

Blind Geometric Distortion Correction on Images Through Deep Learning

Xiaoyu Li¹ Bo Zhang¹ Pedro V. Sander¹ Jing Liao²

¹The Hong Kong University of Science and Technology ²City University of Hong Kong

Abstract

We propose the first general framework to automatically correct different types of geometric distortion in a single input image. Our proposed method employs convolutional neural networks (CNNs) trained by using a large synthetic distortion dataset to predict the displacement field between distorted images and corrected images. A model fitting method uses the CNN output to estimate the distortion parameters, achieving a more accurate prediction. The final corrected image is generated based on the predicted flow using an efficient, high-quality resampling method. Experimental results demonstrate that our algorithm outperforms traditional correction methods, and allows for interesting applications such as distortion transfer, distortion exaggeration, and co-occurring distortion correction.

1. Introduction

Geometric distortion is a common problem in digital imagery and occurs in a wide range of applications. It can be caused by the acquisition system (e.g., optical lens, imaging sensor), imaging environment (e.g., motions of the platform or target, viewing geometry) and image processing operations (e.g., image warping). For example, camera lenses often suffer from optical aberrations, causing barrel distortion (\mathcal{B}), common in wide angle lenses, where the image magnification decreases with distance from the optical axis, and pincushion distortion ($\mathcal{P}i$), where it increases. While lens distortions are intrinsic to the camera, extrinsic geometric distortions like rotation (\mathcal{R}), shear (\mathcal{S}) and perspective distortion (\mathcal{P}) may also arise from the improper pose or the movement of cameras. Furthermore, a wide number of distortion effects, such as wave distortion (\mathcal{W}), can be generated by image processing tools. We aim to design an algorithm that can automatically correct images with these distortions and can be generalized to a wide range of distortions easily (see Figure 2).

Geometric distortion correction is highly desired in both photography and computer vision applications. For exam-

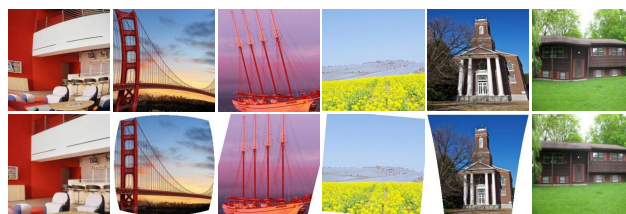


Figure 1. Our proposed learning-based method can blindly correct images with different types of geometric distortion (first row) providing high-quality results (second row).

ple, lens distortion violates the pin-hole camera model assumption which many algorithms rely on. Second, remote sensing images usually contain geometric distortions that cannot be used with maps directly before correction [34]. Third, skew detection and correction is an important pre-processing step in document analysis and has a direct effect on the reliability and efficiency of the segmentation and feature extraction stages [1]. Finally, photos often contain slanted buildings, walls, and horizon lines due to improper camera rotation. Our visual system expects man-made structures to be straight, and horizon lines to be horizontal [24].

Completely blind geometric distortion correction is a challenging problem, which is under-constrained given that the input is only a single distorted image. Therefore, many correction methods have been proposed by using multiple images or additional information. Multiple views methods [4, 15, 23] for radial lens distortion use point correspondences of two or more images. These methods can achieve impressive results. However, they cannot be applied when multiple images under camera motion are unavailable.

To address these limitations, distortion correction from a single image has also been explored. Methods for radial lens distortion based on the plumb line approach [39, 5, 33] assume that straight lines are projected to circular arcs in the image plane caused by radial lens distortion. Therefore, accurate line detection is a very important aspect for the robustness and flexibility of these methods. Correction methods for other distortions [13, 24, 7, 31] also rely on

the detection of special low-level features such as vanishing points, repeated textures, and co-planar circles. But these special low-level features are not always frequent enough for distortion estimation in some images, which greatly restrict the versatility of the methods. Moreover, all of the methods focus on a specific distortion. To our knowledge, there is no general framework which can address different types of geometric distortion from a single image.

In this paper, we propose a learning-based method to achieve this goal. We use the displacement field between distorted images and corrected images to represent a wide range of distortions. The correction problem is then converted to the pixel-wise prediction of this displacement field, or flow, from a single image. Recently, CNNs have become a powerful method in many fields of computer vision and outperform many traditional methods, which motivated us to use a similar network structure for training. The predicted flow is then further improved by our model fitting methods which estimate the distortion parameters. Lastly, we use a modified resampling method to generate the output undistorted image from the predicted flow.

Overall, our learning-based method does not make strong assumptions on the input images while generating high-quality results with few visible artifacts as shown in Figure 1. Our main contribution is to propose the first learning-based methods to correct a wide range of geometric distortions blindly. More specifically, we propose:

1. A single-model network, which implicitly learns the distortion parameters given the distortion type.
2. A multi-model network, which performs type classification jointly with flow regression without knowing the distortion type, followed by an optional model fitting method to further improve the accuracy of the estimation.
3. A new resampling method based on an iterative search with faster convergence.
4. Extended applications that can directly use this framework, such as distortion transfer, distortion exaggeration, and co-occurring distortion correction.

