

S^3 Net: Semantic-Aware Self-supervised Depth Estimation with Monocular Videos and Synthetic Data

Bin Cheng¹, Inderjot Singh Saggu^{2,5}, Raunak Shah³, Gaurav Bansal⁴, and Dinesh Bharadia⁵

¹ Rutgers University, New Brunswick, USA

² Plus.ai, Cupertino, USA **

³ Indian Institute of Technology, Kanpur, India**

⁴ Blue River Technology, Sunnyvale, USA

⁵ University of California, San Diego, USA

cb3974@winlab.rutgers.edu, nderjot.saggu@plus.ai, raunaks@iitk.ac.in, gaurav.bansal@bluerivert.com, dineshb@ucsd.edu

Abstract. Solving depth estimation with monocular cameras enables the possibility of widespread use of cameras as low-cost depth estimation sensors in applications such as autonomous driving and robotics. However, learning such a scalable depth estimation model would require a lot of labeled data which is expensive to collect. There are two popular existing approaches which do not require annotated depth maps: (i) using labeled synthetic and unlabeled real data in an adversarial framework to predict more accurate depth, and (ii) unsupervised models which exploit geometric structure across space and time in monocular video frames. Ideally, we would like to leverage features provided by both approaches as they complement each other; however, existing methods do not adequately exploit these additive benefits. We present S^3 Net, a self-supervised framework which combines these complementary features: we use synthetic and real-world images for training while exploiting geometric, temporal, as well as semantic constraints. Our novel consolidated architecture provides a new state-of-the-art in self-supervised depth estimation using monocular videos. We present a unique way to train this self-supervised framework, and achieve (i) more than 15% improvement over previous synthetic supervised approaches that use domain adaptation and (ii) more than 10% improvement over previous self-supervised approaches which exploit geometric constraints from the real data.

Keywords: Monocular depth prediction, self-supervised learning, domain adaptation, synthetic data, GANs, semantic-aware

1 Introduction

Depth estimation is a fundamental component of 3D scene understanding, with applications in fields such as autonomous driving, robotics and space exploration.

** work done at UCSD

There has been considerable progress in estimating depth through monocular camera images in the last few years, as monocular cameras are inexpensive and widely deployed on many robots. However, building supervised depth estimation algorithms using monocular cameras is challenging, primarily because collecting ground-truth depth maps for training requires a carefully calibrated setup. As an example, many vehicles currently sold in the market have monocular cameras deployed, but there is no trivial way to obtain ground-truth depth information from the images collected from these cameras. Thus, supervised methods for depth estimation suffer due to the unavailability of extensive training labels.

To overcome the lack of depth annotation for monocular camera data, existing work has explored two areas of research: either designing self-supervised/semi-supervised approaches which require minimal labeling, or leveraging labeled synthetic data. Most self-supervised approaches rely on geometric and spatial constraints [44], and have succeeded in reducing the impact of this issue, however they don't always perform well in challenging environments with conditions like limited visibility, object motion, etc. This is because they lack strong training signal from supervision which lets them learn from and generalize to these conditions. In contrast, some effort has been undertaken to use realistic simulated environments to obtain additional synthetic depth data which can be used to compute a supervised training loss.

Synthetic data can be easily generated in different settings with depth labels - for example by varying the lighting conditions, changing the weather, varying object motion, etc. Simply training the original model on synthetic data, however, does not work well in practice as the model does not generalize well to a real-world dataset. To bridge this domain gap between the real-world and synthetic datasets, many domain adaptation techniques have been proposed. Recent works, like [46,27], have found success in using adversarial approaches to address this issue. These solutions typically involve using an adversarial transformation network to align the domains of the synthetic and real-world images, followed by a task network that is responsible for predicting depth. Naturally, we pose the question - can we build a depth estimation network that combines the benefits of information conveyed through real as well as synthetic data?

We present a novel framework S^3 Net that trains the depth network by exploiting these self-supervised constraints (derived from real-world sequential images) and supervised constraints (derived from synthetic data and the respective ground-truth depths). This framework is implemented through several integrated stages which are described below: First, as shown in Fig. 1, we present a novel Generative Adversarial Network (GAN)-based domain adaptation network which exploits geometric constraints across space and time, as well as the semantic consistency between original synthetic images and translated images. These constraints encode additional latent information and thus enhance the quality of domain adaptation. Next, to leverage 'synthetic' supervised cues and 'real' self-supervised cues, we present a novel training approach - weights of the depth estimation network are updated *alternatively* based on the supervised and self-supervised losses. Finally, to impose explicit constraints on object ge-

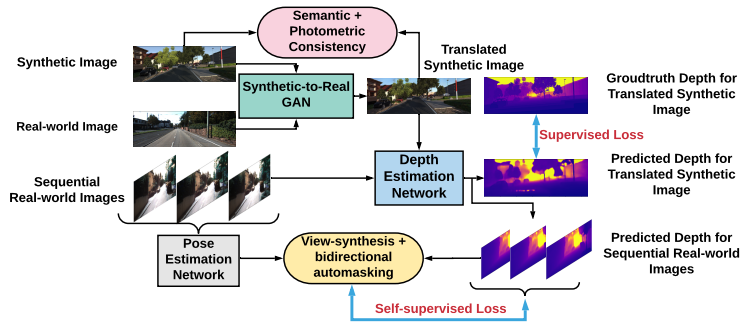


Fig. 1. Overview of our proposed framework: integrating supervised learning on translated synthetic data and self-supervised learning on real videos while imposing spatial, temporal and semantic constraints

ometry we augment the input RGB images with semantic labels, and utilize a bi-directional auto-masking technique to limit pixels which would violate rigid motion constraints.

