

A Matrix Decomposition Perspective on Calibrated Photometric Stereo

Luca Magri¹, Roberto Toldo², Umberto Castellani¹, and Andrea Fusiello³

¹ Dip. Informatica - Università di Verona - Strada Le Grazie 15, Verona, Italy

² 3Dflow SRL, c/o Computer Science Park - Strada Le Grazie 15, Verona, Italy

³ DPIA - Università di Udine - Via Delle Scienze, 208 - Udine, Italy

Abstract. Leveraging on recent advances in robust matrix decomposition, we revisit Lambertian photometric stereo as a robust low-rank matrix recovery problem with both missing and corrupted entries, tailoring Grasta and R-GoDec to normal surface estimation. A method to automatically detect shadows is proposed. The performance of different robust matrix completion techniques are analyzed on the challenging DiLiGENT datasets.

Keywords: Calibrated Photometric Stereo, Robust Matrix Factorization

1 Introduction

Robust matrix decomposition and completion has been an active research topic in recent years, and many methods exploiting low rank and sparsity constraints have sprouted out in several fields of applications, such as pattern recognition, machine learning, and signal processing just to name a few. In this work, we explore the performances of these techniques on calibrated Lambertian photometric stereo [9], *i.e.* the problem of estimating the surface normals of an object by observing several intensity images captured by a fixed camera under different known lighting conditions. In particular, we offer an overview on robust matrix decomposition methods tailored to photometric stereo – using for the first time Grasta and R-GoDec for this scope – and a quantitative experimental evaluation on the recently proposed DiLiGENT dataset. A simple yet effective shadow detection method is also presented.

Notation: Matrix will be indicated in sans serif font $\mathbf{A} = [a_j^i]$, the i -th row of \mathbf{A} is denoted by \mathbf{A}^i , while the j -th column of \mathbf{A} is indicated by \mathbf{A}_j .

2 The geometry of single-light images

Let $I \in \mathbb{R}^p$ be an image composed by p pixels stacked by column. Following [3], under Lambertian assumption the properties of interest of an object Y can be encoded in matrix form as $\text{diag}(R)\mathbf{N}^\top \in \mathbb{R}^{p \times 3}$, where $R = (\rho_1, \dots, \rho_p)^\top$ is

the vector of pixels albedos, and $\mathbf{N} = [\mathbf{N}_1, \dots, \mathbf{N}_p] \in \mathbb{R}^{3 \times p}$ collects the unitary normals of the object. Thus the image of Y illuminated by a distant point-light source $L \in \mathbb{R}^3$, is given by:

$$I = \max(\text{diag}(R)\mathbf{N}^\top L, 0). \quad (1)$$

Varying L , one obtains the so-called *illumination space* of Y defined as $\mathcal{L} = \{I : I = \text{diag}(R)\mathbf{N}L : L \in \mathbb{R}^3\} \subset \mathbb{R}^p$. Clearly $\dim(\mathcal{L}) = \text{rank}(\mathbf{N}^\top L) \leq 3$, therefore, if the normals span \mathbb{R}^3 , the dimension of \mathcal{L} is 3.

Belhumeur and Kriegman also observe that \mathcal{L} intersect at most⁴ $p(p-1)+2$ orthants of \mathbb{R}^p . Let $\mathcal{L}_0 = \mathcal{L} \cap \mathbb{R}_+^p$, where $\mathbb{R}_+^p = \{x \in \mathbb{R}^p : x_i > 0 \forall i\}$, and \mathcal{L}_i be the intersection of \mathcal{L} with the other orthants. By construction, \mathcal{L}_i are convex cones, and correspond to different shading configuration of pixels. As instance, \mathcal{L}_0 corresponds to images having all the pixels illuminated by a lighting source. The space of all possible images of Y is obtained by adding to \mathcal{L}_0 the images where not all the pixels are simultaneously illuminated, i.e. the projection of the cones \mathcal{L}_i , $i \neq 0$, on the boundary of \mathbb{R}_+^p via the map $P : I \mapsto \max(I, 0)$. Therefore, the space of all the images of a convex Lambertian object, varying the direction of a single light source is given by the union of at most $p(p-1)+2$ convex cones.

$$\bigcup_{i=0}^{\nu(\nu-1)+2} P(\mathcal{L}_i). \quad (2)$$

Experimentally, it was demonstrated that this union of cones is “flat” and can be approximated by a linear subspace of dimension 3.

