This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **PRaDA:** Projective Radial Distortion Averaging

Daniil Sinitsyn<sup>1,2</sup>, Linus Härenstam-Nielsen<sup>1,2</sup>, Daniel Cremers<sup>1,2</sup> <sup>1</sup>Technical University of Munich, <sup>2</sup>Munich Center for Machine Learning

daniil.sinitsyn@tum.de, linus.nielsen@tum.de, cremers@tum.de

## Abstract

We tackle the problem of automatic calibration of radially distorted cameras in challenging conditions. Accurately determining distortion parameters typically requires either 1) solving the full Structure from Motion (SfM) problem involving camera poses, 3D points, and the distortion parameters, which is only possible if many images with sufficient overlap are provided, or 2) relying heavily on learning-based methods that are comparatively less accurate. In this work, we demonstrate that distortion calibration can be decoupled from 3D reconstruction, maintaining the accuracy of SfMbased methods while avoiding many of the associated complexities. This is achieved by working in Projective space, where the geometry is unique up to a homography, which encapsulates all camera parameters except for distortion. Our proposed method, Projective Radial Distortion Averaging, averages multiple distortion estimates in a fully projective framework without creating 3d points and full bundle adjustment. By relying on pairwise projective relations, our methods support any feature-matching approaches without constructing point tracks across multiple images.

# 1. Introduction

Having access to an accurate camera model is foundational to essentially all geometric computer vision algorithms, including Structure from Motion (SfM), Simultaneous Localization and Mapping (SLAM), and novel-view synthesis (NVS). However, in many cases, prior accurate intrinsics are not available. For instance, if images are sourced from the internet or if camera parameters have drifted over time due to wear and tear. In such cases, intrinsics must instead be estimated using geometric principles or learned from data. Despite active research, accurately determining intrinsics remains a challenging problem with no established go-to method. Methods based on full SfM and bundle adjustment [48, 49] often fail to converge without a good initialization, and learning-based methods [53] tend to lack accuracy and robustness.

In this work, we introduce Projective Radial Distortion



Figure 1. We average multiple imprecise and mutually inconsistent distortion model estimates into a single consistent model.

Averaging (PRaDA), a simple yet effective approach for calibrating radially distorted cameras from unordered image collections. SfM methods typically operate in Euclidean space, recovering 3D reconstructions and poses up to translation, rotation, and scaling. Projective framework extends this to projective space, where reconstructions are defined up to a homography that encapsulates all non-distortion parameters. This simplifies distortion optimization and enables image relationships to be expressed without 3D points or poses at any stage of the entire algorithm. We observe that standard image matching algorithms, which are typically designed for pinhole cameras, work well even for distorted images. For each image, we start off by computing point correspondences with every other image using wellestablished methods. For each such pair, we then use a robust correspondence-based minimal solver to obtain an estimate of the radial distortion. These individual estimates are then finally averaged to produce a single consistent model per camera. Fig. 1 An overview of the approach can be found in Fig. 3. Our contributions are summarized as follows:

- We introduce a fully projective method for radial distortion averaging. Our method does not require explicit 3D point reconstruction or focal length estimation, significantly simplifying the auto-calibration process.
- 2. We propose a novel distortion averaging technique that fuses inconsistent pairwise estimates into a single, consistent camera model.
- 3. We demonstrate significant improvements in accuracy and robustness over commonly used methods and validate them on challenging datasets.



Figure 2. Epipolar lines for the original distorted fisheye image (top), and the undistorted pinhole-like image (bottom). Points in the left and right images indicate matched points. In the undistorted image, epipolar lines are straight, and point correspondences satisfy Eq. (1).

Our method is agnostic to specific image matching techniques and does not rely on keypoint tracks across images, unlike traditional SfM pipelines.

Several methods exist in the literature for estimating radial distortion from a single pair of camera views (cf. Sec. 3). However, two views often lack sufficient information for a fully accurate distortion estimate across the entire image. The accuracy depends on factors like visual overlap and the quality of pixel correspondences.

# 2. Background

**Epipolar geometry with radial distortion:** The goal of camera distortion calibration is to map a *distorted* image (Fig. 2, top row) to an *undistorted* image (Fig. 2, bottom row), which is consistent with a pinhole camera. For pinhole cameras, the epipolar geometry between two views takes a particularly simple form. Namely, it can be characterized fully by a  $3 \times 3$  fundamental matrix F, which maps points from one image to their corresponding epipolar lines on the other image [17]. Given two pinhole camera images, any 2D point correspondence  $p_i^u$ ,  $q_i^u$  will satisfy the *epipolar constraint*:

$$(\bar{q}_i^u)^\mathsf{T} F \bar{p}_i^u = 0. \tag{1}$$

Here, we use a bar to denote the homogeneous representation of a point p, that is:  $\bar{p} = (p^{\mathsf{T}}; 1)^{\mathsf{T}}$ .

In the presence of distortion, epipolar lines are no longer straight, as shown in the top row of Fig. 2. This means point correspondences  $p_i$ ,  $q_i$  will generally not satisfy Eq. (1) for any single fundamental matrix. However, there is a mapping  $U_{\theta}(p)$ , with parameters  $\theta$ , which maps distorted points p to undistorted points:

$$p_i^u = U_\theta(p_i), \quad q_i^u = U_\theta(q_i), \tag{2}$$

which are consistent with pinhole geometry. In particular, we consider *radial distortion*, where the distortion only changes the distance from the center of the image:

$$U_{\theta}(p) = d_{\theta}(\|p\|)p.$$
(3)

**Radial distortion calibration:** The goal of radial distortion calibration is formulated as follows: determine the parameters  $\theta$  such that the undistorted points  $p_i^u$ ,  $q_i^u$  satisfy Eq. (1) for some fundamental fundamental matrix F. Once the distortion parameters have been estimated, the undistorted points can be used in any 3D reconstruction method designed for pinhole cameras.

