

Self-supervised Pretraining and Finetuning for Monocular Depth and Visual Odometry

Leonid Antsfeld¹ and Boris Chidlovskii¹

Abstract—For the task of simultaneous monocular depth and visual odometry estimation, we propose learning self-supervised transformer-based models in two steps. Our first step consists in a generic pretraining to learn 3D geometry, using cross-view completion objective (CroCo), followed by self-supervised finetuning on non-annotated videos. We show that our self-supervised models can reach state-of-the-art performance ‘without bells and whistles’ using standard components such as visual transformers, dense prediction transformers and adapters. We demonstrate the effectiveness of our proposed method by running evaluations on six benchmark datasets, both static and dynamic, indoor and outdoor, with synthetic and real images. For all datasets, our method outperforms state-of-the-art methods, in particular for depth prediction task.

I. INTRODUCTION

Visual odometry and depth estimation are critical for understanding the scene geometry and camera motion for tasks in robotics and autonomous driving. Supervised learning methods have been applied in many deep neural network frameworks and demonstrated outstanding results in both visual odometry and depth estimation [6], [24], [47]. However, supervised learning requires a significant amount of labeled data for training, thus the research interest moved towards the exploration of unsupervised learning frameworks that can learn scene depth and visual odometry (DVO) simultaneously (see Figure 1) without any ground truth annotations, using geometric constraints between adjacent frames.

Improving the network backbone of depth networks is a particularly effective way to gain in accuracy [12], [50]. Multiple sophisticated convolution-based backbones such as ResNet [12], PackNet [13], HRNet [49], HRDepth [21], CADepth [43] have contributed to advances in the self-supervised monocular depth estimation task. Vision transformers (ViTs) have been deployed in supervised depth estimation as well [19], [29], they were recently combined with convolutions in self-supervised monocular depth estimation [48].

In this paper we argue that sophisticated backbone architectures are not crucial to the performance of DVO models, and they can be replaced with standard designs, such as generic transformers, thus benefiting from architecture unification and best training recipes, while retaining crucial properties of scalability, robustness and efficient transfer. Our approach is inspired by recent advances in self-supervised pre-training to geometric vision tasks [14], [39]. In particular, [39] proposed the pretext task of *cross-view completion* (CroCo), a variant of masked image modelling (MiM) where a partially masked input image is reconstructed given visible patches and an




Fig. 1. Simultaneous depth and visual odometry estimation from a video, with two steps of model pretraining and finetuning.

additional view of the same scene. This pretraining objective is well suited to geometric downstream tasks such as optical flow and stereo matching [40]. The CroCo architecture consists of a ViT encoder to extract features for the non-masked tokens of the first image, as well as for the second image, and a ViT decoder to reconstruct the masked image.

We argue that CroCo is better suited for the DVO task than the existing backbones [12], [21], [28], [35] which commonly use ImageNet-based pretrained models. Due to the object-centric [25] and the balanced [1] nature of ImageNet [30], these models provide semantically rich features and excel on classification tasks, but can fail to capture structural information critical for dense geometric tasks like depth estimation and visual odometry. Instead, used as pretraining task, CroCo can easily reach state-of-the-art results without sophisticated network architecture, by using vanilla ViT blocks and standard geometric constraints on image reconstruction.

In this paper, we show how to efficiently finetune the CroCo pre-trained model on monocular depth and visual odometry, in self-supervised way. This is, to our best knowledge, the first task where *model pre-training and fine-tuning are both self-supervised*, with supervision signals coming from image completion at pretraining step, and from the geometric constraints at finetuning step.

Driven by the concept of foundation models, a large scale pretraining model is trained on large heterogeneous datasets and targets a multitude of downstream tasks [46], [39]. Such an approach represents an efficient alternative to curating a large dataset for one particular task as suggested in [29], [33]. In the case of DVO, we can benefit from the generic pretrained models oriented towards understanding 3D geometry of a scene, and finetune them on non-annotated videos to reach optimal downstream task performance.

We reshape the pretrained model in a way that allows for self-supervised finetuning by sharing one encoder and targeting depth and visual odometry tasks with a proper decoder, by optimizing geometric consistency losses.

We claim the following contributions:

¹Authors are with Naver Labs Europe, Meylan, France. `firstname.lastname@naverlabs.com`