Novel Adversarial Framework: The key idea of our GAN structure is to utilize flow-based photometric consistency and semantic consistency to better guide the image translation and reduce the domain gap. By utilizing the flow and the sequential translated images, the frame at t can be used to reconstruct the frame at $t + 1$. The photometric differences between the reconstructed frame $t + 1$ and the original frame $t + 1$ are primarily due to imperfect image translation. Moreover, the semantic information should remain consistent before and after the image translation. Therefore, we add both a photometric consistency and a semantic consistency loss to create a novel adversarial framework. These offer additional constraints on the domain adaptation and further improve the image translation performance. They also help increase robustness and reduce undesired artifacts in translated images when compared with traditional approaches.

Semantics and Bi-directional Auto-Masking: Inspired by the auto-masking technique proposed in [16], we propose a novel bi-directional auto-masking technique for sequential real-world images, which can filter out the pixels violating the fundamental rigid motion assumption for self-supervised depth learning. The key difference from a single direction mask is that the bi-directional technique fuses the masks learned by reconstructing frame $t + 1$ from frame t and vice versa, which can substantially increase the accuracy of the proposed mask. Moreover, we augment the input images of our model with semantic labels. The semantic labels can provide explicit geometry constraints, which can be beneficial to further boost the performance of the image translation and the depth estimation.

The challenges of training our depth model fall under two major categories: the GAN networks are unstable during training due to the presence of supervised synthetic losses and self-supervised losses, which results in lack of convergence. We address the convergence issue by proposing a two-phase training strategy.

In the first phase, we train the image translation network and depth estimation network with synthetic supervised losses to stabilize the GAN-based image translation network. In the second phase, we freeze the weights of the image translation network and further train the depth estimation network with both supervised and self-supervised losses.

We evaluate our framework on two challenging datasets, i.e., KITTI [26] and Make3D [35]. The evaluation results show that our proposed model can outperform the state-of-the-art approaches in all evaluated metrics. In particular, we show that S^3 Net can outperform both the state-of-the-art synthetic supervised domain adaptation approaches [27] by $\sim 15\%$ and self-supervised approaches [16] by $\sim 10\%$. Moreover, we only require the depth estimation network during inference, so our inference compute requirements are comparable to previous state-of-the-art approaches.

2 Related Work

2.1 Supervised Depth Estimation

Eigen *et al.* [10] proposed the first supervised learning architecture that models multi-scale information for pixel-level depth estimation using direct regression. Inspired by this work, many follow-up works have extended supervised depth estimation in various directions [19,31,22,41,4,32,8,45,40,42,11,39]. However, acquiring these ground-truth depths is prohibitively expensive. Therefore, it is unlikely to obtain a large amount of labelled training data covering various road conditions, weather conditions, etc, which indicates that these approaches may not generalize well.

One promising approach that reduces the labeling cost is to use synthetically generated data. However, models trained on synthetic data typically perform poorly on real-world images due to a large domain gap. Domain adaptation aims to minimize this gap. Recently, GAN-based approaches show promising performance in domain adaptation [13,14,33,3,34]. Atapour *et al.* [1] proposed a CycleGAN-based translator [51] to translate real-world images into the synthetic domain, and then train a depth prediction network using the synthetic labeled data. Zheng *et al.* [47] propose a novel architecture (T^2 Net) where the style translator and the depth estimation network are optimized jointly so that they can improve each other. Despite promising performance, these approaches inherently suffer from mode collapse and semantic distortion due to imperfect synthetic-to-real image translation. Various constraints and techniques have been proposed to improve the quality of the translated images, but image translation ¹ still remains a challenging task.

2.2 Self-supervised Depth Estimation

In addition to supervised solutions, various approaches have been studied to predict depths by extracting disparity and depth cues from stereo image pairs or

¹ “domain adaptation” and “image translation” are used interchangeably

monocular videos. Garg *et al.* [15] introduced a warping loss based on Taylor expansion. An image reconstruction loss with a spatial smoothness constraint was introduced in [30,49,20] to learn depth and camera motion. Recent works [36,50,24,17,16] aim to improve depth estimation by further exploiting geometry constraints. In particular, Godard *et al.* [17] employed epipolar geometry constraints between stereo image pairs and enforced a left-right consistency constraint in training the network. Yin *et al.*[44] proposed GeoNet, which also used depth and pose networks in order to compute rigid flow between sequential images in a video. More specifically, they introduced a temporal, flow-based photometric loss to predict depth for monocular videos in an unsupervised setting. Bian *et al.* [2] used a similar approach along with a self discovered mask to handle dynamic and occluded objects. Gordon *et al.* [18] also addresses these issues in a purely geometric approach. Casser *et al.* [5] adapts a similar framework with an additional online refinement model during inference. Xu *et al.* [43] uses region deformer networks along with the earlier constraints to handle rigid and non-rigid motion. Zhou *et al.* [48] use a dual network attention based model which processes low and high resolution images separately. Godard *et al.* [16] also presented another unsupervised approach which built on their earlier model [17] by modifying the implementation of the unsupervised constraints.

Another set of recently adopted approaches involves using semantic information, which provides additional constraints on object geometry that can potentially boost the accuracy of depth estimation [25,28,7,29]. Meng *et al.* [25] built on top of [44], and proposed several ways to implement a semantic aided network which helped improve performance. Ranjan *et al.* [29] used a competitive collaboration framework to leverage segmentation maps, pose and flow for depth estimation. However, even with various constraints, the self-supervised approaches predict depth primarily based on indirect and weak-supervision depth cues, which can be easily affected by undesired artifacts, such as motion blurring and low visibility.