**Division camera model:** Importantly, the epipolar constraint from Eq. (1) is invariant to multiplication by any non-zero scalar. We can use this fact to rescale the homogeneous points  $\bar{p}_i^u$  and move the dependence on radial distortion to the last homogeneous coordinate:

$$\bar{p}_i^u = \begin{pmatrix} d_\theta(\|p_i\|)p_i \\ 1 \end{pmatrix} \simeq \begin{pmatrix} p_i \\ 1/d_\theta(\|p_i\|) \end{pmatrix}$$
(4)

Where  $\simeq$  denotes equivalence up to a scaling factor. Of particular importance to our work is the one-parameter division camera model [12], given by:  $d_{\lambda}(r) = 1/(1 + \lambda r^2)$ . Which has the undistortion function:

$$U_{\lambda}(\bar{p_i}) \simeq \begin{pmatrix} p_i \\ 1 + \lambda \|p_i\|^2 \end{pmatrix}.$$
 (5)

While the one-parameter division model has wellestablished minimal solvers [12, 28, 30], it is not expressive enough to capture many real-world distortion patterns. Therefore, we also consider the extension to a k-th degree polynomial [31]:

$$U_{\theta}(\bar{p_i}) = \begin{pmatrix} p_i \\ h_{\theta}(\|p_i\|) \end{pmatrix}$$
(6)

where  $h(r) = \sum_{i=0}^{k} \theta_i r^i$ . Which corresponds to  $d(r) = 1/h_{\theta}(r)$ .

### 3. Related work

Accurate distortion estimation plays a crucial role in tasks such as SfM and 3D reconstruction. Various approaches have been proposed to address this challenge, each with different assumptions and requirements.

One category of approaches leverages a known 3D map to estimate camera distortions and focal lengths. These methods can be divided into two subcategories. The first assumes a perfect 3D map, which is the standard setup for camera calibration using patterns like chessboards [37, 47, 61] or learned patterns [58]. The second relies on a 3D map constructed by a SfM framework to register a new image



Figure 3. **Method overview:** Each possible image pair generates separate parameters of the distortion model via a robust solver (*cf*. Sec. 4.1). Then, for each camera  $C_i$ , the models are merged to form a single consistent estimation  $\theta_i$  using *distortion averaging* (*cf*. Sec. 4.4). Each stage is further refined using nonlinear optimization (*cf*. Secs. 4.2 and 4.5)

within it. For an available 3D map, a 3-parameter polynomial distortion model can be obtained from 5 2D-to-3D correspondences [29]. If more correspondences are available, calibration can be estimated using a rational model [31] or implicitly [40]. However, obtaining an accurate 3D map from images with unknown distortion can be a challenging problem.

In the case when a 3D map is not available, existing approaches can be broadly split into three categories: N-point-solvers, bundle-adjustment-based, and learning-based.

N-point solvers work by solving for the unknown fundamental matrix and distortion parameters subject to the constraints from Eq. (2) and Eq. (1). Pioneering work from Fitzgibbon [12] shows that a single-parameter polynomial model can be estimated minimally using 8 correspondences from an image pair from the same camera. The method was extended by Jiang et al. [24] to include focal length estimation. If the two images come from different cameras, a oneparameter (per camera) polynomial model can be estimated minimally with 9 correspondences [8, 28], by 3 variants of solvers requiring 10 correspondences [30], with 12 correspondences [8] or 15 correspondences [4]. SVA [35] goes beyond point correspondences and detects vanishing points by finding the intersection of circular arcs (corresponding to distorted lines). However, because of this reliance on straight lines, it may not work in cases where straight lines are not present. According to performance analysis in [30], the 10-point method F10 gives the best result in terms of performance compared to minimal methods.

**Bundle-adjustment methods** Distortion parameters can also be calculated as part of a SfM pipeline in the bundle adjustment step [14, 41, 48, 49], or NeRF optimization [23], although at the cost of increasing the complexity of the loss function. Global SfM methods such as GLOMAP [41] are promising, as they utilize least squares fitting for each pair of images and then average the results, potentially allowing distortion estimation in the first stage. Calibration-free SfM [32] can avoid the need to estimate radial distortion. This

approach uses clever minimal solvers that require three images in the dataset to have intersecting optical axes or tracks on 4 images [20], which may not always be the case, especially for sparse images. RpOSE [21] solves the projective reconstruction globally in a radial distortion-invariant way to later fit radial distortion. Having a big basin of convergence, this approach requires a set of image points tracked along several images, which can be challenging to obtain without pre-filtering stages that should be carefully handled in the presence of distortion.

Learning-based methods Calibration parameters can also be estimated using learning-based methods. Early works model distortion by learning to map point correspondences directly to 3d points [38] or by learning the residual error induced by the unknown calibration [11, 55]. Rong et al. [43] predicts a single-parameter division model. Trained on synthetically distorted images. DeepFocal [56] predicts focal length, trained on internet images. Hold-Geoffroy et al. [19] Predicts focal length and rotation relative to the horizon line. DeepCalib [7] predicts the focal length and  $\xi$  from the unified spherical model [5, 13, 37], trained on crops from panoramic images. DroidCalib [14] extents DROID-SLAM [52] to optimize unified camera model [37] with Gauss-Newton. GeoCalib [53] regresses a pixel-wise perspective field [25], which can be used to estimate radial distortion with nonlinear optimization. In DeepCalib [7], the authors observe that their solution performs better on images with significant distortion. They explain this by pointing to the non-uniformity of the training data, noting the absence of a dataset containing diverse fisheye distortions. Additionally, they argue that manually generating such data is unfeasible because applying artificial fisheye distortion to an image produces unrealistic results.

