

Single-Stage 3D Geometry-Preserving Depth Estimation Model Training on Dataset Mixtures with Uncalibrated Stereo Data

Nikolay Patakin^{1,2}, Anna Vorontsova¹, Mikhail Artemyev¹, Anton Konushin¹
¹Samsung AI Center, Moscow ²Skolkovo Institute of Science and Technology (Skoltech)
 {n.patakin, a.vorontsova, m.artemyev, a.konushin}@samsung.com

Abstract

Nowadays, robotics, AR, and 3D modeling applications attract considerable attention to single-view depth estimation (SVDE) as it allows estimating scene geometry from a single RGB image. Recent works have demonstrated that the accuracy of an SVDE method hugely depends on the diversity and volume of the training data. However, RGB-D datasets obtained via depth capturing or 3D reconstruction are typically small, synthetic datasets are not photorealistic enough, and all these datasets lack diversity. The large-scale and diverse data can be sourced from stereo images or stereo videos from the web. Typically being uncalibrated, stereo data provides disparities up to unknown shift (geometrically incomplete data), so stereo-trained SVDE methods cannot recover 3D geometry. It was recently shown that the distorted point clouds obtained with a stereo-trained SVDE method can be corrected with additional point cloud modules (PCM) separately trained on the geometrically complete data. On the contrary, we propose GP², General-Purpose and Geometry-Preserving training scheme, and show that conventional SVDE models can learn correct shifts themselves without any post-processing, benefiting from using stereo data even in the geometry-preserving setting. Through experiments on different dataset mixtures, we prove that GP²-trained models outperform methods relying on PCM in both accuracy and speed, and report the state-of-the-art results in the general-purpose geometry-preserving SVDE. Moreover, we show that SVDE models can learn to predict geometrically correct depth even when geometrically complete data comprises the minor part of the training set.

1. Introduction

Single-view monocular depth estimation (SVDE) aims at estimating a dense depth map corresponding to an input RGB image. Being a fundamental problem in computer vision and visual understanding, it has numerous applications

in autonomous driving, robotics navigation and manipulation, and augmented reality. Most applications request for approaches able to process an arbitrary RGB image and estimate a depth map that allows recovering the 3D geometry of a scene (e.g. a point cloud can be reconstructed given camera calibration).

Early SVDE methods were designed to work either in indoor scenes [19] or in autonomous driving scenarios [8], thus being domain-specific. Accordingly, these methods demonstrate poor generalization ability in cross-domain experiments [25]. Overall, it is empirically proved that the generalization ability of SVDE models heavily depends on the diversity of the training data [4, 15, 25, 31, 35].

Recently, efforts have been made to obtain depth data from new sources: from computer simulation [2, 26, 28], via 3D reconstruction [15], or from stereo data. Among them, stereo videos [25, 31] and stereo images [10, 32, 33] collected from the web are the most diverse and accessible data sources. However, the depth data obtained from stereo videos is geometrically incomplete since stereo cameras' intrinsic and extrinsic parameters are typically unknown. Respectively, disparity from a stereo pair can be computed up to unknown shift and scale coefficients (UTSS). Such depth data can be used as a proxy for a ground truth depth map, but is insufficient to restore 3D geometry.

As a result, modern general-purpose SVDE methods trained on stereo data output predictions that cannot be used to recover 3D geometry [16, 25, 33]. Hence, we describe these methods as *not geometry-preserving* and their predictions as *geometrically incorrect*. The only SVDE method that is both general-purpose and *geometry-preserving* is LeReS [35]. It addresses geometry-preserving depth estimation with a multi-stage pipeline that recovers shift and focal length through predicted depth post-processing. This post-processing is performed using trainable point-cloud modules (PCM), thus resulting in significant overhead. Moreover, training PCM requires geometrically complete data with known camera parameters which limits the possible sources of training data.

In this work, we introduce GP², an end-to-end training

scheme and show that SVDE models can learn correct shift themselves and benefit from training on geometrically incomplete data while being geometry-preserving.

To obtain geometrically correct depth predictions, we do not employ additional models like PCM in LeReS [35]. Instead, we use geometrically complete data to encourage an SVDE model to predict geometrically correct depth maps while taking advantage of diverse and large-scale UTSS data. Thus, the proposed training scheme can be applied to make any SVDE model general-purpose and geometry-preserving. To prove this, we train different SVDE models using GP² and evaluate them on unseen datasets.

We demonstrate that the LeReS SVDE model trained with our scheme outperforms the original LeReS+PCM [35] trained on the same data. Then, we improve this result by switching to the LRN-based [21] SVDE model, thus pushing the boundaries of general-purpose geometry-preserving SVDE even further.

Overall, our contribution is three-fold:

- In this paper, we present an end-to-end training scheme GP² that allows training a general-purpose geometry-preserving SVDE method;
- Through multiple zero-shot cross-dataset trials, we show that our training scheme improves both depth estimates and point cloud reconstructions obtained with the existing SVDE methods. As a result, we set a new state-of-the-art in the general-purpose geometry-preserving SVDE with an SVDE model being much faster on inference than competing methods that use PCM;
- In ablation studies, we show that an SVDE model learns to predict geometrically correct depth even when trained mainly on UTSS data with a small portion of the geometrically complete data.

