

FAR: Flexible, Accurate and Robust 6DoF Relative Camera Pose Estimation

Chris Rockwell¹ Nilesh Kulkarni¹ Linyi Jin¹
 Jeong Joon Park¹ Justin Johnson¹ David F. Fouhey²
 University of Michigan¹ New York University²

Abstract

Estimating relative camera poses between images has been a central problem in computer vision. Methods that find correspondences and solve for the fundamental matrix offer high precision in most cases. Conversely, methods predicting pose directly using neural networks are more robust to limited overlap and can infer absolute translation scale, but at the expense of reduced precision. We show how to combine the best of both methods; our approach yields results that are both precise and robust, while also accurately inferring translation scales. At the heart of our model lies a Transformer that (1) learns to balance between solved and learned pose estimations, and (2) provides a prior to guide a solver. A comprehensive analysis supports our design choices and demonstrates that our method adapts flexibly to various feature extractors and correspondence estimators, showing state-of-the-art performance in 6DoF pose estimation on Matterport3D, InteriorNet, StreetLearn, and Map-free Relocalization. Project page: <https://crockwell.github.io/far/>

1. Introduction

Relative camera pose estimation is a fundamental problem in computer vision [23], with applications in augmented reality [27, 35, 40], robotics [48, 65, 78], and autonomous driving [9, 22]. One recent line of work learns to estimate correspondences then solve for pose [18, 41, 49, 63, 67], often offering sub-degree errors. Unfortunately, this framework tends to struggle when faced with large view change (Figure 1, left), and additionally cannot recover scale because it produces the Fundamental or Essential matrix. Another line of work learns to estimate pose directly [5, 10, 39, 60, 72], which is not as precise, but can be more robust and produces translation scale (Figure 1, left and right).

The proposed method builds upon both communities to produce a general method that is no worse than either of the options and often better than both. Critically, it leverages learned correspondence predictions as input, and combines learned pose estimation with a solver to estimate 6DoF pose. For this task, we purposefully select the Transformer,




Figure 1. **Precise and Robust 6DoF Pose Estimation.** Correspondence Estimation + Solver methods (here LoFTR [67], RANSAC [21]) produce precise outputs for moderate rotations, but are not robust to large rotations (left), and cannot produce translation scale. Learning-based methods (here LoFTR with 8-Point ViT [60] head) produce scale (right) and are more robust, but lack precision (left). FAR leverages both for precise and robust prediction, including scale.

which can handle dense features or correspondences as input. Put succinctly, the method is **Flexible**: agnostic to correspondence and feature backbone; **Accurate**: matches the precision of correspondence-based methods; and **Robust**: builds upon the resilience of learned pose methods.

FAR enables learning-based and solver-based methods to improve each other. Learned predictions are more robust than solver output, and are therefore used as a prior to bias the solver. Improved solver output, which tends to be more precise than learned output when it succeeds, is then combined with Transformer predictions to form final output. Predictions are combined via a weighting predicted by the Transformer, meaning the Transformer can learn to rely more upon either method depending on their effectiveness.

Figure 2 analyzes FAR in practice, measuring error as a function of the number of good input correspondences. With many correspondences, the solver is highly accurate, leaving little room for improvement from the prior. As the number of correspondences drop, solver performance degrades, but this can be alleviated meaningfully using the learned prior. The learned weighting also contributes to robustness, and is plotted on the right: the Transformer primarily uses solver output if there are many correspondences, and more heavily uses the regressor when there

are few correspondences (Figure 2, right). The result is a method that does not sacrifice in the case of many correspondences, but has a large gain given few correspondences.

Experiments analyze FAR in detail across a number of scenarios and datasets. First, we analyze theoretical robustness, beginning from ground truth correspondences and procedurally adding (1) noise and (2) outliers. We next evaluate the proposed method on four challenging datasets, spanning both indoors: Matterport3D and InteriorNet, and outdoors: StreetLearn and Map-free Relocalization. Across settings, the proposed method typically outperforms, or occasionally matches, the state of the art. We additionally analyze the components of FAR in ablations, and apply it to a variety of permutations of correspondence and feature estimation backbones. We also study the impact of dataset size upon model behavior.

2. Related Work

Learned Camera Pose Estimation. Learned Camera Pose Estimation has recently made impressive progress. If many views are available, camera pose can be precisely refined during SLAM [14, 69] or Visual Odometry [34, 70, 73]. If fewer views of a scene are available, methods have become increasingly robust to e.g. large rotation [27, 66, 79].

This paper focuses on the wide-baseline two-view setting, which has also strongly progresses [10, 13, 19, 76, 77]. Some of these methods also perform 3D reconstruction [1, 28, 53, 68]. We build off the 8-Point ViT [60], a SOTA method for two-image 6DoF camera pose estimation.

Correspondence Estimation. Correspondence can be learned [12, 17, 18, 26, 33, 49] or attained using classical methods [6, 43, 62], including specialized for wide-baseline stereo [45, 47, 52]. We use recent SOTA methods LoFTR [67] and SuperPoint+SuperGlue [15, 63], but note FAR can readily adapt to alternative estimators.

Camera Pose Estimation from Correspondences. Camera pose estimation from correspondences is a long standing [21] and still active problem [3]. Typically, algorithms use a robust estimator [3, 4, 21] along with the 7-Point [36] or 8-Point [24] algorithm to find the fundamental matrix, if intrinsics are unknown, or 5-Point algorithm [50] to find the essential matrix, if intrinsics are known. F or E can then be decomposed into RT (without translation scale) and estimating direction via triangulation and the chirality check. We assume known intrinsics and use RANSAC with the 5-Point algorithm, but our contributions are orthogonal to estimator.

Recent work incorporates learnable elements into pose estimation from correspondence [55]. Barroso *et al.* [5] learn to select from candidate essential matrices. Roessler and Nießner [61] use a differentiable 8-Point algorithm, while Wei *et al.* [74, 80] use a differentiable robust estimator, to improve F via end-to-end training. DSAC [7]




Figure 2. **Combining Classical and Learned.** Left: Solver output is precise given many inliers, but is poor when few are available; Updated solver output via FAR’s prior improves robustness significantly. FAR’s Transformer is less precise but more robust. The full model fuses prior-guided Solver output and Transformer output for the best of both, giving more weight to the solver when many inliers are available (right).

supervises a Scoring CNN to predict consensus and guide RANSAC. FAR is distinct from these works as its final output is a weighted combination of Solver Pose and Learned Pose. This important difference allows FAR to predict scale and improve robustness to poor or limited correspondences (Figure 2). GRelPose [31] predicts pose from correspondences, but does not use a solver, limiting precision.

