

Single-image estimation of Brown-Conrady distortion in Fringe Projection Profilometry

Osman Ali¹, Dalia Elobaid², Xiangjun Kong¹, Tibebe Yalew¹, Waiel Elmadih³, Samanta Piano¹

¹ University of Nottingham, Nottingham, United Kingdom – osman.ali@nottingham.ac.uk

² Sudanese Materials Scientists & Engineers, Khartoum, Sudan

³ Taraz Metrology Ltd, Nottingham, United Kingdom

Keywords: Lens distortion, Brown-Conrady model, Single-image estimation, Deep learning, Geometric consistency

Abstract

Accurate modelling of lens distortion is essential for reliable geometric reconstruction in computer vision and optical metrology systems. Conventional characterisation approaches estimate distortion parameters from multiple images by leveraging repeated observations of a three-dimensional (3D) target under different viewpoints. However, practical constraints may limit data acquisition to a single usable image, motivating the investigation of reduced-data alternatives. This work investigates the estimation of Brown-Conrady distortion coefficients directly from a single distorted checkerboard image. The problem is formulated as a supervised regression task, in which a convolutional neural network predicts radial and tangential distortion parameters. A synthetic dataset is generated to provide controlled variations in distortion and corresponding ground-truth labels. The results demonstrate that the model is able to extract distortion-related information from image geometry and produce physically plausible parameter estimates. Although the numerical accuracy of the predicted coefficients remains limited, the resulting geometric transformations are consistent and preserve the global distortion structure, with minor local deviations. This behaviour reflects an inherent ambiguity in the estimation problem, where multiple parameter combinations can produce similar image deformations when only a single view is available. Overall, the proposed approach shows that single-image distortion estimation can yield geometrically meaningful predictions, while highlighting the fundamental limitations associated with recovering a unique set of distortion parameters from limited geometric information.

1. Introduction

Lens distortion introduces systematic geometric deviations in captured images, particularly toward the image periphery. If not accurately modelled, these distortions propagate into projection errors, biased measurements, and degraded geometric reconstruction (Hartley and Zisserman, 2003). Accurate estimation of distortion parameters is therefore a fundamental requirement in computer vision applications such as photogrammetry and fringe projection profilometry (FPP), where metric accuracy is critical (Luhmann et al., 2023).

In FPP systems, the quality of the system characterisation directly affects the accuracy of the three-dimensional (3D) reconstruction, making robust distortion modelling a critical requirement (Geng, 2011, Zhang, 2010). Distortion estimation is typically embedded within the characterisation pipeline, where multiple images of a known characterisation target are used to estimate intrinsic, extrinsic and distortion parameters (Zhang, 2000).

The Brown-Conrady model is one of the most widely used parametric formulations for representing lens distortion (Brown, 1966, Conrady, 1919). It describes distortion using a small set of physically interpretable parameters, including radial (k_1, k_2, k_3) and tangential (p_1, p_2) coefficients. In conventional characterisation procedures, these parameters are estimated simultaneously with intrinsic and extrinsic camera parameters using multiple target images, enabling robust optimisation across diverse viewing geometries (Zhang, 2000).

However, acquiring multiple views of the characterisation target is not always feasible. Practical constraints, such as limited system access, time constraints, or the availability of only a single

usable image, motivate the investigation of reduced-data alternatives. Estimating distortion parameters from a single image becomes a meaningful but challenging problem.

This work investigates the feasibility of estimating Brown-Conrady distortion coefficients directly from a single distorted checkerboard image. Rather than producing a visually corrected image, the objective is to estimate the parameter vector

$$\mathbf{d} = [k_1, k_2, p_1, p_2, k_3]^\top,$$

where k_1, k_2, k_3 represent the radial distortion terms and p_1, p_2 represent the tangential ones. These parameters are physically interpretable and can be directly used in standard undistortion routines (e.g., OpenCV), where distorted image coordinates are mapped back to their undistorted positions using the inverse distortion model.

