



UNIVERSITÀ
DEGLI STUDI
DI UDINE

Università degli studi di Udine

Multiple structure recovery with maximum coverage

Original

Availability:

This version is available <http://hdl.handle.net/11390/1123323> since 2018-01-31T17:06:48Z

Publisher:

Published

DOI:10.1007/s00138-017-0883-x

Terms of use:

The institutional repository of the University of Udine (<http://air.uniud.it>) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

Multiple Structure Recovery with Maximum Coverage

Luca Magri · Andrea Fusiello

Received: date / Accepted: date

Abstract We present a general framework for geometric model fitting based on a set coverage formulation, that caters for intersecting structures and outliers in a simple and principled manner. The multi-model fitting problem is formulated in terms of the optimization of a consensus-based global cost function, which allows to sidestep the pitfalls of preference approaches based on clustering and to avoid the difficult trade-off between data fidelity and complexity of other optimization formulations. Two especially appealing characteristics of this method are the ease with which it can be implemented and its modularity with respect to the solver and to the sampling strategy. Few intelligible parameters need to be set and tuned, namely the inlier threshold and the number of desired models. The summary of the experiments is that our method compares favourably with its competitors overall, and it is always either the best performer or almost on par with the best performer in specific scenarios.

Keywords Set cover · Multi-model fitting · Robust estimation · Geometric fitting

1 Introduction

Finding multiple models (or structures) that fit data corrupted by noise and outliers is an omnipresent problem in empirical sciences, including Computer Vision, where organizing unstructured visual data in higher

level geometric structures is a basic and necessary step to derive better descriptions and understanding of a scene. This challenging problem has a chicken-and-egg pattern: in order to estimate models one needs to first segment the data, and in order to segment the data it is necessary to know which structure points belong to.

Moreover, the presence of multiple structures pushes robust estimation to its limit, because, in addition to rogue points, the outliers to a structure of interest include all the inliers to the other structures.

Among the wide range of methods proposed in Computer Vision to address the challenge of multiple models geometric fitting, the analysis of *consensus* together with its counterpart, the analysis of *preferences*, can be recognized as leitmotifs recurring throughout the extensive literature on the subject. The *consensus set* of a model is simply defined as the set of points that are inliers to that model. Dually, the *preference set* of a point is the set of models to which that point is inlier. Most of the multi-model fitting techniques proposed in the literature can be ascribed to one of these two concepts, according to which horn of the chicken-egg-dilemma is addressed first.

Consensus-based algorithms put the emphasis on the estimation part and focus on models that describe as many points as possible. On the other hand, preference approaches concentrate on the segmentation side of the problem, and aim at finding a proper partition of the data, from which model estimation follows.

In this paper we formulate the multi-model fitting problem within the consensus framework, and cast it into a MAXIMUM COVERAGE problem, that caters for intersecting structures and outliers in a simple and principled manner.

A preliminary version of this paper appeared as [1]. In this paper we expanded the introductory material

L. Magri
Department of Mathematics - University of Milan. Via Saldini,
50 - 20133 Milan (Italy) E-mail: magri.luca.l@gmail.com

A. Fusiello
DPIA - University of Udine. Via delle Scienze, 208 - Udine (Italy)
E-mail: andrea.fusiello@uniud.it

with additional insights, and added some new experiments on 3D point clouds.

2 Background

In this section we shall pursue the path that has been followed in the literature to address the challenges of multiple structures recovery, starting from consensus throughout preference analysis.

2.1 Consensus analysis

Consensus analysis stands out as one of the first successful efforts to address robust model estimation. The methods belonging to this category follow a common paradigm. At first the space Θ of all the feasible structures is approximated as a suitable finite hypothesis space Ξ in different ways. Then a voting procedure elects the structures in Ξ that best explain the data in terms of consensus.

The idea of exploiting consensus is at the core of the celebrated Ransac (Random Sample Consensus) [2], aimed at estimating the parameters of a single model in the presence of large amounts of outliers. Ransac maximizes the size of the consensus set of a structure by searching through a pool of putative structures determined by random sampling. In particular at each iteration a *Minimum Sample Set* (MSS) – composed by the minimum number, say ζ , of points necessary to instantiate the free parameters of a model – is drawn. In this way the estimation problem in the continuous domain Θ is converted into a selection problem in a finite discretized subset $\Xi \subset \Theta$. For each estimated model the corresponding consensus set is computed counting the residuals below the inlier threshold. This procedure is repeated until a structure having enough supporting inliers is discovered among the data.

