expOSE: Accurate Initialization-Free Projective Factorization using Exponential Regularization

José Pedro Iglesias[1], Amanda Nilsson[2], Carl Olsson[1,2]

[1]Chalmers University of Technology, Sweden
[2]Lund University, Sweden

Abstract

Bundle adjustment is a key component in practically all available Structure from Motion systems. While it is crucial for achieving accurate reconstruction, convergence to the right solution hinges on good initialization. The recently introduced factorization-based pOSE methods formulate a surrogate for the bundle adjustment error without reliance on good initialization. In this paper, we show that pOSE has an undesirable penalization of large depths. To address this we propose expOSE which has an exponential regularization that is negligible for positive depths. To achieve efficient inference we use a quadratic approximation that allows an iterative solution with VarPro. Furthermore, we extend the method with radial distortion robustness by decomposing the Object Space Error into radial and tangential components. Experimental results confirm that the proposed method is robust to initialization and improves reconstruction quality compared to state-of-the-art methods even without bundle adjustment refinement.[1]

1. Introduction

Factorization is a long-established method in Structure from Motion (SfM). It originates from [38] by Tomasi and Kanade showing how, under the orthographic camera model, structure and motion can be computed simultaneously from an image sequence using singular value decomposition (SVD). The method was later reformulated for affine cameras, including weak perspective projection [32]. Strum and Triggs [36] further extended factorization to projective cameras by accounting for projective depths.

One appeal of these factorization algorithms is they can yield a closed-form solution by using the SVD. It is however only possible to use the SVD if every considered scene point is visible throughout the whole image sequence. In cases of missing data, the SVD can be replaced with iterative methods. Simple splitting methods [4, 8, 22] are able to regularize singular values when computing a proximal operator, but can give rather erroneous solutions because of a low convergence rate close to the optimum. [5, 8] give an idea of convex formulation using the nuclear norm, but are usually too weak for SfM in the presence of noise [19, 30]. The papers [1, 9, 10, 31] suggest different ways to assure that direct bilinear optimization only has a global minimum. However, SfM problems with local minima do not fulfill their required conditions [3].

It was recently shown by Hong et al. [14–17] that direct bilinear estimation of structure and motion can be made robust to local minima in combination with the Variable Projection (VarPro) method. In [15] the objective is exchanged for the Pseudo Object Space Error (pOSE) which is a tradeoff between the object space error and a quadratic regularization term. This was later extended to a radial distortion invariant version RpOSE, presented in [18]. With their bilinear factorization structure and a large basin of convergence when using VarPro, these pOSE models tend to find

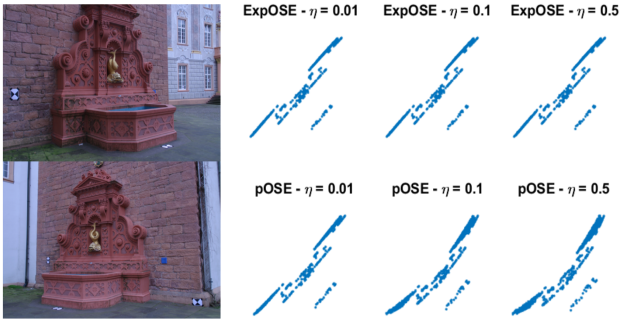

Figure 1. (Left) Examples of two of the images in the Fountain sequence. (Right) Reconstruction obtained with expOSE (top) and pOSE (bottom) for 3 different values of η. Our method achieves the same convergence rate as pOSE while having a higher reconstruction quality and being less dependent on the choice of η.

[1]This work has been funded by the Swedish Research Council (grant no. 2018-05375), the Swedish Foundation for Strategic Research project, Semantic Mapping and Visual Navigation for Smart Robots (grant no. RIT15-0038), and the Wallenberg AI, Autonomous Systems and Software Program (WASP).

a global minimum independently of the initialization. Additionally, both pOSE and RpOSE have in [18] been shown to be local approximations of the reprojection error, enabling iterative refinement to the maximum likelihood solution.

In this paper, we show that the regularization term in the pOSE formulation overly penalizes large positive depths and can thereby limit the range of feasible depths too much to achieve satisfactory solutions. We instead propose regularization with an exponential penalty that is negligible for positive depths. To achieve efficient inference we use a quadratic approximation of the exponential term suitable for optimization with VarPro. Moreover, we extend the method with radial distortion robustness by decomposing the OSE into radial and tangent components.

In short, the main contributions of this paper are:

- We investigate the pOSE models' undesirable penalization of large depths and propose expOSE which has negligible regularization of positive depths;
- We formulate a quadratic approximation of the exponential regularization term in expOSE to make it suitable for optimization with VarPro and show that, with random initialization, the model achieves convergence rates similar to pOSE with significantly higher reconstruction quality;
- We extend expOSE with radial distortion robustness by decomposing the Object Space Error (OSE) into radial and tangent components and propose an SfM pipeline that is able to obtain a complete and accurate Euclidean reconstruction from uncalibrated cameras.

2. Reconstruction Objectives

In this section, we illustrate the problems with direct optimization of reprojection error and discuss how this is addressed using the pOSE model [15]. We then present our exponential regularization and show how this addresses the limitations of the pOSE model.

2.1. Reprojection Error and Cheirality

Bundle adjustment [12, 39] is the standard routine when it comes to solving the Structure-from-Motion problem. Given measured point projections $\boldsymbol{m}_{ij}$ the goal is to attempt to minimize

$$\sum_{ij} \left\| \boldsymbol{m}_{ij} - \frac{\boldsymbol{x}_{ij}}{z_{ij}} \right\|^2, \tag{1}$$

where $\begin{bmatrix} \boldsymbol{x}_{ij} \\ z_{ij} \end{bmatrix} = P_i U_j$. Here $\boldsymbol{x}_{ij}$ is a 2 vector, z_{ij} is a number, referred to as the projective depth, P_i is a 3×4 camera matrix and U_i is a 4×1 vector containing homogeneous coordinates of the projected 3D point. Under the assumption of Gaussian image noise, this gives the maximal likelihood estimate of the camera matrices and 3D points [12].

