

Autocalibration: Finding Infinity in a Projective Reconstruction

Neil Cavan, Paul Fieguth, David A. Clausi

*Vision and Image Processing Lab, Department of Systems Design Engineering
University of Waterloo
Waterloo, Ontario
{nrcavan, pfieguth, dclausi}@uwaterloo.ca*

Abstract—In order to extract accurate 3D models from uncalibrated image data it is necessary to upgrade the generated projective reconstructions to a metric space, a process known as autocalibration. The key challenge associated with autocalibration is the nonlinear optimization of a cost function based on extracting camera intrinsics from a potential upgrading transform, and evaluating fitness with respect to prior knowledge of physical cameras. The nonlinearity of the problem leads, in general, to poor convergence and a failure of the calibration process.

This paper presents a novel autocalibration pipeline that seeks to develop a more robust approach to the nonlinear optimization. After testing a variety of methods, none of which yielded satisfactory solutions, we have developed a strategy combining the best aspects of two methods representing the current state of the art. The former method preconditions the projective space by transformation to a quasi-affine reconstruction with respect to camera centers allowing a naive initialization in the new space, and uses a fitness measure resistant to focal length collapse. The latter method initializes using the best results of an exhaustive search over reasonable values of focal length. Our novel approach, presented here, uses the exhaustive search initialization of the latter combined with the improved fitness measure of the former, producing results that outperform both of its predecessors.

Keywords—3D reconstruction, autocalibration, computer vision

I. INTRODUCTION

There is a large and growing demand for robust, automatic 3D model reconstruction from image data across many industries. Recently, approaches such as Google's StreetView using calibrated camera arrays and assumptions about scene structure have met with great success. 3D Reconstruction also plays a large role in augmented reality applications, which embed virtual objects into video sequences, as shown in nearly every Hollywood movie released over the past few years. No matter the application, algorithms that can proceed automatically without human interaction generally require camera calibration, at the least, and often strong assumptions about the scene, such as calibration objects or an environment dominated by lines and planes.

The literature on 3D reconstruction from uncalibrated image data without such assumptions is less well developed. The ability to reconstruct from uncalibrated images is sometimes necessary, such as the motivation for this research: 2G Robotics Inc., seeks to inspect underwater assets such



Figure 1. Six frames from the “Dinosaur” sequence, a well-conditioned test dataset used in the 3D reconstruction literature

as oil rigs, bridge footings and municipal water systems and compare their current state with their state years or decades ago, in order to measure the rate of corrosion. Although there is a wealth of archival footage, camera calibration is impossible because the cameras used were scrapped years ago.

The problem of uncalibrated reconstruction can be divided into two pieces:

- 1) projective reconstruction, which can be accomplished directly from image data and is largely a solved problem, and
- 2) metric upgrade, which consists of solving for the camera calibration using prior knowledge.

Our research focuses on the latter metric upgrade step, which is an active research area and still something of a black art. A variety of methods exist, but there is no clear “right answer” in the literature that works robustly. The method developed in this paper represents a step in that direction.

II. BACKGROUND AND NOTATION

This paper assumes the reader is familiar with the process of 3D reconstruction from images, projective geometry, and the affine and metric strata. For an excellent and thorough treatment of the topic, the interested reader is referred to [1]. For convenience, important results from the literature which are used in this paper are summarized below.

A. Notation

Because 3D reconstruction makes use of both projective and metric geometry, it is important to distinguish quantities

that live in these very different spaces; in general, quantities are assumed to be projective ($\in \mathbb{P}$). Metric quantities ($\in \mathbb{R}$) will be specifically denoted with a tilde ($\tilde{\cdot}$), and normalized projective quantities with their homogeneous co-ordinate scaled to unity will be denoted with a hat ($\hat{\cdot}$). Additionally, the superscript n will be used to refer to structure points in a reconstruction while the subscript m refers to camera views of the scene; thus \mathbf{u}_m^n is the value of the projective vector \mathbf{u} for the n th structure point in the m th view.