3. Approach

Our goal is to predict relative camera pose, including translation scale, from two overlapping images. This 6DoF pose can be parameterized as $T \in SE(3)$, consisting of $R \in SO(3)$ and $t \in \mathbb{R}^3$. We specifically focus on predicting translation scale, which cannot be solved for from correspondences alone, in order to enable real world applications e.g. 3D reconstruction and neural rendering. The two-view case facilitates these applications on e.g. image collections. We assume known camera intrinsics as they are generally available from modern devices [2].

FAR fuses complimentary strengths from the two lines of pose estimation work: learned correspondence estimation followed by a robust solver, and end-to-end pose estimation. Critically, it produces results that are no worse than either and often better than both. We design the method as flexible to be plugged-in to existing methods with minimal change, showing improved results in a variety of settings and datasets. We outline FAR in Section 3.1, detail the learned network in Section 3.2 and how we apply a prior to the solver in Section 3.3.

3.1. Approach Outline

Figure 3 shows an overview of the proposed approach. At the heart of the approach is the Pose Transformer (Sec 3.2) which takes in dense features and outputs (1) an estimate of 6DoF pose T_t and (2) a relative weight w of this prediction. T_t is then combined with solver-estimated pose T_s using weight w to obtain pose estimate T_1 . T_1 is used as a prior

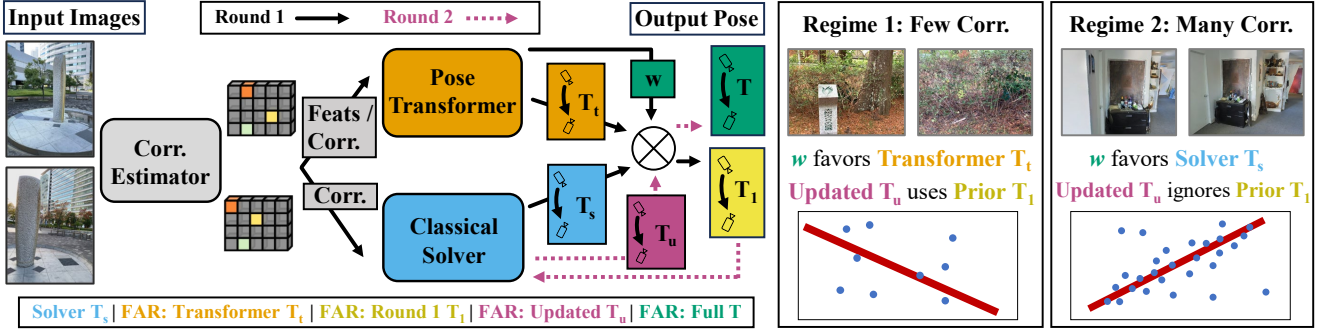


Figure 3. **Overview.** Given dense features and correspondences, FAR’s Transformer produces camera poses (in square boxes \blacksquare) through a transformer (round box \blacksquare) and classical solver (round box \blacksquare). In the first round, the solver produces a pose \mathbf{T}_s . FAR’s pose transformer averages this with its own prediction \mathbf{T}_t via weight \mathbf{w} , to yield the round 1 pose \mathbf{T}_1 . \mathbf{T}_1 pose serves as a prior for the classic solver, which produces an updated pose \mathbf{T}_u . This is combined with an additional estimate of \mathbf{T}_t and weight \mathbf{w} to produce the final result \mathbf{T} . With few correspondences, \mathbf{T}_1 helps solver output, while the network learns to weigh Transformer predictions more heavily; with many correspondences, solver output is often good, so the network relies mostly on solver output.

for the solver, resulting in updated solver output \mathbf{T}_u , which is combined with \mathbf{T}_t via \mathbf{w} to get the final output \mathbf{T} .

This architecture enables the network to learn to behave differently depending on the data regime. In the case of many, high-quality correspondences, classical solvers are typically precise, so the prior has little impact on the solver, while the network learns to heavily rely upon solver output via a low \mathbf{w} . In the case of few, low-quality correspondences, solvers degrade, so the prior is designed to have a strong influence on solver output, while the network relies more heavily on the transformer predictions (high \mathbf{w}).

The approach is agnostic to input features and correspondences. In experiments (Sec 4), we show success with features from three feature estimation methods [2, 60, 67] and correspondences from two correspondence estimators [63, 67]. We use a ViT [16] to handle spatial features, and losses can backpropagate through the backbone. We also explore using only correspondences and descriptors as input.

3.2. Pose Transformer

The goal of the Transformer is to estimate (1) 6DoF relative camera pose \mathbf{T}_t between two wide-baseline images and (2) weight $\mathbf{w} \in [0, 1]$ of its estimate vs. solver estimates from a set of 2D correspondence matches $\mathbb{M} = \{(\mathbf{p}, \mathbf{q}) \mid \mathbf{p}, \mathbf{q} \in \mathbb{R}^2 \text{ and optionally dense 2D image-wise features } f_i, f_j\}$. Given predicted camera pose and weight, the final output is the linear combination of Transformer pose \mathbf{T}_t and solver pose \mathbf{T}_s weighted by \mathbf{w} . We use separate weights for translation w_t and rotation w_r , allowing the Transformer to have different confidences for two subtly different problems.

Two challenges to this setup are that linear combinations of rotations are often not rotation matrices, and solver translation does not have scale. To address the former, we represent pose in the 6D coordinate system of Zhou *et al.* [81], which enables us to combine in 6D space before computing

a rotation matrix using Gram–Schmidt orthogonalization. To address scale-less solver output, we scale translation \mathbf{t}_s by the Transformer predicted translation magnitude $\|\mathbf{t}_t\|$, before linearly combining. We find this stabilizes training compared to first averaging predicted the angles of \mathbf{t}_s and $\mathbf{t}_{t,f}$ and then applying scale to normalized predictions. Our final formula is:

$$\begin{aligned} \hat{\mathbf{R}} &= w_r \mathbf{R}_t + (1 - w_r) \mathbf{R}_s \\ \hat{\mathbf{t}} &= w_t \mathbf{t}_t + (1 - w_t) \|\mathbf{t}_t\| \mathbf{t}_s \end{aligned} \quad (1)$$

Transformer Backbone. We use two distinct architectures to span possible inputs: if features are available, we use a modified ViT. If only correspondences are available, we use a Vanilla Transformer. This means the method can be used with correspondence or regression-based methods producing dense features, while accommodating methods only outputting correspondence. In each case, the Transformer produces features f_o used as input to two MLP heads.

8-Point ViT. This network takes as input pairwise dense features f_i, f_j and produces a feature vector f_o . It consists first of one LoFTR [67] self-attention and cross-attention layer followed by an 8-Point ViT cross-attention layer [16, 44, 60]; both networks are aimed at producing good features for pose estimation. For detailed architectures of each, see the original works.