2. Related Work

Geometric distortion correction. For camera lens distortions, pre-calibration techniques have been proposed for correction with known distortion parameters [32, 10, 35, 18, 44]. However, they are unsuitable for zoom lenses, and the calibration process is usually tedious. On the other hand, auto-calibration methods do not require special calibration patterns and automatically extracts camera parameters from multi-view images [12, 15, 23, 28, 19]. But for many application scenarios, multiple images with different views are

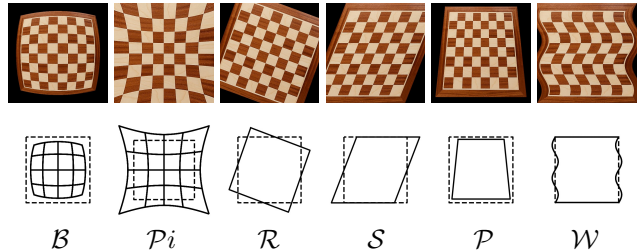


Figure 2. Our system is trained to correct barrel distortion (\mathcal{B}), pincushion ($\mathcal{P}i$), rotation (\mathcal{R}), shear (\mathcal{S}), perspective (\mathcal{P}) and wave distortion (\mathcal{W}).

unavailable. To address these limitations, automatic distortion correction from a single image has gained more research interest recently. Fitzgibbon [12] proposes a division model to approximate the radial distortion curve with higher accuracy and fewer parameters. Wang *et al.* [39] studied the geometry property of straight lines under the division model and proposed to estimate the distortion parameters through arc fitting. Since plumb line methods rely on robust line detection, Aleman-Flores *et al.* [2] used an improved Hough Transform to improve the robustness while Bukhari and Dailey [5] proposed a sampling method that robustly chooses the circular arcs and determines distortion parameters that are insensitive to outliers.

For other distortions, such as rotation and perspective, most of the image correction methods rely on the detection of low-level features such as vanishing point, repeated textures, and co-planar circles [13, 24, 7, 31]. Recently, Zhai *et al.* [43] proposed to use deep convolutional neural networks to estimate the horizon line by aggregating the global image context with the clue of the vanishing point. Workman *et al.* [40] goes further and directly estimates the horizon line in the single image. Unlike these specialized methods, our approach is generalizable for multi-type distortion correction using a single network.

Deformation estimation. There has been recent work on automatic detection of geometric deformation or variations in a single image. Dekel *et al.* [9] use a non-local variations algorithm to automatically detect and correct small deformations between repeating structures from a single image. Wadhwa *et al.* [37] fit parametric models to compute the geometric deviations and exaggerate the departure from ideal geometries. Estimating deformations has also been studied in the context of texture images [22, 16, 27]. None of these techniques are learning-based and are mostly for specialized domains.

Neural networks for pixel-wise prediction. Recently, convolutional neural networks have been used in many pixel-wise prediction tasks from a single image, such as se-

mantic segmentation [25], depth estimation [11] and motion prediction [38]. One of the main problems for dense prediction is how to combine multi-scale contextual reasoning with the full-resolution output. Long *et al.* [25] proposed fully convolutional networks which popularized CNNs for dense predictions without fully connected layers. Some methods focus on dilated or atrous convolution [42, 8] which supports exponential expansion the receptive field and systematically aggregate multi-scale contextual information without losing resolution. Another strategy is to use the encoder-decoder architecture [26, 3, 29]. The encoder gradually reduces the spatial dimension to increase the receptive field of the neuron, while the decoder maps the low-resolution feature maps to full input resolution maps. Noh *et al.* [26] developed deconvolution and unpooling layers for the decoder part. Badrinarayanan *et al.* [3] used pooling indices to connect the encoder and the corresponding decoder, making the architecture more memory efficient. Another popular network is U-net [29], which uses the skip connections to combine the contracting paths with the upsampled feature maps. Our networks use an encoder-decoder architecture with residual connection design and achieve more accurate results.

3. Network Architectures

Geometrically distorted images usually exhibit unnatural structures that can serve as clues for distortion correction. As a result, we presume that the network can potentially recognize the geometric distortions by extracting the features from the input image. We, therefore, propose a network to learn the mapping from the image domain \mathcal{I} to the flow domain \mathcal{F} . The flow is the 2D vector field that specifies where pixels in the input image should move in order to get the corrected image. It defines a non-parametric transformation, thus being able to represent a wide range of distortions. Since the flow is a forward map from the distorted image to the corrected image, a resampling method is needed to produce the final result.

This strategy follows learning methods of other applications which have observed that it is often simpler to predict the transformation from input to output rather than predicting the output directly (e.g., [14, 20]). Thus, we designed our architecture to learn an intermediate flow representation. Additionally, the forward mapping indicates where each pixel with a known color in the distorted image maps to. Therefore, all pixels in the input image learn a distortion flow prediction directly associated to them, which would not be the case if we were attempting to learn a backward mapping, where some input regions could not have correspondences. It can be a serious problem when the distortion changes the image shape greatly. Furthermore, our resampling method that is required to generate the final image is fast and accurate.

We propose two networks by considering whether the user has prior knowledge of the distortion type. Our networks are trained in a supervised manner. Therefore, we first introduce how the paired datasets have been constructed (Section 3.1) and then we introduce our two networks, for single-model and multi-model distortion estimation (Sections 3.2 and 3.3, respectively).

3.1. Dataset construction

We generate the distorted image flow pair by warping an image with a given mapping, thus constructing the distorted image dataset \mathcal{I} and its corresponding distortion flow dataset \mathcal{F} , where $I_j \in \mathcal{I}$ and $F_j \in \mathcal{F}$ are paired.