B. Projective Reconstruction

Projective reconstructions consist of a set of image observations, camera matrices, and structure points: $\{\mathbf{u}_m^n \in \mathbb{P}^2, \mathbf{P}_m \in \mathbb{R}^{3 \times 4}, \mathbf{X}^n \in \mathbb{P}^3\}$. They can be built using only image data from uncalibrated cameras, such as that shown in Figure 1, and the basis of this process is the projective camera model:

$$\mathbf{u}_m^n \simeq \mathbf{P}_m \mathbf{X}^n \quad (1)$$

where \simeq denotes equality up to scale. There are many algorithms for generating projective reconstructions, which are not discussed here in the interest of brevity. This paper uses a robust MSAC algorithm which has a six-point solution for the trifocal tensor \mathcal{T} as its engine [2], [3]. The only salient feature of this method is that the resulting reconstruction does not have image data \mathbf{u}_m^n transformed to a basis of \mathbb{P}^2 , nor transformed anisotropically. This is not true in many reconstruction methods [4], [5]. Some form of normalization is required for numerical stability, and in this work $\{\mathbf{P}_m\}$ are normalized using the image viewport:

$$\mathbf{P}_m = \left(\frac{1}{2} \begin{bmatrix} \sqrt{w^2 + h^2} & 0 & w \\ 0 & \sqrt{w^2 + h^2} & h \\ 0 & 0 & 1 \end{bmatrix} \right)^{-1} \frac{\mathbf{P}_m}{\|\mathbf{P}_m\|} \quad (2)$$

where w and h are the width and height of the image, respectively. This has the effect of scaling image values between -1 and 1 , and places the mean of image values at approximately 0 assuming the true principal point is near the image center. The unit of the resulting image space is the diagonal dimension of the CCD, a physical dimension which is on the order of focal length for most cameras. The lack of anisotropic scaling means this normalization does not invalidate the assumption of unit pixel aspect ratio during autocalibration.

Projective reconstructions are not unique, and can be altered by a 4×4 projective transformation \mathbf{H} without affecting reprojection error:

$$\mathbf{u}_m^n = (\mathbf{P}_m \mathbf{H}^{-1}) (\mathbf{H} \mathbf{x}^n) \quad (3)$$

This implies that convenient projective frames can be freely chosen. A useful choice is to map a reference camera

(usually the first) to canonical form:

$$\begin{aligned} \mathbf{P}_{ref} \mathbf{H}_{can} &= [\mathbf{I} | \mathbf{0}] \\ \mathbf{H}_{can} &= \begin{bmatrix} \mathbf{P}_{ref} \\ \mathbf{c}(\mathbf{P}_{ref})^T \end{bmatrix} \end{aligned} \quad (4)$$

where $\mathbf{c}(\cdot)$, the projective camera center, is computed as follows:

$$\mathbf{c}(\mathbf{P}) = (c_1, c_2, c_3, c_4)^T \text{ with } c_i = (-1)^i \det(\mathbf{P}^{(i)}) \quad (5)$$

with $\mathbf{P}^{(i)}$ defined as the 3×3 matrix obtained by removing the i th column of \mathbf{P} .

The choice of final row for \mathbf{H}_{can} is somewhat arbitrary, and is often chosen to be $[0 \ 0 \ 0 \ 1]$ if the location of the plane at infinity $\hat{\mathbf{l}}_\infty$ must not be changed; using $\mathbf{c}(\mathbf{P}_{ref})^T$ ensures that this row is linearly independent from the rest at the cost of moving $\hat{\mathbf{l}}_\infty$.

C. Autocalibration

The goal of autocalibration is to find the projective transformation $(\mathbf{H}^*)^{-1}$ that puts the reconstruction in a metric space. This is equivalent to finding the plane at infinity $\hat{\mathbf{l}}_\infty = [l_1, l_2, l_3, 1]$ as well as the intrinsics for all cameras, $\{\tilde{\mathbf{K}}_m\}$, with

$$\tilde{\mathbf{K}} = \begin{bmatrix} f_u & s & p_u \\ 0 & f_y & p_v \\ 0 & 0 & 1 \end{bmatrix}$$

Finding the plane at infinity in an arbitrary projective reconstruction is a tricky problem, and many approaches exist based on the Kruppa equations [6], the absolute quadric constraint [7], [8], and nonlinear optimization of $\tilde{\mathbf{K}}$ -based cost functions [9]–[11].

No matter the method, constraints for solving the autocalibration problem derive from prior knowledge that $\tilde{\mathbf{K}}$ represents a physical camera. This requirement for modern CCD cameras can be summarized in three prior conditions:

| | |
|---------------------------------|-----------------|
| square pixels (zero skew) | $s = 0$ |
| unit pixel aspect ratio | $f_u = f_v$ |
| principal point at image center | $p_u = p_v = 0$ |

(6)

Focal length has a much greater impact on projective distortion than the other parameters of $\tilde{\mathbf{K}}$, so an appropriate parameterization of $(\mathbf{H}^*)^{-1}$ is:

$$(\mathbf{H}^*)^{-1} = \begin{bmatrix} \tilde{\mathbf{K}}_{ref}^{-1} & \mathbf{0} \\ \hat{\mathbf{l}}_\infty \end{bmatrix} = \begin{bmatrix} 1/f & 0 & 0 & 0 \\ 0 & 1/f & 0 & 0 \\ 0 & 0 & 1 & 0 \\ l_1 & l_2 & l_3 & 1 \end{bmatrix} \quad (7)$$

This parameterization assumes one of the cameras has been canonified by (4), and that $\tilde{\mathbf{K}}$ for that camera follows (6) exactly.

