

Learning a Depth Covariance Function

Eric Dexheimer and Andrew J. Davison
 Dyson Robotics Lab, Imperial College London
 {e.dexheimer21, a.davison}@imperial.ac.uk

Abstract

We propose learning a depth covariance function with applications to geometric vision tasks. Given RGB images as input, the covariance function can be flexibly used to define priors over depth functions, predictive distributions given observations, and methods for active point selection. We leverage these techniques for a selection of downstream tasks: depth completion, bundle adjustment, and monocular dense visual odometry.

1. Introduction

Inferring the 3D structure of the world from 2D images is an essential computer vision task. In recent years, there has been significant interest in combining principled multiple view geometry with data-driven priors. Learning-based methods that predict geometry provide a prior directly over the latent variables, which avoids the ill-posed configurations of traditional methods. However, direct, overconfident priors may prevent realization of the true 3D structure when multi-view geometry is well-defined. For example, data-driven methods have shown tremendous promise in monocular depth estimation, but often lack consistency when fusing information into 3D space.

Thus far, designing a unified framework that combines the best of learning and optimization methods has proven challenging. Recent data-driven methods have attempted to relax rigid geometric constraints by also predicting low-dimensional subspaces or per-pixel uncertainties. Fixing capacity during training is often wasteful or inflexible at test time, while per-pixel residual distributions typically explain away the limitations of the model instead of the relationship between errors. In reality, the 3D world is anywhere from simple to complex, and an ideal system should explicitly adapt its capacity and correlation based on the scene.

In this paper, rather than directly predicting geometry from images, we propose learning a depth covariance function. Given an RGB image, our method predicts how the depths of any two pixels relate. To achieve this, a neural network transforms color information to a feature space, and

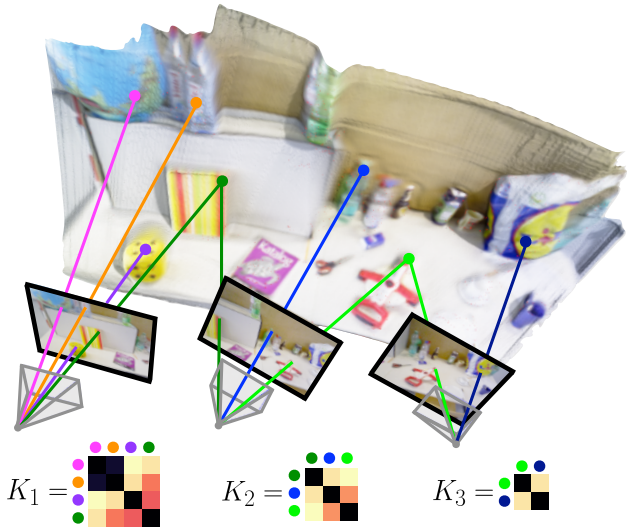


Figure 1. Example monocular reconstruction using the depth covariance for bundle adjustment and dense depth prediction from three seconds (100 frames) of TUM data. Three representative images and the mesh from TSDF fusion of the depth predictions are shown. Each frame leverages the learned covariance function to model geometric correlation between pairs of scene points.

a Gaussian process (GP) models a prior distribution given these features and a base kernel function. The distinction between image processing and the prior enables promoting locality and granting flexible capacity at test time. Locality avoids over-correlating pixels on distinct structures, while adaptive capacity permits tuning the complexity of our subspace to the content of the viewed scene.

Learning this flexible, high-level prior allows for balancing data-driven methods with test-time optimization that can be applied to a variety of geometric vision tasks. Furthermore, the covariance function is agnostic to the 3D representation as it does not directly learn a geometric output. Depth maps may be requested by conditioning on observations, but the prior may also be leveraged for inferring the desired latent 3D representation. In Figure 1, we illustrate depth covariance along with an example of bundle adjustment, dense depth prediction, and multi-view fusion.

In summary, our key contributions are:

- A framework for learning the covariance function by selecting a depth representation, a base kernel function, and a scalable optimization objective
- Application of the prior to depth completion, bundle adjustment, and monocular dense visual odometry

2. Related Work

View-based Priors for Geometric Inference Traditionally, sparse reconstruction methods have ignored per-image correlations [37], [26], while dense methods [27], [13] have used naive priors based on neighboring pixels. With the availability of depth data and the rise of deep learning, a variety of learned depth priors have been proposed. Although monocular depth estimation methods that directly predict geometry [12], [15], [32] have demonstrated great progress, they may predict irreparable mistakes and do not allow for handling ambiguities. Utilizing surface normals may refine these depth estimates [1], but the improvement is still subject to errors from the initial depth prediction. For this reason, we focus on approaches that balance learned priors and test-time optimization. Popular approaches include predicting low-dimensional latent codes [3], [9] and generating a basis of depth maps [42], [16]. While these methods utilize single image information to reduce the dimensionality of depth map inference, they fix capacity during training, produce overly smooth depth maps, and create over-correlated global changes of the depth map. Predicting mesh vertices in the image plane [4] permits distribution of capacity and optimizing depth, but the number of latent variables is still fixed and an explicit mesh representation cannot easily model complex scenes.

