

Low-rank registration of images captured under unknown, varying lighting

Matthieu Pizenberg, Yvain Quéau, Abderrahim Elmoataz

► To cite this version:

Matthieu Pizenberg, Yvain Quéau, Abderrahim Elmoataz. Low-rank registration of images captured under unknown, varying lighting. International Conference on Scale Space and Variational Methods in Computer Vision (SSVM 2021), 2021, Cabourg, France. pp.153-164, 10.1007/978-3-030-75549-2_13. hal-03172399

HAL Id: hal-03172399 https://hal.science/hal-03172399

Submitted on 17 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Low-rank registration of images captured under unknown, varying lighting

Matthieu Pizenberg[⊠], Yvain Quéau, and Abderrahim Elmoataz

Normandie Univ, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France {matthieu.pizenberg,yvain.queau,abderrahim.elmoataz-billah}@unicaen.fr

Abstract. Photometric stereo infers the 3D-shape of a surface from a sequence of images captured under moving lighting and a static camera. However, in real-world scenarios the viewing angle may slightly vary, due to vibrations induced by the camera shutter, or when the camera is hand-held. In this paper, we put forward a low-rank affine registration technique for images captured under unknown, varying lighting. Optimization is carried out using convex relaxation and the alternating direction method of multipliers. The proposed method is shown to significantly improve 3D-reconstruction by photometric stereo on unaligned real-world data, and an open-source implementation is made available.

Keywords: Photometric stereo · registration · shape-from-X.

1 Introduction

Photometric stereo is a 3D-reconstruction technique pioneered by Woodham [17]. It infers the geometry of a surface from a set of images captured under the *same viewing angle*, but varying illumination. The camera is thus assumed to be perfectly still, an assumption which may be violated when the camera is handheld or when it is manually triggered (which induces slight vibrations of the sensor). Consequently, 3D-reconstruction by photometric stereo may be blurry (cf. Fig. 1), and geometric information may even be hallucinated (cf. Fig. 2). To prevent this, the image sequence could be registered prior to 3D-reconstruction.

We focus in this work on quasi-planar scenes, such that the geometric transformation between images can be assumed affine. Despite this simplifying assumption, registration remains arduous due to lighting variations, which may induce cast-shadows or specularities. In the present work, we propose an opensource¹, robust affine registration technique for images captured under unknown, varying lighting, based on low-rank approximation and convex relaxation.

The rest of this paper is organized as follows. After reviewing related works in Sect. 2, we introduce in Sect. 3 a variational approach to image registration under varying lighting. Then, we describe in Sect. 4 an augmented Lagrangian approach for numerically solving this inverse problem. The implementation of this scheme is discussed in Sect. 5, along with experimental evaluation. Our conclusions and perspectives are eventually drawn in Sect. 6.

¹ https://github.com/mpizenberg/lowrr

Matthieu Pizenberg et al.



Fig. 1: Top, left: three (out of m = 13) images of the Bayeux Tapestry, acquired under varying lighting, using a hand-held camera. Top, right: images registered using the proposed low-rank technique. Bottom: estimated reflectance and 3Dshape, using the original (left) or the registered (right) images. Registering images allows one to use a hand-held camera while obtaining a 3D-reconstruction comparable to that obtained using a high-quality tripod (cf. Fig. 3).

2 Related works

Low-rank techniques in photometric stereo Early works on photometric stereo have focused on the ideal case of a Lambertian surface lit by directional sources [17]. In practical setups, these assumptions are however rarely met: lowrank approximation techniques have thus been developed to solve photometric stereo under less restrictive assumptions. Based on earlier investigations on the linear image subspaces spanned by Lambertian surfaces [2], Basri et al. have shown in [3] how to handle general, unknown lighting using spherical harmonics approximation. Non-Lambertian phenomena such as specularities or cast shadows can also be viewed as sparse outliers, and removed from the photometric stereo images by seeking a low-rank observation map sparsely deviating from the input one [18]. Such low-rank approximation techniques for outliers handling also relate to sparsity-promoting estimators used in robust calibrated photometric stereo, see e.g. [11,13]. Low-rank approximation was also considered for solving the uncalibrated case by jointly minimizing the rank and enforcing integrability [15].