The rest of the paper is organized as follows: in Sec. 2, we discuss the relationship between depth data and 3D geometry and give an overview of existing general-purpose SVDE approaches. Our training scheme is introduced in Sec. 3. Sec. 4 is dedicated to experiments: there, we describe our experimental protocol, report quantitative results, and demonstrate visualizations.

2. Related Work

SVDE methods have been actively investigated for decades and have significantly evolved over this period. Early depth estimation methods employed complicated heuristics to process hand-crafted features [9,27]. Recently, deep learning-based approaches have been adopted for solving a wide range of computer vision tasks, including depth estimation.

As in other computer vision problems, deep learning-based SVDE methods rely heavily on the diversity and the amount of training data. Significant work has been done to introduce new datasets, but each of them has its own biases and limitations depending on data acquisition methods. The majority of modern SVDE methods formulate depth estimation as dense pixel-wise labelling in continuous space [3, 5, 13, 20]. Traditionally, depth estimation models were trained using closed-domain sensor-captured datasets with absolute depth measurements, such as NYU [19] and KITTI [8, 18]. Such models overfit to the training domain and do not generalize well [25].

2.1. Depth Data and 3D Geometry

Not all applications that use depth information require absolute depth measurements. For instance, depth ordinal rankings can be sufficient to imitate the bokeh effect for image processing or to estimate occlusion boundaries in AR applications. Ordinal SVDE methods address depth estimation by comparing the depth of each two pixels. Accordingly, the training data might contain only ordinal rankings for several pixels of each depth map; this data can be manually labeled by showing random pixel pairs to an assessor and asking which pixel in a pair is closer [4].

Single-view depth estimation can also be reformulated as predicting inverse depth up to unknown scale and shift (UTSS) [25]. UTSS data can be obtained from a stereo image through stereo matching even when calibration camera parameters are unknown. Hence, learning to estimate depth up to unknown shift and scale opens opportunities for training SVDE models on large-scale and diverse stereo data collected from the web [25, 31, 32]. UTSS predictions are geometrically incomplete and insufficient to restore 3D geometry, yet they can be used in the same scenarios as depth ordinal rankings.

Recovering 3D geometry is necessary for such applications as one-shot 3D photography [12]. Accordingly, these applications require depth estimates that can be either absolute or known up-to-scale (UTS). In the UTS setting, we assume the entire scene to be uniformly scaled by an unknown factor, so estimated UTS depth preserves 3D geometry and keeps the objects shapes unchanged. Ground truth UTS depth maps can be obtained by a combination of structure-from-motion (SfM) and multi-view stereo (MVS) methods given a multi-view image collection [15]. However, this scheme allows processing only static objects: the ground truth depth cannot be obtained for moving objects such as people, cars, or animals.

2.2. General-purpose SVDE

MiDaS. In MiDaS [25], multiple geometrically complete (UTS) and geometrically incomplete (UTSS) datasets are mixed together to obtain diverse and large-scale train-

ing data. Scale-and-shift invariant loss functions were used to address stereo datasets geometric incompleteness. As a result, the MiDaS model outputs UTSS predictions, so 3D geometry cannot be recovered.

LeReS. To the best of our knowledge, the only existing general-purpose and geometry-preserving SVDE method is LeReS [35]. It estimates geometrically correct depth maps through a multi-stage pipeline (see Fig. 1). This pipeline includes an SVDE model and two point cloud modules (PCM). The SVDE model is trained on a mixture of absolute and UTSS data and outputs initial UTSS depth estimates that are refined with PCM. Specifically, one point cloud module corrects depth shift, while another estimates camera focal length. Both modules are essential for estimating geometrically correct depth.

It is worth noting, that the LeReS SVDE model is trained to predict UTSS depth, while stereo data provides UTSS disparity. UTSS disparity cannot be straightforwardly converted to UTSS depth, as adding a constant to a disparity map corresponds to a non-linear depth map transformation. Accordingly, LeReS uses ranking loss functions on uncalibrated stereo data. This leads to an inefficient use of the training data, since metric information is then missed: two ranked points might be either almost equidistant or be far from each other.

UTSS losses computed in disparity space allow to utilize stereo data more efficiently [25]. In this case, UTSS disparities are to be predicted. To reconstruct a point cloud from such an estimate, disparity shift needs to be recovered. Since the LeReS SVDE model outputs UTSS depth, PCM yields a single additive constant in depth domain. Accordingly, it cannot be used to perform the required non-linear depth correction, thus being inapplicable for correcting predictions of stereo-trained SVDE models [24, 25].

Besides aforementioned drawbacks, LeReS requires calibrated data for training PCM. In contrast, our training scheme allows making geometrically correct estimates even if camera parameters are unknown. In other words, GP² imposes weaker restrictions on the training data that opens up opportunities for training on large-scale and diverse data. Furthermore, while LeReS addresses geometry-preserving SVDE with a multi-stage pipeline, we solve the same problem with a single-stage approach.