Similarly to Sarlin et al. [46], we argue that re-learning geometric principles with neural networks is often unnecessary. Instead, these principles can be accurately handled with traditional methods, using deep learning where exact modeling is infeasible. We find that well-established learning-based matchers [10, 34, 36, 45, 51] demonstrate strong robustness to distortion changes [28–31, 40, 41, 48, 49].

**Eucledian vs Projective SfM** In SfM, reconstruction can be performed in either a projective [26, 27] or Euclidean [41, 48] framework [17]. Projective SfM recovers the scene structure up to a homography without requiring information on focal length. The resulting reconstruction can still be optimal regarding reprojection errors and geometric relations but lacks metric accuracy. The point cloud can be arbitrarily skewed. Often, projective reconstruction is utilized as a first step in SfM methods [21, 22, 54] because it requires estimating fewer parameters. [27] is particularly interesting because it does not depend on tracking points across several images.

## 4. Method

We start by estimating 2D-2D correspondences  $(p_i, q_i)_{mn}$  for each pair of images  $\{I_m, I_n\}$ . These correspondences are the basis for estimating the distortion parameters of the camera model for each image. From each correspondence set  $(p_i, q_i)_{mn}$ , we then estimate a fundamental matrix and one-parameter distortion model (Sec. 4.1). This results in multiple initial estimates for each image's camera model. These distortion estimates are further refined by minimizing the Sampson error [15, 44] (Sec. 4.2).

The two-view distortion estimates are typically accurate only in the regions of the image where 2D points are available. So, to obtain a reliable camera model across the full image, we fuse the individual estimates through *distortion averaging* (Sec. 4.4). This step ensures a single consistent camera model per camera, which *optimally* merges the individual models.

As a final step to further improve the distortion estimate, we perform a global refinement (Sec. 4.5) of the Sampson error [15, 44] across all images. The full pipeline, illustrated in Fig. 3, produces a highly reliable distortion estimate without estimation of 3D points or camera poses.

#### 4.1. Two-view initialization

We obtain initial one-parameter distortion models using LO-RANSAC [9] with the F10 solver proposed by Kukelova et al. [30]. This minimal solver supports different camera models for both images. For each image pair  $\{I_m, I_n\}$ , we have a separate estimation of 2 distortion models. We normalize pixel coordinates by the length of the image diagonal to increase the numerical stability of the solver.

The benefits of this stage are two-fold. First off, we obtain initial estimates of the distortion models and fundamental matrices for all cameras from all matched image pairs, respectively. Secondly, RANSAC ensures that outlier correspondences generated by the matcher are filtered out.

#### 4.2. Two-view nonlinear refinement

Next, we refine the initial single-parameter distortion estimates using nonlinear refinement with a higher-order degree-k polynomial model from Eq. (6). Following Scaramuzza et al. [47], we parameterize polynomials such that  $\theta_0$  is fixed to 1 and  $\theta_1$  is fixed to 0. This parameterization ensures the model behaves like a pinhole near the image center. Each higher-order model is initialized by setting  $\theta_0 = 1, \theta_2 = \lambda$ , and other  $\theta_i$  to 0.

Minimal solvers for the higher-order polynomial division model are impractical due to the increasing number of minimal samples required. Therefore, we rely on nonlinear optimization for this stage. For this we use the Sampson error [15, 44]:

$$r_{\text{sampson}}^{2}(p,q,F,\theta_{1},\theta_{2}) = \frac{C^{2}(dp,dq)}{\left\|J_{dp}^{C}\right\|^{2} + \left\|J_{dq}^{C}\right\|^{2}},$$
 (7)

where

$$C(dp, dq) = U_{\theta_1}(q + dq)^{\mathsf{T}} F U_{\theta_2}(p + dp) \tag{8}$$

is the epipolar constraint centered at p, q. Geometrically, the Sampson error approximates the minimum adjustment in pixels required for each point correspondence  $p_i, q_i$  to satisfy the epipolar constraint with respect to F:

$$r_{\text{sampson}}^2 \approx \begin{cases} \min_{dp,dq} \|dp\|_2^2 + \|dq\|_2^2 \\ \text{s.t.} \quad U_{\theta_1}(q_i + dq)^{\mathsf{T}} F U_{\theta_2}(p_i + dp) = 0. \end{cases}$$
(9)

The approximation is obtained by linearizing the epipolar constraint around  $p_i, q_i$ .

Let  $(I_i, I_j)$  be an image pair with m correspondences. Let  $C_i$  and  $C_j$  be the corresponding cameras. The camera parameters estimated during the two-view step (Sec. 4.1) are denoted as  $\theta_{C_i}^{ij}$  and  $\theta_{C_j}^{ij}$ . The superscript indicates that the parameters are specific to this image pair. The optimization problem can then be formulated as:

$$\underset{F_{ij},\theta_{C_i}^{ij},\theta_{C_j}^{ij}}{\operatorname{argmin}} \sum_{l=1}^{m} r_{\operatorname{sampson}}^2 \left( p_l, q_l, F_{ij}, \theta_{C_i}^{ij}, \theta_{C_j}^{ij} \right)$$
(10)

Intuitively, minimizing the Sampson error refines the camera parameters and the fundamental matrices to ensure that the adjustments needed for the points to satisfy the epipolar constraint Eq. (1) are as small as possible. In addition, as Eq. (10) is defined only in terms of the epipolar constraint, it lets us bypass the explicit estimation of 3D points.

