

Self-supervised Learning with Geometric Constraints in Monocular Video

Connecting Flow, Depth, and Camera

Yuhua Chen^{1,2} Cordelia Schmid¹ Cristian Sminchisescu¹

¹Google Research ²ETH Zurich

yuhua.chen@vision.ee.ethz.ch, {cordelias, sminchisescu}@google.com

Abstract

We present **GLNet**, a self-supervised framework for learning depth, optical flow, camera pose and intrinsic parameters from monocular video – addressing the difficulty of acquiring realistic ground-truth for such tasks. We propose three contributions: 1) we design new loss functions that capture multiple geometric constraints (e.g. epipolar geometry) as well as an adaptive photometric loss that supports multiple moving objects, rigid and non-rigid, 2) we extend the model such that it predicts camera intrinsics, making it applicable to uncalibrated video, and 3) we propose several online refinement strategies that rely on the symmetry of our self-supervised loss in training and testing, in particular optimizing model parameters and/or the output of different tasks, thus leveraging their mutual interactions. The idea of jointly optimizing the system output, under all geometric and photometric constraints can be viewed as a dense generalization of classical bundle adjustment. We demonstrate the effectiveness of our method on KITTI and Cityscapes, where we outperform previous self-supervised approaches on multiple tasks. We also show good generalization for transfer learning in YouTube videos.

1. Introduction

One of the fundamental computer vision problems is estimating the 3D geometry of dynamic scenes captured by a moving camera. This encompasses many visual tasks such as depth estimation, optical flow, odometry, *etc.* Robust solutions would support a wide range of applications in autonomous driving, robotics, augmented reality, and scene interaction.

The 3D visual reconstruction problem has been studied extensively. At one end, there are structure from motion systems that leverage sparse, hand-crafted geometric features (e.g. SIFT [29] or SURF [2]), exact linear mathematical relations [17] (e.g. the epipolar constraint and the fundamental matrix for the two view geometry), as well as the

non-linear refinement of the structure and motion outputs under geometric re-projection losses based on bundle adjustment [39]. Impressive real-time processing pipelines of this type are now available, but they still offer sparse reconstructions, are affected by partial occlusion, and face difficulties in cases when the geometry of the scene or the motion is degenerate.

At the other end, recent years have witnessed the development of deep learning techniques where structure and motion estimation is formulated in a supervised setting [10, 24], as a pure prediction problem with little reference to exact geometric relations. This produces dense 3D estimates but comes with the inherent limitation of possible geometric inconsistency, the dependence on large training sets, domain sensitivity, and the difficulty of collecting good quality real-world ground-truth using special equipment like laser or stereo rigs with sensitive calibration [13]. Another option is to use synthetic data [31], however bridging the realism gap can be challenging.

More recently, several authors [15, 51] focused on designing self-supervised systems with training signals resulting from photometric consistency losses between pairs of images. These systems are effective and implicitly embed basic 3D geometry – as points in the first view are back-projected in 3D, displaced by the camera motion and re-projected in the second view –, but still lack the fundamental geometric constraints that link the rigid and apparent motion in multiple views.

In this work we introduce a self-supervised geometric learning framework, **GLNet**, which aims to integrate the advantages of modern deep-learning based self-supervised systems – (a) training without labeled data, (b) offering dense reconstructions where prior knowledge can be automatically incorporated, and (c) leveraging multiple interconnected tasks, with the ones of classical structure from motion – (i) explicitly representing exact mathematical relations (e.g. the epipolar constraint) that always hold for rigid scenes, (ii) being able to jointly refine all outputs, including depth, pose and camera intrinsics, under photometric

and geometric constraints, as in bundle adjustment, and (iii) breaking asymmetries between training and testing.

The proposed framework is extensively evaluated on KITTI and Cityscapes, where we achieve performance gains over the state-of-the-art. Moreover, our framework shows good performance in a transfer learning setting, and is able to robustly learn from uncalibrated video.

2. Related Work

Understanding the geometry and dynamics of 3D scenes usually involves estimating multiple properties such as depth, camera motion and optical flow. *Structure from motion* (SfM) [11, 17, 34, 37, 39, 44] or scene flow estimation [33, 40] are well-established methodologies with a solid record of fundamental and practical progress. Models for structure and motion estimation based on feature matching and geometric verification have produced impressive results [37], but their reconstructions are often sparse and prone to error in textureless or occluded areas.

To address the bottleneck of accurate feature matching, recent work has focused on deep-learning approaches for geometric reasoning. Several methods train networks based on ground-truth labels and have been successfully applied to many tasks, such as monocular depth prediction [9, 10, 27, 28], optical flow [8, 19, 35, 38] and camera pose estimation [22, 23, 24]. To leverage the multiple cues, the tasks can be tackled jointly by fusing boundary detection [20], the estimation of surface normals [28], semantic segmentation [4, 46, 50], *etc.*

To train with ground-truth labels, different authors rely either on specialized equipment (*e.g.* LIDAR) for data collection [13, 14] which is expensive in practice, or on synthetic dataset creation [31] which in many cases leads to domain shift [1, 5, 6]. To reduce the amount of ground-truth labels required for training, unsupervised approaches recently emerged. The core idea is to generate a differentiable warping between two images and use the photometric difference as proxy supervision. This has been applied to stereo images [12, 15, 49] and optical flow [21, 32, 45].