To address this problem, we formulate a supervised regression task in which a convolutional neural network (CNN) maps a single distorted image to its corresponding distortion parameter vector, as shown in Figure 1. The predicted coefficients can then be used within standard undistortion pipelines to demonstrate compatibility with existing photogrammetric toolchains.

Unlike prior learning-based approaches that focus on full camera characterisation from real-world images, this work considers a constrained metrology setting using structured targets. The objective is not to replace classical characterisation methods, but to evaluate whether meaningful distortion parameters can be inferred from a single image under controlled conditions.

The main contributions of this paper are:




Figure 1. Overview of the proposed single-image distortion regression framework.

- A systematic investigation of single-image distortion parameters regression model using a controlled synthetic dataset.
- A learning-based formulation that predicts interpretable Brown-Conrady coefficients compatible with standard characterisation pipelines.
- An experimental analysis highlighting the limitations and challenges of single-image distortion estimation.

2. Related Work

Lens distortion modelling has been extensively studied in computer vision. The Brown-Conrady model provides a parametric representation of radial and tangential distortion and remains widely used in practical characterisation systems (Brown, 1966, Conrady, 1919). In conventional approaches, distortion parameters are estimated jointly with intrinsic and extrinsic parameters using multiple images of a characterisation target (Zhang, 2000, Heikkila and Silvén, 1997). The use of multiple views introduces geometric redundancy, which improves estimation stability and accuracy.

Several methods estimate distortion as part of a full characterisation process using multiple images of a checkerboard (Gasparini et al., 2009, Mei and Rives, 2007), one of the most commonly used characterisation patterns (Chaudhry et al., 2025). While such approaches are generally robust, they are typically limited to controlled acquisition environments (Martins et al., 2020). Targetless methods have been proposed for less controlled scenarios (Barazzetti et al., 2012), although often at the expense of reduced accuracy.

Other approaches aim to estimate distortion from limited data. Line-based methods exploit the constraint that straight lines in the scene should remain straight after distortion correction (Devernay and Faugeras, 2001, Hartley and Kang, 2007). More general formulations estimate distortion jointly with scene geometry under minimal assumptions (Fitzgibbon, 2001). While these methods demonstrate that distortion can be inferred from a single image, they rely on specific geometric structures and accurate feature extraction.

More recently, deep learning has been applied to the regression of camera-related parameters directly from images. Convolutional neural networks have been used to estimate quantities such as camera pose and intrinsic parameters from image content (Kendall et al., 2015, Chaudhry et al., 2025). These approaches suggest that image appearance can encode sufficient information for parameter prediction, particularly when trained on large synthetic or mixed synthetic and real datasets.

In contrast to existing work, this study focuses specifically on the regression of Brown-Conrady distortion coefficients from a single structured characterisation image. The distortion vector consists of radial coefficients (k_1, k_2, k_3), which describe symmetric distortion increasing with distance from the image centre, and tangential coefficients (p_1, p_2), which model asymmetric distortions caused by lens misalignment.

Unlike general vision scenarios, the problem is investigated within the context of fringe projection profilometry (FPP), a metrology technique that relies on precise geometric characterisation for accurate 3D reconstruction. The aim is not to replace full camera characterisation, but to assess whether these interpretable parameters can be recovered from a single image under controlled conditions typical of FPP systems.

3. Methodology

The objective of this work is to estimate Brown-Conrady distortion coefficients from a single distorted checkerboard image using a deep regression model. The formulation combines a standard camera model, a parametric distortion model, and a supervised learning framework.

3.1 Distortion Model

The Brown-Conrady model expresses distortion as a combination of radial and tangential components. The radial distance is defined as

$$r^2 = x^2 + y^2. \quad (1)$$

The radial distortion is given by

$$L(r) = 1 + k_1 r^2 + k_2 r^4 + k_3 r^6, \quad (2)$$

and tangential distortion is defined as

$$\Delta x = 2p_1 xy + p_2(r^2 + 2x^2), \quad (3)$$

$$\Delta y = p_1(r^2 + 2y^2) + 2p_2 xy. \quad (4)$$

The distorted coordinates are obtained as

$$x_d = xL(r) + \Delta x, \quad (5)$$

$$y_d = yL(r) + \Delta y. \quad (6)$$

These normalised distorted coordinates are subsequently mapped to pixel coordinates through the camera intrinsic matrix. Image undistortion corresponds to applying the inverse of this transformation, typically implemented via interpolation-based remapping.