We consider six geometric distortion models in our network. However, the architecture is not specialized to these types of distortion and can potentially be further extended. Each distortion type $\beta = 1, \dots, 6$ has a model \mathcal{M}^β , which defines the mapping from the distorted image lattice to the original one. Bilinear interpolation is used if the corresponding point in the original image is not on the integer grid. The flow $F = \mathcal{M}^\beta(\rho^\beta)$, $F \in \mathcal{F}$, is generated in the meantime to record how the pixel in the distorted image should be moved to the corresponding point in the original image. ρ^β is the distortion parameter that controls the distortion effect. For instance, in the rotation distortion model, ρ^β is the rotation angle while in the barrel and pincushion model, ρ^β represents the parameter in Fitzgibbon’s single parameter division model [12]. All distortion parameters ρ^β in different distortion models are randomly sampled using a uniform distribution within a specified range. As Figure 2 shows, the geometric distortions change the image shapes. Thus we crop the images and flows to remove empty regions.

3.2. Single-model distortion estimation

We first introduce a network \mathcal{N}^β parameterized by θ^β to estimate the flow for distorted images with a known distortion type β . \mathcal{N}^β learns the mapping from \mathcal{I}^β to \mathcal{F}^β with sub-datasets where $\mathcal{I}^\beta \subset \mathcal{I}$ and $\mathcal{F}^\beta \subset \mathcal{F}$ are sub-domains containing the images and flows of distortion type β .

Architecture. A possible architecture choice is to directly regress the distortion flow according to the ground truth with an auto-encoder-like structure. However, the network would only be optimized with the pixel-wise flow error, without taking advantage of global constraints imposed by the known distortion model. Instead, we design a network to first predict the model parameter ρ^β directly. This parameter is then used to generate the flow $F = \mathcal{M}^\beta(\rho^\beta)$ in the network. Though the network should implicitly learn the distortion parameter, there is no explicit constraint for the network to do so exactly.

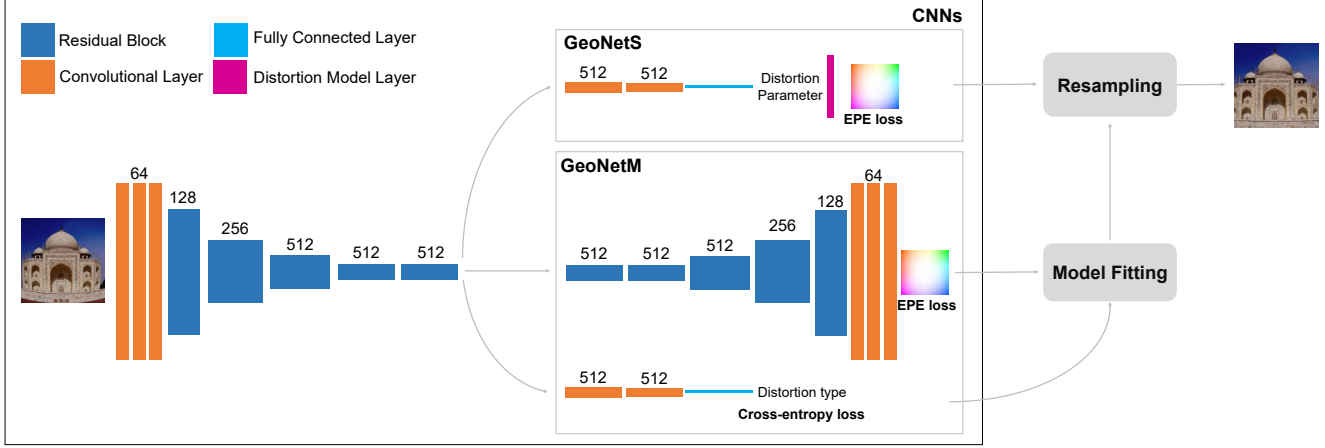


Figure 3. Overview of our entire framework, including the single-model (GeoNetS) and multi-model (GeoNetM) distortion networks (Section 3), and resampling (Section 4). Each box represents some conv layers, with vertical dimension indicating feature map spatial resolution, and horizontal dimension indicating the output channels of each conv layer in the box.

The network architecture, referred to as GeoNetS, is shown in Figure 3. It has three conv layers at the very beginning and five residual blocks ([17]) to gradually downsize the input image and extract the features. Each residual contains two conv layers and has a shortcut connection from input to output. The shortcut connection helps ease the gradient flow, achieving a lower loss according to our experiments. Downsampling in spatial resolution is achieved using conv layers with a stride of 2 and 3×3 kernels. Batch normalization layers and ReLU function are added after each conv layer, which significantly improves training.

After the residual blocks, two conv layers are used to downsize the features further, and a fully-connected layer converts the 3D feature map into a 1D vector ρ^β . With the distortion parameter ρ^β , the corresponding distortion model \mathcal{M}^β analytically generates the distortion flow. The network is optimized with the pixel-wise flow error between the generated flow and the ground truth.

Loss We train the network to minimize the loss \mathcal{L} , which measures the distance between the estimated distortion flow and the ground truth flow:

$$\begin{aligned} \theta^\beta &= \arg \min_{\theta^\beta} \mathcal{L}(\mathcal{N}^\beta(I; \theta^\beta), F) \\ \mathcal{N}^\beta(I; \theta^\beta) &= \mathcal{M}^\beta(n^\beta(I; \theta^\beta)) \end{aligned} \quad (1)$$

where n^β is the sub-network of \mathcal{N}^β , represents the part to regress the distortion parameter implicitly. Here we choose the endpoint error (EPE) as our loss function. The EPE is defined as the Euclidean distance between the predicted flow vector and the ground truth averaged over all pixels. Because the estimated distortion flow is explicitly constrained by the distortion model, it is naturally smooth.

