

Learning Depth with Very Sparse Supervision

Antonio Loquercio¹, Alexey Dosovitskiy² and Davide Scaramuzza¹

Abstract—Motivated by the astonishing capabilities of natural intelligent agents and inspired by theories from psychology, this paper explores the idea that perception gets coupled to 3D properties of the world via interaction with the environment. Existing works for depth estimation require either massive amounts of annotated training data or some form of hard-coded geometrical constraint. This paper explores a new approach to learning depth perception requiring neither of those. Specifically, we propose a novel global-local network architecture that can be trained with the data observed by a robot exploring an environment: images and extremely sparse depth measurements, down to even a single pixel per image. From a pair of consecutive images, the proposed network outputs a latent representation of the camera’s and scene’s parameters, and a dense depth map. Experiments on several datasets show that, when ground truth is available even for just one of the image pixels, the proposed network can learn monocular dense depth estimation up to 22.5% more accurately than state-of-the-art approaches. We believe that this work, in addition to its scientific interest, lays the foundations to learn depth with extremely sparse supervision, which can be valuable to all robotic systems acting under severe bandwidth or sensing constraints.

I. INTRODUCTION

Understanding the three-dimensional structure of the world is crucial for the functioning of robotic systems: for instance, it supports path planning and navigation, as well as motion planning and object manipulation. Animals, including humans, obtain such three-dimensional understanding naturally, without any specialized training. By observing the environment and interacting with it [1], they learn to estimate (possibly non-metric) distances to objects using stereopsis and a variety of monocular cues [2], [3], including motion parallax, perspective, defocus, familiar object sizes. Could a robotic system acquire a metric understanding of its surrounding from a similar feedback?

Endowing robots with such an ability would be valuable for several applications. Consider for example a swarm of nano aerial vehicles, whose task is to explore a previously unseen environment [4], [5]. Constrained by the size and battery life, each robot can only carry limited sensing, *e.g.* a camera and a 1D distance sensor. Learning to estimate dense metric depth maps from such on-board sensors during operation would allow efficient exploration and obstacle avoidance [6].

Classically, multi-view geometry methods are used to reconstruct the 3D coordinates of points given their corresponding

This work was supported by the SNSF-ERC Starting Grant, Intel, and the Swiss National Center of Competence Research (NCCR) Robotics, through the Swiss National Science Foundation.

¹Robotics and Perception Group <http://rpg.ifi.uzh.ch/>, Dep. of Informatics, University of Zurich, and Dep. of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland loquercio@ifi.uzh.ch

²Google Research, Berlin, Germany




Fig. 1. We train a depth perception system with what would be available to a robot interacting with the environment: images and very sparse depth measurements. When trained in such sparse regimes (up to a single pixel per image), our approach learns to perceive depth with higher accuracy than state-of-the-art methods.

projections in multiple images. These geometric approaches, carefully engineered over decades, demonstrate impressive results in a variety of settings and applications [8]. One downside of this class of approaches is that they are using only some of the depth cues (mainly stereo and motion parallax), but typically do not exploit more subtle monocular cues, such as perspective, defocus or known object size. Unsupervised learning approaches to depth estimation [9], [10], [7] combine geometry with deep learning, with the hope that deep networks can learn to utilize the cues not used by the classic methods. In these approaches, depth estimators are trained from monocular or stereo video streams, using photometric consistency between different images as a loss function. Unsupervised learning approaches are remarkably successful in many cases, but they are fundamentally based on hard-coded geometry equations, which makes them potentially sensitive to the camera model and parameters, as well as difficult to tune.

A. Contributions

The present work is motivated by the following question: can a three dimensional perception system be trained with the data that a robot would observe interacting with the environment? To make the problem tractable, we make two assumptions. First, motivated by the extensive evidence from psychology and neuroscience on the fundamental importance of motion perception [11], [12], we provide pre-computed optical flow as an input to the depth estimation system. Optical

flow estimation can be learned either from synthetic data [13], [14] or from real data in an unsupervised fashion [15]. Second, we assume at training time the availability of images and only very sparse depth ground truth (just few pixels per image), similar to what a robot might collect with a 1-D distance sensor while navigating an environment.

To learn a depth estimator from such assumptions, we design a specialized global-local deep architecture consisting of two modules, global and local. The global module takes as input two images and optical flow between them and outputs a compact latent vector of “global parameters”. We expect those parameters to encode information about the observer’s motion, the camera’s intrinsics, and scene’s features, *e.g.*, planarity. The local module then generates a compact fully convolutional network, conditioned on the “global parameters”, and applies it to the optical flow field to generate the final depth estimate. The global and the local modules are trained jointly end-to-end with the available sparse depth labels and without camera’s pose or intrinsics ground truth.