3. Proposed Training Scheme

The key observation behind our scheme is that SVDE models that use UTSS data for training can still learn to make geometrically correct predictions out of UTS samples present in the training mixture. Existing general-purpose methods trained with UTSS data use only UTSS [24, 25, 35] or ranking losses [35], thus missing helpful information about correct depth shifts contained in geometrically complete data. On the contrary, we use this information during

training to enforce the geometrical correctness of depth estimates.

Mixing datasets. To train an SVDE model on the mixture of UTS and UTSS data, we propose using a combination of a scale-invariant (UTS) loss and a shift-and-scale-invariant (UTSS) loss; we use the former whereas possible. More formally, we train SVDE models using the following loss function:

$$\mathcal{L}_{Mixture} = \mathbb{I}_{UTS} \mathcal{L}_{UTS} + \mathcal{L}_{UTSS}, \quad (1)$$

where $\mathbb{I}_{UTS} = 1$ for UTS samples and 0 for UTSS samples. The UTSS loss forces the SVDE model to generalize to the diverse UTSS stereo datasets, while the UTS loss encourages producing geometrically correct estimates.

Choosing predictions space. Uncalibrated stereo data provides inverse-depth, or *disparity*, up to unknown shift and scale; therefore, it cannot be converted into depth straightforwardly. Thus, UTSS SVDE models are typically trained to predict disparities [24, 25]. On the contrary, we aim at training an up-to-scale SVDE model, so we follow best practices for UTS models while choosing predictions space. The advantages of training UTS models to predict depth in logarithmic space are proved empirically [7, 12, 14, 15], and a number of loss functions for log-depth space are proposed [6, 15]. Eventually, we opt for making log-depth predictions.

Shift- and scale-invariant loss. The existing shift- and scale-invariant loss functions [25, 31, 35] are based on L_1 loss function with additional quantile trimmings or *tanh*-based adjustments [35]. In this study, we want to mitigate the effects of using elaborate loss functions, so we validate our training scheme in the most straightforward and demonstrative experimental setting. Specifically, we use a pure L_1 loss function without any additional terms.

Let l denote log-depth estimated by the model, $d = \exp(l)$ – estimated depth, and d^* – ground truth depth (for geometrically complete datasets). We can express disparity as $D = \frac{1}{d} = \exp(-l)$, and ground truth disparity as $D^* = \frac{1}{d^*}$. Before computing loss function, we align the mean and standard deviation of predictions and ground truth values to make it shift- and scale-invariant:

$$\hat{D} = \frac{D - \mu_D}{\sigma_D} \quad \hat{D}^* = \frac{D^* - \mu_{D^*}}{\sigma_{D^*}} \quad (2)$$

Then, the UTSS loss function for a single sample can be formulated as:

$$\mathcal{L}_{UTSS} = \frac{1}{N} \sum_{i=1}^N \left| \hat{D}_i - \hat{D}_i^* \right|, \quad (3)$$

Scale-invariant UTS loss. Such basic functions as L_1 or L_2 , can be used as scale-invariant UTS loss functions. But, a number of more elaborate task-specific scale-invariant

