

SwinDepth: Unsupervised Depth Estimation using Monocular Sequences via Swin Transformer and Densely Cascaded Network

Dongseok Shim and H. Jin Kim*

Abstract—Monocular depth estimation plays a critical role in various computer vision and robotics applications such as localization, mapping, and 3D object detection. Recently, learning-based algorithms achieve huge success in depth estimation by training models with a large amount of data in a supervised manner. However, it is challenging to acquire dense ground truth depth labels for supervised training, and the unsupervised depth estimation using monocular sequences emerges as a promising alternative. Unfortunately, most studies on unsupervised depth estimation explore loss functions or occlusion masks, and there is little change in model architecture in that ConvNet-based encoder-decoder structure becomes a de-facto standard for depth estimation. In this paper, we employ a convolution-free Swin Transformer as an image feature extractor so that the network can capture both local geometric features and global semantic features for depth estimation. Also, we propose a Densely Cascaded Multi-scale Network (DCMNet) that connects every feature map directly with another from different scales via a top-down cascade pathway. This densely cascaded connectivity reinforces the interconnection between decoding layers and produces high-quality multi-scale depth outputs. The experiments on two different datasets, KITTI and Make3D, demonstrate that our proposed method outperforms existing state-of-the-art unsupervised algorithms.

I. INTRODUCTION

Estimating a high-quality depth map from a single RGB image is theoretically an ill-posed problem without additional cues for triangulation. Learning-based methods aim to mitigate such difficulties by training a neural network in a supervised manner with a large amount of RGB images and densely annotated depth labels [1], [2], [3]. However, it is still challenging to collect accurate ground truth depth for supervised learning and, as an alternative, demands for unsupervised training of depth estimation network increase. There are two promising approaches to train depth estimation networks without depth labels which use stereo-paired images [4], [5] or monocular image sequences [6], [7]. In particular, the latter one is more attractive in that it does not require additional sensors such as a stereo camera.

Training only with monocular sequences needs to jointly estimate the depth map and the ego-motion of the camera to calculate the transformation matrix for photometric loss that it requires an additional pose estimation network. Also, it assumes a static scene and a moving camera for rigid motion. Due to these assumptions, the estimation errors such as undesired holes appear in the depth map if there exist some dynamic objects, e.g. pedestrians or cars, which can be a major hurdle for robotics applications [8], [9].

Authors are with Interdisciplinary Program in Artificial Intelligence (IPAI), Seoul National University. E-mail: {tlaehdtjr01, hjinkim}@snu.ac.kr. *Corresponding author




Fig. 1: A single image depth estimation with a hierarchical Transformer and densely cascaded network. From top to bottom, each row indicates input RGB images, estimated depth maps, and self-attention masks from multiple heads in Transformer feature extractor.

To overcome these errors, follow-up studies redesign the photometric loss function or propose masking strategies to prevent occlusions and hide moving objects in the scene, which may improve the quality of estimated depth maps. Unfortunately, there are few studies to improve the architecture of the depth estimation network itself where U-Net structure [10] with Convolutional neural Network (ConvNet)-based feature extractor [11], [12], [13] is dominant both in supervised and unsupervised training.

Recently, Vision Transformer (ViT) [14] shows that a pure Transformer produces better performance in image classification compared to competitive ConvNet. Its follow-up studies also show that ViT is still effective for backbone architecture in various visual recognition tasks such as object detection [15], [16] and semantic segmentation [17].

There exist some approaches trying to apply ViT on depth estimation, but most of them use feature maps generated by ConvNet [18], [19], rather than directly using RGB image as the input of Transformer. Some research such as DPT [20] naively stack multiple ViTs to generate ConvNet-like hierarchical feature maps to leverage Transformer-based encoder for downstream networks with skip connection, but it leads to heavy computational cost and many model parameters.

In this paper, we utilize a convolution-free hierarchical Transformer with a downsizing module and attention window from Swin Transformer [21] to reduce the innate computational burden in calculating feature attention. The resolution of the feature map progressively gets smaller by merging

nearby image patches in deeper layers, and the attention map is calculated in the fixed-size non-overlapping windows. These merits make the hierarchical Transformer encoder become much lighter while keeping the strength of capturing semi-global image features. It produces better performance on depth estimation compared to existing ConvNet-based architecture that has twice as many model parameters.