Since the method is inspired by learning via interaction, we evaluate it on several indoor scenarios. We compare against generic deep architectures, unsupervised depth estimation approaches, and classic geometry-based methods. Standard convolutional networks show good results when trained with dense depth ground truth, but their performance degrades dramatically in the very sparse data regime. Unsupervised learning-methods and classic triangulation approaches are generally strong, but their performance in the challenging monocular two-frame indoor scenario suffers from suboptimal correspondence estimation due to homogeneous surfaces and occlusions. The proposed approach outperforms all these baselines, thanks to its ability to effectively train with very sparse depth labels and its robustness to imperfections in optical flow estimates.

II. RELATED WORK

The problem of recovering the three-dimensional structure of a scene from its two-dimensional projections has been long studied in computer vision [16], [17], [18], [19]. Classic methods are based on multi-view projective geometry [20]. The standard approach is to first find correspondences between images and then use these together with geometric constraints to estimate the camera motion between the images (for instance, with the eight-point algorithm [18]) and the 3D coordinates of the points (*e.g.*, via triangulation [17]). Numerous advanced variations of this basic pipeline have been proposed [21], [22], [23], [24], improving or modifying various its elements. However, key characteristics of these classic methods are that they crucially rely on projective geometry, require laborious hand-engineering, and are not able to exploit non-motion-related depth cues.

To make optimal use of all depth cues, machine learning methods can either be integrated into the classic pipeline, or replace it altogether. The challenge for supervised learning methods is the collection of training data: obtaining ground truth camera poses and geometry for large realistic scenes can be extremely challenging. An alternative is to train on

simulated data, but then generalization to diverse real-world scenes can become an issue. Therefore, while supervised learning methods have demonstrated impressive results [25], [26], [27], [28], it is desirable to develop algorithms that function in the absence of large annotated datasets.

Unsupervised (or self-supervised) learning provides an attractive alternative to the label-hungry supervised learning. The dominant approach is inspired by classic 3D reconstruction techniques and makes use of projective geometry and photometric consistency across frames. Existing works use various depth representations for this task: voxel grids [29], point clouds [30], triangular meshes [31] or depth maps [9], [10], [32], [33]. In this work we focus on the depth map representation. Among the methods for learning depth maps, some operate in the stereo setup (given a dataset of images recorded by a stereo pair of cameras) [9], [32], while others address the more challenging monocular setup, where the training data consists of monocular videos with arbitrary camera motions between the frames [10], [33]. Reprojection-based approaches can often yield good results in driving scenarios, but they crucially rely on geometric equations and precisely known camera parameters (one notable exception being the recent work in [34], which learns the camera parameters automatically) and enough textured views. In contrast, we do not require knowing the camera parameters in advance and are robust in low-textured indoor scenarios.

Several works, similar to ours, aim to learn 3D representations without explicitly applying geometric equations [35], [36], [37]. A scene, represented by one or several images, is encoded by a deep network into a latent vector, from which, given a target camera pose, a decoder network can generate new views of the scene. A downside of this technique is that the 3D representation is implicit and therefore cannot be directly used for downstream tasks such as navigation or motion planning. Moreover, at training time it requires knowing camera pose associated with each image. Our method, in contrast, does not require camera poses, and grounds its predictions in the physical world via very sparse depth supervision. This allows us to learn an explicit 3D representation in the form of depth maps.

III. METHODOLOGY

A. Model architecture

Given two monocular RGB images $\mathbf{I}_1, \mathbf{I}_2$, with unknown camera parameters and relative pose, as well as the optical flow \mathbf{w} between them, we aim to estimate a dense depth map corresponding to the first image. We assume to have an artificial agent equipped with a range sensor, which navigates through an indoor environment. By doing so, it collects a training dataset of image pairs, with depth ground truth \mathbf{d} available only for extremely few pixels. Using this sparsely annotated dataset, we train a deep network $F_{\theta}(\mathbf{I}_1, \mathbf{I}_2, \mathbf{w})$, with parameters θ , that predicts a dense depth map $\hat{\mathbf{d}}$ over the whole image plane. We now describe the network architecture in detail.

An overview of the global-local network architecture is provided in Figure 2. The system operates on an image pair




Fig. 2. Global-local model architecture. An image pair and an estimated flow field are first fed through the global module that estimates the “global parameters” vector \mathbf{g} , representing the camera motion. From these global parameters, the local module generates three convolutional filter banks and applies them to the optical flow field. The output of the local module is then processed by a convolution to generate the final depth estimate.