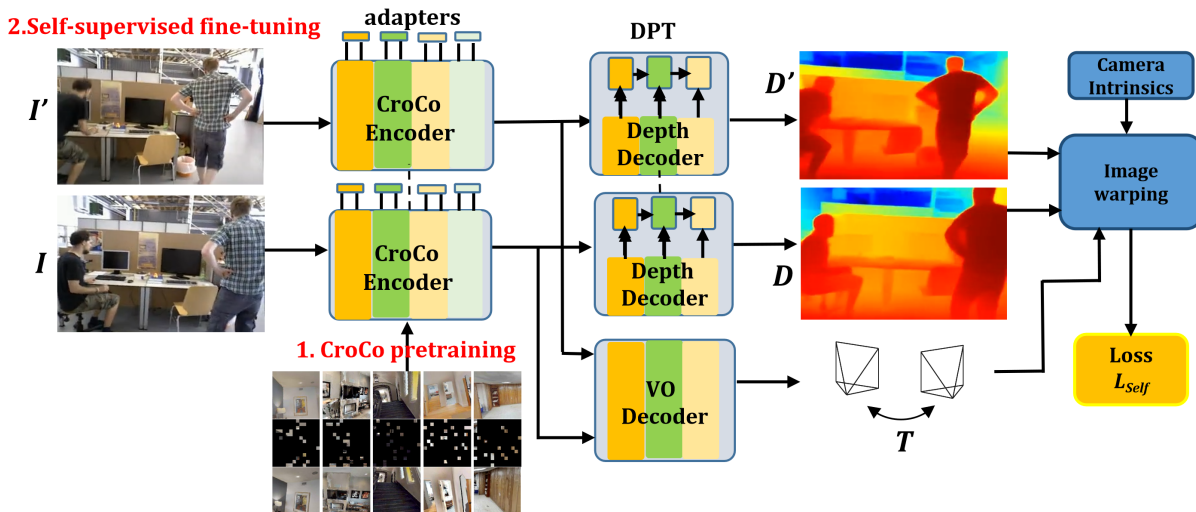


Fig. 2. **CroCo-DVO architecture.** 1. Cross-view completion task is pretrained from a large amount of heterogeneous data. 2. At the finetuning step, given two consecutive frames (I, I'), we estimate their depth maps (D, D') and relative pose T using the network. Then we compute the self-supervised loss L_{Self} which is composed of pixel-wise depth inconsistency between D and D' , geometric consistency loss and a self-discovered mask to handle dynamic objects (see Section III-B).

- 1) We propose a new approach for simultaneous depth and visual odometry estimation, where both pretraining and finetuning steps are self-supervised, with the learning signal coming from image cross completion and constraints on geometric consistency between adjacent frames.
- 2) We reshape the CroCo backbone for self-supervised finetuning to deal with the DVO tasks, and show how to enrich the backbone with recent (and already) standard extensions, such as adapters [8] for efficient transfer, and DPT for depth estimation [28].
- 3) The effectiveness of our proposed method is demonstrated through experiments on six benchmark datasets (NYUv2, KITTI, Tum, Bonn, Ddad and Gibson). We show that the model outperforms the-state-of-the-art methods on indoor and outdoor, static and dynamic, real and synthetic scenes.

II. PRIOR WORK

Self-supervised learning — Self-supervised learning is a way to learn generalizable visual representations often using different pretext tasks for pre-training [17]. Inspired by BERT [9] in NLP, different masked modeling methods (MIM) have been adapted to computer vision and robot perception. MIM pre-training aims at reconstructing masked information from an input image either in the pixel [15], [42] or feature space [2], [38]. Overall, MIM models perform well on denser tasks such as object detection [15] and have been applied to robotic perception [26], [22] when pre-trained on related datasets. Recently, CroCo [39] introduced the pretext task of asymmetric cross-view completion, where a second (non-masked) view of the same scene is added to MIM. This is well suited to geometric downstream tasks: to leverage the second view and improve reconstruction accuracy, the model has to implicitly be aware of the geometry of the scene.

Unsupervised depth and visual odometry — SfM-Learner [51] was first to replace the known stereo baseline with an additional network to regress VO between consecutive frames and to supervise the training only by a stream of monocular images. Later, the explainability mask from SfM-Learner was replaced with uncertainty [18] allowing the network to ignore incorrect predictions. DDVO [36] introduced a differentiable DSO module to refine the VO network prediction. Monodepth2 [12] introduced minimum reconstruction filtering to reduce occlusion artifacts, alongside static pixel automasking via the raw reconstruction loss. D3VO [44] additionally predicted affine brightness transformation parameters for each support frame. SC-SfM-Learners [5] and Sc-Depth-v3 [35] proposed an end-to-end differentiable geometric consistency constraint by synthesizing the support depth view. They included a variant of the absolute relative loss additionally used as auto-masking for the reconstruction loss. Alternatively, the depth models can be distilled [20] and depth values can be discretized, by implicit partitioning the scene into a set of concepts in IDisc [24], or by capturing an input image as a set of orthogonal vertical planes and ground planes in PlaneDepth [37].

Network architectures — Multiple architectures used as backbone showed a significant impact on the performance of monocular depth estimation. First, GeoNet [45] replaced the VGG encoder with a ResNet. PackNet [13] counts on 3D convolutions to learn detail-preserving compression and decompression of features. HRDepth [21] worked on the features decoding by implementing an attention module for multi-scale feature fusion. CADepth [43] extracts the long-range relationships between features via a channel-wise attention-based network to aggregate discriminated features in channel dimensions. Several recent works tackled supervised monocular depth estimation by using Transformer architectures [19], [29]. ViTs were also combined with convolutions in self-supervised monocular depth estimation [48].