Since the geometric distortion models \mathcal{M}^β we consider

are differentiable, the backward gradient of each layer can be computed using the chain rule:

$$\frac{\partial \mathcal{L}}{\partial \theta^\beta} = \frac{\partial \mathcal{L}}{\partial \mathcal{M}^\beta} \frac{\partial \mathcal{M}^\beta}{\partial n^\beta} \frac{\partial n^\beta}{\partial \theta^\beta} \quad (2)$$

Our trained network can estimate the distortion flow blindly from an input image for each distortion type and achieve comparable performance as traditional methods.

3.3. Multi-model distortion estimation

The GeoNetS network is only able to capture a specific distortion type with a distortion model at a time. For a new type, the entire network has to be retrained. Furthermore, the distortion type and model can be unknown in some cases. In view of these limitations, we designed a second network for multi-model distortion estimation. However, since the distortion model and the parameters ρ^β can vary drastically across types, it is impossible to train a multi-model network with the model constraints. We train a network to regress the distortion flow without model constraints and at the same time classify the distortion type. The network is illustrated in Figure 3. The multi-model network \mathcal{N} parameterized by θ is jointly trained for two tasks. The first task estimates the distortion flow, learning the mapping from the image domain \mathcal{I} to the flow domain \mathcal{F} . The second task classifies the distortion type, learning the mapping from image domain \mathcal{I} to type domain \mathcal{T} .

Architecture The entire network adopts an encoder-decoder structure, which includes an encoder part, a decoder part, and a classification part. The input image is fed into an encoder to encode the geometric features and capture the unnatural structures. Then two branches follow: In the first branch, a decoder is used to regress the distortion

flow, while in the second branch a classification subnet is used to classify the distortion type. The encoder part is the same as GeoNetS, and the decoder part is symmetric to the encoder. Downsampling/Upsampling in spatial resolution is achieved using conv/upconv layers with a stride of 2. The classification part also has two conv layers to downsize the features further, and a fully-connected layer converts the 3D feature map into a 1D score vector of each type.

Loss We use the EPE loss $\mathcal{L}_{\text{flow}}$ in the flow regression branch, and a cross entropy loss in the classification branch. The two branches are jointly optimized by minimizing the total loss:

$$*\theta = \arg \min_{\theta} (\mathcal{L}_{\text{flow}} + \lambda \mathcal{L}_{\text{class}}) \quad (3)$$

where the weight λ provides a trade-off between the flow prediction and the distortion type classification.

We observe that jointly learning the distortion type helps reduce the flow prediction error as well. These two branches share the same encoder, and the classification branch helps the encoder learn the geometric features for different distortion types better. Please refer to Section 5 for direct comparisons.

Model fitting Our multi-model network simultaneously predicts the flow and the distortion type from the input image. Based on this information, we can estimate the actual distortion parameters in the model and regenerate the flow to obtain a more accurate result.

The Hough Transform is a widely used technique to extract features in an image. It is robust to noise by eliminating the outliers in the flow using a voting procedure. Moreover, it is a non-iterative approach. Each data point is treated independently, and therefore parallel processing of all points is possible. This makes it more computationally efficient.

For an input image I , given its distortion type β and distortion flow $\mathcal{N}(I; * \theta)$ predicted by our network, we want to fit the corresponding distortion model \mathcal{M}^{β} with the distortion parameter ρ^{β} . In our scenario, we map each data point \mathcal{N}_{ij} in flow $\mathcal{N}(I; * \theta)$ at position (i, j) to a point in the distortion parameter space. The transform is given by

$$\rho_{ij} = \mathcal{M}^{-1}(\mathcal{N}_{ij}) \quad (4)$$

We assume the distortion parameter ρ has a range from ρ_{\min} to ρ_{\max} and split the range into M cells uniformly. All the points ρ_{ij} belong to a cell according to the parameter values. The cell receiving the maximum number of counts determine the best fitting result, and the final result is the average of all the points in this cell. We let $M = 100$ in our experiments.

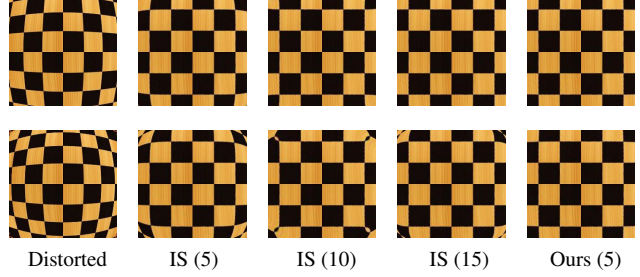


Figure 4. Comparison of the convergence using the traditional iterative search (IS) and our approach. Two examples with different distortion levels are used. Our method converges to good resampling result with 5 iterations.

Once model fitting is completed, we have a refined and smoother flow $F = \mathcal{M}^{\beta}(\rho^{\beta})$. With model fitting, the efficiency in correcting of higher resolution images can be greatly improved. This is because we can estimate the flow and obtain the distortion parameter ρ at a much smaller resolution, and generate the full resolution flow directly according to the distortion parameter.