Nonlinear optimization of the fundamental matrix requires careful handling. While being  $3 \times 3$  real matrix it has 7 degrees-of-freedom [6]. One common way to parameterize it is  $SO(3) \times S^1 \times SO(3)$ . SO(3) is the group of  $3 \times 3$  rotation matrices in  $\mathbb{R}^3$  and has 3 degrees of freedom.  $S^1$  is the 1-dimensional sphere representing the unit circle and can be parameterized as an angle. This parameterization is unique and supports local updates via the corresponding Lie group exponential map [6]

The fundamental matrix is scale-invariant, i.e. it remains unchanged when multiplied by any non-zero scalar, and it has at most two non-zero singular values [17]. Given these properties, the minimal parameterization of the fundamental matrix can be obtained using the following scheme:

1. Compute SVD  $F = U\Sigma V^{\top}$ .

2. Normalize U and V to ensure they have determinant 1:

$$U \leftarrow U \cdot \det(U),$$
  

$$V \leftarrow V \cdot \det(V).$$
(11)

3. Normalize singular values  $\sigma_1$  and  $\sigma_2$ :

$$\sigma' = \frac{(\sigma_1, \sigma_2)}{\|\sigma_1, \sigma_2\|_2}.$$
 (12)

After these steps  $U \in SO(3), V \in SO(3), \sigma' \in S^1$ .

### **4.3.** Polynomial distortion regularization

Since the images are not fully covered by 2D-2D correspondences  $(p_i, q_i)_{mn}$ , the camera model is unconstrained and can behave arbitrarily in the uncovered regions. To address this, we introduce additional regularization to the distortion polynomials.

Following the approach of Hartley and Kang [16], Pan et al. [40] we apply the *local linearity assumption* to the undistortion function, treating it as a function of the radius. This aims to constrain the rate of change of the undistortion function, ensuring it remains monotonic. In the continuous case, this regularization can be formulated as:

$$\min \int_{0}^{R} \left\| \frac{dU_{\theta}(r)}{dr} \right\|^{2} dr \tag{13}$$

This function is not analytically integrable due to the potentially high degree polynomial in the denominator. Instead, we use an efficient numerical approximation of this integral.

### 4.4. Distortion averaging

After the initial pairwise estimation of camera parameters, we collect all the observed parameters for each physical camera model. The goal of this step is to determine the model that best fits the set of estimated camera parameters. For each camera C, we have multiple, potentially mutually inconsistent estimates of the distortion parameters  $\theta$ . Each estimated distortion model is generally consistent only within the regions covered by point correspondences and may be inconsistent in other image regions. To combine these individual estimates into a single, unified distortion polynomial, we propose **distortion averaging**. We solve an optimization problem in functional space so that the solution behaves similarly to each of the estimated camera models within the image.

As a general principle, we can average multiple distortion models  $\hat{U}_i$ , i = 1, ..., n by solving a weighted leastsquares problem within the space of functions on the image plane:

$$\bar{U} = \underset{U}{\operatorname{argmin}} \sum_{i=1}^{n} \omega_i \left\langle U, \ \hat{U}_i \right\rangle^2$$

$$= \underset{U}{\operatorname{argmin}} \sum_{i=1}^{n} \omega_i \int_{p \in I} \left\| U(p) - \hat{U}_i(p) \right\|^2 \mathrm{d}x \,\mathrm{d}y$$
(14)

where p = (x, y) ranges over all pixels of the image, and  $\omega_i$  are weights that sum to 1. Reparametrizing in terms of the radial distortion model U(p) = pd(||p||) and expressing this integral in polar coordinates, we get:

$$\bar{d} = \underset{d}{\operatorname{argmin}} \sum_{i=1}^{n} \omega_i \int_{0}^{R} \left\| d(r) - \hat{d}_i(r) \right\|^2 r^3 \, \mathrm{d}r.$$
(15)

Finally, we constrain the solution to be in the form of an undistortion function of the polynomial division model (See Eq. (6)). Then Eq. (15) changes to:

$$\bar{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^{n} \omega_i \int_{0}^{R} \left\| \frac{1}{h_{\theta}(r)} - \frac{1}{h_{\theta_i}(r)} \right\|^2 r^3 \,\mathrm{d}r \qquad (16)$$

where  $\theta$  corresponds to the coefficients of the polynomial. We solve this numerically, initializing  $\theta$  as a weighted average of  $\theta_i$ :

$$\bar{\theta} = \frac{\sum_{i=1}^{n} \omega_i \theta_i}{\sum_{i=1}^{n} \omega_i}.$$
(17)

Note that our averaging approach can generate polynomials of nearly any degree, making the degree a hyperparameter of the averaging scheme. However, the degree of  $\theta$  from Eq. (17) is limited by the maximum degree of  $\theta$ .

While Eq. (17) is not in general optimal for Eq. (16), it turns out to be an optimal solution of Eq. (15) for the multiplicative distortion model [31] (see Supplementary material).