Instead of predicting geometry directly, we focus on predicting geometric correlation. By decoupling the pipeline into an image processing network and a GP, the network is not responsible for learning geometry directly and we do not need to fix the geometric capacity. The dimensionality of the subspace may be adapted for representing low-rank scenes more compactly and complex geometry with high-fidelity. Lastly, we use a covariance function with locality to bias learning towards local appearance information.

Learning Covariance Functions Choosing covariance functions and performing model selection for GPs is a well-studied topic [34]. While stationary kernels have been explored for merging LiDAR observations with monocular depth estimates [49], more expressive kernels are required for sparse observations. Nonstationary covariance functions using local Gaussian parameterizations [29] have shown potential in robotic terrain mapping [22]. However, since GP model selection occurs per data example, optimizing hyperparameters is challenging, and novel nonstationary kernels have been proposed to limit flexibility [5].

Rather than performing model selection per image, we are inspired by deep kernel learning (DKL) [47], which leverages training data to predict kernel hyperparameters. Alternatively, this may be viewed as meta-learning, where each task is an RGB-depth example [31]. To mitigate over-correlation when using stationary kernels in feature space as proposed in DKL [28], we utilize local information in pixel space [29] to model surface geometry. While deep nonstationary kernel regression has been explored for depth completion [23], GPs balance data fit and model complexity during training, and the uncertainty estimates are conducive to decision-making and inference in optimization frameworks.

Residual Covariances in Computer Vision Uncertainties in machine learning are often divided between two types: model uncertainty and data uncertainty [19]. For vision problems, model uncertainty is often ignored due to the availability of data and for tractable optimization. Residual uncertainty is usually predicted from a network as a per-pixel variance due to high-dimensionality, but the assumption of independence ignores the correlation present in images. For example, variational autoencoders (VAEs) [20] with a diagonal likelihood output overly smooth mean predictions and unnatural samples with salt-and-pepper noise. To introduce correlation in the likelihood, structured uncertainty prediction networks (SUPNs) predict a full information matrix by defining a graph topology consisting of specific neighbors [11]. Since this matrix is sparse, it can be efficiently inverted to obtain the covariance matrix. This has also been leveraged to distill monocular depth prediction ensembles into a single SUPN [40].

Similar to SUPNs, we move beyond diagonal covariance approximations that lack correlation. However, we learn a covariance function which does not require a predefined graph topology and allows for long-range correlations beyond neighboring pixels. Furthermore, the marginal distribution of any set of variables may be examined without constructing the full joint distribution, which is of great interest for compact inference in geometric vision.

3. Learning a Depth Covariance Function

Given an image, we model a distribution over possible depth functions via a Gaussian process (GP) [34]. With the input domain being normalized image coordinates $\mathbf{x}_i = (u_i, v_i)$ where $u_i, v_i \in [-1, 1]$, the outputs for any finite set of inputs are jointly Gaussian. The GP is then defined by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (1)$$

For learning the parameters of this GP, we need to select the representation of depth, the mean and covariance functions, and the optimization objective. We will outline our decisions in the following sections, but there exist many ways to define this prior.

3.1. Depth Representation

When predicting depth, we often want to attenuate errors further from the camera view, so that nearby structures are prioritized. Examples of such representations include inverse depth [7], disparity [32], and log-depth [12]. We select log-depth for two reasons. First, the log-normal distribution is more suitable for skewed distributions [35], as the range of the GP is unbounded. Representing depth or inverse depth with a normal distribution would require *ad hoc* truncation, as depth functions could move behind the camera. Second, we represent depth with relative scale, as the absolute scale ambiguity in monocular images accounts for much of the error in learned depth prediction [12]. We can then focus on learning the relationship between predictions. In the log-depth formulation, a constant mean $m(\mathbf{x}_i) = m$ directly corresponds to a scale variable. Given a log-depth observation $\mathbf{y}_i = \log \mathbf{d}_i$, we can adjust its scale via this mean:

$$e^{\mathbf{y}_i - m} = e^{\mathbf{y}_i} e^{-m} = C e^{\mathbf{y}_i} = C \mathbf{d}_i. \quad (2)$$

In other words, the median of the depth distribution is controlled by the mean of the log-normal distribution. During training, similar to [12], we may find the optimal scale m that minimizes our data loss. In addition, the log-depth representation allows covariance to be in relative scale. At test time, the scale may be fixed if known, or jointly optimized with other variables. Learning a more expressive mean function per image would be useful in the absence of any depth observations, but introduces many degrees of freedom. For GPs, generalization is largely dependent on the covariance function, so we leave the mean to be a constant per image.