4. Resampling

Given the distortion flow, we employ a pixel evaluation algorithm to determine the backward-mapping and resample the final undistorted image. The approach is inspired by the bidirectional iterative search algorithm in [41]. Unlike mesh rasterization approaches, this iterative method runs entirely independently and in parallel for each pixel, fetching the color from the appropriate location of the source image.

The traditional backward mapping algorithm of [41] seeks to find a point p in the source image that maps to q . Since we only have the forward distortion flow, this method essentially inverts this mapping using an iterative search until the location p converges:

$$\begin{aligned} p^{(0)} &= q \\ p^{(i+1)} &= q - f(p^{(i)}) \end{aligned} \quad (5)$$

where $f(p)$ is the computed forward flow from the source pixel p to the undistorted image.

Since the application in this paper often involves large, smooth distortions, we propose a modification that significantly improves the convergence rate and quality. The traditional method initializes p based on the flow at q . More specifically, $p^{(1)} = q - f(p^{(0)})$. If $f(p^{(1)}) \approx f(p^{(0)})$, then the iterative search converges quickly. However, in the presence of large distortions, $\|f(p^{(0)})\|$ is large, $p^{(0)}$ and $p^{(1)}$ are distant, and thus, $f(p^{(1)})$ and $f(p^{(0)})$ can be very different, making it a poor initialization and decreasing conversion speed.

Instead of assuming that the flow in $p^{(0)}$ and $p^{(1)}$ are the same, we compute the local derivative of the flow at $p^{(0)}$ using the finite difference method, and use this derivative to estimate the flow at $p^{(1)}$. We let $f_x(p)$ and $f_y(p)$ represent the horizontal and vertical flow respectively. Formally,

$$\frac{df_x}{dx} = \frac{f_x(p_{nx}^{(0)}) - f_x(p^{(0)})}{(x+1) - x} \quad (6)$$

where $p^{(0)}$ is at coordinates (x, y) , and its horizontal pixel neighbor p_{nx} is at coordinates $(x+1, y)$. We then use this derivative to approximate the flow at $p^{(1)} = (x', y')$:

$$\frac{df_x}{dx} = \frac{f_x(p^{(1)}) - f_x(p^{(0)})}{x' - x} \quad (7)$$

By the definition of forward flow, we have $f_x(p^{(1)}) = x - x'$. Therefore we can compute x' combining Equation 6 and Equation 7:

$$x' = x - \frac{f_x(p^{(0)})}{1 + f_x(p_{nx}^{(0)}) - f_x(p^{(0)})} \quad (8)$$

We compute y' similarly and proceed with the iterative search. Note that we only use this finite difference method in the first iteration to get a coarse initial estimation. The traditional, faster iterative search is used to finetune until convergence.

5. Experiments

In this section, we report the results of our work. We first analyze the results of our proposed networks in Section 5.1. Then we discuss the results of our resampling method in Section 5.2. In Section 5.3, we show qualitative and quantitative comparisons of our approach to previous methods for correcting specific distortion types. In Section 5.4, we show some applications of our method. CNNs training details are given in the supplementary material.

5.1. Networks

To evaluate the performance of GeoNetS, we compare with GeoNetM without classification branch and train these two networks on the same dataset with only single-type distortion. The first two rows in Table 1 show that by explicitly considering the distortion model in the network, GeoNetS achieves better results. Moreover, the flow is globally smooth due to the restriction given by the distortion model.

Second, we examine how the joint learning with classification improves the distortion flow prediction. The third and fourth rows in Table 1 show that GeoNetM with joint learning using the classification branch has more accurate prediction than GeoNetM w/o classification training in the multi-type distortion dataset. The classification accuracy of

GeoNetM is 97.3% for these six kinds of distortion. Table 1 also shows that single-type achieves lower flow error than multi-type since additional information needs to be learned for the multi-type task.

We also examine whether the model fitting method improves prediction accuracy for GeoNetM. The last two rows in Table 1 shows that the Hough transform based model fitting provides more accurate results. More results from GeoNetM are shown in Figure 5. For each example, the distorted image is shown on the left, the corrected output image in the middle, and the three flows (before fitting, after fitting, and ground truth) on the right. More real image results and detailed discussion of model fitting, GeoNetS and GeoNetM are given in the supplementary material.

5.2. Resampling

Next we present the results of our resampling strategy. Figure 4 shows results applied to images with two different distortion levels. Note that, on the top row, it takes roughly 10 iterations to converge using the traditional iterative search approach (IS), whereas 5 iterations suffice when using our initialization. On the second row, with a more severe distortion, even after 15 iterations, the traditional method has not satisfactorily converged, whereas with our initialization the results also converge within 5 iterations.

Figure 6 demonstrates how our method more quickly converges to the ground truth (left), and how the vast majority of pixels already have an endpoint error lower than 1/5 of a pixel after just 5 iterations. A parallel version of our approach has been implemented on the GPU using an Intel Xeon E5-2670 v3 2.3 GHz machine with Nvidia Tesla K80 GPU. It can resample the image under 50 ms.

5.3. Comparison with previous techniques

Next, we compare our distortion correction technique to some existing methods that are specialized to some distortion types. Note that, unlike these methods, our learning-based approach is able to handle different distortion types.