In metric space, camera matrices can be decomposed into intrinsics and extrinsics using RQ decomposition:

$$\tilde{\mathbf{P}} = \tilde{\mathbf{K}}(\tilde{\mathbf{R}} \mid -\tilde{\mathbf{R}}\tilde{\mathbf{c}}) \quad (8)$$

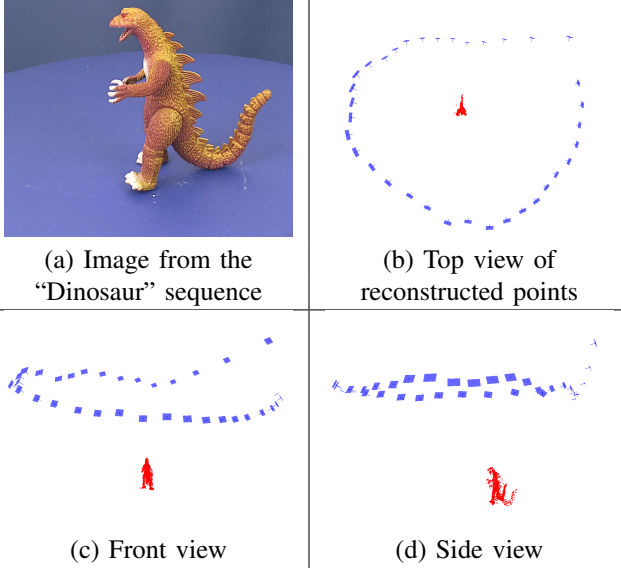


Figure 2. The “Dinosaur” sequence, reconstructed using Method (III)

Require: $\{\mathbf{P}_m\}$ normalized by (2), with $\mathbf{P}_1 \leftarrow [\mathbf{I}|\mathbf{0}]$ by (4)
Require: $(\mathbf{H}^*)^{-1}$ in (7)
 fitness $\leftarrow 0$
for $i = 0$ to m **do**
 $\mathbf{P}_m^* \leftarrow \mathbf{P}_m(\mathbf{H}^*)^{-1}$
 obtain \mathbf{K}_m^* from (8)
 fitness $+= \mathcal{C}(\{\mathbf{K}_m^*\})$ for some fitness measure
end for

Figure 3. The procedure for evaluating a proposed metric upgrade transform \mathbf{H}^* based on the intrinsics of the resulting cameras. This generalized cost function is used by all three methods with varying fitness measures $\mathcal{C}(\{\mathbf{K}_m\})$.

III. METHODOLOGY

Finding the plane at infinity in an arbitrary projective reconstruction can be formulated as a nonlinear optimization problem, and the choice of cost function is of utmost importance. Given the four parameters of $(\mathbf{H}^*)^{-1}$, (7), and a prior for $\tilde{\mathbf{K}}$, (6), the most direct way to evaluate the fitness of the resulting reconstruction is given in Figure 3.

This procedure can be used as a cost function for optimization using local perturbation methods, however the transform-decompose process is highly nonlinear. No matter the fitness measure, optimization is unlikely to converge to the correct solution unless the initial guess is sufficiently close. Three methods for autocalibration are compared in this paper, each following this general process but using different measures of fitness and initialization strategies. Optimization is performed using the Levenberg-Marquadt method [12]. All results were produced using our own implementations; code for the methods introduced by other researchers was requested, but unavailable at the time of submission.