3.2. Covariance Function

We would like to learn the parameters of a depth covariance function using pairs of RGB and depth images. By definition, the covariance function must be positive semidefinite (PSD). In our framework, an RGB image is fed into a convolutional neural network (CNN), which outputs features for a base covariance function. We model the depth observations as being jointly Gaussian so that the CNN and base kernel hyperparameters may be jointly learned.

For the base covariance function, we would like to avoid over-correlation of independent structures. To achieve local influence, we use the family of nonstationary kernels described in [29]. Each point \mathbf{x}_i in an image can be viewed as a Gaussian distribution with a 2D kernel matrix Σ_i . Then, the covariance between two points is the convolution of these two densities over the input domain with a normalizing constant. Alternatively, this can be viewed as the similarity between distributions via the Bhattacharyya kernel [18]. Given an isotropic covariance function $R^S(\sqrt{q})$,




Figure 2. Visualizing our depth covariance function: for every pixel of an input image, the trained network predicts a 2D kernel matrix. Here we show the covariance function between pairs of pixels in both matrix form and as edges in a graph, with the line thickness representing the magnitude of covariance.

the closed-form expression is:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \frac{|\Sigma_i|^{1/4} |\Sigma_j|^{1/4}}{|\Sigma_i + \Sigma_j|^{1/2}} R^S(\sqrt{q}), \quad (3)$$

$$q = (\mathbf{x}_i - \mathbf{x}_j)^T (\Sigma_i + \Sigma_j)^{-1} (\mathbf{x}_i - \mathbf{x}_j), \quad (4)$$

where σ_f^2 is a learnable signal variance. To better handle discontinuities, we select the Matérn function as our isotropic covariance function. With the base covariance function requiring a 2D PSD matrix for each pixel, the CNN outputs three channels (c_1, c_2, c_3), which are parameterized given the positive diagonal and determinant constraints:

$$\Sigma_i = \begin{bmatrix} e^{c_1} & \tanh(c_3) \sqrt{e^{c_1} * e^{c_2}} \\ \tanh(c_3) \sqrt{e^{c_1} * e^{c_2}} & e^{c_2} \end{bmatrix}. \quad (5)$$

An example of five kernel matrices and the marginal covariance in matrix and graph form is shown in Figure 2. We use a UNet architecture [36] and output features at different levels for a multi-scale loss. Each scale also has its own GP hyperparameters, signal variance σ_f^2 and noise variance σ_n^2 . Depth maps are coarsened for lower-levels, so finer levels are weighted higher during the total loss calculation.

While a more global covariance function, such as a squared-exponential over output features, would provide additional flexibility, we avoid this for two reasons. First, changes in depths on one part of the scene may have significant influence on a completely unrelated part, as seen in [3]. Second, deep kernel learning using this setup is unstable and biased towards over-correlating the input domain [28]. By directly using pixel coordinates in our covariance function, we restrict changes in geometry to be local and bias the network to learn local appearance information.

3.3. Optimization Objective

In GP literature, model selection is performed by minimizing the negative log marginal likelihood (NLML)

$$-\log p(\mathbf{y}|X) = \frac{1}{2}(\mathbf{y} - m)^T (K_{\text{ff}} + \sigma_n^2 I)^{-1} (\mathbf{y} - m) + \frac{n}{2} \log 2\pi + \frac{1}{2} \log |K_{\text{ff}} + \sigma_n^2 I|, \quad (6)$$

where the matrix K_{ff} is a PSD matrix with entries defined by the covariance function. The first term is often referred to as the “data fit”, while the third term is the “complexity penalty”. In our scenario, the complexity term is minimized by correlating points, while the data term ensures the ground-truth depth map is plausible given the mean and predicted covariance.

However, the marginal likelihood for a GP requires a $O(N^3)$ matrix inversion. For pixels in an image, this becomes intractable. Since depth images are relatively low-rank, we use a sparse GP approximation [2] during training by randomly sampling inducing points for a Nyström approximation to the full covariance matrix. Given the covariance matrix for the inducing points K_{uu} and the cross-covariance between the inducing points and the entire domain K_{uf} , we have:

$$K_{\text{ff}} \approx \tilde{K}_{\text{ff}} \triangleq K_{\text{fu}} K_{\text{uu}}^{-1} K_{\text{uf}}. \quad (7)$$

Specifically, we use the variational free energy (VFE) framework [45], which defines our training loss as

$$\mathcal{F} = \frac{1}{2}(\mathbf{y} - m)^T (\tilde{K}_{\text{ff}} + \sigma_n^2 I)^{-1} (\mathbf{y} - m) + \frac{n}{2} \log 2\pi + \frac{1}{2} \log |\tilde{K}_{\text{ff}} + \sigma_n^2 I| + \frac{1}{2\sigma_n^2} \text{tr}(K_{\text{ff}} - \tilde{K}_{\text{ff}}). \quad (8)$$

The first three terms are the same as the original NLML, but using the approximate covariance, while the last term penalizes the conditional variances at all inputs given the inducing points, which only requires the diagonal of the full covariance matrix. Note that we also assume homoscedastic observation noise variance σ_n^2 across the entire dataset.