In this work, we focus on learning from monocular video. Many recent methods go along similar lines. Zhou *et al.* [51] couple the learning of monocular depth and ego-motion. Vijayanarasimhan *et al.* [41] additionally learn rigid motion parameters of multiple objects. Subsequent methods further proposed improvements based on various techniques, such as an ICP alignment loss [30], supervision from SfM algorithms [26], optical flow [48, 52], edges [47], modeling multiple rigid motions informed by instance segmentation [3], motion segmentation [36], minimum projection loss [16], *etc.*

3. Methodology

Overview Our Geometric Learning Network (GLNet), for which an overview is presented in fig. 1, solves the inter-related tasks of monocular depth prediction, optical flow, camera pose and intrinsics estimation, by relying on predictors for depth \mathbf{D}_θ , camera \mathbf{C}_α and flow \mathbf{F}_δ within a coupled optimization objective. \mathbf{D}_θ estimates the depth map \mathbf{D} from a single image. \mathbf{C}_α predicts the camera pose (\mathbf{R}, \mathbf{t}) between two adjacent frames, and also the intrinsic parameters \mathbf{K} whenever unknown. \mathbf{F}_δ operates on two image frames and estimates the optical flow. The joint output space of the tasks is denoted as $\mathbf{\Lambda} = \{\mathbf{D}, \mathbf{R}, \mathbf{t}, \mathbf{K}, \mathbf{F}\}$, where the superset of parameters is $\mathbf{\Theta} = \{\theta, \delta, \alpha\}$. The predictors $\{\mathbf{D}_\theta, \mathbf{C}_\alpha, \mathbf{F}_\delta\}$ are implemented as three neural networks with trainable parameters $\mathbf{\Theta}$, referred to as *DepthNet*, *CameraNet* and *FlowNet* respectively.

We formulate an optimization objective which consists of two parts: an adaptive photometric loss that captures the appearance similarity of the static and dynamic structures, and a geometric loss consisting of several components that couple the rigid and apparent motion. In training, the optimization objective is used as a proxy supervision signal to learn the parameters $\mathbf{\Theta}$ of the predictor. During inference, we can further refine the predictions $\mathbf{\Lambda}$ based on the same objective used in training, and we are able to focus on either refining the model parameters $\mathbf{\Theta}$ or the outputs $\mathbf{\Lambda}$, not unlike in geometric bundle adjustment. This is supported by our self-supervised objective and our explicit representation of classical geometric constraints.

3.1. Geometric and Appearance Fundamentals

Consider a source image I^s and a target image I^t , collected by a potentially moving camera with intrinsics \mathbf{K}_α and ego-motion \mathbf{M}_α . The 3D rigid transformation can be represented in homogeneous coordinates, in terms of a rotation matrix and a translation vector as

$$\mathbf{M}_\alpha = \begin{bmatrix} \mathbf{R}_\alpha & \mathbf{t}_\alpha \\ 0 & 1 \end{bmatrix} \quad (1)$$

For camera intrinsics, we fix the optical center at the image center, ignore camera skew and radial distortion, parameterize the focal length along the two optical axes as f^x and f^y , and work with images of resolution $h \times w$. The intrinsics camera matrix writes as

$$\mathbf{K}_\alpha = \begin{bmatrix} f_\alpha^x & 0 & w/2 \\ 0 & f_\alpha^y & h/2 \\ 0 & 0 & 1 \end{bmatrix} \quad (2)$$

Given a pixel \mathbf{p} in the source image I^s , the corresponding 3D scene position in the source camera coordinate system

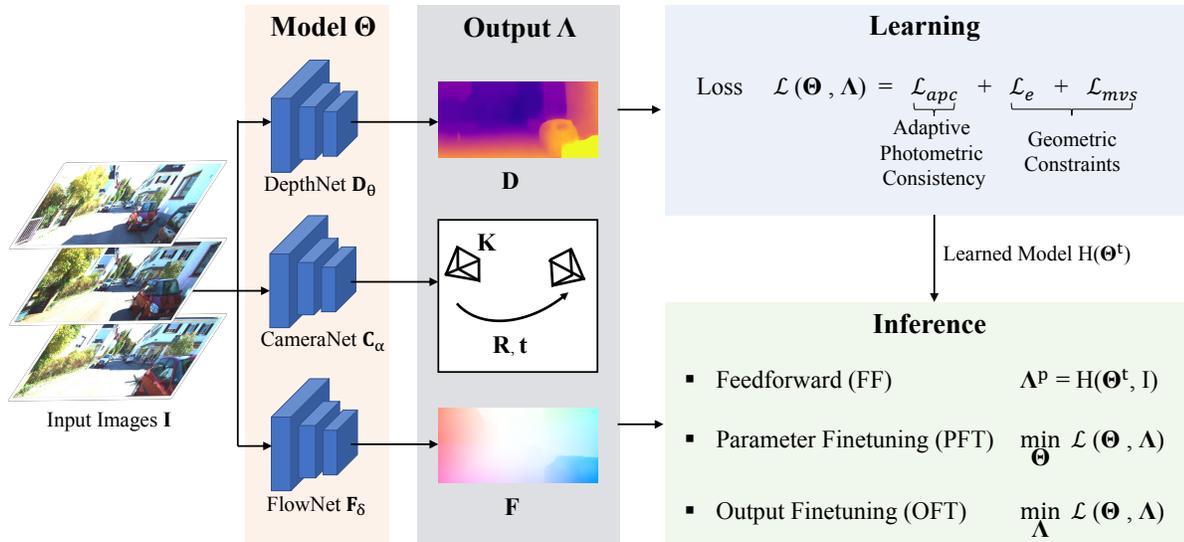


Figure 1. **Overview of the proposed GLNet framework.** Our model can take consecutive image frames as input, solves several tasks including depth, camera and flow estimation, and couples them via loss functions that capture the adaptive photometric and geometric constraints among outputs. An important feature of the proposed architecture is its training-testing symmetry as we learn a model in the training stage and during testing we further refine the parameters and the outputs based on the same optimization objective.

