Robust Multiple Model Fitting with Preference Analysis and Low-rank Approximation

Luca Magri
luca.magri@unimi.it

Andrea Fusiello
andrea.fusiello@uniud.it

Department of Mathematics
Università degli Studi di Milano
Via Saldini, 50 - 20133 Milano, IT
DIEGM
Università degli Studi di Udine
Via delle Scienze, 208 - 33100 Udine, IT

Abstract

This paper deals with the extraction of multiple models from outlier-contaminated data. The method we present is based on preference analysis and low rank approximation. After representing points in a conceptual space, Robust PCA (Principal Component Analysis) and Symmetric NMF (Non negative Matrix Factorization) are employed to reduce the multi-model fitting problem to many single-fitting problems, which in turn are solved with a strategy that resembles MSAC (M-estimator SAmple Consensus). Experimental validation on public, real data-sets demonstrates that our method compares favourably with the state of the art.

1 Introduction

Geometric multi-model fitting aims at extracting parametric models from unstructured data in order to organize and aggregate visual content in suitable higher-level geometric structures. This ubiquitous task can be encountered in many Computer Vision applications, for example in 3D reconstruction, where it is employed to estimate multiple rigid moving objects in order to initialize multi-body structure from motion, or in the processing of 3D point clouds, where planar patches are fitted to produce intermediate geometric interpretations.

In practice, it is necessary to overcome the "chicken-&-egg dilemma" inherent to this problem: in order to estimate models one needs to first segment the data, but in order to segment the data it is necessary to know the models associated with each data point. The presence of multiple structures hinders robust estimation, which has to cope with both *gross* outliers and *pseudo*-outliers.

1.1 Related work.

Among the wide variety of algorithms for robust estimation proposed in Computer Vision, the analysis of consensus together with its dual counterpart, the analysis of preferences,

can be traced as a leitmotiv connecting the extensive literature on geometric fitting. The consensus of a model is defined as the set of data points that fit the model within a certain inlier threshold ε; analogously, the preference set of a point is the set of models having that point as an inlier.

Consensus analysis   Consensus-oriented methods instantiate a pool of provisional model hypotheses by random sampling, hence retain the models that better explain the data by inspecting their consensus sets. This idea is at the core of the celebrated RANSAC (Random Sample Consensus) algorithm and its relaxations, namely MSAC (M-estimator Sample Consensus) and MLESAC (Maximum Likelihood Estimation Sample Consensus) [27]. Many improvements of the RANSAC paradigm have been proposed in the literature, *e.g.* [8, 17], the interest reader can find a nice overview on all these methods in [6]. In the case of multiple models, Multi-RANSAC [34] and its variants [10, 20] rely on the same principle; also the popular Hough transform and the randomized version [30] can be seen as consensus-based algorithms, where models are revealed as peaks obtaining higher consensus in a quantized hypothesis space. Moreover, maximizing the consensus set of models is the foundation of many optimization algorithms [12] designed for geometric fitting.

Preference analysis   In a dual fashion, preference analysis, introduced by Residual Histogram Analysis [33], reverses the role of data and models: rather than considering models and inspecting which points match them, the preference set of individual data points are examined. More precisely, the residuals of each data point are taken into account in order to build a *conceptual space* in which points are portrayed by the *preferences* they have accorded to random provisional models. J-Linkage [26] and T-Linkage [19] belong to this category. In particular, these two algorithms share the same *first-represent-then-segment* scheme: at first data are represented respectively either as characteristic functions or as continuous ones taking values on the hypothesized models, then the conceptual representations are segmented by a greedy bottom-up clustering step exploiting either the Jaccard [13] or the Tanimoto [25] distances in order to measure the agreement between preferences, and using the fact that preferences of inliers from the same structure are correlated.

RCMSA (Random Cluster Model Simulated Annealing) [21] as well takes advantage of this idea representing data points as permutations on a set of tentative models constructed iteratively, using subsets larger than minimal. Point preferences are organized in a weighted graph and the multi-model fitting task is stated as a graph cut problem which is solved efficiently in an annealing framework.

