

FlowMap: High-Quality Camera Poses, Intrinsics, and Depth via Gradient Descent

Cameron Smith*, David Charatan*,
Ayush Tewari, and Vincent Sitzmann

MIT CSAIL

Abstract. This paper introduces FlowMap, an end-to-end differentiable method that solves for precise camera poses, camera intrinsics, and per-frame dense depth of a video sequence. Our method performs per-video gradient-descent minimization of a simple least-squares objective that compares the optical flow induced by depth, intrinsics, and poses against correspondences obtained via off-the-shelf optical flow and point tracking. Alongside the use of point tracks to encourage long-term geometric consistency, we introduce differentiable re-parameterizations of depth, intrinsics, and pose that are amenable to first-order optimization. We empirically show that camera parameters and dense depth recovered by our method enable photo-realistic novel view synthesis on 360° trajectories using Gaussian Splatting. Our method not only far outperforms prior gradient-descent based bundle adjustment methods, but surprisingly performs on par with COLMAP, the state-of-the-art SfM method, on the downstream task of 360° novel view synthesis—even though our method is purely gradient-descent based, fully differentiable, and presents a complete departure from conventional SfM.

1 Introduction

Reconstructing a 3D scene from video is one of the most fundamental problems in vision and has been studied for over five decades. Today, essentially all state-of-the-art approaches are built on top of Structure-from-Motion (SfM) methods like COLMAP [56]. These approaches extract sparse correspondences across frames, match them, discard outliers, and then optimize the correspondences’ 3D positions alongside the camera parameters by minimizing reprojection error [56].

This framework has delivered excellent results which underlie many present-day vision applications, and so it is unsurprising that SfM systems have remained largely unchanged in the age of deep learning, save for deep-learning-based correspondence matching [16, 34, 53, 55].

However, conventional SfM has a major limitation: it is not differentiable with respect to its free variables (camera poses, camera intrinsics, and per-pixel depths). This means that SfM acts as an isolated pre-processing step that cannot be embedded into end-to-end deep learning pipelines. A differentiable, self-supervised SfM method would enable neural networks to be trained self-supervised on internet-scale data for a broad class of multi-view geometry prob-

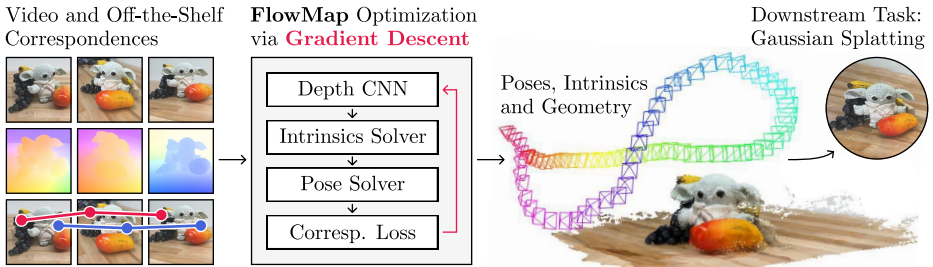


Fig. 1: We present FlowMap, an end-to-end differentiable method that recovers poses, intrinsics, and depth maps of an input video. FlowMap is supervised only with off-the-shelf optical flow and point track correspondences, and optimized per-scene with gradient descent. Gaussian Splats obtained from FlowMap’s reconstructions regularly match or exceed those obtained from COLMAP in quality.

lems. This would pave the way for deep-learning based 3D reconstruction and scene understanding.

In this paper, we present FlowMap, a differentiable and surprisingly simple camera and geometry estimation method whose outputs enable photorealistic novel view synthesis. FlowMap directly minimizes the difference between optical flow that is induced by a camera moving through a static 3D scene and pre-computed correspondences in the form of off-the-shelf point tracks and optical flow. Since FlowMap is end-to-end differentiable, it can naturally be embedded in any deep learning pipeline. Its loss is minimized only via gradient descent, leading to high-quality camera poses, camera intrinsics, and per-pixel depth. Unlike conventional SfM, which outputs sparse 3D points that are each constrained by several views, FlowMap outputs dense per-frame depth estimates. This is a critical advantage in downstream novel view synthesis and robotics tasks. Unlike prior attempts at gradient-based optimization of cameras and 3D geometry [2, 33, 71], we do not treat depth, intrinsics, and camera poses as free variables. Rather, we introduce differentiable feed-forward estimates of each one: depth is parameterized via a neural network, pose is parameterized as the solution to a least-squares problem involving depth and flow, and camera intrinsics are parameterized using a differentiable selection based on optical flow consistency. In other words, FlowMap solves SfM by learning the depth network’s parameters; camera poses and intrinsics are computed via analytical feed-forward modules without free parameters of their own. We show that this uniquely enables high-quality SfM via gradient descent while making FlowMap compatible with standard deep-learning pipelines. Unlike recent radiance-field bundle-adjustment baselines [2, 33], FlowMap does not use differentiable volume rendering, and so it is significantly faster to run, generally reconstructing an object-centric 360° scan in less than 10 minutes.

Through extensive ablation studies, we show that each of FlowMap’s design choices is necessary. On popular, real-world novel view synthesis datasets (Tanks & Temples, Mip-NeRF 360, CO3D, and LLFF), we demonstrate that FlowMap enables photo-realistic novel view synthesis up to full 360° trajectories

using Gaussian Splatting [28]. Gaussian Splats obtained from FlowMap reconstructions far outperform the state-of-the-art gradient-based bundle-adjustment method, NoPeNeRF [2], and those obtained using the SLAM algorithm DROID-SLAM [65], even though both baselines require ground-truth intrinsics. Gaussian Splats obtained from FlowMap are on par with those obtained from COLMAP [56], even though FlowMap only leverages gradient descent, is fully differentiable, and represents a complete departure from conventional SfM techniques.

2 Related Work

Conventional Structure-from-Motion (SfM) and SLAM. Modern SfM methods perform offline optimization using a multi-stage process of descriptor extraction, correspondence estimation, and subsequent incremental bundle adjustment. In bundle adjustment, corresponding 2D pixels are coalesced into single 3D points, and estimated camera parameters are optimized alongside these points’ 3D positions to minimize 3D-to-2D reprojection error. COLMAP [56] is the de-facto standard for accurate, offline camera parameter estimation. Meanwhile, simultaneous localization and mapping (SLAM) usually refers to real-time, online methods. These generally assume that the camera’s intrinsic parameters are known. Similar to SfM, SLAM usually relies on minimizing reprojection error [5, 45, 46, 52], but some methods investigate direct minimization of a photometric error [19, 20]. While deep learning has not fundamentally transformed SfM and SLAM, it has been leveraged for correspondence prediction [12, 38, 43, 48], either via graph neural networks [53] or via particle tracking [17, 24, 79].

FlowMap is a departure from conventional SfM and SLAM techniques. While we rely on correspondence from optical flow and particle tracking, we do not coalesce sets of 2D correspondences into single 3D points. Instead, we use per-frame depth estimates as our geometry representation. Additionally, rather than relying on conventional correspondence matching and RANSAC filtering, we leverage neural point tracking [26] and optical flow estimators [64] to establish correspondence, jointly enabling dense geometry reconstruction without a separate multi-view stereo stage. Finally, FlowMap is end-to-end differentiable and introduces feed-forward estimators of depth, poses, and intrinsics, making it compatible with other learned methods.

Deep-Learning Based SfM. Prior work has attempted to embed the full SLAM pipeline into a deep learning framework [3, 13, 14, 35, 62, 63, 67, 70, 80], usually by training black-box neural networks to directly output camera poses. However, these methods are constrained to short videos of 5 to 10 frames and are not competitive with conventional SLAM and SfM for real-world 3D reconstruction. Bowen et al. [4] elegantly leverage optical flow supervision for self-supervised monocular depth prediction. More recently, DROID-SLAM [65] has yielded high-quality camera poses and depth. However, it requires known intrinsics, is trained fully supervised with ground-truth camera poses, and fails to approach COLMAP on in-the-wild performance and robustness. Concurrent

work to FlowMap explores an end-to-differentiable, point-tracking-based SfM framework [68]. Key differences are that their method is fully supervised with camera poses, point clouds, and intrinsics; requires large-scale, multi-stage training; solves only for sparse depth; and is built around the philosophy of making each part of the conventional SfM pipeline differentiable. Our method is a complete departure from the conventional SfM pipeline—it does not require a training set of known intrinsics, ground-truth poses, or 3D points, and it provides quality gradients for dense depth, poses, and intrinsics. Critically, FlowMap is the first gradient-descent based method to rival the performance of conventional SfM on the novel view synthesis task.

Novel View Synthesis via Differentiable Rendering. Advances in differentiable rendering have enabled photo-realistic novel view synthesis and fine-grained geometry reconstruction using camera poses and intrinsics obtained via SfM [32, 42, 44, 47, 57, 58]. 3D Gaussian Splatting [28] goes further, directly leveraging the 3D points provided by SfM as an initialization. It follows previous methods like [15], which used 3D geometry from depth to supervise neural radiance field (NeRF) reconstructions. We show that when initializing Gaussian Splatting with poses, intrinsics, and 3D points from FlowMap, we generally perform on par with conventional SfM and sometimes even outperform it.

Camera Pose Optimization via Differentiable Rendering. A recent line of work in bundle-adjusting radiance fields [2, 9, 10, 21, 22, 25, 27, 33, 72, 74–76, 78] attempts to jointly optimize unknown camera poses and radiance fields. Jeong et al. [25] additionally solve for camera intrinsic parameters. However, these methods only succeed when given forward-facing scenes or roughly correct pose initializations. More recent work incorporates optical flow and monocular depth priors [2, 36, 40] but requires known intrinsics and only works robustly on forward-facing scenes. Concurrent work [22] somewhat accelerates optimization compared to earlier NeRF-based approaches, but still requires intrinsics and yields significantly worse poses than COLMAP on 360-degree sequences. Further concurrent work proposes real-time SLAM via gradient descent on 3D Gaussians [39], but requires known intrinsics and does not show robustness on a variety of real-world scenes. In contrast, our method is robust and easily succeeds on object-centric scenes where the camera trajectory covers a full 360° of rotation, yielding photo-realistic novel view synthesis when combined with Gaussian Splatting.