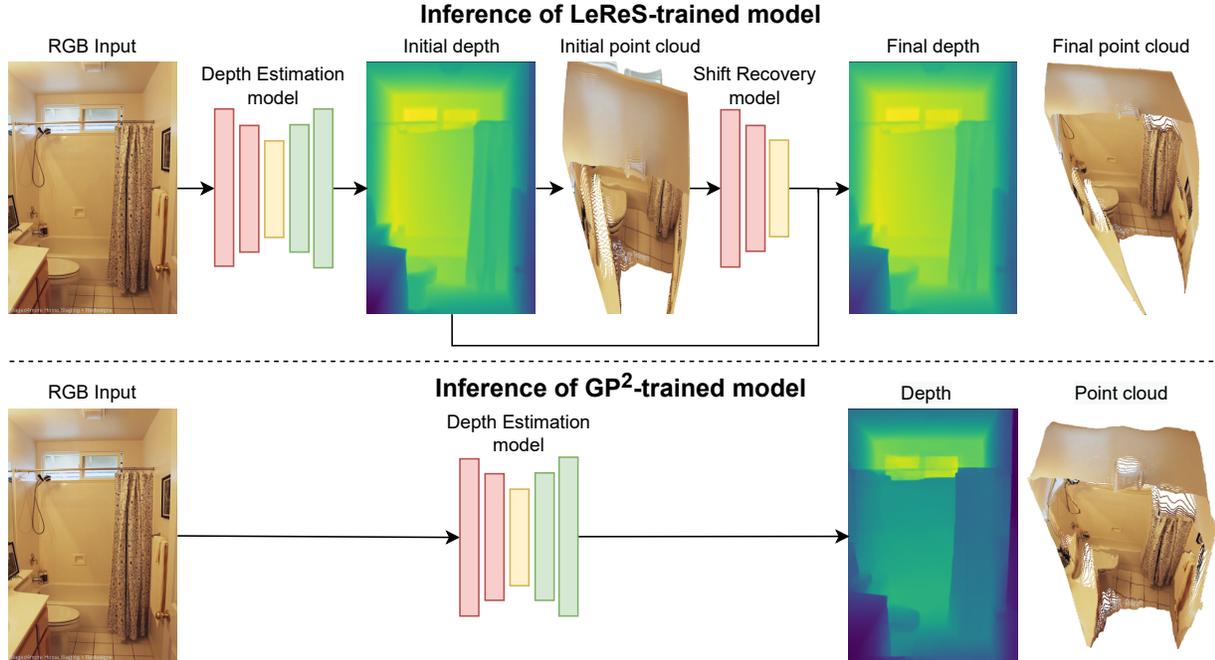


Figure 1. Our approach compared to the only existing general-purpose geometry-preserving depth estimation method. Being trained with our GP² scheme, conventional SVDE models can learn correct shift values themselves and do not require any additional post-processing.

loss functions have been introduced: specifically, Multi-Scale Gradient Loss [15], and Pairwise Loss [6].

In our study, we calculate scale-invariant loss function in log-depth space, beforehand we align median of log-differences:

$$\mu = \text{median}(l - \log d^*) \quad (4)$$

and UTS loss function can be written as:

$$\mathcal{L}_{UTS} = \frac{1}{N} \sum_{i=1}^N |l_i - \log d_i^* - \mu|, \quad (5)$$

Our training scheme does not depend on the choice of a loss function, so other scale-invariant and shift- and scale-invariant losses can be used instead of scale-invariant and shift- and scale-invariant pointwise L_1 .

4. Experiments

This section is dedicated to experiments: we describe our experimental protocol, report quantitative results, and present visualizations. We show the effectiveness of the proposed GP² scheme against multi-stage LeReS approach, as well as against methods that use only UTS data for training. In ablation studies, we show that even when UTS data comprises a minor part of a training set, it is sufficient to train the geometry-preserving SVDE method.

4.1. SVDE Models

First, we conduct experiments with the SVDE model from LeReS [35]. To validate the GP² scheme with different architectures, we train models based on the Lightweight Refine Net [21] meta-architecture. For an efficient light-weight model, we use EfficientNet-Lite0 [30] as a backbone, and for a more powerful model, we choose EfficientNet-B5 [29]. Hereinafter, we refer to these LRN-based SVDE models as to Lite0-LRN and B5-LRN, respectively.

4.2. Dataset Mixtures

To demonstrate versatility of our training scheme, we train SVDE models on two datasets mixtures introduced in MiDaS [25] and LeReS [35], respectively.

Not all the datasets presenting in these mixtures can be used 'as is' due to reproducibility issues. Some of these datasets contain data in a form of web links referring to the sources which are not publicly available [25], or has been deleted ever since [31]; while some other datasets do not contain depth data explicitly, so disparities should be extracted through some unknown stereo-matching protocol [10, 11, 25]. In general, we aim at reproducing dataset mixture as close to the original ones as possible; however, they are non-equivalent.

MiDaS training data. Similar to MiDaS, we train SVDE models on a mixture of four datasets:

Training dataset mixture	Method	Training scheme	TUM	NYU	Sintel	ETH3D	iBims-1	Average rank \downarrow
			$\delta_{1.25} \downarrow$	$\delta_{1.25} \downarrow$	rel \downarrow	rel \downarrow	rel \downarrow	
MD	MegaDepth [15]		0.331	0.344	0.490	0.276	0.389	12
MannequinChallenge	Li et al. [14]		0.226	0.239	0.431	0.249	0.175	10.8
MIX6 [24] \rightarrow NYU [19]	DPT [24]		0.151	0.049*	0.389	0.189	0.101	5.2
DIML _I +MD (only UTS data)	Lite0-LRN	UTS-only	0.170	0.142	0.428	0.191	0.117	9
	B5-LRN	UTS-only	0.145	<u>0.106</u>	0.397	0.174	0.100	4.4
	LeReS	UTS-only	0.146	0.123	0.417	0.174	0.108	6.4
DIML _I +RW+MD+MV	Lite0-LRN	GP ²	0.144	0.142	0.354	0.177	0.115	6.4
	B5-LRN	GP ²	0.130	<u>0.106</u>	<u>0.328</u>	0.154	<u>0.098</u>	2.2
	LeReS	GP ²	<u>0.139</u>	0.117	0.363	0.175	0.110	5.4
LeReS mixture	LeReS+PCM [35]	LeReS	0.193	0.139	0.403	0.172	0.140	7.8
	B5-LRN	GP ²	0.237	0.095	0.320	0.131	0.092	<u>3.2</u>
	LeReS	GP ²	0.192	0.114	0.343	<u>0.146</u>	0.103	4.6

Table 1. Results of SVDE models trained on different dataset mixtures with and without GP². The state-of-the-art results for are marked bold, the second best are underlined.

* This model was fine-tuned on the training subset of NYUv2. Other models had never seen any NYUv2 samples during training.