In addition, a stream of investigations focused on higher order clustering [1, 11, 14, 31] implicitly adopts a preference based approach. In these works higher order similarity tensors are defined between n-tuple of points as the probability that these points are clustered together. In practice this measure is approximated exploiting the residual error of the n points wrt provisional models; this preference information is encapsulated in hypergraphs or multiway order tensors, which are properly reduced to pairwise similarity and fed to spectral clustering-like segmentation algorithms. For example, Sparse Grassmann Clustering (SGC) [14] approximates the multi-way similarity tensor as the Gramian matrix defined by the inner product of points in the preference space. Hence, following the spirit of spectral clustering, which works only with a few eigenvectors of the similarity matrix , the Gramian is projected on its best low rank approximation in a least square sense, using Grouse [3]. At the end data are segmented applying k-means.

1.2   Paper contributions

In this paper we present our attempt to solve the chicken-&-egg dilemma by conceiving an original method aimed at geometric multi model fitting. The main idea, pictorially represented in Fig.1, is to build on the preference analysis exploiting a robust M-estimator and integrating in this approach decomposition techniques, such as Robust Principal Component Analysis (Robust PCA) and Nonnegative Matrix Factorization (NMF).

Loosely speaking, our method can be thought as a sort of "*robust* spectral clustering". It is well known [28] that spectral clustering produces accurate segmentations in two steps: at first data are projected on the space of the first eigenvectors of the Laplacian matrix and then k-means is applied. The shortcoming of this approach is that it is not robust to outliers. We propose to follow the same scheme enforcing robustness: the eigen-decomposition step is replaced by Robust PCA on a pairwise affinity matrix, and Symmetric NMF [16] plays the role of k-means. In this way we are able to reduce the multi-model fitting problem to many single-fitting problems which are solved by scrutinizing the product between the matrix produced by Symmetric NMF and the preference matrix, together with the use of robust statistics.


Figure 1: Our method in a nutshell: data points are shifted in a conceptual space where they are framed as a preference matrix Φ. A similarity matrix K is defined exploiting agreement between preferences. Robust PCA and Symmetric NMF are used to robustly decompose $K = UU^\top + S$, where S is a sparse matrix modelling outliers, and U is a low rank matrix representing the segmentation. Finally, models are extracted inspecting the product of the preference matrix with thresholded U, mimicking the MSAC strategy. (Points are ordered by cluster for visualization purposes)

2   Preference analysis

Along the same line of T-Linkage, our method starts shifting the data points $X = \{x_1, \ldots, x_n\}$, from their ambient space to a conceptual one, where they are portrayed as vectors of preferences. This representation can be formalized by a vectorial map:

$$\phi = (\phi_1, \ldots, \phi_m)^\top : X \to [0,1]^m \ , \tag{1}$$

defined component-wise as the Cauchy weighting function:

$$\phi_j(x) = \frac{1}{1 + \left(\frac{d(x,h_j)}{\theta \sigma_n}\right)^2} \quad j = 1, \ldots, m \ , \tag{2}$$

where d measures the distance between a datum x and a provisional model h_j, σ_n is an estimate of the standard deviation of the residuals and θ is a tuning constant, set to 5.0 in all experiments. Tentative models are instantiated by drawing at random m minimal sample sets (MSS), i.e. the minimum-sized set of data points necessary to estimate a model, as further described below.

The rationale behind this construction is that the map ϕ collects in a vector $\Phi^i = \phi(x_i) \in \mathbb{R}^m$ the preferences granted by x to the models h_j $(j = 1, \ldots, m)$ with a vote in $[0, 1]$ according to its residuals. We indicate with $\Phi = [\Phi^1, \ldots, \Phi^n] \in \mathbb{R}^{m \times n}$ the matrix whose columns are the point coordinates in the preference space. Rows corresponds to provisional models. The agreement between the preferences of two points in the conceptual space reveals the multiple structures hidden in the data: points sharing the same preferences are likely to belong to the same structures. This notion is captured by defining a positive semi-definite kernel matrix $K \in [0, 1]^{n \times n}$ on Φ:

$$K(i, j) = \exp\left(-\tau(i, j)^2\right) \quad \text{where} \quad \tau(i, j) = 1 - \frac{\langle \Phi^i, \Phi^j \rangle}{\|\Phi^i\|^2 + \|\Phi^j\|^2 - \langle \Phi^i, \Phi^j \rangle}. \tag{3}$$

Observe that Φ endowed with this kernel resembles in spirit the conceptual space adopted in [19], where the notions of consensus and preference sets are relaxed to continuous functions and the Tanimoto distance τ is used to measure preference agreement. However some differences can be pointed out. In fact ϕ is a continuous robust weighting function (of the type employed in M-estimators), whereas the embedding proposed in T-Linkage uses the inlier threshold to cut off the preferences of points having distance greater than ε; moreover the kernel K replaces the use of the metric τ.