Learning Priors over Optimization of NeRF and Poses. Our method is inspired by recent methods which learn priors over pose estimation and 3D radiance fields [8, 23, 31, 59]. However, these approaches require known camera intrinsics, are constrained to scenes with simple motion, and do not approach the accuracy of conventional SfM. Like our method, FlowCam [59] uses a pose-induced flow loss and a least-squares solver for camera pose. However, our method has several key differences: we estimate camera intrinsics, enabling optimization on any raw video; we replace 3D rendering with a simple depth estimator, which

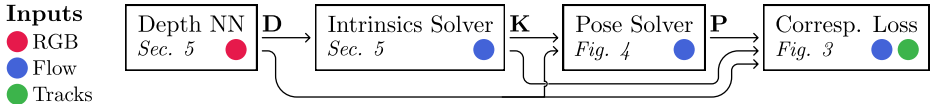


Fig. 2: A FlowMap Forward Pass. Given RGB frames (red), optical flow (blue) and point tracks (green), FlowMap computes dense depth \mathbf{D} , camera poses \mathbf{P} , and intrinsics \mathbf{K} in each forward pass. We obtain depth via a CNN (Sec. 4) and implement differentiable, feed-forward solvers for intrinsics and poses (Sec. 4, Fig. 4). Colored dots indicate which block receives which inputs. FlowMap’s only free parameters are the weights of a depth NN and a small correspondence confidence MLP. These parameters are optimized for each video separately by minimizing a camera-induced flow loss (Fig. 3) via gradient descent, though fully feed-forward operation is possible.

reduces training costs and allows us to reuse pre-trained depth estimators; and we introduce point tracks for supervision to improve global consistency and reduce drift. FlowCam did not approach conventional SfM’s accuracy on real sequences. We demonstrate that optimizing the pose-induced flow objective on a single scene, akin to a test-time optimization, yields pose and geometry estimates which, for the first time, approach COLMAP on full 360° sequences.

3 Supervision via Camera-Induced Scene Flow

Given a video sequence, our goal is to supervise per-frame estimates of depth, intrinsics, and pose using known correspondences. Our method hinges upon the fact that a camera moving through a static scene induces optical flow in image space. Such optical flow can be computed differentiably from any two images’ estimated depths, intrinsics, and relative pose to yield a set of implied pixel-wise correspondences. These correspondences can then be compared to their known counterparts to yield supervision on the underlying estimates.

Consider a 2D pixel at coordinate $\mathbf{u}_i \in \mathbb{R}^2$ in frame i of the video sequence. Using frame i ’s estimated depth \mathbf{D}_i and intrinsics \mathbf{K}_i , we can compute the pixel’s 3D location $\mathbf{x}_i \in \mathbb{R}^3$. Then, using the estimated relative pose \mathbf{P}_{ij} between frames i and j , we can transform this location into frame j ’s camera space. Finally, we can project the resulting point $\mathbf{P}_{ij}\mathbf{x}_i$ onto frame j ’s image plane to yield an implied correspondence $\hat{\mathbf{u}}_{ij}$. This correspondence can be compared to the known correspondence \mathbf{u}_{ij} to yield a loss \mathcal{L} , as illustrated in Fig. 3.

$$\mathcal{L} = \|\hat{\mathbf{u}}_{ij} - \mathbf{u}_{ij}\| \quad (1)$$

Supervision via Dense Optical Flow and Sparse Point Tracks. Our known correspondences are derived from two sources: dense optical flow between adjacent frames and sparse point tracks which span longer windows. Frame-to-frame optical flow ensures that depth is densely supervised, while point tracks minimize drift over time. We compute correspondences from optical flow \mathbf{F}_{ij} via $\mathbf{u}_{ij} = \mathbf{u}_i + \mathbf{F}_{ij}[\mathbf{u}_i]$. Meanwhile, given a query point \mathbf{u}_i , an off-the-shelf point tracker directly provides a correspondence \mathbf{u}_{ij} for any frame j where one exists.

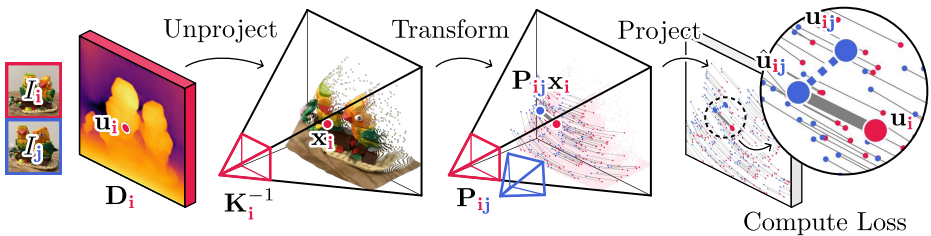


Fig. 3: Camera-Induced Flow Loss. To use a known correspondence u_{ij} to compute a loss \mathcal{L} , we unproject u_i using the corresponding depth map D_i and camera intrinsics K_i , transform the resulting point x_i via the relative pose P_{ij} , reproject the transformed point to yield \hat{u}_{ij} , and finally compute $\mathcal{L} = \|\hat{u}_{ij} - u_{ij}\|$.

Baseline: Pose, Depth and Intrinsics as Free Variables. Assuming one uses standard gradient descent optimization, one must decide how to parameterize the estimated depths, intrinsics, and poses. The simplest choice is to parameterize them as free variables, i.e., to define learnable per-camera intrinsics and extrinsics alongside per-pixel depths. However, this approach empirically fails to converge to good poses and geometry, as shown in Sec. 7.

4 Parameterizing Depth, Pose, and Camera Intrinsics

In this section, we present FlowMap’s feed-forward re-parameterization of depth, pose, and camera intrinsics, which uniquely enables high-quality results when using gradient descent. Later, in Sec. 7, we ablate these parameterizations to demonstrate that they lead to dramatic improvements in accuracy.

Depth Network. If each pixel’s depth were optimized freely, two identical or very similar image patches could map to entirely different depths. We instead parameterize depth as a neural network that maps an RGB frame to the corresponding per-pixel depth. This ensures that similar patches have similar depths, allowing FlowMap to integrate geometry cues across frames: if a patch receives a depth gradient from one frame, the weights of the depth network are updated, and hence the depths of all similar video frame patches are also updated. As a result, FlowMap can provide high-quality depths even for patches which are poorly constrained due to errors in the input flows and point tracks, imperceptibly small motion, or degenerate (rotation-only) motion.

Pose as a Function of Depth, Intrinsics and Optical Flow. Suppose that for two consecutive frames, optical flow, per-pixel depths, and camera intrinsics are known. In this case, the relative pose between these frames can be computed differentiably in closed form. Following the approach proposed in FlowCam [59], we solve for the relative pose that best aligns each consecutive pair of un-projected depth maps. We then compose the resulting relative poses to produce absolute poses in a common coordinate system.

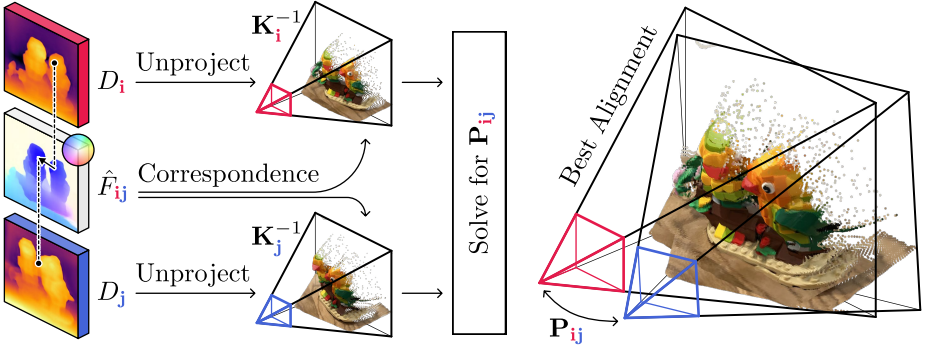


Fig. 4: We solve for the relative poses between consecutive frames using their depth maps, camera intrinsics, and optical flow. To do so, we first unproject their depth maps, then solve for the pose that best aligns the resulting point clouds.

More formally, we cast depth map alignment as an orthogonal Procrustes problem, allowing us to draw upon this problem’s differentiable, closed-form solution [11]. We begin by unprojecting the depth maps D_i and D_j using their respective intrinsics K_i and K_j to generate two point clouds X_i and X_j . Next, because the Procrustes formulation requires correspondence between points, we use the known optical flow between frames i and j to match points in X_i and X_j . This yields X_i^{\leftrightarrow} and X_j^{\leftrightarrow} , two filtered point clouds for which a one-to-one correspondence exists. The Procrustes formulation seeks the rigid transformation that minimizes the total distance between the matched points:

$$P_{ij} = \arg \min_{P \in SE(3)} \|W^{1/2}(X_j^{\leftrightarrow} - PX_i^{\leftrightarrow})\|_2^2 \quad (2)$$

The diagonal matrix W contains correspondence weights that can down-weight correspondences that are faulty due to occlusion or imprecise flow. This weighted least-squares problem can be solved in closed form via a single singular value decomposition [11, 59] which is both cheap and fully differentiable. We further follow FlowCam [59] and predict these weights by concatenating corresponding per-pixel features and feeding them into a small MLP. This MLP’s parameters are the only other free variables of our model. For an overview of the depth map alignment process, see Fig. 4.

Camera Focal Length as a Function of Depth and Optical Flow. We solve for camera intrinsics by considering a set of reasonable candidates K_k , then softly selecting among them. For each candidate, we use our pose solver Eq. 2 to compute a corresponding set of poses, then use the camera-induced flow loss Eq. 1 to compute the loss \mathcal{L}_k implied by K_k and these poses. Finally, we compute the resulting intrinsics K via a softmax-weighted sum of the candidates:

$$K = \sum_k w_k K_k \quad w_k = \frac{\exp(-\mathcal{L}_k)}{\sum_l \exp(-\mathcal{L}_l)} \quad (3)$$

To make this approach computationally efficient, we make several simplifying assumptions. First, we assume that the intrinsics can be represented via a single \mathbf{K} that is shared across frames. Second, we assume that \mathbf{K} can be modeled via a single focal length with a principal point fixed at the image center. Finally, we only compute the soft selection losses on the first two frames of the sequence.

Depth as the Only Free Variable in SfM. FlowMap offers a surprising insight: Given correspondence, SfM can be formulated as solving for per-frame depth maps. FlowMap yields poses and intrinsics in a parameter-free, differentiable forward pass when given correspondences and depths. This means that better initializations of FlowMap’s depth estimator (e.g., from pre-training) will yield more accurate camera parameters (see Fig. 10).