Vanilla Transformer. This network takes as input a set of correspondences $\mathbb{M} = \{(\mathbf{p}, \mathbf{q}) \mid \mathbf{p}, \mathbf{q} \in \mathbb{R}^2 \text{ including associated descriptors, if available, and produces a set of features } f_o\}$. We use a vanilla Transformer encoder with N layers, and map correspondences and descriptors as input. We encode correspondences in a sinusoidal manner with K bands, followed by a linear mapping to a size of c input to the Transformer. If descriptive features of each correspondence point are available, of dimension $d < c$, a Linear layer maps them to $\frac{c}{4}$ and they are concatenated to correspondence locations which are linearly mapped to $\frac{3c}{4}$.

The Vanilla Transformer can also be used with dense features f as input. If networks produce a joint feature encoding for two images, the Transformer can be applied directly to low-resolution features, without positional encoding. This occurs in the Regression model of Arnold *et al.* [2], which we build upon in Table 5.

Regression MLP. This MLP maps Transformer features f_0 to $\mathbf{R} \in \mathbb{R}^6$ and $\mathbf{t} \in \mathbb{R}^3$ using two hidden layers.

Gating MLP. This takes as input Transformer features concatenated with Regression MLP predictions and predictions from the classical solver, along with the number of inlier correspondences in the solver output, using several thresholds. Predictions and number of inliers are normalized then input as scalar features. As we analyzed in Figure 2, the number of inlier correspondences is highly correlated with the performance of solver pose estimate \mathbf{T}_s . The Gating MLP has two hidden layers and ends with a Sigmoid, producing $w_t, w_r \in (0, 1)$.

3.3. Prior-Guided Robust Pose Estimator

Now, having shown how the solver can help learning, how can the learning based methods help the solver? The performance of search-based solver methods like RANSAC [21] is driven by *searching* over a model space by sampling valid hypothesis and then ranking them based on a *scoring function*. The scoring function serves as the measure of probability of data under a hypothesis [71]. It's typical to use such solvers when estimating pose from a set of correspondences, but direct application of these methods can not be robust when correspondence estimation is done with a scarce set. Our key idea is to use the predicted pose estimate, \mathbf{T}_1 , to influence both the *search* and the *scoring* function to help in data scarce scenarios.

We take inspiration from existing lines of work in using learning to better inform sampling and selection in RANSAC-like algorithms [3, 4, 7, 54, 71]. We show that we can recycle estimates from a learning-based model and plug these estimates in simplistically. Specifically, an initial estimated pose, \mathbf{T}_1 , to modify the search function so as to sample more hypothesis close to \mathbf{T}_1 . Secondly, we modify the scoring function to consider the distance to the \mathbf{T}_1 along with inlier count.

RANSAC Preliminaries. The typical approach to pose estimation from correspondences applies random sample consensus (and variants) e.g. RANSAC, USAC [54] or MAGSAC [3, 4] to model fitting e.g. 5, 7, or 8-Point algorithms [24, 37, 42, 50]. These methods use a notion of epipolar distance such as Sampson Error [23] for inlier thresholding (soft and hard). More concretely, given a set of 2D correspondence matches $\mathbb{M} = \{(\mathbf{p}, \mathbf{q}) \mid \mathbf{p}, \mathbf{q} \in \mathbb{R}^2\}$, a minimal subset of points is randomly sampled to fit a model \mathbf{H} via an n-point algorithm. The scoring function counts number of inliers that have Sampson Error less than

a fixed threshold σ . Given, hypothesis \mathbf{H} and set \mathbb{M} of correspondences, $\mathbf{E}(\mathbf{p}, \mathbf{q} \mid \mathbf{H})$, is the Sampson Error between points \mathbf{p} and \mathbf{q} under \mathbf{H} . The scoring function is defined as $\text{score}(\mathbf{H}) = \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathbb{M}} \mathbb{1}(\mathbf{E}(\mathbf{p}, \mathbf{q} \mid \mathbf{H}) < \sigma)$. Sampling repeats up to N times or until stopping heuristics are met for efficiency [54], and the highest scoring model, e.g. the one with the most inliers, is selected. Works like MAGSAC [3], MAGSAC++ [4] have shown that improving scoring functions to show better performance. For simplicity, we continue the exposition with thresholding based function that are popular with classic RANSAC [21].

Limitations in Few-Correspondence Case. The heuristic score of counting inliers typically is not effective especially in the low-correspondence case [5]. When the number of correspondences is only a small multiple of the number of points needed to minimally define a model the algorithm becomes particularly unreliable. Consider the extreme case of doing pose recovery with calibrated cameras from nine points, of which five are inliers. The minimal subset for pose estimation is five points, and so while the true model will have five inliers, so will any other sampled hypothesis. Accordingly, the result will be random hypothesis.

Prior-Guided Estimator. We propose to incorporate a learning based predictions to aid the solver in the case of few or poor correspondences. We operationalize this by incorporating a prior model that estimates the likelihood of hypothesis under the network's prediction using a function $\beta(\cdot \mid \mathbf{T}_1)$. The $\beta(\mathbf{H} \mid \mathbf{T}_1)$ measures the log probability of the hypothesized model \mathbf{H} under \mathbf{T}_1 . We found it difficult to measure probabilities in rotation and translation and weigh them, so as a proxy, we compare how the models transform a fixed set of grid points. In particular, we measure the negative of average squared distance between a fixed set of grid points transformed by \mathbf{T}_1 and the same fixed grid transformed by \mathbf{H} . See Supplemental for details.

Now, we show how this β function can alter the scoring. The modified scoring function measures the likelihood of the hypothesis \mathbf{H} under \mathbf{T}_1 along with measuring the likelihood of data [71] under \mathbf{H} . It is defined as,

$$\text{score}(\mathbf{H}) = \alpha \beta(\mathbf{H} \mid \mathbf{T}_1) + \sum_{\{\mathbf{p}, \mathbf{q}\} \in \mathbb{M}} \mathbb{1}(\mathbf{E}(\mathbf{p}, \mathbf{q} \mid \mathbf{H}) < \sigma) \quad (2)$$

which is the (log) product of probability of the hypothesis given our β prior function and the probability of the the data, \mathbb{M} , under \mathbf{H} . We weigh the prior with a scalar, $\alpha \in \mathbb{R}$. In this setting, the prior tie-breaks ambiguous cases where two hypotheses have similar numbers of inliers, but has diminishing influence as $|\mathbb{M}|$ gets bigger. As $|\mathbb{M}| \rightarrow \infty$, the prior's impact is washed out entirely. This formulation has the desired impact of significant effect when correspondences are few and unbiased hypotheses are poor, and little impact when correspondences are many.




