

From two rolling shutters to one global shutter

Cenek Albl¹ Zuzana Kukelova² Viktor Larsson¹ Michal Polic³ Tomas Pajdla³ Konrad Schindler¹

¹ETH Zurich

²VRG, FEE, CTU in Prague.

³CIIRC, CTU in Prague

Abstract

Most consumer cameras are equipped with electronic rolling shutter, leading to image distortions when the camera moves during image capture. We explore a surprisingly simple camera configuration that makes it possible to undo the rolling shutter distortion: two cameras mounted to have different rolling shutter directions. Such a setup is easy and cheap to build and it possesses the geometric constraints needed to correct rolling shutter distortion using only a sparse set of point correspondences between the two images. We derive equations that describe the underlying geometry for general and special motions and present an efficient method for finding their solutions. Our synthetic and real experiments demonstrate that our approach is able to remove large rolling shutter distortions of all types without relying on any specific scene structure.

1. Introduction

Thanks to low price, superior resolution and higher frame rate, CMOS cameras equipped with rolling shutter (RS) dominate the market for consumer cameras, smartphones, and many other applications. Unlike global shutter (GS) cameras, RS cameras read out the sensor line by line [21]. Every image line is captured at a different time, causing distortions when the camera moves during the image capture. The distorted images not only look unnatural, but are also unsuitable for conventional vision algorithms developed for synchronous perspective projection [13, 3, 28].

There are two main approaches to remove RS distortion. The first is to estimate the distortion and remove it, i.e., synthesize an image with global shutter geometry that can be




Figure 1: When two images are recorded with different rolling shutter directions, their motion-induced distortion is different, and a few point correspondences are enough to recover the motion as well as an undistorted image.

fed to standard vision algorithms [10, 26, 25, 32, 18]. The second approach is to keep the original images and adapt the algorithms to include the RS in the camera model [14, 13, 3, 28, 9]. The latter approach has recently led to RS-aware algorithms for many parts of the 3D vision pipeline, including RS camera calibration [23], RS structure-from-motion reconstruction [13], dense multi-view RS stereo [28], and RS absolute camera pose [1, 20, 3, 5, 29, 16]. Two-view geometry of RS cameras has been studied in [9] and triangulation with a RS stereo rig is discussed in [2].

Recently, more emphasis has been put on explicitly removing RS distortion from the images: in this way, one not only obtains visually appealing images, but can also continue to use the whole ensemble of existing, efficient vision algorithms. For the case of pure camera rotation, the distortion has been modelled as a mixture of homographies [10]. If only a single image is available, some external constraint is needed to recover the distortion, for instance [26] assume a Manhattan world and search for lines, respectively vanishing points. And [18] relaxes the Manhattan assumption and only requires the (curved) images of straight 3D lines to undistort a single RS image, including a RANSAC scheme to filter out false line matches. In [32] an occlusion-aware undistortion method is developed for the specific setting of ≥ 3 RS images with continuous motion and known time delay, assuming a piece-wise planar 3D scene. Others propose learning-based methods where a CNN is trained to warp single RS images to their perspective counterparts [25, 34].

Motivation. Despite more than a decade of research, RS distortion remains a challenge. In fact, when presented with

²FEE - Faculty of Electrical Engineering, Czech Technical University in Prague, ³CIIRC - Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague. TP and MP were supported by the European Regional Development Fund (IMPACT CZ.02.1.01/0.0/0.0/15_003/0000468) and Horizon 2020 projects 856994 and 871245. ZK was supported by OP RDE project International Mobility of Researchers MSCA-IF at CTU Reg. No. CZ.02.2.69/0.0/0.0/17_050/0008025 and OP VVV project Research Center for Informatics Reg. No. CZ.02.1.01/0.0/0.0/16_019/0000765. VL was supported by an ETH Postdoctoral Fellowship. We thank Alexander Wolf for building the camera rig and Nico Lang for help with data acquisition.

a single RS image, it is impossible to remove the distortion unless one makes fairly restrictive assumptions about the scene [26, 18], or about the camera motion [24], (the best-understood case being pure rotation [10, 25]). The same holds true for learned undistortion [25, 34], which only works for the types of scenes it has been trained on. Moreover, it does not guarantee a geometrically correct output that downstream processing steps can digest.

Equations for the generalized epipolar geometry of RS cameras have been derived in [9], however, due to the complexity of the resulting systems there is no practical solution for the RS epipolar geometry. The method of [2] utilizes triangulation and therefore requires non-negligible baseline. Furthermore, their solution is iterative, non-minimal and therefore not suitable for RANSAC-style robust estimation.

Even with multiple views, removing RS distortion either requires strong assumptions, like a piece-wise planar scene observed at a high frame-rate with known shutter timings [18]; or it amounts to full SfM reconstruction [13], thus requiring sufficient camera translation. Moreover, SfM with rolling shutter suffers from a number of degeneracies, in particular it has long been known that constant translational motion along the baseline (e.g., side-looking camera on a vehicle) does not admit a solution [2]. More recently, it has been shown [6] that (nearly) parallel RS read-out directions are in general degenerate, and can only be solved with additional constraints on the camera motion [15].

RS cameras are nowadays often combined into multi-camera systems. Even in mid-range smartphones it is common to have two forward-facing cameras, usually mounted very close to each other. It seems natural to ask whether such an assembly allows one to remove RS distortion.

Contribution. We show that a simple modification of the two-camera setup is sufficient to facilitate removal of RS distortion: the cameras must have known baseline – ideally of negligible length, as in the typical case of smartphones (for scene depth >1 m); and their RS read-out directions must be significantly different – ideally opposite to each other. Finally, the cameras must be synchronized (the offset between their triggers must be known). If those conditions are met, the motion of the cameras can be recovered and a perspective (global shutter) image can be approximated from the two RS images, regardless of image content. The only requirement is enough texture to extract interest points.