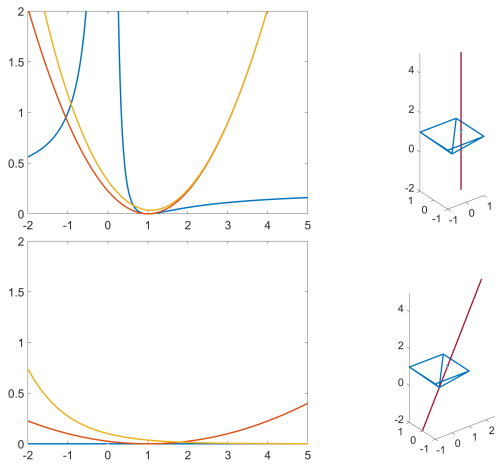

Figure 2. Left: Objective values of the reprojection error (blue), the pOSE error ($\eta = 0.1$, red) and our proposed formulation ($\eta = 0.1$, yellow) on the lines $(1-t)(0.5, 0, -1)+t(0.5, 0, 1)$ (top) and $(1-t)(-0.5, 0, -1)+t(0.5, 0, 1)$ (bottom) when $\boldsymbol{m} = (0.5, 0)$. Note that the reprojection error is undefined at $z = 0$ since this corresponds to the camera center. Right: Corresponding camera and sampling line.

It is well known that optimizing (1) is difficult and requires good initialization to achieve convergence to the right solution. One of the difficulties is the division of $\boldsymbol{x}_{ij}$ by z_{ij}. This creates a barrier of objective values that goes to infinity and needs to be traversed when for example moving from $(\boldsymbol{x}_{ij}, -z_{ij})$ to $(\boldsymbol{x}_{ij}, z_{ij})$. The blue curve of Figure 2 (top) shows a 2D example of this barrier. Here we used $\boldsymbol{m} = (0.5, 0)$ and sampled the function $\left(\boldsymbol{m} - \frac{\boldsymbol{x}}{z}\right)^2$ on the line segment $(\boldsymbol{x}, z) = (1-t)(0.5, 0, -1) + t(0.5, 0, 1)$. The best value over this line is at $t = 1$ which gives $(\boldsymbol{x}, z) = (0.5, 0, 1)$. For comparison, we also plot the corresponding values of the pOSE model [15] (red) and the proposed formulation that we will describe below (yellow). In a calibrated setting the interpretation of z_{ij} is the depth [12] of the observed 3D point. Hence, in practical cases, where observed points are in front of the camera, there is usually no reason to allow solutions with negative z_{ij}. In the uncalibrated case z_{ij} is referred to as a projective depth. It can be shown that when the data is noise free (with sufficiently many visible projections) there is always a solution where the projective depths are all positive [26] if the observed points are in front of the camera. Moreover, any other solution is projectively equivalent to this one, meaning that there is a projective 3D transformation that makes the projective depths positive [25, 26].

2.2. The pOSE Model

In view of the above, constraining the problem to positive depths is no practical restriction. Still finding a good starting solution where all depths are positive is not a trivial

issue. In [15] the objective (1) is exchanged for the object space error (OSE)

$$\ell_{\text{OSE}}(\boldsymbol{x}_{ij}, z_{ij}) = \|z_{ij}\boldsymbol{m}_{ij} - \boldsymbol{x}_{ij}\|^2. \tag{2}$$

Here, the scale-invariant residual of (1) has been replaced with a linear error allowing points to switch from negative to positive projective depths. It can be shown [18] that the OSE residual $z_i\boldsymbol{m}_{ij} - \boldsymbol{x}_{ij}$ is the first order Taylor expansion of the projective residual $\boldsymbol{m}_{ij} - \frac{\boldsymbol{x}_{ij}}{z_{ij}}$ around $(\boldsymbol{x}_{ij}, z_{ij}) = (\boldsymbol{m}_{ij}, 1)$, and it is therefore in some sense the closest linear approximation that we can find. On the downside, the OSE is clearly minimized by the trivial solution $\boldsymbol{x}_{ij} = 0$, $z_{ij} = 0$ for all i, j. Therefore [15] adds the quadratic regularization

$$\ell_{\text{aff}}(\boldsymbol{x}_{ij}) = \|\boldsymbol{x}_{ij} - \boldsymbol{m}_{ij}\|^2, \tag{3}$$

which penalizes the trivial zero solution. Note that (3) and (2) both vanish when $(\boldsymbol{x}_{ij}, z_{ij}) = (\boldsymbol{m}_{ij}, 1)$. The proposed pOSE objective

$$\sum_{ij} \left((1-\eta)\ell_{\text{OSE}}(\boldsymbol{x}_{ij}, z_{ij}) + \eta\ell_{\text{aff}}(\boldsymbol{x}_{ij})\right), \tag{4}$$

where $0 < \eta < 1$, therefore allows arbitrary starting solutions but penalizes projective depths that deviate significantly from 1. The red curve of Figure 2 (top) shows pOSE values (with $\eta = 0.1$) over the line $(1-t)(0.5, 0, -1) + t(0.5, 0, 1)$. In contrast to the reprojection error, the pOSE formulation does not give any barrier at $z = 0$. It is experimentally shown in [15] that when optimized using VarPro [16] this leads to a method that converges to the right solution in the vast majority of cases starting from random initialization (including starting points with negative depths). Note that if we column-stack the camera matrices P_i into a matrix P with 4 columns, and similarly row-stack the 3D points into a matrix U with 4 rows, the resulting product $X = PU$ is a matrix of rank 4. We can therefore formulate the pOSE objective as a low-rank recovery problem

$$\min_{\text{rank}(X)=4} \|\mathcal{A}(X) - b\|^2, \tag{5}$$

where $\mathcal{A}$ is a linear operator. It is well known from compressed sensing that such formulations can often be solved optimally [6, 9–11, 19, 28, 30, 33]. The optimization problem becomes particularly easy for large values of η. On the other hand, the regularization term also introduces an undesirable penalty for large (positive) depths which may constrain the range of feasible depths too much to achieve satisfactory solutions. The bottom images in Figure 2 show the same evaluation as the top ones but over the line $(1-t)(-0.5, 0, -1) + t(0.5, 0, 1)$. All of the points on this line give 0 reprojection error (except at the camera center $(0, 0, 0)$ for which the projection is undefined). The pOSE formulation (red curve) clearly penalizes solutions of small or negative projective depth but its undesirable growth for large positive values is also visible.