Figure 4. **Ground Truth Robustness Study on Matterport3D.** Using true correspondence, the solver is nearly perfect. Adding noise or outliers, it quickly degrades, while prior-guided Updated solver is robust to outliers and the Transformer is robust to noise. FAR matches or beats all methods across settings.

Sampling Good Hypotheses. Randomly sampling points and estimating H is unlikely to lead to hypothesis consistent with the model T_1 . To increase the chance to sampling consistent hypothesis we want to sample a minimal subset that best agrees with the model. We achieve this by weighing the correspondences by their agreement with the model in turn influencing the as $w(\mathbf{p}, \mathbf{q}) = \exp(-\text{Sampson}(\mathbf{p}, \mathbf{q} | T_1) / \tau)$.

In practice, we sample half the hypothesis using biased sampling use uniform sampling for the other half. This improves sample diversity in the case of many correspondences, in which case unbiased sampling is very effective.

3.4. Implementation and Training Details

Across experiments, we use the Adam [32] optimizer. The 8-Pt ViT trains for about 300k iterations or 7 days on 10 GTX 1080Ti; Vanilla TF trains for about 600k iterations or 3 days. We select the checkpoint with the lowest validation mean rotation error, which tends to be marginally more stable than translation error. We represent rotation in 6D coordinates [81] and use L1 loss. Models are trained stage-wise: first we train the Transformer to estimate pose, then to estimate pose jointly with a vanilla solver, then to estimate pose jointly with the prior-based solver. We find this progressive training improves final performance. We implement in PyTorch [51] Lightning [20], using TIMM [75] for the ViT. **8-Point ViT.** We train 8-Point ViT end-to-end with the feature extraction backbone. We found including a self and cross-attention LoFTR layer significantly improved learning capacity, while additional layers did not help further.

Vanilla Transformer. We use a 6-layer encoder with 8 heads and 512 feature size followed by global average pooling. We use $K = 42$ bands for positional encoding, and linearly map them to size 384, concatenating descriptive features linearly mapped to size 128. We found random dropout on correspondences with $p = 0.1$ helps performance. We cache correspondences for fast training. Training speed is ≈ 12 iterations per second on a GTX 1080Ti.

Table 1. **Camera Pose Estimation on Matterport3D.**

Correspondence-based methods have low median but high mean error, and do not produce translation scale. Regression-based methods are less precise but produce scale. FAR builds upon both, resulting in low median and mean error, with translation scale.

Method	Translation (m)			Rotation (°)		
	Med.↓	Avg.↓	≤1m↑	Med.↓	Avg.↓	≤30↑
[58] + [56]	3.34	4.00	8.3	50.98	57.92	29.9
Assoc.3D [53]	2.17	2.50	14.8	42.09	52.97	38.1
Sparse Planes [28]	0.63	1.25	66.6	7.33	22.78	83.4
PlaneFormers [1]	0.66	1.19	66.8	5.96	22.20	83.8
8-Point ViT [60]	0.64	1.01	67.4	8.01	19.13	85.4
NOPE-SAC-Reg [68]	0.52	0.94	73.2	2.77	14.37	89.0
SuperGlue [63]	-	-	-	3.88	24.17	77.8
LoFTR [67]	-	-	-	0.23	9.49	91.4
LoFTR+Reg. Scale	0.85	1.21	56.3	0.26	9.66	91.2
FAR (Vanilla TF)	0.37	0.67	81.9	0.26	6.14	94.2
FAR	0.25	0.49	89.2	0.20	4.93	95.8

Prior-Guided Estimator. We implement in Kornia [59] and use 2k random samples without early stopping and inlier threshold on L2 Sampson Error σ of 3×10^{-7} , finding these results most closely matched OpenCV [8] output using LoFTR settings. For the Prior, we use $\tau = 0.1$ and $\alpha = 3.33$ found through grid search on the validation set.

4. Experiments

We design our experiments to measure the effectiveness of FAR in achieving our stated goals: flexible, accurate and robust 6DoF pose estimation. We first validate robustness by measuring model performance as a function of increasingly perturbed ground truth. We next test precision to moderate view change and robustness to large view change by comparing to the state of the art in wide-baseline relative pose. Having demonstrated accuracy and robustness, we next verify model flexibility to choice (or lack) of dense feature method, correspondence estimation method, and dataset size. Finally, we compare to the state of the art on additional indoor and outdoor datasets.

4.1. Robustness to Correspondence Perturbations

Here, we assume the image correspondences are given and study how variations of our method and baselines perform with varying noise-levels applied to the correspondences.

Dataset. We use image pairs collected using the Habitat [64] embodied simulator upon Matterport3D [11], following the setup of Jin *et al.* [28]. It has 32k train / 5k val / 8k test pairs with small to moderate overlap (average 53° rotation, 2.3m translation, 21% overlap). The variety in view change enables the study of both precision upon moderate cases and robustness to highly challenging cases.

Metrics. Throughout Matterport3D experiments, we report three metrics for rotation and translation: median error, mean error, and percentage of errors within a threshold. These are standard metrics, which identify our two quali-

Table 2. **Ablations on Matterport3D.** (Top) We improve significantly upon LoFTR using a combination of learned and classical. (Middle) This result holds for the case of no input features, where we use the Vanilla TF. (Bottom) Scaling Solver translation is important to FAR performance; selecting separate weights for R and T improves robustness.

	Translation (m)			Rotation ($^{\circ}$)		
	Med. \downarrow	Avg. \downarrow	$\leq 1m\uparrow$	Med. \downarrow	Avg. \downarrow	$\leq 30\uparrow$
<i>Transformer: 8-Point ViT</i>						
LoFTR + Solver + Scale T_s	0.85	1.21	56.3	0.26	9.66	91.2
FAR: Transformer T_t	0.38	0.64	85.4	4.51	9.94	94.2
FAR: One Round T_1	0.25	0.49	<u>89.0</u>	0.20	<u>5.08</u>	<u>95.7</u>
FAR: Updated T_u	0.25	0.50	88.4	0.20	5.35	95.0
FAR: Full T	0.25	0.49	89.2	0.20	4.93	95.8
<i>Transformer: Vanilla TF</i>						
LoFTR + Solver + Scale T_s	0.85	1.21	56.3	<u>0.26</u>	9.66	91.2
FAR: Transformer T_t	0.42	0.75	79.1	3.87	10.8	92.5
FAR: One Round T_1	0.37	0.67	<u>81.8</u>	<u>0.26</u>	<u>6.41</u>	<u>93.8</u>
FAR: Updated T_u	0.37	0.68	81.5	0.25	6.69	93.7
FAR: Full T	0.37	0.67	81.9	<u>0.26</u>	6.14	94.2
<i>Prediction Selection</i>						
Unscaled Solver t_s	0.31	0.55	87.4	0.21	5.09	95.8
One Weight ($w_r = w_t$)	0.25	0.50	88.7	0.20	5.04	95.8
FAR	0.25	0.49	89.2	0.20	4.93	95.8

ties of interest: precision (median) and robustness (mean and percentage). We follow prior work [28, 68] in using rotation threshold of 30° and translation threshold of 1m. For the ground truth study, for brevity we report median error across a variety of settings, which is an indicative summary of performance. Additional results are in Supplemental.