We also show that if the cameras undergo translational motion, depth maps can be computed given optical flow between the images, and that the undistorted sparse features can be used in an SfM pipeline to obtain a reconstruction similar to one from GS images.

In the following, we investigate the geometry of this configuration and propose algorithms to compute the motion parameters needed to remove the RS distortions. We first discuss the solution for the general case of unconstrained

6DOF motion, and develop an efficient solver that can be used in RANSAC-type outlier rejection schemes. We then go on to derive simpler solutions for special motion patterns that often arise in practice, namely pure rotation, pure translation, and translation orthogonal to the viewing direction. These cases are important because in real applications the motion is often restricted or known, e.g., when the camera is mounted on the side of a vehicle. The proposed methods require only point-to-point correspondences between the images to compute the camera motion parameters.

2. Problem formulation

Throughout, we assume two images taken by (perspective) cameras with rolling shutter. In those images, corresponding points have been found such that $\mathbf{u}_i = [u_i \ v_i \ 1]^T$ in the first camera corresponds to $\mathbf{u}'_i = [u'_i \ v'_i \ 1]^T$ in the second camera, and both are projections of the same 3D point $\mathbf{X}_i = [X_{i1}, X_{i2}, X_{i3}, 1]^T$. If neither the cameras nor the objects in the scene move, we can model the projections in both cameras as if they had a global shutter, with projection matrices P_g and P'_g [12] such that

$$\begin{aligned} \lambda_g \mathbf{u}_{gi} &= P_g \mathbf{X}_i &= K [R \mid -RC] \mathbf{X}_i \\ \lambda'_g \mathbf{u}'_{gi} &= P'_g \mathbf{X}_i &= K' [R' \mid -R'C'] \mathbf{X}_i. \end{aligned} \quad (1)$$

where K, K' are the camera intrinsics, R, R' are the camera rotations, C, C' are the camera centers, and λ_g, λ'_g are the scalar perspective depths.

If the cameras move during image capture, we have to generalize the projection function. Several motion models were used to describe the RS geometry [21, 2, 20, 4, 29]. These works have shown that, in most applications, assuming constant rotational and translational velocities¹ during RS read-out is sufficient to obtain the initial motion estimate, which can be further improved using a more complex model [27] in the refinement step. The projection matrices $P(v_i)$ and $P'(v'_i)$ are now functions of the image row, because each row is taken at a different time and hence a different camera pose. We can write the projections as

$$\begin{aligned} \lambda \mathbf{u}_i &= P(v_i) \mathbf{X}_i = K [R_\omega(v_i) R \mid -RC + \tau v_i] \mathbf{X}_i \\ \lambda' \mathbf{u}'_i &= P'(v'_i) \mathbf{X}_i = K' [R'_{\omega'}(v'_i) R' \mid -R'C' + \tau' v'_i] \mathbf{X}_i, \end{aligned} \quad (2)$$

where $R_\omega(v_i), R'_{\omega'}(v'_i)$ are the rotational and τ, τ' the translational motions during image acquisition.

Let us now consider the case when the relative pose of the cameras is known and the baseline is negligibly small (relative to the scene depth). This situation cannot be handled by existing methods [2, 9], but is the one approximately realized in multi-camera smartphones: the typical distance between the cameras is ≈ 1 cm, which means that already at 1 m distance the base-to-height ratio is 100:1 and we can

¹For constant rotational velocity we have $R_\omega(\alpha) = \exp(\alpha[\omega]_\times)$.

safely assume $\mathbf{C} = \mathbf{C}'$ (analogous to using a homography between global shutter images). We point out that the approximation is independent of the focal length, *i.e.*, it holds equally well between a wide-angle view and a zoomed view.

For simplicity, we consider the cameras to be calibrated, $\mathbf{K} = \mathbf{K}' = \mathbf{I}$, and attach the world coordinate system to the first camera, $\mathbf{R} = \mathbf{I}$ and $\mathbf{C} = \mathbf{C}' = \mathbf{0}$, to obtain

$$\begin{aligned} \lambda \mathbf{u}_i &= [\mathbf{R}_\omega(v_i) \mid \mathbf{t} v_i] \mathbf{X}_i \\ \lambda' \mathbf{u}'_i &= [\mathbf{R}'_{\omega'}(v'_i) \mathbf{R}_r \mid \mathbf{t}' v'_i] \mathbf{X}_i, \end{aligned} \quad (3)$$

where \mathbf{R}_r is now the relative orientation of the second camera w.r.t. the first one. Since the cameras are assembled in a rigid rig, their motions are always identical and we have $\mathbf{t}' = \mathbf{R}_r \mathbf{t}$ and $\mathbf{R}'_{\omega'}(v'_i) = \mathbf{R}_r \mathbf{R}_\omega(v_i) \mathbf{R}_r^\top$. This yields even simpler equations in terms of the number of unknowns,

$$\begin{aligned} \lambda \mathbf{u}_i &= [\mathbf{R}_\omega(v_i) \mid \mathbf{t} v_i] \mathbf{X}_i \\ \lambda' \mathbf{u}'_i &= [\mathbf{R}_r \mathbf{R}_\omega(v_i) \mid \mathbf{R}_r \mathbf{t} v_i] \mathbf{X}_i. \end{aligned} \quad (4)$$

From eq. (4) one immediately sees that the further \mathbf{R}_r is from identity, the bigger is the difference between the images and between their RS distortions. Since these differences are our source of information, we want them to be as large as possible and focus on the case of 180° rotation around the z -axis, $\mathbf{R}_r = \text{diag}([-1, -1, 1, \dots])$. In this setting the second camera is upside down and its RS read-out direction is opposite to that of the first one. Note that this is equivalent to $\mathbf{R}_r = \mathbf{I}$ with the shutter direction reversed. If we flip the second image along the x - and y -axes, it will be identical to the first one in the absence of motion, but if the rig moves the RS distortions will be different. This setup is easy to construct in practice: the only change to the standard layout is to reverse the read-out direction of one camera.