Lens distortion. Figure 7 compares our approach with [2] and [30], which are specialized for lens distortion. Note that for cases where the image has obvious distortion (e.g., first row), all methods can correct accurately. However, in cases where the distortion is more subtle or does not exhibit highly distorted lines (e.g., bottom two rows), our approach yields improved results. Figure 8 shows a quantitative comparison based on 50 randomly chosen images from the dataset of [21]. These images include a variety of scene types (e.g., nature, man-made, water) and are distorted with random distortion parameters to generate our synthetic dataset. All of these methods use Fitzgibbon’s single parameter division model [12], therefore we can calculate the relative error of the distortion correction parameters

Configuration		EPE						
Architecture	Training dataset	\mathcal{B}	\mathcal{P}_i	\mathcal{R}	\mathcal{S}	\mathcal{P}	\mathcal{W}	Average
GeoNetS	Single-type	1.43	0.79	2.41	2.19	0.89	1.06	1.46
GeoNetM w/o Clas	Single-type	1.57	1.12	3.01	2.91	1.01	1.32	1.82
GeoNetM w/o Clas	Multi-type	3.07	2.24	3.75	4.99	3.35	1.73	3.19
GeoNetM	Multi-type	2.72	2.03	3.68	3.12	3.29	1.67	2.75
GeoNetM w/ Hou	Multi-type	1.78	1.34	2.77	2.27	2.25	1.22	1.94

Table 1. EPE and classification statistics of our approach using 500 test images per distortion.

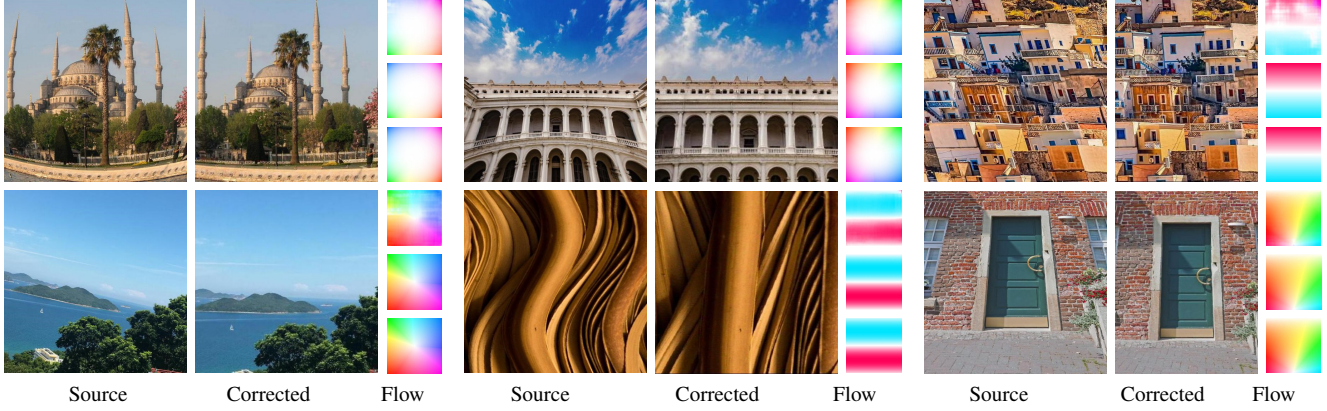


Figure 5. Results of distortions that we considered. Top row: Barrel distortion, pincushion distortion and shear distortion. Bottom row: rotation, wave distortion and perspective distortion. The flows refer to the flow before model fitting (top), after model fitting (middle), and ground truth (bottom).

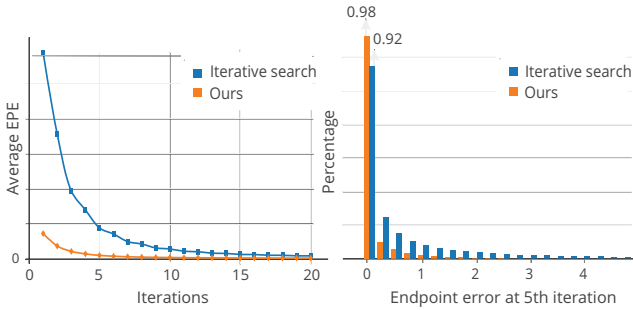


Figure 6. Convergence of EPE (left) and its histogram after five iterations (right) for traditional iterative search and our method. Test on 10 images with different distortion levels.

Table 2. Angle deviation of detected lines with the vertical angle.

method	baseline (input)	[6]	ours
angle deviation	6.44°	3.03°	2.81°

for comparison. Note that, with our approach, the number of sample images (y-axis) that lie within the error thresholds (x-axis) is significantly higher than the other methods.

Perspective distortion. For the perspective distortion, we compare with [6]. Here we use angle deviation as the

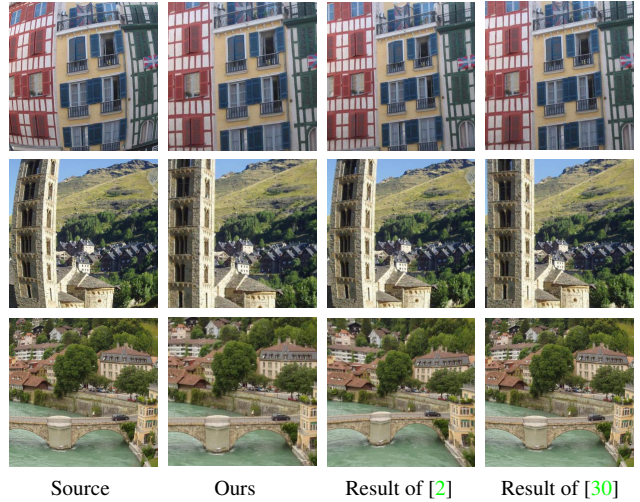


Figure 7. Qualitative comparison with state-of-the-art lens distortion correction methods.