Setup. Beginning with ground truth correspondences, we (1) Apply Gaussian noise with standard deviation from 0 to 32 pixels, upon 480x640 images; (2) Replace true correspondences with randomly sampled coordinates (outliers), with $P(\text{Outliers})$ from 0 to 0.875. Models are trained and evaluated upon each noise and outlier setting independently; e.g. FAR is trained and evaluated four times to make Correspondence Noise Graph in Figure 4, left.

Ablations. We consider the following cases:

- (1) *Solver T_s .* Using LoFTR’s solver, i.e., RANSAC [21]+5-Point Algorithm [50]
- (2) *FAR: Transformer T_t .* Pose Transformer output pose. This is a Vanilla Transformer, since dense features are not available as input; only correspondences.
- (3) *FAR: Updated T_u .* Solver output using FAR’s Prior
- (4) *FAR: Full T.* Full FAR: learned combin. of T_u and T_t

Ablation Results. Figure 4 shows Solver has nearly perfect results on ground truth correspondences but is not robust to noise or many outliers. FAR: Transformer is less precise on ground truth but is more robust as outlier frequency and particularly noise increases. The prior is highly effective at leaving FAR: Updated robust to outliers, showing close to 0° median error even with 87.5% outliers. The full method offers the best of all: precise estimation on ground truth correspondences with the best or equal to best robustness to noise and outliers.

Table 3. **Approach Flexibility to Features and Correspondences.** FAR yields improvement using features from 8-Pt ViT or LoFTR; and correspondences from SuperGlue or LoFTR.

Feats.	Corr.	Pose Est.	Translation (m)			Rotation ($^{\circ}$)		
			Med. \downarrow	Avg. \downarrow	$\leq 1m\uparrow$	Med. \downarrow	Avg. \downarrow	$\leq 30\uparrow$
8-Pt ViT	-	8-Pt ViT	0.64	1.01	67.4	8.01	19.1	85.4
8-Pt ViT	SuperGlue	FAR	0.62	1.01	<u>68.3</u>	7.02	16.6	<u>86.8</u>
8-Pt ViT	LoFTR	FAR	<u>0.63</u>	1.01	68.5	<u>7.06</u>	<u>17.0</u>	86.9
LoFTR	LoFTR	RANSAC+5Pt	-	-	-	<u>0.23</u>	9.49	91.4
LoFTR	LoFTR	FAR (Vanilla TF)	<u>0.37</u>	<u>0.67</u>	<u>81.9</u>	0.26	<u>6.14</u>	<u>94.2</u>
LoFTR	LoFTR	FAR	0.25	0.49	89.2	0.20	4.93	95.8




Figure 5. **Evolving with Dataset Size.** The Transformer learns to rely more heavily upon the solver if data is limited (40% data size), and learns to use Transformer pose estimations as data scales and performance improves (100% data size).

4.2. Wide-Baseline Pose on Matterport3D

In this section, we use the same Matterport3D dataset and the metrics used in Sec. 4.1, but the inputs are images rather than the GT correspondences.

Baselines. We compare against state-of-the-art solver-based and learning-based baselines. For solver-based methods, we choose the popular LoFTR [67] and SuperGlue [63]. In the learned space, we compare to end-to-end classical-estimation-inspired ViT, 8-Point [60]; planar mapping and optimization methods NOPE-SAC [68], PlaneFormers [1] and Sparse Planes [28]; and 3D reconstruction method Associative3D [53]. In this set of experiments our FAR builds upon LoFTR backbone and correspondences. We additionally report results using only correspondence and descriptor as input, as “FAR (Vanilla TF)”.

Results. Table 1 shows the quantitative results on Matterport3D. Among the prior works, end-to-end methods such as 8-Point ViT [60] perform well in absolute Translation, while correspondence-solver methods e.g. LoFTR [67] perform best in rotation. FAR sets a new standard in both metrics, surpassing the best prior baseline (NOPE-SAC-Reg) by a large margin. It reduces the median and mean translation errors by about 50%: from 0.52 to 0.25 and from 0.94 to 0.49, respectively. Additionally, it decreases the mean rotation error by almost 50% compared to the best prior work (LoFTR), from 9.66 to 4.93. Even with only correspondence available as input, “FAR (Vanilla TF)” is typically better than all prior work by a large margin.

Ablations. To investigate the source of FAR’s out-




Figure 6. **Rotation Error on InteriorNet and StreetLearn.** Even when Correspondence + Solver method is relatively poor, we can still leverage it to improve regression results. “Learning-Based”: 8-Point ViT [60]. “Corr. + Solver”: LoFTR [67].

performance, we conduct ablations on Matterport3D using the same setup of Sec. 4.1, except correspondences are predicted. In addition to ablations discussed in Section 4.1, we consider model output after One Round (T_1) to study the impact of the prior-guided solver upon the final model. We further explore performance by using two different Transformer architectures: the dense feature-based Transformer: 8-Point ViT (referred to as FAR in other experiments), and the correspondence-only Transformer: Vanilla TF. Finally, we evaluate the design choices outlined in Section 3.2 under the Prediction Selection category: Unscaled Solver t_s , where Solver translation is not scaled by Transformer magnitude; and One Weight ($w_r = w_t$), which uses equal weights for Transformer translation and rotation prediction.

Ablation Results. As shown in Tab. 2, we can clearly observe the same trends from Fig. 4: the solver is precise in most cases characterized by low median rotation error. However the solver suffers from high mean error (due to outliers) and poor translation errors. Incorporating FAR’s prior significantly improves the solver’s mean rotation error. In contrast, Transformer regression outputs are not nearly as precise, with median rotation error of above 4° , but it reduces the ratios of large errors (those greater than $1m$ or 30°). FAR enhances the best results achieved by both the Transformer and Solver. These patterns hold true for the Vanilla Transformer as well. In *Prediction Selection*, we see predicting Solver translation scale is important for translation performance, while separate weighting for rotation and translation improves robustness.

4.3. Approach Flexibility

We then evaluate the flexibility of FAR in terms of the feature extractor, correspondence estimator, and dataset size.

Dataset and Metrics. We continue using the Matterport3D dataset and metrics as in Section 4.1.