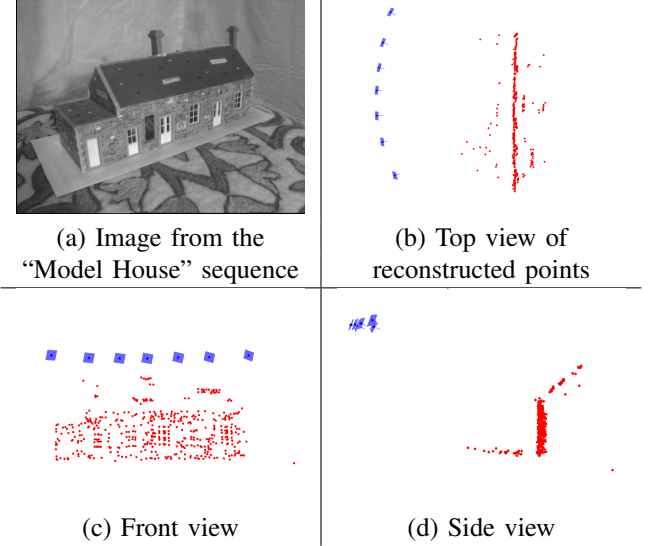


Figure 4. The “Model House” sequence, reconstructed using Method (III)

A. Method (I)

Method (I) [13] attempts to solve the initialization problem by preconditioning the projective reconstruction to ensure that it is quasi-affine (QUARC) with respect to the camera centers. This means that $\hat{\mathbf{l}}_\infty$ does not lie between any two cameras after conditioning. The decomposition of (8) is applicable only to finite cameras, and should local perturbation attempt to move $\hat{\mathbf{l}}_\infty$ across a camera center any cost function will behave erratically as the decomposition fails. The method defines a twist test for two cameras to determine if $\hat{\mathbf{l}}_\infty$ lies between them, and uses this to define a signature for \mathbf{l}_∞ : $\zeta(\mathbf{l}_\infty, \mathbf{P}_{1\dots m})$. The signature is a vector of length m whose entries are either 1 or -1 depending on the results of the twist test. To obtain a QUARC, it is sufficient to map to infinity a plane \mathbf{l}^* that has the same signature as $\hat{\mathbf{l}}_\infty$, and is as far as possible away from all camera centres. This can be expressed as a linear program:

$$\begin{aligned} \max_{\delta} \quad & \mathbf{l}^* \frac{\mathbf{c}(\mathbf{P}_m)}{|\mathbf{c}(\mathbf{P}_m)|} \zeta(\mathbf{l}_\infty, \mathbf{P}_{1\dots m})_m - \delta \geq 0 \quad \forall m \\ \text{s.t.} \quad & -1 \leq \mathbf{l}_i^* \leq 1, \quad i = 1\dots 4 \end{aligned}$$

where $\zeta(\mathbf{l}_\infty, \mathbf{P}_{1\dots m})_m$ is the m th element of the signature, and the bounds on \mathbf{l}_i^* guarantee a unique solution. Transforming the projective reconstruction to QUARC is then accomplished by mapping \mathbf{l}^* to infinity using (3), with

$$\mathbf{H}_{\text{QUARC}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{l}^* & 1 \end{bmatrix} \quad (9)$$

A theoretical framework has been laid showing that transforming the projective space to QUARC is equivalent to placing the naive initialization of (7), $[f, l_1, l_2, l_3] = [1, 0, 0, 0]$ in the correct basin of the cost function, using the fitness measure

$$\mathcal{C}_I(\{\mathbf{K}_m\}) = \sum_m \frac{(f_u - f_v)^2 + s^2 + p_u^2 + p_v^2}{(f_u + f_v)^2} \quad (10)$$

```

 $\mathbf{P}_\alpha \leftarrow \mathbf{P}_1 = [\mathbf{I}|0]$ 
for  $f = 0.3$  to  $3.0$  with logarithmic spacing do
  for  $i = 2$  to  $m$  do
     $\mathbf{P}_\beta \leftarrow \mathbf{P}_i$ 
    form  $\tilde{\mathbf{K}}_\alpha$  and  $\tilde{\mathbf{K}}_\beta$  from  $f$ 
    solve for  $\mathbf{l}_\infty^*$ 
     $\mathbf{P}_m^* \leftarrow \mathbf{P}_m(\mathbf{H}^*)^{-1}$ 
    obtain  $\mathbf{K}_m^*$  from (8)
    fitness =  $\mathcal{C}_{II}(\{\mathbf{K}_m\})$ 
  end for
end for

```

Figure 6. The algorithm used by Method (I) and (III) to perform an exhaustive search for the plane at infinity over reasonable values of focal length.

The penalization for deviation from (6) is evident. Division by the factor $(f_u + f_v)^2$, which is proportional to focal length, has a normalizing effect; a given deviation of aspect ratio or principle point is more plausible at long focal lengths than short ones. It also penalizes the cost function should f approach 0, where the decomposition of \mathbf{P} is undefined. Method (I) can be summarized as QUARC preconditioning followed by nonlinear optimization of focal length and \mathbf{l}_∞ location. For further details, the reader is directed to [13].