4. Leveraging a Depth Covariance Function

Given an image and its corresponding CNN output maps, we may leverage the covariance function by defining a prior over depth functions, conditioning on depth observations to yield a predictive distribution, and actively sampling pixel locations that minimize the variance of depth predictions.

4.1. Gaussian Process Prior

Given a finite set of input points X_* , we may define a Gaussian prior over their log-depths:

$$\mathbf{f}_* \sim \mathcal{N}(m, K(X_*, X_*)). \quad (9)$$

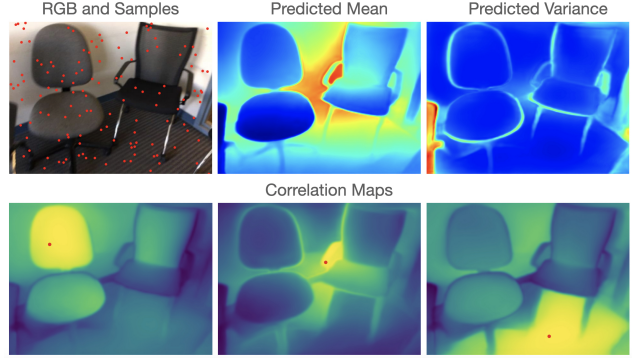


Figure 3. Conditioning example for 128 samples. The posterior variance is high around edges and in areas lacking samples. The columns of K_{fn} , or correlation maps, are shown for select points.

The GP defines a prior over functions, and for any number of input points, we obtain a Gaussian prior. For geometric vision tasks, this prior can be leveraged for anything from sparse to dense methods. Due to the marginalization property of GPs, adding points to the indexing set will not change the marginal distribution of the existing set. Alternatively, the prior may be viewed as an image-conditioned regularizer of any desired latent geometry parameterization.

4.2. Predictive Distribution

In many cases, log-depth observations \mathbf{y} , such as from RGB-D sensors or LiDAR, may be available. We can explicitly condition the prior on these observations to obtain a posterior distribution. The predictive mean \mathbf{f}_* and covariance Σ_* given N samples is:

$$\mathbf{f}_* = m + K_{\text{fn}}(K_{\text{nn}} + \sigma_n^2 I)^{-1} (\mathbf{y} - m), \quad (10)$$

$$\Sigma_* = K_{\text{ff}} - K_{\text{fn}}(K_{\text{nn}} + \sigma_n^2 I)^{-1} K_{\text{nf}}. \quad (11)$$

The predictive mean is a linear function in terms of the observations \mathbf{y} , so test-time inference of latent depths is efficient. If a full covariance is not required, only certain blocks or per-pixel variances need to be calculated. Note that the covariance depends only on the observation coordinates and not the observed values. We visualize the predictive mean, variance, and correlation maps for an example in Figure 3.

4.3. Active Point Selection

For dense reconstruction, it is beneficial to construct a compact representation that can achieve high-fidelity results. As mentioned previously, the predictive covariance depends only on the RGB image and locations of depth observations, but not on the observations themselves. By viewing the CNN as a meta-learned initialization of the non-stationary kernel parameters, we may use 2D observations as a proxy for the complexity of 3D geometry. Active selection of informative pixels up to a desired variance permits




Figure 4. Qualitative comparison of random and active sampling given 32 sample selections. Random sampling misrepresents depth at the top of the image, while active sampling focuses on the chair geometry and avoids redundant samples on the floor.

representing less complex scenes with fewer samples and allocating capacity towards high-frequency geometry.

Inspired by sensor placement literature [17], an entropy-based criterion is used to select informative pixels. In the greedy-case, this simplifies to selecting the input point with the highest conditional variance at each step. Since this requires computing the conditional variance for each newly added point, we leverage incremental updates to the variance and Cholesky factorization of the training covariance matrix [33]. By decoupling the neural network and GP, we do not require any additional network passes.

A qualitative example of active sampling is shown in Figure 4. Random sampling severely misrepresents depth of table and chairs, which all appear at similar depths. Active sampling focus samples around the thin chair edge near the image, while also avoiding oversampling on the floor. Furthermore, active sampling avoids missing the top section of the image, so that the table and chairs are well-represented.

5. Applications

We apply depth covariance to three fundamental geometric vision tasks: depth completion, bundle adjustment, and monocular dense visual odometry (DVO). For depth completion, we use the NYUv2 benchmark [39] and the train-test splits from [25]. For bundle adjustment and DVO, we train the covariance function on ScanNet [10], and evaluate on the TUM RGB-D dataset [41]. For the UNet, we use an input size of 256×192 , 16 channels after the first layer, 5 downsampling steps, and 4 output levels. This results in roughly 9 million parameters.