Alternative Approaches. To assess the versatility of FAR, we explore options that are orthogonal to our core contribution. Specifically, we examine three settings for feature estimation: the recent SOTA methods LoFTR and 8-Point ViT,




Figure 7. **Error on Map-free Relocalization.** FAR leverages Solver output to improve regression results on this highly challenging dataset. “Corr. + Solver + Scale”: LoFTR + DPT [57] trained on KITTI [22]. “Learning-Based”: 6D Reg.

as well as a scenario without dense features. Additionally, we evaluate two recent SOTA settings for correspondence estimation: LoFTR, and SuperPoint [15] + SuperGlue [63].

Results. Table 3 shows FAR improves upon 8-Pt ViT, using either SuperGlue or LoFTR correspondences. Similarly, FAR improves LoFTR, whether it employs both LoFTR features and correspondences or just the correspondences.

Dataset Scaling. We next present the proposed method when trained on a version of the Matterport3D dataset that has been randomly subsampled to 40% of its original data size. In Figure 5, we compare rotation error (left) and solver weight (right) of 40% size and full size. Note that as the training dataset size increases from 40% to 100%, both the solver weight and error decrease. This trend aligns with expectations: as the Transformer’s estimated pose accuracy improves with more training data, it gains a larger influence in the final output, enhancing overall performance. The result suggests a fixed weighting of learned and solver output is not sufficient for best results.

4.4. Wide-Baseline Pose on Additional Datasets

We evaluate our method’s performance on various datasets to assess its versatility. We follow Cai *et al.* [10] and use InteriorNet [38], a synthetic collection of indoor home scenes, and StreetLearn [46], which features outdoor city street photos. Both datasets consist of 90° field-of-view image crops upon panoramas. Image pairs are chosen from different panoramas with varying overlaps, facilitating the assessment of precision and robustness in scenarios with both large ($> 45^\circ$) and small ($< 45^\circ$) overlaps. Additionally, we use Map-free Relocalization [2], a challenging dataset of user-collected videos surrounding outdoor places of interest e.g. sculptures or fountains. SfM has been applied to the videos, so translation with scale can be evaluated.

Metrics. For InteriorNet and StreetLearn, we report rotation error only, in line with prior work [10, 60], using a 10° threshold. For Map-free Relocalization, we calculate median translation and rotation errors per video, then average these. We also include the Virtual Correspondence Repro-

Table 4. **Rotation Performance on InteriorNet and StreetLearn.** Correspondence-based methods (top) struggle to generalize to this data, while regression methods learn helpful features. Building off 8-Point ViT features, we can still utilize LoFTR correspondences to boost performance. Errors were calculated only on successful pairs for SIFT and SuperPoint; gray text indicates failure over 50% of pairs.

Method	InteriorNet						StreetLearn					
	Large Overlap (°)			Small Overlap (°)			Large Overlap (°)			Small Overlap (°)		
	Med.↓	Avg.↓	≤ 10↑	Med.↓	Avg.↓	≤ 10↑	Med.↓	Avg. ↓	≤ 10↑	Med.↓	Avg.↓	≤ 10↑
SIFT* [43]	2.95	7.78	55.5	10.0	18.2	18.5	3.13	18.9	22.4	13.8	38.8	5.7
SuperPoint* [15]	2.79	5.46	65.9	5.82	11.6	11.7	1.79	6.38	16.5	6.85	6.80	1.0
LoFTR [67]	0.54	<u>1.85</u>	97.0	2.64	14.3	70.4	24.8	36.4	31.6	51.2	58.6	19.9
Reg6D [81]	6.91	10.5	67.8	11.4	21.9	44.1	6.02	12.3	69.1	7.59	15.1	63.4
Cai <i>et al.</i> [10]	1.10	2.89	97.6	<u>1.38</u>	10.2	89.8	2.91	9.12	87.5	3.49	13.0	84.2
8-Point ViT [60]	1.83	2.90	<u>97.9</u>	2.38	<u>4.48</u>	96.3	<u>2.43</u>	<u>4.08</u>	<u>90.1</u>	<u>3.25</u>	<u>9.19</u>	<u>87.7</u>
FAR	<u>0.60</u>	1.16	98.5	1.22	3.42	<u>95.4</u>	1.81	3.01	96.7	2.07	7.89	92.4

Table 5. **6DoF Performance on Map-free Relocalization.** We compare against the strongest baselines from [2]; for all comparisons see the original paper. FAR uses the 6D Reg backbone, adding LoFTR or SuperGlue correspondences to beat 6D Reg.

Method	Pose Error				VCRE		
	Avg. Med.↓	Prec.↑	AUC↑	Avg. Med.↓	Prec.↑	AUC↑	
LoFTR	199cm	30.6°	0.15	<u>0.35</u>	168px	0.35	0.61
SuperGlue	188cm	25.4°	<u>0.17</u>	<u>0.35</u>	160px	0.36	0.60
Reg Ang. [2]	210cm	33.7°	0.09	-	200px	0.30	-
6D Reg [2]	168cm	22.5°	0.06	-	147px	0.40	-
FAR (SG)	<u>149cm</u>	17.2°	<u>0.17</u>	<u>0.35</u>	135px	0.44	<u>0.67</u>
FAR (LoFTR)	148cm	<u>18.1°</u>	0.18	0.39	<u>137px</u>	0.44	0.68

jection Error (VCRE) metric to measure reprojection errors (see [2] for details).

Baselines. We compare our method with SOTA correspondence and pose estimation techniques. For InteriorNet and StreetLearn, we compare to Cai *et al.*'s [10] correlation volume-based learning, SuperPoint [15], and the classical SIFT method [43]. We use LoFTR adapted for InteriorNet using Matterport3D, and for StreetLearn using MegaDepth, due to the lack of depth data in these datasets for training correspondences. Since LoFTR cannot be finetuned, we find 8-Point ViT features are more effective, and use these as input to FAR, along with LoFTR correspondences.

Arnold *et al.* [2] train a variety of pose estimation methods on Map-free, including “6D Reg”, which creates a correlation volume [29, 30] and warps accordingly, followed by a ResNet [25], and is supervised upon 6D rotation [81].

Results. Tab. 4 and Fig. 6 show that 8-Point ViT [60] achieves impressive mean errors, under 5°, on InteriorNet, even for small overlap pairs. FAR still adds precision on top of the 8-Point ViT. On the challenging StreetLearn data, FAR significantly outperforms the state of the art, despite LoFTR not generalizing well to StreetLearn.