3.2 Learning Formulation

The distortion estimation problem is formulated as a supervised regression task in which a convolutional neural network (CNN) learns a mapping from an input image to a vector of distortion parameters. The model is defined as

$$\hat{\mathbf{d}} = f_{\theta}(I), \quad (7)$$

where I denotes the input distorted image, $\hat{\mathbf{d}}$ is the predicted distortion vector, and $f_{\theta}(\cdot)$ represents the nonlinear function

implemented by the neural network with learnable parameters θ . The target vector

$$\mathbf{d} = [k_1, k_2, p_1, p_2, k_3]^\top \quad (8)$$

contains the Brown-Conrady distortion coefficients associated with the input image.

Direct regression of the distortion parameters can be challenging due to the varying numerical scales and coefficient distributions. To improve stability, each parameter is transformed using a signed logarithmic mapping

$$\tilde{d}_i = \text{sign}(d_i) \log(1 + |d_i|), \quad (9)$$

where d_i and $\tilde{d}_i$ denote the original and transformed distortion parameters, respectively, which compress large magnitudes while preserving the sign and relative ordering. The network is trained to predict the transformed parameters, and the predictions are mapped back to the original space during inference using the inverse transformation.

The learning objective is defined as a combination of parameter-space and geometric consistency losses. The parameter loss is given by the mean squared error between the predicted and target transformed parameters

$$\mathcal{L}_{\text{param}} = \|\hat{\mathbf{d}} - \tilde{\mathbf{d}}\|_2^2. \quad (10)$$

where $\hat{\mathbf{d}}$ and $\tilde{\mathbf{d}}$ denote the predicted and ground-truth distortion parameters in the transformed space, respectively. To ensure that the predicted parameters correspond to physically meaningful distortions, a geometric loss is introduced based on the Brown-Conrady distortion model. A set of normalised image coordinates (x, y) is sampled over a regular grid, and both the predicted and ground-truth parameters are used to generate corresponding distorted coordinates

$$(x_{\text{pred}}, y_{\text{pred}}), \quad (x_{\text{gt}}, y_{\text{gt}}). \quad (11)$$

where $(x_{\text{pred}}, y_{\text{pred}})$ and $(x_{\text{gt}}, y_{\text{gt}})$ denote the distorted coordinates obtained using the predicted and ground-truth parameters, respectively.

The geometric loss is then defined as the discrepancy between these two mappings over all sampled points

$$\mathcal{L}_{\text{geo}} = \|x_{\text{pred}} - x_{\text{gt}}\|_2^2 + \|y_{\text{pred}} - y_{\text{gt}}\|_2^2. \quad (12)$$

The total loss combines both terms

$$\mathcal{L} = \alpha \mathcal{L}_{\text{param}} + \beta \mathcal{L}_{\text{geo}}, \quad (13)$$

where α and β are weighting factors that balance the contribution of parameter accuracy and geometric consistency. Higher values of α emphasise parameter accuracy, while higher values of β enforce stronger geometric consistency in the predicted distortion mapping.

The network architecture consists of a convolutional backbone for feature extraction, followed by a fully connected regression head that outputs the five distortion coefficients. The convolutional layers learn spatial features related to geometric deformation in the checkerboard pattern, while the regression head maps these features to the parameter space.

This formulation assumes that distortion-related information is implicitly encoded in the spatial deformation of the observed pattern. However, unlike classical calibration methods that rely on multiple views, the model must infer all parameters from a single image, making the regression problem inherently under-constrained.

3.3 Interpretation and Practical Use

Distortion introduces a global deformation pattern in the image. When a structured characterisation pattern such as a checkerboard is observed, deviations from the ideal grid geometry encode information about the underlying distortion parameters, as can be seen in Figure 2.