#### 4.5. Global refinement

Even after the distortion averaging step, the resulting camera models may still lack sufficient accuracy. We refine them together with the fundamental matrix in a global optimization. We minimize a robust Sampson error loss across all images:

$$\underset{\{F_{ij}\},\{\theta_k\}_{l,i,j}}{\operatorname{argmin}} \sum_{\{P_l,i,j\}} \rho\left(r_{\operatorname{sampson}}\left(p_l,q_l,F_{ij},\theta_{C_i},\theta_{C_j}\right)\right)$$
(18)



a) KITTI-360 [33] b) WoodScape [60] Figure 4. Example of images from challenging sequences.

where  $\rho$  is a Cauchy loss function [39] and  $p_l, q_l$  corresponding points. If an image  $I_i$  appears in multiple pairs, such as  $(I_i, I_j)$  and  $(I_i, I_k)$ , the same camera model is used for all pairs, and the optimization is performed jointly over these pairs. Solving Eq. (18) ensures that the camera models remain consistent with the relative geometry across multiple image pairs. In addition, this loss formulation allows the optimization of principal points if there is sufficient coverage across all image pairs.

# 5. Experiments

## 5.1. Metrics

To measure the accuracy of the distortion estimate, we use the Reprojection Error (RE), computed as:

$$\mathbf{RE} = \frac{1}{|\Omega|} \sum_{p \in \Omega} \left\| \pi_{\theta} \left( \pi^{-1} \left( p \right) \right) - p \right\|.$$
(19)

Where  $p \in \Omega$  are the image pixel coordinates,  $\pi$  is the projection function of the ground truth camera model, and  $\pi_{\theta}$  is the projection function of the estimated model with parameters  $\theta$ . MRE is widely used to compare camera models [7, 30, 35, 53]. It provides an interpretable measure of accuracy in units of pixels in the original image.

For a fair comparison of the distortion estimation only, we compute Focal-Adjusted RE (FA-RE), which is obtained by selecting the focal length that minimizes Eq. (19). The error thus reflects the model's ability to represent distortion uniformly across the image. This formulation helps to better capture the true error in highly distorted regions.

We also evaluate the impact of our method on the downstream task of SfM. We measure the deviation in translation direction and rotation angle between the estimated and ground truth relative poses for each possible image pair, as in GLOMAP [41]. This metric is more robust to completely incorrect estimation of a subset of poses than an absolute pose error.

#### 5.2. Datasets

**ETH3D** [50]: A multi-view stereo benchmark dataset comprising video sequences captured by camera rigs and



Figure 5. **ScanNet++ by sequence:** Our method consistently outperforms SfM-based methods Colmap [48, 49] and Glomap [41] in terms of calibration accuracy.

ground-truth calibrations in COLMAP format. We use all frames from all 5 sequences from the "Low-res many-view" category, which includes cameras with significant distortion.

**KITTI 360 [33]:** An extension of the KITTI dataset that includes 360-degree panoramic images created from two 180° fisheye cameras. For evaluation, we use the 600 first frames for each camera from each sequence.

**ScanNet++** [59]: This dataset provides fisheye-lens DSLR images across various indoor scenes. Each scene includes continuous sequences of 200+ images. To evaluate performance, we use all frames from the test set sequences provided by the dataset.

**WoodScape [60]:** WoodScape is a fisheye camera dataset comprising over 10,000 images captured by four 180° fisheye cameras mounted on a vehicle. The images are taken at sparse timestamps. This makes the application of feature-based methods particularly challenging. For DeepCalib and Geocalib, we use all available frames.

KITTI-360 and WoodScape are particularly challenging as the images are heavily distorted; see Fig. 4. We estimate one camera model per distinct physical camera presented in datasets. That is: 2 cameras for ETH3D, 2 for KITTI, one for ScanNet++, and 4 for WoodScape.

#### 5.3. Results

**Baselines:** We compare our method to several common and state-of-the-art autocalibration approaches: **COLMAP** [48, 49], which incrementally estimates the camera model using a dense subset of images; **GLOMAP** [41], which uses global optimization in SfM similar to our approach; **DroidCalib** [14], a SLAM-based method that optimizes camera parameters from video input; **GeoCalib** [53], which jointly refines the camera model and gravity direction from a single image; and **Deep-Calib** [7], which directly regresses camera intrinsic and distortion parameters using Mei's camera model [37]. Additional implementation details are given in supplementary Appendix A.

**Quantitative Results:** The results of the reprojection errors are summarized in Tab. 1. To ensure fairness in the comparison, we run our solution on the same SIFT matches as

	Ś	ScanNet+	+	ET	H3D (car	n 4)	ET	TH3D (car	n 5)	KIT	ГІ-360 (ca	am 2)	KIT	ГІ-360 (ca	am 3)
Method	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
Colmap [48]	0.7	2.0	6.4	13.2	26.0	44.5	18.6	25.1	37.1	117.1	125.5	136.4	97.6	112.4	127.4
Glomap [41]	0.6	1.8	3.4	10.0	18.4	29.3	8.6	19.6	58.0	85.9	122.0	138.3	91.9	113.3	128.7
DroidCalib [14]	0.4	1.2	2.6	19.9	36.3	58.4	18.0	46.4	71.1	98.7	102.2	105.2	120.9	128.1	135.3
GeoCalib [53]	4.19	4.6	4.9	30.1	35.8	43.0	26.6	34.6	123.1	124.4	125.5	63.8	122.0	123.1	124.3
DeepCalib [7]	0.8	10.8	30.7	18.1	20.9	23.6	13.7	18.2	23.5	141.0	160.5	174.0	109.4	153.2	177.1
Ours (w. SIFT)	0.1	0.6	1.7	2.6	5.3	10.4	3.2	14.4	28.8	42.3	44.8	46.2	47.5	50.2	55.1

Table 1. **Performance comparison on challenging datasets:** Minimum, mean, and maximum focal-adjusted reprojection errors (in pixels) are reported for selected baseline methods. The reported errors represent the values for all sequences within each dataset. Our method uses the same SIFT matches as COLMAP and GLOMAP to ensure a fair comparison. For datasets with multiple camera models, errors are reported for each model independently, with the corresponding model in parentheses.