III. ARCHITECTURE

A. Cross-view completion

Cross-View Completion (CroCo) [39] is a generic pretraining task trained on large heterogeneous dataset and which enables a network to perceive low-level geometric cues highly relevant to vision downstream tasks. It is an extension of masked image modeling [15] processing pairs of images (I, I') , which correspond to two different views of the same scene with important overlap. In more details, the images are split into non-overlapping patches. The first input image I is partially masked, the masked patches are discarded, and the remaining ones are fed to an image encoder, which is implemented using a ViT backbone. All patches from the second image are encoded using the same encoder with shared weights. The latent token representations output by the encoder from both images are then fed to a decoder whose goal is to predict the appearance of masked patches. It uses a series of transformer decoder blocks comprising cross-attention layers. This allows non-masked tokens from the first image to attend tokens from the reference image, thus enabling cross-view comparison and reasoning. The model is trained using a pixel reconstruction loss over all masked patches, similar to MAE [3]. The CroCo pretext task has been shown to be applicable to different downstream vision problems; competitive performances were shown for the tasks of supervised monocular depth estimation [39], optical flow and stereo matching [40].

CroCo-DVO — We now adapt CroCo pretrained model to self-supervised finetuning for DVO tasks. Our architecture includes two branches, for depth and visual odometry, which share the same encoder (see Figure 2). When finetuning the CroCo model, a pair of images (I, I') are fed to the encoder, then two decoders process the tokens of both images. For depth, to output a pixel-wise prediction, we initially rely on ViT layers with a linear MLP head per token, in order to output a dense depth map. The depth decoder consists in a series of transformer decoder blocks (self-attention among token features from the first frame, cross-attention with the token features from the second frame, and an MLP). Finally, the second, VO decoder outputs the 6-DoF relative pose, it is designed as a lightweight MLP after patch-wise concatenation.

DPT — As a powerful alternative to ViT depth decoder for dense tasks prediction, we propose to use the Dense Prediction Transformer (DPT) [28], which adapts the standard up convolutions and fusions from multiple layers to vision transformers. This allows to combine features from different blocks by reshaping them to different resolutions and fusing them with convolutional layers. In practice, the DPT uses the features from four blocks, regularly spread over the decoder depth, starting from the last block. In total, using transformers in DPT allows to make more detailed and globally consistent predictions [28].

Adapters — We also extend the CroCo-DVO backbone with adapters [8] which were designed for efficiently transfer large pretrained ViT-based models to downstream tasks. AdaptFormer attains strong transfer learning abilities by only

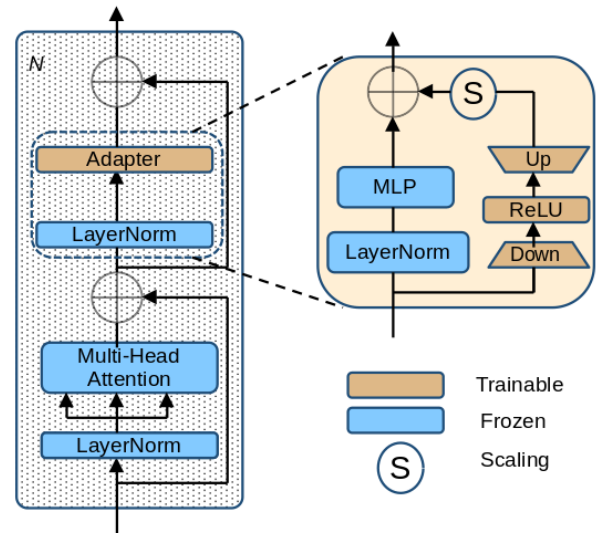


Fig. 3. **Adapters:** Training adapters while freezing the main backbone. AdaptFormer [8] replaces the MLP block in the transformer encoder with AdaptMLP, which is consisted of two sub-branches. The MLP layer in the left branch, identical to the original network, is frozen, right branch introduce a lightweight module for task-specific finetuning.

finetuning a small number of extra parameters. The adapter architecture is shown in Figure 3. Compared to the full finetuning, AdaptFormer keeps the model parameters frozen and trains only additional residual network for each layer. In detail, they replace the MLP block in the transformer encoder with AdaptMLP, which consists of two sub-branches. The MLP layer in the left branch is identical to the original network, while the right branch is an additionally introduced lightweight module for task-specific finetuning. Specifically, the right branch is designed to be a bottleneck structure for limiting the number of parameters purpose, which includes a down-projection layer with parameters $W_{down} \in R^{d \times \hat{d}}$, an up-projection layer with parameters $W_{up} \in R^{\hat{d} \times d}$, where \hat{d} is the bottleneck middle dimension and satisfies $\hat{d} \ll d$. In CroCo backbone, d is 768 for encoder and 512 for decoder, and value of \hat{d} is much smaller, usually 16 or 32. The newly added parameters (right branch in Figure 3) are updated on the finetuning data with the self-supervised losses described in the following section.