in homogeneous coordinates, can be back-projected as

$$\mathbf{x} = \begin{bmatrix} \mathbf{D}_\theta(\mathbf{p})\mathbf{K}_\alpha^{-1}\mathbf{p} \\ 1 \end{bmatrix} \quad (3)$$

Assuming we displace \mathbf{x} rigidly, the corresponding coordinates (projection) \mathbf{p}' in the target image I^t , are

$$\mathbf{p}' = [\mathbf{K}_\alpha|0]\mathbf{M}_\alpha \begin{bmatrix} \mathbf{D}_\theta(\mathbf{p})\mathbf{K}_\alpha^{-1}\mathbf{p} \\ 1 \end{bmatrix} \quad (4)$$

In this case, the displacement field $\mathbf{p}' - \mathbf{p}$ represents the 2D projection of the 3D *scene flow*, *i.e.* the real image motion induced by the true underlying 3D rigid displacement.

For pixels that cannot be explained by a rigid transformation, another key quantity is the measured, apparent motion or *optical flow* \mathbf{F}_δ that provides a dense correspondence field between the source and the target images, respectively

$$\mathbf{p}' = \mathbf{p} + \mathbf{F}_\delta(\mathbf{p}) \quad (5)$$

For simplicity, and in a slight abuse of notation, we also use \mathbf{p} as the 2D (non-homogeneous) image coordinates in the above equation.

3.2. Optimization Objectives

In this section we describe the loss we use for the self-supervised learning of depth, flow, and camera matrix. These tasks are interconnected through photometric and geometric constraints.

Adaptive Photometric Loss The view synthesis loss is widely used in self-supervised learning [51]. The loss measures the photometric difference between the synthesized and the actual image, where synthesis is obtained by a 3D reconstruction in the first frame with backprojected pixel intensities, followed by rigid displacement and perspective projection in the second frame. However, this displacement would only be valid for scene structures consistent with the ego-motion, or moving according to a global rigid displacement. For secondary or non-rigidly moving objects, an adaptive approach is necessary, and we pursue one here.

Given source and target images I^s and I^t , their pixels belong either to regions explained by ego-motion (or global scene motion), or by secondary dynamic objects. For scene structures not explained by the global rigid motion, one can rely on the more flexible optical flow. The main intuition of our adaptive photometric loss is to channel parameter updates towards only those configurations that best explain the displacement, be it produced by either ego-motion or secondary displacements. The term can thus be represented as the minimum photometric error between the two displacements (*i.e.* optical flow and rigid motion), as follows

$$\mathcal{L}_{apc} = \min \left\{ S\left(I^s(\mathbf{p}), I^t([\mathbf{K}_\alpha|0]\mathbf{M}_\alpha \begin{bmatrix} \mathbf{D}_\theta(\mathbf{p})\mathbf{K}_\alpha^{-1}\mathbf{p} \\ 1 \end{bmatrix})\right), S\left(I^s(\mathbf{p}), I^t(\mathbf{p} + \mathbf{F}_\delta(\mathbf{p}))\right) \right\} \quad (6)$$

The term is a per-pixel minimum between two components: the first part represents rigid displacement as in (4), and the second part represents the displacement produced by opti-

cal flow, *c.f.* (5). S is a similarity function between two pixels. As common practice [48, 52], we use a weighted sum of structural similarities (*SSIM*) [43] and L1 components, with trade-off parameter $r (= 0.85)$, as follows

$$S(a, b) = r \frac{1 - SSIM(a, b)}{2} + (1 - r) \|a - b\|_1 \quad (7)$$

Multi-view 3D Structure Consistency Predicting depth independently for each view usually results in inconsistency among the same scene structure in different (*e.g.* consecutive) views. To enforce structural consistency in 3D space, we design a loss component to penalize 3D structural deviations in multiple views.

Given a pixel \mathbf{p} in the source image, and its corresponding pixel in the target image \mathbf{p}' given by (4), their predicted 3D coordinates can be obtained by back-projection, *c.f.* (3). To obtain consistent structure, we penalize deviations between estimates of the same 3D point, once transformed to the common coordinate system of the target camera

$$\begin{aligned} \mathcal{L}_{mvs} &= \|\mathbf{x}' - \mathbf{x}\|_1 = \\ &= \left\| \begin{bmatrix} \mathbf{D}_\theta(\mathbf{p}') \mathbf{K}_\alpha^{-1} \mathbf{p}' \\ 1 \end{bmatrix} - \mathbf{M}_\alpha \begin{bmatrix} \mathbf{D}_\theta(\mathbf{p}) \mathbf{K}_\alpha^{-1} \mathbf{p} \\ 1 \end{bmatrix} \right\|_1 \end{aligned} \quad (8)$$

The loss represents the 3D discrepancy of predictions from two views, and the gradients are used to update *CameraNet* $\{\mathbf{K}_\alpha, \mathbf{R}_\alpha, \mathbf{t}_\alpha\}$ and *DepthNet* \mathbf{D}_θ . Notice that such a loss can immediately generalize to multiple views.