Also note that it is fairly straightforward to extend the algorithms derived below to shutters that are not 180° opposite, as well as to non-zero baseline with known relative orientation, as in [2] (e.g. stereo cameras on autonomous robots). In all these scenarios different RS directions make it possible to remove the distortion.

3. Solutions

In this section we will describe how to solve the problem for typical types of motions.

3.1. General 6DOF motion

To solve for the general motion case, we start from eq. (4). Without loss of generality, we can choose the first camera as the origin of the world coordinate system, such that

$$\begin{aligned} \mathbf{P} &= [\mathbf{I} \mid \mathbf{0}] \\ \mathbf{P}'(v_i, v'_i) &= [\mathbf{R}(v_i, v'_i) \mid v'_i \mathbf{R}_r \mathbf{t} - v_i \mathbf{R}(v_i, v'_i) \mathbf{t}] \end{aligned} \quad (5)$$

where $\mathbf{R}(v_i, v'_i) = \mathbf{R}_r \mathbf{R}_\omega(v'_i) \mathbf{R}_\omega(v_i)^\top$. We can consider $\mathbf{R}(v_i, v'_i)$ and $\mathbf{t}(v_i, v'_i) = v'_i \mathbf{R}_r \mathbf{t} - v_i \mathbf{R}(v_i, v'_i) \mathbf{t}$ as the relative camera orientation and translation for each pair of lines.

This yields one essential matrix

$$\mathbf{E}(v_i, v'_i) = [\mathbf{t}(v_i, v'_i)]_\times \mathbf{R}(v_i, v'_i), \quad (6)$$

for each pair of lines, with six unknowns. The translation can only be determined up to a unknown scale, using 5 correspondences. We next describe how to simplify the equations further and produce an efficient minimal solver.

3.2. Minimal solver for the 6DOF motion

Since both cameras form a rig, the two rotations $\mathbf{R}_\omega(v'_i)$, $\mathbf{R}_\omega(v_i)$ have the same axis and we have

$$\mathbf{R}_\omega(v'_i) \mathbf{R}_\omega(v_i)^\top = \mathbf{R}_\omega(v'_i - v_i). \quad (7)$$

For convenience, let $\mathbf{R}_i = \mathbf{R}_r \mathbf{R}_\omega(v'_i - v_i)$. Then the instantaneous essential matrix for rows v_i and v'_i can be written

$$\begin{aligned} \mathbf{E}(v_i, v'_i) &= [v'_i \mathbf{R}_r \mathbf{t} - v_i \mathbf{R}_i \mathbf{t}]_\times \mathbf{R}_i = \\ &= v'_i [\mathbf{R}_r \mathbf{t}]_\times \mathbf{R}_i - v_i [\mathbf{R}_i \mathbf{t}]_\times \mathbf{R}_i = v'_i [\mathbf{R}_r \mathbf{t}]_\times \mathbf{R}_i - v_i \mathbf{R}_i [\mathbf{t}]_\times \end{aligned} \quad (8)$$

using the identity $(\mathbf{R} \mathbf{u} \times \mathbf{R} \mathbf{v}) = \mathbf{R}(\mathbf{u} \times \mathbf{v})$. As the motion due to the RS effect is small, we linearise it, as often done for RS processing, *e.g.*, [1, 20, 3, 16]. For constant angular velocity we get

$$\mathbf{R}_i \approx \mathbf{R}_r (\mathbf{I}_{3 \times 3} + (v'_i - v_i) [\omega]_\times), \quad (9)$$

where the direction of ω encodes the axis of rotation, and the angular velocity determines its magnitude. Inserting this into (8) we get

$$\begin{aligned} \mathbf{E}(v_i, v'_i) &= v'_i \mathbf{R}_r [\mathbf{t}]_\times (\mathbf{I} + (v'_i - v_i) [\omega]_\times) \\ &\quad - v_i \mathbf{R}_r (\mathbf{I} + (v'_i - v_i) [\omega]_\times) [\mathbf{t}]_\times \end{aligned} \quad (10)$$

Each pair of corresponding points $(\mathbf{u}_i, \mathbf{u}'_i)$ now yields a single equation from the epipolar constraint,

$$\mathbf{u}'_i{}^\top \mathbf{E}(v_i, v'_i) \mathbf{u}_i = 0. \quad (11)$$

This is a quadratic equation in the unknowns ω and \mathbf{t} . Since the scale of the translation is unobservable (eq. (11) is homogeneous in \mathbf{t}), we add a linear constraint $\mathbf{t}_1 + \mathbf{t}_2 = 1$, leading to the parameterisation $\mathbf{t} = (1 - x, x, y)^\top$. Note that this constraint is degenerate for pure forward motion.

From 5 point correspondences we get 5 quadratic equations in 5 unknowns x, y and ω , which we solve with the hidden variable technique [8]. We rewrite the equations as

$$\mathbf{M}(\omega) \begin{bmatrix} x & y & 1 \end{bmatrix}^\top = 0, \quad (12)$$

where $\mathbf{M}(\omega)$ is a 5×3 matrix with elements that depend linearly on ω . This matrix must be rank deficient, thus all 3×3 sub-determinants must vanish, which gives 10 cubic equations in ω , with up to, in general, 10 solutions. Interestingly,

the equations have the same structure as the classic determinant and trace constraints for the essential matrix. To solve them one can thus employ any of the known solutions for that case [22, 11]. We use the solver generator [19].

In a similar fashion, we can produce a solution for the case of a fixed, known baseline between the cameras. Please see the supplementary material for details.