Figure 6. **Qualitative results:** Comparison of estimated camera models on challenging datasets. The models are adjusted so that the center of an image approximately matches the center of the image undistorted by the ground truth model.

Method	FV	RV	MVR	MVL
Ours	51.2	22.2	7.7	9.4
DeepCalib	9.7	14.5	14.5	21.5
GeoCalib	98.0	96.2	93.0	93.9

Table 2. **WoodScape dataset:** The mean focal-adjusted reprojection error on the WoodScape dataset is reported for each camera as follows: **FV** points forward, **RV** points backward, **MVR** represents the right-side camera, and **MVL** represents the left-side camera.

COLMAP and GLOMAP. If either COLMAP or GLOMAP splits the 3D scene into multiple models, we choose the one with more registered images. We run DroidCalib several

times, varying parameters such as image step and number of frames. Geocalib and DeepCalib are run on all images. For DroidCalib, Deepcalib, and Geocalib, we report *the mean error* over all models for each sequence.

Our proposed solution consistently outperforms all baseline methods across all datasets, achieving lower minimum reprojection errors. Despite these improvements, due to the radially symmetric distortion assumption, the proposed method cannot achieve a subpixel reprojection error on datasets with ground-truth models that are not radially symmetric, such as ETH3D [50] or KITTI-360[33].

Fig. 5 shows the histogram of reprojection errors per sequence on the test split of ScanNet++. Our method demonstrates a significantly higher concentration of errors below 1 pixel, outperforming both COLMAP and GLOMAP in this critical accuracy range.

The results on WoodScape are summarized in Tab. 2. We compare our solution to DeepCalib and GeoCalib because other methods do not produce meaningful results. Our solution was applied to the first 60 frames of each camera. Increasing the number of frames had no significant effect on the results. We used LOFTR [51] to build point correspondences. DeepCalib and GeoCalib used all available images for each camera. We report the mean error of all estimated models. As expected, we get worse results for cameras pointing backward and forward. This means that the center of the distortion coincides with or close to the projection of the second camera onto the image, which corresponds to the degenerative case for distortion estimation from correspondences [18, 57]. Intuitively, this is the case when all epipolar lines are straight, regardless of camera distortion.

We compare the accuracy of GLOMAP with and without the distortion model initialization from PRaDA in Tab. 3. We evaluate angular errors for relative poses on the Sparse test images from the test split of ScanNet++ [59], which is a challenging dataset containing extreme baselines with low overlap. The results demonstrate the benefits of a decoupled calibration stage. In contrast, joint estimation of camera calibration and 3D geometry often leads to poor reconstruction performance in non-ideal settings.

**Qualitative results:** Fig. 6 presents qualitative results, showing the visual undistortion of an image from multiple datasets. If the method does not work on the dataset, we show the original image, such as for COLMAP and GLOMAP on WoodScape. Since we are not estimating the focal length, we match it to the one of the ground truth model. To do it, we rely on the observation that a distorted image can be approximated by a pinhole camera near its center. Specifically, for a given focal length f and a small pixel displacement dx, the angular resolution near the image center, denoted by  $\theta$ , can be described as:

$$\frac{dx}{f} \approx \tan(\theta)$$
 (20)

For small angles  $tan(\theta) \approx \theta$ . Then for f:

$$f \approx \frac{dx}{\theta}$$
 (21)

The angle  $\theta$  is defined as the angle between the optical ray corresponding to the principal point and the ray corresponding to the back projection of the displacement dx.

This ensures that all models behave similarly near the center of the image, making it possible to undistort the images defining the same FOV.

	PRaDA + GLOMAP	GLOMAP
$r_{err}$ (deg)	0.18/4.51/44.56	0.25/28.99/118.6
$t_{err}$ (deg)	0.26/8.70/81.07	0.33/27.39/95.76

Table 3. **Sparse ScanNet with different initialization:** Min/Mean/Max angular errors of relative poses for GLOMAP and GLOMAP initialized with PRaDA. These results highlight the benefits of PRaDA for 3D reconstruction.

# 6. Discussion

Similar to COLMAP [48, 49] and GLOMAP [41], our method relies on the performance of the matcher. If the matcher has been trained or optimized primarily for pinhole settings, its errors can noticeably impact our results. This is because we use the Levenberg-Marquardt algorithm as a nonlinear optimizer, which assumes normally distributed errors. This condition may not hold in distorted regions, resulting in reduced accuracy. This can be properly modeled by estimating the error distribution.

Like COLMAP and GLOMAP, which use predefined thresholds for RANSAC and inlier estimation, we also rely on them. We would like to point out that it is possible to overcome the thresholds completely. Using the  $\sigma$  or  $\sigma$ ++ consensus introduced by Barath et al. [2, 3] through all steps of the proposed algorithm makes this possible. However, its application requires thoughtful design. We see this as a promising direction for future research.

#### 7. Conclusion

This paper proposes a new method for estimating radial distortion in a projective setting. Our approach integrates distortion modeling to each stage of the global projective reconstruction pipeline, with model averaging based on pixel errors in the image plane relative to less expressive models. This careful design allows for more accurate distortion correction. We demonstrate that our method outperforms modern techniques, achieving the lowest error on almost all datasets. We show that it can estimate extremely distorted images with a  $180^{\circ}$  lens.