Figure 2. Synthetic example of Brown-Conrady distortion applied to a planar checkerboard. (Left) Undistorted reference image. (Right) Distorted image exhibiting increasing geometric deviation toward the image periphery.

The regression model learns to associate these geometric deviations with the underlying distortion parameters. However, this mapping is not strictly one-to-one, as different parameter combinations can produce visually similar distortion patterns, particularly under moderate distortion levels. As a result, the model aims to recover physically plausible parameter sets rather than an exact match to the ground truth. Moreover, unlike classical characterisation procedures, this approach relies solely on information contained within a single image, making the problem inherently under-constrained.

The predicted coefficients can later be used directly within standard undistortion routines (e.g., OpenCV), confirming compatibility with existing characterisation pipelines.

4. Experiments

This section describes the dataset generation, training setup, and evaluation protocol used to assess the proposed single-image distortion regression framework.

4.1 Dataset Generation

A synthetic dataset was generated to provide full control over distortion parameters and ensure the availability of ground-truth labels. A planar checkerboard pattern was rendered under multiple poses, and distortion was applied using the Brown-Conrady model.

Each dataset sample consists of:




Figure 3. Sample training images with their corresponding distortion coefficients.

- a distorted checkerboard image
- a corresponding distortion parameter vector

$$\mathbf{d} = [k_1, k_2, p_1, p_2, k_3]^T.$$

To avoid bias and improve generalisation, distortion parameters were sampled independently for each image from predefined ranges. These ranges were selected to reflect realistic distortion magnitudes derived from prior characterisations of the FPP imaging system, while maintaining numerical stability during training. Specifically, the coefficients were sampled from uniform distributions with $k_1, k_2 \in [-0.3, 0.3]$, $k_3 \in [-0.1, 0.1]$, and $p_1, p_2 \in [-0.005, 0.005]$.

The dataset of 1000 images was split into training, validation, and testing subsets using a 70/20/10 ratio. Each subset contains images generated from independent distortion samples, ensuring that the model is evaluated on unseen parameter combinations. Figure 3 shows sample images from the dataset along with their corresponding distortion coefficients. The synthetic data simulates multiple checkerboard poses, consistent with characterisation procedures in FPP systems, and includes representative image degradations, either individually or combined, such as Gaussian noise and blur encountered during image acquisition.

4.2 Preprocessing

All images were resized to a fixed resolution of 224×224 pixels prior to training to ensure a consistent input size, reduce computational cost, and maintain sufficient geometric information. To reduce the influence of colour information and encourage the model to focus on geometric structure, images were converted to grayscale. The grayscale images were then replicated across three channels to match the pre-trained convolutional network's input format.

Ground-truth distortion parameters are transformed using a signed logarithmic mapping prior to training. This transformation reduces the effect of large parameter magnitudes while

preserving the sign and relative relationships between coefficients, leading to more stable optimisation across parameters with different numerical scales.

4.3 Training Setup and Evaluation

The distortion regression model is implemented using a convolutional neural network with a ResNet18 backbone, followed by a fully connected regression head that outputs the five distortion coefficients. The choice of a relatively lightweight architecture is motivated by the limited dataset size and the structured nature of the input, where simpler models are sufficient to capture geometric features while reducing the risk of overfitting associated with more complex architectures.

The network is trained in a supervised manner using a combined loss function that incorporates both parameter-space accuracy and geometric consistency. The parameter loss is defined as the mean squared error (MSE) between the predicted and ground-truth distortion parameters in the transformed space. To ensure that the predicted parameters correspond to physically meaningful distortions, a geometric loss is introduced based on the Brown-Conrady model that penalises discrepancies between the distorted coordinate mappings generated by the predicted and ground-truth parameters. The total loss is defined as a weighted combination of these two components. Model performance is evaluated using the mean absolute error (MAE) (Goodfellow et al., 2016) computed in the original parameter space on the validation and test sets.

Training is performed for a fixed number of epochs with early stopping based on validation loss to prevent overfitting. The AdamW optimiser (Loshchilov and Hutter, 2017) is used due to its decoupled weight decay formulation, which improves regularisation compared to standard Adam, and a learning rate scheduler is applied to stabilise convergence. Table 1 demonstrates more details about the training configuration.

