

Fast-Robust PCA

Markus Storer, Peter M. Roth, Martin Urschler, and Horst Bischof

Institute for Computer Graphics and Vision
Graz University of Technology
Inffeldgasse 16/II, 8010 Graz, Austria
{storer,pmroth,urschler,bischof}@icg.tugraz.at

Abstract. Principal Component Analysis (PCA) is a powerful and widely used tool in Computer Vision and is applied, e.g., for dimensionality reduction. But as a drawback, it is not robust to outliers. Hence, if the input data is corrupted, an arbitrarily wrong representation is obtained. To overcome this problem, various methods have been proposed to robustly estimate the PCA coefficients, but these methods are computationally too expensive for practical applications. Thus, in this paper we propose a novel fast and robust PCA (FR-PCA), which drastically reduces the computational effort. Moreover, more accurate representations are obtained. In particular, we propose a two-stage outlier detection procedure, where in the first stage outliers are detected by analyzing a large number of smaller subspaces. In the second stage, remaining outliers are detected by a robust least-square fitting. To show these benefits, in the experiments we evaluate the FR-PCA method for the task of robust image reconstruction on the publicly available ALOI database. The results clearly show that our approach outperforms existing methods in terms of accuracy and speed when processing corrupted data.

1 Introduction

Principal Component Analysis (PCA) [1] also known as Karhunen-Loève transformation (KLT) is a well known and widely used technique in statistics. The main idea is to reduce the dimensionality of data while retaining as much information as possible. This is assured by a projection that maximizes the variance but minimizes the mean squared reconstruction error at the same time. Murase and Nayar [2] showed that high dimensional image data can be projected onto a subspace such that the data lies on a lower dimensional manifold. Thus, starting from face recognition (e.g., [3, 4]) PCA has become quite popular in computer vision¹, where the main application of PCA is dimensionality reduction. For instance, a number of powerful model-based segmentation algorithms such as Active Shape Models [8] or Active Appearance Models [9] incorporate PCA as a fundamental building block.

In general, when analyzing real-world image data, one is confronted with unreliable data, which leads to the need for robust methods (e.g., [10, 11]). Due to

¹ For instance, at CVPR 2007 approximative 30% of all papers used PCA at some point (e.g., [5–7]).

its least squares formulation, PCA is highly sensitive to outliers. Thus, several methods for robustly learning PCA subspaces (e.g., [12–16]) as well as for robustly estimating the PCA coefficients (e.g., [17–20]) have been proposed. In this paper, we are focusing on the latter case. Thus, in the learning stage a reliable model is estimated from undisturbed data, which is then applied to robustly reconstruct unreliable values from the unseen corrupted data.

To robustly estimate the PCA coefficients Black and Jepson [18] applied an M-estimator technique. In particular, they replaced the quadratic error norm with a robust one. Similarly, Rao [17] introduced a new robust objective function based on the MDL principle. But as a disadvantage, an iterative scheme (i.e., EM algorithm) has to be applied to estimate the coefficients. In contrast, Leonardis and Bischof [19] proposed an approach that is based on sub-sampling. In this way, outlying values are discarded iteratively and the coefficients are estimated from inliers only. Similarly, Edwards and Murase introduced adaptive masks to eliminate corrupted values when computing the sum-squared errors.

A drawback of these methods is their computational complexity (i.e., iterative algorithms, multiple hypotheses, etc.), which limits their practical applicability. Thus, we develop a more efficient robust PCA method that overcomes this limitation. In particular, we propose a two-stage outlier detection procedure. In the first stage, we estimate a large number of smaller subspaces sub-sampled from the whole dataset and discard those values that are not consistent with the subspace models. In the second stage, the data vector is robustly reconstructed from the thus obtained subset. Since the subspaces estimated in the first step are quite small and only a few iterations of the computationally more complex second step are required (i.e., most outliers are already discarded by the first step), the whole method is computationally very efficient. This is confirmed by the experiments, where we show that the proposed method outperforms existing methods in terms of speed and accuracy.

This paper is structured as follows. In Section 2, we introduce and discuss the novel fast-robust PCA (FR-PCA) approach. Experimental results for the publicly available ALOI database are given in Section 3. Finally, we discuss our findings and conclude our work in Section 4.