3

Fig. 2: Top, left: a photometric stereo sequence acquired using a tripod, while manually triggering the camera. Manual trigger induces slight image transformations, which are compensated with the proposed registration technique (top, right). Bottom: shape and reflectance estimated by photometric stereo, using either the unregistered (left) or the registered sequence (right). The misalignment causes the estimated reflectance to appear blurry, and induces geometry hallucination (the pictured surface is a perfectly planar book cover).

Image registration in photometric stereo All the works mentioned above assume that the input images have been registered beforehand. When this is not the case, the geometry recovered by using photometric stereo might appear blurry. In some sense, the effect would be similar to observing a translucent object where scattering is observed. In the latter case, the estimated geometry could be deconvolved a posteriori, as suggested in [12]. Yet, pre-registering the images seems more relevant in the absence of scattering. Registration is a long-standing problem in imaging, which has been addressed using e.g., phase correlation [8] or feature matching [9]. On the other hand, photometry-aware criteria have recently proved promising for camera pose optimization [14]. This

Matthieu Pizenberg et al.

invites us to design an image registration technique which is specifically tailored for photometric 3D-reconstruction. To the best of our knowledge, image registration under the photometric stereo perspective has been explored only in [4,10]. However, the former is based on feature detection, hence it may fail when the data lacks texture and geometry variation; and the latter is restricted to perfectly Lambertian surfaces and known illumination.

On the contrary, the low-rank approach discussed in the next section avoids feature detection; it does not require any pre-calibration; and it is robust to sparse deviations from the Lambertian model.

3 Low-rank registration of photometric stereo images

Basri and Jacob have shown that, if the surface is Lambertian, the observation matrix formed by stacking all the graylevels in an $n \times m$ matrix, where $n \ge m$ is the number of pixels and $m \ge 3$ is the number of images, should lie in a low-dimensional subspace [2].

To verify whether this result could be used for image registration, we considered a set of m = 13 real-world photometric stereo images, captured under varying illumination. One (unaligned) sequence was acquired with a hand-held camera (cf. Fig. 1), the other (aligned) with a tripod. We then compared the singular values of each sequence. As can be seen in Fig. 3, the singular values of the aligned sequence are smaller. This invites us to recast photometric stereo images registration as a low-rank optimization problem.



Fig. 3: Left: singular values of the observation matrix for the data in Fig. 1, using either a hand-held camera (without, or with the proposed registration technique) or a tripod ensuring perfect alignment. The singular values of the aligned sequences are smaller than those of the unaligned sequence. Right: reflectance and 3D-shape recovered from the (aligned) tripod sequence, which are comparable with those obtained from the registered, hand-held sequence, cf. Fig. 1-right.

3.1 Displacement parameterization

Let us denote by $u^1, \ldots, u^m \in \mathbb{R}^n$, $n \gg m \ge 3$ the vectorized (graylevel) images (m is the number of images and n is the number of pixels in each image), by $u_j^i := u^i(x_j)$ the graylevel of the *i*-th image at pixel $x_j \in \mathbb{R}^2$, and by $u = \left[u^{1^\top}, \ldots, u^{m^\top}\right]^\top \in \mathbb{R}^{mn}$ the full observation set in vector form.

We are looking for a set $W(u; \theta) \in \mathbb{R}^{mn}$ of registered images, where W is a warping function and $\theta := [\theta^{1\top}, \ldots, \theta^{m\top}]^{\top} \in \mathbb{R}^{mp}$ is the unknown set of warp parameters. Therein, the *i*-th transformation, $i \in \{1, \ldots, m\}$, is characterized by the subset of parameters $\theta^i \in \mathbb{R}^p$. We will focus in this paper on invertible parametric image registration i.e., a closed-form expression for the warping function W and its inverse W^{-1} will be assumed, which both will depend on a "small" number of parameters $(p \leq 2n)$.