B. Method (II)

Method (II) [10] uses the fitness measure:

$$\mathcal{C}_{II}(\{\mathbf{K}_m\}) = \sum_m [\omega_s |s| + \omega_{ar} |f_u - f_v| + \omega_p (|p_u| + |p_v|)] \quad (11)$$

where $\omega_s = 20$, $\omega_{ar} = 2$, and $\omega_p = 1$ are appropriate weights for the terms based on prior expectations - modern CCDs have very nearly zero skew, while unit aspect ratio is less certain and principal point location much less so [14]. Initialization of $(\mathbf{H}^*)^{-1}$ is accomplished using an exhaustive search over reasonable focal lengths. A closed-form solution is given for the location of the plane at infinity given any two cameras $\mathbf{P}_\alpha = [\mathbf{I}|0]$, $\mathbf{P}_\beta = [\mathbf{Q}_\beta|\mathbf{q}_\beta]$ and their intrinsics, $\tilde{\mathbf{K}}_\alpha$, $\tilde{\mathbf{K}}_\beta$. Given that the cameras are normalized by (2), focal length will be on the order of 1 - less for wide-angle lenses and more for telephoto lenses. For most consumer cameras, a sample space of $f = [0.3 \dots 3]$ is reasonable. Making use of (6), all parameters of $\tilde{\mathbf{K}}$ except focal length can be safely ignored during the initialization phase. Method (II) samples the space of focal lengths using the algorithm in Figure III-B. The values of $[f, l_1, l_2, l_3]$ that resulted in the best fitness are used to initialize nonlinear optimization.

C. Method (III)

Method (III) is a combination of the initialization strategy of Method (II) with the fitness measure from Method (I). The




Figure 7. Value of $\mathcal{C}_{III}(\{\mathbf{K}_m\})$ plotted against its parameters, in the solution neighbourhood for the “Dinosaur” sequence. The white circle shows the initialization found by exhaustive search, and the red cross shows the optimized result.

hybrid fitness measure is

$$\mathcal{C}_{III}(\{\mathbf{K}_m\}) = \sum_m \frac{\omega_{sk} s^2 + \omega_{ar} (f_u - f_v)^2 + \omega_{p_u} p_u^2 + \omega_{p_v} p_v^2}{(f_u + f_v)^2} \quad (12)$$

The addition of the weights from (11) to the general form of (10) is intended to increase the use of prior knowledge while maintaining the theoretical rigor of (10) and its resistance to focal length collapse.

IV. RESULTS

The three methods have been compared for two image sequences that have appeared elsewhere in the 3D reconstruction literature [15]: the “Dinosaur” sequence, illustrated in Figure 2, and the “Model House” sequence, illustrated in Figure 4. Both sequences were taken using a stationary camera and the subject on a turntable, with approximately 10° spacing between images. Method (III) clearly outperforms the other two, as shown in Figures 5 and 8, producing results with very little projective distortion that would serve well as an initialization to metric bundle adjustment. To understand the failure of the other two methods it is helpful to look at the behaviour of the cost function in the solution neighbourhood, which have been plotted in Figures 7, 9 and 10. In these plots, it may sometimes appear that the optimized result is worse than the initialization. This is because for each pane two parameters are held constant at the initialization value - the optimized solution is plotted in the initialization space, not its own. Plots of the solution space, omitted here, show that the optimization is indeed converging to a local minimum.

As shown in 7, the optimization of Method (III)’s cost function is initialized within an unambiguous basin. The only other basin is in the $+l_2$ direction, and has a higher cost function value. The optimization step is fairly large in

