

Kick Back & Relax: Learning to Reconstruct the World by Watching SlowTV

Jaime Spencer
University of Surrey
j.spencermartin@surrey.ac.uk

Chris Russell
Oxford Internet Institute
christopher.m.russell@gmail.com

Simon Hadfield
University of Surrey
s.hadfield@surrey.ac.uk

Richard Bowden
University of Surrey
r.bowden@surrey.ac.uk




Figure 1: Zero-shot Generalization. We present the first SS-MDE model capable of generalizing to a wide-range of complex environments. This is achieved by training on the novel large-scale SlowTV dataset. We outperform other existing self-supervised methods and perform on par with recent supervised SotA [47, 46, 72].

Abstract

Self-supervised monocular depth estimation (SS-MDE) has the potential to scale to vast quantities of data. Unfortunately, existing approaches limit themselves to the automotive domain, resulting in models incapable of generalizing to complex environments such as natural or indoor settings.

To address this, we propose a large-scale SlowTV dataset curated from YouTube, containing an order of magnitude more data than existing automotive datasets. SlowTV contains 1.7M images from a rich diversity of environments, such as worldwide seasonal hiking, scenic driving and scuba diving. Using this dataset, we train an SS-MDE model that provides zero-shot generalization to a large collection of indoor/outdoor datasets. The resulting model outperforms all existing SSL approaches and closes the gap on supervised SoTA, despite using a more efficient architecture.

We additionally introduce a collection of best-practices to further maximize performance and zero-shot generalization. This includes 1) aspect ratio augmentation, 2) camera intrinsic estimation, 3) support frame randomization and 4) flexible motion estimation. Code is available at https://github.com/jspenmar/slowtv_monodepth.

1. Introduction

Reliably reconstructing the 3-D structure of the environment is a crucial component of many computer vision pipelines, including autonomous driving, robotics, augmented reality and scene understanding. Despite being an inherently ill-posed task, monocular depth estimation (MDE) has become of great interest due to its flexibility and applicability to many fields.

While traditional supervised methods achieve impressive results, they are limited both by the availability and quality of annotated datasets. LiDAR data is expensive to collect and frequently exhibits boundary artefacts due to motion correction. Meanwhile, Structure-from-Motion (SfM) is computationally expensive and can produce noisy, incomplete or incorrect reconstructions. Self-supervised learning (SSL) instead leverages the photometric consistency across frames to simultaneously learn depth and Visual Odometry (VO) without ground-truth annotations. As only stereo or monocular video is required, SSL has the potential to scale to much larger data quantities.

Unfortunately, existing SS-MDE approaches have relied exclusively on automotive data [18, 13, 24]. The limited diversity of training environments results in models incapable

of generalizing to different scene types (*e.g.* natural or indoors) or even other automotive datasets. Moreover, despite being fully convolutional, these models struggle to adapt to different image sizes. This further reduces performance on sources other than the original dataset.

Inspired by the recent success of supervised MDE [36, 47, 46], we develop an SS-MDE model capable of performing zero-shot generalization beyond the automotive domain. In doing so, we aim to bridge the performance gap between supervision and self-supervision. Unfortunately, most existing supervised datasets are unsuitable for SSL, as they consist of isolated image and depth pairs. On the other hand, existing SSL datasets focus only on the automotive domain.

To overcome this, we make use of SlowTV as an untapped source of high-quality data. SlowTV is a television programming approach originating from Norway consisting of long, uninterrupted shots of relaxing events, such as train or boat journeys, nature hikes and driving. This represents an ideal training source for SS-MDE, as it provides large quantities of data from highly diverse environments, usually with smooth motion and limited dynamic objects.

To improve the diversity of available data for SS-MDE, we have collated the **SlowTV** dataset, consisting of 1.7M frames from 40 videos curated from YouTube. This dataset consists of three main categories—natural, driving and underwater—each featuring a rich and diverse set of scenes. We combine SlowTV with Mannequin Challenge [31] and Kitty [18] to train our proposed models. SlowTV provides a general distribution across a wide range of natural scenes, while Mannequin Challenge covers indoor scenes with humans and Kitty focuses on urban scenes. The resulting models are trained with an order of magnitude more data than any existing SS-MDE approach. Contrary to many supervised approaches [4, 72], we train a single model capable of generalizing to all scene types, rather than separate indoor/outdoor models. This closely resembles the zero-shot evaluation proposed by MiDaS [47] for supervised MDE.

The contributions of this paper can be summarized as:

1. We introduce a novel SS-MDE dataset of SlowTV YouTube videos, consisting of 1.7M images. It features a diverse range of environments including world-wide seasonal hiking, scenic driving and scuba diving.
2. We leverage SlowTV to train zero-shot models capable of adapting to a wide range of scenes. The models are evaluated on 7 datasets unseen during training.
3. We show that existing models fail to generalize to different image shapes and propose an aspect ratio augmentation to mitigate this.
4. We greatly reduce the performance gap w.r.t. supervised models, improving the applicability of SS-MDE to the real-world. We make the dataset, pretrained model and code available to the public.

2. Related Work

Garg *et al.* [17] proposed the first algorithm for SS-MDE, where the target view was synthesized using its stereo pair and predicted depth map. Monodepth [19] greatly improved performance by incorporating differentiable bilinear interpolation [27], an SSIM-weighted reconstruction loss [64] and left-right consistency. SfM-Learner [77] extended SS-MDE into the purely monocular domain by replacing the fixed stereo transform with a trainable VO network. DDVO [63] further refined the predicted motion with a differentiable DSO module [16].

Purely monocular approaches are highly sensitive to dynamic objects, which cause incorrect correspondences. Many works have tried to minimize this impact by introducing predictive masking [77], uncertainty estimation [30, 69, 45], optical flow [70, 48, 35] and motion masks [23, 9, 14]. Monodepth2 [20] proposed the minimum reconstruction loss and static automasking, encouraging the loss to optimize unoccluded pixels and preventing holes in the depth.

Other methods focused on the robustness of the photometric loss. This was achieved through the use of pretrained [74] or learnt [53, 52] feature descriptors and semantic constraints [11, 25, 29]. Mahjourian *et al.* [39] and Bian *et al.* [6] complemented the photometric loss with geometric constraints. ManyDepth [65] additionally incorporated the previous frame’s prediction into a cost volume.

Complementary to these developments, other works proposed changes to the network architecture, including both the encoder [24, 22, 76], and decoder [44, 24, 68, 76, 38, 75]. Akin to supervised MDE developments [4, 5], Johnston *et al.* [28] and Bello *et al.* [21, 22] obtained improvements by representing depth as a discrete volume.