metric. We collect 30 building images under orthographic projection and distort them with different Homography matrices. We control the distorted vertical lines within $[70^\circ, 110^\circ]$. Then we detect straight lines [36] in the correction results using the line segment detector within this range and assume that the angle of these lines should be 90° after

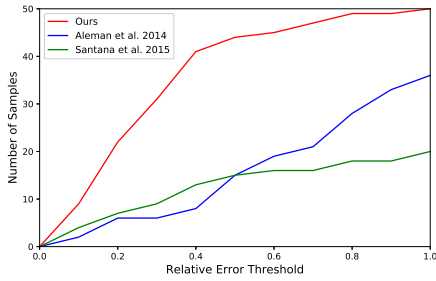


Figure 8. Quantitative comparison on lens distortion correction methods. Given relative error threshold, our method gives more accurate pixels than the other methods.

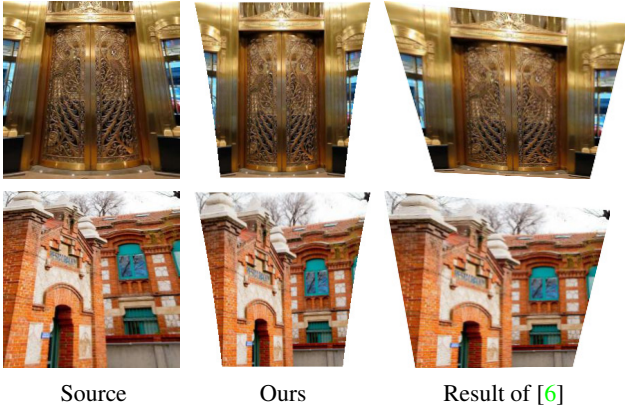


Figure 9. Comparison with previous perspective correction method.



Figure 10. Examples of distortion transfer. Perspective and barrel distortions are used.

correction and calculate their average angle deviation. As shown in Table 2 and Figure 9, our method outperforms the previous approach [6].

5.4. Applications

In addition, we explored applications that can benefit from our distortion correction method directly.



Figure 11. Example of distortion exaggeration.

Distortion transfer. Our system can detect the distortion from a reference image and transfer to a target image. We can estimate the forward flow from the reference image to the corrected version and then directly apply it to the target by bilinear interpolation. Figure 10 shows two examples of transferring distortion from a reference image to a target image, in order to accentuate the perspective of a house photograph (upper row) or apply aggressive barrel distortion to a portrait (lower row).

Distortion exaggeration. To achieve distortion exaggeration, we can reverse the direction of estimated flow field to make the pixels further away from its undistorted position, and use our resampling approach to generate an exaggerated distortion output. Figure 11 shows a building with perspective effect, we can adjust the level of distortion to exaggerate the effect by amplifying or reversing the flow, respectively.

Co-occurring distortion correction. Sometimes an image could have more than one type of distortion. We can correct the distorted image simply by running our correction algorithm twice iteratively. For each iteration, it detects and corrects the most severe type of distortion that it encounters. See the supplementary material for some examples and results.

6. Conclusion

In conclusion, we present the first approach to blindly correct several types of geometric distortions from a single image. Our approach uses a deep learning method trained on several common distortions to detect the distortion flow and type. Our model fitting and parameter estimation approach then accurately predicts the distortion parameters. Finally, we present a fast parallel approach to resample the distortion-corrected images. We compare our techniques to recent specialized methods for distortion correction and present applications such as distortion transfer, distortion exaggeration, and co-occurring distortion correction.

Acknowledgements: This work was partly supported by CityU of Hong Kong Start-up Grant No. 7200607/CS, and Hong Kong GRF Grant No. 16208814.