B. Self-supervised fine-tuning

During the self-supervised finetuning step, depth estimation and visual odometry branches are simultaneously finetuned on a number of monocular videos. Given a pair of consecutive frames (I, I') from a training video, we predict their depths D, D' by using the depth decoder and estimate their relative 6-DoF camera pose T by using the odometry decoder. With the obtained depth and relative camera pose, we follow [4], [35] in computing the warped \hat{D}' by transforming D to 3D space (with known intrinsics) and projecting it to I' using pose T . The inconsistency between warped \hat{D}' and predicted D' is used as geometric consistency loss L_g to supervise the training of the network. The warping between two images using the predicted depth and pose is followed

by synthesizing \hat{I}' by warping I' via bi-linear interpolation. We penalize the color inconsistencies between I and \hat{I}' , and we constrain the geometry consistency between D and D' , which backpropagates the gradients to the network.

The first, *geometry consistency* loss L_g aims to force the predicted depths D and D' to be consistent with each other in 3D space,

$$L_g = \frac{1}{|\mathcal{V}|} \sum_{p \in \mathcal{V}} D_{diff}(p), \quad (1)$$

where \mathcal{V} is a set of valid points successfully projected from I to I' , D_{diff} stands for the pixel-wise depth inconsistency between D and D' , $D_{diff}(p) = \frac{|D'(p) - D(p)|}{D'(p) + D(p)}$. To mitigate the influence of moving objects and occlusions that create geometrical inconsistency across multiple views, we follow [4] in using the self-discovered mask $m_s = 1 - D_{diff}$ which assigns lower weights to dynamic objects and occlusions.

The second, *photometric* loss L_p constrains the warping between I and I' generated by depth D and pose T ,

$$L_p = \frac{1}{|\mathcal{V}|} \sum_{p \in \mathcal{V}} (1 - \lambda) \|I(p) - \hat{I}'(p)\|_1 + \lambda \frac{1 - SSIM_T(p)}{2}, \quad (2)$$

where \hat{I}' is synthesized from I by warping, and SSIM is a widely-used metric to measure image structural similarity, λ is 0.85 by following [12]. Actually, instead of loss L_p in Eq.(2), we use its weighted version L_p^m that counts on the mask m_s to down-weight regions of moving objects and obstacles,

$$L_p^m = \frac{1}{|\mathcal{V}|} \sum_{p \in \mathcal{V}} (m_s(p) \cdot L_p(p)), \quad (3)$$

The third, *edge-aware smoothness* loss regularizes the predicted depth map:

$$L_s = \sum_p (e^{-\Delta I(p)} \cdot \Delta D(p))^2, \quad (4)$$

where Δ is the first derivative along spatial directions, which guides smoothness by image edges.

The overall objective function is a weighted sum of the three loss terms presented above

$$L_{Self} = L_p^m + \beta L_g + \gamma L_s. \quad (5)$$

In practice, we use $\beta = 0.5$ and $\gamma = 0.1$.

C. Implementation details

We build on the publicly available CroCo code ¹. The CroCo-DVO encoder and decoder account for 12 and 8 ViT blocks, respectively. Four DPT layers are pulled from blocks 2, 4, 6, 8 of the CroCo decoder. In both cases, the pixel-wise generation is terminated with 128 heads. The VO decoder accounts for 2-layer MLP for the 6-DoF pose estimation. We add adapters (of dimension $d' = 32$) to both encoder and decoder layers.

IV. EXPERIMENTAL RESULTS

A. Datasets and Metrics

We intensively evaluate our approach on a set of six benchmark datasets, which include indoor and outdoor, static and dynamic scenes, synthetically generated and real images. Outdoor datasets include KITTI [11] and Ddad [13]; Indoor datasets include NYUv2 [32], Bonn [23], Tum [34] and

Gibson [31] datasets. KITTI, NYUv2 and Gibson are mostly static datasets, while Ddad, Tum and Bonn datasets contain people or other dynamic objects. Gibson dataset includes synthetically generated images, others are the real image datasets. All the self-supervised methods are finetuned and tested on each dataset individually for a fair comparison with the state-of-the-art methods.

KITTI. The dataset provides driving videos in urban scenes; it is the most widely-used dataset in self-supervised monocular depth estimation problems. We use Eigen's split [10] with 697 images for testing and the remaining video sequences for training. Depth ranges are capped at 80 meters, and images are resized to the resolution of 832×256 for training networks.

NYUv2. The dataset is widely-used in the computer vision community, it provides a large collection of indoor videos. There are 654 testing images of static scenes for depth evaluation, the remaining videos with 28,565 images are used for training. Images are resized to the resolution of 320×256 .

Ddad. The dataset contains 200 driving videos captured in urban scenes. The standard split includes 150 training scenes (12,650 images) and 50 validation scenes (3,950 images) which are used for evaluation. Depth ranges are capped to at most 200 meters, and images are resized to the resolution of 640×384 for training the network.

Bonn. The dataset contains 26 dynamic indoor videos. 4 video sequences with fast-moving people (1785 images) are used for testing, and with the remaining videos for training. Depth ranges are capped at 10 meters, and images are resized to the resolution of 320×256 .

Tum. This collection consists of 11 sequences. Two sequences that contain moving people (1375 images) are used for testing. The remaining 9 dynamic videos are used for training, and images are resized to the the image resolution of 320×256 .