5.1. Depth Completion

Depth completion is a fundamental task that will be leveraged for additional applications. We may directly condition on sparse observations as described in Section 4.2 to obtain a dense depth map and uncertainties.

5.1.1 Completion Accuracy and Error

Since our method does not predict a specific instance of geometry, we report errors with respect to the GP posterior mean. We compare to the foundational Sparse-to-Dense (S2D) [25], deep kernel regression with (KernelNet+R) and without refinement (KernelNet) [23], and a recent state-of-the-art network RigNet [48]. Results on the validation set with 500 random samples are shown in Table 1.

Method	Error (\downarrow) RMSE	Accuracy (\uparrow)				
		$\delta_{1.02}$	$\delta_{1.05}$	$\delta_{1.10}$	$\delta_{1.25}$	$\delta_{1.25^2}$
S2D [25]	0.204	-	-	-	97.8	99.6
KernelNet [23]	0.198	65.5	82.9	91.6	97.7	99.8
KernelNet+R [23]	0.111	84.8	94.1	97.4	99.3	99.9
RigNet [48]	0.090	-	-	-	99.7	99.9
Ours	0.157	81.2	92.6	97.1	99.4	99.9

Table 1. Depth completion on NYUv2 with 500 sampled points.

KernelNet also predicts three feature maps for 2D covariance parameters, but we achieve better performance under the GP framework. We achieve comparable accuracy to state-of-the-art methods KernelNet+R and RigNet while using fewer parameters. We also do not convert the problem into classification [23] or have complex forward passes with iterative layers [48]. The depth covariance outperforms methods with similar UNet architectures in RMSE. While the error is not state-of-the-art, we use a lightweight network and do not train specifically for the single task of depth completion with the number of samples known *a priori*.

We also explore varying sparsity, as depth completion networks are trained for a specific number of samples. A comparison of RMSE for a varying number of test samples is shown in Table 2. Depth covariance outperforms traditional depth completion methods for sparse inputs, and is competitive with SpAgNet [8] which is designed specifically for these cases and contains an additional non-local spatial propagation layer. We note that for error metrics, we simply use the posterior mean for comparison. However, the GP provides a distribution over depths, not just a single instance, and we will demonstrate additional capabilities such as active decision-making and inferring latent geometry when direct observations are not present.

# Samples	5	50	100	200
CSPN [6]	2.063	0.884	0.388	0.177
NLSPN [30]	1.033	0.423	0.246	0.142
SpAgNet [8]	0.467	0.272	0.209	<u>0.155</u>
Ours	<u>0.717</u>	<u>0.298</u>	<u>0.236</u>	0.193

Table 2. Depth completion RMSE (m) on NYUv2 with a varying number of input samples. The best result is in bold, while the second best is underlined.




Figure 5. Calibration plots of varying posterior marginal covariance block dimensions D on NYUv2 depth completion. The ideal calibration is $y = x$, where the observed confidence matches expected confidence. The region above the line indicates model under-confidence, while the area below signals over-confidence.

5.1.2 Posterior Uncertainty

Calibrated uncertainties are beneficial for balancing multiple views and sensors in optimization. Under-confidence does not account for the full information given by the constraints, while over-confidence may bias the solution. Given that the GP provides uncertainties in addition to the mean, we evaluate the calibration properties for the depth completion setup from Section 5.1.

Since most methods typically estimate per-pixel observation noise, we assess the effect of including off-diagonal terms. Calibration plots measure how well the model’s predicted confidence matches the expected confidence. To extend calibration plots for regression beyond variance [21], we extract marginal covariances of varying dimension D and calculate Mahalanobis distances using this block-diagonal approximation to the full covariance matrix. We then count the frequency of distances below varying chi-square thresholds that define the expected confidence. In the variance case with $D = 1$, residuals do not affect each other. We plot the results in Figure 5 along with the RMSE against the ideal calibration, where the empirical confidence matches the observed confidence. Including more off-diagonal terms improves the model calibration, and significantly reduces the under-confidence of the model.

5.1.3 Active Sampling Evaluation

While the location of depth samples is often not a free variable, such as from LiDAR or depth sensors, in monocular vision, we may wish to actively estimate the depths of certain pixels to improve downstream tasks such as dense reconstruction. We investigate whether the meta-learned covariance parameters may be leveraged for selecting more




Figure 6. Mean depth completion RMSE for random and active sampling with a varying number of samples. Green line shows mean percent improvement of active over random sampling.

informative pixels. We compare randomly sampling pixels uniformly against the greedy conditional variance as described in Section 4.3. The effect on depth error and the mean percentage improvement for a varying number of samples is shown in Figure 6. Note that the greedy active sampling consistently outperforms random sampling. For a large number of samples, the relative improvement of active sampling decreases as there is a sufficient number of observations. For a few samples, only coarse depth structure is retained, and the greedy selection occasionally encourages samples near the image boundaries. Beyond greedy selection, other sampling methods that explicitly minimize uncertainty over the whole domain may demonstrate improved performance at the expense of computational cost.