3 Robust Matrix Completion and Decomposition

The linear property of light superposition inspired the use of matrix completion and robust decomposition techniques to tackle the photometric stereo problem [10, 11]. Given f images of the same object organized as a $p \times f$ matrix $\mathbf{X} = [I_1, \dots, I_f]$, with images stacked as columns, the main intuition is to recover the illumination space as a low-rank matrix \mathbf{A} that models the diffusive Lambertian observations, and to handle the non-Lambertian measurements as outliers. In particular, shadows, i.e. pixels outside \mathcal{L}_0 , are treated as missing entries, whereas a *sparse* error matrix \mathbf{S} accounts for the corruptions produced by strong specularities (highlights).

More formally, the image formation model can be rephrased as

$$\mathbf{X} = P_\Omega(\mathbf{A}) + \mathbf{S} \quad (3)$$

where $\mathbf{A} = \text{diag}(R)\mathbf{N}^\top \mathbf{L}$ is low rank, $\mathbf{L} = [L_1, \dots, L_f]$ collects the known light source vectors, $\Omega = \{(i, j) : \text{where } \mathbf{N}^\top \mathbf{L} \text{ is nonzero}\}$ is the set of observed entries, P_Ω indicates a linear projection of matrices defined component-wise as

⁴ more precisely, \mathcal{L} intersects $\nu(\nu-1)+2$ orthants, where ν is the number of distinct normal in B .




Fig. 1: The factorization of the intensities of an ideal Lambertian object (left), the same factorization in presence of shadow and highlights

$[P_{\Omega}(\mathbf{A})]_j^i = \mathbf{A}_j^i$ if $(i, j) \in \Omega$ and 0 otherwise, and \mathbf{S} is the matrix of sparse error. A graphical representation of this model is depicted in Fig. 1, where the operator P_{Ω} – which is the matrix counterpart of the projection P onto the positive orthant – is represented in an equivalent fashion as the Hadamard (element-wise) product between \mathbf{A} and a matrix \mathbf{V} defined component-wise as

$$\mathbf{V}_j^i = \begin{cases} 1 & \text{if } (i, j) \in \Omega, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

In this way, photometric stereo becomes the problem of recovering a low-rank matrix with both missing entries – the shadows – and corrupted entries corresponding to the unmodelled phenomena (e.g., non Lambertian). The rank of \mathbf{A} may vary according to the image formation model adopted, and for Lambertian photometric stereo $\text{rank}(\mathbf{A}) = 3$. Once the low rank matrix has been recovered, in the calibrated scenario, the normals can be easily estimated in closed form by normalizing the row of $\mathbf{L}^+ \mathbf{A}$, where \mathbf{L}^+ denotes the pseudo-inverse of \mathbf{L} .

Decomposition into low-rank and sparse matrices has been developed in different formulation problems, hereinafter, we briefly review some of them that can be profitably adopted to tackle the Problem (3), namely: Robust Principal Component Analysis and L1-ALM (that have already been tailored to photometric stereo problem), together with Grasta and R-GoDec that we are going to apply to this scenario for the first time.

Robust Principal Component Analysis (RPCA) decomposes \mathbf{X} into a low rank and sparse terms, without being given $\text{rank}(\mathbf{A})$. The cost function is:

$$\arg \min_{\mathbf{A}, \mathbf{S}} \text{rank}(\mathbf{A}) + \lambda \|\mathbf{S}\|_0 \text{ s.t. } \mathbf{X} = P_{\Omega}(\mathbf{A}) + \mathbf{S}. \quad (5)$$

Unfortunately this problem turns to be intractable, therefore, instead of directly minimizing the discontinuous rank function and the ℓ_0 norm, the above objective function is relaxed to its convex surrogate; the rank of \mathbf{A} is replaced with the nuclear norm $\|\mathbf{A}\|_*$ – i.e. the sum of the singular values of \mathbf{A} – and the ℓ_0 norm is substituted for the ℓ_1 norm:

$$\arg \min \|\mathbf{A}\|_* + \lambda \|\mathbf{S}\|_1 \text{ s.t. } \mathbf{X} = P_{\Omega}(\mathbf{A}) + \mathbf{S}. \quad (6)$$

Several techniques can be used to minimize Eq. (6); in [10], e.g., an adaptation of the augmented Lagrange multiplier method is used.