In particular, in our experiments we focus on affine transformations, which combine translation, rotation, scaling and shearing. In this case, p = 6, and the warping depends on parameters $\theta^i := \left[\theta_1^i, \ldots, \theta_6^i\right]^\top \in \mathbb{R}^6$ according to:

$$W(u;\theta) = \operatorname{vec} \begin{bmatrix} u^1 \left(w^1(x_1;\theta^1) \right) \dots u^m \left(w^m(x_1;\theta^m) \right) \\ \vdots & \vdots \\ u^1 \left(w^1(x_n;\theta^1) \right) \dots u^m \left(w^m(x_1;\theta^m) \right) \end{bmatrix}$$

where the i-th warping is locally defined as

$$w^{i}(\cdot;\theta^{i}): \mathbb{R}^{2} \to \mathbb{R}^{2}$$
$$x_{j} \mapsto \begin{bmatrix} 1+\theta_{1}^{i} & \theta_{3}^{i} \\ \theta_{2}^{i} & 1+\theta_{4}^{i} \end{bmatrix} x_{j} + \begin{bmatrix} \theta_{5}^{i} \\ \theta_{6}^{i} \end{bmatrix}$$
(1)

and its inverse by

$$w^{i^{-1}}(\cdot;\theta^{i}): \mathbb{R}^{2} \to \mathbb{R}^{2}$$
$$x_{j} \mapsto \begin{bmatrix} 1+\theta_{1}^{i} & \theta_{3}^{i} \\ \theta_{2}^{i} & 1+\theta_{4}^{i} \end{bmatrix}^{-1} \left(x_{j} - \begin{bmatrix} \theta_{5}^{i} \\ \theta_{6}^{i} \end{bmatrix} \right).$$
(2)

3.2 Low-rank formulation and convex relaxation

Our approach consists in seeking the warping function which minimizes the rank of the warped observation matrix. Formally, this would come down to solving

$$\min_{\theta \in \mathbb{R}^{m_p}} \operatorname{rank} \left(\operatorname{Mat}(W(u; \theta)) \right), \tag{3}$$

where operator Mat is such that $vec \circ Mat = id$: it reorganizes the elements of a vector into a matrix according to

Mat:
$$\mathbb{R}^{mn} \to \mathbb{R}^{n \times m}$$

 $u := \begin{bmatrix} u^{1^{\top}}, \dots, u^{m^{\top}} \end{bmatrix}^{\top} \mapsto \operatorname{Mat}(u) = \begin{bmatrix} u^{1}, \dots, u^{m} \end{bmatrix}^{\cdot}$

However, both the camera small displacements and specular spots in the image may break the low-rank assumption, by generating occlusions and outliers to the Lambertian model. To take such outliers into account, we simultaneously estimate the displacement parameters θ and a restored observation matrix A, which would sparsely deviate from Mat $(W(u; \theta))$, and have minimal rank:

$$\min_{\substack{A \in \mathbb{R}^{n \times m} \\ \theta \in \mathbb{R}^{m_p}}} \operatorname{rank}(A) + \lambda \left\| \operatorname{vec}(A) - W(u;\theta) \right\|_0 \tag{4}$$

with $\lambda \geq 0$ some tuning parameter and $\|\cdot\|_0$ the number of nonzero elements in a vector. Let us remark that if W is the identity transform, then the proposed model comes down to the low-rank image correction technique advocated in [18].

The optimization problem (4) is difficult to solve since both the rank and the 0-"norm" are nonconvex. So we rather optimize a convex relaxation of (4):

$$\min_{\substack{A \in \mathbb{R}^{n \times m} \\ \theta \in \mathbb{R}^{m_p}}} \|A\|_{\star} + \lambda \|\operatorname{vec}(A) - W(u;\theta)\|_{1}$$
(5)

where $\|\cdot\|_{\star}$ denotes the sum of singular values (nuclear norm) and $\|\cdot\|_1$ the sum of absolute values, which are the tightest convex relaxation of rank and $\|\cdot\|_0$, respectively. Next, we propose an algorithm for numerically solving (5).

4 Augmented Lagrangian framework

In this section, we describe an alternating direction method of multipliers (see, e.g., [6] for a thorough presentation of this algorithm) for solving Problem (5).