I_1, I_2 and the optical flow (dense point correspondences) \mathbf{w} between them. In this work, we estimate the flow field with an off-the-shelf optical flow estimation algorithm, which is neither trained nor tuned on our data.

The rest of the model is composed of two modules: a global module G that processes the whole image and outputs a compact vector of “global parameters” and a local module L that applies a compact fully convolutional network, conditioned on the global parameters, to the optical flow field. This design is motivated both by classic 3D reconstruction methods and by machine learning considerations. Establishing an analogy with classic pipelines, the global module corresponds to the relative camera pose estimation, while the local module corresponds to triangulation – estimation of depth given the image correspondences and the camera motion. These connections are described in more detail in the supplement. From the learning point of view, we aim to train a generalizable network with few labels, and therefore need to avoid overfitting. The local module is very compact and operates on a transferable representation – optical flow. The global network is bigger and takes raw images as input, but it communicates with the rest of the model only via the low-dimensional bottleneck of global parameters, which prevents potential overfitting.

The “global module” G is implemented by a convolutional encoder with global average pooling at the end. The network outputs a low-dimensional vector of “global parameters” $\mathbf{g} = G(I_1, I_2, \mathbf{w})$. The idea is that the vector represents the motion of the observer, although no explicit supervision is provided to enforce this behavior. While the optical flow alone is in principle sufficient for ego-motion estimation, we also feed the raw image pair to the network to supply it with additional cues.

The “local module” L takes as input the generated global parameters \mathbf{g} , as well as the optical flow field. First, the global parameter vector is processed by a linear perceptron that outputs several convolutional filters banks, collectively denoted by $\varphi = LP(\mathbf{g})$. Then, these filter banks are stacked

into a small fully convolutional network C_φ that is applied to the optical flow field. We append two channels of x - and y -image coordinates to the input \mathbf{w} of C_φ , as in CoordConv [38]. The output of C_φ is the final depth prediction $\hat{\mathbf{d}} = C_\varphi(\mathbf{w})$.

This design of the local module is motivated by classic geometric methods: for estimating the depth of a point it is sufficient to know its displacement between the two images, its image plane coordinates, and the camera motion. In contrast to this standard formulation of triangulation, we intentionally make the receptive field of the network larger than 1×1 pixel, so that the network has the opportunity to correct for small inaccuracies or outliers in the optical flow input.

B. Loss function

Similarly to previous work [27], [26], we define the loss on the inverse depth $\hat{\mathbf{z}} \doteq \hat{\mathbf{d}}^{-1}$. This is a common representation in computer vision and robotics [39], [22], which allows to naturally handle points and their uncertainty over a large range of depths. We use the L_1 loss on the inverse depth, averaged over the subset P of the pixels that have associated ground truth inverse depth \mathbf{z} :

$$\mathcal{L}_{\text{depth}} = \frac{1}{|P|} \sum_{i \in P} |\hat{z}^i - z^i|. \quad (1)$$

To encourage the local smoothness of the predicted depth maps, we add an L_1 regularization penalty on the gradient $\nabla \hat{\mathbf{z}} = (\partial_x \hat{\mathbf{z}}, \partial_y \hat{\mathbf{z}})$ of the estimated inverse depth. Similarly to classic structure from motion methods and unsupervised depth learning literature [32], we modulate this penalty according to the image gradients ∂I_1 , allowing depth discontinuities to be larger at points with large ∂I_1 :

$$\mathcal{L}_{\text{smooth}} = \frac{1}{|\Omega|} \sum_{i \in \Omega} |\partial_x \hat{z}^i| e^{-|\partial_x I_1^i|} + |\partial_y \hat{z}^i| e^{-|\partial_y I_1^i|}, \quad (2)$$

with Ω representing the full image plane. The full training loss of our network is a weighted sum of these two terms $\mathcal{L}_{\text{total}} = \lambda_p \mathcal{L}_{\text{depth}} + \lambda_s \mathcal{L}_{\text{smooth}}$.

	Abs-Inv	Abs-Rel	S-RMSE
Formula	$\frac{1}{ \Omega } \sum_{\Omega} \left \frac{1}{d} - \frac{1}{\hat{d}} \right $	$\frac{1}{ \Omega } \sum_{\Omega} \frac{ d - \hat{d} }{d}$	$\sqrt{\frac{1}{ \Omega } \sum_{\Omega} E_{\log}^2 - \left(\frac{1}{ \Omega } \sum_{\Omega} E_{\log} \right)^2}$

TABLE I

METRICS FOR QUANTITATIVE EVALUATION OF DEPTH ACCURACY. d IS THE GROUND TRUTH DEPTH, \hat{d} IS THE PREDICTION. FOR CONVENIENCE WE DENOTE $E_{\log} \doteq \log(d/\hat{d}) = \log d - \log \hat{d}$. FOR ALL METRICS, LOWER IS BETTER.