L1-ALM [11] proposes to find a solution to Problem (3) by enforcing exactly the low rank constraint $\text{rank}(\mathbf{A}) = r$, and leverages on the factorization of the matrix $\mathbf{A} = \mathbf{H}\mathbf{K}$ as the product of a $p \times r$ matrix \mathbf{H} and a $r \times f$ matrix \mathbf{K} . As the factorization is defined up to an invertible matrix, in order to shrink the solution space, the matrix \mathbf{H} is enforced to be column-orthogonal, i.e. $\mathbf{H}^\top \mathbf{H} = \mathbf{I}_r$, where \mathbf{I}_r denotes the $r \times r$ identity matrix. The objective L1-ALM tries to minimize is

$$\arg \min_{\mathbf{H}, \mathbf{K}} \|P_\Omega(\mathbf{X} - \mathbf{H}\mathbf{K})\|_1 + \lambda \|\mathbf{K}\|_* \text{ s.t. } \mathbf{H}^\top \mathbf{H} = \mathbf{I}_r. \quad (7)$$

where $P_\Omega(\mathbf{X} - \mathbf{H}\mathbf{K})$ is the ℓ_1 norm of the sparse error matrix $\mathbf{S} = \mathbf{X} - P_\Omega(\mathbf{A})$ and $\|\mathbf{K}\|_* = \|\mathbf{H}\mathbf{K}\|_* = \|\mathbf{A}\|_*$ is a trace-norm regularization term, which, due to the orthogonality of \mathbf{H} , is equivalent to the nuclear norm of \mathbf{A} . This optimization problem is resolved via inexact augmented Lagrange multiplier and Gauss-Seidel iterations

Grassmannian Robust Adaptive Subspace Tracking Algorithm (Grasta) [5] is an online robust subspace tracking algorithm, that works in the presence of corrupted and missing data. Given a sequence of incomplete vectors $\{v_1, \dots, v_t\}$ that lie on a r -dimensional subspace, Grasta estimates this subspace, by minimizing the ℓ_1 error between the recovered subspace and the observed partial vector. This formulation can be casted to the problem of Eq. (3) as

$$\min_{\mathbf{S}} \|\mathbf{S}\|_1 \text{ s.t. } \mathbf{X} = P_\Omega(\mathbf{H}\mathbf{K}) + \mathbf{S} \quad (8)$$

where, similarly to Eq. (7), \mathbf{A} is expressed as the product of two factors \mathbf{H}, \mathbf{K} , the first being an element of the Grassmanian Manifold $G(r, p)$. The problem is iteratively solved for \mathbf{H} and \mathbf{K} separately: fixed \mathbf{H} , \mathbf{K} is update via ADMM, whereas, when \mathbf{K} is fixed, \mathbf{H} is updated performing incremental gradient descent on the Grassmanian manifold. Even if the partial measurements of the matrix \mathbf{X} are required to be exactly fixed, nevertheless, in practice, it was demonstrated that the algorithm is robust to small non sparse additive noise.

Robust Go Decomposition (R-GoDec) [2] proposes a robust approximate matrix completion and decomposition technique that improves GoDec [12]. An additional sparse term \mathbf{S}' that has support on Ω^C – the complementary of Ω – is introduced to account for missing entries. In addition small sparse noise \mathbf{E} is explicitly introduced in the decomposition:

$$\mathbf{X} = \mathbf{A} + \mathbf{S} + \mathbf{S}' + \mathbf{E}. \quad (9)$$

The corresponding minimization problem is

$$\arg \min_{\mathbf{A}, \mathbf{S}} \|\mathbf{X} - \mathbf{A} - \mathbf{S} - \mathbf{S}'\|_F^2 \quad (10)$$

such that $\text{rank}(\mathbf{A}) \leq r$, \mathbf{S} is sparse and \mathbf{S}' has support in Ω^C . This problem is solved using a block-coordinate minimization scheme. At first, the rank- r projection of the incomplete matrix given in input is computed through Bilateral Random Projection – faster than SVD – and assigned to \mathbf{A} . Then, the two sparse terms \mathbf{S} and \mathbf{S}' are updated separately. The outlier term \mathbf{S} is computed via soft-thresholding operator, and \mathbf{S}' is updated as $-P_{\Omega^C}(\mathbf{A})$.