2.3. Exponential Regularization

In this paper, we instead propose to regularize the depth using an exponential function (yellow curves in Figure 2). Specifically, we replace the affine term (3) with

$$\ell_{\text{exp}}(\boldsymbol{x}_{ij}, z_{ij}) = e^{-\left(\frac{\boldsymbol{m}_{ij}\boldsymbol{x}_{ij} + z_{ij}}{\sqrt{\|\boldsymbol{m}_{ij}\|^2+1}}\right)}. \tag{6}$$

The term $\frac{\boldsymbol{m}_{ij}\boldsymbol{x}_{ij} + z_{ij}}{\sqrt{\|\boldsymbol{m}_{ij}\|^2+1}}$ is the length (with sign) of the projection of the vector $(\boldsymbol{x}_{ij}, z_{ij})$ onto $(\boldsymbol{m}_{ij}, 1)$. Note that its sign is negative when the angle between $(\boldsymbol{x}_{ij}, z_{ij})$ and $(\boldsymbol{m}_{ij}, 1)$ is larger than $90°$. The exponential function will penalize such values heavily. Still, the penalty is finite for all values making it is possible to use start the optimization from anywhere. On the other hand for positive growing values the exponential function tends to 0 and therefore does not restrict the feasible projective depths as the affine term (3) does.

The proposed expOSE objective is then

$$\ell_{\text{expOSE}} = \sum_{ij} (1-\eta)\ell_{\text{OSE}}(\boldsymbol{x}_{ij}, z_{ij}) + \eta\ell_{\text{exp}}(\boldsymbol{x}_{ij}, z_{ij}). \tag{7}$$

At first glance it may seem as if replacing (3) with (6) will yield an ill-posed problem since large depths are hardly penalized by (6). Adding a small penalty for these values to ensure a well-posed problem may therefore be warranted. Note, however, that unless there is an exact solution (with zero reprojection errors) the OSE term is not scale invariant but has a weak shrinking bias. In practice, we empirically observe that this bias is generally enough for our proposed algorithm to converge well from random starting solutions.

We conclude this section by noting that our proposed method is much less sensitive to parameter selection than the original pOSE model [15]. Since the shrinking bias of the OSE term is relatively weak an increased regularization cost, due to a change of parameters, can often be compensated for by changing the scale of the reconstruction. In contrast, the choice of η in the original pOSE model is crucial. Figure 1 shows how η affects the reconstruction (more details about this figure are provided in Section 3.2).

3. Optimization with VarPro

One of the main benefits of the pOSE formulation [15] is that it is quadratic in the elements of X. Therefore, given values for camera matrices P the optimal 3D points $U^*(P)$ can be computed in closed form using a pseudo inverse. The VarPro method [16, 27, 35, 41] solves the reduced problem

$$\min_P \|\mathcal{A}\left(PU^*(P)\right) - b\|^2, \tag{8}$$

using the Levenberg-Marquardt method [12, 39]. In contrast to standard Gauss-Newton type methods that optimize

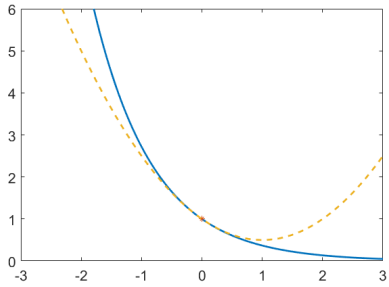

Figure 3. The exponential function and its Taylor approximation.

locally over both U and P, the main benefit of the elimination of U is that dampening only needs to be applied to P. This has been shown empirically to greatly improve convergence [14, 16]. The intuition is that small changes in P will sometimes result in large changes in U, but this is prevented by a dampening term which causes the algorithm to stall.

Since the exponential regularization term is not quadratic VarPro is not directly applicable to our formulation. We, therefore, employ an iterative approach that locally approximates (6) with a quadratic function. Consider the 2nd order Taylor expansion of $e^{-\boldsymbol{a}^T\boldsymbol{y}}$ at a point $\bar{\boldsymbol{y}}$ given by

$$e^{-\boldsymbol{a}^T\boldsymbol{y}} \approx e^{-\boldsymbol{a}^T\bar{\boldsymbol{y}}}\left(1-\boldsymbol{a}^T(\boldsymbol{y}-\bar{\boldsymbol{y}})+\frac{1}{2}(\boldsymbol{a}^T(\boldsymbol{y}-\bar{\boldsymbol{y}}))^2\right). \tag{9}$$

Completing squares gives the expression

$$e^{-\boldsymbol{a}^T\boldsymbol{y}} \approx \frac{e^{-\boldsymbol{a}^T\bar{\boldsymbol{y}}}}{2}\left(\boldsymbol{a}^T(\boldsymbol{y}-\bar{\boldsymbol{y}})-1\right)^2 + e^{-\boldsymbol{a}^T\bar{\boldsymbol{y}}}. \tag{10}$$

Note that when minimizing with respect to $\boldsymbol{y}$ the last term is constant and can be ignored. Since the exponential function is positive the result is a weighted linear least squares term in the unknown $\boldsymbol{y}$. With $\boldsymbol{y} = \begin{bmatrix}\boldsymbol{x}_{ij}\\ z_{ij}\end{bmatrix}$ and $\boldsymbol{a} = \frac{1}{\sqrt{\|\boldsymbol{m}_{ij}\|^2+1}}\begin{bmatrix}\boldsymbol{m}_{ij}\\ 1\end{bmatrix}$ we get our approximation

$$\tilde{\ell}_{\text{exp}}(\boldsymbol{x}_{ij}, z_{ij}) \approx \frac{\ell_{\text{exp}}(\bar{\boldsymbol{x}}_{ij}, \bar{z}_{ij})}{2}\left(\frac{\boldsymbol{m}_{ij}^T\Delta\boldsymbol{x}_{ij}+\Delta z_{ij}}{\sqrt{\|\boldsymbol{m}_{ij}\|^2+1}}-1\right)^2, \tag{11}$$

where $\Delta\boldsymbol{x}_{ij} = \boldsymbol{x}_{ij} - \bar{\boldsymbol{x}}_{ij}$ and $\Delta z_{ij} = z_{ij} - \bar{z}_{ij}$. To the left in Figure 3 we show e^{-ay} with $a = 1$ (blue curve), and the Taylor approximation at $\bar{y} = 0$ (orange dashed curve). In the supplementary material, we compare level sets of the expOSE objective, its approximation, and pOSE.

3.1. The expOSE Model

Replacing the exponential regularization in (7) with the quadratic approximation (11) at $\bar{\boldsymbol{y}}_{ij}$ results in a quadratic loss that can be written as $\|\mathcal{A}(PU)-b\|^2$, which can be optimized using VarPro as described in Algorithm 1. The linear operator $\mathcal{A}$ and the vector b can be computed in each iteration based on the image measurements $\boldsymbol{m}_{ij}$, the current estimations $\bar{\boldsymbol{y}}_{ij}$ and η. For the initial approximation of the regularization, we use $\bar{\boldsymbol{y}}_{ij} = (\boldsymbol{m}_{ij}, 1)$.