MegaDepth [15] (denoted as MD for brevity; contains UTS data), DIML Indoor [11] (DIML_I, UTS), RedWeb [32] (RW, UTSS), and Stereo Movies (MV, UTSS). We do not use WSVD dataset [31] from the original mixture. Differently, we use 49 stereo movies (listed in Supplementary) against 23 used in the original MiDaS. Overall, our Stereo Movies dataset contains \approx 140k samples against 75k samples in [25]. Hereinafter, we refer to the obtained dataset mixture as to DIML_I+RW+MD+MV.

LeReS training data. The LeReS dataset mixture contains Taskonomy [36] (absolute depths), 3D Ken Burns [22] (absolute), DIML Outdoor [11] (absolute), HoloPix [10] (UTSS), and HRWSI [33] (UTSS). Below, we denote this dataset mixture as the LeReS mixture.

We evaluate all SVDE models on several datasets unseen during training: the NYU test set, the split of TUM RGB-D proposed by Li et al. [14], the synthetic Sintel dataset, ETH3D, and iBims-1. For the ETH3D dataset, we follow MiDaS, rendering ground truth depth maps from reconstructed point clouds.

4.3. Experimental Protocol

Implementation details. We train LRN-based models with Radam [17] optimizer with LookAhead [37]. The learning rate is 0.001, and each batch consists of 24 samples. We select samples from each training dataset with equal probability, as proposed in MiDaS [25], which helps to avoid domain biases caused by unequal datasets sizes. The training takes 40 epochs with 10,000 training steps per epoch. During training, we use the same set of augmentations: resizing, padding, taking random crops to obtain samples of size 384×384 , geometrical transformations, and color distortions. During inference, we resize input images so their smaller side becomes 384 pixels.

We reproduce the LeReS training procedure using the official code for a fair comparison. Specifically, we train the SVDE model using SGD with a batch size of 40, an initial learning rate of 0.02, and a learning rate decay of 0.1. Input images are resized to 448×448 , and flipped horizontally with a probability of 0.5.

All methods are implemented with Python and PyTorch [23]. For training and evaluation, we used Nvidia Tesla P40, while performance tests were run on a Nvidia GeForce GTX 1080 Ti.

Point cloud reconstruction. To obtain a point cloud from a depth map, camera parameters should be known. We assume that the RGB camera used to capture data can be described in terms of a pinhole model. Accordingly, the unprojection from 2D coordinates and depth to 3D points can be performed as:

$$x = \frac{u - u_0}{f}d, y = \frac{v - v_0}{f}d, z = d \quad (6)$$

where (u_0, v_0) are the camera optical center, f is the focal length, and d is the depth. The focal length affects the point cloud shape as it scales x and y coordinates, but not z . We assume that the camera optical center is the geometric center of an image, so only focal distance should be estimated. Following original LeReS, we estimate it via PCM.

Metrics. We follow the MiDaS evaluation protocol, using $\delta_{1.25}$ for NYU and TUM RGB-D and rel for Sintel, ETH3D, and iBims-1. We measure the quality of point cloud reconstruction with pointwise RMSE instead of LSIV proposed in LeReS. We argue that it is a more transparent and straightforward approach since it does not require any additional modalities such as segmentation maps in LSIV.

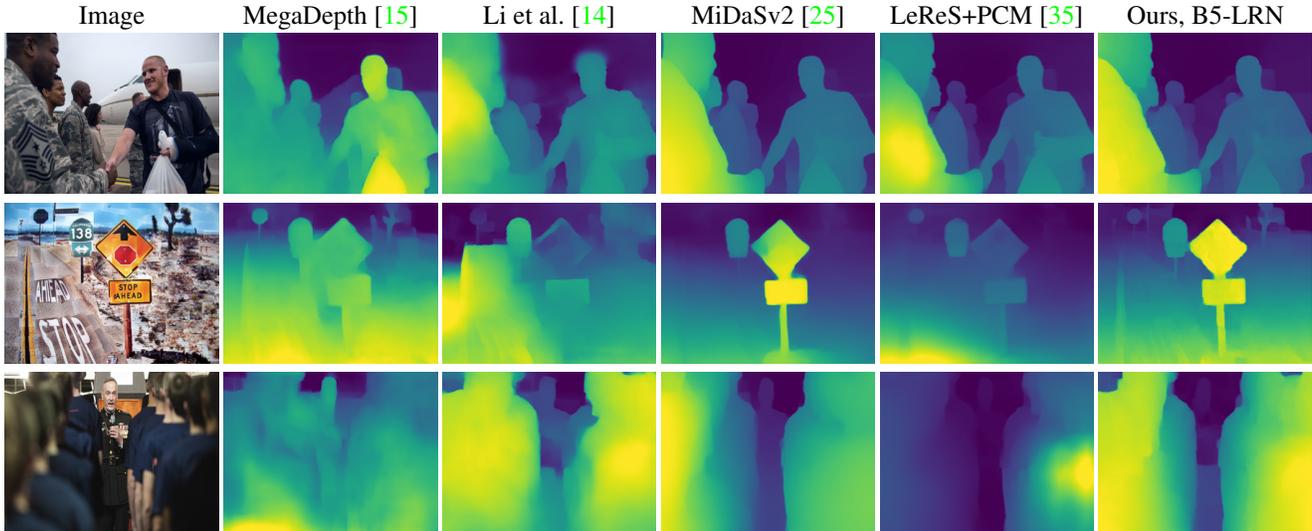


Figure 2. Depth maps estimated with existing SVDE methods and our GP²-trained B5-LRN model.