5.2. Bundle Adjustment

Bundle adjustment is the foundation for many vision pipelines, and when successful, it produces very accurate camera poses. However, the standard formulation of assuming independence between observations proves challenging for monocular SLAM systems, which may fail during initialization or when translation is negligible compared to rotation. Many realistic scenarios require robust and fast initialization with little motion. Therefore, we evaluate the use of the depth covariance in small baseline scenarios. We also show that the depth covariance is not limited to 2D depth map inference, but can also be used to infer 3D structure.

Bundle adjustment jointly optimizes camera poses and point landmarks given pixel correspondences. Traditionally, the cost involves the sum of reprojection errors, as well as pose and scale priors on the first pose to fix gauge freedom which we omit for brevity. We also add our depth prior factor per camera, so for landmarks \mathbf{P} in the world frame, and poses T_{cW} for $c = 1, \dots, C$, we have:

$$\begin{aligned}
 E = & \sum_c \sum_i \|\pi(\mathbf{T}_{cW}, \mathbf{P}_i) - \mathbf{x}_{c,i}\|_{\sigma_i^2 I}^2 \\
 & + \sum_c \|[T_{cW} \mathbf{P}_c]_z - m_c\|_{K_c}^2, \quad (12)
 \end{aligned}$$




Figure 7. Factor graphs for bundle adjustment and monocular dense visual odometry. Pose and scale priors to constrain gauge freedom are omitted for simplicity.

where each frame also has a scale variable m_c and depth covariance K_c for vectorized observed landmarks \mathbf{P}_c in that frame. The projection function π projects 3D landmarks into the image plane. To reduce the influence of outliers, we use the Huber cost function for reprojection errors. The least-squares cost is optimized using Levenberg-Marquardt via GTSAM [14]. The corresponding factor graph is shown on the left of Figure 7.

We divide all of the Freiburg 1 sequences from the TUM RGB-D dataset into small baseline 5-frame windows. This yields 1544 sequences, emulating initializations to visual odometry systems. For the frontend, Shi-Tomasi corners [38] are detected and tracked using Lucas-Kanade tracking [24], and outliers are filtered out via essential matrix RANSAC. Poses are initialized using the motion capture data closest to the current RGB frame timestamps, while landmarks are triangulated if sufficiently constrained. We optimize each sequence with and without the GP depth prior to evaluate its effect. Using the corresponding depth frames, we compare the error of the reprojected sparse landmarks where valid depths occur, as well as the dense depth map error by conditioning on the sparse landmarks. To handle monocular scale ambiguity, we align depth maps with the optimal scale before computing the error. A boxplot of errors between the two methods is shown in Figure 8. By exploiting the correlation between observations, significantly more coherent geometry can be generated, as visualized by an example in Figure 9. Note that the prior is able to achieve consistent geometry despite the low baseline and only optimizing sparse landmarks. The shadows reveal the complete chair legs and table. We also show a longer example with over three seconds of data fused into a TSDF in Figure 1.

5.3. Monocular Dense Visual Odometry

We propose a monocular dense visual odometry system based on the depth covariance and conditioning on per-




Figure 8. Boxplot of sparse and dense depth errors across all 5-frame windows, with and without the GP depth prior included in the optimization. The median RMSE is written next to the plots.




Figure 9. Qualitative example of small baseline bundle adjustment with and without the GP depth factor. Both examples fuse 5 predicted depth maps densified from the optimized landmarks into a TSDF, and are visualized with identical lighting.

frame latent depth points. We perform sliding window optimization over a fixed number of keyframes, where poses, scales, and inducing depths are optimized per-frame. To regularize the depths, we may condition on the current estimates of the inducing depths as in Section 4.2, project depths into neighboring keyframes, and apply relative photometric error constraints. The photometric error for pixels n across all directed edges E is:

$$E = \sum_{i,j \in E} \sum_n \|I_i(\mathbf{x}_n) - I_j(w(\mathbf{x}_n, \mathbf{T}_i, \mathbf{T}_j, \mathbf{y}_i, m_i))\|_{\sigma_I^2}^2, \quad (13)$$

where I is the image intensity, \mathbf{x} are pixel coordinates, w is the warping function that queries a depth map and projects the corresponding point into the neighboring image, \mathbf{T} are camera poses, \mathbf{y} are estimated log-depth observations, and m is the mean log depth. The corresponding factor graph with the depth priors included is shown in the right of Figure 7. To initialize a new keyframe, we first run the RGB through the network to get the covariance feature maps and sample inducing points. We perform active sampling since this is independent of whether log-depths are known as mentioned in 4.3, and stop sampling if a maximum variance threshold is achieved. The log-depths \mathbf{y} of a new keyframe are initialized by minimizing the least-squares cost of the