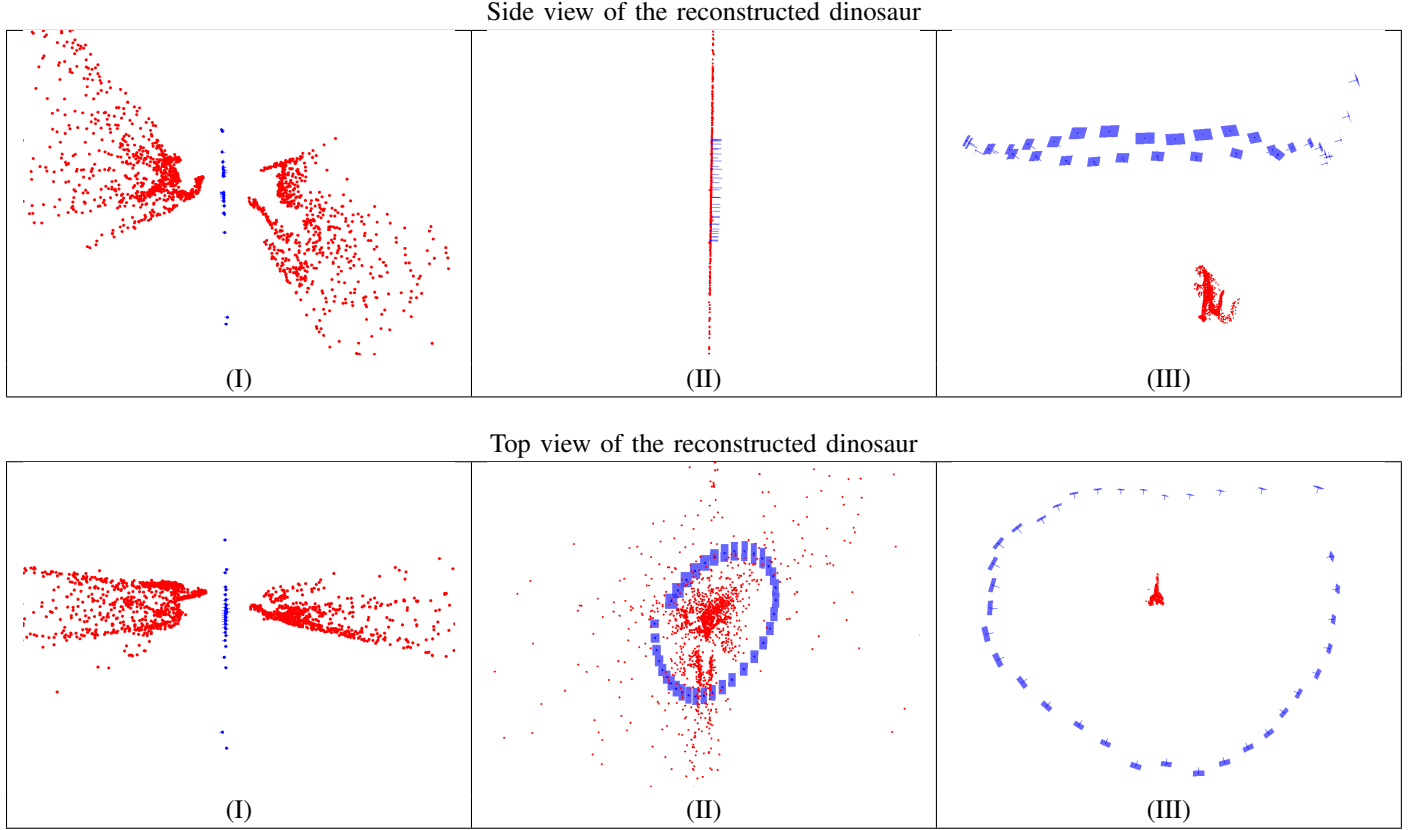


Figure 5. A comparison of the three metric reconstruction methods for the “Dinosaur” sequence. The plane at infinity still intersects the reconstruction after applying Method (I). The extreme proximity of the cameras to the scene in the side view of Method (II) indicates focal length collapse. The proposed Method (I) produces high-quality results with virtually no projective distortion.




Figure 9. Value of $C_I(\{\mathbf{K}_m\})$ plotted against its parameters, in the solution neighbourhood for the “Dinosaur” sequence. The white circle shows the naive initialization in the QUARC frame, and the red cross shows the optimized result. Note the other prominent basin in the cost function in the $-l_2$ direction. This contains the true solution for l_∞ .

the $+f$ direction due to the true solution lying above of the $f = [0.3...3]$. Although the convergence from $f = 3$ is acceptable, an extended sample space of $f = [0.3...10]$ will be used in future implementations. The mild increase




Figure 10. Value of $C_{II}(\{\mathbf{K}_m\})$ plotted against its parameters, in the solution neighbourhood for the “Dinosaur” sequence. The white circle shows initialization found by exhaustive search, and the red cross shows the optimized result. Note the large values compared to the other two methods. This basin is a numerical artifact generated by focal length collapse.

in computational cost during the exhaustive search will be more than offset by faster convergence during the nonlinear optimization.