Our model architecture is influenced by various previous work, e.g. approaches in [16,47,27,9]. But, compared to these approaches, our S³Net cooperatively combines both supervised depth prediction on synthetic data and self-supervised depth prediction on sequential images, such that the two strategies can complement each other in a mutually beneficial setting.

3 Proposed Methods

We propose a joint framework for monocular depth estimation that is trained on translated synthetic images in a supervised manner and further fine-tuned on sequences of real-world images in a self-supervised fashion. Our proposed framework can be broken down into two main components: a) Synthetic-to-Real Translation and Task Prediction ($G_{S \rightarrow R}$, D_R , f_T), and, b) View-synthesis guided self-supervised fine-tuning ($Pose$, f_T)

Novel GAN Architecture Models trained on synthetic data do not generalize well to real-world data because of domain shift. To address this problem, we

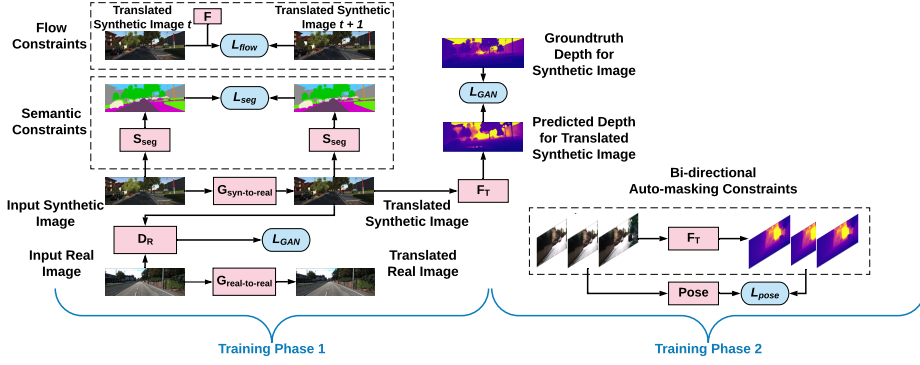


Fig. 2. Our detailed architecture, a) Semantic and photometric consistent GAN for synthetic supervised depth estimation, b) self-supervised architecture trained on sequence of real-world images with warping view-synthesis loss.

build upon the work of T^2 Net [46] for supervised depth estimation on translated synthetic images.

Adversarial Constraints The goal of our generator is to translate a synthetic image X_s to the real domain X_r . To achieve this, a discriminator D_R and a transformer architecture $G_{S \rightarrow R}$ are trained jointly such that discriminator tries to predict if the image is real or synthetic. This accounts for our GAN loss, \mathcal{L}_{GAN} as shown in Fig. 2:

$$\mathcal{L}_{GAN} = \mathbb{E}_{x_r \sim X_r} [\log D_R(x_r)] + \mathbb{E}_{x_s \sim X_s} [\log(1 - D_R(G_{S \rightarrow R}(x_s)))] \quad (1)$$

Identity Constraints To improve the quality of translated images, T^2 Net imposes a identity constraint such that if a real image x_r is given as input, the generator network ($G_{S \rightarrow R}$)’s output should be the identical real image x_r . This additional constraint is incorporated as an identity loss \mathcal{L}_r in Fig. 2:

$$\mathcal{L}_r = \|G_{S \rightarrow R}(x_r) - x_r\|_1 \quad (2)$$

Semantic Consistency While the identity constraint improves upon a vanilla GAN architecture, we observe that the translated images had artifacts as shown in Fig. 4, which cause imperfect domain translation and subsequently hurt depth prediction. To address this we introduce a semantic consistency loss \mathcal{L}_{seg} (Fig. 2) based on the idea that given a semantic segmentation model S_{seg} trained on the source domain, x_s and $G_{S \rightarrow R}(x_s)$ should have identical semantic segmentation maps. This is intuitive as domain translation shouldn’t affect the semantic structure of the image. We enforce this by treating $S_{seg}(x_s)$ as a ground truth label for pixel-wise prediction scores $S_{seg}(G_{S \rightarrow R}(x_s))$. These are used to compute a cross-entropy loss function over semantic labels:

$$\mathcal{L}_{seg} = - \sum_{pixels} \sum_{labels} S_{seg}(x_s) \log(S_{seg}(G_{S \rightarrow R}(x_s))) \quad (3)$$

However, because of domain shift we cannot expect S_{seg} trained on synthetic images to generalize well to the translated image domain, and hence we also continue training S_{seg} while training our GAN architecture so that it can learn features that are generalized to both domains.

Photometric Consistency with Ground-truth Flow In addition to semantic constraints we introduce a flow guided photometric loss[44] to exploit the temporal structure in translated image sequences. By applying ground-truth flow, a frame t can be used to reconstructed the frame $t + 1$. We represent this as a transformation \mathcal{F} . In Eq. 4 below, $G_{S \rightarrow R}(x_{s,t})$ represents the translated image of a synthetic frame at time t and $\mathcal{F}(G_{S \rightarrow R}(x_{s,t}))$ indicates the reconstructed frame $t + 1$ based on frame t . $L_{pe}(\ast)$ computes the photometric differences between the reconstructed frame and the true frame. This photometric loss provides an indirect supervision on the synthetic-to-real image translation.

$$\mathcal{L}_{flow} = L_{pe}(\mathcal{F}(G_{S \rightarrow R}(x_{s,t})), G_{S \rightarrow R}(x_{s,t+1})) \quad (4)$$

Incorporating the above constraints in our GAN framework results in an improved domain translator that is largely devoid of artifacts and preserves semantic structure as shown in Fig. 4.