A number of efforts have been made to improve the RANSAC paradigm. For example, MSAC (M-estimator Sample Consensus) and MLESAC (Maximum Likelihood Estimation Sample Consensus) [3] propose to increase the robustness of the Ransac paradigm incorporating the use of M-estimator techniques. A lot of other refinements in terms of both accuracy and efficiency have been made [4], for example different sampling strategies have been proposed in the literature to reduce the number of iterations necessary to recovery an inlier structure. A nice survey on all these advancements can be found in [5] or in the more comprehensive overview of recent researches presented in [6] where USAC (Universal Framework for Random Sample Consensus) is derived.

A straightforward generalization to multiple models is Sequential Ransac, an iterative, greedy algorithm that executes Ransac many times and removes the found inliers from the data as each structure is detected. Zuliani et al. [7] pointed out some drawbacks of this greedy *estimate-and-remove* approach, which in fact may happen to be sub-optimal since the quality of the attained solution can be affected by inaccurate estimation of the initial structures. In order to correct this behaviour [7] introduced Multi-Ransac, that replaces the sequential scheme with a parallel approach. Rather than looking for a single structure having the largest consensus, κ models having maximal support are searched simultaneously at each iteration. This is done by updating iteratively a collection of κ models with κ new sampled structures using a fusion procedure that enforces explicitly the disjointness of the obtained consensus sets. However, as demonstrated experimentally in [8], this method falls short in the presence of intersecting structures.

The Hough transform and its randomized version [9] can be regarded as well as consensus-oriented algorithms. In these approaches the parameter space Θ is approximated as a quotient space $\Xi = \Theta/\sim$ in which models are represented as equivalence classes of similar structures. The space Ξ is hence employed to build an accumulator collecting data votes: every point adds a vote to the bins representing the structures it belongs to. After voting is complete, the accumulator is analysed to locate the maxima that individuate the desired structures. Differently from Ransac, where Ξ is a discrete sampled version of Θ , in Hough transform the elements of the hypothesis space provide an exhaustive representation of the parameter space, and tentative models are all considered simultaneously. This however comes at the cost of defining a proper quantization of the space, which rapidly becomes intractable as the degrees of freedom of the models increase. The Randomized Hough Transform instead of considering the votes of all the points, exploits random sampling to approximate the accumulator for votes, thereby reducing the computational load.

A more general approach consists in finding modes directly in Θ (*e.g.* [10]). In this way the difficulties of the quantization step are alleviated by mapping the data into the parameter space through random sampling and then by seeking the modes of the distribution (*e.g.* with mean-shift [11]).

In all these consensus based methods, the approximation of Θ is a recurring and tricky issue. The crucial point is that, when multiple structures are hidden in the data, consensus oriented algorithms have to disambiguate between genuine structures and redundant

ones, *i.e.* instances of the same model with slightly different parameter. This issue is addressed by enforcing several disjointness criteria, either explicitly or implicitly by different approximations of the solution space.

For instance, Hough transform handles redundancy by capturing similar structures in the same equivalence class via the quantization of Θ . Along the same line, the bandwidth used in mean shift can be thought as a (softer) way to localize and aggregate redundant models. Also Sequential Ransac and Multi-Ransac enforce disjointness by avoiding to sample similar models. In Sequential Ransac, this is implemented by the iterative removal of the discovered inliers and the subsequent sampling of the hypotheses on the remaining data. In Multi-Ransac this is more evident, since this algorithm explicitly search for the best collection of k disjoint models.

2.2 Preference analysis

In order to overcome some of the drawbacks inherent to consensus methods, the problem has been tackled from a different point of view, where the role of data and models are switched: rather than representing models and inspecting which points support them, the preference sets of individual data points are examined.

This idea can be traced back to [12] where the residuals distribution of points, with respect to a set of putative structures randomly sampled, is analysed. In particular, the most significant structures are revealed as peaks in the histograms of the residuals. In addition, the number of models is automatically determined by the median number of modes found over all data points. In practice, the mode-finding step of this strategy has low accuracy and depends critically on the bin size adopted. However this method has the merit to reformulate the model-estimation task in a space where points are described by their residuals.

Building on this idea, the J-Linkage algorithm [8, 13] was the first successful application of a preference-based representation of data. A two steps *first-represent-then-cluster* scheme is implemented: data are represented by the votes they grant to a set of model hypotheses, then a greedy agglomerative clustering is performed to obtain a partition of the data.

Several elements in common with previous methods can be recognized: an inlier threshold ϵ is used as in Ransac and the idea of casting votes echoes Randomize Hough Transform. Despite that, J-Linkage does not rely on a quantized space, which causes the shortcoming of Hough Transform, nor on the residual space, which

leads to the difficulties of modes estimation, but explicitly introduces a *conceptual* space [14] where points are portrayed by the *preferences* they have accorded to random provisional models. The changes of perspective entailed by preference analysis results in a different approach to the chicken-&-egg dilemma. Structures are recognized as cluster of points in this *preference space* therefore the emphasis is shifted from the estimation to the segmentation part of the problem. T-Linkage [15] extends this idea by using soft votes to depict points preference more accurately.