Epipolar Constraint Loss for Optical Flow The epipolar constraint is widely used in classical geometric methods [17], in order to compute a closed-form solution to initialize non-linear bundle adjustment procedures. Notably absent from existing deep learning-based structure and motion prediction systems, the epipolar constraint is an algebraic relationship that couples 3D scene projections, *e.g.* \mathbf{p} and \mathbf{p}' in two views, via a fundamental or essential matrix, that embeds the geometry of the camera displacement \mathbf{M} and its intrinsic parameters \mathbf{K} . An alternative view of the epipolar constraint, that we leverage in our model, is as a verification equation for putative correspondences, as provided, *e.g.* by optical flow, *c.f.* (5).

Optical flow has been traditionally formulated as an optimization of a matching and smoothness loss, or more recently as per-pixel regression problem within deep-learning methods [8, 19]. However, results computed that way may be inconsistent with any epipolar geometry. To endow the learning process with geometric awareness, we incorporate the epipolar constraint as a penalty over dense correspondences computed by optical flow. The resulting epipolar constraint loss writes

$$\mathcal{L}_e = \mathbf{p}^\top \mathbf{K}_\alpha^{-\top} \mathbf{R}_\alpha [\mathbf{t}_\alpha]_\times \mathbf{K}_\alpha^{-1} (\mathbf{p} + \mathbf{F}_\delta(\mathbf{p})) \quad (9)$$

where $[\mathbf{t}]_\times$ is the skew-symmetric matrix corresponding to the translation vector \mathbf{t} . The loss enforces a global epipolar constraint over the dense correspondences from optical flow, and the associated gradients are used to update *CameraNet* $\{\mathbf{K}_\alpha, \mathbf{R}_\alpha, \mathbf{t}_\alpha\}$ and *FlowNet* \mathbf{F}_δ .

Regularization We also rely on standard regularization terms \mathcal{L}_r for our learning framework. Specifically, we use separate spatial smoothness terms over depth \mathbf{D}_θ , and optical flow \mathbf{F}_δ respectively, as well as forward-backward flow consistency constraints.

Total Loss The total loss is a weighted sum over losses previously introduced, each densely summing over image pixels, with weighting dropped for simplicity

$$\mathcal{L}(\Lambda, \Theta) = \mathcal{L}_{apc} + \mathcal{L}_{mvs} + \mathcal{L}_e + \mathcal{L}_r \quad (10)$$

The model is trained jointly in an end-to-end manner.

3.3. Online Optimization

During inference, most existing approaches produce results by running trained models feed-forward. This can be problematic as the structural constraints between multiple outputs, or between inputs and outputs, *e.g.* image intensity, flow, camera motion and depth, are no longer preserved.

Our solution is to operate with – and optimize over – similar losses during both learning and inference, *c.f.* (10). This is possible as our objective is self-supervised, hence it eliminates the asymmetry between training and testing. Therefore an online refinement process, either for model parameters, or for model outputs, is possible, and can leverage task dependencies for optimal performance and seamless model adaptation to new environments.

Formally, we optimize our model on the training set in order to obtain initial estimates Θ^t . Given a test image pair from a monocular video, initial predictions are computed, feed-forward, in the standard way, as $\Lambda^p = \{\mathbf{D}^p, \mathbf{F}^p, \mathbf{R}^p, \mathbf{t}^p, \mathbf{K}^p\}$. The (self-supervised) objective in (10) can be computed, as no ground-truth label is needed. To prevent overfitting to the optimization objective, based on a single image pair, we further enforce a regularizer on the output space $\|\Lambda - \Lambda^p\|$ in order to penalize large deviation from the original prediction of the learned model. The full optimization objective writes as

$$\{\Lambda^*, \Theta^*\} = \arg \min_{\Theta, \Lambda} \mathcal{L}(\Theta, \Lambda) + \|\Lambda - \Lambda^p\| + \|\Theta - \Theta^t\| \quad (11)$$

To update the output, we design two strategies, which can be used together or independently. In **Parameter Finetuning (PFT)** we update the parameters of the predictor Θ based on the loss. In **Output Finetuning (OFT)** we optimize the output directly without recomputing the network but with the trained prediction as prior. This is implemented

by considering the outputs Λ as free variables initialized as the predictions of the trained neural network, and using the gradient of the loss to refine them. Besides initialization, the computational graph associated to the deep model is no longer used, and only one forward pass through the deep model is needed. Moreover, the number of output optimization variables is typically much lower than the number of network parameters. Thus OFT offers the benefit of increased speed – in practice one order of magnitude.

3.4. Network Architecture

The focus of this work is on different loss functions and fine-tuning options, hence our network design mostly aligns with existing self-supervised learning components for geometric processing. Here we briefly cover the architectures used in this work.

DepthNet maps an input image to a per-pixel depth map. It is based on a fully convolutional encoder-decoder structure. The encoder is based on ResNet18 [18]. The decoder relies on DispNet [31], and consists of several deconvolutional layers. Skip connections are used to provide spatial context, and depth is predicted at four different scales.

CameraNet takes two adjacent image frames as input, and regresses the 6DOF camera pose, represented as a translation vector and a relative rotation matrix parameterized in terms of three Euler angles. When learning from uncalibrated video, the network also predicts camera intrinsics. We use the model of [51] which is a small network with 8 convolutional layers. A global pooling on the last layer is used for the final prediction.

FlowNet predicts the optical flow \mathbf{F} between two adjacent image frames. We use the same architecture described in [48], which is an encoder-decoder with a ResNet backbone.