Algorithm 1: VarPro for solving expOSE (7)

```
Normalize image measurements by removing the mean and
  dividing by 3 standard deviations;
Select the inputs η, and randomly initialize elements of P from a
  normal distribution of unit std ;
Set ȳ_ij = [m_ij^T, 1]^T;
Set up A and b by approximating the exponential regularization
  by a quadratic form around each ȳ_ij;
Compute U by minimizing (7) with P fixed;
Set do update = 0 if scheduling update of regularization is
  considered, otherwise do update = 1;
while true do
    Compute the Jacobians
    J_P = A(U^T ⊗ I); J_U = A(I ⊗ P);
    and the residuals r = Avec(PU) − b;
    Compute P_new and U_new from J_P, J_U, and r as
      P_new = P + ΔP and U_new = U + ΔU, with
      ΔP = (J_P^T(I − J_U J_U^†)J_P + λI)^{-1} J_P^T r, and
      ΔU = −J_U^†(r + J_P ΔP);
    Evaluate the loss ℓ_new;
    if ℓ_new < ℓ_best then
        ℓ_best = ℓ_new ;
        P ← P_new; and U ← U_new;
        if do update then
            Set ȳ_ij = P_i U_j;
            Set up A and b by approximating the regularization
              by a quadratic form around each ȳ_ij;
        end
    end
    if stopping criterion then
        if do update then
            break;
        else
            do update = 1
        end
    end
end
```

Regularization update scheduling: In order to improve the convergence of the algorithm, we propose to keep the initial quadratic approximation of the regularization (11) either for a fixed number of iterations or until convergence of the initial approximation. This delays the approximation of the exponential regularization in each iteration until a stable initial solution with positive depths is found. In Section 3.2 we show empirically the advantage of doing so.

Data normalization: Since our regularization term is geometrically motivated and our approach replaces reprojection error with OSE it is important to use normalization of the image data to achieve a well-conditioned formulation [13]. Here we follow standard approaches: We first subtract the image center from all image points, then divide them with the resulting standard deviation over the image.

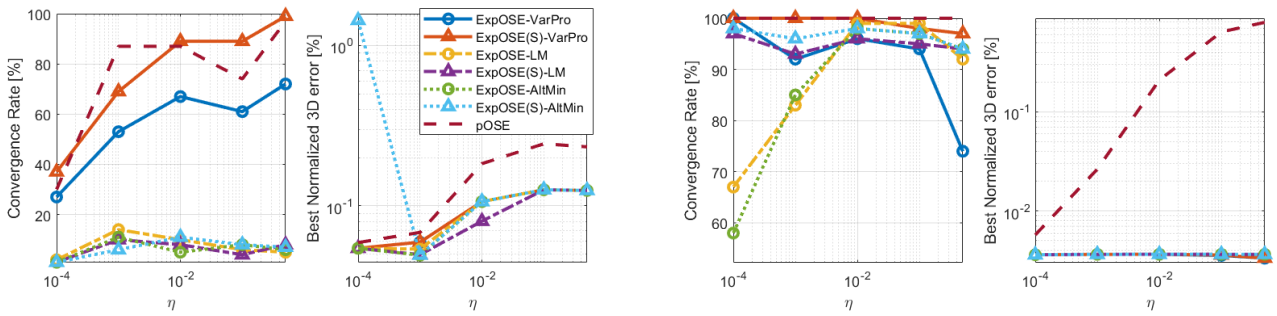

(a) Dino (Small) (b) Fountain

Figure 4. Comparison of convergence rate and normalized 3D error of different methods on the Dino (a) and Fountain (b) datasets. The metrics are obtained by running 100 instances starting from random initializations. In dashed we should the metrics for the pOSE baseline.

3.2. Performance evaluation of expOSE

Before presenting our model for radial distortion we evaluate the effects of using exponential regularization with the standard OSE. We use the Dino (Small) [3] (36 cameras, 319 points, 77% missing data) and Fountain [34] (11 cameras, 1167 points, 23% missing data) datasets to evaluate the performance of expOSE with varying parameters - the weight η and scheduling of regularization update-, and optimization strategies - VarPro, Levenberg-Marquardt (LM), and Alternating Minimization (AltMin) [7].

The metrics used for the comparisons are convergence rate of the algorithm and relative 3D error to GT. The convergence rate is calculated by counting the number of times the algorithms converged to the lowest loss over 100 problem instances starting from random initializations (a threshold of 2% above the smallest loss value is used). The 3D error is computed as $e_{3D} = \frac{\|U' - U_{\text{GT}}\|}{\|U_{\text{GT}}\|_F}$ where U' is the result of performing projective registration of the factor U to the ground-truth point cloud U_{GT}. In this way, we are able to measure the quality of the factors U that are outputted by each method. For a fair comparison, we compute the 3D errors for solutions that converged to the desired optimum.

The methods are implemented in MATLAB, and we let each method perform a maximum of 500 iterations. For the case of regularization update scheduling, which we call expOSE(S), we delay the update of the regularization quadratic approximation by 250 iterations or until the initial optimization converges - whichever occurs first.

Effect of η and scheduling: The performance of expOSE is evaluated for multiple values of η ranging from 10^{-4} to 0.5. The results are plotted in Figure 4. We show that expOSE is significantly more robust to η than pOSE in terms of 3D errors (see also Figure 1). We also show that delaying the update of the quadratic approximation of the regularization results in a significant boost in convergence rate, allowing us to achieve rates similar to pOSE.

Comparison with other optimization strategies: We compare the performance of expOSE (with and without scheduling) when using VarPro, LM and AltMin. The results confirm that, just like with pOSE, VarPro is the most reliable method for expOSE, while LM and AltMin achieve poor convergence rates.

4. Robustness to Radial Distortion

In the previous sections, we considered modifications to the original pOSE model which assumes a regular pinhole camera. In [18] the RpOSE model which instead uses a radial camera [20, 21, 23, 24, 29, 37, 40] is presented. This model is invariant to radial distortion which the standard pOSE model does not handle. We note however that the radial model requires more data for parameter estimation since it essentially only measures errors in one direction of the image. To address this issue we introduce an intermediate model by decomposing the reprojection error into a tangential and a radial component. By down-weighting the tangential error we obtain a model that is more robust to radial distortion than the pinhole camera but less sensitive to missing data than the radial camera. We then introduce an exponential regularization term for this model.

4.1. Decoupling Tangential and Radial Errors

When working with the radial camera model it is typically assumed that the principal point and the distortion center are the center of the image and have coordinates $(0, 0)$. We make the same assumption here.