Finally, several works have complemented self-supervision with proxy depth regression. These are typically obtained from SLAM [30, 49], synthetic data [37] or hand-crafted disparity estimation [60, 65]. In particular, DepthHints [65] improved the proxy depth robustness by generating estimates with multiple hyperparameters.

The works described here train exclusively on automotive data, such as Kitty [18], CityScapes [13] or DDAD [24]. Recent benchmark studies [56, 54] have shown that this lack of variety limits generalization to out-of-distribution domains, such as forests, natural or indoor scenes. We propose to greatly increase the diversity and scale of the training data by leveraging unlabelled videos from YouTube, without requiring manual annotation or expensive pre-processing.

3. SlowTV Dataset

SlowTV is a style of TV programming featuring uninterrupted shots of long-duration events. Our dataset consists of 40 curated videos ranging from 1–8 hours and a total of 135 hours.



Figure 2: SlowTV. Sample images from the proposed dataset, featuring diverse scenes for hiking, driving and scuba diving. The dataset consists of 40 videos curated from YouTube, totalling to 1.7M frames. Diversifying the training data allows our SS-MDE models to generalize to unseen datasets.

Table 1: Datasets Comparison. The top half shows commonly used SS-MDE training datasets. The proposed SlowTV greatly diversifies training environments and scales to much larger quantities. The bottom half summarizes the testing datasets used in our *zero-shot generalization* evaluation.

	Urban	Natural	Scuba	Indoor	Depth	Acc	Density	#Img
Kitti [18, 61] [†]	✓	✗	✗	✗	LiDAR	High	Low	71k
DDAD [24]	✓	✗	✗	✗	LiDAR	Mid	Low	76k
CityScapes [13]	✓	✗	✗	✗	Stereo	Low	Mid	88k
Mannequin [31] [†]	✓	✗	✗	✓	SfM	Mid	Mid	115k
SlowTV (Ours)[†]	✓	✓	✓	✗	✗	✗	✗	1.7M
Kitti [18, 61]	✓	✗	✗	✗	LiDAR	High	Low	652
DDAD [24]	✓	✗	✗	✗	LiDAR	Mid	Low	1k
Sintel [7]	✗	✓	✗	✗	Synth	High	High	1064
SYNS-Patches [1, 56]	✓	✓	✗	✓	LiDAR	High	High	775
DIODE [62]	✓	✗	✗	✓	LiDAR	High	High	771
Mannequin [31]	✓	✗	✗	✓	SfM	Mid	Mid	1k
NYUD-v2 [41]	✗	✗	✗	✓	Kinect	Mid	High	654
TUM-RGBD [57]	✗	✗	✗	✓	Kinect	Mid	High	2.5k

[†] Datasets used to train our networks.

We focus on three categories: hiking, driving and scuba diving. Hiking videos target natural settings, including forests, mountains or fields, which are non-existent in current datasets. These videos were collected in a diverse set of locations and conditions. This includes the USA, Canada, the Balkans, Eastern Europe, Indonesia and Hawaii, and conditions such as rain, snow, autumn and summer.

Existing automotive datasets tend to focus on urban driving in densely populated cities [18, 13, 26, 10, 24, 71, 8]. Our SlowTV dataset features complementary data in the form of long drives in scenic routes, such as mountain and natural trails. Finally, underwater is an otherwise unused domain, which increases the diversity of the training data and prevents overfitting to purely urban scenes. Figure 2 shows the variability of the proposed dataset, with additional examples and details in Appendix A.

Videos were downloaded at HD resolution (720×1280) and extracted at 10 FPS to reduce storage, while still providing smooth motion and large overlap between adjacent frames. To make the dataset size tractable and reduce self-similarity, only 100 consecutive frames out of every 250 were retained. Despite this, the final training dataset con-

sists of a total of 1.7M images, composed of 1.1M natural, 400k driving and 180k underwater. Table 1 compares existing datasets with those used in this publication.

Since our dataset targets self-supervised methods, the only annotations required are the camera intrinsic parameters. We apply COLMAP [51] to a sub-sequence to estimate the intrinsics for each video. However, as discussed in Section 4.2, it is possible to let the network jointly optimize camera parameters alongside depth and motion. This improves performance and results in a truly self-supervised perception and navigation framework, requiring only monocular video to learn how to reconstruct.

4. Methodology

MDE is an alternative to traditional depth estimation techniques, such as stereo matching and cost volumes. Rather than relying on multi-view images, these depth networks take only a single image as input. From this image, a disparity or inverse depth map is estimated as $\hat{\mathbf{D}}_t = \Phi_D(\mathbf{I}_t)$, where Φ_D represents a trainable DNN, \mathbf{I}_t is the target image at time-step t and $\hat{\mathbf{D}}_t$ the predicted sigmoid disparity.

As SlowTV contains only monocular videos, we adopt a fully monocular pipeline [77], whereby our framework also estimates the relative pose $\hat{\mathbf{P}}_{t+k}$ between the target \mathbf{I}_t and support frames \mathbf{I}_{t+k} , where $k = \pm 1$ is the offset between adjacent frames. This is represented as $\hat{\mathbf{P}}_{t+k} = \Phi_P(\mathbf{I}_t \oplus \mathbf{I}_{t+k})$, where \oplus is channel-wise concatenation. Pose is predicted as a translation and axis-angle rotation.

4.1. Losses

The correspondences required to warp the support frames and compute the photometric loss are given by back-projecting the depth and re-projecting onto each support frame. This process is summarized as

$$\mathbf{p}'_{t+k} = \mathbf{K}\hat{\mathbf{P}}_{t+k}\mathbf{D}_t(\mathbf{p}_t)\mathbf{K}^{-1}\mathbf{p}_t, \quad (1)$$

where \mathbf{K} are the camera intrinsic parameters, \mathbf{D}_t is the inverted and scaled disparity prediction $\hat{\mathbf{D}}_t$, \mathbf{p}_t are the 2-D pixel coordinates in the target frame and \mathbf{p}'_{t+k} are the re-projected coordinates in the support frame. We omit the transformation to homogeneous coordinates for simplicity.

The warped support frames are then given by $\mathbf{I}'_{t+k} = \mathbf{I}_{t+k} \langle \mathbf{p}'_{t+k} \rangle$, where $\langle \cdot \rangle$ represents differentiable bilinear interpolation [27]. These warped frames are used to compute the photometric loss w.r.t. the original target frame. As is common, we use the weighted combination of SSIM+L₁ [19], given by

$$\mathcal{L}_{ph}(\mathbf{I}, \mathbf{I}') = \lambda \frac{1 - \mathcal{L}_{ssim}(\mathbf{I}, \mathbf{I}')}{2} + (1 - \lambda) \mathcal{L}_l(\mathbf{I}, \mathbf{I}'), \quad (2)$$

where $\lambda = 0.85$ is the loss balancing weight.