Fig. 3. Qualitative comparison of the depth maps generated with the baselines and our approach. Overall, Struc2Depth’s predictions are generally poor in homogeneous and repetitive regions, while DispNet tends to over-smooth depth maps. In contrast, our method can predict fine details of the scene geometry.

C. Model details

In all our experiments the input images have resolution 256×192 pixels. Unless mentioned otherwise, ground truth depth is provided for a single pixel of each image, but we also experiment with denser ground-truth signals. We use a pre-trained PWC-Net [14] for optical flow estimation.

We use the Leaky ReLU non-linearity in all networks. The global module is implemented by a 5-layer convolutional encoder with the number of channels growing from 16 to 256, with stride 2 in the first 4 layers. The last 256-channel hidden layer is followed by a convolution with 6 channels and global average pooling, resulting in the 6-dimensional predicted global parameter vector \mathbf{g} . The local module consists of a single linear perceptron which transforms \mathbf{g} linearly to a 3.9K vector. Empirically, we did not find an advantage in utilizing a multi-layer non-linear perceptron in place of the linear operation. The resulting vector is split into three parts, which are reshaped into filter banks with kernel size 3×3 and number of output channels 20, 10, 20, respectively. These filter banks, with Leaky ReLUs in between, constitute the compact fully-convolutional depth estimation network C_{φ} . The 20-channel network output is then processed by a single 3×3 convolutional layer to shrink the channels to 1. The design of this compact fully convolutional network has been inspired by the refinement layer used by previous works on supervised depth estimation [27], [26].

We train the model with the Adam optimizer [40] with an initial learning rate of 10^{-4} for a total of approximately

94K iterations with a mini-batch size of 16. We apply data augmentation during training. Further details are provided in the supplement.

IV. EXPERIMENTS

Navigating a physical agent in the real world to collect interaction data is challenging due to problem spanning perception, planning, and control. Therefore, in order to isolate the contribution of our proposed method, we simulate distance observations from a navigation agent, *e.g.* nano-drones [5], by masking out all depth ground truth except for a single one on several depth estimation benchmarks. Specifically, we test the approach on three datasets collected in cluttered indoor environments, either real or simulated: Scenes11, Sun3D, and RGB-D. Scenes11 [27] is a large synthetic dataset with randomly generated scenes composed of objects from ShapeNet [43] against diverse backgrounds composed of simple geometric shapes. SUN3D [44] is a large collection of RGB-D indoor videos collected with a Kinect sensor. RGB-D SLAM [45] is another RGB-D dataset collected with Kinect in indoor spaces.

For all datasets, we use the splits proposed by [27]. As commonly done in two-view depth estimation methods and in structure-from-motion methods [8], [27], we resolve the inherent scale ambiguity by normalizing the depth values such that the norm of the translation vector between the two views is equal to 1. To quantitatively evaluate the generated depth maps, we adopt three standard error metrics summarized in Table I. We compare to both standard convolutional

Method	Scenes11			SUN3D			RGB-D		
	Abs-Inv	Abs-Rel	S-RMSE	Abs-Inv	Abs-Rel	S-RMSE	Abs-Inv	Abs-Rel	S-RMSE
Eigen [25]	0.045	0.57	0.77	0.072	0.82	0.38	0.046	0.54	0.37
DispNet [41]	0.038	0.51	0.70	0.041	0.49	0.33	0.038	0.45	0.36
FCRN [42]	0.041	0.52	0.74	0.047	0.44	0.30	0.042	0.45	0.35
Small Enc-Dec	0.046	0.66	0.83	0.064	0.73	0.45	0.049	0.58	0.46
Struct2Depth [7]	0.058	0.95	0.81	0.037	0.44	0.27	0.037	0.44	0.48
Struct2Depth [7] + Flow	0.056	0.94	0.79	0.036	0.42	0.27	0.035	0.42	0.45
Ours	0.031	0.43	0.61	0.035	0.37	0.25	0.033	0.37	0.33

TABLE II

IN THE SPARSE TRAINING REGIME, OUR METHOD CAN EFFICIENTLY LEARN TO PREDICT DEPTH FROM SINGLE POINT SUPERVISION, OUTPERFORMING SIGNIFICANTLY BOTH STANDARD ARCHITECTURES AND UNSUPERVISED DEPTH ESTIMATION SYSTEMS. FOR ALL ERROR METRICS, LOWER IS BETTER.

neural networks, unsupervised methods, and classic structure from motion methods. Additional quantitative and qualitative comparisons are provided in the supplementary material.