Along the same line of J-Linkage, Kernel Fitting (KF) [16], Robust Preference Analysis [17] (RPA) and Random Cluster Model Simulated Annealing (RCMSA) [18] exploits points preferences.

KF and RPA first derive a kernel matrix to measure agreement between preferences, then a (different) transformation is applied in order to detect and remove outliers. Then the cleaned kernel matrix is used by KF to over-segment the remaining inliers and reassemble the structures with a merging scheme. RPA performs symmetric non negative factorization of the cleaned kernel matrix, in order to extract the most representative models. Robust statistic is then employed to assign the points to the recovered structures.

RCMSA [18] organizes point preferences 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.

Finally, we can ascribe to preference analysis also all the approaches based on higher order clustering [19, 20, 21, 22], where higher order similarity tensors are defined between n -tuple of points as the probability of points to be clustered together measured in terms of residual errors with respect to provisional models. In this way preferences give rise to a hypergraph whose hyperedges encode the existence of a structure able to explain the incident vertices. The problem of multi-model fitting is hence reduced to finding highly connected component in this preference hypergraph. In practice, this similarity tensor is usually flattened into a pairwise similarity matrix and fed to a clustering method.

It goes without saying that the state-of-the art on multi-model fitting can be also described along other dimensions. The interested reader is referred to [23] for a review of multi-model fitting from the perspective of the optimization of a global energy functional.

2.3 A unified view

If one abstracts from the subtle differences between the implementations of the various multi-model fitting

techniques based on either consensus or preferences, a unified view can be readily achieved by looking at the consensus/preference matrix P defined as

$$P(i, j) = \begin{cases} 1 & \text{if } \text{err}(x_i, \theta_j) < \epsilon \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $x_i \in X$ are data points, $\theta_j \in \Xi$ are tentative structures, $\text{err}: X \times \Theta \rightarrow \mathbb{R}^+$ is a suitable error function that associates to every point-model pair $(x, \theta) \in X \times \Theta$ the corresponding deviation or residual $\text{err}(x, \theta)$, and ϵ is the inlier threshold.

The binary matrix P can be interpreted in several ways. It can be regarded as the incidence matrix of an hyper-graph where rows correspond to vertices and columns represent hyper-edges. Its rows, identified with preference sets, can be interpreted as representations of data in high dimensional spaces. In both these cases multi-model fitting boils down to cluster analysis.

Changing the perspective, if columns are taken into account they can be regarded as a collection of consensus sets. Ransac and Sequential Ransac simply aim at finding the column or the k columns respectively having greatest sums. Multi-Ransac seeks for the k “most orthogonal” columns having maximal consensus.

The binary preference matrix can be also viewed as the biadjacency matrix of a bipartite graph, where one set of vertices represents points and the other one represents structures, an edge links two vertices if the corresponding points belong to the consensus set of the related structures. Maximal bicliques in this graph correspond to biclusters of points and structures. The problem of finding maximal edge biclique can be relaxed to a continuous formulation that in turn is solved by NMF [24]. However, strictly speaking in multi-model fitting we are not interested in maximizing the number of edges of a biclique, as explained in Figure 1. Nevertheless this interpretation traces an interesting connection with RPA which relies also on symNMF, a particular instance of NMF.

2.4 Preference vs Consensus

All the preference-oriented techniques lean towards the segmentation side of the multi-model fitting problem. Indubitably, they have the great advantage of casting specific multi-model fitting problems in the general clustering framework. Nevertheless it has been largely recognized by the research community that the segmentation/clustering problem is ill-posed, and that a “*no free lunch theorem*” [25] holds, which states that a given clustering method can be optimal only with respect to some specific type of datasets.




Fig. 1: Biclique interpretation. The binary preference matrix can be interpreted as a bi-adjacency matrix of a bipartite graph. The maximal edge biclique induces a segmentation of the data in which a unique cluster is determined in correspondence of the intersection of the two lines. Two points remain unassigned, whereas maximizing consensus leaves a single point unexplained.

Moreover, Kleinberg [26] confirms that clustering techniques are inherently fraught with ambiguities: he conceives an axiomatic theory in which he defines three desirable properties that a clustering scheme ought to satisfy, namely scale-invariance, a richness condition that all partitions are achievable, and a consistency requirement on the shrinking and stretching of distances. In this setting an “*impossibility theorem*” is derived, demonstrating that there is no clustering function satisfying simultaneously all the three properties. For example, linkage clustering with a distance-based stopping condition enjoys the nice theoretical properties of scale-invariance and richness, but not consistency.