2 Fast-Robust PCA

Given a set of n high-dimensional data points $\mathbf{x}_j \in \mathbb{R}^m$ organized in a matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{m \times n}$, then the PCA basis vectors $\mathbf{u}_1, \dots, \mathbf{u}_{n-1}$ correspond to the eigenvectors of the sample covariance matrix

$$\mathbf{C} = \frac{1}{n-1} \hat{\mathbf{X}} \hat{\mathbf{X}}^\top, \quad (1)$$

where $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n]$ is the mean normalized data with $\hat{\mathbf{x}}_j = \mathbf{x}_j - \bar{\mathbf{x}}$. The sample mean $\bar{\mathbf{x}}$ is calculated by

$$\bar{\mathbf{x}} = \frac{1}{n} \sum_{j=1}^n \mathbf{x}_j. \quad (2)$$

Given the PCA subspace $\mathbf{U}_p = [\mathbf{u}_1, \dots, \mathbf{u}_p]$ (usually only p , $p < n$, eigenvectors are sufficient), an unknown sample $\mathbf{x} \in \mathbb{R}^m$ can be reconstructed by

$$\tilde{\mathbf{x}} = \mathbf{U}_p \mathbf{a} + \bar{\mathbf{x}} = \sum_{j=1}^p a_j \mathbf{u}_j + \bar{\mathbf{x}}, \quad (3)$$

where $\tilde{\mathbf{x}}$ denotes the reconstruction and $\mathbf{a} = [a_1, \dots, a_p]$ are the PCA coefficients obtained by projecting \mathbf{x} onto the subspace \mathbf{U}_p .

If the sample \mathbf{x} contains outliers, Eq. (3) does not yield a reliable reconstruction; a robust method is required (e.g., [17–20]). But since these methods are computationally very expensive (i.e., they are based on iterative algorithms) or can handle only a small amount of noise, they are often not applicable in practice. Thus, in the following we propose a new fast robust PCA approach (FR-PCA), which overcomes these problems.

2.1 FR-PCA Training

The training procedure, which is sub-divided into two major parts, is illustrated in Figure 1. First, a standard PCA subspace \mathbf{U} is generated using the full available training data. Second, N sub-samplings \mathbf{s}_n are established from randomly selected values from each data point (illustrated by the red points and the green crosses). For each sub-sampling \mathbf{s}_n , a smaller subspace (sub-subspace) \mathbf{U}^n is estimated, in addition to the full subspace.

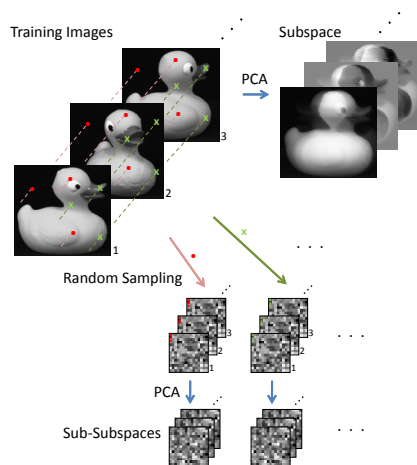


Fig. 1. FR-PCA training: A global PCA subspace and a large number of smaller PCA sub-subspaces are estimated in parallel. Sub-subspaces are derived by randomly subsampling the input data.

2.2 FR-PCA Reconstruction

Given a new unseen test sample \mathbf{x} , the robust reconstruction $\tilde{\mathbf{x}}$ is estimated in two stages. In the first stage (*gross outlier detection*), the outliers are detected based on the reconstruction errors of the sub-subspaces. In the second stage (*refinement*), using the thus estimated inliers, a robust reconstruction $\tilde{\mathbf{x}}$ of the whole sample is generated.