Gibson. The Gibson dataset [41] is extensively used in indoor navigation tasks [7], [27], where the navigation goal is given by a point on a map, an image or an instance of object class and trajectories are traced by an expert. We use the Gibson train split (72 scenes) for finetuning, and Gibson-val (14 scenes) to evaluate performance; image resolution is set to 320×256 .

We start from the pretrained CroCo model and finetune it for 50 epochs for NYUv2, Tum and Bonn datasets using AdamW optimizer and batch size of 12, 8, 8, respectively; the number of epochs is 25 for KITTI, Gibson and Ddad datasets, all with batch size of 4. The initial learning rate is 10^{-4} for all datasets except for KITTI where the initial learning rate is 10^{-5} .

Evaluation Metrics — We use standard depth evaluation metrics [5], [51] that include a) mean absolute relative error (AbsRel), b) root mean squared error (RMSE), and c)-e) the accuracy under threshold ($\delta_i < 1.25^i, i = 1, 2, 3$). Besides, we follow [5], [35] in rescaling the predicted depth maps to match the median of the ground truth for evaluation, $s = \text{median}(D_{gt}) / \text{median}(D_{pred})$.

¹<https://github.com/naver/croco>.

Method	AbsRel ↓	RMSE ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
PackNet[13]	0.109	4.696	0.884	0.961	0.981
SC-SfM-Learners[4]	0.118	4.803	0.866	0.958	0.981
Sc-Depth-v3[35]	0.118	4.709	0.864	0.960	0.984
CADepth-Net[43]	0.105	4.535	0.892	0.964	0.983
DiffNet[49]	0.102	4.483	0.896	0.965	0.983
MonoViT[48]	0.099	4.372	0.900	0.967	0.984
CroCo-DVO	0.106	4.346	0.882	0.966	0.986
CroCo-DVO+Ad	0.104	4.314	0.886	0.967	0.986
CroCo-DVO+DPT	0.098	4.210	0.900	0.969	0.986

TABLE I

SELF-SUPERVISED DEPTH ESTIMATION ON KITTI DATASET.

Method	AbsRel ↓	RMSE ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [12]	0.169	0.614	0.745	0.946	0.987
MonoIndoor [16]	0.134	0.526	0.823	0.958	0.989
SC-SfM-Learners [4]	0.138	0.532	0.796	0.948	0.986
Sc-Depth-v3 [35]	0.123	0.486	0.848	0.963	0.991
CroCo-DVO	0.106	0.410	0.885	0.970	0.994
CroCo-DVO+Ad	0.101	0.398	0.892	0.976	0.995
CroCo-DVO+DPT+Ad	0.095	0.377	0.905	0.980	0.995

TABLE II

SELF-SUPERVISED DEPTH ESTIMATION ON NYUV2 DATASET.

B. Evaluation results

In this section we evaluate various configurations of CroCo-DVO and extensions for the finetuning step. First, **CroCo-DVO** is our baseline, it reuses and finetunes CroCo encoder and depth decoder with no change. Second, **CroCo-DVO+Ad** adds and trains adapters in both encoder and decoder layers by freezing the large part of the network parameters. Third, **CroCo-DVO+DPT** extends the baseline’s depth decoder with four DPT layers for the accurate depth prediction. Lastly, **CroCo-DVO+DPT+Ad** combines the two extensions, adapters and DPT, in one network.

We use six datasets presented in the previous section to evaluate our models; the quantitative depth estimation results are reported in Tables I-VI. For all datasets except Gibson, we compare our method with the state-of-the-art methods, including Monodepth2 [12], PackNet [13], CADepth-Net[43], DiffNet[49], SC-SfM-Learners [4], Sc-Depth-v3 [35] and MonoViT [48]. In each table we report the performance of our baseline and two other CroCo-DVO configurations. As results show, our method outperforms previous methods by a large margin in in four or all five metrics, thus demonstrating the efficacy of our CroCo-DVO architecture.

Self-supervised finetuning the CroCo backbone on DVO tasks already improves the state-of-the-art on two datasets, Bonn and NYUv2. It also shows competitive results on KITTI, Tum and Ddad datasets. Adapters and DPT plays a key role in the performance and improves the state-of-the-art. Adding extensions to the CroCo backbone reports the best results

Method	AbsRel ↓	RMSE ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [12]	0.312	1.408	0.474	0.793	0.905
SC-SfM-Learners [4]	0.223	0.282	0.643	0.862	0.932
Sc-depth-v3 [35]	0.163	0.265	0.797	0.882	0.937
CroCo-DVO	0.180	0.259	0.773	0.906	0.972
CroCo-DVO+Ad	0.182	0.252	0.752	0.903	0.978
CroCo-DVO+DPT+Ad	0.180	0.251	0.756	0.905	0.978

TABLE III

SELF-SUPERVISED DEPTH ESTIMATION ON TUM DATASET.