4.1 Splitting

Both terms in the objective function of (5) are convex, yet simultaneously optimizing them remains challenging. We therefore split the optimization problem over both terms by turning it into the following equivalent, constrained one:

$$\begin{array}{ll} \min_{\substack{A \in \mathbb{R}^{m \times m} \\ e \in \mathbb{R}^{m p} \\ \theta \in \mathbb{R}^{m p}}} & \|A\|_{\star} + \lambda \|e\|_{1}, \\ \text{(6)} \end{array}$$

Let us consider the augmented Lagrangian associated with this constrained optimization problem:

$$\mathcal{L}^{\sharp}_{\rho}(A, e, \theta, y) := \left\|A\right\|_{\star} + \lambda \left\|e\right\|_{1} + \left\langle y|W(u; \theta) + e - \operatorname{vec}(A)\right\rangle + \frac{\rho}{2} \left\|W(u; \theta) + e - \operatorname{vec}(A)\right\|^{2},$$
(7)

where $y \in \mathbb{R}^{mn}$ is the set of Lagrange multipliers, and $\rho > 0$ is a penalty parameter arbitrarily set to 0.1 in all our experiments. Problem (6) can be solved

iteratively, by alternating minimizations of (7) wrt the primal variables A, e and θ , and a dual ascent step over the dual variable y, yielding the following sequence:

$$A^{(k+1)} = \underset{A}{\operatorname{argmin}} \ \mathcal{L}^{\sharp}_{\rho}(A, e^{(k)}, \theta^{(k)}, y^{(k)}), \tag{8}$$

$$e^{(k+1)} = \underset{e}{\operatorname{argmin}} \ \mathcal{L}_{\rho}^{\sharp}(A^{(k+1)}, e, \theta^{(k)}, y^{(k)}), \tag{9}$$

$$\theta^{(k+1)} = \underset{\theta}{\operatorname{argmin}} \ \mathcal{L}^{\sharp}_{\rho}(A^{(k+1)}, e^{(k+1)}, \theta, y^{(k)}), \tag{10}$$

$$y^{(k+1)} = y^{(k)} + \rho \left(W(u; \theta^{(k+1)}) + e^{(k+1)} - \operatorname{vec}(A^{(k+1)}) \right),$$

starting from the initial estimate $A^{(0)} = \text{Mat}(u), e^{(0)}, \theta^{(0)}, y^{(0)} \equiv 0$ (i.e., the initial solution is set to the unregistered data), and until the relative residual $\|\text{vec}(A^{(k+1)} - A^{(k)})\| / \|\text{vec}(A^{(k)})\|$ falls below some threshold (this threshold is set to 10^{-3} in our experiments).

Next, we detail the resolution of each sub-problem in the above ADMM algorithm.

4.2 Updating the corrected observation matrix A

Update (8) can be computed in closed-form. Indeed:

$$\begin{split} A^{(k+1)} &= \underset{A}{\operatorname{argmin}} \ \mathcal{L}_{\rho}^{\sharp}(A, e^{(k)}, \theta^{(k)}, y^{(k)}) \\ &= \underset{A}{\operatorname{argmin}} \ \|A\|_{\star} + \frac{\rho}{2} \left\| W(u; \theta^{(k)}) + e^{(k)} - \operatorname{vec}(A) + y^{(k)}/\rho \right\|_{F}^{2} \\ &= \underset{A}{\operatorname{argmin}} \ \|A\|_{\star} + \frac{\rho}{2} \left\| A - \operatorname{Mat}(W(u; \theta^{(k)}) + e^{(k)} + y^{(k)}/\rho) \right\|_{F}^{2} \\ &= \underset{d}{\operatorname{prox}}_{\frac{1}{\rho} \|\cdot\|_{\star}} (\operatorname{Mat}(W(u; \theta^{(k)}) + e^{(k)} + y^{(k)}/\rho)), \end{split}$$

where $W(u; \theta^{(k)})$ can be evaluated by interpolation (we used linear interpolation in our experiments) of the unregistered images, and the proximity operator of the nuclear norm admits a closed-form expression as shrinkage of the singular values. More precisely, denoting by $B = U \operatorname{diag}(\sigma^1, \ldots, \sigma^m) V^{\top}$ the singular value decomposition of some matrix $B \in \mathbb{R}^{n \times m}$, $m \leq n$:

$$\operatorname{prox}_{\frac{1}{\rho}\|\cdot\|_{\star}}(B) = U\operatorname{diag}\left(\operatorname{shrink}(\sigma^{1}, 1/\rho), \dots, \operatorname{shrink}(\sigma^{m}, 1/\rho)\right) V^{\top},$$

using the following definition of the shrinkage (soft-thresholding) operator:

$$\operatorname{shrink}(\cdot, 1/\rho) = \operatorname{sign}(\cdot) \max\left\{ |\cdot| - 1/\rho, 0 \right\}.$$