In the *gross outlier detection*, first, N sub-samplings \mathbf{s}_n are generated according to the corresponding sub-subspaces \mathbf{U}^n , which were estimated as described in Section 2.1. In addition, we define the set of ‘‘inliers’’ \mathbf{r} as the union of all selected pixels: $\mathbf{r} = \mathbf{s}_1 \cup \dots \cup \mathbf{s}_N$, which is illustrated in Figure 2(a) (green points). Next, for each sub-sampling \mathbf{s}_n a reconstruction $\tilde{\mathbf{s}}_n$ is estimated by Eq. (3), which allows to estimate the error-maps

$$\mathbf{e}_n = |\mathbf{s}_n - \tilde{\mathbf{s}}_n|, \quad (4)$$

the mean reconstruction error \bar{e} over all sub-samplings, and the mean reconstruction errors \bar{e}_n for each of the N sub-samplings.

Based on these errors, we can detect the outliers by local and global thresholding. The local thresholds (one for each sub-sampling) are defined by $\theta_n = \bar{e}_n w_n$, where w_n is a weighting parameter and the global threshold θ is set to the mean error \bar{e} . Then, all points $s_{n,(i,j)}$ for which

$$e_{n,(i,j)} > \theta_n \quad \text{or} \quad e_{n,(i,j)} > \theta \quad (5)$$

are discarded from the sub-samplings \mathbf{s}_n obtaining $\hat{\mathbf{s}}_n$. Finally, we re-define the set of ‘‘inliers’’ by

$$\mathbf{r} = \hat{\mathbf{s}}_1 \cup \dots \cup \hat{\mathbf{s}}_q, \quad (6)$$

where $\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_q$ indicate the first q sub-samplings (sorted by \bar{e}_n) such that $|\mathbf{r}| \leq k$; k is the pre-defined maximum number of points. The thus obtained ‘‘inliers’’ are shown in Figure 2(b).

The *gross outlier detection* procedure allows to remove most outliers. Thus, the obtained set \mathbf{r} contains almost only inliers. To further improve the final result in the *refinement* step, the final robust reconstruction is estimated similar to [19]. Starting from the point set $\mathbf{r} = [r_1, \dots, r_k]$, $k > p$, obtained from the *gross outlier detection*, repeatedly reconstructions $\tilde{\mathbf{x}}$ are computed by solving an over-determined system of equations minimizing the least squares reconstruction error

$$E(\mathbf{r}) = \sum_{i=1}^k \left(x_{r_i} - \sum_{j=1}^p a_j \mathbf{u}_{j,r_i} \right)^2. \quad (7)$$

Thus, in each iteration those points with the largest reconstruction errors can be discarded from \mathbf{r} (selected by a reduction factor α). These steps are iterated until a pre-defined number of remaining points is reached. Finally, an outlier-free subset is obtained, which is illustrated in Figure 2(c).




Fig. 2. Data point selection process: (a) data points sampled by all sub-subspaces, (b) occluded image showing the remaining data points after applying the sub-subspace procedure, and (c) resulting data points after the iterative refinement process for the calculation of the PCA coefficients. This figure is best viewed in color.

A robust reconstruction result obtained by the proposed approach compared to a non-robust method is shown in Figure 3. One can clearly see that the robust method considerably outperforms the standard PCA. Note, the blur visible in the reconstruction of the FR-PCA is the consequence of taking into account only a limited number of eigenvectors.




Fig. 3. Demonstration of the insensitivity of the robust PCA to noise (i.e., occlusions): (a) occluded image, (b) reconstruction using standard PCA, and (c) reconstruction using the FR-PCA.

In general, the robust estimation of the coefficients is computationally very efficient. In the *gross outlier detection* procedure, only simple matrix operations have to be performed, which are very fast; even if hundreds of sub-subspace reconstructions have to be computed. The computationally more expensive part is the *refinement* step, where repeatedly an overdetermined linear system of equations has to be solved. Since only very few refinement iterations have to be performed due to the preceding *gross outlier detection*, the total runtime is kept low.

3 Experimental Results

To show the benefits of the proposed fast robust PCA method (FR-PCA), we compare it to the standard PCA (PCA) and the robust PCA approach presented in [19] (R-PCA). We choose the latter one, since it yields superior results among the presented methods in the literature and our refinement process is similar to theirs.

In particular, the experiments are evaluated for the task of robust image reconstruction on the "Amsterdam Library of Object Images (ALOI)" database [21]. The ALOI database consists of 1000 different objects. Over hundred images of each object are recorded under different viewing angles, illumination angles and illumination colors, yielding a total of 110,250 images. For our experiments we arbitrarily choose 30 categories (009, 018, 024, 032, 043, 074, 090, 093, 125, 127, 135, 138, 151, 156, 171, 174, 181, 200, 299, 306, 323, 354, 368, 376, 409, 442, 602, 809, 911, 926), where an illustrative subset of objects is shown in Figure 4.