Φ can also be thought as a weighted hyper-graph: hyper-edges represent sampled models and weights correspond to point preferences. In this interpretation the role of the kernel matrix K is to reduce the multi-way similarity to a pairwise affinity.

Biased Random Sampling. Different sampling strategies can be used to generate the pool of provisional models. If uniform sampling is employed, a large number of trials is required for reaching a reasonable probability of hitting at least a *pure* (i.e., outlier free) MSS per model, as explained in [26]. Hence, many strategies have been proposed in order to guide sampling towards promising models both in the case of one model [7, 9] and in the multiple models scenario [5]. With localized sampling in the data space [15] neighbouring points are selected with higher probability, thereby reducing the number of hypotheses that have to be generated. However the assumption that points of the same structure are close in the ambient space is not satisfied in every applicative scenario (e.g in homography fitting, points belonging to the same plane may be far apart in the image) and introducing a local bias requires additional prior knowledge on the data distributions in the ambient space.

On the contrary in the preference space it is possible to exploit the Tanimoto distance, which always ranges in the interval $[0, 1]$ in order to promote the extraction of inliers of the same structure using conditional sampling. In this way the probability of sampling outliers is considerably reduced. In fact outliers can be recognized as the most separate points in

the conceptual space, because their preferences deviate significantly from the rest of the data resulting in higher Tanimoto distances.

For this reason, following the approach proposed in [5], we sample the hypotheses directly in our conceptual space. This can be easily done performing a preliminary uniform sampling of hypotheses, hence representing the data in the preference space according to these provisional models and then drawing MSS by biased sampling in the conceptual space.

In particular if a point x_i has already been selected, then a point x_j such that $i \neq j$ has the following probability of being drawn:

$$P(x_i|x_j) = \frac{1}{Z} \exp\left(\frac{\tau(i,j)}{\alpha}\right)^2 . \tag{4}$$

where Z is a normalization constant and α controls the local bias in the conceptual space. Setting α to the median of point-to-point Tanimoto distances we obtained results comparable to localized sampling, without the need of a fine tuning.

The kernel matrix K is then updated on the fly based on the hypotheses already sampled.

3 Clustering

We shall now describe how the affinity matrix can be exploited to segment the data. Consider an ideal affinity $n \times n$ matrix F which encodes point membership to the same segment: $F_{i,j} = 1$ if x_i and x_j are clustered together and $F_{i,j} = 0$ otherwise. If data belonging to the same segment are arranged as consecutive points, the matrix F exhibits a block structure and therefore has rank k equal to the number of clusters in the data.

As described in [31, 32] the problem of partitioning a set of data points in k segments starting from a positive semi-definite affinity matrix K is equivalent to approximating K in a least square sense by means of an ideal affinity matrix F. In formulae, denoting by $\|\cdot\|_F$ the Frobenius norm of a matrix, we are interested in:

$$\min_F \|K - F\|_F^2 , \tag{5}$$

under conditions on F to be further specified. This problem is usually formulated by introducing a matrix $U \in \mathbb{R}^{n \times k}$ such that $F = UU^\top$, which represents a soft segmentation of the data: the element U_{ij} measures the probability that the i-th point belongs to the j-th segment.

According to the constraints imposed on U, the solution of (5) corresponds to different classical clustering algorithms, such as spectral clustering or k-means. These constraints are: $U \geq 0$, $\text{rank}(U) = k$, $U^\top U = I$ and $UU^\top$ is doubly stochastic. Hard-clustering assignment implies orthogonality; being doubly stochastic represents a balancing condition on the sizes of the clusters; the non negativity of U ensures physical meaning of the entry of U which can be interpreted of probability of points to belong to a given segment.