4.3 Updating the error vector e

Update (9) can also be computed in closed-form:

$$\begin{split} e^{(k+1)} &= \underset{e}{\operatorname{argmin}} \ \mathcal{L}_{\rho}^{\sharp}(A^{(k+1)}, e, \theta^{(k)}, y^{(k)}) \\ &= \underset{e}{\operatorname{argmin}} \ \lambda \left\| e \right\|_{1} + \frac{\rho}{2} \left\| e - \left(\operatorname{vec}(A^{(k+1)}) - W(u; \theta^{(k)}) - y^{(k)} / \rho \right) \right\|^{2} \\ &= \underset{d}{\operatorname{prox}}_{\frac{\lambda}{\rho} \| \cdot \|_{1}} \left(\operatorname{vec}(A^{(k+1)}) - W(u; \theta^{(k)}) - y^{(k)} / \rho \right) \\ &= \operatorname{shrink}(\operatorname{vec}(A^{(k+1)}) - W(u; \theta^{(k)}) - y^{(k)} / \rho, \lambda / \rho) \end{split}$$

where shrinkage is to be understood component-wise.

4.4 Updating the displacement coefficients θ

Let us now consider the update (10), which writes as

$$\theta^{(k+1)} = \underset{\theta}{\operatorname{argmin}} \left\| W(u;\theta) - \left(\operatorname{vec}(A^{(k+1)}) - e^{(k+1)} - y^{(k)} / \rho \right) \right\|^2.$$
(11)

This is a classic least-squares warping problem resembling the Kanade-Lucas-Tomasi (KLT) problem, and we refer the interested reader to [1] for an indepth presentation. Therein, several iterative algorithms are discussed, which have different preconditions but are shown to be equivalent if the set of warps forms a group and if warps are differentiable - as it is the case for affine warping. Among the various possibilities, we chose the forward compositional algorithm for the simpler expression of the Jacobian in general.

Let us denote $v^{(k)} := \operatorname{vec}(A^{(k+1)}) - e^{(k+1)} - y^{(k)}/\rho$. Instead of solving (11), one step of the forwards compositional algorithm computes $\theta^{(k+1)}$ as follows:

$$\delta\theta^{(k+1)} = \underset{\delta\theta}{\operatorname{argmin}} \left\| W(W(u;\delta\theta);\theta^{(k)}) - v^{(k)} \right\|^2,$$
(12)
$$W(\cdot;\theta^{(k+1)}) = W(\cdot;\theta^{(k)}) \circ W(\cdot;\delta\theta^{(k+1)}).$$

This update can be carried out independently for each image u^i . We denote $\theta^{i(k)} \in \mathbb{R}^p$ the warp parameters for image *i* at step *k*, such that $\theta^{(k)} \in \mathbb{R}^{mp}$ is the concatenation of all $\theta^{i(k)}$ at step *k*. Similarly, we denote $v^{i(k)} \in \mathbb{R}^n$ the part of $v^{(k)} \in \mathbb{R}^{mn}$ corresponding to the *i*-th image. As explained in [1], performing a first-order Taylor expansion of the expression in (12) gives

$$\delta\theta^{i(k+1)} = \underset{\delta\theta}{\operatorname{argmin}} \sum_{j} \left[W(u;\theta^{(k)})_{j}^{i} + \nabla W(u;\theta^{(k)})_{j}^{i} \frac{\partial W}{\partial \theta} \Big|_{x_{j}} \delta\theta - v_{j}^{i(k)} \right]^{2}, \quad (13)$$

where $(\nabla W(u; \theta^{(k)})_j^i)^\top \in \mathbb{R}^2$ is the warped image gradient at pixel j and $\frac{\partial W}{\partial \theta}\Big|_{x_j} \in \mathbb{R}^{2 \times p}$ are the partial derivatives of the warp function relative to warp parameters. The solution to this least-squares optimization problem is

$$\delta \theta^{i(k+1)} = H^{-1} \left(J^{\top} [v^{i(k)} - W(u; \theta^{(k)})^i] \right),$$

where $J \in \mathbb{R}^{n \times p}$ is the Jacobian matrix and $H = J^{\top}J$ is the Gauss-Newton approximation of the Hessian matrix. The rows of J are computed as

$$J_j = \nabla W(u; \theta^{(k)})_j^i \frac{\partial W}{\partial \theta} \Big|_{x_j}$$