4 Detecting shadows

This section is aimed at estimating the set Ω^C of shadowed pixels, in order to treat them as missing data and to reduce their influence in the low-rank matrix recovery. For this purpose, it becomes necessary to reason about the *visibility* of light source with respect to each image pixel in order to recognize which lights shine on which points and to discard the pixels in shadow.

To this end, a commonly employed solution is intensity-based thresholding: Pixels whose intensity lies below a certain threshold are considered in shadow. While this heuristic in some cases is enough to recover the light-visibility information, in general, the intensity of individual pixels depends on the variations of the *unknown* albedo of the object, thus, the brightness alone turns to be an unreliable cue.

In order to overcome this pitfall, other techniques have been proposed. For example, [4] adopts a graph cuts based method to estimate light visibility in a Markov Random Field formulation, where a per-pixels error, based on photometric stereo, is balanced by a smoothness constraints on shadows, aimed at promoting spatial coherence. Sunkavalli et al. in [8] avoid to enforce spatial coherence on shadows and present a method that works both in the calibrated and uncalibrated scenario leveraging on subspace clustering. Pixels sharing the same visibility configuration lie on linear subspaces, termed visibility subspaces, that are extracted using Sequential RANSAC. Once these subspaces are recovered and the object surface is segmented accordingly, the set of lights that shine on each region are identified analysing the magnitude of the subspace lighting obtained via SVD.

The visibility information can be encoded in the $n \times f$ *visibility matrix* \mathbf{V} defined as in Eq. (4). Each row of \mathbf{V} can be seen as the indicator function of the subset of lights visible by each pixel.

In our calibrated scenario, we want to recover \mathbf{V} given the intensity matrix \mathbf{X} and the lighting directions \mathbf{L} . To this end, assuming that there are at least $f \geq 4$ images, we propose a simple approach based on LMEDS [6].

The main idea is to approximate at first the space of the possible visibility configuration by randomly sampling triplets ω of lights. Fixed a pixel i , a tentative normal vector is estimated via least square regression for every lighting triplets. Hence, the normal $\hat{\mathbf{N}}_i$ which minimise the median of squared residuals is retained as a solution. By scrutinising the residual vector $I^i - \max(0, \hat{\mathbf{N}}_i^\top \mathbf{L})$, a binary weighting vector w_i is defined setting its j -th entry equals 1, if the j -th error is smaller of $2.5\hat{\sigma}$, and 0 otherwise, where $\hat{\sigma}$ is a robust estimate of the variance of the per pixels residuals defined by:

$$\hat{\sigma} = 1.4826(1 + 5/(f - 3))\sqrt{\text{median } r_{\hat{\omega}}^2} \quad (11)$$

At the end, the normal estimate $\hat{\mathbf{N}}_i$ is refined using iteratively reweighted least squares (IRLS) on the set of lights $\{\mathbf{L}_j : w_i^j = 1\}$.

The matrix $\mathbf{W} = [w_1, \dots, w_p]^\top$, composed by the weight-vectors arranged by row, could be used as a proxy for the visibility matrix, however here we take

light directions into account, and we obtain a visibility matrix \hat{V} setting:

$$\hat{V}_j^i = \begin{cases} 1 & \text{if } \hat{N}_i^\top L_j > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

5 Evaluation the DiLiGenT dataset

The methods presented in the previous section are here challenged on the *Directional Lightings*, objects of *General* reflectance, and ground *Truth* shapes datasets (DiLiGenT) [7], a recently proposed benchmark of ten objects shown in Fig. 2.