Training dataset mixture	Method	Training scheme	$\delta_{1.25} \downarrow$	rel \downarrow	Point cloud RMSE \downarrow
LeReS mixture	B5-LRN+PCM	GP ²	0.127	0.113	0.363
	LeReS+PCM	GP ²	0.161	0.137	0.395
	LeReS+PCM	LeReS	0.216	0.164	0.429

Table 2. Results of depth estimation and point cloud reconstruction on the 2D-3D-S dataset for methods trained with and without GP². The GP²-trained methods use pretrained PCM from LeReS only for camera focal length estimation.

4.4. Quantitative Comparison

In this section, we demonstrate the efficiency of our training scheme for depth estimation and point cloud reconstruction.

Depth estimation. We present the results of our SVDE models trained on DIML₁+RW+MD+MV and LeReS mixtures in Tab. 1.

According to average ranks, the GP²-trained LeReS SVDE model outperforms the original LeReS+PCM while using the same training data. Moreover, GP² allows making UTS predictions with a simple single-stage pipeline without any post-processing such as PCM.

As one might see, training on the UTSS part of DIML₁+RW+MD+MV mixture via GP² is profitable for all examined SVDE models as compared to the baselines trained on the UTS data only.

Also, we claim our GP²-trained B5-LRN to set a new state-of-the-art in the general-purpose geometry-preserving SVDE. Being trained with GP² on the LeReS mixture, B5-LRN is the most accurate on NYU, Sintel, ETH3D, and iBims-1. At the same time, our B5-LRN model trained on the DIML₁+RW+MD+MV mixture via GP² has the best average rank among all models.

Note that since our SVDE models are trained to estimate

UTS depth, we compare them to other UTS SVDE methods. Such well-known stereo-trained methods as DPT [24] (trained on MIX6 [24] only) and MiDaS [25] predict UTSS disparity, so the shift of these estimates cannot be corrected via shift recovery, as discussed in Sec. 2.2. Accordingly, we do not perform quantitative evaluation against SVDE methods predicting UTSS disparity, yet we present the qualitative comparison in Fig. 2 and supplementary. Authors of DPT [24] released several models, including the one fine-tuned on the training subset of NYUv2 to estimate absolute depth. In Tab. 1, we report the results of the fine-tuned DPT model.

There, we visualize depth estimates obtained by several existing SVDE methods (Li et. al [14], MegaDepth [15], MiDaS [25], LeReS+PCM [35]) and our GP²-trained B5-LRN model.

Point cloud reconstruction. It is worth noticing that geometry-preserving SVDE implies predicting depth maps that allow recovering 3D geometry, but it does not imply recovering 3D geometry itself. Accordingly, we do not propose any point cloud reconstruction techniques and use the LeReS PCM for focal length estimation instead. Importantly, our SVDE pipeline does not require PCM. We reconstruct point clouds only for a comprehensive comparison with LeReS+PCM.