Method	AbsRel ↓	RMSE ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
MonoDepth2 [12]	0.565	2.337	0.352	0.591	0.728
SC-SfM-Learners [4]	0.211	0.619	0.714	0.873	0.936
Sc-depth-v3 [35]	0.126	0.379	0.889	0.961	0.980
CroCo-DVO	0.117	0.366	0.882	0.971	0.987
CroCo-DVO+Ad	0.115	0.356	0.886	0.972	0.989
CroCo-DVO+DPT+Ad	0.113	0.343	0.901	0.973	0.988

TABLE IV

SELF-SUPERVISED DEPTH ESTIMATION ON BONN DATASET.

Method	AbsRel ↓	RMSE ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [12]	0.239	18.392	0.752	0.899	0.949
SC-SfM-Learners [35]	0.169	16.290	0.773	0.905	0.961
Sc-depth-v3 [35]	0.142	15.868	0.813	0.922	0.963
CroCo-DVO	0.156	14.509	0.789	0.929	0.971
CroCo-DVO+Ad	0.151	14.151	0.815	0.935	0.973
CroCo-DVO+DPT	0.142	13.323	0.836	0.946	0.979

TABLE V

SELF-SUPERVISED DEPTH ESTIMATION ON Ddad DATASET.

by all metric except AbsRel on Tum and Ddad datasets. Our results is particularly important for KITTI and NYUv2 datasets, which constitute the most difficult and competitive benchmarks.

The analysis shows that DPT systematically improves the depth prediction, in particular it makes blurry regions sharper. The advantage of using adapters is more subtle. It is very consistent on NYUv2, Tum, and Bonn datasets while yielding more limited benefits on Kitti and Ddad datasets. We hypothesize that outdoor driving scenes which compose these two datasets are not well represented in the large-scale dataset the CroCo backbone was pretrained on. Training a small adapter network is insufficient to capture the specific street geometry. As result, CroCo-DVO-DPT that finetunes the entire network performs better.

C. CroCo-DVO ablations

We now ablate components we described in Section III to design the CroCo-DVO architecture. First, in Table VII we compare the CroCo-DVO baseline to CroCo-DVO+Ad with the frozen backbone and adapters dimension between 8 and 128, while finetuning on Bonn dataset. The table reports the optimal values for AbsRel, RMSE and δ_3 metrics when the adapter dimension is between 16 and 64. In dataset evaluations, the dimension of adapters is fixed to 32.

Second, we ablate more configurations of the CroCo backbone extended with adapters and DPT. Adapters are used in all layers by freezing the main backbone and training the small adapter network. Adding adapters increases the number of parameters by 2% but reduces the finetuning time by 65%. DPT layers, as presented in Section 2, are pulled from the depth decoder. In addition, we test a configuration where DPT is branched to the encoder, by pulling from blocks 3, 6, 9, 12, and skipping the decoder. Table VIII presents six different configurations and reports their performance on Bonn dataset. As the table shows, contributions brought by the DPT and adapters are complementary. Used jointly, adapters and DPT with the depth decoder bring the most important improvement to the CroCo-DVO baseline.




Fig. 4. Qualitative depth estimation for the state-of-the-art and our methods.

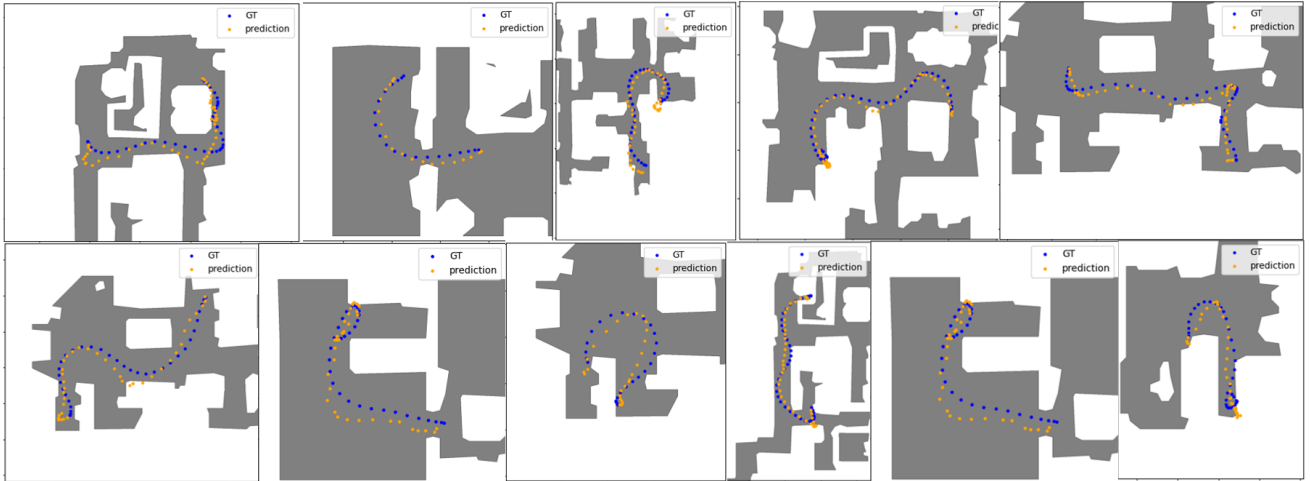


Fig. 5. Visual odometry estimation: Estimated Gibson trajectories (orange) vs GT trajectories (blue).