5 Implementation and Optimization Details

FlowMap is optimized on each specific scene, achieving convergence between 500 and 5,000 steps using the Adam [29] optimizer. Though per-scene optimization is key to achieving high accuracy, we find that exploiting FlowMap’s feed-forward nature for pre-training yields an initialization that leads to improved convergence and accuracy, as shown in Fig. 10. We use RAFT [64] and CoTracker V1 [26] to compute the optical flow and point tracks that FlowMap uses as input.

Focal Length Regression. While our soft selection approach robustly yields near-correct focal lengths, its performance is slightly worse compared to well-initialized direct regression. We therefore switch to focal length regression after 1,000 steps, using our softly selected focal length as initialization.

Memory and Time Requirements. FlowMap’s complexity in time and memory is linear with the number of input video frames. During each optimization step, FlowMap recomputes depth for each frame, then derives poses and intrinsics from these depths to generate gradients. In practice, FlowMap optimization for a 150-frame video takes about 20 minutes, with a peak memory usage of about 36 GB. Precomputing point tracks and optical flow takes approximately 2 minutes. Note that FlowMap’s runtime could be reduced by early stopping, and its memory usage could be reduced by performing backpropagation on video subsets during each step, but we leave these optimizations to future work.

Sequence Length and Drift. Since adjacent frames in typical 30 FPS videos usually contain redundant information, we run FlowMap on subsampled videos. We perform subsampling by computing optical flow on the whole video, then selecting frames so as to distribute the overall optical flow between them as evenly as possible. With this strategy, we find that an object-centric, full 360° trajectory as is common in novel view synthesis papers is covered by about 90 frames. We note that FlowMap does not have a loop closure mechanism. Rather, point tracks provide long-range correspondences that prevent the accumulation of drift in long sequences.




Fig. 5: View Synthesis Results. FlowMap’s camera parameters and geometry produce near-photorealistic 3D Gaussian Splatting results on par with COLMAP’s.

6 Results

We benchmark FlowMap via the downstream task of 3D Gaussian reconstruction [28]. This allows us to measure the quality of the camera parameters and geometry (depth maps) it outputs *without having access to ground-truth scene geometry and camera parameters*.

Baselines. We benchmark FlowMap against several baselines. First, we evaluate against COLMAP [56], the state-of-the-art structure-from-motion (SfM)

	MipNeRF 360 (3 scenes)					LLFF (7 scenes)				
Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time (min.) \downarrow	ATE	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time (min.) \downarrow	ATE
FlowMap	29.84	0.916	0.073	19.8	0.00055	27.23	0.849	0.079	7.5	0.00209
COLMAP	29.95	0.928	0.074	4.8	N/A	25.73	0.851	0.098	1.1	N/A
COLMAP (MVS)	31.03	0.938	0.060	42.5	N/A	27.99	0.867	0.072	13.4	N/A
DROID-SLAM*	29.83	0.913	0.066	0.6	0.00017	26.21	0.818	0.094	0.3	0.00074
NoPE-NeRF*	13.60	0.377	0.750	1913.1	0.04429	17.35	0.490	0.591	1804.0	0.03920

	Tanks & Temples (14 scenes)					CO3D (2 scenes)				
Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time (min.) \downarrow	ATE	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time (min.) \downarrow	ATE
FlowMap	27.00	0.854	0.101	22.3	0.00124	31.11	0.896	0.064	22.1	0.01589
COLMAP	26.74	0.848	0.130	5.5	N/A	25.17	0.750	0.190	12.6	N/A
COLMAP (MVS)	27.43	0.863	0.097	51.4	N/A	25.35	0.762	0.175	52.0	N/A
DROID-SLAM*	25.70	0.824	0.133	0.8	0.00122	25.97	0.790	0.139	0.8	0.01728
NoPE-NeRF*	13.38	0.449	0.706	2432.9	0.03709	14.97	0.400	0.770	2604.9	0.03648

Table 1: Camera parameter and geometry initializations from FlowMap produce 3D Gaussian reconstruction results that far outperform prior gradient-based baselines and are generally on par with those produced by COLMAP. Methods marked with an asterisk require ground-truth intrinsics. We report ATE with respect to COLMAP’s pose estimates for reference, since no ground-truth trajectories exist for common view synthesis datasets. We exclude scenes where COLMAP or FlowMap fail entirely; each fails on 4 scenes. See the supplementary document for more details.

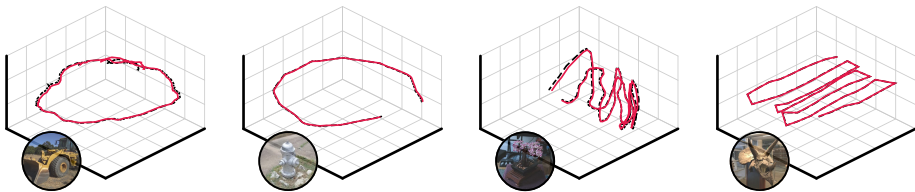


Fig. 6: Qualitative Pose Estimation Comparison. FlowMap (solid red) recovers camera poses that are very close to those of COLMAP (dotted black).

method. Given a collection of images, COLMAP outputs per-image camera poses and intrinsics alongside a sparse 3D point cloud of the underlying scene. 3D Gaussian Splatting, which was designed around COLMAP’s SfM outputs, is initialized using this point cloud. Second, we evaluate against COLMAP multi-view stereo (MVS), which enhances COLMAP’s output with a much denser 3D point cloud. When initialized using this denser point cloud, 3D Gaussian Splatting produces slightly better results. However, note that COLMAP MVS is rarely used in practice because it can be prohibitively time-consuming to run. Third, we evaluate against DROID-SLAM, a neural SLAM system trained on a synthetic dataset of posed video trajectories. Finally, we evaluate against NoPE-NeRF, an method that jointly optimizes a neural radiance field and unknown camera poses. Note that unlike FlowMap and COLMAP, both DROID-SLAM and NoPE-NeRF require camera intrinsics as input.

Datasets. We analyze FlowMap on four standard novel view synthesis datasets: MipNeRF-360 [1], Tanks & Temples [30], LLFF [41], and CO3D [51]. Because FlowMap runs on video sequences, we restrict these datasets to just the video-like sequences they provide.




Fig. 7: Point Clouds Reconstructed by FlowMap. Unprojecting FlowMap depths using FlowMap’s intrinsics and poses yields dense and consistent point clouds.

Methodology. We run FlowMap and the baselines using images that have been rescaled to a resolution of about 700,000 pixels. We then optimize 3D Gaussian scenes for all methods except NoPE-NeRF, since it provides its own NeRF renderings. We use 90% of the available views for training and 10% for testing. During 3D Gaussian fitting, we follow the common [61] practice of fine-tuning the initial camera poses and intrinsics. Such refinement is beneficial because the camera poses produced by SfM algorithms like COLMAP are generally not pixel-perfect [33, 49]. We use the 3D points provided by COLMAP, DROID-SLAM, and FlowMap as input to 3D Gaussian Splatting. For FlowMap, we combine the output depth maps, poses, and intrinsics to yield one point per depth map pixel.

6.1 Novel View Synthesis Results

Tab. 1 reports rendering quality metrics (PSNR, SSIM, and LPIPS) on the held-out test views, and Fig. 5 shows qualitative results. Qualitatively, FlowMap facilitates high-quality 3D reconstructions with sharp details. Quantitatively, FlowMap performs slightly better than COLMAP SfM and significantly outperforms DROID-SLAM and NoPE-NeRF. Only COLMAP MVS slightly exceeds FlowMap in terms of reconstruction quality. As noted previously, COLMAP MVS is rarely used for 3D Gaussian Splatting, since it is very time-consuming to run on high-resolution images.

6.2 Camera Parameter Estimation Results

Since the datasets we use do not provide ground-truth camera parameters, they cannot be used to directly evaluate camera parameter estimates. Instead, Tab. 1 reports the average trajectory error (ATE) of FlowMap, DROID-SLAM, and NoPE-NeRF with respect to COLMAP. Since COLMAP’s poses are not perfect [49], this comparison is not to be understood as a benchmark, but rather as an indication of how close these methods’ outputs are to COLMAP’s state-of-the-art estimates. We find that DROID-SLAM and FlowMap both recover poses

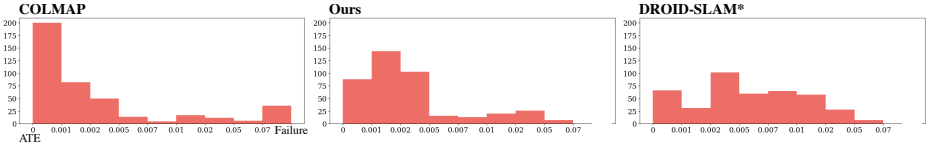


Fig. 8: Large-scale Robustness Study. We run FlowMap and DROID-SLAM on 420 CO3D scenes across 10 categories and plot mean ATEs with respect to CO3D’s COLMAP-generated pose metadata. We also re-run COLMAP on the same data. Compared to DROID-SLAM, which requires ground-truth intrinsics, FlowMap produces notably lower ATEs. FlowMap’s ATE distribution is similar to one obtained by re-running COLMAP, with most ATEs falling under 0.005 in both cases.

that are close to COLMAP’s, while NoPE-NeRF’s estimated poses are far off. When computing ATEs, we normalize all trajectories such that $\text{tr}(XX^T) = 1$, where X is an n -by-3 matrix of camera positions.

Fig. 6 plots trajectories recovered by FlowMap against those recovered by COLMAP, showing that they are often nearly identical. Fig. 7 shows point clouds derived from FlowMap’s estimated depth maps and camera parameters, illustrating that FlowMap recovers well-aligned scene geometry.

6.3 Large-Scale Robustness Study

We study FlowMap’s robustness by using it to estimate camera poses for 420 CO3D scenes from 10 categories. We compare these trajectories to CO3D’s pose annotations, which were computed using COLMAP. Since the quality of CO3D’s ground-truth trajectories varies between categories, we focus on categories that have been used to train novel view synthesis models [6, 66, 73], where pose accuracy is expected to be higher. We find that FlowMap’s mean ATE (0.0056) is lower than DROID-SLAM’s (0.0082) and similar to the mean ATE obtained by re-running COLMAP and comparing the results to the provided poses (0.0038). This demonstrates that FlowMap consistently estimates poses which are close to COLMAP’s. We note that COLMAP failed to estimate poses for 36 scenes, possibly because we ran it at a sparser frame rate to be consistent with our method or because the original annotations were generated using different COLMAP settings; we exclude COLMAP’s failures from the above mean ATE. See Fig. 8 for distributions of ATE values with respect to CO3D’s provided camera poses.