While Mannequin Challenge consists almost exclusively of static scenes, Kitti and SlowTV contain dynamic objects, such as vehicles, hikers, and wild marine life. Rather than introducing motion masks [23, 9, 14], commonly requiring semantic segmentation, we opt for the minimum reconstruction loss [20]. This loss reduces the impact of occluded pixels by optimizing only the pixels with the smallest loss across all support frames and is computed as

$$\mathcal{L}_{rec} = \sum_{\mathbf{p}} \min_k \mathcal{L}_{ph}(\mathbf{I}_t, \mathbf{I}'_{t+k}), \quad (3)$$

where \sum indicates averaging over a set.

Finally, automasking [20] helps remove holes of infinite depth caused by static frames and objects moving at similar speeds to the camera. Automasking simply discards pixels where the photometric loss for the *unwarped* target frame is lower than the loss for the synthesized view, given by

$$\mathbb{M} = \left[\left[\min_k \mathcal{L}_{ph}(\mathbf{I}_t, \mathbf{I}'_{t+k}) < \min_k \mathcal{L}_{ph}(\mathbf{I}_t, \mathbf{I}_{t+k}) \right] \right], \quad (4)$$

where $[\cdot]$ represents the Iverson brackets. Additional results showing the effectiveness of the minimum reconstruction loss and automasking can be found in Appendix E.

This reconstruction loss is complemented by the common edge-aware smoothness regularization [19]. These networks and losses constitute the core baseline required to train the desired zero-shot depth estimation models. To improve existing performance and generalization, we incorporate several new components into the pipeline.

4.2. Learning Camera Intrinsic

As discussed in Section 3, we use COLMAP to estimate camera intrinsic for each dataset video. Whilst this is significantly less computationally demanding than obtaining full reconstructions, it introduces additional pre-processing requirements. Eliminating this step would simplify dataset collection and allow for even easier scale-up.

We take inspiration from [23, 12] and predict camera intrinsic using the pose network Φ_p . This is achieved by adding two decoder branches with the same architecture used to predict pose. The modified network is defined as




Figure 3: Generalizing to Image Shapes. The same model, at different resolutions, can produce significantly different predictions. Distorting the image (and resizing the prediction) can improve performance, despite introducing artefacts. Note the improved boundary sharpness in (d).

$\hat{\mathbf{P}}_{t+k}, \mathbf{f}_{xy}, \mathbf{c}_{xy} = \Phi_P(\mathbf{I}_t \oplus \mathbf{I}_{t+k})$, where \mathbf{f}_{xy} and \mathbf{c}_{xy} are the focal lengths and principal point.

Both quantities are predicted as normalized and scaled by the image shape prior to combining them into \mathbf{K} . The focal length decoder uses a softplus activation to guarantee a positive output. The principal point instead uses a sigmoid, under the assumption that it will lie within the image. All parameters—depth, pose and intrinsic—are optimized simultaneously, as they all establish the correspondences across support frames, given by (1).

4.3. Aspect Ratio Augmentation

The depth network is commonly a fully convolutional network that can process images of any size. In practice, these networks can overfit to the training size, resulting in poor out-of-dataset performance. Figure 3 shows this effect, where resizing to the training resolution improves results, despite introducing stretching or squashing distortions.

Since both SlowTV and Mannequin Challenge were sourced from YouTube, they feature the common widescreen aspect ratio (16:9). However, the objective is to train a model that can be easily applied to real-world settings in a zero-shot fashion. To this end, we propose an aspect ratio augmentation (AR-Aug) that randomizes the image shape during training, increasing the data diversity.

AR-Aug has two components: centre cropping and resizing. The cropping stage uniformly samples from a set of predefined aspect ratios. A random crop is generated using this aspect ratio, covering 50-100% of the original height or width. By definition, the sampled crop will be smaller than the original image and of different shape. The crop is therefore resized to match the number of pixels in the original image. Appendix B details the full set of aspect ratios used and shows training images obtained using this procedure.

AR-Aug has the effect of drastically increasing the distribution of image shapes, aspect ratios and object scales seen by the network during training. As shown in Section 5.4, this greatly increases performance, especially when evaluating on datasets with different image sizes.

5. Results

We evaluate the proposed models in a variety of settings and datasets, including in-distribution and zero-shot. Since the trained model is purely monocular, the predicted depth is in arbitrary units. Instead of using traditional median alignment [77, 20], we follow MiDaS [47] and estimate scale and shift alignment parameters based on a least-squares criterion. We apply the same strategy to every baseline. Results using median alignment are shown in Appendix F. Note that datasets with SfM ground-truth (e.g. Mannequin Challenge) are also scaleless and would require this step even for techniques that predict metric depth.

5.1. Implementation Details

The proposed models are implemented in PyTorch [43] using the baselines from the Monodepth Benchmark [56]. The depth network uses a pretrained ConvNeXt-B backbone [33, 66] and a DispNet decoder [40, 19]. The pose network instead uses ConvNeXt-T for efficiency. Each model variant is trained with three random seeds and we report average performance. This improves the reliability of the results and reduces the impact of non-determinism.

The final models were trained on a combination of SlowTV (1.7M), Mannequin Challenge (115k) and Kitti Eigen-Benchmark (71k). To make the duration of each epoch tractable and balance the contribution of each dataset, we fix the number of images per epoch to 30k, 15k and 15k, respectively. The subset sampled from each dataset varies with each epoch to ensure a high data diversity.

The models were trained for 60 epochs using AdamW [34] with weight decay 10^{-3} and a base learning rate of 10^{-4} , decreased by a factor of 10 for the final 20 epochs. Empirically, we found that linearly warming up the learning rate for the first few epochs stabilized learning and prevented model collapse. We use a batch size of 4 and train the models on a single NVIDIA GeForce RTX 3090.

SlowTV and Mannequin Challenge use a base image size of 384×640 , while Kitti uses 192×640 . As is common, we apply horizontal flipping and colour jittering augmentations with 50% probability. AR-Aug is applied with 70% probability, sampling from 16 predefined aspect ratios. The full set of aspect ratios can be found in Appendix B.

Since existing models are trained exclusively on automotive data, most of the motion occurs in a straight-line and forward-facing direction. It is therefore common practice to force the network to always make a forward-motion prediction by reversing the target and support frame if required. Handheld videos, while still primarily featuring forward motion, also exhibit more complex motion patterns. As such, removing the forward motion constraint results in a more flexible model that improves performance.