3.3. Pure rotation

Next, let us consider the case where the cameras only rotate around the center of projection. We now have $\mathbf{t} = 0$ and $\mathbf{R}_\omega(\alpha) \neq \mathbf{I}$ for $x \in \mathbb{R} \setminus 0$. Equations (4) become

$$\begin{aligned} \lambda \mathbf{u}_i &= [\mathbf{R}_\omega(v_i) \mid 0] \mathbf{X}_i \\ \lambda' \mathbf{u}'_i &= [\mathbf{R}_r \mathbf{R}_\omega(v'_i) \mid 0] \mathbf{X}_i. \end{aligned} \quad (13)$$

and we can express the relationship between \mathbf{u}_i and \mathbf{u}'_i as

$$\lambda' \mathbf{u}'_i = \mathbf{R}_\omega(v'_i) \mathbf{R}_r \mathbf{R}_\omega^\top(v_i) \lambda \mathbf{u}_i. \quad (14)$$

This resembles a homography between GS images, except that the matrix $\mathbf{H} = \mathbf{R}_\omega(v'_i) \mathbf{R}_r \mathbf{R}_\omega^\top(v_i)$ changes for every correspondence. To get rid of λ, λ' we divide (14) by λ , then left-multiply with the skew-symmetric $[\mathbf{u}'_i]_\times$ to obtain

$$0 = [\mathbf{u}'_i]_\times \mathbf{R}_\omega(v'_i) \mathbf{R}_r \mathbf{R}_\omega^\top(v_i) \mathbf{u}_i \quad (15)$$

For constant angular velocity we now have three unknown parameters for the rotation $\mathbf{R}_\omega(\alpha)$. Each correspondence yields two equations, so the solution can be found from 1.5 correspondences. If we further linearise $\mathbf{R}_\omega(\alpha)$ via first-order Taylor expansion, as in [1, 20, 3], we get

$$0 = [\mathbf{u}'_i]_\times (\mathbf{I} + v'_i [\boldsymbol{\omega}]_\times) \mathbf{R}_r (\mathbf{I} - v_i [\boldsymbol{\omega}]_\times) \mathbf{u}_i, \quad (16)$$

where $\boldsymbol{\omega}$ is the angle-axis vector. This is a system of three 2nd-order equations in three unknowns, which can be solved efficiently with the *e3q3* solver [17].

3.4. Translation

Next, let us consider a general translation with constant velocity and no rotation, $\mathbf{R}_\omega(\alpha) = \mathbf{I}$. Substituting $\mathbf{R}_\omega(\alpha) = \mathbf{I}$ in equations (4) and subtracting the second equation from the first one we obtain (for details see supplementary)

$$\begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} v_i - \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} v'_i = \begin{bmatrix} u_i \\ v_i \\ 1 \end{bmatrix} \lambda_i - \begin{bmatrix} -u'_i \\ -v'_i \\ 1 \end{bmatrix} \lambda'_i. \quad (17)$$

Each correspondence adds three equations and two unknowns λ_i, λ'_i . Two correspondences give us 6 linear homogeneous equations in 7 unknowns $t_x, t_y, t_z, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2$, which allow us to find a solution up to scale, *i.e.*, relative to one depth (e.g., λ_1) or to the magnitude of the translation.

Translation in the xy -plane We also consider the case of translation orthogonal to the viewing direction, $\mathbf{t} =$

$[t_x \ t_y \ 0]$. The system (17) now lacks t_z in the 3rd equation and we find that $\lambda_i = \lambda'_i$, *i.e.* the perspective depth of a point \mathbf{X}_i is the same in both cameras. By solving this system for t_x and t_y , we can express them in terms of λ_i ,

$$t_x = \frac{u_i + u'_i}{v_i - v'_i} \lambda_i, \quad t_y = \frac{v_i + v'_i}{v_i - v'_i} \lambda_i, \quad (18)$$

and obtain an equivalent global shutter projection as

$$\mathbf{u}_{gi} = \left[\frac{u_i v'_i - u'_i v_i}{v_i - v'_i}, \frac{-2v_i v'_i}{v_i - v'_i}, 1 \right]^\top. \quad (19)$$

Translation along x -axis Finally, let us assume a translation only along the camera x -axis, such as for a side-looking camera on a moving vehicle, or when observing passing cars. In this case the global shutter projection satisfies $\mathbf{u}_{gi} = \left[\frac{u_i + u'_i}{2}, v_i, 1 \right]^\top$ (see supplementary for a detailed derivation), which implies that for constant velocity along the x -axis we can obtain GS projections by simply interpolating between the x -coordinates of corresponding points in the two RS images. Analogously we interpolate the y -coordinate for translation along y -axis.

4. Refinement with advanced motion model

Our minimal solutions resort to simplified motion models to support real-time applications and RANSAC. After obtaining an initial solution with one of the minimal solvers, we can improve the motion estimates through a non-linear refinement with a more complex model of camera motion.

For the rotation case, we can obtain the cost function from eq. (14) and sum the residuals of all correspondences

$$\sum_{i=0}^N \left\| \begin{bmatrix} u_i \\ v_i \end{bmatrix} - \begin{bmatrix} \mathbf{h}_1^i \mathbf{u}'_i / \mathbf{h}_3^i \mathbf{u}'_i \\ \mathbf{h}_2^i \mathbf{u}'_i / \mathbf{h}_3^i \mathbf{u}'_i \end{bmatrix} \right\|, \quad (20)$$

where $\mathbf{H}_i = [\mathbf{h}_1^i \mathbf{h}_2^i \mathbf{h}_3^i]^\top = \mathbf{R}_\omega(v'_i) \mathbf{R}_r \mathbf{R}_\omega^\top(v_i)$, and $\mathbf{R}_\omega(\alpha)$ is now parametrised via the Rodrigues formula.