Table 1. The model training configuration.

Parameter	Value
Backbone	ResNet18
Input size	224×224
Batch size	8
Epochs	50
Optimiser	AdamW
Learning rate	1×10^{-4}
Weight decay	1×10^{-4}
Loss function	Combined (parameter + geometric)
Scheduler	ReduceLROnPlateau

Model performance is evaluated using MAE computed on the test set, reported as an overall average across all coefficients and individually per coefficient. The overall MAE summarises the average prediction error over the full distortion vector, while the per-coefficient MAE highlights the estimation accuracy of each parameter separately. MAE is selected for evaluation as it provides a direct interpretation of the prediction error in the original parameter units.

To assess whether the model learns meaningful image-dependent information, a baseline predictor is introduced: a mean-coefficient model that outputs the average distortion vector computed from the training set for all inputs. This baseline

represents a scenario in which no image information is used, and therefore serves as a lower bound on learning performance. All reported MAE values, including the baseline, are computed in the original parameter space to ensure consistent and fair comparison.

The model is implemented in Python using a standard deep learning framework and trained on a workstation equipped with an NVIDIA Quadro P5000 GPU. All experiments are conducted under consistent training settings, including network architecture, input resolution, preprocessing, optimiser configuration, and training schedule, to ensure fair comparison across configurations. The specific parameter values are summarised in Table 1. These settings were selected based on standard deep learning practices and preliminary experiments to ensure stable and reproducible training behaviour.

5. Results and Discussion

This section presents the quantitative and qualitative evaluation of the proposed single-image distortion regression framework. Results are reported in original coefficient units and compared against a baseline predictor.

5.1 Quantitative Performance

During training, the model exhibits a steady decrease in the training mean squared error (MSE), which serves as the optimisation objective and converges to low values in the final epochs. The MSE is used during optimisation due to its sensitivity to large deviations, encouraging the network to reduce significant errors in the predicted parameters.

In contrast, the validation mean absolute error (MAE) improves primarily during the early stages of training and then stabilises around 0.076, remaining above the baseline MAE of 0.066, which indicates limited generalisation. As shown in Figure 4, this behaviour reflects that improvements in the optimisation loss do not directly translate into improved performance in the original parameter space.

While MSE is used to guide the optimisation process, MAE is reported for validation and testing as it provides a more interpretable measure of prediction accuracy in the original parameter units. Both metrics are shown to illustrate the difference between optimisation behaviour and generalisation performance.




Figure 4. Training and validation error curves over 50 epochs.

Regression accuracy varies across distortion coefficients, as shown in Table 2. The radial parameters (k_1, k_2, k_3) exhibit larger errors compared to the tangential parameters (p_1, p_2), reflecting the increased difficulty of modelling radial distortion

from limited geometric information. In particular, k_1 and k_2 show the highest errors (0.1737 and 0.1128, respectively), indicating that the dominant components of radial distortion are not reliably recovered from a single checkerboard image. In contrast, the tangential coefficients exhibit lower absolute errors (on the order of 10^{-3}), with $p_1 = 0.0142$ and $p_2 = 0.0191$, although their contribution to overall distortion is comparatively smaller.

Table 2. Per-parameter mean absolute error (MAE) on the test set.

Distortion parameter	MAE
k_1 (radial)	0.1737
k_2 (radial)	0.1128
p_1 (tangential)	0.01419
p_2 (tangential)	0.01911
k_3 (radial)	0.0549

The results indicate that, while the model captures general trends in the distortion parameter distribution, it struggles to infer precise parameter values from a single image.

5.2 Qualitative Performance

A qualitative assessment was performed by applying the predicted distortion parameters to unseen undistorted images using standard distortion routines. Figure 5 presents an example of the resulting transformation.



Figure 5. Qualitative comparison of distortion effects. (a) Undistorted reference image. (b) distortion using ground-truth parameters. (c) distortion using predicted parameters. The predicted distortion produces a similar global deformation with slight boundary shifts.