Also, we propose a simple but effective feature aggregation model, Densely Cascaded Multi-scale Network (DCMNet), to estimate multi-scale depth maps using hierarchical image features from Transformer encoder. Inspired by Dense Convolutional Network (DenseNet) [22] and Feature Pyramid Network (FPN) [23], we leverage the dense concatenation and the top-down cascade addition to improve the interconnection between features from decoding layers. We densely connect the hierarchical feature maps directly with one another that all the coarse low-dimension features are exploited to generate fine high-resolution depth outputs without any attenuation by passing through decoding layers.

In short, the contributions of our paper can be summarized as follows:

- We show that pure Transformer-based hierarchical feature extractor outperforms competitive ConvNet-based architecture in depth estimation.
- We propose DCMNet which improves the interconnection between decoding layers with a top-down cascade pathway and dense concatenation to produce high-quality multi-resolution depth outputs.
- We demonstrate the effectiveness of our approach on KITTI [24] and Make3D dataset [25] compared to existing state-of-the-art unsupervised depth estimation.

II. RELATED WORK

In this section, we summarize the previous research and the applications on Vision Transformer, and review learning-based monocular depth estimation algorithms which are categorized into supervised and unsupervised approaches.

A. Vision Transformer

Transformer architecture with self-attention [26] first suggested for machine translation has become one of the most widely used models in natural language processing (NLP).

To leverage the advantage of Transformer in NLP, Vision Transformer (ViT) [14] divides an RGB image into several patches and flatten them to treat the image as sequential data for Transformer input. Convolution-free Transformer encoder with patch embedding achieves a huge success and outperforms state-of-the-art ConvNets in image classification. Later, DeiT [27] introduces an additional distillation token to alleviate the dependency on the large-scale dataset for pre-training.

SETR [17] utilizes ViT as the backbone network and the feature aggregation module for semantic segmentation. DPT [20] stacks multiple ViTs to generate hierarchical feature maps for dense prediction. As stacking ViT modules without downsizing makes the network too large, Swin Transformer [21] proposes a patch merging and attention window to

reduce the computational and memory burden of Transformer block. It produces better performance as a backbone architecture for image classification, object detection and semantic segmentation compared to competitive ConvNet architecture, ViT and its variants.

In this paper, we also utilize the patch merging strategy and attention window to formulate a hierarchical Transformer feature extractor, hypothesizing that they will provide the same advantage shown in other visual recognition tasks for unsupervised depth estimation as well.

B. Learning-based Depth Estimation

It is challenging to generate high-quality depth outputs with a single RGB image and the development of deep learning relaxes difficulties in depth estimation with deep neural networks and a large-scale training dataset.

Eigen [1] proposes an end-to-end supervised training algorithm which consists of two-scale ConvNets each generating low-dimension coarse and high-dimension fine depth outputs. Laina [2] utilizes ResNet [13] with its novel up-projection network for fast and efficient training, and DORN [3] formulates depth estimation as a discrete problem by proposing an ordinal regression loss.

Due to the poor availability of depth labels for supervised learning, unsupervised depth estimation emerges as a promising alternative especially using consecutive monocular frames. SfMLearner [6] jointly estimates depth and the ego-motion of the camera with two different networks. The photometric loss is calculated by projecting one monocular frame to the other camera coordinate via Transformation matrix from the outputs of the pose estimation network and estimated depth values. The problem is that it assumes a rigid motion and the presence of moving objects leads to poor performance both in pose and depth estimation. To alleviate the performance degradation, Monodepth2 [7] leverages a minimum operation in calculating the photometric reprojection loss to prevent occlusion and proposes an auto-generated mask to hide moving objects in the scene.

We adopt the unsupervised depth estimation strategy suggested in Monodepth2 to estimate both depth values and relative pose with monocular image sequences.