6D Reg is the strongest overall baseline on Map-free Relocalization, so we use its features for FAR, taking correspondences from LoFTR or SuperGlue. In both cases, FAR improves upon 6D Reg and other methods (see Tab. 5 and

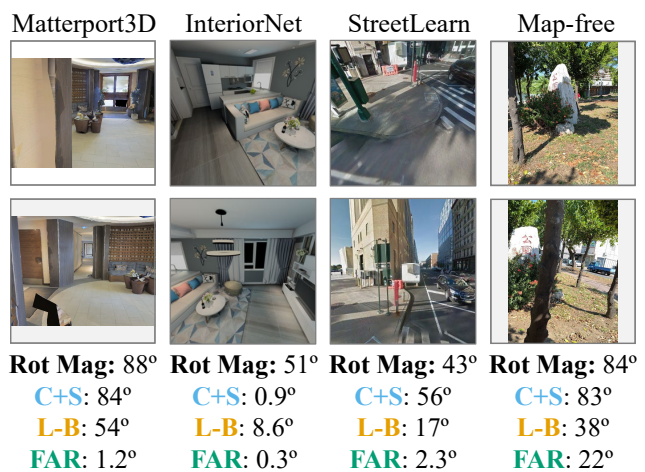


Figure 8. **Success Cases.** For some challenging wide-baseline image pairs, our method often dramatically outperforms the baselines. “Learning-Based”: LoFTR [67] with 8-Pt. ViT [60] head (Matterport3D), 8-Pt. ViT (InteriorNet, StreetLearn), 6D Reg [2] (Map-free). “Corr. + Solver”: LoFTR.

Fig. 7). These results across datasets show FAR’s adaptability to different backbones and its robustness to sub-optimal correspondence estimates, highlighted in Fig. 8.

5. Conclusion

In this work, we address wide-baseline 6DoF relative camera pose estimation. Our proposed method is simple yet potent, merging the best aspects of correspondence and learning-based methods. This results in precise and robust outcomes, adaptable to various backbones and solvers.

Limitations and Societal Impact. FAR consists of several components and implements Prior-Guided RANSAC in Kornia, slowing inference speed to 3.3 it/sec on 10 1080Ti GPUs; analysis vs. other methods appears in Supplemental. Training upon affluent homes of Matterport3D can result in worse performance on more typical residences.

Acknowledgments. Toyota Research Institute provided funds to support this work.

References

- [1] Samir Agarwala, Linyi Jin, Chris Rockwell, and David F. Fouhey. PlaneFormers: From sparse view planes to 3d reconstruction. In *ECCV*, 2022. 2, 5, 6
- [2] Eduardo Arnold, Jamie Wynn, Sara Vicente, Guillermo Garcia-Hernando, Aron Monszpart, Victor Prisacariu, Daniyar Turmukhambetov, and Eric Brachmann. Map-free visual relocalization: Metric pose relative to a single image. In *ECCV*, 2022. 2, 3, 4, 7, 8
- [3] Daniel Barath, Jiri Matas, and Jana Noskova. MAGSAC: marginalizing sample consensus. In *CVPR*, 2019. 2, 4
- [4] Daniel Barath, Jana Noskova, Maksym Ivashechkin, and Jiri Matas. MAGSAC++, a fast, reliable and accurate robust estimator. In *CVPR*, 2020. 2, 4
- [5] Axel Barroso-Laguna, Eric Brachmann, Victor Adrian Prisacariu, Gabriel J Brostow, and Daniyar Turmukhambetov. Two-view geometry scoring without correspondences. In *CVPR*, 2023. 1, 2, 4
- [6] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In *ECCV*, 2006. 2
- [7] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC-differentiable RANSAC for camera localization. In *CVPR*, 2017. 2, 4
- [8] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. 5
- [9] Guillaume Bresson, Zayed Alsayed, Li Yu, and Sébastien Glaser. Simultaneous localization and mapping: A survey of current trends in autonomous driving. *T-IV*, 2017. 1
- [10] Ruojin Cai, Bharath Hariharan, Noah Snavely, and Hadar Averbuch-Elor. Extreme rotation estimation using dense correlation volumes. In *CVPR*, 2021. 1, 2, 7, 8
- [11] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. In *3DV*, 2017. 5
- [12] Hongkai Chen, Zixin Luo, Lei Zhou, Yurun Tian, Mingmin Zhen, Tian Fang, David Mckinnon, Yanghai Tsin, and Long Quan. ASpanFormer: Detector-free image matching with adaptive span transformer. In *ECCV*, 2022. 2
- [13] Kefan Chen, Noah Snavely, and Ameesh Makadia. Wide-baseline relative camera pose estimation with directional learning. In *CVPR*, 2021. 2
- [14] Jan Czarnowski, Tristan Laidlow, Ronald Clark, and Andrew J Davison. DeepFactors: Real-time probabilistic dense monocular SLAM. *RA-L*. 2
- [15] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-supervised interest point detection and description. In *CVPRW*, 2018. 2, 7, 8
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 3
- [17] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In *CVPR*, 2019. 2
- [18] Johan Edstedt, Ioannis Athanasiadis, Mårten Wadenbäck, and Michael Felsberg. DKM: Dense kernelized feature matching for geometry estimation. In *CVPR*, 2023. 1, 2
- [19] Sovann En, Alexis Lechervy, and Frédéric Jurie. RpNet: An end-to-end network for relative camera pose estimation. In *ECCVW*, 2018. 2
- [20] William Falcon and The PyTorch Lightning team. PyTorch Lightning, 2019. 5
- [21] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981. 1, 2, 4, 6
- [22] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 1, 7
- [23] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1, 4
- [24] Richard I Hartley. In defense of the eight-point algorithm. *TPAMPI*, 19(6):580–593, 1997. 2, 4
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8
- [26] Dihe Huang, Ying Chen, Yong Liu, Jianlin Liu, Shang Xu, Wenlong Wu, Yikang Ding, Fan Tang, and Chengjie Wang. Adaptive assignment for geometry aware local feature matching. In *CVPR*, 2023. 2
- [27] Hanwen Jiang, Zhenyu Jiang, Kristen Grauman, and Yuke Zhu. Few-view object reconstruction with unknown categories and camera poses. *arXiv preprint arXiv:2212.04492*, 2022. 1, 2
- [28] Linyi Jin, Shengyi Qian, Andrew Owens, and David F Fouhey. Planar surface reconstruction from sparse views. In *ICCV*, 2021. 2, 5, 6
- [29] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, 2017. 8
- [30] Sameh Khamis, Sean Fanello, Christoph Rhemann, Adarsh Kowdle, Julien Valentin, and Shahram Izadi. StereoNet: Guided hierarchical refinement for real-time edge-aware depth prediction. In *ECCV*, 2018. 8
- [31] Fadi Khatib, Yuval Margalit, Meirav Galun, and Ronen Basri. GRelPose: Generalizable end-to-end relative camera pose regression. *arXiv preprint arXiv:2211.14950*, 2022. 2
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 5
- [33] Dominik A Kloepper, Joao F Henriques, and Dylan Campbell. Scenes: Subpixel correspondence estimation with epipolar supervision. 2024. 2
- [34] Lei Lai, Zhongkai Shangguan, Jimuyang Zhang, and Eshed Ohn-Bar. Xvo: Generalized visual odometry via cross-modal self-training. In *ICCV*, 2023. 2
- [35] Zihang Lai, Sifei Liu, Alexei A Efros, and Xiaolong Wang. Video autoencoder: self-supervised disentanglement of static 3d structure and motion. In *ICCV*, 2021. 1