A. Learning from very sparse ground truth

We compare the proposed global-local architecture to strong generic deep models – the encoder-decoder architecture of Eigen et al. [25], the popular fully convolutional architecture DispNet [41], and the multi-scale encoder-decoder of Laina et al. (FCRN) [42]. For a fair comparison with our method, we provide all the baselines with both the image pair and the optical flow field. Specifically, we generate input samples by concatenating the image sequence and the flow on the last dimension. We additionally tune the models to reach best performance on our task. The details of the tuning process are reported in the appendix. We also compare to a reduced-sized DispNet [41] (Small Enc-Dec), that has a number of parameters similar to our model (including both the global and the local module). Its encoder consists of 4 convolutions with (16, 32, 54, 128) filters, with sizes of (7, 5, 3, 3), and stride 2. Its decoder is composed of 4 up-convolutions with (128, 64, 32, 16) filters of size 3 and stride 1. Encoder and decoder layers are connected through skip connections. Finally, we compare against Struct2Depth [7], current state-of-the-art system for unsupervised depth estimation.

As shown in Table II, our approach outperforms all the baselines in the sparse supervision regime. Specifically, we outperform the architecture of Eigen et al. [25] on average by 53%, the architecture of Laina et al. [42] by 22.5%, and the fully convolutional DispNet by 20%. Indeed, due to over-parametrization, these baselines tend to overfit to the training points, failing to generalize to unobserved images and locations.

This is empirically demonstrated in Fig. 4, where we plot the depth loss on training points as a function of the number of iterations. Decreasing the size of the architecture to address overfitting does not however solve the problem: the Small Enc-Dec, with number of parameters similar to our network, achieves poor results, mainly due to its limited capacity.

Our approach also achieves on average 24% better error than the unsupervised depth estimation baseline [7] over all datasets and metrics. Indeed, the considered datasets represent




Fig. 4. For large networks, the loss on training points (solid lines) is significantly higher than the validation loss (dashed lines). In contrast, our global-local architecture learns generalizable representations.

a challenge for geometry-based methods given the presence of large homogeneous regions, occlusions, and small baselines between views, which are typical factors encountered in indoor scenes. Noticeably, the performance of Struct2Depth on the SUN3D dataset is relatively good, boosted by the larger baseline between views and the abundance of features. Interestingly, providing optical flow to the unsupervised baseline only increases performance of 3.5% on average. We hypothesize that this is due to the fact that unsupervised methods already estimate correspondences internally, and providing flow as input gives redundant information.

Fig. 5 analyzes the performance of our and the DispNet architectures (our strongest baseline) as a function of the number of observed ground-truth pixels per image. Unsurprisingly, both methods learn to predict accurate depth maps when dense annotations (D) are available. Decreasing the amount of supervision obviously leads to performance drops. However, for our method the error increases on average by only 5% when going to sparser supervision, compared to 12% for the baseline, which leads to a large advantage over the baseline in the single-pixel supervision regime. This shows that the global-local architecture provides an appropriate inductive bias for learning from extremely sparse depth ground-truth.

B. Robustness to Dynamically Changing Camera Parameters

In practical applications camera internal parameters, such as focal length, may change through time. Indeed, environmental




Fig. 5. Depth estimation errors with increasing number of training pixels per image and dense supervision (D). When supervision gets sparser, our method’s performance degrades more gracefully than the baseline.

Method	Scenes11			SUN3D			RGB-D		
	Abs-Inv	Abs-Rel	S-RMSE	Abs-Inv	Abs-Rel	S-RMSE	Abs-Inv	Abs-Rel	S-RMSE
Struct2Depth [7]	0.062	2.19	0.87	0.045	0.52	0.25	0.050	0.54	0.46
DispNet [41]	0.039	0.57	0.70	0.041	0.46	0.27	0.046	0.56	0.38
Ours	0.034	0.51	0.61	0.034	0.43	0.24	0.036	0.40	0.33

TABLE III

DEPTH ESTIMATION ERRORS WITH CAMERA INTRINSICS VARYING UP TO 20% OF THEIR NOMINAL VALUE BETWEEN VIEWS. BASED ON PROJECTIVE GEOMETRY, UNSUPERVISED METHODS SUFFER THE MOST FROM PARAMETER UNCERTAINTY.

changes like temperature, humidity and pressure could cause severe variations to their nominal value. Due to these variations, methods based on projective geometry, which are sensitive to the accuracy of calibration parameters, can experience large performance drops. Although the problem could be alleviated by automatic re-calibration, these changes would have to be detected in the first place and would require either collecting multiple views of an object [46], [47] or additional sensing [48].