III. METHOD

In this section, we elaborate the architectural details and key ideas of hierarchical Transformer and our proposed Densely Cascaded Multi-scale Network (DCMNet). Then, we describe the training algorithm for unsupervised depth estimation using monocular image sequences. An overview of our depth estimation network and the training algorithm are depicted in Figure 2.

A. Architecture

1) *Hierarchical Transformer*: Similar to Vision Transformer (ViT) [14], a hierarchical Transformer first splits the input RGB image into several non-overlapping patches (2×2). Image patches are then flattened and mapped to C -dimensional feature space by trainable linear projection in the patch embedding module.

where \mathbf{z}'_l and \mathbf{z}_l denote the features from (S)W-MSA module and the MLP module in the l^{th} layer of Transformer block respectively.

Thanks to these advantages of hierarchical Transformer, the depth estimation network can leverage hierarchical feature maps to generate dense predictions via skip connection strategy with low computational costs and model parameters.

2) *Densely Cascaded Multi-scale Network*: Existing depth estimation networks mostly adopt U-Net style depth decoder [10] which utilizes the ConvNet followed by bilinear interpolation and improve the connection with the hierarchical encoder via skip connection. The problem is that the sole skip connection cannot improve the connectivity between the features inside the depth decoder. To overcome this limited connection, we leverage the dense concatenation [22] and the top-down addition [23] to reinforce the interconnection of decoding layers, simultaneously taking advantage of the improved connection between the encoder and the depth decoder using skip connection.

Hierarchical image features of 4 different resolutions (1/2, 1/4, 1/8, 1/16) from the encoder are passed to each stage of the decoding network. Each feature is first projected to the same C' -dimensional space and the features from different stages are subsequently connected by element-wise addition followed by bilinear interpolation via top-down *cascade* pathway,

$$\mathbf{x}_s = \sum_{k=1}^s (\mathbf{f}_k)_{\uparrow} \quad (2)$$

where \mathbf{f}_k and $(\cdot)_{\uparrow}$ denote projected features from k^{th} stage and bilinear interpolation respectively. The connected features are mapped to the latent space with the same dimension C' with a non-linear feature mapping module in order to alleviate the aliasing effect by bilinear upsampling. For much denser connectivity in the decoder, we directly connect all the subsequent features from the feature mapping module $\mathbf{x}'_1, \dots, \mathbf{x}'_s$ by concatenation to mitigate the attenuation of the low-resolution image features by passing through deep decoding layers. Consequently, the depth decoder of the s^{th} cascade stage $H_s(\cdot)$ receives the feature maps of all the preceding stages as the input:

$$\hat{D}_s = H_s([\mathbf{x}'_1, \dots, \mathbf{x}'_s]_{\uparrow \text{cat}}) \quad (3)$$

where $[\cdot]_{\uparrow \text{cat}}$ and \hat{D}_s denote the channel-wise concatenation of all the elements after bilinear interpolation and the estimated depth map in the s^{th} stage respectively. The estimated depth maps have different spatial size according to the cascade stage of DCMNet so that each stage produces depth outputs with different ratios (1/1, 1/2, 1/4, 1/8) of the input resolution.

Additionally, we employ a Pyramid Pooling Module (PPM) from PSPNet [46] on the first cascade stage and replace it with the regular projection module. As the hierarchical Transformer backbone cannot capture the fully-global representation, the PPM head helps the encoder to understand the entire scene information by providing a global contextual prior. It applies multiple average pooling with different

sizes (1, 2, 3, 6) after the last layer of the Transformer encoder that it brings effective global prior representations for downstream depth estimation decoder.

The depth decoder leverages both 3×3 and 1×1 convolution layers, whereas the projection module and the feature mapping module respectively utilize a single 3×3 and 1×1 convolution layer. Dense features are $2 \times$ upsampled before the first layer of the depth decoder, and no convolution operation changes the spatial size of the features. The architecture of DCMNet is illustrated in Figure 2(b).