- [36] Viktor Larsson, Magnus Oskarsson, Kalle Astrom, Alge Wallis, Zuzana Kukelova, and Tomas Pajdla. Beyond grobner bases: Basis selection for minimal solvers. In *CVPR*, 2018. 2
- [37] Hongdong Li and Richard Hartley. Five-point motion estimation made easy. In *ICPR*, 2006. 4
- [38] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. InteriorNet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. In *BVMC*, 2018. 7
- [39] Amy Lin, Jason Y Zhang, Deva Ramanan, and Shubham Tulsiani. RelPose++: Recovering 6d poses from sparse-view observations. *arXiv preprint arXiv:2305.04926*, 2023. 1
- [40] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: Bundle-adjusting neural radiance fields. In *ICCV*, 2021. 1
- [41] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. 1
- [42] H Christopher Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133–135, 1981. 4
- [43] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*. 2, 8
- [44] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 32, 2019. 3
- [45] Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*, 22(10):761–767, 2004. 2
- [46] Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, Denis Teplyashin, Karl Moritz Hermann, Mateusz Malinowski, Matthew Koichi Grimes, Karen Simonyan, Koray Kavukcuoglu, Andrew Zisserman, et al. The StreetLearn environment and dataset. *arXiv preprint arXiv:1903.01292*, 2019. 7
- [47] Dmytro Mishkin, Jiri Matas, Michal Perdoch, and Karel Lenc. WxBS: Wide baseline stereo generalizations. *BMVC*, 2015. 2
- [48] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: a versatile and accurate monocular SLAM system. *T-RO*. 1
- [49] Junjie Ni, Yijin Li, Zhaoyang Huang, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. PATS: Patch area transportation with subdivision for local feature matching. In *CVPR*, 2023. 1, 2
- [50] David Nistér. An efficient solution to the five-point relative pose problem. *TPAMI*, 2004. 2, 4, 6
- [51] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32, 2019. 5
- [52] Phil Pritchett and Andrew Zisserman. Wide baseline stereo matching. In *ICCV*, 1998. 2
- [53] Shengyi Qian, Linyi Jin, and David F Fouhey. Associative3D: Volumetric reconstruction from sparse views. In *ECCV*, 2020. 2, 5, 6
- [54] Rahul Raguram, Ondrej Chum, Marc Pollefeys, Jiri Matas, and Jan-Michael Frahm. USAC: A universal framework for random sample consensus. *IEEE transactions on pattern analysis and machine intelligence*, 2012. 4
- [55] René Ranftl and Vladlen Koltun. Deep fundamental matrix estimation. In *ECCV*, 2018. 2
- [56] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *TPAMI*, 2020. 5
- [57] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 7
- [58] Carolina Raposo, Miguel Lourenço, Michel Antunes, and Joao Pedro Barreto. Plane-based odometry using an RGB-D camera. In *BMVC*. 5
- [59] E. Riba, D. Mishkin, D. Ponsa, E. Rublee, and G. Bradski. Kornia: an open source differentiable computer vision library for PyTorch. In *WACV*, 2020. 5
- [60] Chris Rockwell, Justin Johnson, and David F Fouhey. The 8-point algorithm as an inductive bias for relative pose prediction by ViTs. In *3DV*, 2022. 1, 2, 3, 5, 6, 7, 8
- [61] Barbara Roessle and Matthias Nießner. End2End multi-view feature matching with differentiable pose optimization. In *ICCV*, 2023. 2
- [62] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *ICCV*, 2011. 2
- [63] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 1, 2, 3, 5, 6, 7
- [64] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied AI research. In *ICCV*, 2019. 5
- [65] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 1
- [66] Samarth Sinha, Jason Y Zhang, Andrea Tagliasacchi, Igor Gilitschenski, and David B Lindell. SparsePose: Sparse-view camera pose regression and refinement. In *CVPR*, 2023. 2
- [67] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. In *CVPR*, 2021. 1, 2, 3, 5, 6, 7, 8
- [68] Bin Tan, Nan Xue, Tianfu Wu, and Gui-Song Xia. NOPE-SAC: Neural one-plane RANSAC for sparse-view planar 3d reconstruction. *TPAMI*, 2023. 2, 5, 6
- [69] Zachary Teed and Jia Deng. DROID-SLAM: Deep visual SLAM for monocular, stereo, and RGB-D cameras. *NeurIPS*, 2021. 2
- [70] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *arXiv preprint arXiv:2208.04726*, 2022. 2
- [71] Philip HS Torr and Andrew Zisserman. MLESAC: A new robust estimator with application to estimating image geom-

- etry. *Computer vision and image understanding*, 78(1):138–156, 2000. 4
- [72] Jianyuan Wang, Christian Rupprecht, and David Novotny. PoseDiffusion: Solving pose estimation via diffusion-aided bundle adjustment. In *ICCV*, 2023. 1
- [73] Wenshan Wang, Yaoyu Hu, and Sebastian Scherer. TartanVO: A generalizable learning-based VO. In *CoRL*, 2021. 2
- [74] Tong Wei, Yash Patel, Alexander Shekhovtsov, Jiri Matas, and Daniel Barath. Generalized differentiable RANSAC. In *ICCV*, 2023. 2
- [75] Ross Wightman. PyTorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. 5
- [76] Zhenpei Yang, Jeffrey Z Pan, Linjie Luo, Xiaowei Zhou, Kristen Grauman, and Qixing Huang. Extreme relative pose estimation for RGB-D scans via scene completion. In *CVPR*, 2019. 2
- [77] Zhenpei Yang, Siming Yan, and Qixing Huang. Extreme relative pose network under hybrid representations. In *CVPR*, 2020. 2
- [78] Chao Yu, Zuxin Liu, Xin-Jun Liu, Fugui Xie, Yi Yang, Qi Wei, and Qiao Fei. DS-SLAM: A semantic visual SLAM towards dynamic environments. In *IROS*, 2018. 1
- [79] Jason Y Zhang, Deva Ramanan, and Shubham Tulsiani. Rel-Pose: Predicting probabilistic relative rotation for single objects in the wild. In *ECCV*, 2022. 2
- [80] Chen Zhao, Yixiao Ge, Feng Zhu, Rui Zhao, Hongsheng Li, and Mathieu Salzmann. Progressive correspondence pruning by consensus learning. In *ICCV*, 2021. 2
- [81] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. 3, 5, 8