7 Ablations and Analysis

We perform ablations to answer the following questions:

- Question 1: Are FlowMap’s reparameterizations of depth, pose, and intrinsics necessary, or do free variables perform equally well?
- Question 2: Are point tracks critical to FlowMap’s performance?
- Question 3: Does self-supervised pre-training of the depth estimation and correspondence weight neural networks improve performance?

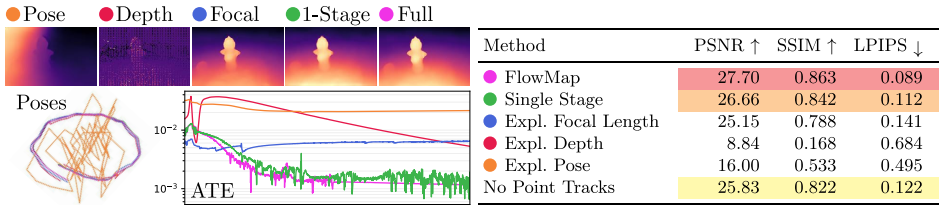


Fig. 9: Ablations. We ablate the proposed feed-forward re-parameterizations of depth, pose, and intrinsics across all datasets. We find that these reparameterizations are not only critical for high-quality downstream 3D Gaussian Splatting, but also lead to dramatically accelerated convergence, where FlowMap generally converges to high quality poses within a fraction of the optimization steps required for the ablated variants. We further find that point tracks lead to a significant boost over optical flow alone (right). See the supplemental document for more ablations.

Parameterizations of Depth, Pose, and Camera Intrinsics (Q1) We compare the reparameterizations described in Sec. 4 to direct, free-variable optimization of pose, depth, and intrinsics. Fig. 9 shows qualitative results and quantitative results averaged across 33 scenes. We find that free-variable variants of FlowMap produce significantly worse reconstruction results and converge much more slowly, confirming that FlowMap’s reparameterizations are crucial.

It is worth noting that often, explicitly optimizing a focal length produces high-quality results, as indicated by the relatively high performance of the “Expl. Focal Length” ablation. In fact, given a good initialization, direct focal length regression produces slightly better results than the proposed focal length reparameterization alone on about 80 percent of scenes. However, on about 20 percent of scenes, this approach falls into a local minimum and reconstruction fails catastrophically. This justifies the approach FlowMap uses, where the first 1,000 optimization steps use a reparameterized focal length, which is then used to initialize an explicit focal length used for another 1,000 optimization steps.

We further highlight that FlowMap’s reparameterizations are necessary to estimate poses and intrinsics in a single forward pass, which is crucial for the generalizable (pre-training) setting explored in Q3.

Point Tracking (Q2) While optical flow is only computed between adjacent frames, point track estimators can accurately track points across many frames. In Fig. 9, we show that FlowMap’s novel view synthesis performance drops moderately when point tracks are disabled. Qualitatively, we find that point tracks reduce drift for longer sequences, such as object-centric 360° scenes. This suggests that FlowMap will benefit from further improvements in point tracking methods. We note that FlowMap’s loss formulation is compatible with conventional correspondence methods (e.g. SIFT [37] with RANSAC) and learned correspondences [54], which can be treated identically to point tracks. FlowMap could also be extended to use conventional loop closure mechanisms, which would further reduce drift.

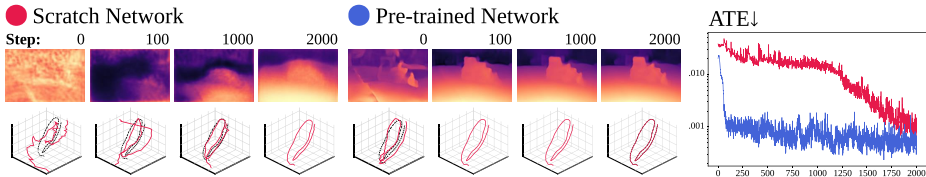


Fig. 10: Effects of pretraining. While a randomly initialized FlowMap network often provides accurate poses after optimization, pre-training leads to faster convergence and slightly improved poses. Here we plot depth estimates at specific optimization steps (left) as well as pose accuracy with respect to COLMAP during optimization (right). Randomly initialized FlowMap networks often require more than 20,000 steps to match the accuracy of a pre-trained initialization at 2,000 steps.

Pre-training Depth and Correspondence Networks (Q3) Since FlowMap is differentiable and provides gradients for any depth-estimating neural network, it is compatible with both randomly initialized neural networks and pre-trained priors. Learned priors can come from optimization on many scenes, from existing depth estimation models, or from a combination of the two. In practice, starting with a pre-trained prior leads to significantly faster convergence, as illustrated in Fig. 10. Note that pre-training and generalization are uniquely enabled by the proposed feed-forward reparameterizations of depth, focal length, and poses.

8 Conclusion

We have introduced FlowMap, a simple, robust, and scalable first-order method for estimating camera parameters from video. Our model outperforms existing gradient-descent based methods for estimating camera parameters. FlowMap’s depth and camera parameters enable subsequent reconstruction via Gaussian Splatting of comparable quality to COLMAP. FlowMap is written in PyTorch and achieves runtimes of 3 minutes for short sequences and 20 minutes for long sequences, and we anticipate that concerted engineering efforts could accelerate FlowMap by an order of magnitude. Perhaps most excitingly, FlowMap is fully differentiable with respect to per-frame depth estimates. FlowMap can thus serve as a building block for a new generation of self-supervised monocular depth estimators, deep-learning-based multi-view-geometry methods, and methods for generalizable novel view synthesis [7, 18, 60, 66, 69, 77], unlocking training on internet-scale datasets of unposed videos.

Acknowledgements. This work was supported by the National Science Foundation under Grant No. 2211259, by the Singapore DSTA under DST00OEC120300823 (New Representations for Vision and 3D Self-Supervised Learning for Label-Efficient Vision), by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) under 140D0423C0075, by the Amazon Science Hub, and by IBM. The Toyota Research Institute also partially supported this work. The views and conclusions contained herein reflect the opinions and conclusions of its authors and no other entity.

References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021) [10](#)
2. Bian, W., Wang, Z., Li, K., Bian, J., Prisacariu, V.A.: Nope-nerf: Optimising neural radiance field with no pose prior (2023) [2](#), [3](#), [4](#)
3. Bloesch, M., Czarnowski, J., Clark, R., Leutenegger, S., Davison, A.J.: Codeslam—learning a compact, optimisable representation for dense visual slam. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2560–2568 (2018) [3](#)
4. Bowen, R.S., Tucker, R., Zabih, R., Snavely, N.: Dimensions of motion: Monocular prediction through flow subspaces. In: Proceedings of the International Conference on 3D Vision (3DV). pp. 454–464. IEEE (2022) [3](#)
5. Campos, C., Elvira, R., Rodríguez, J.J.G., M. Montiel, J.M., D. Tardós, J.: Orbslam3: An accurate open-source library for visual, visual-inertial, and multimap slam. Transactions on Robotics (6), 1874–1890 (2021) [3](#)
6. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aitala, M., De Mello, S., Karras, T., Wetzstein, G.: Generative novel view synthesis with 3d-aware diffusion models. Proceedings of the International Conference on 3D Vision (3DV) (2023) [12](#)
7. Charatan, D., Li, S., Tagliasacchi, A., Sitzmann, V.: pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023) [14](#)
8. Chen, Y., Lee, G.H.: Dbarf: Deep bundle-adjusting generalizable neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 24–34 (June 2023) [4](#)
9. Cheng, Z., Esteves, C., Jampani, V., Kar, A., Maji, S., Makadia, A.: Lu-nerf: Scene and pose estimation by synchronizing local unposed nerfs. arXiv preprint arXiv:2306.05410 (2023) [4](#)
10. Chng, S.F., Ramasinghe, S., Sherrah, J., Lucey, S.: Garf: gaussian activated radiance fields for high fidelity reconstruction and pose estimation. arXiv e-prints pp. arXiv–2204 (2022) [4](#)
11. Choy, C., Dong, W., Koltun, V.: Deep global registration. In: Proc. CVPR (2020) [7](#)
12. Choy, C.B., Gwak, J., Savarese, S., Chandraker, M.: Universal correspondence network. Advances in neural information processing systems **29** (2016) [3](#)
13. Clark, R., Bloesch, M., Czarnowski, J., Leutenegger, S., Davison, A.J.: Learning to solve nonlinear least squares for monocular stereo. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 284–299 (2018) [3](#)
14. Czarnowski, J., Laidlow, T., Clark, R., Davison, A.J.: Deepfactors: Real-time probabilistic dense monocular SLAM. Computing Research Repository (CoRR) (2020) [3](#)
15. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised NeRF: Fewer views and faster training for free. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2022) [4](#)
16. DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 224–236 (2018) [1](#)