3.1 Combining Supervised and Self-supervised Depth Estimation

Supervised Depth Estimation on Synthetic Data With the ground-truth depth labels for synthetic data, we formulate the depth estimation on synthetic data as a regression problem. In Eq. 5 below, $f_T(G_{S \rightarrow R}(x_{s,t}))$ is the estimated depth map for frame t and $y_{s,t}$ is the correspond ground-truth label for synthetic frame t .

$$\mathcal{L}_{task} = ||f_T(G_{S \rightarrow R}(x_{s,t})) - y_{s,t}||_1 \quad (5)$$

In accordance with our base network (T^2 Net) we add an edge consistency/awareness loss which penalizes discontinuity (or inconsistency) in the edges between the image x_s and its depth map $f_T(x_s)$.

$$\mathcal{L}_s = |\partial_x f_T(x_r)|e^{-|\partial_x x_r|} + |\partial_y f_T(x_r)|e^{-|\partial_y x_r|} \quad (6)$$

Our training is divided into two phases. In the first phase, we train a GAN-based image transfer network ($(G_{S \rightarrow R}$ and $D_R)$) and a depth estimation network (f_T). A detailed explanation of our training methodology follows in Section 3.4. The first loss objective is a weighted combination of the above constraints (Sections 3.1, 3.2):

$$\mathcal{L}_{phase1} = \mathcal{L}_{GAN} + \alpha_r \mathcal{L}_r + \alpha_{seg} \mathcal{L}_{seg} + \alpha_{flow} \mathcal{L}_{flow} + \alpha_{task} \mathcal{L}_{task} + \alpha_s \mathcal{L}_s \quad (7)$$

where α_r , α_{seg} , α_{flow} , α_{task} , and α_s are hyper-parameters.

Self-supervised Depth Estimation on Monocular Videos In addition to supervised depth prediction on translated synthetic images, we also perform

self-supervised depth estimation on monocular videos. The corresponding pixel coordinates of one rigid object in two consecutive frames follows the relationship

$$\mathbf{p}_{t+1} = \mathbf{K}\mathbf{T}_{t \rightarrow t+1}\mathbf{f}_T(\mathbf{p}_t)\mathbf{K}^{-1}\mathbf{p}_t \quad \mathbf{p}_{t+1} = \mathcal{W}^{-1}\mathbf{p}_t \quad (8)$$

where \mathbf{p}_t and \mathbf{p}_{t+1} are the positions of a pixel in frames at time t and $t + 1$, \mathbf{K} denotes the camera intrinsic parameters, $\mathbf{T}_{t \rightarrow t+1}$ represents the relative camera pose from frame t to from $t + 1$ and \mathcal{W} is the equivalent warping transformation. By sampling these pixels at \mathbf{p}_{t+1} in frame $t + 1$, one can construct frame t from frame $t + 1$. The photometric difference (denoted by L_{pe}) between the constructed image and the true image of frame t provides self-supervision for both the depth network (f_T) and pose estimation networks (Fig. 2).

$$\mathcal{L}_{pose} = L_{pe}(\mathcal{W}(x_{r,t+1}), x_{r,t}) \quad (9)$$

Bi-directional Auto-Masking Inspired by the auto-masking method proposed in [16], we compute the photometric loss for different sequential image pairs, e.g., from frame $t + 1$ to frame t and from frame $t - 1$ to frame t , and then aggregate these photometric losses by extracting their minimum value, i.e., $\min_{t' \in \{\dots t-1, t+1 \dots\}} L_{pe}(\mathcal{W}(x_{r,t'}), x_{r,t})$. Additionally, the pixels satisfying

$$mask = \min_{t' \in \{\dots t-1, t+1 \dots\}} L_{pe}(x_{r,t'}, x_{r,t}) > \min_{t' \in \{\dots t-1, t+1 \dots\}} L_{pe}(\mathcal{W}(x_{r,t'}), x_{r,t})$$

are selected for further loss computation. It is because the discarded pixels are more likely to belong to moving objects with a similar moving speed as the moving camera, or stationary objects captured by a stationary camera. For a more complete loss computation, we consider a bi-directional warping transformation, i.e., from frame t to frame t' as well as from frame t' to frame t .

$$\begin{aligned} \mathcal{L}_{mask} = & mask_{t' \rightarrow t} \circ \min_{t' \in \{\dots t-1, t+1 \dots\}} L_{pe}(\mathcal{W}(x_{r,t'}), x_{r,t}) \\ & + mask_{t \rightarrow t'} \circ \min_{t' \in \{\dots t-1, t+1 \dots\}} L_{pe}(\mathcal{W}(x_{r,t}), x_{r,t'}) \end{aligned} \quad (10)$$

In the second phase of our training, we train the depth and pose networks (f_T and $Pose$) with a combination of sequential real images and GAN translated synthetic images. Our total loss objective which includes training the depth network f_T with supervised loss \mathcal{L}_{task} from synthetic data, can be written as:

$$\mathcal{L}_{phase2} = \alpha_{pose}\mathcal{L}_{pose} + \alpha_{mask}\mathcal{L}_{mask} + \alpha_{task}\mathcal{L}_{task} \quad (11)$$

where α_{pose} , α_{mask} and α_{task} are hyper-parameters.

3.2 Semantic Augmentation

Semantic labels provide important information about object shape and geometry. We believe such information helps improve the accuracy of depth estimation

by imposing additional constraints. For example, on 2D images, the pixels on the object boundaries can have very different depths. Semantic information can help regulate the pixels belonging to certain objects and facilitate the learning process of depth estimation. In this work, to utilize semantic information, we augment the input RGB images with additional semantic labels. We also experimented with augmenting RGB images with semantic labels during synthetic-to-real image translation and obtained substantial improvements in the quality of our translated images.