In addition, two other main issues are not satisfactorily handled by clustering techniques. In first instance, classical clustering approaches are designed to yield a *partition* of the data, hence they are not suitable for explicitly dealing with intersecting structures. As a result, they are either ignored or handled indirectly with some *ad hoc* post processing on the obtained segmentations.

In second place, the treatment reserved to outliers is not completely sound. For estimation purposes, gross outliers ought to fall in a special group of points, but clustering treats all the segments in the same way. This is the reason why partitional clustering schemes are not able to enforce robustness by simply throwing in one additional model with the hope that outliers will be clustered together. Hierarchical methods in practice are more resilient to rogue points, but in principle outliers do not have a specific treatment during the clustering phase: for example T-Linkage relies on a *posteriori* specific heuristics to ensure robustness.

Consensus-oriented methods focus more on the estimation part of the problem. Classical algorithms such as Ransac, Sequential Ransac and Multi-Ransac in prac-

tice do not operate on the whole matrix P as is the case with preference analysis. The subtle and decisive difference is that these consensus based techniques generate the columns of P iteratively to save computational efforts. This principle of parsimony has no drawback in case of single model estimation, but in multiple model scenarios it causes the shortcomings of Sequential Ransac. Once a structure of inlier is detected, its supporting points are removed and successive hypothesis are sampled exploiting only the remaining of the data, as a consequence inaccurate detection at early stage of the algorithm can heavily deteriorate the results. In addition, points in the intersections do not contribute to the sampling of subsequent structures and this greedy strategy is inherently prone to achieve a suboptimal segmentation (see Figure 2).




Fig. 2: Shortcomings of greediness on a fitting line example. Data points (on the left): three lines supporting each 2, 4 and 8 points are intersected by two lines collecting half the total number of points each (7 points). Clearly these two lines fit all the data points, therefore the optimal value for the MC problem is $\text{opt} = 2$, the greedy algorithm will pick the 3 remaining lines.

This *estimate-and-remove* approach is tantamount of enforcing disjointness on the attained structures, which is essentially meant to distinguish between genuine structures and redundant ones.

In practice, however, using consensus as the only criterion seems short-sighted, for true models can have mutual intersections greater than redundant ones; in this case the algorithm would fail in discerning authentic structures, as shown in Figure 3.

2.5 Our approach

In this work we sidestep the pitfalls of clustering and focus on the goal of maximizing consensus, promoting a well motivated global objective function to the detriment of less controllable and more fragile segmentation methods. In doing this, we shall, at the same time, counteract the greediness and poor handling of intersecting models typical of consensus-based approaches. For these reasons, we substitute partitions with the notion of *covers*, so to restate the multi-model fitting task




Fig. 3: Redundant structures (left) may happen to have fewer points in common than genuine intersecting ones (right), for this reason consensus based method struggle in the presence of multiple-structures.

in terms of covering problems, namely SET COVER and MAXIMUM COVERAGE. The main advantages of this approach are that

- points lying in model intersections are allowed to belong to multiple structures,
- if the number of desired model is specified, robustness to outlier is easily achieved in a neat way,
- the typical shortcoming of sequential approaches can be avoided leveraging on optimal solvers
- the discrete fra of covering problems eludes, to some extent, the tricky trade-off between data-fidelity and model complexity of some optimization-based multi-model fitting method.

3 Set Cover Formulation

In what follows we will assume a discrete setting in which point votes are expressed by binary values, as in Ransac or J-linkage. A consensus matrix P has to be instantiated via random sampling, but we are agnostic here on the specific sampling strategy (*e.g.* [27, 28, 17]).

At first we concentrate on the case in which all the points are inliers (the case of outliers will be dealt with later on). A natural requirement is to ask that all the points are explained by some structures, in other words, the true structures determine, by means of their consensus sets, a *cover* of the data, *i.e.* a collection of sets whose union contains X :

$$F = \{S_j : j \in J\} \quad \text{such that} \quad X \subseteq \bigcup_{j \in J} S_j, \quad (2)$$

Note that we are not requiring that these sets are disjoint, so we are not limited to partitions and we can properly handle the case of intersecting models. By invoking the Occam's principle, a straightforward formulation is therefore to ask for a cover consisting of a minimal number of consensus sets. In this way we are

implicitly discouraging redundant models. Thus we are naturally led to a SET COVER problem.

Definition 1 (SET COVER) Given a ground set X and $F = \{S_1, \dots, S_m\}$ a cover of X , select the minimum number of subsets from F that covers X .

