
- [43] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [44] D. Sanchez-Escobedo, X. Lin, J. R. Casas, and M. Pardas, "Hybridnet for depth estimation and semantic segmentation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1563–1567.
- [45] M. Schön, M. Buchholz, and K. Dietmayer, "Mgnet: Monocular geometric scene understanding for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 15 804–15 815.
- [46] I. Lopes, T.-H. Vu, and R. de Charette, "Cross-task attention mechanism for dense multi-task learning," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2329–2338.
- [47] I. Misra, A. Shrivastava, A. Gupta, and M. Hebert, "Cross-stitch networks for multi-task learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3994–4003.
- [48] S. Vandenhende, S. Georgoulis, and L. Van Gool, "Mti-net: Multi-scale task interaction networks for multi-task learning," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*. Springer, 2020, pp. 527–543.
- [49] D. Brüggemann, M. Kanakis, A. Obukhov, S. Georgoulis, and L. Van Gool, "Exploring relational context for multi-task dense prediction," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 15 869–15 878.
- [50] H. Ye and D. Xu, "Inverted pyramid multi-task transformer for dense scene understanding," in *European Conference on Computer Vision*. Springer, 2022, pp. 514–530.