The predicted coefficients produce a geometrically consistent deformation that closely resembles the ground-truth distortion pattern. The overall structure of the checkerboard and the global deformation trend are preserved, indicating that the model captures the dominant distortion characteristics. However, subtle differences can be observed near the image boundaries, where a slight global shift is introduced, resulting in an asymmetric distribution of background regions. This behaviour suggests a small misalignment in the estimated distortion centre, likely arising from inaccuracies in the tangential components. Despite these deviations, the resulting transformation remains physically plausible and visually consistent with the ground-truth distortion.

To further illustrate the relationship between the predicted and ground-truth distortion parameters, a grid-based visual comparison is provided in Figure 6. Unlike the checkerboard pattern, the ideal grid more clearly reveals geometric deviations across the entire image domain.

Consistent with the checkerboard example, the predicted distortion produces a global deformation pattern that closely matches the ground truth. However, the grid representation makes local




Figure 6. Grid-based visualisation of distortion. (a) Undistorted reference grid. (b) Grid distorted using ground-truth parameters. (c) Grid distorted using predicted parameters. The predicted distortion preserves the global deformation pattern with slight deviations in alignment and boundary behaviour.

discrepancies more apparent, particularly in the form of slight misalignments and boundary shifts.

Although these differences are subtle, the overall distortion patterns remain visually similar. This supports the observation that multiple parameter combinations can produce comparable geometric effects, reinforcing the inherent ambiguity of the single-image distortion regression problem.

These observations suggest that the model captures meaningful distortion-related features from the input image, even though it does not recover the exact ground-truth coefficients. Consequently, the model tends to converge towards solutions that yield geometrically consistent corrections rather than exact parameter recovery.

Overall, the qualitative results complement the quantitative findings by demonstrating that, while precise coefficient estimation remains challenging, the proposed framework is capable of producing physically coherent distortion representations from a single image.

5.3 Discussion

The results demonstrate that distortion-related information can be inferred from a single image, with the model consistently producing physically plausible parameter estimates. While the numerical accuracy of the predicted coefficients remains limited, the resulting geometric transformations are coherent and preserve the global distortion structure.

The behaviour observed in both quantitative and qualitative results highlights an inherent ambiguity in the estimation problem. Multiple combinations of distortion coefficients can produce similar image deformations, particularly when only a single view is available. As a result, the regression task does not admit a unique solution, and the model tends to converge towards parameter sets that generate geometrically consistent, rather than exact, reconstructions.

This effect is evident in the qualitative analysis, where the predicted distortion closely matches the global deformation pattern of the ground truth, while exhibiting slight misalignment and boundary shifts. These deviations, although subtle, indicate differences in the estimated distortion centre and emphasise the sensitivity of the mapping to parameter variations. The grid-based visualisation further exposes these discrepancies, revealing local geometric differences that are less apparent in the checkerboard representation.

The variation in performance across distortion coefficients reflects the uneven observability of the parameters. Radial distortion components, which dominate the global deformation

pattern, exhibit higher prediction errors, while tangential components, associated with local asymmetric effects, contribute to subtle spatial shifts and misalignments. This suggests that different parameters influence distinct geometric characteristics in the image.

It is also important to note that the evaluation is conducted on synthetic data, where distortion is applied under controlled conditions. While this enables precise ground-truth comparison, the behaviour of the model under real imaging conditions remains to be investigated. Factors such as sensor noise, illumination variations, and optical imperfections may introduce additional complexity that affects parameter estimation.

Overall, the findings indicate that single-image distortion regression can produce geometrically meaningful predictions, but the ambiguity of the inverse mapping limits the recovery of exact parameter values. This highlights the importance of considering geometric consistency, rather than parameter accuracy alone, when evaluating such approaches.

6. Conclusion

This work investigated the estimation of Brown-Conrady distortion coefficients from a single distorted checkerboard image using a deep learning approach. The problem was formulated as a regression task, where a convolutional neural network predicts physically interpretable distortion parameters that can be directly integrated into standard undistortion pipelines.