Method	AbsRel↓	RMSE↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
CroCo-DVO	0.144	0.307	0.846	0.952	0.983
CroCo-DVO+DPT	0.125	0.246	0.870	0.971	0.990
CroCo-DVO+DPT+Ad	0.123	0.250	0.880	0.961	0.985

TABLE VI

SELF-SUPERVISED DEPTH ESTIMATION ON GIBSON DATASET.

Metrics	No ads	dim=8	dim=16	dim=32	dim=64	dim=128
δ_3 ↑	0.987	0.988	0.989	0.989	0.988	0.983
AbsRel↓	0.117	0.120	0.116	0.116	0.115	0.130
RMSE↓	0.366	0.359	0.358	0.356	0.356	0.374

TABLE VII

ABLATION OF THE ADAPTER DIMENSIONALITY ON BONN DATASET.

D. Visual Odometry Evaluation

Beyond the depth estimation, we use CroCo-DVO to estimate the visual odometry on the Gibson dataset, where the ground truth 6-DoF poses are available for all trajectories in the validation set. This allows us to evaluate the translation and rotation errors and visually compare the estimated and ground truth trajectories. Taking into account that robots navigate on 2D-plane (which is not the case of other datasets) we project predicted trajectories on the xy-plane, and use the mean square error (MAE) to estimate the mean translation error (in meters) and the mean angular error is (in degrees) for all trajectories in the validation set (see Table IX). As the table shows, better depth estimation with extended CroCo-DVO versions is not necessarily translated in better odometry estimation. All three versions of CroCo-DVO report close values of translation (0.178-0.183 m) and rotation (3.65-3.71°) errors. Figure 5 completes this section; it proposes qualitative comparison between predictions and ground truth for a subset

Method	Ads	DPT	AbsRel↓	RMSE↓	δ_1	δ_2	δ_3
DVO			0.117	0.366	0.882	0.971	0.987
DVO+Ad	✓		0.115	0.356	0.886	0.972	0.989
DVO-DPT(E)		Enc	0.117	0.362	0.886	0.972	0.987
DVO+DPT		Dec	0.115	0.353	0.888	0.971	0.987
DVO+DPT(E)+Ad	✓	Enc	0.115	0.357	0.886	0.971	0.988
DVO+DPT+Ad	✓	Dec	0.113	0.343	0.901	0.973	0.988

TABLE VIII

SIX CROCO-DVO CONFIGURATIONS ON BONN DATASET.

Method	Translation Error(m)	Rotation Error(°)
CroCo-DVO	0.183	3.65
CroCo-DVO+Ad	0.181	3.71
CroCo-DVO+DPT+Ad	0.178	3.68

TABLE IX

TRANSLATION AND ROTATION ERRORS FOR GIBSON DATASET.

of trajectories from Gibson validation set.

V. CONCLUSION

We proposed a ViT-based model for the self-supervised estimation of monocular depth and visual odometry. It benefits from the generic pretraining oriented towards understanding 3D geometry of any scene followed by self-supervised finetuning on non-annotated streams of images. We show that the model, standalone or completed with standard extensions like adapters or DPT, can easily reach the state-of-the-art performance 'without bells and whistles'. The effectiveness of our proposed method is demonstrated through experiments on six benchmark datasets, where our method outperforms the state-of-the-art methods in the depth estimation, both indoor and outdoor, static and dynamic, real and synthetic image scenes.