4 Experiments

In this section, we first present the implementation details of our framework and then compare our framework with other state-of-the-art work on the KITTI and Make3D benchmarks. Finally, we study the importance of each component in our framework through various ablation experiments.

4.1 Implementation Details

Network Implementation Our framework mainly consists of two main sub-modules: (i) the synthetic-to-real image translation network which translates synthetic images to real-style images, and (ii) the depth estimation network which predicts depth maps for both translated synthetic images and real videos. For the synthetic-to-real image translation network, we build our network on top of the T^2 Net architecture [46], with added constraints that use the synthetic ground-truth labels for semantic and optical flow. For the semantic consistency loss we use DeepLab v3+ with a MobileNet backbone as our S_{seg} model. We pre-trained the model on vKITTI, achieving a median IoU of 0.898 on the validation set. We tested the U-Net, VGGNet, and ResNet50 architectures for the depth estimation network and selected the U-Net architecture due to its best performance. A VGG-based architecture was used to estimate the relative camera poses between sequential images. These depth maps and camera poses are subsequently used to compute the self-supervised loss.

Data Pre-processing We use vKITTI [12] and KITTI [26] as the synthetic dataset and the real-world dataset, respectively while training the synthetic-to-real image translation network. The training dataset consists of 20470 images from vKITTI and 41740 images from KITTI. The training images of the KITTI dataset are further divided into small sequences. We use 697 images from KITTI as our test dataset as per the eigen split [10]. The input images are resized to 640×192 (width \times height) during both training and testing. The ground-truth depth information, semantic labels, and optical flow information from the synthetic vKITTI dataset are also used during training. Similar to [46], we clip the synthetic ground truth depths to the range of [0.01, 80] meters. For the real data, the input sequential images are augmented by the corresponding semantic labels, which are generated by DeepLab v3+ model [6].

Model Training Training our model has two major challenges: (i) the training of GAN-based networks is well-known to be unstable; (ii) the depth estimation in our model consists of two components and thus the weights of the depth estimation network are updated by two separate loss functions, which can lead to a convergence issue. To tackle this problem, we design a two-phase training strategy. In the first phase, we pre-train the synthetic-to-real image translation network along with synthetic supervised depth estimation to provide a stable initialization for the image translation network. In phase 2, we freeze the weights of the image translation network and train the depth estimation network using both synthetic supervised and self-supervised losses. We primarily tested two training methods to harmonize the two sources of losses: 1) *weighted sum training*: updating the weights of the depth estimation network based on a weighted sum of the two sources of losses; 2) *alternating training*: alternatively updating the weights of the depth estimation networks by the two sources of losses. We find that the two training methods are resulting in comparable evaluation results but the alternating training provides another control knob to optimize the model training, and generalize well to both the data sources and therefore to unseen datasets. Due to space limitation, we show the results using alternating training only in this paper.

Further, we use the Adam optimizer [21], with initial learning rate of $2e^{-5}$ for the image translation network, $5e^{-5}$ for the depth estimation network, and $5e^{-5}$ for the camera pose estimation network.

Our network was trained on a RTX 2080Ti GPU and the training took 2.6 hours per epoch. On average, our depth estimation network can process 33 frames per second during inference.

4.2 Monocular Depth Estimation on KITTI Dataset

We follow the procedure defined in [46] when evaluating on the KITTI dataset. First, the ground truth depths are generated by projecting 3D LiDAR points to the image plane and then the depth predictions are clipped at a distance of 80m and 50m. The evaluation results are listed in Table 1, where all metrics are computed according to the evaluation strategy proposed in [10]. As shown in the table, our framework shows the best performance across all metrics. We believe this is because our model can combine the benefits of synthetic supervised depth estimation and self-supervised depth estimation. Typical supervised synthetic approaches train models by using low-cost synthetic ground-truth depths, but these approaches also suffer from unstable and inconsistent image translation, leading to less accurate translated images with low resolution. On the other hand, self-supervised approaches can learn the depth from high resolution sequential images; however, these depths are learned from indirect cues which are sensitive to in-view object movements, blockages, etc. Training the model with modified supervised and self-supervised constraints in our consolidated framework ensures that we exploit the best of both worlds, which ultimately leads to better prediction results.

Method	Dataset	Error-related metrics				Accuracy-related metrics		
		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
depth capped at 80m								
Zhou <i>et al.</i> [49]	K	0.183	1.595	6.709	0.270	0.734	0.902	0.959
Yin <i>et al.</i> [44]	K	0.155	1.296	5.857	0.233	0.793	0.931	0.973
Wang <i>et al.</i> [38]	K	0.151	1.257	5.583	0.228	0.810	0.936	0.974
Ramirez <i>et al.</i> [28]	K	0.143	2.161	6.526	0.222	0.850	0.939	0.972
Casser <i>et al.</i> [5]	K	0.141	1.026	5.290	0.215	0.816	0.945	0.979
Ranjan <i>et al.</i> [29]	K	0.140	1.070	5.326	0.217	0.826	0.941	0.975
Xu <i>et al.</i> [43]	K	0.138	1.016	5.352	0.216	0.823	0.943	0.976
Meng <i>et al.</i> [25]	K	0.133	0.905	5.181	0.208	0.825	0.947	0.981
Godard <i>et al.</i> [16] ²	K	0.132	1.044	5.142	0.210	0.845	0.948	0.977
Zheng <i>et al.</i> [46]	K + V	0.174	1.410	6.046	0.253	0.754	0.916	0.966
Mou <i>et al.</i> [27]	K + V	0.145	1.058	5.291	0.215	0.816	0.941	0.977
Ours	K + V	0.124	0.826	4.981	0.200	0.846	0.955	0.982
depth capped at 50m								
Yin <i>et al.</i> [44]	K	0.147	0.936	4.348	0.218	0.810	0.941	0.977
Zheng <i>et al.</i> [46]	K + V	0.168	1.199	4.674	0.243	0.772	0.912	0.966
Mou <i>et al.</i> [27]	K + V	0.139	0.814	3.995	0.203	0.830	0.949	0.980
Ours	K + V	0.118	0.615	3.710	0.187	0.862	0.962	0.984

