Locally-weighted Homographies for Calibration of Imaging Systems

Pradeep Ranganathan and Edwin Olson

***Abstract*— A homography is traditionally formulated as a linear transformation and is used in multiple-view geometry as a linear map between projective planes (or images). Analogous to the use of homography-based techniques to calibrate a pin-hole camera, non-linear homographies extend the pin-hole camera model to deal with non-linearities such as lens distortion.**

In this work, we propose a novel *non-parametric* non-linear homography technique. Unlike a parametric non-linear mapping that can have inherent biases, this technique automatically adjusts model complexity to account for non-linearities in observed data. With this technique, we demonstrate non-parametric estimation of lens distortion from a single calibration image.

We evaluate this technique on real-world lenses and show that this technique can improve the stability of camera-calibration. Furthermore, the non-parametric nature of our technique allows rectification of arbitrary sources of lens distortion.

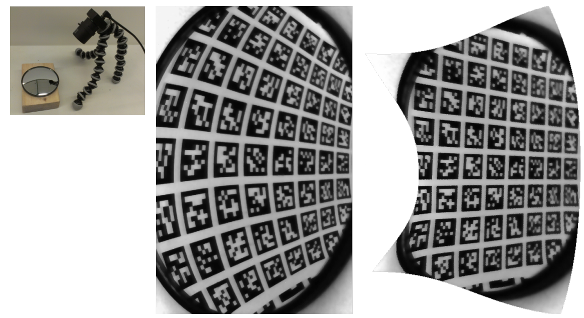

Fig. 1: Original distorted image and undistorted image from an ad hoc setup involving a camera and an automotive blind-spot mirror. The proposed method can non-parametrically estimate and correct arbitrary distortion from just a single image of a planar target.

I. Introduction

The term homography means *similar drawing*. Mathematically, it is a matrix that linearly transforms a point from one projective plane onto a point in another projective plane. It is an invaluable tool in multiple-view geometry for reasoning about multiple views of the same object.

Zhang's landmark paper [1] makes use of the homography between a planar calibration target and the camera image plane to calibrate the intrinsics of a camera. This method and its extensions represent some of the most popular methods for camera calibration.

Practical cameras do not conform to the linear transformation expected of an ideal pin-hole camera; the most obvious source of non-linearity being distortion caused by the lens. Thus, the linear projective camera must be augmented with a model of lens distortion in order to model real-world cameras.

Recent works by Claus and Fitzgibbon [2], Barreto et. al. [3] and Gasparini et. al. [4], seek to model non-linearity in cameras using a *lifted homography*. This formulation projects the input coordinates into a higher-dimensional feature space (*lifting*) and then constructs a homography in this higher-dimensional space, thus producing a non-linear mapping. By construction, lifted homographies are parametric with respect to the size of the higher-dimensional space, and they are geometrically unintuitive. In this work, we propose a novel alternative non-linear homography technique that is also non-parametric. The technique retains geometric intuition because its mapping action can be interpreted in terms of a locally-linear homography. The contributions of this work include:

1) A novel non-parametric non-linear homography formulation that can map lines on the source plane to arbitrary smooth curves on the target plane,
2) An application of this technique to build a model of lens distortion; we compare the use of both parametric polynomial models and non-parametric Gaussian process models of lens distortion,
3) Improved stability of camera calibration using this independent non-parametric estimate of lens distortion, and
4) An evaluation of the performance of both the non-linear homography technique and the resulting lens distortion model on real-world lenses.

This formulation of a non-parametric non-linear homography is a novel idea to the best of the authors' knowledge. Its ability is not limited by a choice of parametric model for lens distortion, and it has the capability to model arbitrary sources of distortion (e.g. Fig. 1). More interestingly, it produces an independent estimate of the distortion without relying on any prior knowledge of the intrinsics or extrinsics of the camera used to acquire the image.

In essence, the *straight lines are straight* rectification procedure as described by Devernay and Faugeras [5] serves a similar purpose to our method. However, our method operates in terms of homographies; therefore, it implicitly encodes projective constraints.