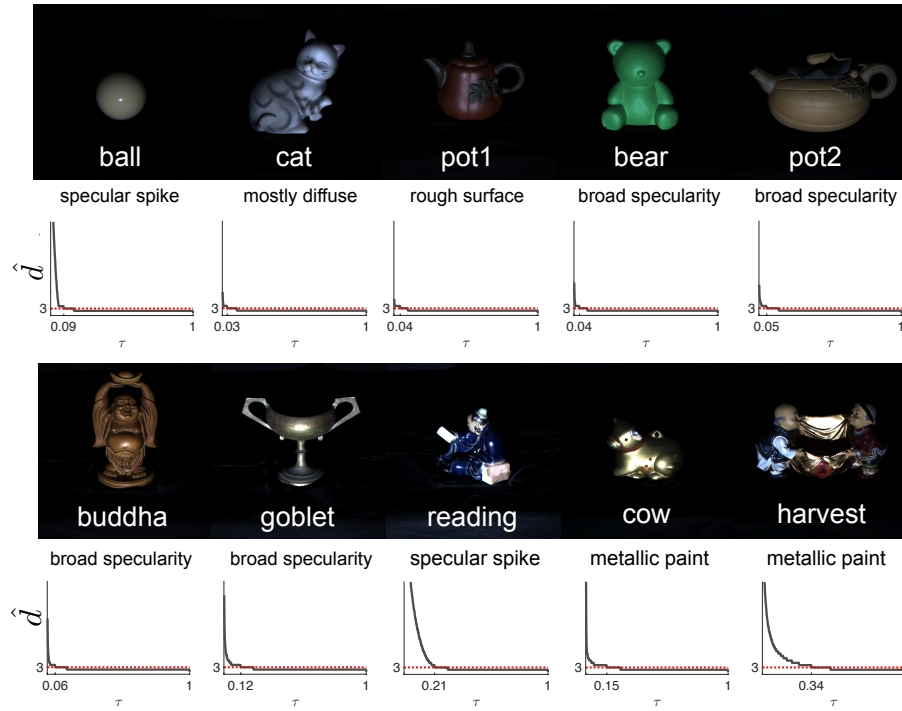


Fig. 2: The ten object of the DiLiGenT datasets with the respective \hat{d} index.

This collection offers a great variety in terms of materials, appearances, geometries and type of deviations from the Lambertian model – from sparse specular spikes to broad specular lobes. This miscellany of non-diffusive phenomena can be captured analyzing the behavior of the index

$$\hat{d} = \min_d \left\{ \frac{\sum_{i=d+1}^f \sigma_i^2}{\sum_{i=1}^f \sigma_i^2} < \tau \right\} \quad (13)$$




Fig. 3: Example of visibility masks on the *ball* and on the *pot1* dataset. Visibility patterns are color coded: same colors correspond to the same shadowing configurations.

with respect to τ which represents the smallest number such that the fraction of information discarded by the corresponding rank approximation is less than a threshold τ : the last three objects are the ones that deviates more from the rank-3 Lambertian approximation.

Visibility mask. Sample results attained by this method on the DILIGENT dataset are shown in Fig. 3, where it can be appreciated that \hat{V} well approximates the ground truth visibility – computed as in Eq. 4 using the ground truth normals. As a reference, we also compare the LMEDS approach with the one based on visibility subspaces [8]. Some differences can be pointed out. First, the extraction of visibility subspaces requires two parameters, namely the inlier threshold of RANSAC and a threshold on the magnitude of light. The inlier threshold is not always an educated guess, as the subspace estimation may be strained by the presence of highlighted pixels whose intensity profiles follow a different distribution with respect to shadowed points. LMEDS, on the contrary, is parameter-free and avoids this difficulties.

Second, LMEDS estimates the visibility configuration locally per pixels, visibility subspace, on the other hand, are estimated globally and pixels that lie in the intersections of multiple subspaces are not properly handled. Third, the random sampling performed to extract the visibility subspace acts on pixels, therefore the dimension of possible samples is $\binom{p}{3}$, which, as usually $p > f$ is higher than the upper bound on the number of samples of LMEDS $\binom{f}{3}$. Finally, LMEDS procedure can be parallelized in a straightforward way.

The visibility masks estimated by LMEDS are fed to the matrix completion algorithms to reduce the influence of shadowed pixels on the low rank estimation step.

Normal estimation. We randomly chose 9 different lighting configurations for each dataset, and we compare the estimated normals with the ground truth ones, averaging the results on 10 trials per dataset. The rank was fixed to 3 for Grasta and R-GoDec, and the regularization parameter to $\lambda = 1/\sqrt{p}$, whereas for L1-ALM we used $\lambda = 10^{-3}$ as suggested in the authors implementation [1]. The performances of the matrix factorization methods are recorded in Tab. 1 where the mean, the median and the standard deviation of angular errors were reported for each method. As a reference we also detailed the errors attained by Least Square and LMEDS. When the accuracy of a method is worse than Least