Table 1. Monocular depth estimation on KITTI dataset with Eigen *et al.* [10] split. The highlighted scores mark the best performance among selected models. In “Dataset” column, “K” and “V” stands for the KITTI and the vKITTI dataset, respectively.

In Fig. 3 we compare qualitative depth estimation results of purely self-supervised GeoNet [44], purely synthetic supervised T^2 Net [46] and our proposed framework. Purely self-supervised approaches results in depth maps which are blurred and do not model depth discontinuity at object boundaries well. On the other hand purely synthetic supervised approach results in sharper depth maps but because of imperfect domain translation it fails to predict depth for surfaces with multiple textures. For example, in the first row of Fig 3, T^2 Net predicts incorrect depth values for the wall on the right because of the window on the wall adding additional texture. These defects severely limit the real-world application of purely self-supervised and synthetic supervised techniques. Our S^3 Net on the other hand generates sharper depth maps than GeoNet and doesn’t suffer from the problems discussed for T^2 Net depth, further proving our point about combining best features from both.

In Fig.4 we compare syn-to-real translated images for T^2 Net GAN and our semantic consistent GAN. Without the presence of a specific task loss, e.g. a depth estimation loss, the resulting objective drives the image translator to generate a realistic interpretation of synthetic images. However, when a task loss is introduced the main objective is shifted to project synthetic images to a space that is optimized for the task. Therefore, some of these differences might not be visually perceivable but can lead to a large gain in performance. Our approach

² For fair comparison, we selected the results for the model without pre-training on the ImageNet dataset




Fig. 3. Qualitative Depth Prediction Results: Column (a) real-world images from KITTI, Column (b), (c), (d) are results for GeoNet [44], T^2 Net [46], and our S^3 Net framework, respectively.

significantly reduces artifacts and successfully retains the semantic structure across synthetic and translated images.




Fig. 4. Translated images: Column (a) input synthetic images; Column (b) T^2 Net; Column (c) our S^3 Net GAN with semantic constraints.

We also evaluate our model on the KITTI odometry dataset. We use 3 sequential images as input and follow the same evaluation strategy as given in [49]. As shown in Table 2, our model outperforms the state-of-the-art approaches by a convincing margin.

4.3 Generalization Study on Make3D Dataset

To show the generalization capability, we also test our framework on the Make3D dataset [35]. We use the model trained using the KITTI dataset and the vKITTI dataset, and evaluate the model on the Make3D test dataset, following the evaluation strategy in [17]. Table 3 shows that our S^3 Net outperforms other self-supervised approaches by a considerable margin. It is because our framework gains knowledge from both synthetic and real datasets. Therefore, our framework sees more scenarios, which further leads to a better generalization performance.

4.4 Ablation Study

In this subsection, we perform a set of ablation experiments on the KITTI dataset to discuss how each individual component in our framework contributes to the final performance. The evaluation results are reported in Table 4.

Method	# of snippets	Seq.09	Seq.10
ORB-SLAM (full)	5	0.014 ± 0.008	0.012 ± 0.011
ORB-SLAM (short)	5	0.064 ± 0.141	0.064 ± 0.130
DDVO (Wang <i>et al.</i> [37])	3	0.045 ± 0.108	0.033 ± 0.074
SfmLearner (Zhou <i>et al.</i> [49])	5	0.021 ± 0.017	0.020 ± 0.015
SfmLearner [49] updated	5	0.016 ± 0.009	0.013 ± 0.009
GeoNet (Yin <i>et al.</i> [44])	5	0.012 ± 0.007	0.012 ± 0.009
MonoDepth2* (Godard <i>et al.</i> [16])	2	0.017 ± 0.008	0.015 ± 0.010
EPC++ (Luo <i>et al.</i> [23])	3	0.013 ± 0.007	0.012 ± 0.008
Ours	3	0.0097 ± 0.0046	0.0099 ± 0.0071

Table 2. Absolute Trajectory Error (ATE) on the KITTI odometry dataset.

Method	Error-related metrics			
	Abs Rel	Sq Rel	RMSE	RMSE log
MonoDepth (Godard <i>et al.</i> [17])	0.544	10.940	11.760	0.193
SfmLearner (Zhou <i>et al.</i> [49])	0.383	5.321	10.47	0.478
T ² Net (Zheng <i>et al.</i> [46])	0.508	6.589	8.935	0.574
MonoDepth2 (Godard <i>et al.</i> [16])	0.322	3.589	7.417	0.163
TCDA (Mou <i>et al.</i> [27])	0.384	3.885	7.645	0.181
Ours (no semantic augmentation)	0.372	5.699	7.844	0.176
Ours (with semantic augmentation)	0.322	3.238	7.187	0.164

Table 3. Error metrics for depth estimation on the Make3D dataset.