Figure 10. Example depth maps of keyframes while running the dense odometry on ScanNet validation sequences and TUM.

prior at the sampled locations and a predictive data term over the reprojected depths \mathbf{d} from the latest keyframe:

$$\mathbf{y}_{\text{init}} = \arg \min_{\mathbf{y}} \|\mathbf{y} - \mathbf{m}\|_K^2 + \|\log \mathbf{d} - \mathbf{f}_*\|_{\text{diag}[\Sigma_*]}. \quad (14)$$

When the window size is full, the last keyframe is dropped, and scale and pose priors are added to the current estimates of the oldest keyframe still in the window. Initialization of the system is achieved via two-frame SfM, where the first frame’s depth map is optimized along with the relative pose to another frame. Tracking is performed using coarse-to-fine photometric tracking against the last keyframe. We also optimize affine brightness parameters for varying exposure.

For quantitative evaluation, we compare our odometry to other learning-based dense systems on Freiburg 1 sequences in the TUM dataset. Since our method is purely a sliding-window odometry system, we compare against the comparable monocular methods evaluated in [44] that lack global optimization and bundle adjustment. We omit the floor sequence which drops a significant number of frames. The absolute trajectory errors (ATE) are shown in Table 3. Since our dense odometry does not have global scale, we align to the ground-truth trajectory’s scale.

Despite the simplicity of our system, it is first or second in 6 of 8 sequences, and achieves the best mean error. We only use photometric error compared to the additional depth and keypoint factors in DeepFactors. While we lack many features in our simple sliding window formulation, we believe that including additional components will further improve the results. Qualitative examples of inferred depth maps from the odometry on ScanNet validation sequences and TUM are shown in Figure 10. Note that our representation can predict a variety of thin structures and regularize flat ones, even in the case of high-frequency texture.

Odometry runs at an average of 7-10 Hz on 256×192 images with 8 keyframes on a single RTX 3080. The code is not optimized and alternates between tracking and mapping in a single thread, so we believe significant speed-ups are possible given additional optimization or compute.

Sequence	TartanVO [46]	DeepV2D [43]	DeepFactors [9]	Ours
360	0.178	0.243	<u>0.159</u>	0.128
desk	<u>0.125</u>	0.166	0.170	0.056
desk2	<u>0.122</u>	0.379	0.253	0.048
plant	0.297	0.203	0.305	<u>0.261</u>
room	0.333	0.246	0.364	<u>0.257</u>
rpy	<u>0.049</u>	0.105	0.043	0.052
teddy	<u>0.339</u>	0.316	0.601	0.475
xyz	0.062	0.064	0.035	<u>0.056</u>
mean	<u>0.188</u>	0.215	0.241	0.167

Table 3. Absolute trajectory error (m) on TUM Freiburg1. The best result is in bold, while the second best is underlined.

6. Conclusion and Future Work

We have proposed learning a depth covariance function, and provided selections of the depth representation, the base covariance function, and the optimization objective. While we selected a specific nonstationary kernel function, exploration of alternatives is a valuable direction. There is great potential to draw additional connections to the literature on Bayesian methods and GPs. For example, geometry is low-rank and spatially local, so active selection and sparse kernels will permit scalability to higher resolutions and larger capacities. Lastly, we demonstrated the utility of the covariance function on a selection of geometric vision tasks: depth completion, bundle adjustment, and monocular dense visual odometry. Beyond these applications, the prior can be tightly integrated into existing geometric vision frameworks and serve as a foundation for novel formulations.

7. Acknowledgements

Research presented in this paper has been supported by Dyson Technology Ltd. We would like to thank Tristan Laidlow, Hidenobu Matsuki, Riku Murai, and members of the Dyson Robotics Lab for insightful discussions.