For the translation case, we can minimise the Sampson error, as in [9], which leads to the cost function

$$\sum_{i=0}^N \frac{(\mathbf{u}_i^\top \mathbf{E}_i \mathbf{u}_i)^2}{(\mathbf{E}_i \mathbf{u}_i)_1^2 + (\mathbf{E}_i \mathbf{u}_i)_2^2 + (\mathbf{E}_i^\top \mathbf{u}'_i)_1^2 + (\mathbf{E}_i \mathbf{u}'_i)_2^2} \quad (21)$$

with $\mathbf{E}_i = v'_i [\mathbf{R}_r \mathbf{t}]_\times \mathbf{R}_i - v_i \mathbf{R}_i [\mathbf{t}]_\times$ as defined in equation 8. Again \mathbf{R}_i is defined via the Rodrigues formula.

It has been shown [27] that a uniform motion model across the entire image may not be sufficient, instead using three different motions defined at their respective "knots" worked better for handheld footage. If desired, this extension is also applicable in our refinement step. To that end, we simply define intermediate poses for the camera system by taking the initial parameters $\boldsymbol{\omega}, \mathbf{t}$ from the minimal solver, and then minimize either (20) or (21), with the instantaneous rotation and translation of the cameras interpolated between those intermediate poses.

5. Undistorting the image

Once the motion parameters have been estimated, we can choose between two approaches to undistort the images, based on whether we treat pure rotation or also translation.

Rotation and image warping. Under pure rotation, creating an image in global shutter geometry is simple. For each pixel, we have a forward mapping $\lambda \mathbf{u}_{gi} = \mathbf{R}_\omega(v_i)^\top \mathbf{u}_i$ from the first image and $\lambda \mathbf{u}'_{gi} = \mathbf{R}_\omega(v'_i)^\top \mathbf{u}'_i$ from the second image into the virtual GS image plane. This mapping depends only on the row index of the respective RS image, which defines the time and therefore the pose. No pixel-wise correspondences are needed. We can also use the backward mapping $\lambda \mathbf{u}_i = \mathbf{R}_\omega(v_i) \mathbf{u}_{gi}$, in that case v_i appears on both sides of the equation, possibly in non-linear form, depending on how $\mathbf{R}_\omega(v_i)$ is parameterised. In either case, we can iteratively solve for v_i , starting at the initial value v_{gi} . One can transform both RS inputs to the same virtual GS image for a more complete result, see supplementary.

Translation and dense undistortion. Translational RS distortion poses a greater challenge, because the transformation of each pixel depends on its correspondence in the other RS image. Consequently, one must recover dense correspondence to obtain a dense, visually pleasing result [26, 33, 32]. Empirically, we found optical flow methods to work best, particularly PWC-net [31] consistently gave the best results. Given dense correspondences, we can convert them to a depth map and use it to transfer image content from the RS images to the virtual GS image, using some form of z -buffering to handle occlusions. Note that rows near the middle of the image have been captured at (nearly) the same time and pose in both RS views, thus depth is not observable there. Fortunately, this barely impacts the undistortion process, because also the local RS distortions are small in that narrow strip, such that they can be removed by simply interpolating between the inputs. For details, see the supplementary material.

6. Experiments

We have tested the proposed algorithms both with synthetic and with real data. Synthetic correspondences serve to quantitatively analyze performance. Real image pairs were acquired with a rig of two RS cameras, see Fig. 2, and undistorted into global shutter geometry to visually illustrate the high quality of the corrected images across a range of motion patterns and scenes.

Synthetic data. To generate synthetic data in a realistic manner, we started from the GS images of [13] and performed SfM reconstruction. We then placed virtual RS pairs at the reconstructed camera locations, to which we applied various simulated motions with constant translational and angular velocities, and reprojected the previously re-

constructed 3D structure points. The angular velocity was gradually increased up to 30 degrees/frame which means that the camera has rotated by 30 degrees during the acquisition of one frame. The translational velocity was increased up to 1/10 of the distance of the closest part of the scene per frame. Gaussian noise with $\mu = 0$ and $\sigma = 0.5pix$ was added to the coordinates of the resulting correspondences. We generated around 1.4K images this way.

We test four different solvers, two for pure translations, one for pure rotation, and one for the full 6DOF case, see Fig. 2. Additionally, we run a simple baseline that just averages the image coordinates of the two corresponding points (*err-interp*). We do not consider the single axis translation solvers \mathbf{tx} and \mathbf{ty} , since they are covered by \mathbf{txy} , which also requires only a single correspondence.

Three different variants are tested for each solver: (*v1*) individually fits a motion model at each correspondence, sampling the minimal required set of additional correspondences at random and undistorting the individual coordinate. This local strategy can handle even complex global distortions with a simpler motion model, by piecewise approximation. The downside is that there is no redundancy, hence no robustness against mismatches. (*v2*) is a robust approach that computes a single, global motion model for the entire image and uses it to undistort all correspondences. The solver is run with randomly sampled minimal sets, embedded in a LO-RANSAC loop [7] that verifies the solution against the full set of correspondences and locally optimizes the motion parameters with non-linear least squares. (*v3*) explores a hybrid LO-RANSAC method that uses one of the simpler models to generate an initial motion estimate, but refines it with the full model with parameters $\{t_x, t_y, t_z, \omega_x, \omega_y, \omega_z\}$.

The results provide some interesting insights, see Fig. 3. As expected, the full $\omega \mathbf{t}$ performs best, followed by the rotation-only model ω . The translational solvers, including the simple \mathbf{txy} , work well when used locally per correspondence (*v1*), moreover the error has low variance. This means that the rotational distortion component can be well approximated by piece-wise translations \mathbf{txy} , whose estimation is more reliable than that of both \mathbf{txyz} and ω .