We empirically study the robustness of our method and the baselines to dynamically changing camera intrinsics. In particular, we randomly change, for each image pair, the horizontal and vertical focal lengths, as well as the center of projection, by up to 20% of their nominal value. The unsupervised depth estimation baseline suffers the most from the uncertainty in the camera intrinsics. Indeed, its estimation error increases on average by 26% with respect to the case in which camera parameters are correctly set. In contrast, as our approach does not explicitly rely on projective geometry, it does not exhibit such sensitivity to the camera parameters. Indeed, it experiences only a small decrease in performance, of approximately 5% with respect to the case where the intrinsics are fixed, since the learning problem becomes more challenging.

C. Global parameters and the camera motion

According to the intuition behind our model, the global parameters should have information about the observer’s ego-

motion between the frames, and as such should be related to the actual metric camera motion. Here we study this relation empirically, by training a camera pose predictor on the output of our global module, in supervised fashion. Note that this is done for analysis purposes only, after our full model has been trained: at training time the model has no access to the ground truth camera poses. Specifically, we add a small two-layer MLP with 256 hidden units on top of the global module that is either pre-trained with our method or randomly initialized. We then either train the full network or only the appended small MLP to predict the camera motion in supervised fashion (details of the training process are provided in the supplement).

Results in Table IV show that the global parameters indeed contain information about the camera pose. In both training setups pre-trained network substantially outperforms the random initialization: 17% to 64% error reduction across datasets and metrics when only tuning the MLP and up to 11% error reduction when training the full system. Interestingly, our method is also competitive against classic state-of-the-art baselines for motion estimation [27].

D. Robustness to Optical Flow Outliers

Optical flow estimation plays a central role in our learning procedure. In this section, we study the impact of correspondences’ errors to the quality of the predicted depth map. Specifically, we compute the per-pixel Abs-Inv metric as a function of the optical flow’s error, normalized for

Method	Scenes11		SUN3D		RGB-D	
	rot	trans	rot	trans	rot	trans
Scratch-MLP	1.3	74.4	3.6	55.5	5.3	78.4
Pretrained-MLP	0.9	26.7	2.7	32.5	4.4	51.5
Scratch-Full	0.7	10.3	1.8	25.0	3.2	30.5
Pretrained-Full	0.7	9.2	1.7	22.4	3.2	28.7
KLT [27]	0.9	14.6	5.9	32.3	12.8	49.6
8-point FF [27]	1.3	19.4	3.7	33.3	4.7	46.1

TABLE IV
ESTIMATION OF CAMERA MOTION BASED ON THE GLOBAL
PARAMETERS ESTIMATED BY OUR MODEL.

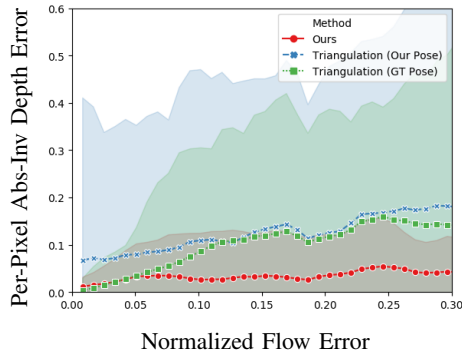


Fig. 6. Relation between error in correspondences and depth error for the local-global network and classic triangulation, either with perfect pose (GT Pose) or with the pose provided by our fine-tuned global network (see Sec. IV-C). Standard deviations of the errors are shadowed. Our approach learns to filter out errors in correspondences by exploiting its receptive field larger than one and regularities of those errors in the data.

the image size. The results of this study are shown in Fig. 6 for the SUN3D dataset. Results for other datasets are available in the appendix. Our approach (red line) can correct for outliers in the correspondences significantly better than the traditional triangulation approach. Interestingly, the main reason behind this behaviour does not consist in the precision of the camera motion estimation. Indeed, doing triangulation with either ground truth camera motion or the motion estimated from the global parameters (see Sec. IV-C) result in approximately the same performance. Conversely, we hypothesize that our approach can cope with outliers in correspondences because of the information encoded in the global parameters. Such information not only includes the relative motion between views, but also specifics about the scene (e.g. planarity) and the camera intrinsics. To retrieve this information, the global network uses monocular cues like perspective, focus, or parallax. Conditioned on this scene knowledge and exploiting regularities in the data, the local network can adapt to the observed scene and filter out outliers. Such high-level reasoning is absent in traditional triangulation.

E. Ablation study

Our architecture is based on several design choices that we now validate through an ablation study. In particular, we ablate the following components: (i) the use of optical flow as an intermediate representation, (ii) the estimation of global

variables to generate convolutional filters, (iii) the use of coordinate convolution in the fully convolutional network and (iv) the use of the image pair, in addition to optical flow, for the estimation of global parameters.