REFERENCES

- [1] Mahmoud Assran, Randall Balestriero, Quentin Duval, Florian Bordes, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, and Nicolas Ballas. The hidden uniform cluster prior in self-supervised learning. In *ICLR*, 2023.
- [2] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, 2022.
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, pages 456–473, 2022.
- [4] Jia-Wang Bian, Huangying Zhan, Naiyan Wang, Tat-Jun Chin, Chunhua Shen, and Ian D. Reid. Auto-rectify network for unsupervised indoor depth estimation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(12):9802–9813, 2022.
- [5] Jia-Wang Bian, Huangying Zhan, Naiyan Wang, Zhichao Li, Le Zhang, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth learning from video. *Int. J. Comput. Vis.*, 129(9):2548–2564, 2021.
- [6] Wenjie Chang, Yueyi Zhang, and Zhiwei Xiong. Transformer-based monocular depth estimation with attention supervision. In *British Machine Vision Conference (BMVC)*, page 136. BMVA Press, 2021.
- [7] Devendra Singh Chaplot, Dhiraj Gandhi, Abhinav Gupta, and Ruslan Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *NeurIPS*, 2020.
- [8] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. In *NeurIPS*, 2022.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL HLT*, 2019.
- [10] David Eigen and Job Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *European Conference on Computer Vision (ECCV)*, pages 2650–2658, 2015.
- [11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *Int. J. Robotics Res.*, 32(11):1231–1237, 2013.
- [12] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. In *European Conference on Computer Vision (ECCV)*, pages 3827–3837. IEEE, 2019.
- [13] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2482–2491, 2020.
- [14] Agrim Gupta, Jiajun Wu, Jia Deng, and Li Fei-Fei. Siamese masked autoencoders. *CoRR*, abs/2305.14344, 2023.
- [15] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022.
- [16] Pan Ji, Runze Li, Bir Bhanu, and Yi Xu. Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments. In *European Conference on Computer Vision (ECCV)*, pages 12767–12776, 2021.
- [17] Longlong Jing and Yingli Tian. Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey. *IEEE Trans. PAMI*, 2021.
- [18] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: Learning SFM from SFM. In *ECCV*, volume 11214, pages 713–728. Springer, 2018.
- [19] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *Mach. Intell. Res.*, 20(6):837–854, 2023.
- [20] Zhong Liu, Ran Li, Shuwei Shao, Xingming Wu, and Weihai Chen. Self-supervised monocular depth estimation with self-reference distillation and disparity offset refinement. *IEEE Trans. Circuits Syst. Video Technol.*, 33(12):7565–7577, 2023.
- [21] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu and Yong Liu, Xinxin Chen, and Yi Yuan. Hr-depth: High resolution self-supervised monocular depth estimation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*, pages 2294–2301, 2021.
- [22] Mirko Nava, Luca Maria Gambardella, and Alessandro Giusti. State-consistency loss for learning spatial perception tasks from partial labels. *IEEE Robotics Autom. Lett.*, 6(2):1112–1119, 2021.
- [23] Emanuele Palazzolo, Jens Behley, Philipp Lottes, Philippe Giguère, and Cyrill Stachniss. Refusion: 3d reconstruction in dynamic environments for RGB-D cameras exploiting residuals. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7855–7862, 2019.
- [24] Luigi Piccinelli, Christos Sakaridis, and Fisher Yu. idisc: Internal discretization for monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [25] Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. In *NeurIPS*, 2020.
- [26] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. *CoRL*, 2022.
- [27] Santhosh Kumar Ramakrishnan, Devendra Singh Chaplot, Ziad Al-Halah, Jitendra Malik, and Kristen Grauman. PONI: Potential Functions for ObjectGoal Navigation with Interaction-free Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [28] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *European Conference on Computer Vision (ECCV)*, pages 12159–12168, 2021.
- [29] 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 Trans. Pattern Anal. Mach. Intell.*, 44(3):1623–1637, 2022.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sathesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.
- [31] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchammi, Micael Tchammi, Kent Vainio, Josiah Wong, Li Fei-Fei, and Silvio Savarese. iGibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *IROS*, pages 7520–7527, 2021.
- [32] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGB-D images. In *ECCV*, pages 746–760, 2012.
- [33] Jaime Spencer, Simon Hadfield, Chris Russell, and Richard Bowden. Kick Back & Relax: Learning to reconstruct the world by watching SlowTV. In *European Conference on Computer Vision (ECCV)*, pages 15722–15733, 2023.
- [34] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of RGB-D SLAM systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 573–580. IEEE, 2012.
- [35] Libo Sun, Jia-Wang Bian, Huangying Zhan, Wei Yin, Ian D. Reid, and Chunhua Shen. Sc-depthv3: Robust self-supervised monocular depth estimation for dynamic scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(1):497–508, 2024.
- [36] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2022–2030, 2018.
- [37] Ruoyu Wang, Zehao Yu, and Shenghua Gao. Planedepth: Plane-based self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [38] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked Feature Prediction for Self-Supervised Visual Pre-Training. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [39] P. Weinzaepfel, V. Leroy, T. Lucas, R. Brégier, Y. Cabon, V. Arora, L. Antsfeld, B. Chidlovskii, G. Csurka, and J. Revaud. CroCo: Self-Supervised Pretraining for 3D Vision Tasks by Cross-View Completion. In *NeurIPS*, 2022.
- [40] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *European Conference on Computer Vision (ECCV)*, 2023.
- [41] Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for

- embodied agents. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9068–9079, 2018.
- [42] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A Simple Framework for Masked Image Modeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [43] Jiaying Yan, Hong Zhao, Penghui Bu, and Yusheng Jin. Channel-wise attention-based network for self-supervised monocular depth estimation. In *Proc. 3D Vision Conference (3DV)*, 2021.
- [44] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3VO: deep depth, deep pose and deep uncertainty for monocular visual odometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1278–1289, 2020.
- [45] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1983–1992, 2018.
- [46] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision. In *arXiv:2111.11432*, 2021.
- [47] Huangying Zhan, Chamara Saroj Weerasekera, Jia-Wang Bian, and Ian Reid. Visual odometry revisited: What should be learnt? In *International Conference on Robotics and Automation (ICRA)*, pages 4203–4210. IEEE, 2020.
- [48] 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. In *International Conference on 3D Vision (3DV)*, sep 2022.
- [49] Hang Zhou, David Greenwood, and Sarah Taylor. Self-supervised monocular depth estimation with internal feature fusion. In *British Machine Vision Conference (BMVC)*, page 378, 2021.
- [50] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6612–6619, 2017.
- [51] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.