Similarly, existing models are trained with a fixed set of support frames—usually previous and next. Since SlowTV

and Mannequin Challenge are mostly composed of handheld videos, the change from frame-to-frame is greatly reduced. We make the model more robust to different motion scales and appearance changes by randomizing the separation between target and support frames. In general, we sample such that handheld videos use a wider time-gap between frames, while automotive has a small time-gap to ensure there is significant overlap between frames. As shown later, this leads to further improvements and greater flexibility.

5.2. Baselines

We use the SSL baselines from [56], trained on Kitti Eigen-Zhou with a ConvNeXt-B backbone. We minimize architecture changes and training settings w.r.t. the baselines to ensure models are comparable and improvements are solely due to the contributions from this paper.

We also report results for recent State-of-the-Art (SotA) supervised MDE approaches, namely MiDaS [47], DPT [46] and NeWCRFs [72]. MiDaS and DPT were trained on a large collection of supervised datasets that do not overlap with our testing datasets (unless otherwise indicated). As such, these models are also evaluated in a zero-shot fashion. We use the pre-trained models and pre-processing provided by the PyTorch Hub. NeWCRFs provides separate indoor/outdoor models, trained on Kitti and NYUD-v2 respectively. We evaluate the corresponding model in a zero-shot manner depending on the dataset category. Despite predicting metric depth, we apply scale and shift alignment to ensure results are comparable.

5.3. Evaluation Metrics

We report the following metrics per dataset:

Rel. Absolute relative error (%) between target y and prediction \hat{y} as $\text{Rel} = \sum |y - \hat{y}| / y$.

Delta. Prediction threshold accuracy (%) as $\delta_{.25} = \sum (\max(\hat{y}/y, y/\hat{y}) < 1.25)$.

F. Pointcloud reconstruction F-Score [42] (%) as $F = (2PR) / (P + R)$, where P and R are the Precision and Accuracy of the 3-D reconstruction with a correctness threshold of 10cm.

Table 2: Model Complexity. Supervised SotA approaches make use of computationally expensive transformer backbones. Despite being of equivalent complexity to the SSL baselines [56], our model closes the gap to supervised performance.

	Backbone	MParam↓	FPS↑
KBR (Ours)	ConvNeXt-B [33]	92.65	61.50
MiDaS [47]	ResNeXt-101 [67]	105.36	51.38
DPT [46]	ViT-L [15]	344.06	14.54
DPT [46]	BEiT-L [3]	345.01	9.60
NeWCRFs [73]	Swin [32]	270.44	21.61

We additionally compute multi-task metrics to summarize the performance across all datasets:

Rank. Average ordinal ranking order across all metrics as $\text{Rank} = \sum_m r_m$, where m represents each available metric and r is the ordinal rank.

Improvement. Average relative performance increase (%) across all metrics as $\Delta = \sum_m (-1)^{l_m} (M_m - M_m^0) / M_m^0$, where $l_m = 1$ if lower is better, M_m is the performance for a given metric and M_m^0 is the baseline’s performance.

5.4. Ablation

We perform a “leave-one-out” ablation study, whereby a single component is removed per-experiment from the full model. This helps to understand the impact of each proposed contribution. We report this ablation on Kitti Eigen-Zhou, Mannequin Challenge and SYNS-Patches.

As shown in Table 3, the full model with all contributions performs best. *Fwd P* represents a network forced to always predict forward-motion. $k = \pm 1$ uses fixed support frames, instead of the randomization in Section 5.1. *Fixed K* removes the learnt intrinsics from Section 4.2, while *No AR-Aug* removes the aspect ratio augmentation. It is worth noting that none of these contributions increase the number of depth network parameters. Learning the intrinsics results in a negligible increase in the pose network, which is not required for inference. Despite this, each contribution significantly improves accuracy and generalization.

5.5. In-distribution

We compare our best approach—Kick Back & Relax (KBR)—against existing SotA on the two training datasets with ground-truth: Kitti and Mannequin Challenge. This represents the most common evaluation, where the test data is sampled from the same distribution as the training data.

As shown in Table 4 (*In-Distribution*), all variants of the proposed models outperform the improved SSL baselines from [56]. Even more surprising, our models also outperform most *supervised* baselines on Kitti, despite DPT-BEiT

Table 3: Leave-one-out Ablation. We study the contribution of each proposed component. Randomizing the support frames, learning camera parameters and augmenting the image shape all contribute to improving overall performance.

	Multi-task		Kitti Eigen-Zhou			Mannequin			SYNS (Val)		
	R↓	Δ↑	Rel↓	F↑	δ.25↑	Rel↓	F↑	δ.25↑	Rel↓	F↑	δ.25↑
Full	2.20	0.00	6.16	57.60	95.52	14.39	17.67	82.23	20.34	17.08	69.88
Fwd <i>P</i>	2.60	-0.08	6.18	57.47	95.47	14.36	17.52	82.22	20.24	17.20	69.52
$k = \pm 1$	2.30	-0.60	6.03	58.23	95.67	14.17	17.92	82.43	21.04	16.04	68.33
Fixed <i>K</i>	4.00	-1.46	6.30	56.93	95.38	14.95	17.11	81.00	20.46	17.11	69.56
No AR-Aug	4.50	-4.72	7.42	52.89	93.99	14.32	17.87	82.16	21.32	16.10	67.64
None	5.40	-5.52	7.47	51.83	94.19	14.62	17.01	81.29	21.21	16.72	67.03

Highlighted cells indicate zero-shot results.

being trained on it. NeWCRFs is the only supervised model to outperform ours by a slight margin. This may be due to the additional automotive data from SlowTV, which increases the variety and improves generalization. Finally, our model outperforms even the supervised SotA on Mannequin Challenge F-Score.

5.6. Zero-shot Generalization

The core of our evaluation takes place in a *zero-shot* setting, *i.e.* models are not fine-tuned. This demonstrates the capability of our model to generalize to previously unseen environments. While several existing SS-MDE approaches provide zero-shot evaluations, this is usually limited to CityScapes [13] and Make3D [50]. These datasets provide low-quality ground-truths and focus exclusively on urban environments similar to Kitti. We instead opt for a collection of challenging datasets, constituting a mixture of urban, natural, synthetic and indoor scenes.

Outdoor. These results can be found in Table 4 (*Outdoor*), where all evaluated models are zero-shot. Once again, our models outperform the SSL baselines in every metric, across all datasets. NeWCRFs is capable of generalizing to other automotive datasets and provides good performance on DDAD. However, our model adapts better to complex synthetic (Sintel) and natural (SYNS-Patches) scenes. Despite being fully self-supervised and requiring no depth annotations during training, our model outperforms MiDaS and DPT-ViT. DPT leverages expensive transformer-based backbones and additional datasets to improve performance.