Notice that our framework is agnostic to the particular choice of each component network, and other options are possible. Hence we can benefit from higher performing networks for individual tasks.

4. Experiments

In this section, we validate **GLNet** through extensive experiments on depth, optical flow, and camera pose estimation. We first introduce the datasets and the parameter settings used in experiments.

Dataset Our experiments are mainly conducted on KITTI [13, 14] and Cityscapes [7]. KITTI is a widely used dataset for benchmarking geometric understanding tasks such as depth estimation, odometry, and optical flow. Images are captured using cameras mounted on cars. We evaluate using the ground-truth labels provided with the official KITTI dataset. We additionally train the framework

on Cityscapes [7], and study how well the proposed models transfer across datasets. Similar with KITTI, Cityscapes mainly contains data collected by cars driving in European cities.

Parameter Settings In training, we use Adam [25] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The initial learning rate is set to 2×10^{-4} and batch size to 4. The images are resized to 128×416 resolution. Each training sample is a snippet consisting of three consecutive frames. Additionally, random resizing, cropping, flipping and color jittering are used during training, as common practice, for data augmentation. The backbone network is initialized with ImageNet weights, and we optimize the network for maximum 30 epochs although convergence usually occurs earlier.

For online refinement, we initialize the models with weights learned in the training stage. We use a batch size of 1. The batch consists of the test image and its two adjacent frames. Online refinement is performed for 50 iterations on one test sample with the same hyper-parameter introduced before. No data augmentation is used in the inference phase.

In all experiments, the model with only standard (non-adaptive) photometric loss is used as the baseline.

4.1. Depth Estimation

We begin with the evaluation of depth estimation. As common in protocols previously used [51], we report results of depth estimation using the Eigen [10] split of the raw KITTI dataset [13], which consists of 697 test images. Frames that are similar to the test scenes are removed from the training set. We compare the performance of the proposed framework with the baseline, as well as recent state-of-the-art works in the same setting [3, 30, 42, 47, 48, 51, 52].

Main Results As shown in table 1, our method achieves significant gains over the baseline, as well as the other competing methods, when using ground-truth camera intrinsics (as typical). Qualitative results are shown in fig. 2, where we observe a clear improvement in the visual quality of the depth map.

We notice that the online refinement is not used in other methods except [3]. Thus to facilitate the comparison, we also report the results of **GLNet** without online refinement, denoted as **GLNet(-ref.)**. Still, this yields better performance than the competing methods without refinement [30, 42, 48, 51, 52], demonstrating the effectiveness of our proposed loss functions.

In the uncalibrated scenario, where no ground-truth camera intrinsics are given, we use our *CameraNet* to predict them from input images. The resulting performance matches the calibrated setting. This is a sanity check but not entirely surprising as the camera used to collect the test set has a similar setup with the ones used in training.

	Method	Calib.	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25$ ↑	$\delta < 1.25^2$ ↑	$\delta < 1.25^3$ ↑
Trained on KITTI	Zhou <i>et al.</i> * [51]	✓	0.183	1.595	6.709	0.270	0.734	0.902	0.959
	Mahjourian <i>et al.</i> [30]	✓	0.163	1.240	6.220	0.250	0.762	0.916	0.968
	GeoNet [48]	✓	0.155	1.296	5.857	0.233	0.793	0.931	0.973
	Wang <i>et al.</i> [42]	✓	0.151	1.257	5.583	0.228	0.810	0.936	0.974
	LEGO [47]	✓	0.162	1.352	6.276	0.252	-	-	-
	DF-Net [52]	✓	0.150	1.124	5.507	0.223	0.806	0.933	0.973
	Casser <i>et al.</i> [3]	✓	0.109	0.825	4.750	0.187	0.874	0.957	0.982
	<i>Baseline</i>	✓	0.156	1.450	5.913	0.228	0.861	0.931	0.955
	GLNet(-ref.)	✓	0.135	1.070	5.230	0.210	0.841	0.948	0.980
GLNet	×	0.100	0.811	4.806	0.189	0.875	0.958	0.982	
GLNet	✓	0.099	0.796	4.743	0.186	0.884	0.955	0.979	
Trained on CS	GeoNet [48]	✓	0.210	1.723	6.595	0.281	0.681	0.891	0.960
	Casser <i>et al.</i> [3]	✓	0.153	1.109	5.557	0.227	0.796	0.934	0.975
	<i>Baseline</i>	✓	0.206	1.611	6.609	0.281	0.682	0.895	0.959
	GLNet	×	0.144	1.492	5.473	0.219	0.831	0.932	0.967
	GLNet	✓	0.129	1.044	5.361	0.212	0.843	0.938	0.976

Table 1. **Results of depth estimation on the KITTI Eigen Split.** We report models trained on KITTI and Cityscapes. The best result in each setting is marked in bold. 'GLNet(-ref.)' denotes the results without online refinement. *The result of Zhou *et al.* [51] is based on their updated Github version.