The last constraint is the most important according to [31, 32], where it is claimed the key ingredients for solving Problem (5) are the low-rank nature of both the affinity matrix and U (since since $k \ll n$), together with the non-negativeness of U.

Symmetric NMF (SymNMF) [16], that recently stands out in the clustering literature, enforces exactly these two proprieties. The idea at the basis of SymNMF is to rephrase (5) in the equivalent formulation

$$\min_{U \in \mathbb{R}_+^{n \times k}} \|K - UU^\top\|_F^2 \tag{6}$$

and hence to find U minimizing (6) using an improved Newton-like algorithm that exploits the second-order information efficiently.

When data is contaminated by gross outliers K has no longer low rank. For this reason, before applying SymNMF, we search *robustly* for the lowest-rank matrix L and the column-sparsest matrix S such that the data matrix can be decomposed as

$$K = L + S \ . \tag{7}$$

This Robust PCA step mimics in a outlier-resilient way the projection of data on the space of k eigenvectors of the similarity matrix performed in spectral clustering. The decomposition is obtained with the Augmented Lagrangian Method (ALM) [18], which solves the problem

$$\arg\min \|L\|_* + \lambda \|S\|_1 \qquad \text{s.t} \qquad K = L + S \ , \tag{8}$$

where $\|\cdot\|_*$ denotes the nuclear norm and $\|\cdot\|_1$ the sum of absolute values. The parameter λ has a provable optimal value [4] at $\lambda = \frac{1}{\sqrt{n}}$, where n is the dimension of the square matrix K.

Please note that this approximation differs from the one adopted by SGC [14] where a low rank space is fit to a Gramian matrix in a least square sense. We depart from this model because a least squares fit is reliable as long as the sampled hypotheses are pure, but this property can not be ensured in presence of outliers.

We can now apply the SymNMF machinery to L (instead of K) in order to find a completely positive rank-k factorization $L = UU^\top$. A segmentation is obtained from U by considering the matrix B with the same dimension of U that has a one in each row where U achieves its row-maximum, and zero otherwise: $B_{i,j} = 1$ means that point i belongs to segment j. This last step is similar to the customary k-means that comes at the end of spectral clustering.

At this point, the matrix B represents a provisional segmentation of the points into k segments containing outliers. The goal of the next section is to refine this segmentation and prune outliers, by solving, within each segment, a robust *single model* fitting.

4 Pruning outliers

Model extraction. Let us first observe that $\Phi \mathbb{1}$ (where $\mathbb{1}$ is a vector of ones) is the sum of the preference vectors of all the points in Φ , so its entries are the votes obtained by each model. Hence finding the maximal entry of $\Phi \mathbb{1}$ is equivalent to doing a sort of MSAC (M-estimator SAmple and Consensus) with the Cauchy weighting function (Eq. 2).

We have seen that columns of $B = [B^1, \ldots, B^k]$ can be regarded as indicators of the segments. Hence ΦB^i is the sum of the preference vectors of the points in the segment i, and its maximal entry represents the most preferred model in that segment. Therefore, the maximum over the columns of ΦB are the indices of the models in Φ that achieve maximum consensus in each segment. According to the observation above, this is equivalent to running a MSAC within each segment i with preference matrix $(\Phi \operatorname{diag}(B^i))$. The above reasoning can be extended to the matrix $U \circ B$ with entries in $[0,1]$, that corresponds to a soft segmentation in which outliers are under-weighted ($\circ$ denotes the component-wise or Hadamard product).

We found beneficial, prior to this step, to augment Φ with some pure models by random sampling and to remove "spurious" ones, according to the segmentation represented by B. In particular, we relax the concept of "spurious" to those models that are not contained in a

single segment with at least 50% of their points; in other words, we label the points in Φ according to the segmentation given by B and we remove the columns where no label occurs more than 50% of the times. The new sampling is implemented by drawing random MSS within each segment i with probabilities given by the non-zero entries of $(U \circ B)^i$.

In summary, the maximal entry in each column of $\Phi(U \circ B)$ corresponds to the index of the most preferred model by the points of the segment, hence we choose it as the model that represents the segment. This could be a final result if the goal was to find the correct models. However, having recognized the entangled nature of model fitting and segmentation problems, we will unravel it by iterating between refining the model and updating the segmentation.