Therein, for any point $x_j := [x_{j,1}, x_{j,2}]^\top \in \mathbb{R}^2$:

$$\frac{\partial W}{\partial \theta}\Big|_{x_j} = \begin{bmatrix} x_{j,1} & 0 & x_{j,2} & 0 & 1 & 0 \\ 0 & x_{j,1} & 0 & x_{j,2} & 0 & 1 \end{bmatrix},$$

since we focus on affine transformations as defined in (1),

We now have all the ingredients for solving the low-rank registration problem (5). In the next section, we discuss the practical implementation of the proposed scheme, and evaluate it on synthetic and real-world data.

5 Implementation

The numerical solution comes down to iterating steps (8) to (10). Yet, convergence of these steps can be established only towards a locally optimum solution, which may be far away from the global optimum if the displacements are important. In order to ease convergence towards a reasonable solution, as well as to accelerate the algorithm, we embed this algorithm inside a multi-scale scheme.

5.1 Multi-scale scheme

The first-order Taylor expansion done in (13) is only viable in the case of small displacements, when the image gradients and the residuals $v^{(k)} - W(u; \theta^{(k)})$ are correlated. Yet at full resolution, warps may generate displacements of dozens of pixels. One workaround, commonly used in KLT-related problems [5], is to consider a multi-scale approach, by generating a pyramid of images as in Fig. 4. Each level halves/doubles the resolution of the previous one, and the only parameter to tune is the number of levels of the pyramid. At the lowest resolution, every pixel covers a much larger area, therefore artificially reducing the image displacement, and improving the conditions for the first-order Taylor approximation. Once the algorithm has converged at one level, we adapt the resulting registration parameters to the next level of the pyramid and use those as initialization. In the case of the affine warp, θ_1^i to θ_4^i stay unchanged, and θ_5^i and θ_6^i are doubled at a level transition.

5.2 Using a sparse subset of pixels

Engel et al. showed in [7] that direct image alignment is possible with a sparse subset of points as long as they are well distributed in the image, and located at pixels with higher gradient magnitude. This can be used as a trick to accelerate the proposed algorithm.

Matthieu Pizenberg et al.

If the scene is highly textured, as in the datasets of Figs. 1 and 2, it is possible to restrict our attention to pixels where the gradient magnitude is maximal. At the coarsest level, all pixels can be considered. Then, for each selected pixel, only one or two of the four subpixels can be selected at the next pyramid level, as illustrated in Fig. 4. Let us emphasize, however, that this feature-based speed up is only a way to accelerate the process, unlike e.g. in [10] where the whole registration builds upon feature detection, hence cannot be employed for unsufficiently textured scenes.



Fig. 4: Top row: pyramid of resolutions for one image of the sequence in Fig. 1. The resolution doubles at each level. Bottom row: sparse subset of pixels (indicated in red) which can be selected at each level, in order to accelerate the process.

5.3 Empirical evaluation

We first quantitatively evaluated our method on simulated image transformations. We considered two challenging real-world photometric stereo datasets from [16], namely the "Buddha" and "Harvest" sequences. Both consist of m =96 images of a shiny surface (exhibiting significant deviations from the Lambertian assumption) taken from the same viewing angle under varying lighting. For each object, we generated random synthetic translations, with the parameters uniformly taken between 0 and 1% of the image width. We then evaluated the accuracy of different algorithms by measuring their mean displacement error. This process was repeated 40 times per object image sequence. The results in Fig. 5 show that the proposed method estimates the correct displacement in most cases. It outperforms standard image registration algorithms which are not designed for handling illumination changes. This can be seen by comparing with Matlab's implementation of similarity optimization (imregtform) or phase correlation (imregcorr).



Fig. 5: Mean displacement error (in pixels) of every warp estimation for our approach (lowrr), and the imregtform and imregcorr Matlab's algorithms, on two sequences ("Buddha" and "Harvest") of the DiLiGenT dataset [16].