In this formulation, $X = \{x_1, \dots, x_n\}$ contains the data points and the collection $F = \{S_1, \dots, S_m\}$ is composed by the consensus sets of the sampled models $\theta_1, \dots, \theta_m \in \Xi$: *i.e.* $S_j = \{x \in X : \text{err}(x, \theta_j) < \epsilon\}$ instantiated on minimal sample sets as in Ransac. The property that F is a cover of X can be easily enforced by requiring that every points of X is sampled at least once. SET COVER can be rephrased rigorously using the matrix P in the constraints formulation and introducing m binary variables $z_j \in \{0, 1\}$ for each subset S_j . If S_j is selected in the solution then $z_j = 1$, otherwise $z_j = 0$. In this way SET COVER can be shown to be equivalent to an Integer Linear Programming (ILP) problem:

$$\begin{aligned} & \text{minimize} && \sum_{j=1}^m z_j \\ & \text{subject to} && Pz \geq \mathbf{1}. \end{aligned} \quad (3)$$

The constraint can be expanded as

$$\sum_{j: S_j \ni x_i} z_j \geq 1 \quad \forall x_i \in X \quad (4)$$

to emphasize that it is meant to ensure that the solution $\{S_j\}_{j: z_j=1}$ is a cover of X .

Empirically, we found beneficial to preprocess the input as follows. First of all, keeping in mind that our aim is to maximize consensus, we refit a structure to each consensus set via least squares, and, if the consensus has increased, we update the structure and its supporting points. The remaining sets are hence ordered by cardinality $|S_1| \geq |S_2| \geq \dots \geq |S_k|$ and a set S_j is discarded if

$$S_j \subseteq \bigcup_{i=1}^{j-1} S_i. \quad (5)$$

The rationale of this choice is to keep only those structures that cover at least one point that otherwise would be left uncovered by the union of larger ones. Please note that in particular we are also removing subsets that are contained in one larger set.

SET COVER is long known to be NP-hard [29]: not surprisingly, since the inherent complexity of multi-model fitting does not disappear by simply rephrasing it in different terms. Nevertheless, this optimization problem is among the oldest, most studied and widespread ones

in the mathematical programming literature. Therefore we can reap the outcomes of the efforts made by the scientific community in addressing this issues, and enjoy the fruits of several studies focused on approximating the solutions of this problem.

For example, the greedy strategy – hereinafter Greedy-RansaCov– which keeps choosing the set that covers most uncovered points until they all are covered, embodies the spirit of Sequential Ransac with the only differences that the hypothesis space is not sampled iteratively and, instead of returning a partition, intersecting segments are allowed. It has been demonstrated by Feige [30] that this greedy strategy is the best possible in terms of approximation ratio. More precisely an approximation of $h(n)$ holds in the case of SET COVER problem, where $h(n)$ denotes the n -th harmonic number. This result applies effortlessly to Greedy-RansaCov giving a provable quality measure of the solution.

Another straightforward strategy consist in solving a relaxed Linear Programming (LP) problem and converting the solution by rounding up all non-zero variables to 1. In this case [31] shows that the solution achieves an approximation guarantee of a factor equal to the maximal point frequency, where the frequency of a point is the number of sets that cover that point. Interestingly enough, our preprocessing not only reduces the number of sets to be considered, but also improves the approximation factor of the relaxed LP solution, for it actually reduces the maximal frequency of the points.

In practice, more sophisticated strategies are used by ILP solvers, but the relaxed LP solution is a good starting point. In our algorithm – henceforth dubbed ILP-RansaCov– we use the `intlinprog` Matlab function, which attempts to tighten the LP relaxation with several heuristics and falls back to branch and bound in case of failure.

4 Maximum Coverage formulation

If X is corrupted by rogue points we can integrate them in the formulation of the problem at the cost of introducing an additional parameter k equal to the desired number of structures. Requiring some extra information to deal with outliers seems to be unavoidable. In this respect, k is a more guessable parameter than others.

Instead of trying to find the smallest number of sets that cover all elements, we search for the largest number of points that can be covered by k sets, possibly leaving some points (the outliers) uncovered. This leads to the so called MAXIMUM COVERAGE problem.

Definition 2 (MAXIMUM COVERAGE) Given a ground set X , $F = \{S_1, \dots, S_m\}$ a collection of subsets of X

and an integer k , select from F at most k subsets that cover the maximum number of points in X .

This problem is translated in an ILP one thanks to a collection of n auxiliary variables y_i , such that $y_i = 1$ if x_i belongs to the returned subsets, 0 otherwise:

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^n y_i \\ & \text{subject to} && \sum_{j=1}^m z_j \leq k \\ & && \sum_{j: S_j \ni x_i} z_j \geq y_i \quad \forall x_i \in X \\ & && 0 \leq y_i \leq 1, \quad z_j \in \{0, 1\}. \end{aligned} \quad (6)$$

The first condition enforces that no more than k sets are picked and the second constraint ensures that if $y_i \geq 0$ then at least one set $S_j \ni x_i$ is selected.