In the following section, we present a brief overview of necessary background. Our method and evaluation follows in section III. We end this work with a discussion of its contributions and future work in section IV.

The authors are with the Computer Science and Engineering Department, University of Michigan, Ann Arbor, MI 48104, USA. {rpradeep,ebolson}@umich.edu

II. BACKGROUND

Our formulation of the non-linear homography estimation builds on the methods of classic planar homography estimation and uses Gaussian process regression to provide predictions of distortion from observed estimates. We provide brief reviews of both these topics in this section.

A. Planar homography estimation

Mathematically, a planar perspective homography $\mathbf{H}$ defines an invertible linear transformation between two perspective planes. It maps a point $\mathbf{p}$ in one plane onto a point $\mathbf{q}$ in another plane such that $\mathbf{p} \sim \mathbf{Hq}$. The point $\mathbf{p}$ is similar and not equal to $\mathbf{Hq}$ because of the universal scale ambiguity in perspective projection. Further, it maps a line on the source plane to a line on the target plane. When a planar target is viewed by a perspective camera, points on the planar target are related to their image on the camera plane by a homography $\mathbf{H}$.

In camera-based view geometry, the homography acquires a special interpretation $\mathbf{H} = \mathbf{KE}$, where $\mathbf{K}$ is the camera perspective (intrinsics) matrix and $\mathbf{E}$ is the Euclidean transformation (extrinsics) matrix that describes the pose of the camera when viewing the target. For an extensive treatment of the concept of a homography, we refer the reader to [6]. In the following paragraph, we briefly explain the process (adapted from [7]) of estimating a homography that relates two sets of corresponding points.

Consider a pair of corresponding points $\mathbf{p} = (x_1, y_1, z_1)^T$ and $\mathbf{u} = (x_2, y_2, z_2)^T$ related by a homography $\mathbf{H}$:

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} \sim \begin{bmatrix} h_{11} & h_{12} & h_{13} & h_{14} \\ h_{21} & h_{22} & h_{23} & h_{24} \\ h_{31} & h_{32} & h_{33} & h_{34} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ 1 \end{bmatrix}$$

In normalized coordinates: $(x_2^{'}, y_2^{'})^T = (x_2/z_2, y_2/z_2)^T$, i.e.,

$$x_2^{'} = \frac{h_{11}x_1 + h_{12}y_1 + h_{13}z_1 + h_{14}}{h_{31}x_1 + h_{32}y_1 + h_{33}z_1 + h_{34}}$$

$$y_2^{'} = \frac{h_{21}x_1 + h_{22}y_1 + h_{23}z_1 + h_{24}}{h_{31}x_1 + h_{32}y_1 + h_{33}z_1 + h_{34}}$$

Thus, each correspondence $\mathbf{p} \rightleftarrows \mathbf{u}$, results in two linear equations in the unknowns $\mathbf{h} = (h_{11}, h_{12}, \ldots, h_{34})^T$. With multiple correspondences, we collect multiple pairs of such linear constraints to obtain a coefficient matrix $\mathbf{A}$. We then obtain a least-squares solution for $\mathbf{h}$ by solving

$$(\mathbf{A}^T\mathbf{A})\mathbf{h} = \mathbf{0} \quad (1)$$

The solution for $\mathbf{h}$ is obtained as the eigenvector corresponding to the smallest singular value of $\mathbf{A}^T\mathbf{A}$.

B. Gaussian process regression

For Gaussian process (GP) regression on a set of observations t at a set of input locations x, we assume that the data was obtained from a zero-mean, stationary GP with covariance function $k(x_n, x_m)$. Let β denote the precision of the observed target values t. We then define $\mathbf{C}$ as the covariance matrix with elements

$$\mathbf{C}(x_n, x_m) = k(x_n, x_m) + \beta^{-1}\delta_{nm}$$

When we require a prediction for a new input x_{N+1}, we first construct the covariance matrix C_{N+1} and partition it as follows:

$$\mathbf{C}_{N+1} = \begin{bmatrix} \mathbf{C}_N & \mathbf{k} \\ \mathbf{k}^T & c \end{bmatrix}$$

Then the mean and variance of the predicted value at x_{N+1} are given by

$$m(x_{N+1}) = \mathbf{k}^T\mathbf{C}_N^{-1}t \quad (2)$$

$$\sigma^2(x_{N+1}) = c - \mathbf{k}^T\mathbf{C}_N^{-1}\mathbf{k} \quad (3)$$

We refer the reader to [8], [9] for a detailed mathematical derivation of the above results. Instead, we provide an intuitive summary of GP regression in the following paragraphs.