Segmentation. The models computed from $\max_{\text{cols}}(\Phi(U \circ B))$ define a new tentative segmentation by assigning points to the nearest model. Within this segmentation, outliers are singled-out as points with a residual higher than a threshold $T = \theta \, \widehat{\sigma}$ where $\widehat{\sigma}$ is an estimate of the standard deviation of the residuals *of the inliers* and θ is the same tuning constant as in Eq.1 (set to 5.0 in our experiments).

The value of $\widehat{\sigma}$ can be obtained in several ways: it can be user provided ($\widehat{\sigma} = \sigma_n$) or can be computed from the residuals themselves, in a robust way. The second solution is to be preferred, as it leaves the choice of σ_n a noncritical step and makes the threshold T data-adaptive. We preferred the S_n estimator proposed in [23]:

$$S_n = c \, \text{med}_i(\text{med}_j(|r_i - r_j|)), \tag{9}$$

(where r_i, $i = 1, \ldots, n$ denotes the residual between the data x_i and the considered model) as a valid alternative to the more common median absolute deviation (MAD), which is aimed at symmetric distributions, and has a low (37%) Gaussian efficiency, . S_n instead copes with skewed distributions, has the same breakdown as MAD but a higher Gaussian efficiency (58%).

The factor c can be set to 1.1926 for consistency with a normal distribution, but other distributions require different values (see [23] for details). In our experiments it has been tuned heuristically by analysing the distribution of the residuals of inliers given by the ground-truth. Values are reported in Tab.1.

We noticed that in some cases most of the outliers are assigned to a single segment, resulting in a contamination greater than 50% that inevitably skews S_n. As a guard against this, S_n is computed only on the residuals smaller than $5.0\sigma_n$.

The model is then refined with a least-squares fit on the inliers, and the threshold T is computed again to determine the final segmentation.

5 Experimental evaluation

In this section, we assess experimentally the effectiveness of our algorithm, henceforth dubbed RPA. All the code is written in Matlab and is available for download[1]. We used the inexact ALM code [2], whereas the SymNMF implementation is taken from [16].

Accurately quantifying the performance of multi-model fitting is a difficult task. As mentioned *en passant* in the introduction, the problem has a dual nature, depending on whether one considers as its output the models (that fit the data) or the segmentation (induced by the

---

[1] http://www.diegm.uniud.it/fusiello/demo/rpa/

models). The evaluation metric changes accordingly. The definition of a similarity measure between estimated models and ground-truth ones can turn to be an elusive task, for this reason this problem is usually tackled from the classification point of view and metrics based on hard assignment are usually adopted. A popular choice is the misclassification error (ME), which measures the percentage of misclassified point with respect to a ground-truth segmentation in models and outliers.

Experiment	σ_n	c	s	β
Motion segmentation	0.005	1.53	0.005	100
Planar segmentation	0.013	2.11	0.005	10

Table 1: Parameters used in the experiments. σ_n is the overall standard deviation of the residuals of the inliers, as computed from ground-truth (units refer to normalized image coordinates). c is the value in Eq.(9) that experimentally provides the best estimate of σ_n from S_n. Parameters s and β refer to [21] and the values are the ones provided by the authors in their implementation.