17. Doersch, C., Yang, Y., Vecerik, M., Gokay, D., Gupta, A., Aytar, Y., Carreira, J., Zisserman, A.: Tapir: Tracking any point with per-frame initialization and temporal refinement. *arXiv preprint arXiv:2306.08637* (2023) [3](#)
18. Du, Y., Smith, C., Tewari, A., Sitzmann, V.: Learning to render novel views from wide-baseline stereo pairs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2023) [14](#)
19. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence* **40**(3), 611–625 (2017) [3](#)
20. Engel, J., Schöps, T., Cremers, D.: Lsd-slam: Large-scale direct monocular slam. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 834–849. Springer (2014) [3](#)
21. Fu, H., Yu, X., Li, L., Zhang, L.: Cbarf: Cascaded bundle-adjusting neural radiance fields from imperfect camera poses (2023) [4](#)
22. Fu, Y., Liu, S., Kulkarni, A., Kautz, J., Efros, A.A., Wang, X.: Colmap-free 3d gaussian splatting (2023) [4](#)
23. Fu, Y., Misra, I., Wang, X.: Mononerf: Learning generalizable nerfs from monocular videos without camera poses (2023) [4](#)
24. Harley, A.W., Fang, Z., Fragkiadaki, K.: Particle video revisited: Tracking through occlusions using point trajectories. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2022) [3](#)
25. Jeong, Y., Ahn, S., Choy, C., Anandkumar, A., Cho, M., Park, J.: Self-calibrating neural radiance fields. In: *Proceedings of the International Conference on Computer Vision (ICCV)*. pp. 5846–5854 (2021) [4](#)
26. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Rupprecht, C.: CoTracker: It is better to track together (2023) [3](#), [8](#)
27. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. *arXiv preprint arXiv:2312.02126* (2023) [4](#)
28. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)* **42**(4), 1–14 (2023) [3](#), [4](#), [9](#)
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv:1412.6980* (2014) [8](#)
30. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (TOG)* **36**(4) (2017) [10](#)
31. Lai, Z., Liu, S., Efros, A.A., Wang, X.: Video autoencoder: self-supervised disentanglement of static 3d structure and motion. In: *Proceedings of the International Conference on Computer Vision (ICCV)*. pp. 9730–9740 (2021) [4](#)
32. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2023) [4](#)
33. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: *Proceedings of the International Conference on Computer Vision (ICCV)*. pp. 5741–5751 (2021) [2](#), [4](#), [11](#)
34. Lindenberger, P., Sarlin, P.E., Pollefeys, M.: LightGlue: Local Feature Matching at Light Speed. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2023) [1](#)
35. Liu, C., Gu, J., Kim, K., Narasimhan, S.G., Kautz, J.: Neural rgb (r) d sensing: Depth and uncertainty from a video camera. In: *Proceedings of the IEEE Con-*

- ference on Computer Vision and Pattern Recognition (CVPR). pp. 10986–10995 (2019) [3](#)
36. Liu, Y.L., Gao, C., Meuleman, A., Tseng, H.Y., Saraf, A., Kim, C., Chuang, Y.Y., Kopf, J., Huang, J.B.: Robust dynamic radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13–23 (2023) [4](#)
 37. Lowe, D.: Object recognition from local scale-invariant features. In: Proceedings of the Seventh IEEE International Conference on Computer Vision. pp. 1150–1157 vol.2 (1999) [13](#)
 38. Luo, Z., Shen, T., Zhou, L., Zhu, S., Zhang, R., Yao, Y., Fang, T., Quan, L.: Geodesc: Learning local descriptors by integrating geometry constraints. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 168–183 (2018) [3](#)
 39. Matsuki, H., Murai, R., Kelly, P.H.J., Davison, A.J.: Gaussian Splatting SLAM. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2024) [4](#)
 40. Meuleman, A., Liu, Y.L., Gao, C., Huang, J.B., Kim, C., Kim, M.H., Kopf, J.: Progressively optimized local radiance fields for robust view synthesis. In: CVPR (2023) [4](#)
 41. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* **38**(4), 1–14 (2019) [10](#)
 42. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 405–421 (2020) [4](#)
 43. Mishchuk, A., Mishkin, D., Radenovic, F., Matas, J.: Working hard to know your neighbor’s margins: Local descriptor learning loss. *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017) [3](#)
 44. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)* **41**(4), 102:1–102:15 (2022) [4](#)
 45. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: Orb-slam: a versatile and accurate monocular slam system. *Transactions on Robotics* (5), 1147–1163 (2015) [3](#)
 46. Mur-Artal, R., Tardós, J.D.: Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *Transactions on Robotics* **33**(5), 1255–1262 (2017) [3](#)
 47. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3504–3515 (2020) [4](#)
 48. Ono, Y., Trulls, E., Fua, P., Yi, K.M.: Lf-net: Learning local features from images. *Advances in Neural Information Processing Systems (NeurIPS)* **31** (2018) [3](#)
 49. Park, K., Henzler, P., Mildenhall, B., Barron, J.T., Martin-Brualla, R.: Camp: Camera preconditioning for neural radiance fields. *ACM Transactions on Graphics (TOG)* **42**(6), 1–11 (2023) [11](#)
 50. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020) [22](#)

51. Reizenstein, J., Shapovalov, R., Henzler, P., Sbordone, L., Labatut, P., Novotny, D.: Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In: *Proceedings of the International Conference on Computer Vision (ICCV)*. pp. 10901–10911 (2021) [10](#)
52. Rosinol, A., Abate, M., Chang, Y., Carlone, L.: Kimera: an open-source library for real-time metric-semantic localization and mapping. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. pp. 1689–1696. IEEE (2020) [3](#)
53. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4938–4947 (2020) [1](#), [3](#)
54. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: SuperGlue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020) [13](#)
55. Sarlin, P.E., Unagar, A., Larsson, M., Germain, H., Toft, C., Larsson, V., Pollefeys, M., Lepetit, V., Hammarstrand, L., Kahl, F., Sattler, T.: Back to the Feature: Learning Robust Camera Localization from Pixels to Pose. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) [1](#)
56. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4104–4113 (2016) [1](#), [3](#), [9](#)
57. Sitzmann, V., Thies, J., Heide, F., Nießner, M., Wetzstein, G., Zollhöfer, M.: Deepvoxels: Learning persistent 3d feature embeddings. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) [4](#)
58. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems (NeurIPS)* (2019) [4](#)
59. Smith, C., Du, Y., Tewari, A., Sitzmann, V.: Flowcam: Training generalizable 3d radiance fields without camera poses via pixel-aligned scene flow. *Advances in Neural Information Processing Systems (NeurIPS)* (2023) [4](#), [6](#), [7](#), [22](#)
60. Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Generalizable patch-based neural rendering. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer (2022) [14](#)
61. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Kerr, J., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., McAllister, D., Kanazawa, A.: Nerfstudio: A modular framework for neural radiance field development. In: *ACM Transactions on Graphics (TOG)* (2023) [11](#)
62. Tang, C., Tan, P.: BA-Net: Dense bundle adjustment network. *arXiv preprint arXiv:1806.04807* (2018) [3](#)
63. Teed, Z., Deng, J.: Deepv2d: Video to depth with differentiable structure from motion. *arXiv preprint arXiv:1812.04605* (2018) [3](#)
64. Teed, Z., Deng, J.: RAFT: Recurrent all-pairs field transforms for optical flow. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2020) [3](#), [8](#)
65. Teed, Z., Deng, J.: Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in Neural Information Processing Systems (NeurIPS)* **34** (2021) [3](#)
66. Tewari, A., Yin, T., Cazenavette, G., Rezchikov, S., Tenenbaum, J.B., Durand, F., Freeman, W.T., Sitzmann, V.: Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *Advances in Neural Information Processing Systems (NeurIPS)* (2023) [12](#), [14](#)

67. Ummenhofer, B., Zhou, H., Uhrig, J., Mayer, N., Ilg, E., Dosovitskiy, A., Brox, T.: Demon: Depth and motion network for learning monocular stereo. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5038–5047 (2017) [3](#)
68. Wang, J., Karaev, N., Rupprecht, C., Novotny, D.: Visual geometry grounded deep structure from motion. arXiv preprint arXiv:2312.04563 (2023) [4](#)
69. Wang, Q., Wang, Z., Genova, K., Srinivasan, P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snively, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [14](#)
70. Wang, W., Hu, Y., Scherer, S.: Tartanvo: A generalizable learning-based vo. In: Conference on Robot Learning. pp. 1761–1772. PMLR (2021) [3](#)
71. Wang, Z., Wu, S., Xie, W., Chen, M., Prisacariu, V.A.: NeRF—: Neural radiance fields without known camera parameters. arXiv:2102.07064 (2021) [2](#)
72. Wang, Z., Wu, S., Xie, W., Chen, M., Prisacariu, V.A.: Nerf-: Neural radiance fields without known camera parameters. arXiv preprint arXiv:2102.07064 (2021) [4](#)
73. Wewer, C., Raj, K., Ilg, E., Schiele, B., Lenssen, J.E.: latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. In: arXiv (2024) [12](#)
74. Wu, X., Xu, J., Zhang, X., Bao, H., Huang, Q., Shen, Y., Tompkin, J., Xu, W.: Scannerf: Scalable bundle-adjusting neural radiance fields for large-scale scene rendering. ACM Transactions on Graphics (TOG) (2023) [4](#)
75. Xia, Y., Tang, H., Timofte, R., Van Gool, L.: Sinerf: Sinusoidal neural radiance fields for joint pose estimation and scene reconstruction. arXiv preprint arXiv:2210.04553 (2022) [4](#)
76. Yen-Chen, L., Florence, P., Barron, J.T., Rodriguez, A., Isola, P., Lin, T.Y.: inerf: Inverting neural radiance fields for pose estimation. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1323–1330. IEEE (2021) [4](#)
77. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelNeRF: Neural radiance fields from one or few images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [14](#)
78. Yugay, V., Li, Y., Gevers, T., Oswald, M.R.: Gaussian-slam: Photo-realistic dense slam with gaussian splatting. arXiv preprint arXiv:2312.10070 (2023) [4](#)
79. Zhao, W., Liu, S., Guo, H., Wang, W., Liu, Y.J.: Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 523–542. Springer (2022) [3](#)
80. Zhou, H., Ummenhofer, B., Brox, T.: Deeptam: Deep tracking and mapping. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 822–838 (2018) [3](#)

Supplemental Material

S1 Additional Ablation Studies

In Tab. S1, we include three additional ablations. The “Random Init.” ablation uses a randomly initialized CNN for FlowMap training with 2,000 steps of optimization. The “Random Init. (20k)” ablation is identical, but runs for 20,000 optimization steps. The “No Corresp. Weights” ablation removes the correspondence weights used in FlowMap’s Procrustes-solving step. We note that the “Random Init. (20k)” ablation’s performance almost matches FlowMap’s, indicating that although pre-training helps FlowMap converge much more quickly, it is not necessary for accuracy. In Tab. S2, we report per-scene ablation results. Finally, Fig. S1 compares convergence between FlowMap and the free-variable parameterization variants on more scenes.

S2 Additional Results

S2.1 Pre-Trained Depth vs. Fine-Tuned Depth vs. High-Resolution Fine-Tuned Depth.