The reprojection error is obtained by taking the length of the error vector $\boldsymbol{e}(\boldsymbol{x}, z) = \frac{\boldsymbol{x}}{z} - \boldsymbol{m}$. The coordinates of this vector are given in w.r.t. the canonical image basis $(1, 0)$ and $(0, 1)$ of the image and can be interpreted as errors in the x- and y-directions respectively. For a point $\boldsymbol{m}$ we are interested in measuring the error in the radial direction $\frac{\boldsymbol{m}}{\|\boldsymbol{m}\|}$ and the tangential direction $\frac{\boldsymbol{m}_\perp}{\|\boldsymbol{m}\|}$, where $\boldsymbol{m}_\perp$ is the orthogonal vector to $\boldsymbol{m}$ (see Figure 5). We, therefore, write the error vector as a linear combination of these. It is not diffi-

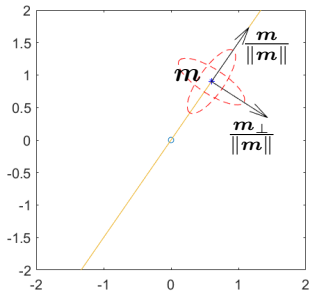

Figure 5. Levelsets (red ellipses) of ℓ_{wOSE} for $\alpha = 0.1$ and 0.9. Here $\boldsymbol{m} = (0.6, 0.9)$ and $z = 1$.

cult to verify that

$$\frac{\boldsymbol{x}}{z} - \boldsymbol{m} = \left(\frac{\boldsymbol{m}^T\boldsymbol{x}}{z\|\boldsymbol{m}\|} - \|\boldsymbol{m}\|\right)\frac{\boldsymbol{m}}{\|\boldsymbol{m}\|} + \frac{\boldsymbol{m}_\perp^T\boldsymbol{x}}{\|\boldsymbol{m}\|z}\frac{\boldsymbol{m}_\perp}{\|\boldsymbol{m}\|}. \quad (12)$$

In the basis $\frac{\boldsymbol{m}}{\|\boldsymbol{m}\|}, \frac{\boldsymbol{m}_\perp}{\|\boldsymbol{m}\|}$ the error vector can be written as

$$\boldsymbol{e}(\boldsymbol{x}, z) = \frac{1}{\|\boldsymbol{m}\|}\begin{bmatrix}\boldsymbol{m}^T \\ \boldsymbol{m}_\perp^T\end{bmatrix}\frac{\boldsymbol{x}}{z} - \begin{pmatrix}\|\boldsymbol{m}\| \\ 0\end{pmatrix}. \quad (13)$$

Independently of the basis chosen, the reprojection error is nonlinear due to the division by z, making it unsuitable for optimization. The OSE in the new basis is obtained by rescaling the reprojection error $\boldsymbol{e}(\boldsymbol{x}, z)$ by the depth z. The expression for OSE error in the new basis is therefore

$$\|z\boldsymbol{e}(\boldsymbol{x}, z)\|^2 = \left(\frac{\boldsymbol{m}^T}{\|\boldsymbol{m}\|}\boldsymbol{x} - \|\boldsymbol{m}\|z\right)^2 + \left(\frac{\boldsymbol{m}_\perp^T}{\|\boldsymbol{m}\|}\boldsymbol{x}\right)^2. \quad (14)$$

4.2. Reweighting the Error Components

Radial distortion is usually modeled by modifying the projection according to

$$\kappa_r(\boldsymbol{m})\boldsymbol{m} = \frac{\boldsymbol{x}}{z} \quad (15)$$

where κ_r is a scalar that depends on the distance to the distortion center. It is clear that the second term of (14) vanishes when inserting $(\boldsymbol{x}, z)$ fulfilling (15) for any κ_r, but not the first term. To handle radial distortion we could incorporate the additional parameter κ_r in (14) and explicitly estimate it. Unfortunately, this results in a more complex model (with trilinear interactions) making optimization difficult. Alternatively, to achieve robustness to radial distortion we can remove the first term, as in [18]. The downside of doing this is that it removes roughly half of the data (one out of two coordinates for each projection) available for use in inference. Therefore we here propose to compensate for the unknown radial distortion by down-weighting the first term or equivalently allowing a larger standard deviation in the radial direction.

Let σ_r^2 and σ_t^2 denote the uncertainties of the reprojection error $\epsilon = s\boldsymbol{x}/z - \boldsymbol{m}$ along the radial and tangential direction, respectively, and where s is an unknown positive scalar that models radial distortion effects and focal length scaling. Assuming the reprojection error ϵ is sampled from a 2D normal distribution $\mathcal{N}(0, \Sigma)$, the probability of the model $\{\boldsymbol{x}, z\}$ given $\boldsymbol{m}$ is

$$P(\boldsymbol{x}, z|\boldsymbol{m}) = \frac{1}{2\pi\det(\Sigma)^{1/2}}e^{-s^2(\frac{1}{s}\boldsymbol{m}-\boldsymbol{x}/z)^T\Sigma^{-1}(\frac{1}{s}\boldsymbol{m}-\boldsymbol{x}/z)}. \quad (16)$$

Maximizing the likelihood (16) w.r.t. $\{\boldsymbol{x}, z\}$ is equivalent to minimizing

$$\frac{s^2}{\sigma_r^2}\left(\frac{\boldsymbol{m}^T}{\|\boldsymbol{m}\|}\frac{\boldsymbol{x}}{z} - \frac{1}{s}\|\boldsymbol{m}\|\right)^2 + \frac{s^2}{\sigma_t^2}\left(\frac{\boldsymbol{m}_\perp^T}{\|\boldsymbol{m}\|}\frac{\boldsymbol{x}}{z}\right)^2, \quad (17)$$

where $\Sigma = R^T\text{diag}(\sigma_r^2, \sigma_t^2)R$, and R is a rotation matrix that aligns the coordinate axis with $\boldsymbol{m}/\|\boldsymbol{m}\|$ and $\boldsymbol{m}_\perp/\|\boldsymbol{m}\|$. While the second term quadratic term of (17) is not affected by s, in the first term $\|\boldsymbol{m}\|$ is weighted by $1/s$, which is undesirable as previously motivated. We propose to approximate (17) by

$$\underbrace{\frac{1}{\sigma_r^2}}_{(1-\alpha)}\left(\frac{\boldsymbol{m}^T}{\|\boldsymbol{m}\|}\frac{\boldsymbol{x}}{z} - \|\boldsymbol{m}\|\right)^2 + \underbrace{\frac{1}{\sigma_t^2}}_{\alpha}\left(\frac{\boldsymbol{m}_\perp^T}{\|\boldsymbol{m}\|}\frac{\boldsymbol{x}}{z}\right)^2. \quad (18)$$