	k	%out	T-lnkg	RCMSA	RPA		k	%out	T-lnkg	RCMSA	RPA
biscuitbookbox	3	37.21	**3.10**	16.92	3.88	unionhouse	5	18.78	48.99	**2.64**	10.87
breadcartoychips	4	35.20	14.29	25.69	**7.50**	bonython	1	75.13	**11.92**	17.79	15.89
breadcubechips	3	35.22	**3.48**	8.12	5.07	physics	1	46.60	29.13	48.87	**0.00**
breadtoycar	3	34.15	9.15	18.29	**7.52**	elderhalla	2	60.75	10.75	29.28	**0.93**
carchipscube	3	36.59	**4.27**	18.90	6.50	ladysymon	2	33.48	**24.67**	39.50	24.67
cubebreadtoychips	4	28.03	9.24	13.27	**4.99**	library	2	56.13	**24.53**	40.72	31.29
dinobooks	3	44.54	20.94	23.50	**15.14**	nese	2	30.29	7.05	46.34	**0.83**
toycubecar	3	36.36	15.66	13.81	**9.43**	sene	2	44.49	7.63	20.20	**0.42**
biscuit	1	57.68	16.93	14.00	**1.15**	napiera	2	64.73	28.08	31.16	**9.25**
biscuitbook	2	47.51	**3.23**	8.41	3.23	hartley	2	62.22	21.90	37.78	**17.78**
boardgame	1	42.48	21.43	19.80	**11.65**	oldclassicswing	2	32.23	**20.66**	21.30	25.25
book	1	44.32	3.24	4.32	**2.88**	barrsmith	2	69.79	49.79	**20.14**	36.31
breadcube	2	32.19	19.31	9.87	**4.58**	neem	3	37.83	25.65	41.45	**19.86**
breadtoy	2	37.41	5.40	3.96	**2.76**	elderhallb	3	49.80	31.02	35.78	**17.82**
cube	1	69.49	7.80	8.14	**3.28**	napierb	3	37.13	**13.50**	29.40	31.22
cubetoy	2	41.42	**3.77**	5.86	4.04	johnsona	4	21.25	34.28	36.73	**10.76**
game	1	73.48	**1.30**	5.07	3.62	johnsonb	7	12.02	24.04	**16.46**	26.76
gamebiscuit	2	51.54	9.26	9.37	**2.57**	unihouse	5	18.78	33.13	**2.56**	5.21
cubechips	2	51.62	6.14	7.70	**4.57**	bonhall	6	6.43	21.84	**19.69**	41.67
mean			9.36	12.37	**5.49**	mean			24.66	28.30	**17.20**
median			7.80	9.87	**4.57**	median			23.38	29.40	**17.53**

Table 2: Misclassification error (ME %) for motion segmentation (left) and planar segmentation (right). k is the number of models and % out is the percentage of outliers. All figures are the average of the middle 3 out of 5 runs.

We deal with two applications of geometric multi model fitting on real data: motion segmentation and plane segmentation. In the *motion segmentation* experiments, given two images of the same scene composed by several objects moving independently, the aim is to fit fundamental matrices to subsets of point matches. In *plane segmentation* scenario, given two uncalibrated views of a scene, the aim is to recover the multi-planar structures by fitting homographies to point correspondences. The experiments are carried on the AdelaideRMF [29] dataset, composed of 38 image pairs (19 for motion segmentation and 19

for plane segmentation) with matching points corrupted by gross outliers. The ground-truth segmentations are also available.

We compared RPA with other two state-of-the-art multi-structure geometric fitting approaches, namely T-Linkage (available at [24]), which uses preference analysis and agglomerative clustering, and RCMSA (available at [22]), which relies on an efficient graph cut clustering based on a label optimization framework.

| toycubecar (ME=9.43) | boardgame (ME=11.65) | dinobooks (ME=15.14) |

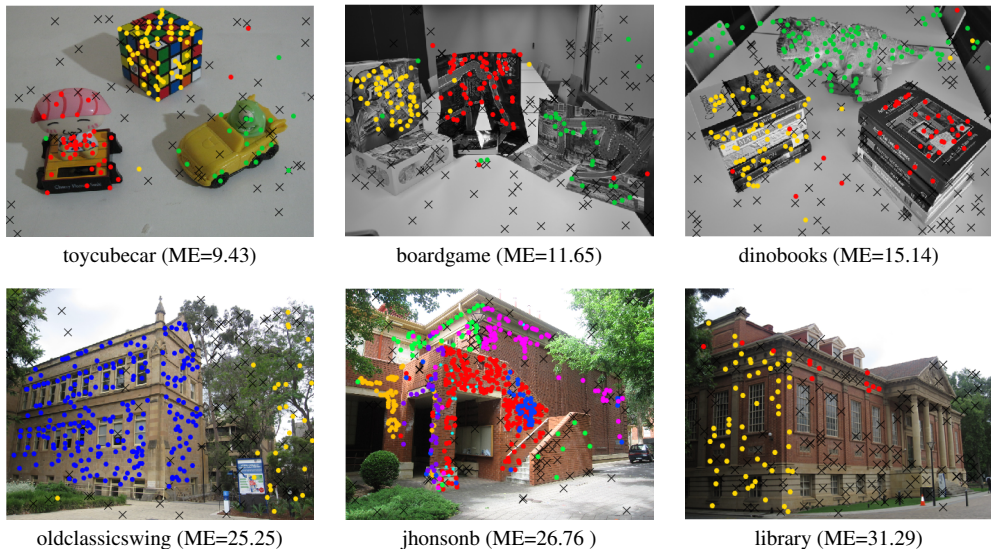| oldclassicswing (ME=25.25) | jhonsonb (ME=26.76) | library (ME=31.29) |