Method (I) fails due to being initialized in the wrong basin,

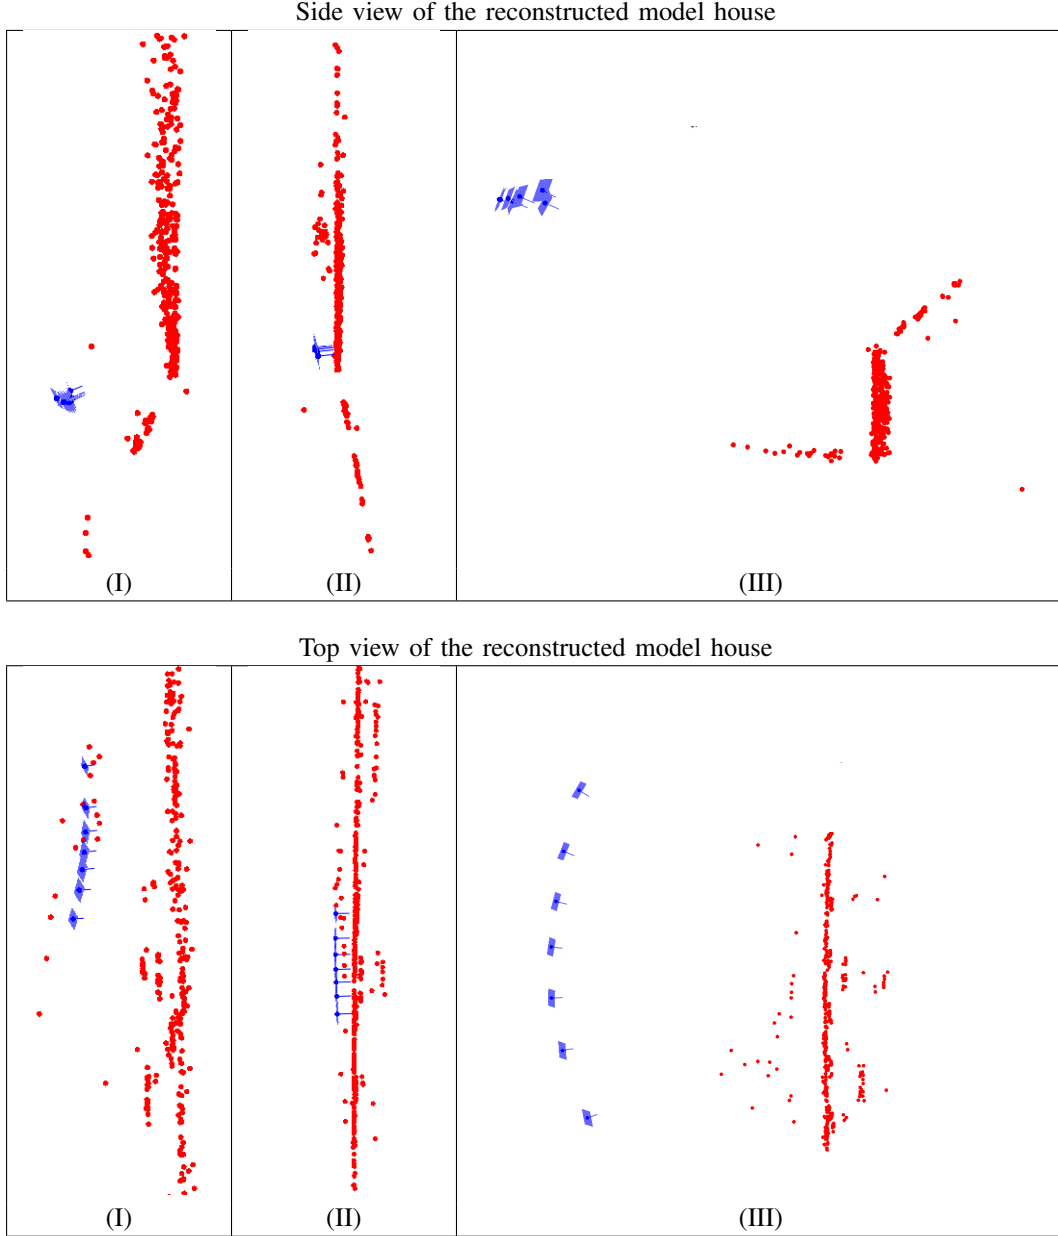


Figure 8. A comparison of the three metric reconstruction methods for the “Model House” sequence. Methods (II) and (III) exhibit severe projective distortion, but the separate planes of the floor, wall, and roof of the house can be identified with some difficulty. By contrast, the proposed Method (I) produces a result with only slight projective distortion, illustrated by the slightly less than 90° angle between the floor and the wall.