In Fig. S4, we compare the depths produced by FlowMap’s initialization to the depths produced after FlowMap optimization. We additionally compare these results to a MiDaS CNN fine-tuned at a significantly higher resolution. We find that per-scene fine-tuning leads to high-quality depth predictions. This is illustrated by Fig. S3, which demonstrates FlowMap’s ability to generate high-quality, consistent depths. However, it is worth noting that FlowMap’s off-the-shelf depths are slightly blurry. To investigate whether this is a limitation of our loss or the architecture of the depth-predicting CNN, we also perform optimization at a higher resolution. We find that this leads to crisp depth maps, demonstrating that blurry depth maps are a result of insufficient capacity of the MiDaS backbone and not a limitation of our camera-induced flow loss. Notably, the poses barely change in this fine-tuning stage. It is likely that replacing the MiDaS depth predictor with a more powerful depth backbone would lead to sharper depth without high-resolution fine-tuning.

S2.2 Additional Point Clouds and Qualitative Pose Reconstructions

In Fig. S3, we display 12 additional point clouds plus estimated camera poses across popular datasets and scenes across the LLFF, Tanks and Temples, MipNeRF 360, and CO3D datasets. FlowMap robustly recovers camera poses and scene geometry across these diverse, challenging, and real-world sequences.

S2.3 Failure Cases

While running FlowMap, we observed failures on several scenes. These include the Tanks-and-Temples Auditorium scene (our model struggles with rotation-dominant trajectories), the LLFF Leaves scene (our model falls into a “hollow-face minimum”), and the Tanks-and-Temples Lighthouse scene (this video features a large lens flare which degrades the optical flow). Future extensions to FlowMap could use an occlusion-aware formulation to avoid hollow-face minima.

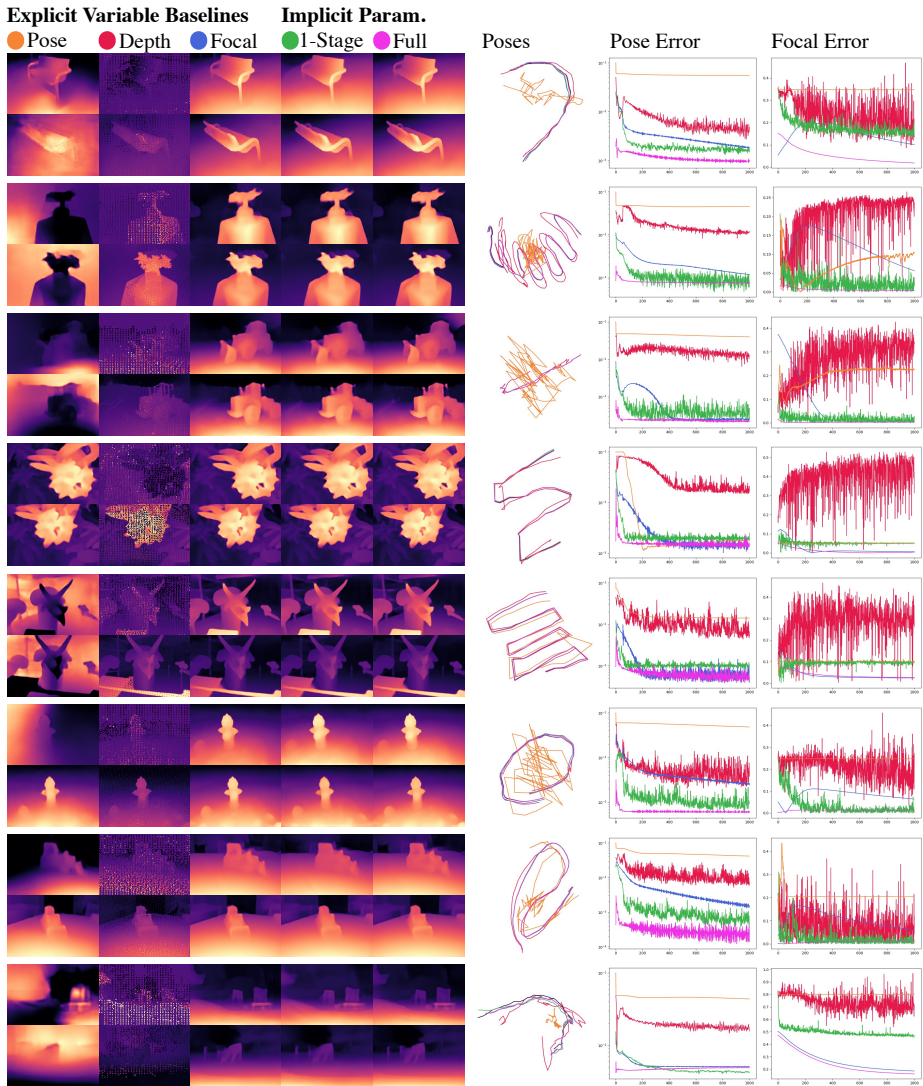


Fig. S1: Pose and Geometry Convergence for Free-Variable vs. Proposed Parameterizations. We plot poses, depths, focal lengths, and pose error (ATE) obtained with our proposed parameterizations (“Full”) vs. those obtained with free-variable parameterizations at various optimization steps. With our proposed reparameterizations (“Full”) as a baseline, we ablate either depth, focal length, or poses as free-variable optimizations and plot the resulting optimizations’ pose and depth estimates. For instance, “Depth” corresponds to making the depth an explicit free-variable in the optimization. Using pose-as-variable and depth-as-variable often lead to “hollow-face” geometry, where the geometry is effectively inverted but still mostly satisfies the optical flow constraints. We also show results from a single-stage FlowMap pipeline, which only uses the implicit parameterization of intrinsics rather than switching to regressed intrinsics halfway through optimization. Note that the plotted lines for “Full” are initialized with the results of “1-Stage” and represent the second stage (explicit focal length) of FlowMap optimization.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
FlowMap	27.70	0.863	0.089
Single Stage	26.66	0.842	0.112
Expl. Focal Length	25.15	0.788	0.141
Expl. Depth	8.84	0.168	0.684
Expl. Pose	16.00	0.533	0.495
No Tracks	25.83	0.822	0.122
Random Init.	25.54	0.808	0.129
Random Init. (20k)	27.25	0.850	0.101
No Corresp. Weights	24.18	0.765	0.168

Table S1: Additional Ablations. We report additional ablation results averaged across all scenes alongside the ablations found in the main paper.

S3 Implementation Details

S3.1 Procrustes Solver Details

Our pose solver is the one introduced in FlowCam [59]; see [59] for details. The only difference is that instead of selecting 1000 random points for the Procrustes estimation, we fix the points (uniformly spaced throughout the image) when performing per-scene overfitting. We find that fixing the points used for the pose solver allows the network to better overfit confidence weights and subsequently yields better poses.

S3.2 Intrinsic Solver Details

For the intrinsic solver, we assume a pinhole camera estimate and discretize a set of 60 candidate focal lengths between .5 and 2 (in resolution-independent units). We use a softmax on the flow error maps, as discussed in the main paper. We scale the error maps by a temperature factor of 10 and weight the error maps by the flow confidence weights. See Fig. S6 for illustration.

S3.3 Depth NN (MiDaS) details

For our depth network, we use the lightweight CNN version of MiDaS [50], pretrained with the publicly available weights trained on relative-depth estimation.

S3.4 Correspondence Weight MLP

The correspondence weight MLP is a three-layer MLP with ReLU activations and 128 hidden units per layer. It takes as input two corresponding image features and outputs a per-correspondence weight between 0 and 1 via a sigmoid activation. Here we use intermediate feature maps from the depth network as the image features. These weights are used in the weighted Procrustes pose solver.



Fig. S3: Additional Point Clouds Here we plot additional point clouds across the Tanks and Temples, LLFF, Mip-NeRF360, and CO3D datasets.

S4 Experiment Details

S4.1 Image Resolution

To manage computational cost (our current implementation loads the entire video into memory), we compute optical flow and point tracks at a resolution of around 700,000 pixels, then perform FlowMap optimization at 1/16th the resolution.

S4.2 Hyperparameters

We train for 2000 steps using Adam and use a learning rate of $3e-5$. For the pose-as-variable experiments, we choose Euler angles as the parameterization of the rotation matrix.

S4.3 Pre-Training Details

Before performing per-scene fine-tuning, we found it useful to learn a large-scale prior for better initialization. We use the same FlowMap loss formulation but train it on datasets of videos (instead of optimizing on a single scene). We use videos from CO3D, Real Estate 10K, and KITTI for pretraining. Note that we only use the raw videos from these datasets (no intrinsics, poses, or sparse geometry).




Fig. S4: Depth Estimates Before and After Optimization. The depth prediction neural network can either be randomly initialized or pre-trained, though pre-trained depth networks lead to much faster convergence. In the second row, we show the output of the depth prediction neural network after pre-training it on a dataset consisting of CO3D, KITTI, and RealEstate10k. These estimates converge to high-quality depth within only a few hundred FlowMap optimization steps. We see that the quality of the initial, pre-trained depth predictions is not critical to achieve accurate reconstructions. Although we estimate geometry at a lower resolution during optimization to manage memory constraints, we can quickly fine-tune at high-resolution for more detailed depth maps if necessary (bottom row).

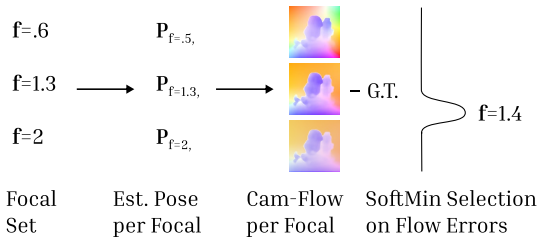
S5 Limitations

While our method is much faster than MVS COLMAP, it is about 30 percent slower than COLMAP at its highest quality setting (on long sequences, about 20 minutes for our method vs. 14 minutes for COLMAP). It additionally requires significantly more GPU memory than COLMAP does. Our method’s pose and intrinsics predictions are less accurate and robust than COLMAP’s, as measured by ATE, though after Gaussian Splatting with fine-tuning of camera parameters, we often perform on par with COLMAP.

Our method further depends on correspondences estimated by point tracks and optical flow. While existing methods for computing point tracks and optical flow are robust, failures sometimes occur, and these failures can affect FlowMap’s accuracy if they are significant. On the other hand, FlowMap will directly improve alongside advancements in these domains.

Finally, our method is constrained to work on frame sequences with significant overlap (i.e., videos) and fails when input sequences contain significant scene motion. The latter limitation is shared with COLMAP, though we hope that our method may serve as a step towards novel methods that address this shortcoming.