B. Unsupervised Depth Estimation

We use monocular image sequences to train a depth estimation network in an unsupervised manner. We jointly train depth estimation and pose estimation network to predict the depth map of the image and the ego-motion of the camera. We build a depth estimation network as a hierarchical Transformer followed by DCMNet and pose estimation networks as ResNet-18 [13] with 6-dimension outputs. As our depth estimation network estimates multi-scale depth maps, we explain the training algorithm with a single scale depth map and its corresponding resized target image I_t for simplicity. The source image $I_{t'}$ for estimating ego-motion is sampled from the temporally adjacent frames of the monocular video $t' \in \{t-1, t+1\}$.

IV. EXPERIMENT

In this section, we demonstrate the effectiveness of our proposed method on KITTI 2015 [24] and Make3D dataset [25]. We also validate that (1) the hierarchical Transformer outperforms depth estimation performance compared to previous ResNet [13] and ViT/DeiT [14], [27] backbone with much smaller model parameters, and (2) how each component of DCMNet contributes to the network performance.

A. Dataset

We train and evaluate our method on KITTI with train/test data split suggested in Eigen [1]. Following SfMLearner [6] and Monodepth2 [7], we remove static frames in monocular sequences, which results in 39,810 images for training and 4,424 images for validation. The resolution of the image is set to 640×192 , which is most widely adopted to evaluate depth estimation performance on KITTI dataset. We also evaluate our method on 134 test images from Make3D dataset without any fine-tuning.

B. Implementation Details

We implement our methods on the public deep learning platform PyTorch [47] and train them on 4 Nvidia RTX 3090 GPUs. We adopt Adam [48] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a batch size of 12 for 40 epochs. We set the initial learning rate as 10^{-4} which is decayed after the first 15 epochs by factor of 10. Similar to [34], [7], we train the model with weights pre-trained on ImageNet-1k [49] to reduce the training time and improve the overall model performance.

We build the hierarchical Transformer and DCMNet to obtain the model size similar to or smaller than the existing

TABLE I: Quantitative results. Comparison of our proposed method to existing methods on KITTI 2015 [24] using train/test split in Eigen [1]. Best results are presented in **bold** and the second best are underlined. All the results are presented without post-processing suggested in [5]. Train column indicates available data during training as; D - Depth supervision, D* - Auxiliary depth supervision, M - Monocular self-supervision, S - Stereo self-supervision and † indicates newer results from github. R18 and R50 denote ResNet-18 and ResNet-50 backbone respectively. The error metrics (Abs Rel, Sq Rel, RSME, RMSE log) indicate *lower is better* and the accuracy metrics ($\delta < 1.25^i$) indicate *higher is better*.