Fig. 4. Illustrative examples of ALOI database objects [21] used in the experiments.

In our experimental setup, each object is represented in a separate subspace and a set of 1000 sub-subspaces, where each sub-subspace contains 1% of data points of the whole image. The variance retained for the sub-subspaces is 95% and 98% for the whole subspace, which is also used for the standard PCA and the R-PCA. Unless otherwise noted, all experiments are performed with the parameter settings given in Table 1.

Table 1. Parameters for the FR-PCA (a) and the R-PCA (b) used for the experiments.

(a)		(b)	
FR-PCA		R-PCA	
Number of initial points k	130 p	Number of initial hypotheses H	30
Reduction factor α	0.9	Number of initial points k	48 p
		Reduction factor α	0.85
		K_2	0.01
		Compatibility threshold	100

A 5-fold cross-validation is performed for each object category, resulting in 80% training- and 20% test data, corresponding to 21 test images per iteration. The experiments are accomplished for several levels of spatially coherent occlusions and several levels of salt & pepper noise. Quantitative results for the

root-mean-squared (RMS) reconstruction-error per pixel for several levels of occlusions are given in Table 2. In addition, in Figure 5 we show box-plots of the RMS reconstruction-error per pixel for different levels of occlusions. Analogously, the RMS reconstruction-error per pixel for several levels of salt & pepper noise is presented in Table 3 and the corresponding box-plots are shown in Figure 6.

Table 2. Comparison of the reconstruction errors of the standard PCA, the R-PCA and the FR-PCA for several levels of occlusion showing RMS reconstruction-error per pixel given by mean and standard deviation.

Occlusion	Error per Pixel											
	0%		10%		20%		30%		50%		70%	
	mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
PCA	10.06	6.20	21.82	8.18	35.01	12.29	48.18	15.71	71.31	18.57	92.48	18.73
R-PCA	11.47	7.29	11.52	7.31	12.43	9.24	22.32	21.63	59.20	32.51	94.75	43.13
FR-PCA	10.93	6.61	11.66	6.92	11.71	6.95	11.83	7.21	26.03	23.05	83.80	79.86

Table 3. Comparison of the reconstruction errors of the standard PCA, the R-PCA and the FR-PCA for several levels of salt & pepper noise showing RMS reconstruction-error per pixel given by mean and standard deviation.

Salt&Pepper Noise	Error per Pixel									
	10%		20%		30%		50%		70%	
	mean	std	mean	std	mean	std	mean	std	mean	std
PCA	11.77	5.36	14.80	4.79	18.58	4.80	27.04	5.82	36.08	7.48
R-PCA	11.53	7.18	11.42	7.17	11.56	7.33	11.63	7.48	15.54	10.15
FR-PCA	11.48	6.86	11.30	6.73	11.34	6.72	11.13	6.68	14.82	7.16

From Table 2 and Figure 5 it can be seen – starting from an occlusion level of 0% – that all subspace methods exhibit nearly the same RMS reconstruction-error. Increasing the portion of occlusion, the standard PCA shows large errors whereas the robust methods are still comparable to the non-disturbed (best feasible) case, where our novel FR-PCA presents the best performance. In contrast, as can be seen from Table 3 and Figure 6, all methods can generally cope better with salt & pepper noise. However, also for this experiment FR-PCA yields the best results.

Finally, we evaluated the runtime¹ for the applied different PCA reconstruction methods, which are summarized in Table 4. It can be seen that for the given setup compared to R-PCA for a comparable reconstruction quality the robust reconstruction can be speeded up by factor of 18! This drastic speed-up can be explained by the fact that the refinement process is started from a set of data points mainly consisting of inliers. In contrast, in [19] several point sets (hypotheses) have to be created and the iterative procedure has to be run for every

¹ The runtime is measured in MATLAB using an Intel Xeon processor running at 3GHz. The resolution of the images is 192x144 pixels.