(a) Intrinsic Estimation via Best-Explaining Focal
Cam-Induced Flow per Candidate Focal *Choose Focal with Flow Closest to GT*



(b) Depth Estimation via CNN

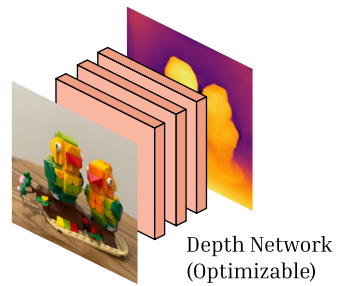


Fig. S6: In (a) we illustrate our implicit focal length formulation, which considers a set of candidate focal lengths, assigns each one an error score, and softly selects the focal length with the lowest error. To calculate the error score for a focal length, we use that focal length to estimate a pose, and then compare the resulting pose-induced optical flow to the ground truth optical flow. In (b) we illustrate that we parameterize depth via the output of a monocular depth prediction CNN.

Method	Fern (LLFF)					Flower (LLFF)					Fortress (LLFF)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	23.70	0.801	0.096		4.8	0.00233	29.07	0.877	0.084	6.6	0.00079	31.13	0.906	0.060	7.8
Single Stage	23.64	0.797	0.098		4.9	0.00294	29.06	0.877	0.086	6.7	0.00065	31.05	0.908	0.058	7.9
Expl. Focal Length	23.07	0.787	0.119		4.6	0.00296	29.13	0.874	0.079	6.4	0.00293	29.82	0.891	0.062	7.6
Expl. Depth	4.71	0.001	0.785		2.9	0.00785	8.04	0.007	0.839	3.6	0.00666	2.60	0.001	0.774	4.1
Expl. Pose	4.71	0.001	0.785		4.5	0.01118	15.51	0.569	0.428	6.3	0.00192	16.49	0.577	0.594	7.4
No Tracks	23.58	0.796	0.099		4.3	0.00316	29.29	0.879	0.084	5.5	0.00337	30.92	0.906	0.059	6.8
Random Init.	22.68	0.756	0.113		4.8	0.00371	28.47	0.864	0.084	6.6	0.00303	30.96	0.904	0.059	7.8
Random Init. (20k)	23.50	0.791	0.098		44.2	0.00312	27.31	0.880	0.083	59.4	0.00054	31.04	0.911	0.057	69.7
No Corresp. Weights	23.27	0.784	0.104		4.4	0.00311	27.61	0.844	0.090	6.0	0.00554	24.05	0.709	0.138	7.1
Method	Horns (LLFF)					Orchids (LLFF)					Room (LLFF)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	28.35	0.903	0.071		10.6	0.00049	19.16	0.615	0.132	5.5	0.00127	32.93	0.958	0.037	7.8
Single Stage	28.19	0.899	0.071		10.7	0.00051	19.33	0.623	0.129	5.6	0.00120	32.75	0.959	0.037	7.9
Expl. Focal Length	24.57	0.823	0.153		10.5	0.00100	18.96	0.606	0.151	5.3	0.00184	32.13	0.953	0.040	7.6
Expl. Depth	5.94	0.002	0.788		5.3	0.00603	6.25	0.003	0.886	3.2	0.00647	5.78	0.004	0.616	4.1
Expl. Pose	14.66	0.506	0.685		10.0	0.01230	12.80	0.279	0.429	5.2	0.01496	16.92	0.767	0.466	7.3
No Tracks	28.32	0.900	0.071		8.2	0.00173	19.21	0.616	0.133	4.8	0.00195	28.98	0.922	0.068	6.3
Random Init.	23.93	0.729	0.172		10.6	0.00486	18.85	0.594	0.146	5.4	0.00188	29.19	0.920	0.067	7.8
Random Init. (20k)	28.33	0.900	0.068		94.2	0.00054	19.40	0.629	0.126	49.5	0.00112	31.92	0.949	0.043	69.4
No Corresp. Weights	28.17	0.893	0.072		9.6	0.00169	18.76	0.597	0.148	5.1	0.00307	31.81	0.952	0.041	7.1
Method	Trex (LLFF)					Bonsai (MipNeRF 360)					Kitchen (MipNeRF 360)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	26.27	0.880	0.075		9.7	0.00655	32.24	0.950	0.047	24.2	0.00048	30.47	0.936	0.049	10.9
Single Stage	26.65	0.886	0.073		9.8	0.00653	32.21	0.951	0.048	24.3	0.00046	30.26	0.925	0.051	11.1
Expl. Focal Length	24.57	0.852	0.107		9.5	0.00766	31.26	0.869	0.231	24.1	0.00184	29.88	0.966	0.054	10.1
Expl. Depth	5.56	0.002	0.719		4.8	0.04406	10.46	0.045	0.633	11.3	0.02830	5.08	0.010	0.753	5.3
Expl. Pose	15.09	0.540	0.544		9.2	0.02277	13.12	0.425	0.577	22.8	0.01407	14.18	0.887	0.587	10.3
No Tracks	24.36	0.831	0.110		7.8	0.00217	32.53	0.863	0.115	18.0	0.00291	25.48	0.794	0.112	8.5
Random Init.	24.84	0.835	0.099		9.7	0.00598	18.38	0.585	0.342	24.2	0.01380	24.94	0.764	0.113	10.9
Random Init. (20k)	26.45	0.882	0.076		86.5	0.00991	18.75	0.600	0.342	214.8	0.01433	31.69	0.945	0.044	96.8
No Corresp. Weights	25.33	0.859	0.094		8.8	0.01083	25.59	0.841	0.118	21.8	0.00141	24.62	0.742	0.118	9.9
Method	Counter (MipNeRF 360)					Barn (Tanks & Temples)					Caterpillar (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	26.80	0.862	0.121		24.2	0.00076	27.10	0.872	0.090	22.3	0.00048	28.25	0.830	0.113	22.3
Single Stage	26.65	0.857	0.124		24.3	0.00083	26.28	0.857	0.104	22.4	0.00102	28.32	0.834	0.110	22.4
Expl. Focal Length	21.82	0.697	0.237		24.1	0.00355	27.04	0.873	0.086	21.9	0.00085	27.12	0.789	0.133	22.1
Expl. Depth	8.80	0.029	0.719		11.4	0.01003	17.01	0.591	0.489	10.7	0.00923	9.15	0.016	0.732	10.7
Expl. Pose	14.40	0.596	0.554		11.4	0.01541	18.09	0.625	0.455	20.9	0.02160	17.57	0.491	0.554	10.7
No Tracks	23.91	0.788	0.183		18.1	0.00240	25.41	0.837	0.122	16.1	0.00363	27.33	0.807	0.134	16.1
Random Init.	26.05	0.847	0.131		24.2	0.00068	26.24	0.864	0.100	22.2	0.00079	26.27	0.750	0.169	22.2
Random Init. (20k)	26.88	0.867	0.115		214.3	0.00064	26.80	0.871	0.091	197.3	0.00049	28.01	0.823	0.122	197.3
No Corresp. Weights	17.93	0.575	0.391		21.8	0.01099	24.53	0.820	0.133	20.3	0.00244	25.93	0.734	0.174	20.1
Method	Church (Tanks & Temples)					Courthouse (Tanks & Temples)					Family (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	28.29	0.883	0.074		22.4	0.00061	27.51	0.911	0.055	22.2	0.00129	27.96	0.889	0.067	22.1
Single Stage	27.60	0.875	0.079		22.3	0.00061	27.67	0.914	0.054	22.3	0.00150	27.65	0.880	0.075	22.4
Expl. Focal Length	27.29	0.856	0.088		22.0	0.00093	26.86	0.897	0.069	22.1	0.00082	27.10	0.873	0.082	22.0
Expl. Depth	16.21	0.518	0.474		10.7	0.02314	3.52	0.001	0.745	10.7	0.00718	4.09	0.001	0.773	10.8
Expl. Pose	17.66	0.582	0.457		21.0	0.00807	19.68	0.726	0.251	21.0	0.00511	15.79	0.562	0.507	21.1
No Tracks	26.93	0.851	0.100		16.1	0.00259	25.27	0.858	0.108	16.1	0.00442	27.00	0.869	0.088	16.1
Random Init.	27.45	0.858	0.089		22.2	0.00112	26.55	0.894	0.071	22.2	0.00314	26.36	0.858	0.093	22.2
Random Init. (20k)	28.67	0.886	0.074		197.4	0.00030	27.62	0.911	0.054	197.8	0.00101	28.07	0.892	0.066	196.9
No Corresp. Weights	27.86	0.875	0.081		20.3	0.00086	25.62	0.868	0.086	20.1	0.00264	19.01	0.629	0.313	20.2
Method	Francis (Tanks & Temples)					Horse (Tanks & Temples)					Ignatius (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	31.90	0.903	0.080		22.4	0.00058	28.35	0.917	0.064	22.4	0.00054	24.54	0.773	0.131	22.4
Single Stage	32.80	0.905	0.078		22.4	0.00054	28.12	0.935	0.046	22.3	0.00359	24.48	0.764	0.131	22.4
Expl. Focal Length	32.49	0.884	0.104		22.0	0.00040	27.82	0.905	0.074	22.0	0.00101	23.12	0.723	0.157	22.0
Expl. Depth	7.42	0.006	0.631		10.8	0.01956	2.67	0.000	0.691	10.7	0.02555	5.68	0.006	0.867	10.7
Expl. Pose	11.19	0.639	0.464		20.9	0.03102	14.60	0.661	0.468	21.0	0.03918	12.48	0.314	0.640	20.9
No Tracks	30.72	0.887	0.100		16.1	0.00113	25.50	0.882	0.101	16.1	0.00241	23.54	0.727	0.163	16.1
Random Init.	29.44	0.862	0.122		22.3	0.00289	25.07	0.871	0.119	22.2	0.00380	23.50	0.737	0.159	22.1
Random Init. (20k)	31.56	0.899	0.085		197.7	0.00138	28.16	0.915	0.067	197.0	0.00066	24.47	0.771	0.138	197.5
No Corresp. Weights	28.92	0.850	0.130		20.1	0.00397	25.32	0.871	0.100	20.2	0.02075	21.89	0.655	0.197	20.2
Method	M60 (Tanks & Temples)					Museum (Tanks & Temples)					Panther (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	23.23	0.805	0.190		22.4	0.00838	28.48	0.862	0.078	22.2	0.00070	27.50	0.882	0.105	22.3
Single Stage	23.30	0.803	0.187		22.3	0.00832	28.62	0.866	0.076	22.4	0.00058	27.31	0.881	0.104	22.4
Expl. Focal Length	19.65	0.696	0.278		22.1	0.01400	28.15	0.850	0.092	22.0	0.00124	21.86	0.737	0.239	22.1
Expl. Depth	13.59	0.527	0.547		16.2	0.01514	25.75	0.856	0.489	16.2	0.01314	16.17	0.619	0.456	16.2
Expl. Pose	14.37	0.566	0.546		21.1	0.01454	16.12	0.489	0.524	21.1	0.02338	16.41	0.633	0.469	20.9
No Tracks	23.17	0.805	0.195		16.1	0.00674	27.63	0.844	0.096	16.1	0.00150	24.69	0.833	0.161	16.1
Random Init.	21.81	0.781	0.206		22.3	0.01008	25.57	0.849	0.094	22.3	0.00080	25.25	0.845	0.141	22.2
Random Init. (20k)	23.61	0.817	0.171		197.7	0.00869	27.84	0.868	0.075	196.9	0.00067	27.61	0.884	0.105	198.1
No Corresp. Weights	9.44	0.545	0.651		20.2	0.02315	29.74	0.849	0.086	20.1	0.00122	26.00	0.850	0.138	20.0
Method	Playground (Tanks & Temples)					Train (Tanks & Temples)					Truck (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓	PSNR ↑	SSIM ↑	LPIS ↓	Time (min.) ↓	ATE ↓
FlowMap	24.29	0.727	0.192		22.2	0.00096	26.22	0.870	0.077	22.2	0.00082	24.34	0.828	0.098	22.3
Single Stage	23.99	0.710	0.209		22.4	0.00102	26.11	0.863	0.082	22.3	0.00010	24.21	0.826	0.102	22.3
Expl. Focal Length	18.55	0.525	0.380		21.9	0.00621	25.82	0.847	0.092						