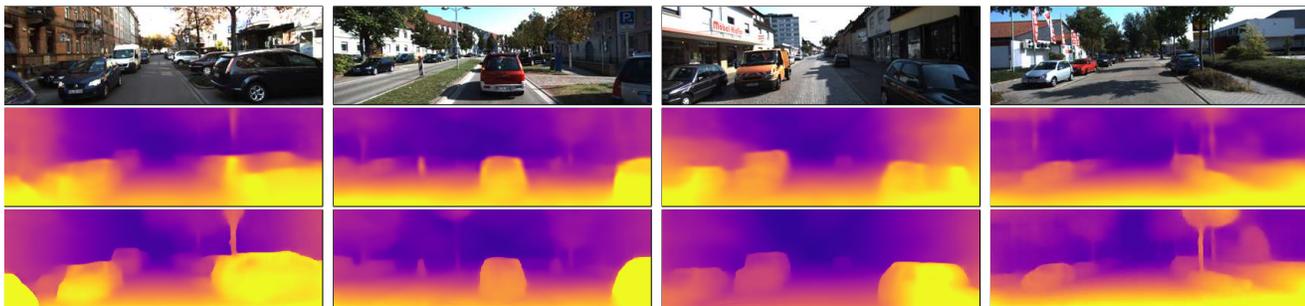


Figure 2. **Qualitative results of depth estimation.** *Top Row:* Input images, *Middle Row:* Baseline results, *Bottom Row:* GLNet results. The proposed framework offers sharper predictions.

The framework is also tested in a transfer learning setting. We train the models self-supervised on Cityscapes, then apply them on KITTI. Again, we observe performance gains when using the proposed components, and overall better results compared to competing methods. This shows the ability of the proposed framework to generalize to new environments. It also demonstrates that geometric constraints are powerful in closing domain gaps, as such mathematical relations are always valid. Interestingly, our method still displays competitive performance when no ground-truth camera calibration is given, even though the camera intrinsics of Cityscapes are different from those of KITTI.

Generalization to Videos in the Wild To further verify the generalization ability of the proposed method, we test GLNet on a set of videos collected from YouTube, where

the camera intrinsic parameters are unknown. We present qualitative results of GLNet in fig. 3. As can be seen, the proposed method is able to estimate depth from uncalibrated video frames across a variety of objects, structures and scenes. This supports the claim that our method can generalize to uncalibrated videos.

Ablation Study on Losses To analyze the individual impact of each loss component, we provide an ablation study over different combinations of losses. As shown in table 2, our adaptive photometric loss achieves a healthy improvement over the standard photometric loss. The performance is further improved by using the geometric losses, especially when coupled with online refinement. Out of the three losses proposed, the multi-view structure consistency seems most effective for depth estimation, which is under-

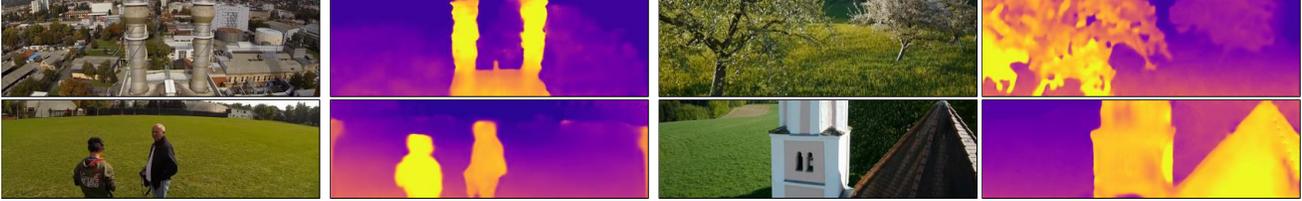


Figure 3. **Qualitative results on YouTube videos, where camera intrinsics are unknown.** *Left:* Input images, *Right:* Depth estimation results of **GLNet**. The proposed method can robustly estimate depth from uncalibrated video.

\mathcal{L}_{apc}	\mathcal{L}_{mvs}	\mathcal{L}_e	refinement	Abs Rel
				0.156
✓				0.144
✓	✓			0.138
✓	✓	✓		0.135
			✓	0.137
✓			✓	0.130
✓	✓		✓	0.103
✓	✓	✓	✓	0.099

Table 2. **Ablation study on losses.** We evaluate ablated version of the proposed method on the Eigen split. Notice the improvements offered by the various components.

standable as it is directly linked to 3D and integrates information over multiple views.

Online Refinement We also conduct an ablation study over refinement strategies. We use the normal parameter finetuning without regularizer as baseline, as in [3]. We compare the baseline with the two proposed refinement strategies, *i.e.* OFT and PFT. As shown in fig. 4, using standard finetuning we achieve some performance gain, however the downside is rapid overfitting to the test sample. The regularizer can effectively prevent the model from overfitting, for both PFT and OFT. The improvement from OFT is comparable to PFT. However, it is much faster compared to the other options as OFT only updates output variables, whose dimensionality is, for this problem, in the number of pixels $\sim 10k$. PFT needs to update the parameter of the neural network, whose size is around $\sim 1M$. At runtime, OFT takes about 2 seconds for 50 iterations, whereas PFT typically runs for around 40 seconds for the same number of iterations, which is *one order of magnitude slower*. Combining PFT and OFT brings some improvement over PFT alone, and achieves the best overall performance.

We also note online refinement is especially effective in the transfer learning setting. As illustrated in fig. 5, the model trained on KITTI produces erroneous depth prediction, when running the standard feed-forward component in a new environment. The prediction results can be significantly improved by both PFT and OFT. PFT delivers results that appear slightly sharper, visually, than OFT, which is

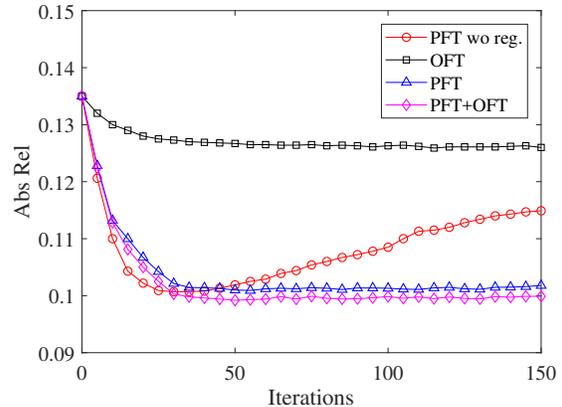


Figure 4. **Study of refinement strategies.** We report the Abs Rel evaluation metric as a function of iteration. Our method achieves the best result and does not overfit to the test sample as refinement proceeds.

consistent with the quantitative results.