GP regression can be viewed as kernelization of Bayesian linear regression to accommodate non-linearity (see [9]). The covariance function plays the role of the kernel function here; the terms are often used interchangeably in the context of GP regression. Intuitively, GP regression predicts new outputs using a locally weighted sum of the observed targets, t, where the local weighting is specified indirectly through the covariance function.

A common choice for the covariance function is the squared-exponential (SE) covariance function (or the Gaussian kernel function).

$$k(x_m, x_n) = \theta_1^2 \exp\left\{-0.5\frac{(x_m - x_n)^2}{\theta_0^2}\right\} + \delta_{mn}\beta^{-1} \quad (4)$$

The parameter θ_0 controls the smoothness of the resulting GP. The parameter θ_1 controls the scale of the similarity and also affects the extrapolation power of the GP.

The parameters θ_0, θ_1 and β are considered hyper-parameters of the GP. For effective GP regression, one must estimate optimal values for these hyper-parameters as described in [8].

III. METHOD

In this section, we present our formulation of a non-linear homography and use it to obtain an estimate of lens distortion. We then learn predictive (polynomial and Gaussian process) models from these distortion estimates and then show how these models can aid and improve the process of camera calibration via non-linear optimization.

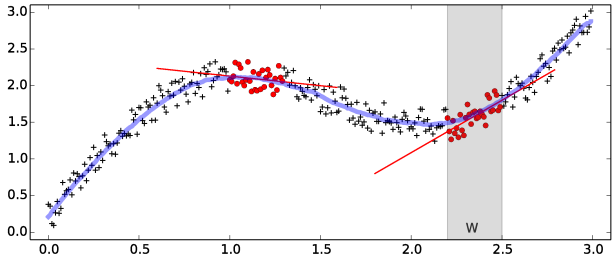

Fig. 2: Local regression on non-linear data.

A. Locally-weighted homography estimation

Consider the plot in Fig. 2 with data (shown in black) exhibiting a non-linear trend. This is a typical candidate for non-linear regression. However, in a sufficiently small window W (shaded gray), the data is mostly linear and a line is, locally, a good fit as shown by the red line through the dots.

Now, let us consider sliding the window W, over the entire domain of data and computing local regression estimates at each location of the window. To enforce smoothness between estimates in adjacent window locations, we substitute the use of a discrete neighborhood with a fuzzy neighborhood: we weight points using a weighting scheme that weights closer points higher than points further away. This regression technique is termed *locally-weighted linear regression*[1] [10].

This regression estimate is *dynamic* in the sense that it changes with the choice of window location and the weighting function (for example the other set of red points to the left produce a different locally-valid, linear regressor). When we systematically perform multiple local regressions over the domain of the data, we obtain a non-linear locally-weighted regressor (blue curve in the plot).

Homography estimation is also a regression problem, and this idea of locally-weighted regression can be extended to homography estimation by weighting the estimate derived in (1). This gives us:

$$(\mathbf{A}^T\mathbf{W}\mathbf{A})\mathbf{h} = \mathbf{0},$$

where the weight matrix $\mathbf{W}$ is a diagonal matrix, with elements $k(\mathbf{q}, \mathbf{x}_i)$ on its diagonal; $\mathbf{q}$ is the location of the current estimation window and $\mathbf{x}_i$ are the points on the source plane used to estimate the homography. $k(.)$ is the kernel or weighting function:

$$k(\mathbf{q}, \mathbf{x}_i) = \exp\left\{-\frac{\mathbf{q}^T.\mathbf{x}_i}{2\tau^2}\right\} \quad (5)$$

The resulting non-linear homography is directly interpretable as a continuous mosaic of smoothly varying planar homographies, which relates each point on the source to a unique corresponding point on the target. Analogous to how the planar homography maps source lines onto target lines, the non-linear homography maps source lines onto smooth curves on the target plane, as shown in Fig. 3.