Indoor. Table 4 (*Indoor*) shows results for all indoor datasets. Note that NeWCRFs was trained exclusively on NYUD-v2, while DPT-BEiT used it as part of its training collection. As such, this subset of results is *not* zero-shot. As with the outdoor evaluations, our model provides significant improvements over all existing SSL approaches. This is due to the focus on Kitti and the lack of indoor training data, highlighting the need for more varied training sources. However, the supervised models still provide improvements over our method, likely due to the additional indoor datasets used for training. Once again, we emphasize that our model is *fully self-supervised*. Despite this, we close the performance gap on complex supervised models.

Visualizations. We visualize the network predictions in Figure 4. As seen, the proposed model clearly outperforms the best SSL baseline. This is most noticeable in indoor settings, where the baseline treats human faces as background. In many cases, our self-supervised model provides similar or better depth maps than the supervised baselines. Once again, these rely on ground-truth annotations and expensive transformer-based backbones. Meanwhile, our model simply requires curated collections of freely-available monocular YouTube videos, without even camera intrinsics.

Failure Cases. Our approach does not explicitly use explicit

Table 4: Results. *Outdoor* and *Indoor* represent zero-shot evaluations. We outperform all SS-MDE baselines [56] (top block). In many cases, our model performs on par with supervised SotA (bottom block), without requiring ground-truth depth annotations for training.

	Train	Multi-task		In-Distribution				Outdoor						Indoor							
				Kitti		Mannequin		DDAD		DIODE		Sintel		SYNS		DIODE		NYUD-v2		TUM	
		Rank↓	Δ↑	Rel↓	F↑	Rel↓	F↑	Rel↓	F↑	Rel↓	δ _{.25} ↑	Rel↓	F↑	Rel↓	F↑	Rel↓	δ _{.25} ↑	Rel↓	δ _{.25} ↑	Rel↓	δ _{.25} ↑
Garg [17]	S	7.58	-38.52	<u>7.65</u>	<u>53.28</u>	27.63	<u>9.08</u>	26.93	7.80	39.60	44.15	39.41	<u>31.93</u>	26.05	15.17	19.18	70.54	22.49	59.60	23.53	62.82
Monodepth2 [20]	MS	7.74	-38.34	7.90	50.50	27.44	7.97	24.31	8.25	39.53	44.71	40.09	29.49	25.31	14.83	19.40	70.42	22.41	60.09	23.50	62.36
DiffNet [76]	MS	7.05	-36.84	7.98	49.60	27.46	7.76	<u>23.03</u>	9.43	<u>38.87</u>	<u>46.14</u>	39.93	28.77	25.09	14.64	19.11	70.94	21.82	<u>61.30</u>	23.21	63.08
HR-Depth [38]	MS	5.95	-35.16	7.70	51.49	<u>27.01</u>	8.39	23.13	<u>9.94</u>	39.09	45.60	<u>38.82</u>	30.90	<u>25.07</u>	<u>15.48</u>	<u>18.93</u>	<u>71.19</u>	<u>21.74</u>	61.18	<u>23.18</u>	<u>63.50</u>
KBR (Ours)	M	3.37	0.00	6.84	56.17	14.39	17.67	12.63	20.21	33.49	57.08	33.34	40.81	22.40	18.50	14.91	80.77	11.59	87.23	15.02	80.86
MiDaS [47]	D	4.89	-11.84	13.71	33.44	16.96	12.62	16.00	15.41	32.72	59.04	30.95	39.55	26.94	14.69	10.71	88.42	10.48	89.59	14.43	82.35
DPT-ViT [46]	D	3.32	-1.74	10.98	40.56	<u>15.52</u>	14.46	15.49	18.25	<u>32.59</u>	<u>59.82</u>	<u>25.53</u>	<u>43.57</u>	<u>23.24</u>	<u>17.44</u>	<u>9.60</u>	<u>91.38</u>	10.10	90.10	<u>12.68</u>	<u>86.25</u>
DPT-BEiT [46]	D	1.84	11.12	<u>9.45</u>	<u>44.22</u>	<u>13.55</u>	<u>16.58</u>	<u>10.70</u>	<u>22.63</u>	<u>31.08</u>	<u>61.51</u>	<u>21.38</u>	<u>46.46</u>	<u>21.47</u>	<u>17.73</u>	<u>7.89</u>	<u>93.34</u>	5.40	96.54	10.45	89.68
NeWCRFs [72]	D	<u>3.26</u>	<u>1.03</u>	<u>5.23</u>	<u>59.20</u>	18.20	<u>15.17</u>	<u>9.59</u>	<u>23.02</u>	37.01	49.66	39.25	32.43	24.28	16.76	14.05	84.95	<u>6.22</u>	<u>95.58</u>	14.63	82.95

Highlighted cells are NOT zero-shot results. S=Stereo, M=Monocular, D=Ground-truth Depth.