With a single, global RANSAC fit (*v2*) the residual errors of the undistorted points are generally higher (except for $\omega \mathbf{t}$), due to the more rigid global constraint. The drop is strongest for \mathbf{txy} and \mathbf{txyz} , i.e., a global translation model cannot fully compensate rotational distortions. The hybrid solution (*v3*) is not shown since it does not improve over the global one (*v2*), suggesting that the general model gets stuck in local optima when initialised with a restricted solver.

Zero-baseline assumption. Next, we evaluate the approximation error due to our assumption of negligible baseline between the two RS cameras. We start with a synthetic experiment, with the same setup as in Fig. 3, but this time with



Alg.	# corr.	Param.	Runtime
t_{xy}	1	t_x, t_y	≈ 0
t_{xyz}	2	t_x, t_y, t_z	≈ 0
ω	2	$\omega_x, \omega_y, \omega_z$	$4 \mu s$
ωt	5	$t_x, t_y, t_z, \omega_x, \omega_y, \omega_z$	$40 \mu s$

Figure 2: Rig with two RS and one GS camera (left); Solvers used in the experiments (right).

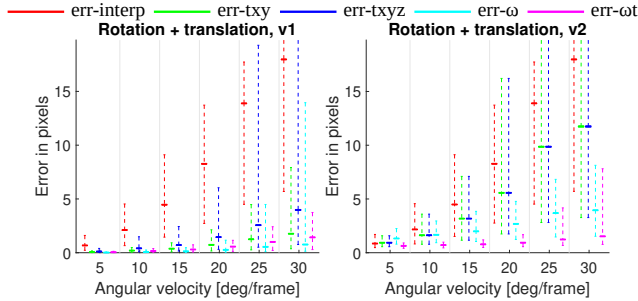


Figure 3: Results on synthetic RS pairs. Fitting a separate motion model per correspondence (left), and fitting a single model per image (right). Plots show the error of undistorted correspondences w.r.t. the ground truth GS image.

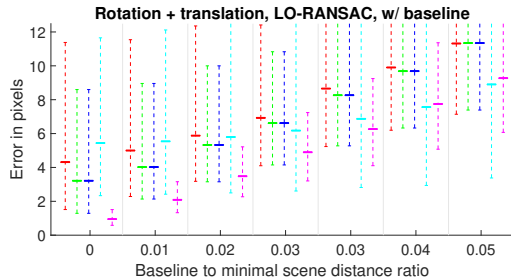


Figure 4: The impact of baselines >0 is negligible up to base-to-depth ratio $\approx 1:100$, and remains low up to $\approx 1:30$.

baselines $\neq 0$. We use the global model fit (v2), with angular velocity $15^\circ/\text{frame}$. The baseline was increased from 0 to 5% of the smallest scene depth. As shown in Fig. 4, the zero-baseline assumption has negligible effect for base-to-depth ratios up to 1:100, i.e., at a typical smartphone baseline of at most 2 cm for objects ≥ 2 m from the camera. Even at an extreme value of 1:20 (40 cm from the camera) the approximation error remains below 10 pixels.

A further experiment with real data supports the claim that the zero baseline assumption is insignificant in the case of rotational motion, unless most correspondences are on the closest scene parts. Fig. 7 shows rotational distortion removal with objects as close as $10\times$ the baseline. As long as enough correspondences are found also on more distant scene parts, RANSAC chooses those and finds a mapping that is valid also for close objects. For translational motion, the closest correspondences carry more information about the translation than the distant ones. In our experiments the scenes always contained enough correspondences close

enough to estimate the motion, but far enough to neglect the baseline. For scenes with mostly low depth (relative to the baseline) we recommend the 6pt fixed-baseline solver, described in the supplementary material.

Real images. We built a rig with two RS cameras, mounted ≈ 3 cm apart and pointing in the same direction, see Fig. 2. Their rolling shutters run in opposite directions, with ≈ 30 ms readout time for a complete frame, a typical value for consumer devices. The image resolution is 3072×2048 pix. Additionally, we added a GS camera to the rig with resolution 1440×1080 pix (despite the weaker specs, that camera cost almost as much as the two RS cameras together). All cameras are triggered synchronously.

We captured images of various scenes, with significant RS distortions from a variety of different camera motions. The angular velocity in the rotation experiments was between 8 and 15 degrees per frame or $240\text{-}450^\circ/\text{s}$. In the translation experiments, the car was moving with 30-50 km/h and, since the camera was hand-held, there was also a non-negligible rotational component. The correspondences between images were either matched SIFT features, or taken from the optical flow in the case of translational motion, see Sec. 5. Although the proposed camera configuration is new and there are no other algorithms suited to handle such data, it is interesting to see the results of existing RS undistortion algorithms. See examples in Fig. 9 for rotation and Fig. 10 for translation, where we compare our undistorted GS images with those of most recent competing methods [26, 25, 18, 32, 33].

We used RANSAC with a fixed number of 200 iterations, which proved to be enough due to the small minimal set for the ω solver, respectively the dense optical flow with low outlier fraction for the ωt solver. Note that we compare each method only in relevant scenarios, e.g., [25, 18] work under the assumption of pure rotation and therefore are not able to handle translation properly; [33] requires a baseline and thus handles only translation; [26, 32, 34] should be able to handle both cases, the results of [32, 25] were unsatisfactory for rotation, so we do not present them.

Compared to existing methods, our results demonstrate robustness to the motion type as well as the scene content. The proposed camera setup allows us to correct various types of distortion with small residual error compared to the corresponding GS image. For rotation [18] in some cases provides satisfactory results (Fig. 9, rows 1 and 3), but it fails when there are too few straight lines in the scene (row 6). [26] almost always fails and [34], although trained on 6DOF motion data, only produced usable result in rows 3, 5 and 6 with rotation around the y -axis.