Table 1: Angular error on the DiLiGENT dataset

	Least Square			Lmeds			RPCA			L1-ALM			Grasta			R-GoDec		
	mean	med	std	mean	med	std	mean	med	std	mean	med	std	mean	med	std	mean	med	std
ball	4.38	2.53	0.74	3.59	2.40	0.86	5.23	3.77	0.53	5.20	2.53	0.87	2.70	2.28	0.92	2.86	2.35	0.80
cat	9.09	6.78	0.67	8.50	6.45	0.84	10.05	7.28	1.22	10.85	7.64	6.39	8.24	6.23	1.27	8.24	6.28	1.24
pot1	9.46	6.83	0.70	8.93	6.43	0.93	9.33	6.68	0.93	9.41	6.69	0.99	8.69	6.11	0.95	8.87	6.43	0.87
bear	10.49	8.16	1.15	9.90	7.77	1.21	10.13	8.28	0.99	10.32	8.08	1.31	9.39	7.50	1.50	9.36	7.55	1.48
pot2	15.89	12.02	1.12	15.46	11.35	1.13	12.37	10.21	0.88	15.88	11.77	1.03	15.38	10.88	1.05	15.57	11.41	1.10
buddha	15.45	10.25	0.87	14.34	9.46	1.00	15.37	10.96	1.30	14.28	9.33	1.35	14.37	9.20	1.39	14.19	9.24	1.28
goblet	19.43	15.72	0.71	18.65	14.76	0.75	17.26	13.96	0.91	20.82	16.21	5.51	18.43	14.03	0.81	18.39	14.05	0.83
reading	20.20	12.55	1.48	18.16	10.84	1.36	23.12	19.79	0.92	29.44	21.62	7.33	37.08	33.77	4.04	19.38	12.17	1.57
cow	26.48	26.87	0.75	25.63	25.75	0.84	15.00	13.95	1.37	33.02	28.16	8.78	31.83	31.15	2.43	26.66	26.58	0.61
harvest	31.19	25.59	0.61	30.34	24.16	0.59	27.74	22.37	1.61	35.75	29.57	3.18	33.88	27.59	2.55	32.31	24.73	0.96

Square, we colored the corresponding cell with gray. Other colors are used to highlights the best results achieving the minimum error.

Grasta performed well on those datasets that manifest a clear diffusive component corrupted by local and sparse non-Lamberitan effects, whereas it worsened the results of Least Square estimation with respect to the last three sequences.

On the contrary RPCA achieved less accurate results on the first sequences and performed better on those challenging datasets characterized by board specularly and complex BRDF (pot2, goblet, cow and harvest have a metal appearance). R-GoDec behavior is similar to Grasta as can be sensed, looking at Fig. 4 – where the mean angular error is plotted for each sequence of the dataset.

Sample results of attained normals are shown in Fig. 5. One can also note that LMEDS always improved the performance of LS.




Fig. 4: Mean angular error on the diligent dataset (9 images, average on 10 trials)

Regularization parameter. We recall that λ is a weighting parameter that is used by L1-ALM, R-GoDec and RPCA to balance between the low-rank and the sparsity terms. In all the above experiments, this parameter was fixed. However, with better choice, it is possible to correct larger amount of outliers, enhancing the performance of the algorithms. Here we demonstrated the effect of λ on L1-




Fig. 5: Sample normal maps obtained on cat, buddha and harvest.

ALM, R-GoDec and RPCA with respect to different number of input images we performed normal estimation on 4, 6, 12, 18, 24, 30, 36 randomly drawn images using $C \frac{1}{\sqrt{p}}$ with $C \in \{0.05, 0.1, 0.2, 0.4, 0.8, 1, 1.2, 1.4, 1.6\}$. The corresponding mean angular errors are shown in Fig. 6, where it can be appreciated that L1-ALM and R-GoDec benefit of the prior knowledge of rank being less sensitive to the number of images and the choice of λ . The minimum mean angular error per each datasets are reported in Tab. 2, where, for completeness, we also added the performance of Grasta varying only the number of images (λ is not required).

Table 2: Minimum mean error in degree varying λ and the number of images

	ball	cat	pot1	bear	pot2	buddha	goblet	reading	cow	harvest	mean	median
RPCA	2.68	7.41	7.42	6.38	10.40	11.99	14.56	15.04	10.80	25.58	11.23	10.60
L1-ALM	2.11	7.14	7.89	6.10	12.74	12.41	16.65	14.30	24.09	29.49	13.29	12.58
Grasta	2.11	7.13	7.95	6.11	12.78	12.45	16.64	20.81	25.03	29.60	14.06	12.61
R-godec	2.11	7.08	7.90	6.09	12.74	12.41	16.65	14.29	23.96	29.51	13.27	12.58

Distribution of light directions. In this experiment, we studied the effects of the distribution of light sources. We considered three different light configurations depicted in Fig. 7b: A) 9 lights are randomly selected; B) we choose a central light and the remaining 8 are those maximizing their distance from it;




Fig. 6: Average angular error on the whole DILIGENT dataset varying the number of images and the regularization parameter λ . (the scale of the colorbar is different for each methods)

C) we select 9 neighboring light sources. We run all the methods on the ball dataset, which is the only convex object and therefore results are less affected by the actual light orientations.