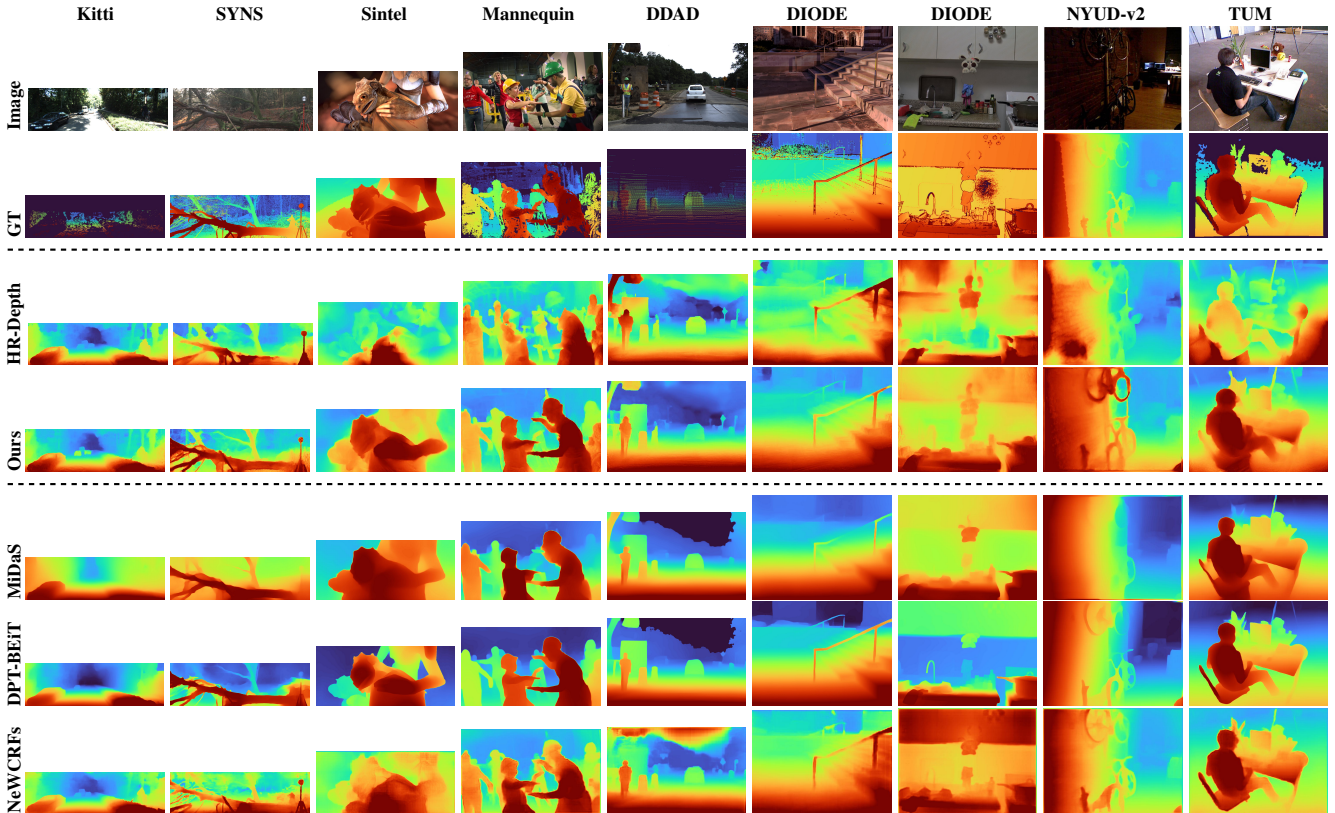


Figure 4: Zero-shot SS-MDE. The proposed model adapts to a wide range of datasets and environments. It greatly outperforms the updated self-supervised baselines from [56, 38] and performs on-par with SotA supervised baselines [47, 46, 73], whilst being more efficient. *Middle=Self-Supervised – Bottom=Supervised.*

motion masks to handle dynamic objects. Instead, we rely only on the minimum reconstruction loss and automasking [20]. Whilst this improves the robustness, it can be seen how dynamic objects such as cars can cause incorrect predictions (e.g. Kitti or DDAD). This represents one of the

most important avenues for future research. Further discussions regarding these failure cases and additional visualizations can be found in Appendix G.

MDEC-2. The Monocular Depth Estimation Challenge [54, 55] tested zero-shot generalization on SYNS-Patches. We

RESULTS										
#	User	Entries	Date of Last Entry	F-Score \uparrow	F-Score (Edges) \uparrow	MAE \uparrow	RMSE \uparrow	AbsRel \uparrow	Edge Accuracy \uparrow	Edge Completion \uparrow
1	jspenmar2	1	03/07/23	17.9528 (1)	9.5412 (2)	4.7760 (3)	8.2655 (2)	26.2810 (4)	3.5964 (18)	6.7604 (1)
2	KaiCheng	9	03/15/23	17.5085 (2)	8.8041 (6)	4.5159 (1)	8.7180 (5)	24.3172 (1)	3.2177 (6)	21.6526 (14)
3	xiangmochu	1	03/15/23	16.9423 (3)	9.6338 (1)	4.7103 (2)	7.9952 (1)	25.3457 (3)	3.5645 (14)	19.9490 (12)
4	cv_challenge	7	03/15/23	16.7012 (4)	9.3625 (3)	4.9111 (4)	8.6327 (4)	24.3265 (2)	3.0178 (1)	18.0730 (11)

Figure 5: MDEC-2 [55]. Our submission (*jspenmar2*) was top of the leaderboard in F-Score reconstruction. The challenge evaluated zero-shot performance on SYNS-Patches for both supervised and self-supervised approaches.

compare our model to all submissions from the latest edition (CVPR2023). As seen in Figure 5, our method (*jspenmar2*) achieves the highest F-Score reconstruction and is top-3 in all metrics except AbsRel and Edge-Accuracy. Once again, this illustrates the benefits of SlowTV, which contains large quantities of natural data not present in other datasets.

5.7. Map-Free Relocalization

Map-free relocalization is the task of localizing a target image using a single reference image. This is contrary to traditional pipelines, which require large image collections to first build a scene-specific map, such as SfM or training a CNN. Recent work [59, 2] has shown the benefit of incorporating metric MDE into feature matching pipelines to resolve the ambiguous scale of the predicted pose.

We evaluate all depth models on the MapFreeReloc benchmark [2] validation split, serving as an example real-world task. The feature-matching baseline [2] consists of LoFTR [58] correspondences, a PnP solver and DPT [46] fine-tuned on either Kitti or NYUD-v2. Since this benchmark requires metric depth but does not provide ground-truth, we align all models to the baseline fine-tuned DPT predictions using least-squares. We report the metrics provided by the benchmark authors. This includes translation (meters), rotation (deg) and reprojection (px) errors. Pose Precision/AUC were computed with an error threshold of 25 cm & 5°, while Reprojection uses a threshold of 90px.

As shown in Table 5, our method has the best performance across all SS-MDE approaches by a large margin. Our performance is on par with the supervised SotA, without requiring ground-truth supervision. This further demonstrates the benefits of the proposed SlowTV dataset and its applicability to real-world scenarios. Interestingly, we find that the original DPT models perform better than their fine-tuned counterparts, despite using these as the metric scale reference. This suggests that the fine-tuning procedure of [2] may provide metric scale at the cost of generality. However, this highlights the need for models that predict accurate metric depth, rather than only relative depth.

Table 5: Map-free Relocalization [2]. We incorporate KBR into a feature-matching pipeline for single-image relocalization. We once again outperform the SS-MDE baselines in every metric and perform on par with supervised SotA.

	Train	Pose				VCRE		
		Trans \downarrow	Rot \downarrow	P \uparrow	AUC \uparrow	Error \downarrow	P \uparrow	AUC \uparrow
Garg [17]	S	2.96	52.57	5.43	17.15	188.20	24.84	51.61
Monodepth2 [19]	MS	2.95	52.92	5.50	17.22	189.67	24.38	50.63
DiffNet [76]	MS	2.97	53.19	5.65	17.71	188.80	24.78	51.24
HR-Depth [38]	MS	2.94	52.95	5.67	17.95	187.83	25.06	51.52
KBR (Ours)	M	2.63	49.01	11.54	32.02	181.21	29.96	58.89
MiDaS [47]	D	2.60	46.92	11.39	30.44	180.64	30.45	59.72
DPT-ViT [46]	D	2.56	45.62	11.27	30.92	181.34	30.60	60.03
DPT-BEiT [46]	D	2.49	44.99	12.56	32.48	181.67	32.46	62.03
NeWCRFs [72]	D	2.89	51.92	6.69	20.77	184.63	25.89	52.93
DPT-NYUD [2]	D+FT	2.67	47.66	9.17	26.46	184.53	28.68	56.87
DPT-Kitti [2]	D+FT	2.66	49.21	10.86	29.99	178.49	28.37	56.86

Trans=meters, *Rot*=deg, *VCRE*=px, *Precision*=%, *AUC*=%.