| Method | Train | Abs Rel | Sq Rel | RMSE | RMSE log | $\delta < 1.25$ | $\delta < 1.25^2$ | $\delta < 1.25^3$ |
|--------------------|-------|--------------|--------------|--------------|--------------|-----------------|-------------------|-------------------|
| Eigen [1] | D | 0.203 | 1.548 | 6.307 | 0.282 | 0.702 | 0.890 | 0.890 |
| Liu [28] | D | 0.201 | 1.584 | 6.471 | 0.273 | 0.680 | 0.898 | 0.967 |
| Klodt [29] | D*M | 0.166 | 1.490 | 5.998 | - | 0.778 | 0.919 | 0.966 |
| AdaDepth [30] | D* | 0.167 | 1.257 | 5.578 | 0.237 | 0.771 | 0.922 | 0.971 |
| Kuznetsov [31] | DS | 0.113 | 0.741 | 4.621 | 0.189 | 0.862 | 0.960 | 0.986 |
| DVSO [32] | D*S | 0.097 | 0.734 | 4.442 | 0.187 | 0.888 | 0.958 | 0.980 |
| SVSM FT[33] | DS | <u>0.094</u> | <u>0.626</u> | 4.252 | 0.177 | 0.891 | 0.965 | 0.984 |
| Guo [34] | DS | 0.096 | 0.641 | 4.095 | <u>0.168</u> | <u>0.892</u> | <u>0.967</u> | <u>0.986</u> |
| DORN [3] | D | 0.072 | 0.307 | 2.727 | 0.120 | 0.932 | 0.984 | 0.994 |
| SfMLearner [6] † | M | 0.183 | 1.595 | 6.709 | 0.270 | 0.734 | 0.902 | 0.959 |
| Yang [35] | M | 0.182 | 1.481 | 6.501 | 0.267 | 0.725 | 0.906 | 0.963 |
| Mahjourian [36] | M | 0.163 | 1.240 | 6.220 | 0.250 | 0.762 | 0.916 | 0.968 |
| GeoNet [37] † | M | 0.149 | 1.060 | 5.567 | 0.226 | 0.796 | 0.935 | 0.975 |
| DDVO [38] | M | 0.151 | 1.257 | 5.583 | 0.228 | 0.810 | 0.936 | 0.974 |
| DF-Net [39] | M | 0.150 | 1.124 | 5.507 | 0.223 | 0.806 | 0.933 | 0.973 |
| LEGO [40] | M | 0.162 | 1.352 | 6.276 | 0.252 | - | - | - |
| Ranjan [41] | M | 0.148 | 1.149 | 5.464 | 0.226 | 0.815 | 0.935 | 0.973 |
| EPC++ [42] | M | 0.141 | 1.029 | 5.350 | 0.216 | 0.816 | 0.941 | 0.976 |
| Struct2depth [43] | M | 0.141 | 1.026 | 5.291 | 0.215 | 0.816 | 0.945 | 0.979 |
| Monodepth2 [7] R18 | M | 0.115 | 0.903 | 4.863 | 0.193 | 0.877 | 0.959 | 0.981 |
| Monodepth2 [7] R50 | M | 0.110 | 0.831 | 4.642 | 0.187 | 0.883 | <u>0.962</u> | 0.982 |
| PackNet-SfM [44] | M | 0.111 | <u>0.785</u> | <u>4.601</u> | 0.189 | 0.878 | 0.960 | 0.982 |
| Lyu [45] | M | 0.109 | 0.792 | 4.632 | <u>0.185</u> | <u>0.884</u> | <u>0.962</u> | <u>0.983</u> |
| Ours | M | 0.106 | 0.739 | 4.510 | 0.182 | 0.890 | 0.964 | 0.984 |

unsupervised depth estimation network. We set the window size of the hierarchical Transformer w as 4 and the dimension of the hidden layers in the first block C as 64. The projected dimension C' of DCMNet is set to 128.

C. KITTI Results

We compare our model with other competitive depth estimation algorithms on KITTI 2015 with the image resolution of 640×192 at maximum using both standard error and accuracy metrics. During evaluation, depths are capped to 80 m and per-image median ground truth scaling [6] is applied for monocular self-supervised methods. As shown in Table I, our proposed method outperforms existing state-of-the-art unsupervised approaches. Especially, our method achieves better results compared to recent larger networks, i.e. Monodepth2 [7] with ResNet-50 backbone (35M) and PackNet-SfM [44] (120M), by leveraging relatively small model parameters (25M). Qualitative results are shown in Figure 3 and it demonstrates that our method produces distinct boundaries between objects and background compared to blurry edges shown in competitive unsupervised methods. The estimated depth maps from our methods are capable of recognizing thin, small, but important objects for autonomous driving such as traffic lights, signals and the guardrail, simultaneously recovering the details of the foliage of roadside trees. We attribute such advantages to the window attention-based semi-global feature mapping by hierarchical Transformer which captures both local geometric and global

semantic information of the environment.

D. Ablation Study

For better understanding of our proposed methods, we perform an ablation study in Table II by changing (a) the feature extractor and (b) components of the dense decoder of the depth estimation network. We also compare our method with other hierarchical Transformers (c) for dense prediction.

To explore the effectiveness of the feature extractor in the depth estimation, we compare the Swin Transformer-based hierarchical feature extractor (Swin) with the standard ConvNets, i.e. ResNet, and the previous state-of-the-art Transformer networks, i.e. ViT and DeiT, in Table II(a). We employ U-Net [10] as the baseline architecture and the evaluation is done by replacing only the backbones with all the network kept the same. The feature extractors are pre-trained on ImageNet-1k except for ViT-B/16; ViT-B/16 is pre-trained on ImageNet-21k [53] because it cannot produce high-quality outputs without pre-training on a large-scale dataset.