	Abs-Inv	Abs-Rel	S-RMSE
Full Model	0.033	0.43	0.61
– Image Pair	0.033	0.45	0.62
– CoordConv	0.038	0.52	0.71
– Glob. Mod.	0.041	0.55	0.73
– Flow	0.052	0.73	0.81

TABLE V
ABLATION STUDY ON THE SCENES11 DATASET.

The results in Table V show that all components are important and some have larger impact than others. The use of optical flow and coordinate convolution are crucial since they both provide essential cues for depth estimation. However, a basic encoder-decoder architecture (*i.e.* without global variables or coordinate convolutions) underperforms even when provided with optical flow. Unsurprisingly, the least important factor is providing the image pair to the global module, since, when camera parameters are fixed, the optical flow is a sufficient statistics of the observer’s ego-motion.

V. DISCUSSION

Motivated by the way natural agents learn to predict depth, we propose an approach for training a dense depth estimator from two unconstrained images given only very sparse supervision at training time and without the explicit use of geometry. We show that in cluttered indoor environments our global-local model outperforms state-of-the-art architectures for depth estimation by up to 22.5% in the sparse data regime.

Our methodology comes with some advantages and limitations with respect to previous work. One of the strongest advantage with respect to learning-based methods consists in its ability to learn from extremely sparse data. In addition, our method performs well in the cases where camera parameters are unknown or corrupted by noise. However, one limitation that our approach shares with supervised and unsupervised methods is the lack of generalization between visually different environments. This limitation can be softened by training in multiple indoor and outdoor environments. With respect to traditional geometry, our method can better cope with outliers in correspondences, given its ability to adapt to the scene characteristics. However, similarly to classic methods, our approach suffers in the case of pure rotation, where correspondences are not informative for depth. This limitation can be overcome by adding memory to the neural network through a recurrent connection.

In the future, we plan to address the aforementioned limitations to increase the network performance. In addition, we believe that a very exciting venue for future work to be the extension of our algorithm to a more strictly interactive procedure on a physical platform.