6. Conclusion

This paper has presented the first approach to SS-MDE capable of generalizing across many datasets, including a wide range of indoor and outdoor environments. We demonstrated that our models significantly outperform existing self-supervised models, even in the automotive domain where they are currently trained. By leveraging the large quantity and variety of data in the new SlowTV dataset, we are able to close the gap between supervised and self-supervised performance. Additional components, such as the novel AR-Aug, randomized support frames and more flexible pose estimation, further improve the performance and zero-shot generalization of the proposed models.

Future work should explore alternative sources of data to incorporate even more scene variety. In particular, additional indoor data may significantly reduce the remaining gap between self-supervised and supervised approaches. Another key direction is improving the accuracy in dynamic scenes. A promising approach would be using optical flow to refine the estimated correspondences. This could be incorporated in a self-supervised manner, without requiring semantic segmentation or motion masks. However, it introduces additional costs due to the increased computational requirements from the new network.

Developing models capable of predicting metric depth would further increase their applicability to real-world applications. Finally, as the diversity of training environments increases, it will become crucial to further diversify the benchmarks used to evaluate these models.

Acknowledgements

This work was partially funded by the EPSRC under grant agreements EP/S016317/1 & EP/S035761/1.

References

- [1] Wendy J Adams, James H Elder, Erich W Graf, Julian Leyland, Arthur J Lugtigheid, and Alexander Murry. The Southampton-York Natural Scenes (SYNS) dataset: Statistics of surface attitude. *Scientific Reports*, 6(1):35805, 2016. [3](#)
- [2] Eduardo Arnold, Jamie Wynn, Sara Vicente, Guillermo Garcia-Hernando, Áron Monszpart, Victor Adrian Prisacariu, Daniyar Turmukhambetov, and Eric Brachmann. Map-free visual relocalization: Metric pose relative to a single image. In *ECCV*, 2022. [8](#)
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEit: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022. [5](#)
- [4] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. [2](#)
- [5] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Localbins: Improving depth estimation by learning local distributions. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 480–496. Springer, 2022. [2](#)
- [6] Jiawang Bian, Zhichao Li, Naiyan Wang, Huangying Zhan, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised Scale-consistent Depth and Ego-motion Learning from Monocular Video. In *Advances in Neural Information Processing Systems*, volume 32, 2019. [2](#)
- [7] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part VI 12*, pages 611–625. Springer, 2012. [3](#)
- [8] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. [3](#)
- [9] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8001–8008, 2019. [2](#), [4](#)
- [10] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [3](#)
- [11] Po-Yi Chen, Alexander H Liu, Yen-Cheng Liu, and Yu-Chiang Frank Wang. Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 2624–2632, 2019. [2](#)
- [12] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7063–7072, 2019. [4](#)
- [13] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [1](#), [2](#), [3](#), [6](#)
- [14] Qi Dai, Vaishakh Patil, Simon Hecker, Dengxin Dai, Luc Van Gool, and Konrad Schindler. Self-supervised object motion and depth estimation from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. [2](#), [4](#)
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [5](#)
- [16] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct Sparse Odometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3):611–625, 2018. [2](#)
- [17] Ravi Garg, Vijay Kumar, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. In *European Conference on Computer Vision*, pages 740–756, 2016. [2](#), [7](#), [8](#)
- [18] A Geiger, P Lenz, C Stiller, and R Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research*, 32(11):1231–1237, 2013. [1](#), [2](#), [3](#)
- [19] Clement Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation with Left-Right Consistency. *Conference on Computer Vision and Pattern Recognition*, pages 6602–6611, 2017. [2](#), [4](#), [5](#), [8](#)
- [20] Clement Godard, Oisín Mac Aodha, Michael Firman, and Gabriel Brostow. Digging Into Self-Supervised Monocular Depth Estimation. *International Conference on Computer Vision*, 2019-Octob:3827–3837, 2019. [2](#), [4](#), [5](#), [7](#)
- [21] Juan Luis Gonzalez Bello and Munchurl Kim. Forget About the LiDAR: Self-Supervised Depth Estimators with MED Probability Volumes. In *Advances in Neural Information Processing Systems*, volume 33, pages 12626–12637, 2020. [2](#)
- [22] Juan Luis Gonzalez Bello and Munchurl Kim. PLADE-Net: Towards Pixel-Level Accuracy for Self-Supervised Single-View Depth Estimation with Neural Positional Encoding and Distilled Matting Loss. In *Conference on Computer Vision and Pattern Recognition*, pages 6847–6856, 2021. [2](#)
- [23] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8977–8986, 2019. [2](#), [4](#)
- [24] Vitor Guizilini, Ambrus Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3D packing for self-