[1] A reader with a background in signal processing or time-series modeling will notice similarities between this technique, convolution and moving averages.

Furthermore, this non-linear homography is also *non-parametric* because it can model non-linearities of any complexity using a single hyper-parameter τ, the weighting bandwidth, in (5). Optimal estimates of τ are obtained by optimizing a cross-validation score: one (or more) points are *held-out* from the set of known correspondences, and τ is chosen as the value that minimizes mapping error of the held-out set of correspondences.

B. Lens distortion estimation

The distortion caused by lenses is predominantly radial (see Fig. 3 for an example). However, manufacturing errors or shock can introduce slight tangential distortion components. Because of its radial nature, it is also generally true that lens distortion is minimal at the center of the image. Lens distortion is modeled as a distortion function $f_d(.)$ that acts on a point on the image plane and displaces it to a new point.

Our distortion estimation technique consists of the following steps. Note that a locally-weighted homography is specified by the mapping direction (source, target) and point at which it is estimated:

1) Obtain a set of pairwise correspondences between image points $\mathbf{x}_i$ and world points $\mathbf{x}_w$ such that $\mathbf{x}_w^k \rightleftarrows \mathbf{x}_i^k$.
2) Estimate $\mathbf{H}_{i\to w}$, the locally-weighted homography from image to planar target, at the center of the image $\mathbf{c}$. Use this to find $\mathbf{p}$, the corresponding point on the planar target, where $\mathbf{p} = \mathbf{H}_{i\to w}\mathbf{c}$.
3) Now, find a locally-weighted homography in the other direction $\mathbf{H}_{w\to i}$, from planar target to image, at the point $\mathbf{p}$. By construction, $\mathbf{p}$ and $\mathbf{c}$ are related by $\mathbf{H}_{i\to w}$, so $\mathbf{H}_{w\to i} = (\mathbf{H}_{i\to w})^{-1}$. This mapping is not affected by lens distortion since we have assumed that there is minimal distortion at the center of the image.
4) Transform all other $\mathbf{x}_w^k$ by $\mathbf{H}_{w\to i}$, to obtain $\mathbf{y}_i^k$. The estimate of the distortion at $\mathbf{x}_i^k$ is then $\mathbf{d}_k = (\mathbf{y}_i^k - \mathbf{x}_i^k)$

In other words, we estimate the (world-to-image) homography at the center of the image, $\mathbf{H}_{w\to i}$, and project all the world points through this homography. The distortion is then the displacement between what was observed and what is predicted by $\mathbf{H}_{w\to i}$. Displacement of observed points from predicted image points gives us the distortion estimate while its inverse, the displacement of predicted image points from observed points, gives us the undistortion function:

$$\begin{aligned} \mathbf{d}_k &= \mathbf{y}_i^k - \mathbf{x}_i^k \\ \mathbf{u}_k &= \mathbf{x}_i^k - \mathbf{y}_i^k \end{aligned}$$

For reasons that will be apparent in the following sections, we chose to directly model the undistortion function from its estimates $\mathbf{u}_k$.