as shown in Figure 9. There is another strong basin in the $-l_2$ direction, and when the optimization is initialized with $[f, l_1, l_2, l_3] = [1, 0, -2, 0]$ instead of $[1, 0, 0, 0]$, the result is identical to that of Method (I). Very convincing results were reported for this method using many image sequences [13], which we have been unable to reproduce. It is possible that our QUARC preconditioning is not being carried out in precisely the same way, or that some detail of normalization has been missed. Even if this is the case, the failure to implement it after careful review of the

literature and consultation with the author mark QUARC preconditioning as a tricky business, best avoided. Method (II) fails due to focal length collapse. The exhaustive search picks $f = 0.3$, which has a cost similar to but slightly lower than those in the neighbourhood of $f = 3$. The cost function makes no distinction between these cases, though a physical camera with a short focal length is expected to have much smaller deviations from unit aspect ratio and principal point location than one with a long focal length. Because the cost function has no mechanism to penalize extremely

low focal lengths, the optimization takes it into the $f \approx 0$ region, where the decomposition of \mathbf{P} is ill-defined and cost function values cannot be trusted.

V. CONCLUSION

The research of this paper has been motivated by a lack of robust convergence for published autocalibration methods. In a variety of experimental tests, convergence was highly initialization dependent, even for noise-free synthetic data. This paper presents a substantial effort to deliver reliable performance on synthetic and real data sets, a novel autocalibration method combining the best aspects of two recent state-of-the-art methods. Experimental results show that a far more robust autocalibration is achieved.

ACKNOWLEDGMENTS

The authors would like to thank Bill Triggs, Riccardo Gherardi, and David Nister for their valuable guidance on the autocalibration process, as well as Manmohan Chandraker for providing a copy of David Nister's projective reconstruction dataset. This research has been sponsored by 2G Robotics Inc., the Canadian Institute for Photonic Innovations (CIPI), the Natural Sciences and Engineering Research Council (NSERC) of Canada and the Ontario Centres of Excellence.

REFERENCES

- [1] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge Univ Press, 2003.
- [2] F. Schaffalitzky, A. Zisserman, R. Hartley, and P. Torr, "A six point solution for structure and motion," *ECCV 2000*, pp. 632–648, 2000.
- [3] K. Mierle. (2008) Open source 3D reconstruction. Master's Thesis. [Online]. Available: https://tspace.library.utoronto.ca/bitstream/1807/10437/1/Mierle_Keir_B_200801_MASc_thesis.pdf
- [4] R. I. Hartley, "In defence of the 8-point algorithm," in *ICCV 1995*, 1995, pp. 1064–1070.
- [5] L. Quan, "Invariants of six points and projective reconstruction from three uncalibrated images," *PAMI*, vol. 17, no. 1, pp. 34–46, 2002.
- [6] O. Faugeras, Q. Luong, and S. Maybank, "Camera self-calibration: Theory and experiments," in *ECCV 1992*. Springer, 1992, pp. 321–334.
- [7] R. Hartley, "Euclidean reconstruction from uncalibrated views," *Applications of invariance in computer vision*, pp. 235–256, 1994.
- [8] M. Pollefeys, R. Koch, and L. V. Gool, "Self-calibration and metric reconstruction in spite of varying and unknown intrinsic camera parameters," *IJCV*, vol. 32, no. 1, pp. 7–25, 1999.
- [9] M. Chandraker, S. Agarwal, F. Kahl, D. Nister, and D. Kriegman, "Autocalibration via rank-constrained estimation of the absolute quadric," in *CVPR 2007*, 2007, pp. 1–8.
- [10] R. Gherardi and A. Fusiello, "Practical autocalibration," *ECCV 2010*, pp. 790–801, 2010.
- [11] D. Nister, "Reconstruction from uncalibrated sequences with a hierarchy of trifocal tensors," *Computer Vision-ECCV 2000*, pp. 649–663, 2000.
- [12] M. A. Lourakis and A. Argyros, "SBA: A Software Package for Generic Sparse Bundle Adjustment," *ACM Trans. Math. Software*, vol. 36, no. 1, pp. 1–30, 2009.
- [13] D. Nister, "Untwisting a projective reconstruction," *IJCV*, vol. 60, no. 2, pp. 165–183, 2004.
- [14] M. Pollefeys, F. Verbiest, and L. Van Gool, "Surviving dominant planes in uncalibrated structure and motion recovery," *ECCV 2002*, pp. 613–614, 2002.
- [15] A. Fitzgibbon and A. Zisserman, "Automatic camera recovery for closed or open image sequences," *ECCV 1998*, pp. 311–326, 1998.