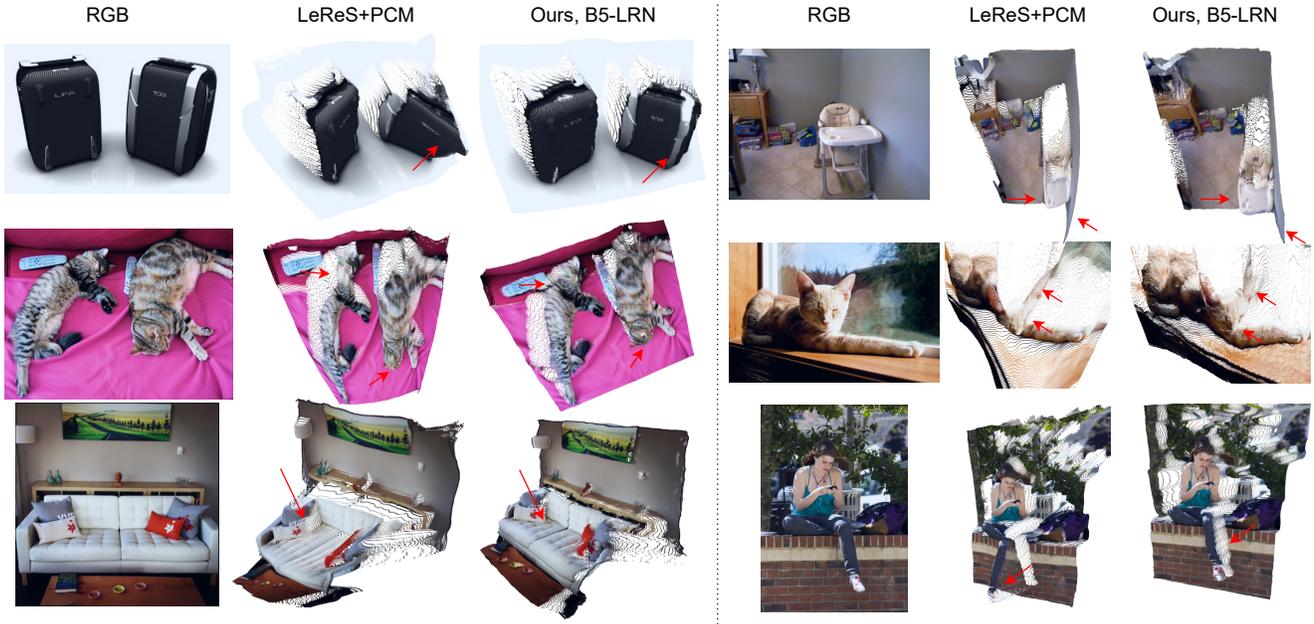


Figure 3. Point clouds reconstructed from depth maps estimated by the original LeReS+PCM and our GP²-trained B5-LRN using PCM for focal recovery. Red arrows are pointing on areas which appear to be difficult for LeReS+PCM, while our method recovers 3D geometry nicely in these areas.

In Tab. 2, we report point cloud accuracy on the 2D-3D-S dataset [1] unseen during training. Since both LeReS and our GP²-trained SVDE model rely on the same PCM, the observed RMSE improvement should be attributed to the more accurate depth estimation.

The examples of reconstructed point clouds are visualized in Fig. 3. We identify the areas where the original LeReS+PCM outputs erroneous predictions while our GP²-trained B5-LRN estimates 3D geometry adequately (depicted with red arrows).

4.5. Ablation Study

UTS / UTSS ratio. In ablation studies, we show that even when UTS data comprises a minor part of a training set, it is sufficient to train the geometry-preserving SVDE method. To prove that, we conduct experiments with our B5-LRN model on the NYU Raw dataset [19]. Since it contains absolute depth maps, we can make them either UTS or UTSS. To convert absolute depth to UTS, we multiply it by a random positive coefficient. To obtain UTSS data, we multiply the inverse depth by a random scale and then shift it by a random additive value. In this study, we divide the data into two parts and make them UTS and UTSS, respectively. To obtain comprehensive results, we run experiments with different ratios of UTS data. We compare the model trained with GP² on the UTS and UTSS data against the same model trained only on the UTS part of data.

According to Fig. 4, the more UTS training data we use,

the more accurate are the depth estimates. At the same time, the method trained on a mixture of UTS and UTSS data shows similar results for all UTS ratios; in other words, it performs as if it was trained on the dataset consisting of UTS data only. In practice, that means that the full geometrically correct dataset is not required since the proposed GP² allows training as accurate SVDE model with only 10-20% of UTS data.

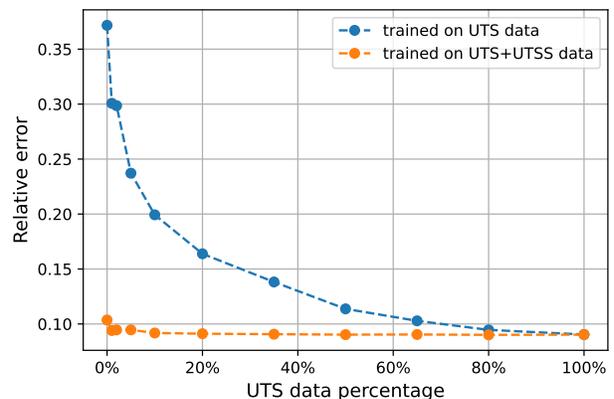


Figure 4. Results of the ablation study on the ratio of UTS and UTSS data from NYU Raw.

Model	Backbone	Inference time, ms	Average rank↓
LeReS+PCM [35]	ResNeXt101 [34]	201.9	7.8
LeReS SVDE model	ResNeXt101 [34]	45.1	4.6
B5-LRN (ours)	EfficientNet-B5 [29]	35.4	2.2
Lite0-LRN (ours)	EfficientNet-lite0 [30]	11.1	6.4

Table 3. Average GPU latencies for different SVDE methods.

4.6. Performance

Tab. 3 shows inference time for different SVDE models with and without PCM. The reported timings were measured on a single Nvidia GeForce GTX 1080 Ti GPU with a batch size of 1. We infer RGB images in 384×384 resolution and report the latency averaged over 100 runs. To prove that we do not achieve better performance at the cost of depth estimation accuracy, we also report the average dataset rank for each model (same as in Tab. 1).

Here, we do not consider the CPU time required to project an estimated depth map into a point cloud, which is necessary to run PCM. Nevertheless, the use of PCM results in a significant overhead even if only model GPU inference time is taken into account.

Hence, if computational resources are limited, fast and efficient SVDE models that do not rely on PCM should be preferred. Accordingly, we can conclude that GP²-trained SVDE models outperform LeReS+PCM in both depth estimation accuracy and inference speed.

5. Limitations

Through extensive evaluation, we demonstrate that GP²-trained SVDE models perform well on both indoor and outdoor datasets. However, these datasets do not cover all possible use cases, so we cannot state that our method will be as accurate if given arbitrary inputs.