To continue with our goal of not making any radial or tangential assumptions, we model the undistortion as a direct

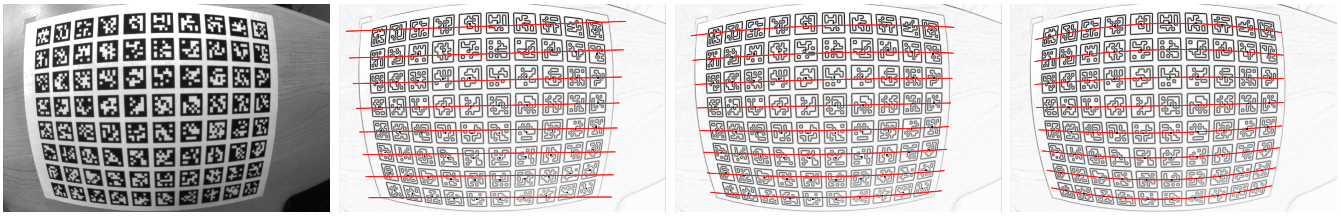

Fig. 3: Locally-weighted homography estimation. Left-most image is the camera image of the planar target (see [11], [12] for information on the fiducial-based planar target). The following figures show the mapping of lines through row-centers on the target onto the image plane, for weighting bandwidth $\tau = 1{,}0.005$ and 0.0005, respectively. Note how the line mappings are progressively more curved. At $\tau = 0.0005$, the mapping is accurate to within 0.1 pixels.

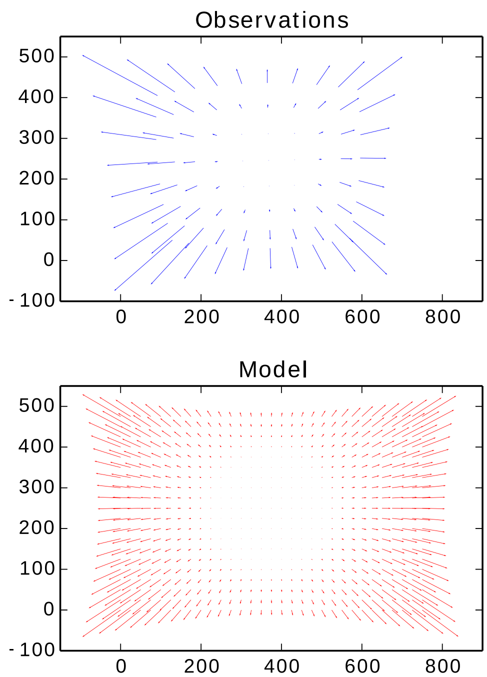

Fig. 4: Distortion model using a locally-weighted homography at the center of the image. The quiver plot at the top shows the distortion observations for the image in Fig. 3. At the bottom is a quiver plot of the GP undistortion model inferred from these observations.

2-D function of pixel coordinates. A reasonable choice is to use a 2-D polynomial mapping to model the undistortion. However, we observed that 2-D polynomial regression exhibited over-fitting on our datasets, with polynomial orders greater than five causing an increase in reprojection error.

One could solve the problem of over-fitting by regularization. Instead, we chose to model the undistortion using a 2-D GP. This implicitly takes care of regularization with the additional advantage of making the undistortion function non-parametric.

In our implementation[2], the 2-D GP consists of two independent GPs, one modeling the undistortion u_x in the x-direction and another modeling u_y in the y-direction. The distortion components can be modeled separately because

[2] available on our website: `http://april.eecs.umich.edu`.

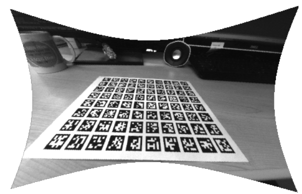

(a) Tamron 2.2 mm lens

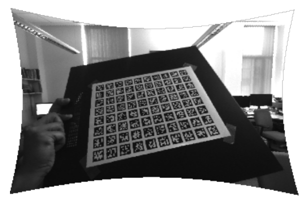

(b) Tamron 2.8 mm lens

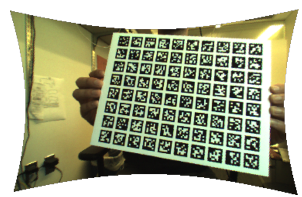

(c) Tokina 3.3 mm lens

Fig. 5: Images undistorted using a GP model estimated from a single image.