Swin Transformer-based hierarchical Transformer achieves the best results in most of the metrics compared to other backbones, including both ConvNet-based and Transformer-based, where ResNet-101 with the dilated convolution [54] leverages model parameters in double, and ViT-B/16 and DeiT-B in quadruple.

ViT-B/16 and DeiT-B cannot take advantage of the skip connection strategy as they do not produce hierarchical

TABLE II: Ablation. A comparison between different variants of our model with monocular self-supervision on KITTI 2015 [24] using the Eigen split. (a) The baseline model is U-Net and replace its backbone with ConvNet-based, ViT-based networks, and Swin Transformer. (b) Three components of DCMNet (PPM, Top-down Addition, and Dense Concatenation) are ablated one by one to validate the effectiveness of each component. (c) Our SwinDepth produces on-par performance compared to recent depth estimation using hierarchical Transformers. * indicates no ImageNet pre-training as its pre-trained weights are not provided in the official repository.

| | Method | Param | Skip Connect | PPM | Top-down Addition | Dense Concat | Abs Rel | Sq Rel | RMSE | RMSE log | $\delta < 1.25$ | $\delta < 1.25^2$ | $\delta < 1.25^3$ |
|-----|----------------------------|-----------|--------------|-----|-------------------|--------------|--------------|--------------|--------------|--------------|-----------------|-------------------|-------------------|
| (a) | ResNet-18 + U-Net | 14.8M | ✓ | | | | 0.115 | 0.903 | 4.863 | 0.193 | 0.877 | 0.959 | 0.981 |
| | ResNet-101 Dilated + U-Net | 51.6M | ✓ | | | | 0.110 | 0.876 | 4.853 | 0.189 | 0.879 | 0.961 | 0.982 |
| | ViT-B/16 + U-Net | 88.9M | N/A | | | | 0.132 | 1.298 | 5.494 | 0.216 | 0.851 | 0.947 | 0.974 |
| | DeiT-B + U-Net | 88.9M | N/A | | | | 0.127 | 1.126 | 5.313 | 0.210 | 0.857 | 0.949 | 0.976 |
| | Swin + U-Net | 23.0M | ✓ | | | | 0.111 | 0.864 | 4.671 | 0.189 | 0.887 | 0.962 | 0.982 |
| (b) | Swin + DCMNet w/o PPM | 23.9M | ✓ | | ✓ | ✓ | 0.108 | 0.795 | 4.554 | 0.184 | 0.890 | 0.964 | 0.983 |
| | Swin + DCMNet w/o TA | 25.2M | ✓ | ✓ | | ✓ | 0.107 | 0.747 | 4.551 | 0.183 | 0.888 | 0.964 | 0.984 |
| | Swin + DCMNet w/o DC | 24.3M | ✓ | ✓ | ✓ | | 0.107 | 0.763 | 4.592 | 0.184 | 0.888 | 0.963 | 0.983 |
| | Swin + DCMNet (full) | 25.2M | ✓ | ✓ | ✓ | ✓ | 0.106 | 0.739 | 4.510 | 0.182 | 0.890 | 0.964 | 0.984 |
| | (c) | DPT* [20] | 123M | | | | | 0.210 | 0.847 | 4.554 | 0.184 | 0.847 | 0.948 |
| | HRFormer [50] | 56.0M | | | | | 0.108 | 0.812 | 4.634 | 0.185 | 0.891 | 0.963 | 0.982 |




Fig. 3: Qualitative results on KITTI. Our method performs better on small objects such as traffic lights, signs, and guardrails compared to competitive unsupervised algorithms; Monodepth2 [7] and PackNet-SfM [44].