Figure 2: Some of the worst results obtained by RPA on motion segmentation (top row) and planar segmentation (bottom row). Model membership is colour coded, black crosses ($\times$) are outliers.

RPA and T-Linkage shared the same biased sampling in the conceptual space: we drew $3n$ hypotheses by uniform sampling and we used them to instantiate other $3n$ MSSs according to Equation (4). In all the experiments α was set to the median of all the Tanimoto distances between data points.

We provided T-Linkage with the inlier thresholds computed from the ground-truth segmentation for each single image pair, and we retained as inliers the largest k clusters, k being the correct number of models according to ground-truth. The input parameters of RCMSA and RPA are reported in Table 1 and have been kept fixed and equal for all the image pairs in each experiment.

Results are reported in Tab. 2, and demonstrate that our method outperforms its competitors, obtaining the lowest ME in most cases and the best mean and median results overall.

Some of the worst cases for RPA are reported in Fig. 2. The top row shows the results of motion segmentation that achieve the highest ME: the quality of the segmentation is nevertheless acceptable. The situation is different in the bottom row – corresponding to homography fitting – where three defective segmentations are shown, and the ME is indeed higher.

In "jhonsonb" the fault is of Symmetric NMF, which fails in finding a correct segmentation of the data, whereas in "library" and in "oldclassicswing" it is the value of σ_n that is respectively too low (over-segmentation) and too high (under-segmentation). While there are

no remedies for the first case, the last two can be cured by a better choice of σ_n: for example, the ME drops to 24.53% for "library" and to 0.55% for "oldclassicswing" after assigning to σ_n the standard deviation of the residuals of the inliers *for that specific image pair*.

6   Conclusions

In this paper we argued that preference analysis combined with robust matrix decompositions provides a versatile tool for robust geometric fitting. We proposed an approach similar in spirit to classic spectral clustering, with the advantage of being robust to outliers. Our strategy was to reduce the multi-model fitting task to many single robust model estimation problems. In particular, we conceived three levels of protection against outliers. The first one is the adoption of the Cauchy function to model points preferences. The second level appears in the robust low rank approximation, which gives rise to a soft segmentation where outliers are under-weighted. Robust extraction of models in a MSAC-like framework, together with outlier rejection based on robust scale estimates is our third guard against outliers. The value of σ_n and the number of models k are the only inputs required from the user.

Experiments have provided evidence that our method compares favourably with state of the art competing algorithms.

References

[1] Sameer Agarwal, Jongwoo Lim, Lihi Zelnik-manor, Pietro Perona, David Kriegman, and Serge Belongie. Beyond pairwise clustering. In *Computer Vision and Pattern Recognition*, pages 838–845, 2005.

[2] ALM.   http://perception.csl.illinois.edu/matrix-rank/sample_code.html.

[3] Laura Balzano, Robert Nowak, and Benjamin Recht. Online identification and tracking of subspaces from highly incomplete information. In *Proceedings of Allerton*, September 2010.

[4] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3), 2011.

[5] Tat-Jun Chin, Jin Yu, and David Suter. Accelerated hypothesis generation for multi-structure data via preference analysis. *Trans. Pattern Anal. Mach. Intell.*, pages 533–546, 2012.

[6] Sunglok Choi, Taemin Kim, and Wonpil Yu. Performance evaluation of ransac family. In *British Machine Vision Conference (BMVC)*, 2009.

[7] O. Chum and J. Matas. Randomized ransac with $T_{d,d}$ test. In *Image and Vision Computing*, volume 22, pages 837–842, 2002.

[8] Ondřej Chum and Jiří Matas. Matching with prosac-progressive sample consensus. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 220–226. IEEE, 2005.