The summary of the experiment is that Grasta and R-GoDec preferred random and spread distribution, whereas RPCA and L1-ALM take advantage of the redundancy provided by the dense configuration.

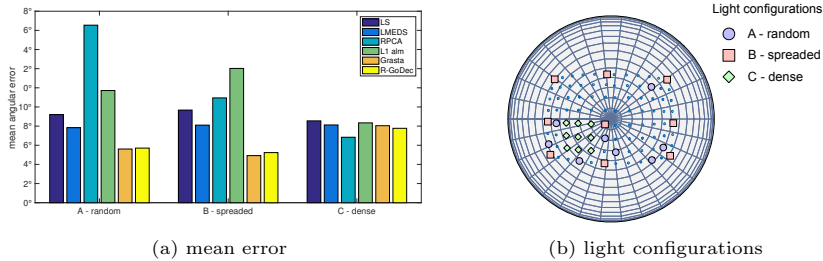


Fig. 7: Varying the light configurations on the *ball* dataset

6 Conclusion

In this work, we tackle the problem of photometric stereo leveraging on robust matrix factorization techniques. We showed that the proposed shadow estimation based on LMEDS is able to produce accurate results, that, in turn, can be profitably fed to matrix completion algorithms. Experiments on a challenging datasets demonstrate that, if the object of interest is mostly Lambertian with strong and sparse non diffusive phenomena, it is advisable to adopt matrix approximation method with fixed rank. In this situation Grasta, followed by R-GoDec, performs better than L1-ALM. On the other side, if one is interested in recovering the normals of a surface that does not exhibit a strong diffusive behavior, all the methods suffer of low precision, but RPCA attains the more accurate results.

References

1. <https://sites.google.com/site/yinqiangzheng/home/RegL1-ALM.zip?attredirects=0>
2. Arrigoni, F., Magri, L., Rossi, B., Fragneto, P., Fusiello, A.: Robust absolute rotation estimation via low-rank and sparse matrix decomposition. In: 3D Vision (3DV), 2014 2nd International Conference on. vol. 1, pp. 491–498. IEEE (2014)
3. Belhumeur, P.N., Kriegman, D.J.: What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision* 28(3), 245–260 (1998)
4. Chandraker, M., Agarwal, S., Kriegman, D.: Shadowcuts: Photometric stereo with shadows. In: *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*. pp. 1–8. IEEE (2007)
5. He, J., Balzano, L., Szelam, A.: Incremental gradient on the grassmannian for online foreground and background separation in subsampled video. In: *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. pp. 1568–1575. IEEE (2012)
6. Leroy, A.M., Rousseeuw, P.J.: *Robust regression and outlier detection*. Wiley Series in Probability and Mathematical Statistics, New York: Wiley, 1987 (1987)
7. Shi, B., Wu, Z., Mo, Z., Duan, D., Yeung, S.K., Tan, P.: A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)
8. Sunkavalli, K., Zickler, T., Pfister, H.: Visibility subspaces: Uncalibrated photometric stereo with shadows. In: *European Conference on Computer Vision*. pp. 251–264. Springer (2010)
9. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19(1), 191139–191139– (1980), <http://dx.doi.org/10.1117/12.7972479>
10. Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., Ma, Y.: Robust photometric stereo via low-rank matrix completion and recovery. In: *Asian Conference on Computer Vision*. pp. 703–717. Springer (2010)
11. Zheng, Y., Liu, G., Sugimoto, S., Yan, S., Okutomi, M.: Practical low-rank matrix approximation under robust l_1 -norm. In: *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. pp. 1410–1417. IEEE (2012)
12. Zhou, T., Tao, D.: Godec: Randomized low-rank & sparse matrix decomposition in noisy case. In: Getoor, L., Scheffer, T. (eds.) *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. pp. 33–40. ICML '11, ACM, New York, NY, USA (June 2011)