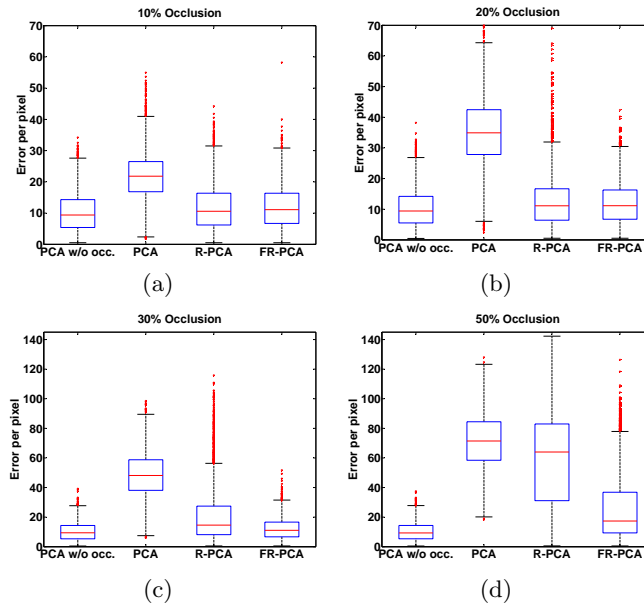


Fig. 5. Box-plots for different levels of occlusions for the RMS reconstruction-error per pixel. PCA without occlusion is shown in every plot for the comparison of the robust methods to the best feasible reconstruction result.

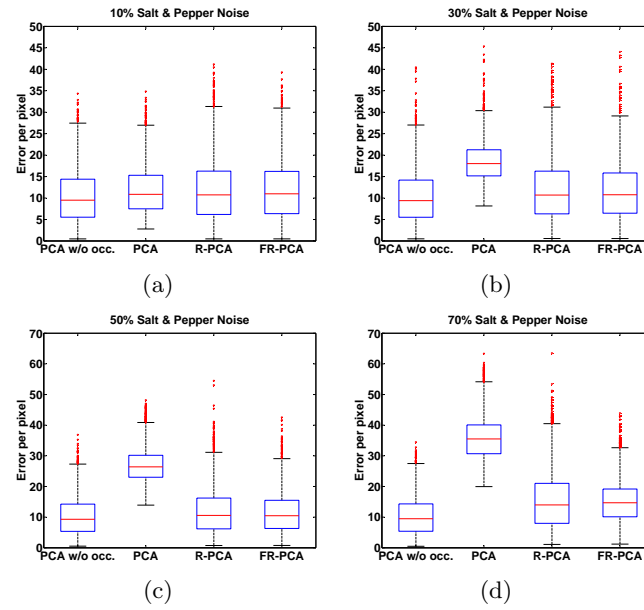


Fig. 6. Box-plots for different levels of salt & pepper noise for the RMS reconstruction-error per pixel. PCA without occlusion is shown in every plot for the comparison of the robust methods to the best feasible reconstruction result.

set resulting in a poor runtime performance. Reducing the number of hypotheses or the number of initial points would decrease the runtime, but, however, the reconstruction accuracy gets worse. In particular, the runtime of our approach only depends slightly on the number of starting points, thus having nearly constant execution times. Clearly, the runtime depends on the number and size of used eigenvectors. Increasing one of those values, the gap between the runtime for both methods is even getting larger.

Table 4. Runtime comparison. Compared to R-PCA, FR-PCA speeds-up the computation by a factor of 18.

Occlusion	Mean Runtime [s]					
	0%	10%	20%	30%	50%	70%
PCA	0.006	0.007	0.007	0.007	0.008	0.009
R-PCA	6.333	6.172	5.435	4.945	3.193	2.580
FR-PCA	0.429	0.338	0.329	0.334	0.297	0.307

4 Conclusion

We developed a novel fast robust PCA (FR-PCA) method based on an efficient two-stage outlier detection procedure. The main idea is to estimate a large number of small PCA sub-subspaces from a subset of points in parallel. Thus, for a given test sample, those sub-subspaces with the largest errors are discarded first, which reduce the number of outliers in the input data (gross outlier detection). This set – almost containing inliers – is then used to robustly reconstruct the sample by minimizing the least square reconstruction error (refinement). Since the gross outlier detection is computationally much cheaper than the refinement, the proposed method drastically decreases the computational effort for the robust reconstruction. In the experiments, we show that our new fast robust PCA approach outperforms existing methods in terms of speed and accuracy. Thus, our algorithm is applicable in practice and can be applied for real-time applications such as robust Active Appearance Model (AAM) fitting [22]. Since our approach is quite general, FR-PCA is not restricted to robust image reconstruction.