REFERENCES

- [1] J. J. Gibson, *The Ecological Approach to Visual Perception*. Houghton Mifflin, 1979.
- [2] I. P. Howard and B. J. Rogers, *Seeing in depth, Vol. 2: Depth perception*. University of Toronto Press, 2002.
- [3] E. B. Goldstein and J. Brockmole, *Sensation and perception*. Cengage Learning, 2016.
- [4] D. Palossi, A. Loquercio, F. Conti, E. Flamand, D. Scaramuzza, and L. Benini, "A 64mw dnn-based visual navigation engine for autonomous nano-drones," *IEEE Internet of Things Journal*, pp. 8357–8371, 2019.
- [5] K. N. McGuire, C. De Wagter, K. Tuyls, H. J. Kappen, and G. de Croon, "Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment," *Science Robotics*, vol. 4, no. 35, p. eaaw9710, 2019.
- [6] T. Cieslewski, E. Kaufmann, and D. Scaramuzza, "Rapid exploration with multi-rotors: A frontier selection method for high speed flight," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Sep. 2017, pp. 2135–2142.
- [7] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova, "Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos," in *Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, 2019, pp. 8001–8008.
- [8] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 30, no. 2, pp. 328–341, Feb. 2008.
- [9] R. Garg, B. G. V. Kumar, G. Carneiro, and I. D. Reid, "Unsupervised CNN for single view depth estimation: Geometry to the rescue," in *Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 740–756.
- [10] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017, pp. 6612–6619.
- [11] K. Nakayama and J. Loomis, "Optical velocity patterns, velocity-sensitive neurons, and space perception: a hypothesis," *Perception*, vol. 3, no. 1, pp. 63–80, 1974.
- [12] J. J. Koenderink, "Optic flow," *Vision research*, vol. 26, no. 1, pp. 161–179, 1986.
- [13] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017, pp. 2462–2470.
- [14] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 8934–8943.
- [15] S. Meister, J. Hur, and S. Roth, "Unflow: Unsupervised learning of optical flow with a bidirectional census loss," in *AAAI Conf. Artificial Intell.*, 2018, pp. 146–157.
- [16] H. C. Longuet-Higgins, "A computer algorithm for reconstructing a scene from two projections," *Nature*, no. 293, pp. 133–135, 1981.
- [17] R. I. Hartley and P. F. Sturm, "Triangulation," *Computer Vision and Image Understanding*, vol. 68, no. 2, pp. 146–157, 1997.
- [18] R. I. Hartley, "In defense of the eight-point algorithm," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, no. 6, pp. 580–593, 1997.
- [19] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon, "Bundle adjustment – a modern synthesis," in *Vision Algorithms: Theory and Practice*, ser. LNCS, W. Triggs, A. Zisserman, and R. Szeliski, Eds., vol. 1883. Springer Verlag, 2000, pp. 298–372.
- [20] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003, second Edition.
- [21] H. Hirschmuller and D. Scharstein, "Evaluation of stereo matching costs on images with radiometric differences," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 31, no. 9, pp. 1582–1599, 2009.
- [22] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, "DTAM: Dense tracking and mapping in real-time," in *Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 2320–2327.
- [23] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [24] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016, pp. 4104–4113.
- [25] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *Neural Inf. Process. Syst. (NeurIPS)*, 2014, pp. 2366–2374.
- [26] R. Ranftl, K. Lasinger, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *arXiv:1907.01341*, 2019.
- [27] B. UmmeHofer, H. Zhou, J. Uhrig, N. Mayer, E. Ilg, A. Dosovitskiy, and T. Brox, "Demon: Depth and motion network for learning monocular stereo," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017, pp. 5038–5047.
- [28] H. Zhou, B. UmmeHofer, and T. Brox, "Deeptam: Deep tracking and mapping," in *Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 822–838.
- [29] S. Tulsiani, A. A. Efros, and J. Malik, "Multi-view consistency as supervisory signal for learning shape and pose prediction," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 2897–2905.
- [30] C. Lin, C. Kong, and S. Lucey, "Learning efficient point cloud generation for dense 3d object reconstruction," in *AAAI Conf. Artificial Intell.*, 2018, pp. 7114–7121.
- [31] H. Kato, Y. Ushiku, and T. Harada, "Neural 3d mesh renderer," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 3907–3916.
- [32] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017, pp. 6602–6611.
- [33] R. Mahjourian, M. Wicke, and A. Angelova, "Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 5667–5675.
- [34] A. Gordon, H. Li, R. Jonschkowski, and A. Angelova, "Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019, pp. 8977–8986.
- [35] M. Tatarchenko, A. Dosovitskiy, and T. Brox, "Multi-view 3d models from single images with a convolutional network," in *Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 322–337.
- [36] D. J. Rezende, S. M. A. Eslami, S. Mohamed, P. Battaglia, M. Jaderberg, and N. Heess, "Unsupervised learning of 3d structure from images," in *Neural Inf. Process. Syst. (NeurIPS)*, 2016, pp. 4997–5005.
- [37] S. M. A. Eslami, D. Jimenez Rezende, F. Besse, F. Viola, A. S. Morcos, M. Garnelo, A. Ruderman, A. A. Rusu, I. Danihelka, K. Gregor, D. P. Reichert, L. Buesing, T. Weber, O. Vinyals, D. Rosenbaum, N. Rabinowitz, H. King, C. Hillier, M. Botvinick, D. Wierstra, K. Kavukcuoglu, and D. Hassabis, "Neural scene representation and rendering," *Science*, vol. 360, no. 6394, pp. 1204–1210, 2018.
- [38] R. Liu, J. Lehman, P. Molino, F. P. Such, E. Frank, A. Sergeev, and J. Yosinski, "An intriguing failing of convolutional neural networks and the CoordConv solution," in *Neural Inf. Process. Syst. (NeurIPS)*, 2018, pp. 9605–9616.
- [39] J. Civera, A. Davison, and J. Montiel, "Inverse depth parametrization for monocular slam," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 932–945, 2008.
- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Int. Conf. Learning Representations (ICLR)*, 2015.
- [41] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016, pp. 4040–4048.
- [42] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab, "Deeper depth prediction with fully convolutional residual networks," in *(3DV)*. IEEE, 2016, pp. 239–248.
- [43] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su *et al.*, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.
- [44] J. Xiao, A. Owens, and A. Torralba, "Sun3D: A database of big spaces reconstructed using sfm and object labels," in *Int. Conf. Comput. Vis. (ICCV)*, 2013, pp. 1625–1632.
- [45] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2012, pp. 573–580.
- [46] M. Pollefeys, R. Koch, and L. Van Gool, "Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters," *International Journal of Computer Vision*, vol. 32, no. 1, pp. 7–25, 1999.
- [47] J. I. Ronda, A. Valdés, and G. Gallego, "Line geometry and camera autocalibration," *Journal of Mathematical Imaging and Vision*, vol. 32, no. 2, pp. 193–214, 2008.
- [48] M. Hwangbo, J.-S. Kim, and T. Kanade, "Imu self-calibration using factorization," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 493–507, 2013.