TABLE III: Make3D results. All results of the monocular self-supervision (M) benefit from median-scaling.

| | Type | Abs Rel | Sq Rel | RMSE | \log_{10} |
|----------------|------|--------------|--------------|--------------|--------------|
| Karsch [51] | D | 0.428 | 5.079 | 8.389 | 0.149 |
| Liu [52] | D | 0.475 | 6.562 | 10.05 | 0.165 |
| Laina [2] | D | 0.204 | 1.840 | 5.683 | 0.084 |
| Monodepth [5] | S | 0.544 | 10.94 | 11.760 | 0.193 |
| SfMLearner [6] | M | 0.383 | 5.321 | 10.470 | 0.478 |
| DDVO [38] | M | 0.387 | 4.720 | 8.090 | 0.204 |
| Monodepth2 [7] | M | 0.322 | 3.589 | 7.417 | 0.163 |
| Ours | M | 0.295 | 3.042 | 6.917 | 0.150 |

feature maps, rather they produce a single resolution feature output. Both show the worst results with the largest model parameters, and it indicates the importance of generating hierarchical feature maps and skip connection in determining the depth estimation performance.

To evaluate how each part of DCMNet contributes to the model performance, we remove three components of the network one by one and report the performance of each ablated model in Table II(b). We observe that the ablation of the top-down cascade pathway via element-wise addition leads to the performance degradation without any parameter loss. A simple top-down addition strategy on the early layer of DCMNet reinforces the feature connections of all the following layers, and leads to the performance improvement in depth estimation. Also, PPM head and dense concatenation improve the performance across all the metrics

with additional model parameters, and our method with all the components combined together achieves the best results.

We compare our SwinDepth to recent hierarchical Transformers for dense prediction. SwinDepth produces depth estimation results on par with state-of-the-art architectures with smaller parameters (DPT [20] (123M), HRFormer [50] (56.0M)).

V. CONCLUSION

In this paper, we present that the convolution-free hierarchical Transformer outperforms existing ConvNet-based architectures in monocular depth estimation. Also, we propose DCMNet to improve the model performance by reinforcing the interconnection of the decoding layers with top-down addition and dense concatenation. We achieve state-of-the-art performance compared to competitive unsupervised algorithms using monocular sequences on the KITTI and Make3D dataset.

VI. ACKNOWLEDGEMENT

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)] and by a research grant from Hyundai Motor Company.