References

- [1] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Iron-Depth: Iterative refinement of single-view depth using surface normal and its uncertainty. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2022. 2
- [2] Matthias Bauer, Mark van der Wilk, and Carl Edward Rasmussen. Understanding probabilistic sparse Gaussian process approximations. In *Neural Information Processing Systems (NeurIPS)*, 2016. 4
- [3] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J. Davison. CodeSLAM - learning a compact, optimisable representation for dense visual SLAM. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3
- [4] Michael Bloesch, Tristan Laidlow, Ronald Clark, Stefan Leutenegger, and Andrew Davison. Learning meshes for dense visual SLAM. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [5] Weizhe Chen, Roni Kharon, and Lantao Liu. AK: Attentive kernel for information gathering. In *Proceedings of Robotics: Science and Systems (RSS)*, 2022. 2
- [6] Xinjing Cheng, Peng Wang, and Ruigang Yang. Depth estimation via affinity learned with convolutional spatial propagation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 5
- [7] Javier Civera, Andrew J. Davison, and J. M. Martínez Montiel. Inverse depth parametrization for monocular SLAM. *IEEE Transactions on Robotics (T-RO)*, 2008. 3
- [8] Andrea Conti, Matteo Poggi, and Stefano Mattoccia. Sparsity agnostic depth completion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023. 5
- [9] Jan Czarnowski, Tristan Laidlow, Ronald Clark, and Andrew J. Davison. DeepFactors: Real-time probabilistic dense monocular SLAM. *IEEE Robotics and Automation Letters (RA-L)*, 2020. 2, 8
- [10] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [11] Garoe Dorta, Sara Vicente, Lourdes Agapito, Neill D. F. Campbell, and Ivor Simpson. Structured uncertainty prediction networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [12] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Neural Information Processing Systems (NeurIPS)*, 2014. 2, 3
- [13] Jakob Engel, Thomas Schöps, and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 2
- [14] Frank Dellaert et al. borglab/gtsam, May 2022. 7
- [15] Clement Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [16] Benjamin Graham and David Novotny. RidgeSfM: Structure from motion via robust pairwise matching under depth uncertainty. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2020. 2
- [17] Carlos Guestrin, Andreas Krause, and Ajit Paul Singh. Near-optimal sensor placements in Gaussian processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2005. 5
- [18] Tony Jebara, Risi Kondor, and Andrew Howard. Probability product kernels. *The Journal of Machine Learning Research*, 2004. 3
- [19] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [20] Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 2
- [21] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018. 6
- [22] Tobias Lang, Christian Plagemann, and Wolfram Burgard. Adaptive non-stationary kernel regression for terrain modeling. In *Proceedings of Robotics: Science and Systems (RSS)*, 2007. 2
- [23] Lina Liu, Yiyi Liao, Yue Wang, Andreas Geiger, and Yong Liu. Learning steering kernels for guided depth completion. *IEEE Transactions on Image Processing*, 2021. 2, 5
- [24] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1981. 7
- [25] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 5
- [26] Raúl Mur-Artal and Juan D. Tardós. ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras. *IEEE Transactions on Robotics (T-RO)*, 2017. 2
- [27] Richard A. Newcombe, Steven J. Lovegrove, and Andrew J. Davison. DTAM: Dense tracking and mapping in real-time. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2011. 2
- [28] Sebastian W. Ober, Carl E. Rasmussen, and Mark van der Wilk. The promises and pitfalls of deep kernel learning. In *Conference on Uncertainty in Artificial Intelligence*, 2021. 2, 3
- [29] Christopher Paciorek and Mark Schervish. Nonstationary covariance functions for Gaussian process regression. In *Neural Information Processing Systems (NeurIPS)*, 2003. 2, 3
- [30] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 5

- [31] Massimiliano Patacchiola, Jack Turner, Elliot J. Crowley, Michael O' Boyle, and Amos J Storkey. Bayesian meta-learning for the few-shot setting via deep kernels. In *Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [32] Rene Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2022. 2, 3
- [33] Ananth Ranganathan, Ming-Hsuan Yang, and Jeffrey Ho. Online sparse Gaussian process regression and its applications. *IEEE Transactions on Image Processing*, 2011. 5
- [34] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2005. 2
- [35] Jean-Michel M. Rendu. Normal and lognormal estimation. *Journal of the International Association for Mathematical Geology*, 1979. 3
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2015. 3
- [37] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [38] J. Shi and C. Tomasi. Good features to track. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994. 7
- [39] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. 5
- [40] Ivor J. A. Simpson, Sara Vicente, and Neill D. F. Campbell. Learning structured Gaussians to approximate deep ensembles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [41] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of RGB-D SLAM systems. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2012. 5
- [42] Chengzhou Tang and Ping Tan. BA-Net: Dense bundle adjustment networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 2
- [43] Zachary Teed and Jia Deng. DeepV2D: Video to depth with differentiable structure from motion. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 8
- [44] Zachary Teed and Jia Deng. DROID-SLAM: Deep visual SLAM for monocular, stereo, and RGB-D cameras. In *Neural Information Processing Systems (NeurIPS)*, 2021. 8
- [45] Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2009. 4
- [46] Wenshan Wang, Yaoyu Hu, and Sebastian Scherer. TartanVO: A generalizable learning-based VO. In *Conference on Robot Learning (CoRL)*, 2021. 8
- [47] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016. 2
- [48] Zhiqiang Yan, Kun Wang, Xiang Li, Zhenyu Zhang, Jun Li, and Jian Yang. RigNet: Repetitive image guided network for depth completion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 5
- [49] Sungho Yoon and Ayoung Kim. Balanced depth completion between dense depth inference and sparse range measurements via KISS-GP. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2020. 2