KITTI-360 [33] and WoodScape [60] serve as strong examples of the robustness of the proposed solution, as they use cameras with an exact 180° field of view. This demonstrates the ability of the proposed solution to effectively handle extreme wide-angle fisheye cameras.

### Acknowledgments

This work was supported by the ERC Advanced Grant "SIMULACRON" (agreement #884679), the GNI Project "AI4Twinning", and the DFG project CR 250/26-1 "4D YouTube".

### References

- [1] Sameer Agarwal, Keir Mierle, and The Ceres Solver Team. Ceres Solver, 2023. 1
- [2] Daniel Barath, Jiri Matas, and Jana Noskova. Magsac: marginalizing sample consensus. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10197–10205, 2019. 8
- [3] Daniel Barath, Jana Noskova, Maksym Ivashechkin, and Jiri Matas. Magsac++, a fast, reliable and accurate robust estimator. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 1304–1312, 2020. 8
- [4] J.P. Barreto and K. Daniilidis. Fundamental matrix for cameras with radial distortion. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, pages 625–632 Vol. 1, 2005. 3
- [5] João P. Barreto. A unifying geometric representation for central projection systems. *Computer Vision and Image Understanding*, 103(3):208–217, 2006. Special issue on Omnidirectional Vision and Camera Networks. 3
- [6] Adrien Bartoli and Peter Sturm. Nonlinear estimation of the fundamental matrix with minimal parameters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(3): 426–432, 2004. 4, 5
- [7] Oleksandr Bogdan, Viktor Eckstein, Francois Rameau, and Jean-Charles Bazin. Deepcalib: A deep learning approach for automatic intrinsic calibration of wide field-of-view cameras. In *Proceedings of the 15th ACM SIGGRAPH European Conference on Visual Media Production*, pages 1–10, 2018. 3, 6, 7
- [8] Martin Byrod, Zuzana Kukelova, Klas Josephson, Tomas Pajdla, and Kalle Astrom. Fast and robust numerical solutions to minimal problems for cameras with radial distortion. In 2008 IEEE Conference on Computer Vision and Pattern Recognition, pages 1–8, 2008. 3
- [9] Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized ransac. In Pattern Recognition: 25th DAGM Symposium, Magdeburg, Germany, September 10-12, 2003. Proceedings 25, pages 236–243. Springer, 2003. 4, 1
- [10] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 4
- [11] Yongtae Do. Application of neural networks for stereocamera calibration. In IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No. 99CH36339), pages 2719–2722. IEEE, 1999. 3
- [12] Andrew W Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Proceedings* of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001, pages I–I. IEEE, 2001. 2, 3
- [13] Christopher Geyer and Kostas Daniilidis. A unifying theory for central panoramic systems and practical implications. In *Computer Vision — ECCV 2000*, pages 445–461, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg. 3

- [14] Annika Hagemann, Moritz Knorr, and Christoph Stiller. Deep geometry-aware camera self-calibration from video. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3438–3448, 2023. 3, 6, 7
- [15] Matthew Harker and Paul O'Leary. First order geometric distance (the myth of sampsonus). In *BMVC*, pages 87–96, 2006. 4, 1
- [16] Richard Hartley and Sing Bing Kang. Parameter-free radial distortion correction with center of distortion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelli*gence, 29(8):1309–1321, 2007. 5
- [17] Richard Hartley and Andrew Zisserman. Multiple View Geometry in Computer Vision. Cambridge University Press, New York, NY, USA, 2 edition, 2003. 2, 4, 5
- [18] Jose Henrique Brito, Roland Angst, Kevin Koser, and Marc Pollefeys. Radial distortion self-calibration. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1368–1375, 2013. 8
- [19] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-François Lalonde. A perceptual measure for deep single image camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2354–2363, 2018. 3
- [20] Petr Hruby, Viktor Korotynskiy, Timothy Duff, Luke Oeding, Marc Pollefeys, Tomas Pajdla, and Viktor Larsson. Four-view geometry with unknown radial distortion. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 8990–9000, 2023. 3
- [21] José Pedro Iglesias and Carl Olsson. Radial distortion invariant factorization for structure from motion. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 5906–5915, 2021. 3, 4
- [22] José Pedro Iglesias, Amanda Nilsson, and Carl Olsson. expose: Accurate initialization-free projective factorization using exponential regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8959–8968, 2023. 4
- [23] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5846–5854, 2021. 3
- [24] Fangyuan Jiang, Yubin Kuang, Jan Erik Solem, and Kalle Åström. A minimal solution to relative pose with unknown focal length and radial distortion. In *Computer Vision–ACCV* 2014: 12th Asian Conference on Computer Vision, Singapore, Singapore, November 1-5, 2014, Revised Selected Papers, Part II 12, pages 443–456. Springer, 2015. 3
- [25] Linyi Jin, Jianming Zhang, Yannick Hold-Geoffroy, Oliver Wang, Kevin Blackburn-Matzen, Matthew Sticha, and David F Fouhey. Perspective fields for single image camera calibration. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 17307– 17316, 2023. 3
- [26] Yoni Kasten, Amnon Geifman, Meirav Galun, and Ronen Basri. Algebraic characterization of essential matrices and

their averaging in multiview settings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5895–5903, 2019. 4