[9] Ondrej Chum and Jiri Matas. Matching with PROSAC - progressive sample consensus. In *Computer Vision and Pattern Recognition*, pages 220–226, 2005.

[10] Orazio Gallo, Roberto Manduchi, and Abbas Rafii. Cc-ransac: Fitting planes in the presence of multiple surfaces in range data. *Pattern Recognition Letters*, 32(3):403–410, 2011.

[11] Venu M. Govindu. A Tensor Decomposition for Geometric Grouping and Segmentation. *Computer Vision and Pattern Recognition*, 1:1150–1157, 2005. ISSN 1063-6919.

[12] Hossam Isack and Yuri Boykov. Energy-based geometric multi-model fitting. *International Journal of Computer Vision*, 97(2):123–147, 2012.

[13] Paul Jaccard. *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. Impr. Corbaz, 1901.

[14] Suraj Jain and Venu Madhav Govindu. Efficient higher-order clustering on the grassmann manifold. In *International Conference on Computer Vision*, 2013.

[15] Yasushi Kanazawa and Hiroshi Kawakami. Detection of planar regions with uncalibrated stereo using distribution of feature points. In *British Machine Vision Conference*, pages 247–256, 2004.

[16] Da Kuang, Sangwoon Yun, and Haesun Park. Symnmf: nonnegative low-rank approximation of a similarity matrix for graph clustering. *Journal of Global Optimization*, pages 1–30, 2014.

[17] Karel Lebeda, Jirı Matas, and Ondrej Chum. Fixing the locally optimized ransac–full experimental evaluation. In *British Machine Vision Conference*, pages 1–11, 2012.

[18] Zhouchen Lin, Minming Chen, and Yi Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*, 2010.

[19] Luca Magri and Andrea Fusiello. T-linkage: A continuous relaxation of j-linkage for multi-model fitting. In *Computer Vision and Pattern Recognition*, pages 3954–3961, June 2014.

[20] Kai Ni, Hailin Jin, and Frank Dellaert. Groupsac: Efficient consensus in the presence of groupings. In *International Conference on Computer Vision*, pages 2193–2200, 2009.

[21] Trung-Thanh Pham, Tat-Jun Chin, Jin Yu, and David Suter. The random cluster model for robust geometric fitting. *Pattern Analysis and Machine Intelligence*, 36(8):1658–1671, 2014.

[22] RCMSA. http://cs.adelaide.edu.au/~trung/lib/exe/fetch.php?media=rcmsa_robust_fitting.zip.

[23] Peter J. Rousseeuw and Christophe Croux. Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424):1273–1283, 1993. doi: 10.1080/01621459.1993.10476408.

[24] T-linkage. http://www.diegm.uniud.it/fusiello/demo/jlk/Tlinkage_beta.zip.

[25] T Tanimoto. Technical report, IBM Internal Report, 17th Nov. 1957.

[26] R. Toldo and A. Fusiello. Robust multiple structures estimation with J-Linkage. In *European Conference on Computer Vision*, volume 5302, pages 537–547, 2008.

[27] P. H. S. Torr and A. Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78:2000, 2000.

[28] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4): 395–416, 2007.

[29] H. S. Wong, T.-J. Chin, J. Yu, and D. Suter. Dynamic and hierarchical multi-structure geometric model fitting. In *International Conference on Computer Vision*, 2011.

[30] L. Xu, E. Oja, and P. Kultanen. A new curve detection method: randomized Hough transform (RHT). *Pattern Recognition Letters*, 11(5):331–338, 1990.

[31] Ron Zass and Amnon Shashua. A unifying approach to hard and probabilistic clustering. In *International Conference on Computer Vision*, volume 1, pages 294–301, 2005.

[32] Ron Zass and Amnon Shashua. Doubly stochastic normalization for spectral clustering. In *Neural Information Processing Systems*, pages 1569–1576, 2006.

[33] Wei Zhang and Jana Kosecká. Nonparametric estimation of multiple structures with outliers. In *European Conference on Computer Vision*, volume 4358, pages 60–74, 2006.

[34] M. Zuliani, C. S. Kenney, and B. S. Manjunath. The multiRANSAC algorithm and its application to detect planar homographies. In *International Conference on Image Processing*, 2005.