The MAXIMUM COVERAGE formulation can be straightforwardly generalized to a weighted version, where the objective function is not the number of the points covered but the sum of their weights. This allows to take advantage of any kind of available prior information to promote or discourage the coverage of certain points by weighting them.

Definition 3 (WEIGHTED MAXIMUM COVERAGE) Given a ground set X , $F = \{S_1, \dots, S_m\}$ a collection of subsets of X , weights $c_i \geq 0$ associated to the elements of X and an integer k , select from F at most k subsets that maximize the sum of the weight of covered points in X .

The ILP formulation of the problem changes accordingly, replacing the objective function of MAXIMUM COVERAGE expressed in Equation (6) with the sum of weights of the covered points:

$$\max \sum_{i=1}^n c_i y_i. \quad (7)$$

Finally, please note that if k is not available, a practical solution is to solve SET COVER and detect outliers a-posteriori as points belonging to structures with the smallest consensus sets, as in J-linkage [15].

5 Relationship to FACILITY LOCATION.

The closest methods to our in the literature are those casting multi-model fitting as a FACILITY LOCATION problem: provided a set of potential facilities (which corresponds to the pool of tentative structures), FACILITY

LOCATION selects an optimal subset of facilities and assigns customers (*i.e.* data points) to one facility each, so as to minimize the sum of facility opening costs and the distances between customers to their assigned facilities. This leads to the optimization of a cost function composed by two terms, mimicking the classical MAP-MRF objectives: a modelling error (customers-facility distances) which can be interpreted as a likelihood term, and a penalty term (the cost to open the facilities) that encodes model complexity. Some authors solve it with ILP [32,33,34,35] while others propose different combinatorial optimization techniques [23,36,37,38]. Although SET COVER and FACILITY LOCATION are related (the first can be rephrased as a special case of the second) and ILP has been used to solve both, ILP-RansaCov differs from previous work based on FACILITY LOCATION in many respects.

In first instance, FACILITY LOCATION methods need to strike the correct trade-off between data fidelity and model complexity, in order to achieve the proper balance between over and under fitting. For example PeARL [37,23] regularizes the modeling-fitting error, expressed in terms of residual, by introducing a label costs that penalizes the number of different structures, whereas FLOSS [35], aimed at fitting subspace of different dimensions to outlier-free data, exploits a penalty term encoding subspace dimension. In contrast, ILP-RansaCov elude this thorny trade-off: in the outlier-free scenario SET COVER regularizes the solution invoking the minimality of cover, while, if outliers are present, MAXIMUM COVERAGE plainly requires the maximum number of models as a clear, intelligible parameter, instead of balancing two incommensurable quantity in the cost function.

Second, FACILITY LOCATION methods minimize the fitting error on the continuum of residuals, in the same spirit of MLE estimators, while ILP-RansaCov gains resiliency to outliers by maximizing the consensus *à la* Ransac. This comes at the price of assuming that all the structures have the same error scale, while MLE-like estimators can compute the scale along the parameters of each model.

In our formulation the rogue points will be simply left uncovered by MAXIMUM COVERAGE, whereas FACILITY LOCATION methods cope with outliers either by introducing a special additional model for which a constant fidelity measure has to be manually tuned [23], or an upper bound to the total number of outliers has to be defined [33]

Finally, FACILITY LOCATION approaches enforce hard-membership constraints, producing a partition of the data, whereas ILP-RansaCov inherently caters for intersecting solutions.

6 Experiments on simulated data

In this section we investigate the performance of ILP-RansaCov with respect to Greedy-RansaCov, which emulates Sequential Ransac, J-Linkage [8] and T-Linkage [15] on synthetic data, using *the same sampling* and the same inlier threshold for all the methods (or, equivalently, the same P matrix). We obtained the implementations of J-Linkage and T-Linkage from [39]. The MATLAB code of ILP-RansaCov is available on the web¹.

Performances are measured by the misclassification errors (ME) – *i.e.* the percentage of misclassified points – which is computed as follows. First, of all the map between ground-truth labels and estimated ones that minimize standard ME is found (as in [32]), then a point is deemed as correct if at least one of its label is in correspondence with the correct ground-truth one.

The data sets (Figure 4) consist of segments in several configurations and circles, as in [8]. Each structure consists of 50 inliers points, contaminated by Gaussian noise and outlying points in different percentages (reported in Table 1). All the methods have been provided with the correct number of structures k ; in the case of J-Linkage and T-Linkage, among the structures produced by the algorithms the k ones with more supporting points are kept as inliers.

	outliers	J-Lnkg	T-Lnkg	Grdy-RansaCov	ILP-RansaCov
Stair4	50%	10.20	10.00	39.20	12.00
Star5	60%	15.20	14.40	10.40	3.80
Star11	50%	35.00	33.09	32.36	25.18
Circle4	50%	26.50	23.00	30.25	11.25
mean		20.12	20.12	28.05	13.06

Table 1: Misclassification error (ME %) on **simulated data**.