REFERENCES

- [1] Eigen, David, Puhrsch, Christian, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in Neural Information Processing Systems*, vol. 27, pp. 2366–2374, 2014.
- [2] 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.
- [3] 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.
- [4] Garg, Ravi, Bg, V. Kumar, Carneiro, Gustavo, Reid, and Ian, "Unsupervised cnn for single view depth estimation: Geometry to the rescue," in *European conference on computer vision*. Springer, 2016, pp. 740–756.
- [5] Godard, Clément, M. Aodha, Oisín, Brostow, and G. J., "Unsupervised monocular depth estimation with left-right consistency," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 270–279.
- [6] Zhou, Tinghui, Brown, Matthew, Snavely, Noah, Lowe, and D. G., "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.
- [7] Godard, Clément, M. Aodha, Oisín, Firman, Michael, Brostow, and G. J., "Digging into self-supervised monocular depth estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3828–3838.
- [8] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [9] C. Yu, Z. Liu, X.-J. Liu, F. Xie, Y. Yang, Q. Wei, and Q. Fei, "Dslam: A semantic visual slam towards dynamic environments," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1168–1174.
- [10] Ronneberger, Olaf, Fischer, Philipp, Brox, and Thomas, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2015.
- [12] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [13] 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.
- [14] 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," in *International Conference on Learning Representations*, 2020.
- [15] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European Conference on Computer Vision*. Springer, 2020, pp. 213–229.
- [16] J. Beal, E. Kim, E. Tzeng, D. H. Park, A. Zhai, and D. Kislyuk, "Toward transformer-based object detection," 2020.
- [17] S. Zheng, J. Lu, H. Zhao, X. Zhu, Z. Luo, Y. Wang, Y. Fu, J. Feng, T. Xiang, P. H. Torr *et al.*, "Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6881–6890.
- [18] Bhat, S. Farooq, Alhashim, Ibraheem, Wonka, and Peter, "Adabins: Depth estimation using adaptive bins," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4009–4018.
- [19] Z. Li, X. Liu, N. Drenkow, A. Ding, F. X. Creighton, R. H. Taylor, and M. Unberath, "Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers," 2021.
- [20] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," 2021.
- [21] 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," 2021.
- [22] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [23] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [24] 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.
- [25] Saxena, Ashutosh, Sun, Min, and A. Y. Ng, "Make3d: Learning 3d scene structure from a single still image," *IEEE transactions on pattern analysis and machine intelligence*, vol. 31, no. 5, pp. 824–840, 2008.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [27] 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. 10 347–10 357.
- [28] F. Liu, C. Shen, G. Lin, and I. Reid, "Learning depth from single monocular images using deep convolutional neural fields," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 10, pp. 2024–2039, 2015.
- [29] M. Klodt and A. Vedaldi, "Supervising the new with the old: learning sfm from sfm," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 698–713.
- [30] J. N. Kundu, P. K. Uppala, A. Pahuja, and R. V. Babu, "Adadepth: Unsupervised content congruent adaptation for depth estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2656–2665.
- [31] Kuznetsov, Yevhen, Stuckler, Jorg, Leibe, and Bastian, "Semi-supervised deep learning for monocular depth map prediction," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 6647–6655.
- [32] N. Yang, R. Wang, J. Stuckler, and D. Cremers, "Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 817–833.
- [33] Y. Luo, J. Ren, M. Lin, J. Pang, W. Sun, H. Li, and L. Lin, "Single view stereo matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 155–163.
- [34] X. Guo, H. Li, S. Yi, J. Ren, and X. Wang, "Learning monocular depth by distilling cross-domain stereo networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 484–500.
- [35] Z. Yang, P. Wang, W. Xu, L. Zhao, and R. Nevatia, "Unsupervised learning of geometry from videos with edge-aware depth-normal consistency," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [36] Mahjourian, Reza, Wicke, Martin, Angelova, and Anelia, "Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5667–5675.
- [37] Z. Yin and J. Shi, "Geonet: Unsupervised learning of dense depth, optical flow and camera pose," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1983–1992.
- [38] C. Wang, J. M. Buenaposada, R. Zhu, and S. Lucey, "Learning depth from monocular videos using direct methods," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2022–2030.
- [39] Zou, Yuliang, Luo, Zelun, Huang, and Jia-Bin, "Df-net: Unsupervised joint learning of depth and flow using cross-task consistency," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 36–53.
- [40] Z. Yang, P. Wang, Y. Wang, W. Xu, and R. Nevatia, "Lego: Learning edge with geometry all at once by watching videos," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 225–234.

- [41] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 240–12 249.
- [42] C. Luo, Z. Yang, P. Wang, Y. Wang, W. Xu, R. Nevatia, and A. Yuille, "Every pixel counts++: Joint learning of geometry and motion with 3d holistic understanding," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2624–2641, 2019.
- [43] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova, "Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 8001–8008.
- [44] 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.
- [45] X. Lyu, L. Liu, M. Wang, X. Kong, L. Liu, X. Liu, and Y. Yuan, "Hr-depth: High resolution self-supervised monocular depth estimation," in *AAAI*, 2021.
- [46] Zhao, Hengshuang, Shi, Jianping, Qi, Xiaojuan, Wang, Xiaogang, Jia, and Jiaya, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [47] Paszke, Adam, Gross, Sam, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in PyTorch," in *NeurIPS-W*, 2017.
- [48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [49] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [50] Y. Yuan, R. Fu, L. Huang, W. Lin, C. Zhang, X. Chen, and J. Wang, "Hrformer: High-resolution vision transformer for dense predict," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [51] Karsch, Kevin, Liu, Ce, Kang, and S. Bing, "Depth transfer: Depth extraction from video using non-parametric sampling," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 11, pp. 2144–2158, 2014.
- [52] Liu, Miaomiao, Salzmann, Mathieu, He, and Xuming, "Discrete-continuous depth estimation from a single image," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 716–723.
- [53] T. Ridnik, E. Ben-Baruch, A. Noy, and L. Zelnik-Manor, "Imagenet-21k pretraining for the masses," 2021.
- [54] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017.