Acknowledgments

This work has been funded by the Biometrics Center of Siemens IT Solutions and Services, Siemens Austria. In addition, this work was supported by the FFG project AUTOVISTA (813395) under the FIT-IT programme, and the Austrian Joint Research Project Cognitive Vision under projects S9103-N04 and S9104-N04.

References

1. Jolliffe, I.T.: *Principal Component Analysis*. Springer (2002)
2. Murase, H., Nayar, S.K.: Visual learning and recognition of 3-d objects from appearance. *Intern. Journal of Computer Vision* **14**(1) (1995) 5–24
3. Kirby, M., Sirovich, L.: Application of the karhunen-loeve procedure for the characterization of human faces. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **12**(1) (1990) 103–108
4. Turk, M., Pentland, A.: Eigenfaces for recognition. *Journal of Cognitive Neuroscience* **3**(1) (1991) 71–86
5. Wang, Y., Huang, K., Tan, T.: Human activity recognition based on r transform. In: *Proc. CVPR*. (2008)
6. Tai, Y.W., Brown, M.S., Tang, C.K.: Robust estimation of texture flow via dense feature sampling. In: *Proc. CVPR*. (2007)
7. Lee, S.M., Abbott, A.L., Araman, P.A.: Dimensionality reduction and clustering on statistical manifolds. In: *Proc. CVPR*. (2007)
8. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. *Computer Vision and Image Understanding* **61** (January 1995) 38–59
9. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **23**(6) (2001) 681–685
10. Huber, P.J.: *Robust Statistics*. John Wiley & Sons (2004)
11. Hampel, F.R., Ronchetti, E.M., Rousseeuw, P.J., Stahel, W.A.: *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons (1986)
12. Xu, L., Yuille, A.L.: Robust principal component analysis by self-organizing rules based on statistical physics approach. *IEEE Trans. on Neural Networks* **6**(1) (1995) 131–143
13. Torre, F.d., Black, M.J.: A framework for robust subspace learning. *Intern. Journal of Computer Vision* **54**(1) (2003) 117–142
14. Roweis, S.: EM algorithms for PCA and SPCA. In: *Advances in Neural Information Processing Systems*. (1997) 626–632
15. Tipping, M.E., Bishop, C.M.: Probabilistic principal component analysis. *Journal of the Royal Statistical Society B* **61** (1999) 611–622
16. Skočaj, D., Bischof, H., Leonardis, A.: A robust PCA algorithm for building representations from panoramic images. In: *Proc. European Conf. on Computer Vision*. Volume IV. (2002) 761–775
17. Rao, R.: Dynamic appearance-based recognition. In: *Proc. CVPR*. (1997) 540–546
18. Black, M.J., Jepson, A.D.: Eigenttracking: Robust matching and tracking of articulated objects using a view-based representation. In: *Proc. European Conf. on Computer Vision*. (1996) 329–342
19. Leonardis, A., Bischof, H.: Robust recognition using eigenimages. *Computer Vision and Image Understanding* **78**(1) (2000) 99–118
20. Edwards, J.L., Murase, J.: Coarse-to-fine adaptive masks for appearance matching of occluded scenes. *Machine Vision and Applications* **10**(5–6) (1998) 232–242
21. Geusebroek, J.M., Burghouts, G.J., Smeulders, A.W.M.: The Amsterdam Library of Object Images. *International Journal of Computer Vision* **61**(1) (January 2005) 103–112
22. Storer, M., Roth, P.M., Urschler, M., Bischof, H., Birchbauer, J.A.: Active appearance model fitting under occlusion using fast-robust PCA. In: *Proc. International Conference on Computer Vision Theory and Applications (VISAPP)*. Volume 1. (February 2009) 130–137