The results are collected in Figure 4 while Table 1 reports the misclassification errors (ME) – *i.e.* the percentage of misclassified points.

First of all we can notice that in the **Stair4** experiment (originally used in [7] to criticize Sequential Ransac), Greedy-RansaCov performs poorly. The shortcomings of this greedy strategy are here afoot: the incorrect selection of the first structure compromises the subsequent interpretation of the data. A greedy approach to the MAXIMUM COVERAGE problem yields a suboptimal segmentation also on the **Circle4** dataset, where one of the four structures is over-segmented by Greedy-RansaCov at the expense of the smaller circle in the centre.

On **Star11** J-linkage misses a ground truth segment. During the merging process some inliers are incorrectly

aggregated to spurious models, hence the recovered segment that actually corresponds to a ground truth structure collects fewer inliers, to the point that it falls outside the first k largest models and is deemed as outlier. In general the tendency of loosing inliers during the segmentation step affects J-Linkage (and T-Linkage) also in the other datasets, *e.g.* it is particularly evident on **Circle4**,

Even when the discovered inliers are enough to recover the corresponding structures, this behaviour has a detrimental effect on the model estimate, for it increases the variance.

ILP-RansaCov yields reliable segmentations in all the experiments, and it achieves the best average ME. The reason can be ascribed to the non-greedy minimization strategy (w.r.t. Greedy-RansaCov) and to the departure from the partitioning paradigm (w.r.t. J-Linkage and T-Linkage). As a matter of fact, when models do not intersect, as in **Stair4**, the performance of J-Linkage and T-Linkage are in the same range of ILP-RansaCov.

6.1 Computational complexity.

We run an experiment to probe of how the running time scales with the input dimension and where the time is spent. To this end, we run ILP-RansaCov on a line fitting problem extracted from *Star11* with variable number of sampled models and number of points. The execution times, broken down for each step, are reported in Figure 7.

The instantiation of the preference matrix (sampling) dominates the complexity for moderate point number, whereas ILP takes over when the number of points increases. Also, while the dependence from the number of sampled models appears to be polynomial, the execution time grows exponentially with the number of points, in accordance with theoretical prediction.

The impact of the preprocessing step, related to Equation (5), is negligible in terms of the running time, but it improves the quality of the solution: *e.g.*, with reference to Table 4 (F), the mean ME of ILP-RansaCov *without this refinement* raises to 11.44.

7 Experimental comparison on real data

In this section, we demonstrate the performance of ILP-RansaCov on three classical Computer Vision applications, namely: i) vanishing point detection, ii) video motion segmentation, and iii) two-views segmentation. In all these scenarios we compare ILP-RansaCov with J-Linkage [8], T-linkage [15] and RPA [17], whose implementation is taken from [40].

¹ www.diegm.uniud.it/fusiello/demo/cov/




Fig. 4: Comparison on **simulated data** (outliers marked as x).



Fig. 5: A sample of the worst ILP-RansaCov results on YorkUrbanDB (**vanishing point detection**). Line membership is color-coded.

In addition, one reference method have been added to the comparison for each specific scenario, namely: MFIGP [38] in the vanishing point experiments, SSC [41] for video motion segmentation and RCMSA [18] for two-views segmentation. These methods have been selected because i) they are among the best performers, ii) the original code from the authors is available (MFIGP [42], RCMSA [42], SSC [43]), and iii) they have been tested on the same respective datasets, and this allows also indirect comparison with other methods considered in the respective papers [38, 41, 18].

MFIGP and RCMSA are considered only in one scenario out of three because the authors provided the tuning parameters only for that particular application (vanishing point detection and two-views segmentation, respectively). SSC instead is tailored specifically to subspace segmentation, hence it cannot be applied in the other two cases (where models are not linear or affine subspaces).

All the algorithms but SSC and RCMSA were provided with the same pool of putative models, sampled as in [17]. Since none of the competing algorithm had special provisions for the outlier-free case, the **MAXIMUM**




Fig. 6: Results on YorkUrbanDB. (a) is the cumulative distributions of the errors per sequence; (b) shows the area above the curve (the smaller the better).




Fig. 7: Execution time of ILP-RansaCov on simulated data w.r.t. the dimensions of the problem.

COVERAGE version of ILP-RansaCov has been used in this experiment, for the sake of fairness.