In Fig. 1 we show a sample with very strong, real RS effect. Even in this situation our method produces a near-perfect GS image, whereas competing methods fail. Furthermore, in Fig. 8 we demonstrate that even using a sub-

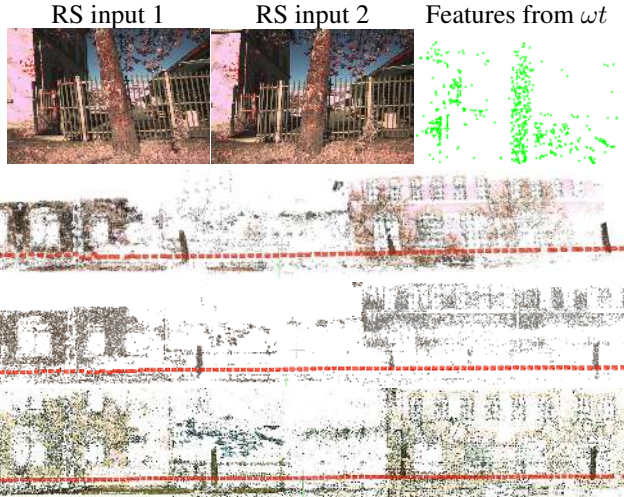


Figure 5: GS-equivalent sparse features from RS images. The corrected features can be further used, e.g. feeding them to an SfM pipeline yields a better reconstruction (middle) than feeding in the raw RS features (top). At the bottom is the reconstruction from a real GS camera.




Figure 6: Depth maps. The top row shows the two input images (left) and the resulting undistorted image (right). The bottom row shows the depth maps created from both input images (left) and the final fused depth map (right).

window of the second image, one is able to recover the motion correctly and undistort the full image. This suggests that a standard smartphone configuration with a wide-angle and a zoom camera can be handled.

For translational motion, undistorting the entire image is in general a hard problem [32], as it requires pixel-wise correspondences as well as occlusion handling, c.f. Sec. 5 and supplementary material. The results in Fig. 10 show that with our method translation distortion, including occlusions, can be compensated as long as good pixel-wise correspondences can be obtained. On the contrary, [26] struggled to compensate the building (row 2) and the tree (row 3); [33] works in some situations (row 1) and fails in others (row 2 and 3); [32] often provides good results (row 1 and 3), but sometimes compensates the motion only partially, and also exhibits post-processing artefacts (row 2). We also tried the recent method [34], but the authors do not provide



Figure 7: Example of rotation undistortion (ω solver) for close-range scenes. The distance to the nearest scene points is $\approx 10\times$ the baseline.



Figure 8: Example of rotation undistortion (ω solver). Correspondences are displayed in red. One camera has narrower FOV. Although parts of the wide-angle view have no correspondences, they are undistorted correctly.

the code or the trained model, and our re-implementation trained on 6DOF data provided worse results than all other methods in all cases, so we do not show them. Note that [33, 32] require two, respectively three consecutive frames for which the exact (relative) trigger times as well as the exact shutter read-out speed must be known.

We show further outputs that can be obtained with our method, besides undistorted images. One interesting possibility is to output undistorted (sparse) feature points, which can be fed into an unmodified SfM pipeline. Figure 5 shows an example of SIFT features extracted from RS images, recorded from a car in motion and corrected with the model ωt . Feature point on both the background and the tree were successfully included as inliers and undistorted. Figure 5 (top row) shows the result of sparse SfM reconstruction with COLMAP [30]. One can clearly see the difference between using corrected or uncorrected feature points, especially on the trees in the foreground.

As an intermediate product during translation correction, we obtain depth maps, see Fig. 6. While the camera rig undergoes translation, we obtain depth with lower latency than from consecutive frames, since we are using the inter-row baseline rather than the inter-frame baseline. The price to pay is that our stereo baseline diminishes towards the central image rows, such that depth estimates become less accurate (and are impossible for the exact middle row).

Choice of solver. The ω solver is well suited for slow-moving (e.g., hand-held) devices where translation is insignificant, and for distant scenes with shallow depth range (e.g., drones). The t_{xy} solver suits scenarios with pure, fast translation (e.g., side-looking cameras on straight roads or rails), and can in the presence of small rotations still be used, trading some accuracy for higher robustness. For fast, general motion we found the 6DOF solver to perform significantly better, as expected.




Figure 9: Camera undergoing a rotational motion. Significant RS distortion can be successfully removed using the proposed camera setup. Competing methods for RS image correction [26, 25, 18] provide visibly worse results.



Figure 10: Camera undergoing translational motion. Images are undistorted pixel-wise, using optical flow obtained with [31].

Another possible scenario which we unfortunately did not manage to test is the case where the camera rig stands still and there are objects moving in the scene, as in surveillance cameras. There, especially the t_{xy} solver could provide a fast and efficient RS correction and depth estimation.

7. Conclusion

We present a novel configuration of two RS cameras that is easy to realise in practice, and makes it possible to accurately remove significant RS distortions from the images:

by mounting two cameras close to each other and letting the shutters roll in opposite directions, one obtains different distortion patterns. Based on those, we derive algorithms to compute the motion of the two-camera rig and undistort the images. Using the corrected geometry, we can perform SfM and compute depth maps equivalent to a GS camera. Our derivations show that similar constructions are in principle also possible when there is a significant baseline between the cameras. Hence, conventional stereo rigs, for instance on robots or vehicles, could in the future also benefit from opposite shutter directions.