This approximation of the first term adds a bias to the obtained solution based on the unknown shift $\left(\frac{1}{s} - 1\right)\|\boldsymbol{m}\|$. We regulate the effect of this bias - and thus the robustness to radial distortion - by controlling the relative weight of the first quadratic term (biased) versus the second quadratic term (unbiased) through the value of $\alpha \in [0, 1]$. For the extreme case of $\alpha = 1$ the radial component of the error is completely dropped resulting in the loss presented in [18]. Linear residuals can be obtained by replacing (18) with its component-weighted OSE counterpart

$$\ell_{\text{wOSE}} = (1-\alpha)\left(\frac{\boldsymbol{m}^T}{\|\boldsymbol{m}\|}\boldsymbol{x} - \|\boldsymbol{m}\|z\right)^2 + \alpha\left(\frac{\boldsymbol{m}_\perp^T}{\|\boldsymbol{m}\|}\boldsymbol{x}\right)^2. \quad (19)$$

Figure 5 shows an example of level sets (in the image plane $z = 1$) for $\alpha = 0.1$ and 0.9.

Note that the same approach can be used to handle unknown focal lengths. If we assume that the intrinsic calibration matrix of the camera is $\mathbf{K} = \text{diag}(f, f, 1)$, the relation between the reprojected point and the image measurement is $\frac{\kappa_r}{f}\boldsymbol{m} = \frac{\boldsymbol{x}}{z}$ and therefore the re-weighted formulation can be applied to this setting as well. An unknown/varying focal length f is however modeled by the standard pOSE model in contrast to κ_r which depends on the distance between the projection and the principal point and thus cannot be included in a factorization algorithm without adding extra variables.

4.3. Regularization for radial distortion invariance

Weighting differently the radial and tangential of the OSE does not change, in general, the exponential regularization described in Section 2. However, one must note that for the extreme case $\alpha = 1$, for a given $X = PU$ the variables in every third row of X and P vanish from the OSE. In other words, decreasing the total loss will always be possible by increasing z through the third row of P, and consequently decreasing the $e^{\frac{z}{\sqrt{\|\boldsymbol{m}\|^2+1}}}$ part of the exponential regularization. To avoid such undesirable behavior, we proposed an alternative exponential regularization for the particular case $\alpha = 1$ acting only on x and y, i.e.,

$$\ell_{\text{exp}} = e^{-\frac{\boldsymbol{m}^T}{\|\boldsymbol{m}\|}\boldsymbol{x}} \tag{20}$$

This alternative regularization enforces the reprojection $\boldsymbol{x}$ according to the 1D radial camera model $\boldsymbol{m} = \lambda\boldsymbol{x}$ to have positive scale $\lambda > 0$, canceling out the shrinking bias of the OSE as in the general case.

The expOSE loss for weighted radial and tangent components of the OSE can then be approximated as

$$\ell_{\text{expOSE}} = \sum_{ij}(1-\eta)\ell_{\text{wOSE}}(\boldsymbol{x}_{ij}, z_{ij}) + \eta\tilde{\ell}_{\text{exp}}(\boldsymbol{x}_{ij}, z_{ij}) \tag{21}$$

with $\tilde{\ell}_{\text{exp}}$ defined as

$$\begin{cases} \frac{\ell_{\text{exp}}(\bar{\boldsymbol{x}}_{ij}, \bar{z}_{ij})}{2}\left(\frac{\boldsymbol{m}_{ij}^T\Delta\boldsymbol{x}_{ij}+\Delta z_{ij}}{\sqrt{\|\boldsymbol{m}_{ij}\|^2+1}} - 1\right)^2, & \alpha \in [0, 1[\\ \frac{\ell_{\text{exp}}(\bar{\boldsymbol{x}}_{ij})}{2}\left(\frac{\boldsymbol{m}_{ij}^T\Delta\boldsymbol{x}_{ij}}{\|\boldsymbol{m}_{ij}\|} - 1\right)^2, & \alpha = 1 \end{cases}. \tag{22}$$

This radial distortion robust version of expOSE can be optimized following Algorithm 1 nonetheless since both the component-weighted OSE and the quadratic approximation of the regularization can still be written as $\|\mathcal{A}(PU) - b\|^2$.

5. Outline of Full Reconstruction Pipeline

We propose to use expOSE as a solution to uncalibrated and radial distortion invariant Structure-from-Motion. A few Bundle Adjustment steps can be performed for further refinement. The pipeline takes as input 2D image measurements of points tracked along multiple views, just like any other factorization-based SfM pipeline. The proposed radial distortion-invariant pipeline can be decomposed into the following sequential modules:

1. expOSE factorization: Given a set of image points tracked along several images, we use Algorithm 1 to obtain estimations of the uncalibrated camera matrix, and the 3D points, up to projective ambiguity.

2. Radial distortion estimation (and camera matrix completion): Using the solution obtained with expOSE, the distortion parameters and, for $\alpha = 1$, the third row of the uncalibrated camera matrix are estimated from the equations in (15). Note that by assuming a Brown-Conrady radial distortion model [2] with $\kappa(\boldsymbol{m}) = \sum_j k_j\|\boldsymbol{m}\|^{2j}$, for each camera a system of equations of the form
$$M_i\begin{bmatrix} p_i^{(3)} \\ \mathbf{k} \end{bmatrix} = b_i \tag{23}$$
can be obtained, where $p_i^{(3)}$ is the third row of the ith camera matrix, and $\mathbf{k}$ is a vector of the distortion parameters. Here we use a distortion model with three parameters, $k_j, j = 1, ..., 3$. Assuming that the distortion model is constant along all views, the overall system of equations can be written as $M[{p^{(3)}}^T, \mathbf{k}^T]^T = b$, with p being a $4\times$ #views vector with all third rows of the camera matrices. For $\alpha = 1$ both $p^{(3)}$ and $\mathbf{k}$ are unknowns and are estimated in this step. For $\alpha \neq 1$, the system can be simplified to $M\mathbf{k} = b - Mp^{(3)}$ since $p^{(3)}$ is already estimated by expOSE. If it is assumed that there is no radial distortion and $\alpha \neq 1$, then this step can be completely skipped.