The results demonstrate that the model is able to extract distortion-related information from image geometry and produce physically plausible parameter estimates. Although the numerical accuracy of the predicted coefficients remains limited, the resulting geometric transformations are consistent and preserve the global distortion structure, with only minor local deviations.

The findings highlight an inherent ambiguity in the estimation problem, where multiple parameter combinations can produce similar image deformations. Consequently, recovering a unique set of distortion coefficients from a single image is fundamentally challenging. This limitation reflects the under-constrained nature of the problem rather than shortcomings in model capacity or optimisation.

Overall, the proposed framework demonstrates that single-image distortion estimation can yield geometrically meaningful predictions within a standard calibration context. Future work will focus on incorporating additional constraints, such as multi-view information or alternative parameter representations, as well as validating the approach on real-world data to assess its applicability in practical metrology systems, including Fringe Projection Profilometry.

Acknowledgements

The authors would like to thank the UK Engineering and Physical Science Research Council (EPSRC) and Taraz Metrology Ltd for funding this work. Thanks are extended to the Sudanese Materials Scientists & Engineers (SMSE) for their collaboration and assistance in draft reviewing, experimental work, and analysis.

References

- Barazzetti, L., Mussio, L., Remondino, F., Scaioni, M., 2012. Targetless camera calibration. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38, 335–342.
- Brown, D., 1966. Decentering distortion of lenses. *Photogrammetric engineering*, 32(3), 444–462.
- Chaudhry, F. M., Ralli, J., Leudet, J., Sohrab, F., Pakdaman, F., Corbani, P., Gabbouj, M., 2025. Deep-brownconrady: prediction of camera calibration and distortion parameters using deep learning and synthetic data. *IEEE Transactions on Automation Science and Engineering*.
- Conrady, A., 1919. Lens-systems, decentered. *Monthly Notices of the Royal Astronomical Society*, Vol. 79, p. 384–390.
- Devernay, F., Faugeras, O., 2001. Straight lines have to be straight. *Machine vision and applications*, 13(1), 14–24.
- Fitzgibbon, A. W., 2001. Simultaneous linear estimation of multiple view geometry and lens distortion. *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, 1, IEEE, I–I.
- Gasparini, S., Sturm, P., Barreto, J. P., 2009. Plane-based calibration of central catadioptric cameras. *2009 IEEE 12th International Conference on Computer Vision*, IEEE, 1195–1202.
- Geng, J., 2011. Structured-light 3D surface imaging: a tutorial. *Advances in optics and photonics*, 3(2), 128–160.
- Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y., 2016. *Deep learning*. 1Number 2, MIT press Cambridge.
- Hartley, R., Kang, S. B., 2007. Parameter-free radial distortion correction with centre of distortion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(8), 1309–1321.
- Hartley, R., Zisserman, A., 2003. *Multiple view geometry in computer vision*. Cambridge University Press.
- Heikkilä, J., Silvén, O., 1997. A four-step camera calibration procedure with implicit image correction. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE, 1106–1112.
- Kendall, A., Grimes, M., Cipolla, R., 2015. PoseNet: A convolutional network for real-time 6-dof camera relocalisation. *Proceedings of the IEEE International Conference on Computer Vision*, 2938–2946.
- Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularisation. *arXiv preprint arXiv:1711.05101*.
- Luhmann, T., Robson, S., Kyle, S., Boehm, J., 2023. *Close-range photogrammetry and 3D imaging*. Walter de Gruyter GmbH & Co KG.
- Martins, P. F., Costelha, H., Bento, L. C., Neves, C., 2020. Monocular camera calibration for autonomous driving—a comparative study. *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, IEEE, 306–311.
- Mei, C., Rives, P., 2007. Single view point omnidirectional camera calibration from planar grids. *Proceedings 2007 IEEE International Conference on Robotics and Automation*, IEEE, 3945–3950.
- Zhang, S., 2010. Recent progresses on real-time 3D shape measurement using digital fringe projection techniques. *Optics and lasers in engineering*, 48(2), 149–158.
- Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11), 1330–1334.