$u_x(i,j)$ is conditionally independent of $u_y(i,j)$ given the pixel coordinates (i,j). This GP that tracks two independent input variables (i,j) uses a 2D SE kernel function:

$$k(\mathbf{x}_m, \mathbf{x}_n) = \theta_1^2 \exp\left\{-\frac{1}{2}\mathbf{x}_m^T \mathbf{\Sigma}^{-1}\mathbf{x}_n\right\} + \delta_{mn}\beta^{-1}$$

which has extra hyper-parameters corresponding to the entries in the matrix $\mathbf{\Sigma}$.

As an evaluation of the local homography-based undistortion technique, in Table I, we undistort correspondences from three real-world lenses and report pixel deviation from straightness. Overall, we observe deviations of less than one pixel. However, for the set of lenses used in our evaluation, there is no significant difference in the performance of the polynomial (up to fifth order) and GP models. Fig. 5 has some examples of images undistorted using a non-parametric GP model.

C. Single image calibration

Combining the technique of section III-B with camera intrinsics estimation using orthogonal vanishing points (see [13], [14]), one can obtain useful estimates of both the camera intrinsics and lens distortion simultaneously. This can be done with just one calibration image as follows:

1) Obtain an image of the planar target, such that the target covers most of the image plane while producing

Pix Err→		Dataset 1		Dataset 2		Dataset 3		Dataset 4		Dataset 5	
		AVG	MAX	AVG	MAX	AVG	MAX	AVG	MAX	AVG	MAX
Tamron 2.2	POLY	0.09	0.37	0.11	0.54	0.09	0.39	0.05	0.17	0.09	0.41
	GP	0.09	0.36	0.12	0.51	0.08	0.44	0.04	0.16	0.09	0.38
Tamron 2.8	POLY	0.07	0.24	0.11	0.39	0.11	0.56	0.16	0.65	0.13	0.76
	GP	0.07	0.24	0.11	0.39	0.10	0.37	0.15	0.64	0.13	0.78
Tokina 3.3	POLY	0.04	0.21	0.05	0.32	0.11	0.65	0.06	0.29	0.08	0.43
	GP	0.04	0.32	0.05	0.31	0.12	0.67	0.06	0.33	0.08	0.42

TABLE I: Mean absolute deviation from straightness after undistorting test datasets using the estimated lens distortion model. Overall, we observe that the deviations (mean and max) are within 1 pixel. For the lenses used in this evaluation, there is no significant difference between the polynomial and GP undistortion models.

two vanishing points. A target that covers the image plane helps in a confident distortion estimate and at least two vanishing points are required for recovering the camera matrix $\mathbf{K}$.

2) Estimate the undistortion as described in section III-B, and undistort the image.
3) Estimate vanishing points from the undistorted image and use the technique of [13] to estimate the camera matrix. Note that with only two vanishing points, we must assume that the principal point (c_x, c_y) is at the center of the image.

This estimate of camera intrinsics and distortion from a single image can be further refined by non-linear optimization.

D. Non-linear optimization of camera intrinsics

We can use the lens distortion estimate obtained by our method to simplify classic camera calibration. This is done by undistorting the calibration images first, and then optimizing just the camera intrinsics and extrinsics. This results in an optimization that optimizes a smaller set of parameters.

However, we must note that by undistorting calibration images we are committing to a single point-estimate of the undistortion and ignoring any uncertainty. As an empirical approximation, we can integrate out the uncertainty parameters from the optimization by sampling from the GP lens distortion estimate, and then using these samples to expand the calibration image set. We list out the steps in this *augmented camera calibration* algorithm below:

1) Obtain a model of undistortion using the technique listed in section III-B.
2) Obtain a set of calibration images.
3) For each image in the calibration set, sample the GP undistortion model multiple times, and undistort the image using the obtained samples. This results in an expanded calibration set that empirically accounts for the uncertainty in the distortion estimate.
4) Perform non-linear least-squares optimization of the camera intrinsic and extrinsic parameters on the undistorted images, as in the classic calibration method.

In Fig. 7, we compare the performance of the classic and augmented calibration methods and present histograms of errors obtained on multiple testing datasets. We report both the RMSE and max pixel errors and find that the classic calibration method has significantly more outliers.