3. Bundle adjustment: We perform local optimization of
$$\sum_{ij}\left\|\boldsymbol{m}_{ij} - (1+\kappa(\boldsymbol{m}_{ij}))\frac{\boldsymbol{x}_{ij}}{z_{ij}}\right\|^2 \tag{24}$$
starting from the estimations of P, X, and $\mathbf{k}$ found with the previous steps. The optimization is solved using Levenberg–Marquardt algorithm. If there is no radial distortion then the parameters $\mathbf{k}$ can be set to zero and kept constant during optimization. For expOSE initialization, we observe that usually only a few steps are needed (5-10 steps).

4. Euclidean update: Finally we estimate the projective transformation $H \in \mathbb{R}^{4\times4}$ such that the factorization $\{PH, H^{-1}X\}$ is a Euclidean reconstruction. This is done by estimating the dual absolute conic as described in [12].

5.1. Experiments

The performance of the proposed pipeline is evaluated on 3 sequences from [24] with radial distortion: Grossmunster (19 cam., 1874 pts, 41% missing data), Kirchenge (30 cam., 1158 pts, 60% missing data), and Munterhof (20 cam., 2108 pts, 42% missing data). We compare the performance when using either expOSE ($\eta = 0.01$), pOSE, or RpOSE (both with $\eta = 0.001$) in step 1 of the pipeline. We use expOSE with scheduling for regularization update, as described in Section 3.2. Refinement of the solutions is done by performing up to 50 iterations of BA.

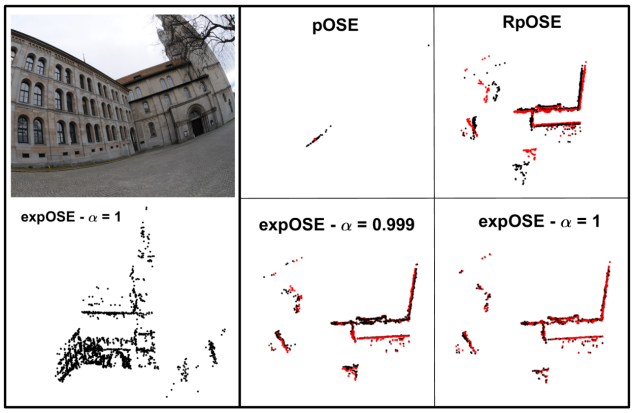

Figure 6. Visualization of reconstructions on the Grossmunster sequence. (Left) An example of one of the images on the sequence. At the bottom, we show a view of the 3D reconstruction of expOSE for $\alpha = 1$. (Right) Comparison between the top view reconstructions (black) obtained with pOSE, RpOSE and expOSE. In red we show the ground-truth 3D point cloud. All reconstructions shown here were not refined with bundle adjustment.

The metrics used are convergence rate (similarly to the experiments in Section 2), 2D reprojection error, rotation error, and 3D error. In order to compute the last two, we perform Euclidean registration on the output of the pipeline, i.e. after the Euclidean update, to the ground-truth 3D point cloud. The inverse of that Euclidean transformation is applied to the camera matrices. Rotation error is then computed as $e_{\text{rot}} = \text{acos}\left(\left(\text{trace}\left(R_i^{GT} R_i^T\right) - 1\right)/2\right)$ and the 3D error as the median of all $\|X_j - X_j^{GT}\|$. The values presented in Table 1 correspond to the average over all instances that converged to the desired optimum. The chosen metrics are evaluated at two points of the pipeline: after the radial distortion estimation (step 2), and after the bundle adjustment (step 3). At both stages, a metric update is performed in order to obtain a Euclidean reconstruction.

The results show that expOSE clearly outperforms both pOSE and RpOSE. The difference in performance is even more evident when looking at the output of the factorizations, where expOSE was able to achieve reprojection errors that almost match the refined solution with BA. Note that in many cases expOSE even got better rotation and 3D errors than its refined counterpart. A visualization for the Grossmunter sequence is shown in Figure 6. It is also possible to notice the impact of using the regularization for radial distortion invariance as described in Section 4.3. For $\alpha = 0.999$ the method has slow convergence, leading to poor solutions as can be seen by the high rotation and reprojection errors. Additional results for other values of α and sequences are presented in the supplementary material.

In practice, as seen in these experiments, we notice that $\alpha = 1$ achieves the best results for images with radial distortion. In the supplementary material we provide additional experiments that show the benefit of using values $1/2 < \alpha < 1$ in particular problem instances where data availability is too low for the stability of a pure radial model (e.g. few viewpoints and/or points per camera available).

Table 1. Results on the Grossmunster, Kirchenge, and Munsterhof datasets (over 10 instances). For each method two rows are presented: the first consists of the results for the output of the factorization method; the second of the output of the Bundle Adjustment (+BA). In green, we show the best results for each metric.

		Conv. Rate	Rot. [deg]	3D [unit]	2D [pix]
Grossmunster					
pOSE		50%	148.25	0.762	18.48
	+ BA	50%	27.61	0.293	1.50
RpOSE		90%	2.24	0.082	2.91
	+ BA	90%	0.53	0.011	1.48
ExpOSE	α=0.999	100%	44.74	0.227	41.51
	α=0.999+BA	100%	0.43	0.007	1.48
	α=1	100%	0.18	0.004	1.86
	α=1+BA	100%	0.42	0.006	1.48
Kirchenge					
pOSE		100%	160.38	6.844	14.95
	+ BA	100%	0.72	0.024	1.22
RpOSE		90%	0.98	0.062	1.94
	+ BA	90%	1.06	0.031	1.22
ExpOSE	α=0.999	60%	24.71	0.022	45.28
	α=0.999+BA	80%	1.19	0.021	1.22
	α=1	80%	0.51	0.026	1.57
	α=1+BA	80%	2.92	0.050	1.22
Munsterhof					
pOSE		100%	14.01	0.230	12.08
	+ BA	100%	0.44	0.027	1.70
RpOSE		60%	1.00	0.071	11.96
	+ BA	60%	0.44	0.027	1.70
ExpOSE	α=0.999	100%	20.13	0.021	47.71
	α=0.999+BA	100%	0.47	0.029	1.70
	α=1	80%	0.12	0.013	3.43
	α=1+BA	90%	0.45	0.030	1.70

6. Conclusions

In this paper, we propose the use of exponential regularization on projective factorization problems as a way to enforce Cheirality conditions on the reconstruction. Radial distortion robustness is achieved by weighting differently the radial and tangential components of the object space error. We show that the proposed regularization results in higher reconstruction quality (that matches bundle adjustment refined solutions) while keeping the same convergence properties as state-of-the-art factorization methods and being less sensitive to the choice of the weight η of the regularization.