6. Broader Impact

In this paper, we aim to build a method that would push the boundaries of general-purpose geometry-preserving SVDE and become a new state-of-the-art. Hence, future research in SVDE might benefit from this paper. Our method may be potentially used in self-driving perception and navigation systems, thus improving the safety of autonomous vehicles. At the same time, there is a concern on privacy issues of the widespread deployment of modern computer vision systems. Like many computer vision algorithms that can be used for surveillance purposes, our method may harm privacy if exploited by malicious actors. Nevertheless, we believe that the positive effect on safety is more significant than a hypothetical negative effect on privacy.

7. Conclusion

In this paper, we presented GP², a novel scheme of training an arbitrary SVDE model so as it becomes general-purpose and geometry-preserving. Experiments with different SVDE models and multiple dataset mixtures proved that our training scheme improves both depth estimates and point cloud reconstructions obtained with these models. Furthermore, with an SVDE model trained with GP², we set a new state-of-the-art in the general-purpose geometry-preserving SVDE. Moreover, we showed that a small amount of UTS data in the mixture is sufficient to train a geometry-preserving SVDE model; this opens new opportunities for using large-scale and diverse yet uncalibrated stereo data for geometry-preserving SVDE. Through performance tests, we demonstrated that GP²-trained SVDE models outperform the competitors not only in depth estimation accuracy but also in inference speed.

References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 7
- [2] Daniel Butler, Jonas Wulff, Garrett Stanley, and Michael Black. A naturalistic open source movie for optical flow evaluation. pages 611–625, 10 2012. 1
- [3] Yuanzhouhan Cao, Zifeng Wu, and Chunhua Shen. Estimating depth from monocular images as classification using deep fully convolutional residual networks. *IEEE*, 2017. 2
- [4] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 730–738. Curran Associates, Inc., 2016. 1, 2
- [5] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2650–2658, 2015. 2
- [6] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network, 2014. 3, 4
- [7] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation, 2018. 3
- [8] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark

- suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 1, 2
- [9] Derek Hoiem, A.A. Efros, and M. Hebert. Geometric context from a single image. volume 1, pages 654–661 Vol. 1, 11 2005. 2
- [10] Yiwen Hua, Puneet Kohli, Pritish Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. Holopix50k: A large-scale in-the-wild stereo image dataset. In *CVPR Workshop on Computer Vision for Augmented and Virtual Reality, Seattle, WA, 2020.*, June 2020. 1, 4, 5
- [11] Youngjung Kim, Hyungjoo Jung, Dongbo Min, and Kwanghoon Sohn. Deep monocular depth estimation via integration of global and local predictions. *IEEE Transactions on Image Processing*, PP:1–1, 05 2018. 4, 5
- [12] Johannes Kopf, Kevin Matzen, Suhib Alsisan, Ocean Quigley, Francis Ge, Yangming Chong, Josh Patterson, Jan-Michael Frahm, Shu Wu, Matthew Yu, Peizhao Zhang, Zijian He, Peter Vajda, Ayush Saraf, and Michael Cohen. One shot 3d photography. 39(4), 2020. 2, 3
- [13] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 239–248. IEEE, 2016. 2
- [14] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T. Freeman. Learning the depths of moving people by watching frozen people. *CoRR*, abs/1904.11111, 2019. 3, 5, 6
- [15] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 4, 5, 6
- [16] Julian Lienen and Eyke Hullermeier. Monocular depth estimation via listwise ranking using the plackett-luce model. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14590–14599, 2021. 1
- [17] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond, 2019. 5
- [18] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [19] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 1, 2, 5, 7
- [20] Vladimir Nekrasov, Thanuja Dharmasiri, Andrew Spek, Tom Drummond, Chunhua Shen, and Ian Reid. Real-time joint semantic segmentation and depth estimation using asymmetric annotations. *arXiv preprint arXiv:1809.04766*, 2018. 2
- [21] Vladimir Nekrasov, Chunhua Shen, and Ian Reid. Light-weight refinenet for real-time semantic segmentation, 2018. 2, 4
- [22] Simon Niklaus, Long Mai, Jimei Yang, and Feng Liu. 3d ken burns effect from a single image. *ACM Transactions on Graphics*, 38(6):184:1–184:15, 2019. 5
- [23] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5
- [24] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, October 2021. 3, 5, 6
- [25] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer, 2019. 1, 2, 3, 4, 5, 6
- [26] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1
- [27] A. Saxena, M. Sun, and A. Y. Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):824–840, 2009. 2
- [28] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. *Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1
- [29] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2019. 4, 8
- [30] Allen Wang. Efficientnet-lite. <https://github.com/tensorflow/tpu/tree/master/models/official/efficientnet/lite>, 2020. 4, 8
- [31] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes, 2019. 1, 2, 3, 4, 5
- [32] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2, 5
- [33] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 5
- [34] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2016. 8
- [35] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (CVPR)*, 2021. 1, 2, 3, 4, 5, 6, 8
- [36] Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 5

- [37] Michael R. Zhang, James Lucas, Geoffrey Hinton, and Jimmy Ba. Lookahead optimizer: k steps forward, 1 step back, 2019. [5](#)