- supervised monocular depth estimation. *Conference on Computer Vision and Pattern Recognition*, pages 2482–2491, 2020. 1, 2, 3
- [25] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. *arXiv preprint arXiv:2002.12319*, 2020. 2
- [26] Xinyu Huang, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. The apolloscape open dataset for autonomous driving and its application. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2702–2719, 2019. 3
- [27] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial Transformer Networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015. 2, 4
- [28] Adrian Johnston and Gustavo Carneiro. Self-Supervised Monocular Trained Depth Estimation Using Self-Attention and Discrete Disparity Volume. In *Conference on Computer Vision and Pattern Recognition*, pages 4755–4764, 2020. 2
- [29] Hyunyoung Jung, Eunhyeok Park, and Sungjoo Yoo. Fine-grained semantics-aware representation enhancement for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12642–12652, 2021. 2
- [30] Maria Klodt and Andrea Vedaldi. Supervising the New with the Old: Learning SFM from SFM. In *European Conference on Computer Vision*, pages 713–728, 2018. 2
- [31] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Mannequin-challenge: Learning the depths of moving people by watching frozen people. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12):4229–4241, 2020. 2, 3
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [33] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 5
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [35] Chenxu Luo, Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia, and Alan Yuille. Every pixel counts ++: Joint learning of geometry and motion with 3d holistic understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(10):2624–2641, 2020. 2
- [36] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Trans. Graph.*, 39(4), aug 2020. 2
- [37] Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, and Liang Lin. Single View Stereo Matching. *Conference on Computer Vision and Pattern Recognition*, pages 155–163, 2018. 2
- [38] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu, Yong Liu, Xinxin Chen, and Yi Yuan. HR-Depth: High Resolution Self-Supervised Monocular Depth Estimation. *AAAI Conference on Artificial Intelligence*, 35(3):2294–2301, 2021. 2, 7, 8
- [39] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Un-supervised Learning of Depth and Ego-Motion from Monocular Video Using 3D Geometric Constraints. *Conference on Computer Vision and Pattern Recognition*, pages 5667–5675, 2018. 2
- [40] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. *Conference on Computer Vision and Pattern Recognition*, pages 4040–4048, 2016. 5
- [41] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 3
- [42] Evin Pinar Örnek, Shristi Mudgal, Johanna Wald, Yida Wang, Nassir Navab, and Federico Tombari. From 2D to 3D: Re-thinking Benchmarking of Monocular Depth Prediction. *arXiv preprint*, 2022. 5
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [44] Sudeep Pillai, Rareş Ambrus, and Adrien Gaidon. SuperDepth: Self-Supervised, Super-Resolved Monocular Depth Estimation. In *International Conference on Robotics and Automation*, pages 9250–9256, 2019. 2
- [45] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the Uncertainty of Self-Supervised Monocular Depth Estimation. In *Conference on Computer Vision and Pattern Recognition*, pages 3224–3234, 2020. 2
- [46] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, October 2021. 1, 2, 5, 7, 8
- [47] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 1, 2, 5, 7, 8
- [48] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael 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*, pages 12240–12249, 2019. 2
- [49] Rui, Stückler Jörg, Cremers Daniel Yang Nan, and Wang. Deep Virtual Stereo Odometry: Leveraging Deep Depth Prediction for Monocular Direct Sparse Odometry. In *European Conference on Computer Vision*, pages 835–852, 2018. 2
- [50] Ashutosh Saxena, Min Sun, and Andrew Y. Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence*, 31(5):824–840, 2009. 6
- [51] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [52] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. Feature-Metric Loss for Self-supervised Learning of Depth and Egomotion. In *European Conference on Computer Vision*, pages 572–588, 2020. 2
- [53] Jaime Spencer, Richard Bowden, and Simon Hadfield. DeFeat-Net: General monocular depth via simultaneous unsupervised representation learning. In *Conference on Computer Vision and Pattern Recognition*, pages 14390–14401, 2020. 2
- [54] Jaime Spencer, C Stella Qian, Chris Russell, Simon Hadfield, Erich Graf, Wendy Adams, Andrew J Schofield, James H Elder, Richard Bowden, Heng Cong, et al. The monocular depth estimation challenge. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 623–632, 2023. 2, 7
- [55] Jaime Spencer, C. Stella Qian, Michaela Trescakova, Chris Russell, Simon Hadfield, Erich Graf, Wendy Adams, Andrew J. Schofield, James Elder, Richard Bowden, and Others. The second monocular depth estimation challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2023. 7, 8
- [56] Jaime Spencer, Chris Russell, Simon Hadfield, and Richard Bowden. Deconstructing self-supervised monocular reconstruction: The design decisions that matter. *Transactions on Machine Learning Research*, 2022. Reproducibility Certification. 1, 2, 3, 5, 6, 7
- [57] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012. 3
- [58] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 8
- [59] Carl Toft, Daniyar Turmukhambetov, Torsten Sattler, Fredrik Kahl, and Gabriel J Brostow. Single-image depth prediction makes feature matching easier. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 473–492. Springer, 2020. 8
- [60] Fabio Tosi, Filippo Aleotti, Matteo Poggi, and Stefano Mattoccia. Learning monocular depth estimation infusing traditional stereo knowledge. *Conference on Computer Vision and Pattern Recognition*, 2019-June:9791–9801, 2019. 2
- [61] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity Invariant CNNs. *International Conference on 3D Vision*, pages 11–20, 2018. 3
- [62] Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z Dai, Andrea F Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R Walter, et al. Diode: A dense indoor and outdoor depth dataset. *arXiv preprint arXiv:1908.00463*, 2019. 3
- [63] Chaoyang Wang, Jose Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning Depth from Monocular Videos Using Direct Methods. *Conference on Computer Vision and Pattern Recognition*, pages 2022–2030, 2018. 2
- [64] Zhou Wang, A C Bovik, H R Sheikh, and E P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 2
- [65] Jamie Watson, Michael Firman, Gabriel Brostow, and Daniyar Turmukhambetov. Self-supervised monocular depth hints. *International Conference on Computer Vision*, 2019-Octob:2162–2171, 2019. 2
- [66] Ross Wightman. PyTorch Image Models, 2019. 5
- [67] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *Conference on Computer Vision and Pattern Recognition*, 2017-Janua:5987–5995, 2017. 5
- [68] Jiaying Yan, Hong Zhao, Penghui Bu, and YuSheng Jin. Channel-Wise Attention-Based Network for Self-Supervised Monocular Depth Estimation. In *International Conference on 3D Vision*, pages 464–473, 2021. 2
- [69] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3VO: Deep Depth, Deep Pose and Deep Uncertainty for Monocular Visual Odometry. In *Conference on Computer Vision and Pattern Recognition*, pages 1278–1289, 2020. 2
- [70] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1983–1992, 2018. 2
- [71] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. 3
- [72] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Neural window fully-connected crfs for monocular depth estimation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3906–3915, 2022. 1, 2, 5, 7, 8
- [73] Weihao Yuan, Yazhan Zhang, Bingkun Wu, Siyu Zhu, Ping Tan, Michael Yu Wang, and Qifeng Chen. Stereo matching by self-supervision of multiscopic vision. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5702–5709. IEEE, 2021. 5, 7
- [74] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian M. Reid. Unsupervised Learning of Monocular Depth Estimation and Visual Odometry with Deep Feature Reconstruction. *Conference on Computer Vision and Pattern Recognition*, pages 340–349, 2018. 2
- [75] Chaoqiang Zhao, Youmin Zhang, Matteo Poggi, Fabio Tosi, Xianda Guo, Zheng Zhu, Guan Huang, Yang Tang, and Stefano Mattoccia. Monovit: Self-supervised monocular depth estimation with a vision transformer. *International Conference on 3D Vision*, 2022. 2

- [76] Hang Zhou, David Greenwood, and Sarah Taylor. Self-Supervised Monocular Depth Estimation with Internal Feature Fusion. In *British Machine Vision Conference*, 2021. [2](#), [7](#), [8](#)
- [77] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised Learning of Depth and Ego-Motion from Video. *Conference on Computer Vision and Pattern Recognition*, pages 6612–6619, 2017. [2](#), [3](#), [5](#)