4.2. Optical Flow

To benchmark optical flow, we use the KITTI 2015 stereo/flow training set [33] which has 200 training images. Similarly with previous work [48, 52], we use the training set for evaluation, as the proposed framework is self-supervised. We report the performance using the average end-point error (EPE) over non-occluded regions (Noc) and overall regions (All).

As shown in table 3, the proposed adaptive loss and multi-view structure consistency loss don't show much advantage over the baseline, possibly due to their indirect influence on the flow prediction. However, the epipolar constraint loss achieves a considerable performance gain over the baseline. Results can be further improved by online refinement. Qualitative results are provided in fig. 6, where we observe that geometric constraints substantially improve the flow quality for rigidly moving scene regions. In an uncalibrated setting, the performance is very close to the calibrated one. This is understandable, as the optical flow prediction does not rely on precise intrinsics but can benefit from geometric epipolar corrections which are however calibration-sensitive.

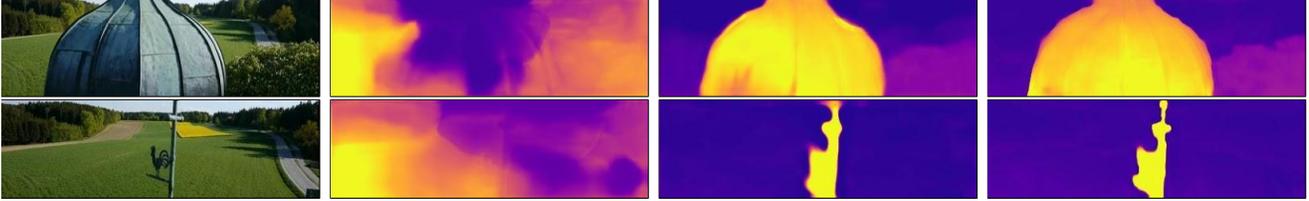


Figure 5. **Results of online refinement on YouTube videos.** From left to right: input images, feed-forward (FF) results, OFT results, PFT results. The results of online refinement, OFT, are noticeably better than the feed-forward model – comparable to PFT but much faster.

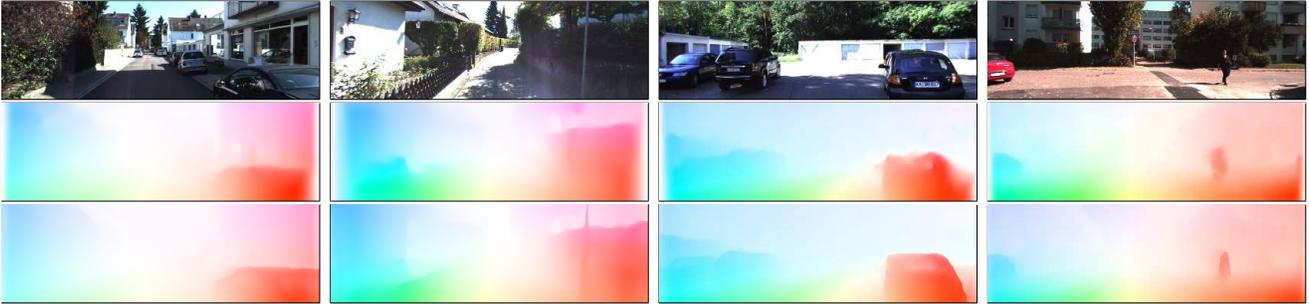


Figure 6. **Qualitative results of flow estimation.** Top Row: Input images, Middle Row: without geometric constraints, Bottom Row: with geometric constraints. The geometric constraints significantly improve the quality of the predicted rigid component flow.

Method	Noc	All
FlowNetS [8]	8.12	14.19
FlowNet2 [19]	4.93	10.06
GeoNet [48]	8.05	10.81
DF-Net [52]	-	8.98
<i>Baseline</i>	6.80	12.28
\mathcal{L}_{apc}	6.78	12.26
$\mathcal{L}_{apc} + \mathcal{L}_{mvs}$	6.77	12.20
$\mathcal{L}_{apc} + \mathcal{L}_{mvs} + \mathcal{L}_e$	5.40	8.95
GLNet	4.86	8.35
GLNet (uncalibrated)	4.90	8.41

Table 3. **Evaluation of optical flow.** We report average end-point error (EPE) on the KITTI 2015 flow training set over non-occluded regions (Noc) and overall regions (All). The first two methods are supervised, and other methods are trained on KITTI unsupervisedly. The best result is marked in bold.

4.3. Pose Estimation Results

We also evaluate the performance of our **GLNet** on the official KITTI visual odometry benchmark. As in the standard setting, we use the sequences 00-08 for training, and sequences 09-10 for testing. The pose estimation results are summarized in table 4, showing improvement over existing methods, as well as our baseline.