References

[1] Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Global optimality of local search for low rank matrix recovery. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3873–3881. Curran Associates, Inc., 2016. 1

[2] Dean Brown. Decentering distortion of lenses. 1966. 7

[3] A. M. Buchanan and A. W. Fitzgibbon. Damped newton algorithms for matrix factorization with missing data. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 1, 5

[4] Alessio Del Bue, João M. F. Xavier, Lourdes Agapito, and Marco Paladini. Bilinear modeling via augmented lagrange multipliers (BALM). *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(8):1496–1508, 2012. 1

[5] R. Cabral, F. De la Torre, J. P. Costeira, and A. Bernardino. Unifying nuclear norm and bilinear factorization approaches for low-rank matrix decomposition. In *International Conference on Computer Vision (ICCV)*, 2013. 1

[6] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009. 3

[7] I. Csiszár and G. Tusnády. Information Geometry and Alternating Minimization Procedures. 5

[8] Y. Dai, H. Li, and M. He. Projective multiview structure and motion from element-wise factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2238–2251, 2013. 1

[9] Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. *arXiv preprint*, arxiv:1704.00708, 2017. 1, 3

[10] Rong Ge, Jason D. Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2016. 1, 3

[11] Christian Grussler, Anders Rantzer, and Pontus Giselsson. Low-rank optimization with convex constraints. *IEEE Transactions on Automatic Control*, 63(11):4000–4007, 2018. 3

[12] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, USA, 2 edition, 2003. 2, 3, 7

[13] Richard I. Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(6):580–593, 1997. 4

[14] Je Hyeong Hong and Andrew Fitzgibbon. Secrets of matrix factorization: Approximations, numerics, manifold optimization and random restarts. In *Int. Conf. on Computer Vision*, 2015. 1, 4

[15] Je Hyeong Hong and Christopher Zach. pose: Pseudo object space error for initialization-free bundle adjustment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2, 3

[16] J. H. Hong, C. Zach, and A. Fitzgibbon. Revisiting the variable projection method for separable nonlinear least squares problems. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5939–5947, 2017. 1, 3, 4

[17] Je Hyeong Hong, Christopher Zach, Andrew W. Fitzgibbon, and Roberto Cipolla. Projective bundle adjustment from arbitrary initialization using the variable projection method. In *European Conf. on Computer Vision*, 2016. 1

[18] Jose Iglesias and Carl Olsson. Radial distortion invariant factorization for structure from motion. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 1, 2, 3, 5, 6

[19] José Pedro Iglesias, Carl Olsson, and Marcus Valtonen Örnhag. Accurate optimization of weighted nuclear norm for non-rigid structure from motion. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 3

[20] Jae-Hak Kim, Yuchao Dai, Hongdong li, Xin Du, and Jonghyuk Kim. Multi-view 3d reconstruction from uncalibrated radially-symmetric cameras. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1896–1903, 12 2013. 5

[21] Z. Kukelova, M. Bujnak, and T. Pajdla. Real-time solution to the absolute pose problem with unknown radial distortion and focal length. In *2013 IEEE International Conference on Computer Vision*, pages 2816–2823, 2013. 5

[22] Suryansh Kumar. Non-rigid structure from motion: Prior-free factorization method revisited. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 51–60. IEEE, 2020. 1

[23] Viktor Larsson, Torsten Sattler, Zuzana Kukelova, and Marc Pollefeys. Revisiting radial distortion absolute pose. In *International Conference on Computer Vision (ICCV)*. IEEE, September 2019. 5

[24] Viktor Larsson, Nicolcas Zobernig, Kasim Taskin, and Marc Pellefeys. Calibration-free structure-from-motion with calibrated radial trifocal tensors. In *European Conference of Computer Vision*, 2020. 5, 7

[25] Ludovic Magerand and Alessio Del Bue. Practical projective structure from motion (p2sfm). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 39–47, 2017. 2

[26] Behrooz Nasihatkon, Richard I. Hartley, and Jochen Trumpf. A generalized projective reconstruction theorem and depth constraints for projective factorization. *Int. J. Comput. Vis.*, 115(2):87–114, 2015. 2

[27] Takayuki Okatani and Koichiro Deguchi. On the wiberg algorithm for matrix factorization in the presence of missing components. *International Journal of Computer Vision*, 72(3):329–337, 2007. 3

[28] Carl Olsson, Daniele Gerosa, and Marcus Carlsson. Relaxations for non-separable cardinality/rank penalties. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, IEEE International Conference on Computer Vision Workshops, pages 162–171, 2021. 3

[29] Carl Olsson, Viktor Larsson, and Fredrik Kahl. A quasiconvex formulation for radial cameras. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14571–14580, 2021. 5

[30] Marcus Valtonen Ornhag, Carl Olsson, and Anders Heyden. Bilinear parameterization for differentiable rank-regularization. *2020 IEEE/CVF Conference on Computer*

Vision and Pattern Recognition Workshops (CVPRW), Jun 2020. 1, 3

[31] Dohyung Park, Anastasios Kyrillidis, Constantine Carmanis, and Sujay Sanghavi. Non-square matrix sensing without spurious local minima via the Burer-Monteiro approach. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 65–74, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR. 1

[32] Conrad J. Poelman and Takeo Kanade. A paraperspective factorization method for shape and motion recovery. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(3):206–218, 1997. 1

[33] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, Aug. 2010. 3

[34] C. Strecha, W. von Hansen, L. Van Gool, P. Fua, and U. Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 5

[35] D. Strelow, Q. Wang, L. Si, and A. Eriksson. General, nested, and constrained wiberg minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1803–1815, 2016. 3

[36] Peter F. Sturm and Bill Triggs. A factorization based algorithm for multi-image projective structure and motion. In *Proceedings of the 4th European Conference on Computer Vision-Volume II - Volume II*, ECCV '96, page 709–720, Berlin, Heidelberg, 1996. Springer-Verlag. 1

[37] SriRam Thirthala and Marc Pollefeys. Radial multi-focal tensors. *International Journal of Computer Vision - IJCV*, 96, 06 2012. 5

[38] Carlo Tomasi and Takeo Kanade. Shape and motion from image streams under orthography: A factorization method. *International Journal of Computer Vision*, 9(2):137–154, 1992. 1

[39] Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment - a modern synthesis. In *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice*, ICCV '99, pages 298–372. Springer-Verlag, 2000. 2, 3

[40] R. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal on Robotics and Automation*, 3(4):323–344, August 1987. 5

[41] T. Wiberg. Computation of principal components when data are missing. In *Proceedings of the Second Symposium of Computational Statistics*, page 229–326, 1976. 3