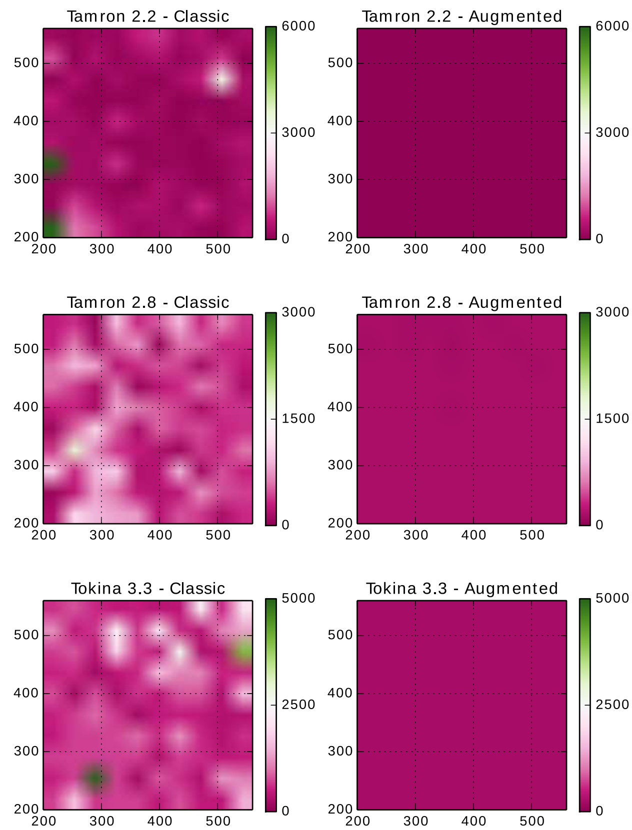

Fig. 6: Convergence of classic camera calibration vs. augmented camera calibration for three different lenses. The augmented calibration method uses the non-parametric undistortion model estimate to undistort the image before calibrating camera intrinsics and extrinsics. The x-axis and y-axis are the values of (f_x, f_y) used to initialize the optimization. (c_x, c_y) was initialized to the image center. The colormapped value at each point corresponds to the Euclidean distance of the final calibration from a nominal reference calibration. We find that the augmented camera calibration is very flat showing that it is more consistent. For reference, the range of values for the augmented calibration plots are $(2.0, 6.9)$, $(3.3, 44.8)$ and $(2.12, 12.9)$, top to bottom. This improvement in convergence suggests that the augmented method results in a more stable optimization.

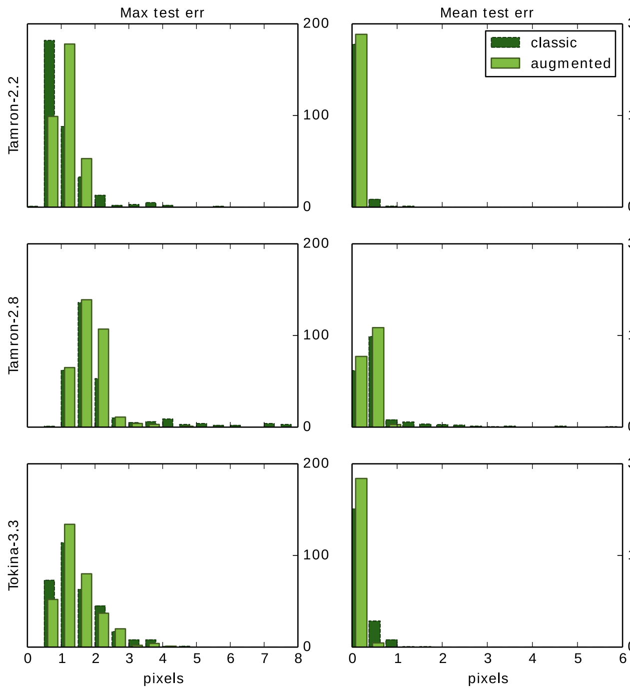

Fig. 7: Testing errors for the classic and augmented calibration methods. We observe that both the classic and augmented methods have a very similar distribution of errors except for a significant tail of outliers for the classic method. This tail of outliers suggests that the classic method is prone to over-fitting.