Method	Fern (LLFF)					Flower (LLFF)					Fortress (LLFF)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	23.70	0.801	0.096	4.8	0.00233	29.07	0.877	0.084	6.6	0.00079	31.13	0.906	0.060	7.8	0.00049
COLMAP	24.04	0.818	0.133	0.4	N/A	29.60	0.884	0.090	2.5	N/A	25.69	0.892	0.087	1.4	N/A
COLMAP (MVS)	24.33	0.826	0.094	6.7	N/A	29.82	0.888	0.085	11.3	N/A	30.97	0.909	0.059	13.9	N/A
DROID-SLAM*	23.13	0.762	0.125	0.1	0.00089	28.48	0.860	0.079	0.2	0.00162	30.05	0.856	0.065	0.3	0.00038
NoPE-NeRF*	19.33	0.520	0.580	1227.0	0.01470	19.63	0.540	0.470	1777.7	0.02581	21.00	0.530	0.510	533.4	0.02068
Method	Horns (LLFF)					Orchids (LLFF)					Room (LLFF)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	28.35	0.903	0.071	10.6	0.00049	19.16	0.615	0.132	5.5	0.00127	32.93	0.958	0.037	7.8	0.00274
COLMAP	27.95	0.912	0.092	1.5	N/A	19.33	0.636	0.126	0.7	N/A	25.69	0.927	0.096	0.3	N/A
COLMAP (MVS)	28.68	0.902	0.067	20.5	N/A	19.79	0.657	0.117	8.7	N/A	33.43	0.963	0.033	14.6	N/A
DROID-SLAM*	28.37	0.881	0.064	0.5	0.00045	18.44	0.555	0.179	0.2	0.00072	27.63	0.924	0.078	0.3	0.00051
NoPE-NeRF*	11.37	0.300	0.820	2597.7	0.07315	13.11	0.270	0.620	1377.9	0.05492	17.79	0.650	0.590	2500.5	0.03714
Method	Trex (LLFF)					Bonsai (MipNeRF 360)					Kitchen (MipNeRF 360)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	26.27	0.880	0.075	9.7	0.00655	32.24	0.950	0.047	24.2	0.00048	30.47	0.936	0.049	10.9	0.00041
COLMAP	27.52	0.888	0.095	1.1	N/A	32.64	0.949	0.053	6.9	N/A	28.82	0.936	0.056	0.6	N/A
COLMAP (MVS)	28.92	0.922	0.049	18.4	N/A	33.14	0.957	0.045	52.2	N/A	31.33	0.948	0.045	22.4	N/A
DROID-SLAM*	27.36	0.898	0.067	0.3	0.00062	31.96	0.947	0.045	0.9	0.00016	29.75	0.933	0.054	0.4	0.00015
NoPE-NeRF*	18.71	0.550	0.550	2614.1	0.04796	13.49	0.370	0.770	2615.2	0.04475	14.86	0.370	0.710	516.3	0.05471
Method	Counter (MipNeRF 360)					Barn (Tanks & Temples)					Caterpillar (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	26.80	0.862	0.121	24.2	0.00076	27.10	0.872	0.090	22.3	0.00048	28.25	0.830	0.113	22.3	0.00030
COLMAP	28.39	0.899	0.107	4.1	N/A	27.18	0.874	0.108	3.5	N/A	28.05	0.825	0.134	6.6	N/A
COLMAP (MVS)	28.61	0.909	0.089	52.9	N/A	27.91	0.889	0.075	51.5	N/A	28.52	0.839	0.103	51.1	N/A
DROID-SLAM*	27.78	0.890	0.099	0.7	0.00019	27.03	0.877	0.082	0.8	0.00029	28.13	0.820	0.108	0.9	0.00020
NoPE-NeRF*	12.44	0.390	0.770	2607.8	0.03342	13.06	0.460	0.710	2608.4	0.03761	16.42	0.390	0.680	2469.9	0.03112
Method	Church (Tanks & Temples)					Courthouse (Tanks & Temples)					Family (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	28.29	0.883	0.074	22.4	0.00061	27.51	0.911	0.055	22.2	0.00129	27.96	0.889	0.067	22.1	0.00039
COLMAP	27.93	0.866	0.107	6.3	N/A	27.79	0.916	0.056	5.9	N/A	27.13	0.878	0.092	5.0	N/A
COLMAP (MVS)	28.71	0.890	0.068	50.9	N/A	28.36	0.926	0.044	51.9	N/A	28.40	0.897	0.062	50.9	N/A
DROID-SLAM*	27.79	0.899	0.084	0.8	0.00065	27.94	0.916	0.051	0.9	0.00034	27.78	0.873	0.081	0.8	0.00040
NoPE-NeRF*	12.91	0.400	0.700	2575.8	0.02752	14.92	0.510	0.590	2599.3	0.03462	12.87	0.470	0.700	2597.4	0.03232
Method	Francis (Tanks & Temples)					Horse (Tanks & Temples)					Ignatius (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	31.90	0.903	0.080	22.4	0.00058	28.35	0.917	0.064	22.4	0.00054	24.54	0.773	0.131	22.4	0.00037
COLMAP	31.85	0.896	0.124	3.6	N/A	27.34	0.903	0.097	3.4	N/A	24.95	0.781	0.153	5.6	N/A
COLMAP (MVS)	32.73	0.913	0.069	51.1	N/A	28.82	0.926	0.062	53.2	N/A	24.93	0.795	0.119	51.2	N/A
DROID-SLAM*	22.23	0.753	0.275	0.9	0.00041	27.61	0.909	0.069	0.8	0.00051	24.28	0.750	0.142	0.8	0.00025
NoPE-NeRF*	17.27	0.570	0.640	524.9	0.02569	9.87	0.590	0.700	2587.4	0.04710	10.90	0.260	0.780	2583.2	0.04241
Method	M60 (Tanks & Temples)					Museum (Tanks & Temples)					Panther (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	23.23	0.805	0.190	22.4	0.00838	28.48	0.862	0.078	22.2	0.00070	27.50	0.882	0.105	22.3	0.00112
COLMAP	22.04	0.803	0.219	6.2	N/A	28.94	0.863	0.100	5.3	N/A	27.32	0.882	0.129	5.0	N/A
COLMAP (MVS)	21.75	0.791	0.221	51.9	N/A	29.05	0.874	0.070	50.6	N/A	27.96	0.891	0.101	52.2	N/A
DROID-SLAM*	22.66	0.792	0.195	0.7	0.00667	27.74	0.833	0.096	0.8	0.00088	27.48	0.878	0.106	0.8	0.00150
NoPE-NeRF*	12.67	0.490	0.720	2485.1	0.04258	14.26	0.430	0.800	2606.9	0.03224	13.71	0.500	0.690	2591.0	0.03854
Method	Playground (Tanks & Temples)					Train (Tanks & Temples)					Truck (Tanks & Temples)				
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE
FlowMap	24.29	0.727	0.192	22.2	0.00096	26.22	0.870	0.077	22.2	0.00082	24.34	0.828	0.098	22.3	0.00078
COLMAP	22.24	0.684	0.292	7.4	N/A	26.09	0.857	0.104	8.4	N/A	25.57	0.848	0.104	4.9	N/A
COLMAP (MVS)	22.92	0.693	0.230	51.6	N/A	27.43	0.888	0.063	51.4	N/A	26.39	0.864	0.080	50.4	N/A
DROID-SLAM*	21.11	0.642	0.301	0.7	0.00284	26.51	0.872	0.069	0.8	0.00088	21.48	0.739	0.208	0.8	0.00127
NoPE-NeRF*	13.53	0.360	0.770	2613.1	0.04120	13.18	0.440	0.670	2614.8	0.04052	11.71	0.410	0.740	2603.3	0.04583
Method	Bench (CO3D)					Hydrant (CO3D)									
	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE	PSNR ↑	SSIM ↑	LPiPS ↓	Time (min.) ↓	ATE					
FlowMap	33.17	0.927	0.045	22.0	0.03094	29.05	0.865	0.083	22.1	0.00083					
COLMAP	19.87	0.600	0.309	17.2	N/A	30.46	0.900	0.070	8.0	N/A					
COLMAP (MVS)	20.00	0.616	0.292	53.2	N/A	30.70	0.908	0.057	50.8	N/A					
DROID-SLAM*	22.48	0.699	0.206	0.9	0.03433	29.46	0.880	0.073	0.7	0.00024					
NoPE-NeRF*	13.20	0.500	0.750	2604.0	0.03432	16.74	0.300	0.790	2605.8	0.03864					

Table S3: Results for all individual scenes on all datasets.