Synthetic Supervised Depth Estimation Due to a large domain gap between the synthetic and the real domain, a model that is only trained on synthetic data typically generates unacceptable depth predictions on the real-world data. Synthetic-to-real image translation is one of the most effective remedies for this issue. Even with a native image translation network as proposed in [46], the depth predictions on the real-world data can be improved by about 40%. Additionally, the flow-guided photometric consistency and semantic consistency constraints further regulates the image translation and improves our depth prediction accuracy by another 8.4% and 8%, respectively. Continuing training S_{seg} gives better performance compared to freezing network parameters. This is because S_{seg} is trained on synthetic semantic labels and cannot generalize to translated domain if not trained further.

Synthetic Supervised + Semantic Augmentation We investigate the importance of semantic augmentation to our model by (i) only augmenting the input images for the depth estimation network, but keeping RGB images for the image translation network, and (ii) augmenting the input images for both the depth estimation network and the image translation networks. Compared with the first augmentation strategy, the second strategy introduces a larger improvement. It is because the semantic information can impose additional constraints

Method	Error-related metrics				Accuracy-related metrics		
	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Synthetic Translated Supervised							
Without synthetic-to-real image translation	0.278	3.216	6.268	0.322	0.681	0.854	0.929
Native synthetic-to-real image translation	0.168	1.199	4.674	0.243	0.772	0.912	0.966
With flow-guided photometric consistency	0.1539	0.993	4.4492	0.2241	0.7986	0.9356	0.9752
With semantic consistency (frozen S_{seg})	0.1555	0.9680	4.7412	0.2324	0.7773	0.9245	0.9721
With semantic consistency	0.1544	0.9633	4.7422	0.2322	0.7786	0.9241	0.9727
Synthetic Translated Supervised with Semantic Augmentation							
Semantic augmentation for depth estimation network input only	0.1532	0.9631	4.3872	0.2275	0.7945	0.9325	0.9738
Semantic augmentation for both the image translation networks & the depth estimation networks	0.1455	0.8869	4.2177	0.2154	0.8133	0.9411	0.9773
Synthetic Translated Supervised + Real-world Self-Supervised							
With self-supervised depth estimation	0.1292	0.6969	3.8399	0.1964	0.8428	0.9554	0.9826
With auto-masking	0.1198	0.6671	3.7696	0.1921	0.8637	0.9583	0.9819
With semantic augmentation	0.1183	0.6150	3.7105	0.1876	0.8620	0.9622	0.9844

Table 4. Performance gain in depth estimation from different model components. The predicted depth is capped at 50m

on object geometry and these constraints are useful for regulating the shape of objects and determining the depth prediction on the boundary of objects.

Synthetic Supervised + Self-Supervised By adding photometric losses to real-world sequential images and jointly training the synthetic supervised and self-supervised depth estimation, the depth estimation accuracy on real-world dataset is further improved by 11%. The real-world sequential images are typically clear and accurate, which can compensate for the shortcomings in the imperfect translated images. However, the photometric losses for real-world sequential images are computed based on an assumption that the displacement of pixels is purely caused by movement of the camera. Such an assumption does not always hold. The direct supervision on the synthetic translated images can help alleviate the negative effect of violating this assumption. But, our study indicates that by selecting valid pixels and filtering out the pixels that violate the assumption, our model is further improved by a noticeable margin.

5 Conclusion and Next Steps

In this paper, we present a framework for monocular depth estimation which combines the features of both synthetic images and real video frames in a novel semantic-aware, self-supervised setting. The complexity of our model does not affect its scalability, as we only require a depth network during inference time. We outperform all existing approaches on the KITTI benchmark as well as on our generalization to the new Make3D dataset. These factors contribute to the increased accuracy, scalability, and robustness of our framework as compared to other existing approaches. Our framework extends typical dataset-specific models to improve generalization performance, making it more relevant for real world applications. In the future, we plan to explore strategies which can apply similar frameworks to other related tasks in visual perception.

References

1. Atapour-Abarghouei, A., Breckon, T.P.: Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2800–2810 (2018) [4](#)
2. Bian, J., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M.M., Reid, I.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. In: Advances in Neural Information Processing Systems. pp. 35–45 (2019) [5](#)
3. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3722–3731 (2017) [4](#)
4. Cao, Y., Wu, Z., Shen, C.: Estimating depth from monocular images as classification using deep fully convolutional residual networks. *IEEE Transactions on Circuits and Systems for Video Technology* **28**(11), 3174–3182 (2017) [4](#)
5. Casser, V., Pirk, S., Mahjourian, R., Angelova, A.: Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos (2018) [5](#), [11](#)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018) [9](#)
7. Chen, P.Y., Liu, A.H., Liu, Y.C., Wang, Y.C.F.: Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2624–2632 (2019) [5](#)
8. Chen, W., Fu, Z., Yang, D., Deng, J.: Single-image depth perception in the wild. In: Advances in Neural Information Processing Systems pp. 730–738 (2016) [4](#)
9. Cherian, A., Sullivan, A.: Sem-gan: Semantically-consistent image-to-image translation. CoRR [abs/1807.04409](#) (2018), <http://arxiv.org/abs/1807.04409> [5](#)
10. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: Advances in neural information processing systems. pp. 2366–2374 (2014) [4](#), [9](#), [10](#), [11](#)
11. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [4](#)
12. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4340–4349 (2016) [9](#)
13. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. arXiv preprint [arXiv:1409.7495](#) (2014) [4](#)
14. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* **17**(1), 2096–2030 (2016) [4](#)
15. Garg, R., VijayKumarB., G., Reid, I.D.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: European Conference on Computer Vision (ECCV) (2016) [5](#)
16. Godard, C., Aodha, O.M., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3828–3838 (2019) [3](#), [4](#), [5](#), [8](#), [11](#), [13](#)

17. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017) [5](#), [12](#), [13](#)
18. Gordon, A., Li, H., Jonschkowski, R., Angelova, A.: Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras (2019) [5](#)
19. He, L., Wang, G., Hu, Z.: Learning depth from single images with deep neural network embedding focal length. *IEEE Transactions on Image Processing* **27**(9), 4676–4689 (2018) [4](#)
20. Jason, J.Y., Harley, A.W., Derpanis, K.G.: Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In: European Conference on Computer Vision (ECCV). pp. 3–10 (2016) [5](#)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
22. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence* **38**(10), 2024–2039 (2016) [4](#)
23. Luo, C., Yang, Z., Wang, P., Wang, Y., Xu, W., Nevatia, R., Yuille, A.: Every pixel counts++: Joint learning of geometry and motion with 3d holistic understanding. arXiv preprint arXiv:1810.06125 (2018) [13](#)
24. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [5](#)
25. Meng, Y., Lu, Y., Raj, A., Sunarjo, S., Guo, R., Javidi, T., Bansal, G., Bharadia, D.: Signet: Semantic instance aided unsupervised 3d geometry perception. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9810–9820 (2019) [5](#), [11](#)
26. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3061–3070 (2015) [4](#), [9](#)
27. Mou, Y., Gong, M., Fu, H., Batmanghelich, K., Zhang, K., Tao, D.: Learning depth from monocular videos using synthetic data: A temporally-consistent domain adaptation approach. arXiv preprint arXiv:1907.06882 (2019) [2](#), [4](#), [5](#), [11](#), [13](#)
28. Ramirez, P.Z., Poggi, M., Tosi, F., Mattoccia, S., Stefano, L.D.: Geometry meets semantics for semi-supervised monocular depth estimation (2018) [5](#), [11](#)
29. Ranjan, A., Jampani, V., Balles, L., Kim, K., Sun, D., Wulff, J., Black, M.J.: Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation (2018) [5](#), [11](#)
30. Ren, Z., Yan, J., Ni, B., Liu, B., Yang, X., Zha, H.: Unsupervised deep learning for optical flow estimation. In *AAAI* **3** (2017) [5](#)
31. Repala, V.K., Dubey, S.R.: Dual cnn models for unsupervised monocular depth estimation. arXiv preprint arXiv:1804.06324 (2018) [4](#)
32. Roy, A., Todorovic, S.: Monocular depth estimation using neural regression forest. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5506–5514 (2016) [4](#)
33. Sankaranarayanan, S., Balaji, Y., Castillo, C.D., Chellappa, R.: Generate to adapt: Aligning domains using generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8503–8512 (2018) [4](#)
34. Sankaranarayanan, S., Balaji, Y., Jain, A., Nam Lim, S., Chellappa, R.: Learning from synthetic data: Addressing domain shift for semantic segmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [4](#)

35. Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence* **31**(5), 824–840 (2008) [4](#), [12](#)
36. Vijayanarasimhan, S., Ricco, S., Schmid, C., Sukthankar, R., Fragkiadaki, K.: Sfmnet: Learning of structure and motion from video. (2017), preprint [5](#)
37. Wang, C., Miguel Buenaposada, J., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2022–2030 (2018) [13](#)
38. Wang, C., Miguel Buenaposada, J., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018) [11](#)
39. Xian, K., Shen, C., Cao, Z., Lu, H., Xiao, Y., Li, R., Luo, Z.: Monocular relative depth perception with web stereo data supervision. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018) [4](#)
40. Xu, D., Ricci, E., Ouyang, W., Wang, X., Sebe, N.: Multiscale continuous crfs as sequential deep networks for monocular depth estimation. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (July 2017) [4](#)
41. Xu, D., Wang, W., Tang, H., Liu, H., Sebe, N., Ricci, E.: Structured attention guided convolutional neural fields for monocular depth estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3917–3925 (2018) [4](#)
42. Xu, D., Wang, W., Tang, H., Liu, H., Sebe, N., Ricci, E.: Structured attention guided convolutional neural fields for monocular depth estimation. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018) [4](#)
43. Xu, H., Zheng, J., Cai, J., Zhang, J.: Region deformer networks for unsupervised depth estimation from unconstrained monocular videos (2019) [5](#), [11](#)
44. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1983–1992 (2018) [2](#), [5](#), [7](#), [11](#), [12](#), [13](#)
45. Zhang, Z., Schwing, A.G., Fidler, S., Urtasun, R.: Monocular object instance segmentation and depth ordering with cnns. *2015 IEEE International Conference on Computer Vision (ICCV)* pp. 2614–2622 (2015) [4](#)
46. Zheng, C., Cham, T.J., Cai, J.: T2net: Synthetic-to-realistic translation for solving single-image depth estimation tasks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 767–783 (2018) [2](#), [6](#), [9](#), [10](#), [11](#), [12](#), [13](#)
47. Zheng, C., Cham, T.J., Cai, J.: T2net: Synthetic-to-realistic translation for solving single-image depth estimation tasks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 767–783 (2018) [4](#), [5](#)
48. Zhou, J., Wang, Y., Qin, K., Zeng, W.: Unsupervised high-resolution depth learning from videos with dual networks (2019) [5](#)
49. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *Conference, I. (ed.) IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. page 7 (2017) [5](#), [11](#), [12](#), [13](#)
50. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *CVPR* (2017) [5](#)
51. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2223–2232 (2017) [4](#)