In Fig. 6, we compare the convergence of classic non-linear least-squares-based camera calibration with augmented camera calibration. The augmented calibration results in an optimization problem that has a smaller number of parameters,3 and hence it has more consistent convergence.

An alternative (and perhaps more principled) approach for incorporating the lens distortion estimate in the non-linear optimization is to use the GP distortion estimate as a GP prior over lens distortion. We refer the reader to [15] for an explanation of this technique. We intend to explore this approach as future work.

IV. Discussion & Conclusion

In this work, we have described a novel non-parametric non-linear homography technique; unlike a planar homography, it is capable of mapping lines from a source plane onto arbitrary smooth curves in the target plane. We then use this technique to estimate lens distortion as the observed deviation from the homography at the center of the image and non-parametrically model lens distortion using a Gaussian process.

We then use this distortion model to undistort real-world images. Classic camera calibration involves simultaneous estimation of lens distortion and camera intrinsic parameters. We show that this classic calibration technique can be augmented with our independent lens distortion estimation technique to improve the stability of camera calibration.

Our technique can be visualized as building a continuous mosaic of homographies from the source plane to the target plane. As an extension, we can decompose the homographies on the mosaic and interpret the resulting image as the result of multiple appropriately placed pinhole cameras. This lets us interpret this technique as building a non-parametric, non-linear mapping from 2-D image points to rays in 3-D.

The ability to interpret an acquired image as the result of a camera locus might have implications in estimating the caustic of a catadioptric (mirror+lens) camera system [16]. We intend to explore this direction as future work.

Acknowledgements

This work was funded by U.S. DoD Grant FA2386-11-1-4024.

References

[1] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000. [Online]. Available: http://dx.doi.org/10.1109/34.888718

[2] D. Claus and A. W. Fitzgibbon, "A rational function lens distortion model for general cameras," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2005, pp. 213–219.

[3] J. Barreto, J. Roquette, P. Sturm, and F. Fonseca, "Automatic Camera Calibration Applied to Medical Endoscopy," in *BMVC 2009 - 20th British Machine Vision Conference*, Sept. 2009. [Online]. Available: http://hal.inria.fr/inria-00524388

[4] S. Gasparini, P. Sturm, and J. Barreto, "Plane-based calibration of central catadioptric cameras," in *Computer Vision, 2009 IEEE 12th International Conference on*, Sept 2009, pp. 1195–1202.

[5] F. Devernay and O. Faugeras, "Straight lines have to be straight: Automatic calibration and removal of distortion from scenes of structured enviroments," *Mach. Vision Appl.*, vol. 13, no. 1, pp. 14–24, Aug. 2001. [Online]. Available: http://dx.doi.org/10.1007/PL00013269

[6] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, ISBN: 0521540518, 2004.

[7] D. Kriegman. Homography estimation. [Online]. Available: http://cseweb.ucsd.edu/classes/wi07/cse252a/homography_estimation/homography_estimation.pdf

[8] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[9] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[10] W. S. Cleveland, "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association*, vol. 74, pp. 829–836, 1979.

[11] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2011, pp. 3400–3407.

[12] A. Richardson, J. Strom, and E. Olson, "AprilCal: Assisted and repeatable camera calibration," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, November 2013.

[13] P. A. Beardsley and D. W. Murray, "Camera calibration using vanishing points," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 1992, pp. 43.1–43.10, doi:10.5244/C.6.43.

[14] R. Cipolla, T. Drummond, and D. Robertson, "Camera calibration from vanishing points in image of architectural scenes," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 1999, pp. 38.1–38.10, doi:10.5244/C.13.38.

[15] P. Ranganathan and E. Olson, "Gaussian process for lens distortion modeling," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2012.

[16] R. Swaminathan, M. Grossberg, and S. Nayar, "Caustics of Catadioptric Cameras," in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2, Jul 2001, pp. 2–9.