5. Conclusions

We have presented **GLNet**, a geometrically-inspired learning framework to jointly learn depth, flow, camera pose

Method	Seq.09	Seq.10
ORB-SLAM (full)	0.014 ± 0.008	0.012 ± 0.011
ORB-SLAM (short)	0.064 ± 0.141	0.064 ± 0.130
Zhou <i>et al.</i> [51]	0.016 ± 0.009	0.013 ± 0.009
Mahjourian <i>et al.</i> [30]	0.013 ± 0.010	0.012 ± 0.011
GeoNet [48]	0.012 ± 0.007	0.012 ± 0.009
DF-Net [52]	0.017 ± 0.007	0.015 ± 0.009
Casser <i>et al.</i> [3]	0.011 ± 0.006	0.011 ± 0.010
<i>Baseline</i>	0.013 ± 0.007	0.012 ± 0.010
GLNet	0.011 ± 0.006	0.011 ± 0.009

Table 4. **Evaluation on camera pose estimation.** Absolute trajectory error (ATE) on KITTI odometry dataset is reported as the evaluation metrics. Best result is shown in bold.

and intrinsic parameters from monocular video. The model is self-supervised and combines novel photometric and geometric loss functions, some based on fundamental relations like the epipolar constraint – and relying on continuously estimated correspondences from optical flow–, others on structure consistency over time. We’ve also introduced new parameter and output finetuning methods that generalize bundle adjustment, break the existing asymmetries between training and testing, and offer online adaption speed-ups of up to one order of magnitude. Given its ability to predict camera intrinsics, the model can be applied to uncalibrated video, and exhibits consistent performance across different training and testing domains. This supports the conclusion that geometric constraints represent a valuable regularizer in a transfer learning setting.

References

- [1] Amir Atapour-Abarghouei and Toby P Breckon. Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In *CVPR*, 2018. 2
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In *ECCV*, 2006. 1
- [3] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. *arXiv:1811.06152*, 2018. 2, 5, 6, 7, 8
- [4] Yuhua Chen, Wen Li, Xiaoran Chen, and Luc Van Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *CVPR*, 2019. 2
- [5] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster R-CNN for object detection in the wild. In *CVPR*, 2018. 2
- [6] Yuhua Chen, Wen Li, and Luc Van Gool. ROAD: reality oriented adaptation for semantic segmentation of urban scenes. In *CVPR*, 2018. 2
- [7] 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 *CVPR*, 2016. 5
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 2, 4, 8
- [9] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. 2
- [10] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014. 1, 2, 5
- [11] Yasutaka Furukawa, Brian Curless, Steven M Seitz, and Richard Szeliski. Towards internet-scale multi-view stereo. In *CVPR*, 2010. 2
- [12] Ravi Garg, Vijay Kumar BG, Gustavo Carneiro, and Ian Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016. 2
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1, 2, 5
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, 2012. 2, 5
- [15] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 1, 2
- [16] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel Brostow. Digging into self-supervised monocular depth estimation. *arXiv preprint arXiv:1806.01260*, 2018. 2
- [17] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1, 2, 4
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [19] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 2, 4, 8
- [20] Eddy Ilg, Tonmoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In *ECCV*, 2018. 2
- [21] J Yu Jason, Adam W Harley, and Konstantinos G Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *ECCV*, 2016. 2
- [22] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *ICRA*, 2016. 2
- [23] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, 2017. 2
- [24] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In *ICCV*, 2015. 1, 2
- [25] Diederik P Kingma and Jimmy Ba. ADAM: A method for stochastic optimization. *arXiv:1412.6980*, 2014. 5
- [26] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning sfm from sfm. In *ECCV*, 2018. 2
- [27] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016. 2
- [28] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical CRFs. In *CVPR*, 2015. 2
- [29] David G Lowe et al. Object recognition from local scale-invariant features. In *ICCV*, 1999. 1
- [30] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *CVPR*, 2018. 2, 5, 6, 8
- [31] 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. In *CVPR*, 2016. 1, 2, 5
- [32] Simon Meister, Junhwa Hur, and Stefan Roth. UnFlow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 2018. 2
- [33] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, 2015. 2, 7
- [34] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. DTAM: Dense tracking and mapping in real-time. In *ICCV*, 2011. 2

- [35] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, 2017. 2
- [36] 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 *CVPR*, 2019. 2
- [37] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 2
- [38] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 2
- [39] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment: a modern synthesis. In *International workshop on vision algorithms*, pages 298–372. Springer, 1999. 1, 2
- [40] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *ICCV*, 1999. 2
- [41] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. SfM-Net: Learning of structure and motion from video. *arXiv:1704.07804*, 2017. 2
- [42] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *CVPR*, 2018. 5, 6
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4
- [44] Changchang Wu. VisualSFM: A visual structure from motion system. 2011. 2
- [45] Jonas Wulff and Michael J Black. Temporal interpolation as an unsupervised pretraining task for optical flow estimation. *arXiv:1809.08317*, 2018. 2
- [46] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. PAD-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *CVPR*, 2018. 2
- [47] Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, and Ram Nevatia. LEGO: Learning edge with geometry all at once by watching videos. In *CVPR*, 2018. 2, 5, 6
- [48] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 2, 4, 5, 6, 7, 8
- [49] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *CVPR*, 2018. 2
- [50] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Zequn Jie, Xiang Li, and Jian Yang. Joint task-recursive learning for semantic segmentation and depth estimation. In *ECCV*, 2018. 2
- [51] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 1, 2, 3, 5, 6, 8
- [52] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. DF-Net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018. 2, 4, 5, 6, 7, 8