References

- [1] Omar Ait-aider, Nicolas Andreff, Jean Marc Lavest, and Philippe Martinet. Simultaneous object pose and velocity computation using a single view from a rolling shutter camera. In *ECCV*, pages 56–68, 2006. 1, 3, 4
- [2] O. Ait-Aider and F. Berry. Structure and kinematics triangulation with a rolling shutter stereo rig. In *ICCV*, 2009. 1, 2, 3
- [3] Cenek Albl, Zuzana Kukelova, Viktor Larsson, and Tomas Pajdla. Rolling Shutter Camera Absolute Pose. *PAMI*, 2019. 1, 3, 4
- [4] Cenek Albl, Zuzana Kukelova, and Tomas Pajdla. R6p - rolling shutter absolute pose problem. In *CVPR*, 2015. 2
- [5] Cenek Albl, Zuzana Kukelova, and Tomas Pajdla. Rolling shutter absolute pose problem with known vertical direction. In *CVPR*, 2016. 1
- [6] Cenek Albl, Akihiro Sugimoto, and Tomas Pajdla. Degen-eracies in rolling shutter sfm. In *ECCV*, 2016. 2
- [7] Ondřej Chum, Jifí Matas, and Josef Kittler. Locally Opti-mized RANSAC. In *DAGM*. 2003. 5
- [8] David A. Cox, John Little, and Donal O’Shea. *Using Alge-braic Geometry*. Graduate Texts in Mathematics. Springer-Verlag, 2005. 3
- [9] Yuchao Dai, Hongdong Li, and Laurent Kneip. Rolling shut-ter camera relative pose: Generalized epipolar geometry. In *CVPR*, 2016. 1, 2, 4
- [10] M. Grundmann, V. Kwatra, D. Castro, and I. Essa. Calibration-free rolling shutter removal. In *ICCP*, 2012. 1, 2
- [11] R. Hartley and Hongdong Li. An efficient hidden variable approach to minimal-case camera motion estimation. *PAMI*, 34(12), 2012. 4
- [12] Richard Hartley and Andrew Zisserman. *Multiple view ge-ometry in computer vision*. Cambridge university press, 2003. 2
- [13] Johan Hedborg, Per-Erik Forssén, Michael Felsberg, and Erik Ringaby. Rolling shutter bundle adjustment. In *CVPR*, 2012. 1, 2, 5
- [14] Johan Hedborg, Erik Ringaby, Per-Erik Forssen, and Michael Felsberg. Structure and motion estimation from rolling shutter video. In *ICCV Workshops*, 2011. 1
- [15] Eisuke Ito and Takayuki Okatani. Self-Calibration-Based Approach to Critical Motion Sequences of Rolling-Shutter Structure From Motion. In *CVPR*, 2017. 2
- [16] Zuzana Kukelova, Cenek Albl, Akihiro Sugimoto, and Tomas Pajdla. Linear solution to the minimal absolute pose rolling shutter problem. In *ACCV*, 2018. 1, 3
- [17] Zuzana Kukelova, Jan Heller, Martin Bujnak, Andrew Fitzgibbon, and Tomas Pajdla. Efficient solution to the epipolar geometry for radially distorted cameras. In *ICCV*, 2015. 4
- [18] Yizhen Lao and Omar Ait-Aider. A robust method for strong rolling shutter effects correction using lines with automatic feature selection. In *CVPR*, 2018. 1, 2, 6, 8
- [19] V. Larsson, Kalle Åström, and Magnus Oskarsson. Efficient Solvers for Minimal Problems by Syzygy-Based Reduction. In *CVPR*, 2017. 4
- [20] Ludovic Magerand, Adrien Bartoli, Omar Ait-Aider, and Daniel Pizarro. Global optimization of object pose and mo-tion from a single rolling shutter image with automatic 2d-3d matching. In *ECCV*, 2012. 1, 2, 3, 4
- [21] Marci Meingast, Christopher Geyer, and Shankar Sastry. Ge-ometric Models of Rolling-Shutter Cameras. *Computing Re-search Repository*, 2005. 1, 2
- [22] David Nistér. An efficient solution to the five-point relative pose problem. *PAMI*, 26(6), 2004. 4
- [23] Luc Oth, Paul Timothy Furgale, Laurent Kneip, and Roland Siegwart. Rolling Shutter Camera Calibration. In *CVPR*, 2013. 1
- [24] Pulak Purkait and Christopher Zach. Minimal solvers for monocular rolling shutter compensation under ackermann motion. In *WACV*, 2018. 2
- [25] Vijay Rengarajan, Yogesh Balaji, and A. N. Rajagopalan. Unrolling the Shutter: CNN to Correct Motion Distortions. In *CVPR*, 2017. 1, 2, 6, 8
- [26] Vijay Rengarajan, A. N. Rajagopalan, and Rangarajan Ar-avind. From Bows to Arrows: Rolling Shutter Rectification of Urban Scenes. In *CVPR*, 2016. 1, 2, 5, 6, 7, 8
- [27] Erik Ringaby and Per-Erik Forssén. Efficient Video Rectifi-cation and Stabilisation for Cell-Phones. *IJCV*, 96(3), 2012. 2, 4
- [28] Olivier Saurer, Kevin Köser, Jean-Yves Bouguet, and Marc Pollefeys. Rolling Shutter Stereo. In *ICCV*, 2013. 1
- [29] Olivier Saurer, Marc Pollefeys, and Gim Hee Lee. A minimal solution to the rolling shutter pose estimation problem. In *IROS*, 2015. 1, 2
- [30] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion Revisited. In *CVPR*, 2016. 7
- [31] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. In *CVPR*, 2017. 5, 8
- [32] Subeesh Vasu, Mahesh Mohan M. R., and A. N. Ra-jagopalan. Occlusion-Aware Rolling Shutter Rectification of 3d Scenes. In *CVPR*, 2018. 1, 5, 6, 7, 8
- [33] Bingbing Zhuang, Loong-Fah Cheong, and Gim Hee Lee. Rolling-Shutter-Aware Differential SfM and Image Rectifi-cation. In *ICCV*, 2017. 5, 6, 7, 8
- [34] Bingbing Zhuang, Quoc-Huy Tran, Pan Ji, Loong-Fah Cheong, and Manmohan Chandraker. Learning structure-and-motion-aware rolling shutter correction. In *CVPR*, 2019. 1, 2, 6, 7, 8