- [27] Yoni Kasten, Amnon Geifman, Meirav Galun, and Ronen Basri. Gpsfm: Global projective sfm using algebraic constraints on multi-view fundamental matrices. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3264–3272, 2019. 4
- [28] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Automatic generator of minimal problem solvers. In *Computer Vision – ECCV 2008*, pages 302–315, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. 2, 3, 4
- [29] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Realtime solution to the absolute pose problem with unknown radial distortion and focal length. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2816– 2823, 2013. 3
- [30] Zuzana Kukelova, Jan Heller, Martin Bujnak, Andrew Fitzgibbon, and Tomas Pajdla. Efficient solution to the epipolar geometry for radially distorted cameras. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2, 3, 4, 6, 1
- [31] Viktor Larsson, Torsten Sattler, Zuzana Kukelova, and Marc Pollefeys. Revisiting radial distortion absolute pose. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 1062–1071, 2019. 2, 3, 4, 5
- [32] Viktor Larsson, Nicolas Zobernig, Kasim Taskin, and Marc Pollefeys. Calibration-free structure-from-motion with calibrated radial trifocal tensors. In *Computer Vision–ECCV* 2020: 16th European Conference, Glasgow, UK, August 23– 28, 2020, Proceedings, Part V 16, pages 382–399. Springer, 2020. 3
- [33] Yiyi Liao, Jun Xie, and Andreas Geiger. KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *Pattern Analysis and Machine Intelligence (PAMI)*, 2022. 6, 7, 8
- [34] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. Lightglue: Local feature matching at light speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17627–17638, 2023. 4
- [35] Yaroslava Lochman, Oles Dobosevych, Rostyslav Hryniv, and James Pritts. Minimal solvers for single-view lensdistorted camera auto-calibration. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 2887–2896, 2021. 3, 6
- [36] David G Lowe. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 4
- [37] Christopher Mei and Patrick Rives. Single view point omnidirectional camera calibration from planar grids. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3945–3950. IEEE, 2007. 2, 3, 6
- [38] Márcio Mendonça, Ivan N Da Silva, and José EC Castanho. Camera calibration using neural networks. 2002. 3
- [39] Thamsanqa Mlotshwa, Heinrich van Deventer, and Anna Sergeevna Bosman. Cauchy loss function: Robustness under gaussian and cauchy noise. In Southern

African Conference for Artificial Intelligence Research, pages 123–138. Springer, 2022. 6

- [40] Linfei Pan, Marc Pollefeys, and Viktor Larsson. Camera pose estimation using implicit distortion models. In *Proceed*ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12819–12828, 2022. 3, 4, 5
- [41] Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L Schönberger. Global structure-from-motion revisited. In *European Conference on Computer Vision (ECCV)*, 2024. 3, 4, 6, 7, 8, 1
- [42] Chuck Pheatt. Intel® threading building blocks. 23(4):298, 2008. 1
- [43] Jiangpeng Rong, Shiyao Huang, Zeyu Shang, and Xianghua Ying. Radial lens distortion correction using convolutional neural networks trained with synthesized images. In Computer Vision–ACCV 2016: 13th Asian Conference on Computer Vision, Taipei, Taiwan, November 20-24, 2016, Revised Selected Papers, Part III 13, pages 35–49. Springer, 2017. 3
- [44] Felix Rydell, Angélica Torres, and Viktor Larsson. Revisiting sampson approximations for geometric estimation problems. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4990–4998, 2024. 4, 1
- [45] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4938–4947, 2020. 4
- [46] Paul-Edouard Sarlin, Ajaykumar Unagar, Mans Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, et al. Back to the feature: Learning robust camera localization from pixels to pose. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3247–3257, 2021. 3
- [47] Davide Scaramuzza, Agostino Martinelli, and Roland Siegwart. A toolbox for easily calibrating omnidirectional cameras. In 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5695–5701. IEEE, 2006. 2, 4
- [48] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 3, 4, 6, 7, 8
- [49] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 3, 4, 6, 8
- [50] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with highresolution images and multi-camera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017. 6, 7
- [51] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching

with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 4, 8

- [52] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. Advances in neural information processing systems, 34:16558–16569, 2021.
   3
- [53] Alexander Veicht, Paul-Edouard Sarlin, Philipp Lindenberger, and Marc Pollefeys. GeoCalib: Single-image Calibration with Geometric Optimization. In *ECCV*, 2024. 1, 3, 6, 7
- [54] Simon Weber, Je Hyeong Hong, and Daniel Cremers. Power variable projection for initialization-free large-scale bundle adjustment. arXiv preprint arXiv:2405.05079, 2024. 4
- [55] J Wen and G Schweitzer. Hybrid calibration of ccd cameras using artificial neural nets. In [Proceedings] 1991 IEEE International Joint Conference on Neural Networks, pages 337–342. IEEE, 1991. 3
- [56] Scott Workman, Connor Greenwell, Menghua Zhai, Ryan Baltenberger, and Nathan Jacobs. Deepfocal: A method for direct focal length estimation. In 2015 IEEE International Conference on Image Processing (ICIP), pages 1369–1373. IEEE, 2015. 3
- [57] Changchang Wu. Critical configurations for radial distortion self-calibration. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 25–32, 2014. 8
- [58] Wenqi Xian, Aljaž Božič, Noah Snavely, and Christoph Lassner. Neural lens modeling. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8435–8445, 2023. 2
- [59] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023. 6, 8
- [60] Senthil Yogamani, Ciaran Hughes, Jonathan Horgan, Ganesh Sistu, Padraig Varley, Derek O'Dea, Michal Uricar, Stefan Milz, Martin Simon, Karl Amende, Christian Witt, Hazem Rashed, Sumanth Chennupati, Sanjaya Nayak, Saquib Mansoor, Xavier Perrotton, and Patrick Perez. Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 6, 8
- [61] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000. 2