References

- [1] A. M. Al-Shatnawi and K. Omar. Skew detection and correction technique for arabic document images based on centre of gravity. *Journal of Computer Science*, 5(5):363, 2009. [1](#)
- [2] M. Alemán-Flores, L. Alvarez, L. Gomez, and D. Santana-Cedr s. Automatic lens distortion correction using one-parameter division models. *Image Processing On Line*, 4:327–343, 2014. [2](#), [6](#), [7](#)
- [3] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*, 2015. [3](#)
- [4] J. P. Barreto and K. Daniilidis. Fundamental matrix for cameras with radial distortion. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 1, pages 625–632. IEEE, 2005. [1](#)
- [5] F. Bukhari and M. N. Dailey. Automatic radial distortion estimation from a single image. *Journal of mathematical imaging and vision*, 45(1):31–45, 2013. [1](#), [2](#)
- [6] K. Chaudhury, S. DiVerdi, and S. Ioffe. Auto-rectification of user photos. In *IEEE Image Processing (ICIP)*, pages 3479–3483, 2014. [7](#), [8](#)
- [7] K. Chaudhury and S. J. DiVerdi. Automatic rectification of distortions in images, June 23 2015. US Patent 9,064,309. [1](#), [2](#)
- [8] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018. [3](#)
- [9] T. Dekel, T. Michaeli, M. Irani, and W. T. Freeman. Revealing and modifying non-local variations in a single image. *ACM Transactions on Graphics (TOG)*, 34(6):227, 2015. [2](#)
- [10] C. B. Duane. Close-range camera calibration. *Photogramm. Eng.*, 37(8):855–866, 1971. [2](#)
- [11] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. [3](#)
- [12] A. W. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–I. IEEE, 2001. [2](#), [3](#), [6](#)
- [13] A. C. Gallagher. Using vanishing points to correct camera rotation in images. In *Computer and Robot Vision, 2005. Proceedings. The 2nd Canadian Conference on*, pages 460–467. IEEE, 2005. [1](#), [2](#)
- [14] M. Gharbi, Y. Shih, G. Chaurasia, J. Ragan-Kelley, S. Paris, and F. Durand. Transform recipes for efficient cloud photo enhancement. *ACM Transactions on Graphics (TOG)*, 34(6):228, 2015. [3](#)
- [15] R. Hartley and S. B. Kang. Parameter-free radial distortion correction with center of distortion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(8):1309–1321, 2007. [1](#), [2](#)
- [16] J. Hays, M. Leordeanu, A. A. Efros, and Y. Liu. Discovering texture regularity as a higher-order correspondence problem. In *European Conference on Computer Vision*, pages 522–535. Springer, 2006. [2](#)
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#)
- [18] J. Heikkil  and O. Silven. A four-step camera calibration procedure with implicit image correction. In *Computer Vision and Pattern Recognition, 1997. Proceedings., 1997 IEEE Computer Society Conference on*, pages 1106–1112. IEEE, 1997. [2](#)
- [19] J. Henrique Brito, R. Angst, K. Koser, and M. Pollefeys. Radial distortion self-calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1368–1375, 2013. [2](#)
- [20] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017. [3](#)
- [21] H. Jegou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *European conference on computer vision*, pages 304–317. Springer, 2008. [6](#)
- [22] V. G. Kim, Y. Lipman, and T. A. Funkhouser. Symmetry-guided texture synthesis and manipulation. *ACM Trans. Graph.*, 31(3):22–1, 2012. [2](#)
- [23] Z. Kukeleva and T. Pajdla. A minimal solution to radial distortion autocalibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2410–2422, 2011. [1](#), [2](#)
- [24] H. Lee, E. Shechtman, J. Wang, and S. Lee. Automatic upright adjustment of photographs. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 877–884. IEEE, 2012. [1](#), [2](#)
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [3](#)
- [26] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015. [3](#)
- [27] M. Park, K. Brocklehurst, R. T. Collins, and Y. Liu. Deformed lattice detection in real-world images using mean-shift belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(10):1804–1816, 2009. [2](#)
- [28] S. Ramalingam, P. Sturm, and S. K. Lodha. Generic self-calibration of central cameras. *Computer Vision and Image Understanding*, 114(2):210–219, 2010. [2](#)
- [29] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [3](#)
- [30] D. Santana-Cedr s, L. Gomez, M. Alem n-Flores, A. Salgado, J. Esclar n, L. Mazarra, and L. Alvarez. Invertibility and estimation of two-parameter polynomial and division lens distortion models. *SIAM Journal on Imaging Sciences*, 8(3):1574–1606, 2015. [6](#), [7](#)

- [31] D. Santana-Cedr s, L. Gomez, M. Alem n-Flores, A. Salgado, J. Esclar n, L. Mazorra, and L. Alvarez. Automatic correction of perspective and optical distortions. *Computer Vision and Image Understanding*, 161:1–10, 2017. 1, 2
- [32] J.-P. Tardif, P. Sturm, M. Trudeau, and S. Roy. Calibration of cameras with radially symmetric distortion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(9):1552–1566, 2009. 2
- [33] T. Thorm hlen, H. Broszio, and I. Wassermann. Robust line-based calibration of lens distortion from a single view. *Mirage 2003*, pages 105–112, 2003. 1
- [34] T. Toutin. Geometric processing of remote sensing images: models, algorithms and methods. *International journal of remote sensing*, 25(10):1893–1924, 2004. 1
- [35] 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, 1987. 2
- [36] R. G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall. Lsd: A fast line segment detector with a false detection control. *IEEE transactions on pattern analysis and machine intelligence*, 32(4):722–732, 2010. 7
- [37] N. Wadhwa, T. Dekel, D. Wei, F. Durand, and W. T. Freeman. Deviation magnification: revealing departures from ideal geometries. *ACM Transactions on Graphics (TOG)*, 34(6):226, 2015. 2
- [38] J. Walker, A. Gupta, and M. Hebert. Dense optical flow prediction from a static image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2443–2451, 2015. 3
- [39] A. Wang, T. Qiu, and L. Shao. A simple method of radial distortion correction with centre of distortion estimation. *Journal of Mathematical Imaging and Vision*, 35(3):165–172, 2009. 1, 2
- [40] S. Workman, M. Zhai, and N. Jacobs. Horizon lines in the wild. *arXiv preprint arXiv:1604.02129*, 2016. 2
- [41] L. Yang, Y.-C. Tse, P. V. Sander, J. Lawrence, D. Nehab, H. Hoppe, and C. L. Wilkins. Image-based bidirectional scene reprojection. *ACM Trans. Graph.*, 30(6):150:1–150:10, 2011. 5
- [42] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. 3
- [43] M. Zhai, S. Workman, and N. Jacobs. Detecting vanishing points using global image context in a non-manhattanworld. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 5657–5665. IEEE, 2016. 2
- [44] Z. Zhang. Flexible camera calibration by viewing a plane from unknown orientations. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 1, pages 666–673. Ieee, 1999. 2