Next, we considered a real-world hand-held sequence of m = 13 images of a medieval embrodery (cf. Fig. 1), exhibiting fine-scale geometric variations. We registered the images with the proposed method, and then observed the impact of image registration on the results of a recent photometric stereo algorithm [13]. The proposed approach allows to reach results comparable with those obtained from perfectly aligned images acquired with a tripod (compare with Fig. 3). Lastly, we carried out a similar experiment, on a perfectly planar surface with spatially-varying reflectance (cf. Fig. 2). In this experiment, the camera stands on a tripod, yet the camera is manually triggered, inducing tiny displacements (a few pixels). Taking into account these displacements improves the sharpness of the estimated reflectance, and avoids the hallucination of geometric details.

6 Conclusions and perspectives

We have presented an affine registration technique for images captured under unknown, varying lighting. The proposed approach seeks the registration parameters which minimize the rank of the warped observations, while robustness is ensured by sparsity enforcement. An efficient solution based on the augmented Lagrangian has also been introduced. Experiments on real-world data show that the proposed approach significantly improves 3D-reconstruction by photometric stereo on unaligned data, and even allows using hand-held images. In future work, we will extend the proposed formulation so as to handle non-affine transformations, by inserting optical flow constraints inside the variational problem. Acknowledgements This work was supported by the RIN project "Guide Muséal", and by the ANR grant "Inclusive Museum Guide" (ANR-20-CE38-0007). The authors would like to thank C. Berthelot at the Bayeux Tapestry Museum for supervising the image acquisition campaign.

References

- Baker, S., Matthews, I.: Lucas-Kanade 20 years on: a unifying framework. IJCV 56(3), 221–255 (2004)
- Basri, R., Jacobs, D.W.: Lambertian reflectance and linear subspaces. PAMI 25(2), 218–233 (2003)
- Basri, R., Jacobs, D.W., Kemelmacher, I.: Photometric stereo with general, unknown lighting. IJCV 72, 239–257 (2007)
- 4. Berkiten, S., Rusinkiewicz, S.: Alignment of images captured under different light directions. Tech. Rep. TR-974-14 (2014)
- 5. Bouguet, J.Y.: Pyramidal implementation of the affine Lucas-Kanade feature tracker: description of the algorithm. Tech. rep., Intel corporation (2001), tech. note
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Founds. and Trends in Mach. Learn. 3, 1–122 (2010)
- Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. PAMI 40(3), 611–625 (2017)
- Foroosh, H., Zerubia, J., Berthod, M.: Extension of phase correlation to subpixel registration. TIP 11(3), 188–200 (2002)
- Förstner, W.: A feature based correspondence algorithm for image matching. In: ISP Comm. III. pp. 150–166 (1986)
- Harrison, A.P., Joseph, D.: Translational photometric alignment of single-view image sequences. CVIU 116(6), 765–776 (2012)
- 11. Ikehata, S., Wipf, D., Matsushita, Y., Aizawa, K.: Robust photometric stereo using sparse regression. In: CVPR. pp. 318–325 (2012)
- Inoshita, C., Mukaigawa, Y., Matsushita, Y., Yagi, Y.: Surface normal deconvolution: Photometric stereo for optically thick translucent objects. In: ECCV. pp. 346–359 (2014)
- Quéau, Y., Wu, T., Lauze, F., Durou, J.D., Cremers, D.: A non-convex variational approach to photometric stereo under inaccurate lighting. In: CVPR. pp. 99–108 (2017)
- Schmitt, C., Donne, S., Riegler, G., Koltun, V., Geiger, A.: On joint estimation of pose, geometry and svBRDF from a handheld scanner. In: CVPR. pp. 3493–3503 (2020)
- Sengupta, S., Zhou, H., Forkel, W., Basri, R., Goldstein, T., Jacobs, D.: Solving uncalibrated photometric stereo using fewer images by jointly optimizing low-rank matrix completion and integrability. JMIV 60(4), 563–575 (2018)
- Shi, B., Mo, Z., Wu, Z., Duan, D., Yeung, S., Tan, P.: A benchmark dataset and evaluation for non-Lambertian and uncalibrated photometric stereo. PAMI 41(2), 271–284 (2019)
- 17. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. Opt. Eng. **19**(1), 134–144 (1980)
- Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., Ma, Y.: Robust photometric stereo via low-rank matrix completion and recovery. In: ACCV. pp. 703–717 (2010)

12