7.1 Vanishing point detection.

In this experiment we compare the performances of ILP-RansaCov with MFIGP on vanishing point detection using the York Urban Line Segment Database [44], or York Urban DB in short, a collection of 102 images of architectural Manhattan-like environments (*i.e.* scenes dominated by two or three mutually orthogonal vanishing directions). Annotated line-segments that match with the 3-d orthogonal frame of the urban scene are provided with the ground-truth, no outliers are present in the data. The aim is to group the supplied segments in order to recover two or three orthogonal vanishing points.

MFIGP (Model Fitting with Interacting Geometric Priors) is a recently proposed method [38] that builds on [37] adding high-level geometric priors. In particular, in this application, an additional term expressing interaction between vanishing points is included into the FACILITY LOCATION formulation, to promote the extraction of orthogonal vanishing points. The global input parameters recommended in the original paper have been optimized individually for each single image to enhance the results.

Figure 5 shows three images where ILP-RansaCov achieved the worst ME, which are nevertheless qual-

	J-Lnkg	T-Lnkg	RPA	MFIGP	Grdy-RansaCov	ILP-RansaCov
Mean	2.85	1.44	1.08	3.51	2.38	0.19
Med	1.80	0.00	0.00	0.16	0.00	0.00

Table 2: Misclassification error (ME %) on YorkUrbanDB.

itatively correct. Figure 6(a) reports the cumulative distribution of the ME per sequence, *i.e.* the value on the ordinate corresponds to the percentage number of sequences where the algorithm achieved a ME lower than the abscissa. The differences among the methods can be better appreciated by plotting the area above the cumulative distribution of ME (Figure 6(b)) or by analysing the average and median ME, collated in Table 2. These quantitative results confirms that ILP-RansaCov achieves the most accurate performance, followed by RPA. As MFIGP enhances PeARL, figures in Table 2 indirectly corroborate the advantage of ILP-RansaCov over PeARL. Please note also that Greedy-RansaCov, a proxy of the vilified Sequential Ransac, performs better than other sophisticated methods, in this task.

7.2 Video motion segmentation.

In this experiments we considered Sparse Subspace Clustering [41] a state-of-the-art technique that exploits a sparse representation to build an affinity matrix, which in turns is segmented by spectral clustering.

We use the 51 *real* video sequences from the Hopkins 155 dataset [45], each containing two or three moving objects, with no outliers. The input data is a set of features trajectories across a video taken by a moving camera, and the aim is to recover the different rigid-bodies. Motion segmentation can be seen as a subspace segmentation problem under the assumption of affine cameras. In this case all feature trajectories associated with a single moving object lie in a linear subspace of dimension at most 4 in \mathbb{R}^{2f} (where f is the number of video frames). For this reason feature trajectories of a dynamic scene containing n rigid motion lie in the union of n low dimensional subspace of \mathbb{R}^{2f} and segmentation can be reduced to clustering data points in a union of subspaces. Following [46], in order to deal with degenerate motions, we project the data onto an affine 4-dimensional space where the rigid-body segmentation is translated in a 3-d plane fitting problem.

Figure 8 reports some sample results, in particular three sequences belonging to `Traffic2` and `Others3` subsets, respectively, where ILP-RansaCov achieves sub-optimal segmentations. Figure 9 and Table 3 provide a comparison of the performances in terms of ME: ILP-




Fig. 8: A sample of the worst ILP-RansaCov results on Hopkins155 (video motion segmentation). Point membership is color-coded.




Fig. 9: Results on Hopkins155. (a) is the cumulative distributions of the errors per sequence; (b) shows the area above the curve (the smaller the better).

		J-Lnkg	T-Lnkg	RPA	SSC	Grdy-RansaCov	ILP-RansaCov
Traffic3	Mean	1.58	0.48	0.19	0.76	28.65	0.35
	Med	0.34	0.19	0.00	0.00	1.53	0.19
Traffic2	Mean	1.75	1.31	0.14	0.06	7.48	0.54
	Med	0.00	0.00	0.00	0.00	0.00	0.00
Others3	Mean	6.91	5.32	9.11	2.13	14.89	2.13
	Med	6.91	5.32	9.11	2.13	14.89	2.13
Others2	Mean	5.32	6.47	4.41	3.95	8.57	2.40
	Med	1.30	2.38	2.44	0.00	0.20	1.30
All	Mean	2.70	2.47	1.42	1.08	10.91	0.98
	Med	0.00	0.00	0.00	0.00	0.00	0.00

Table 3: Misclassification error (ME %) on Hopkins155.

RansaCov places in the same range of SSC and achieves the best overall results. In this case the advantage of solving the MAXIMUM COVERAGE problem with a global approach is afoot, since the greedy strategy of Greedy-RansaCov, sampling being equal, fails in recovering accurate segmentations.

Please note that FLOSS [35] has been evaluated on the Hopkins 155 dataset, therefore an indirect comparison with our method can be derived. Even if FLOSS handles a variety of models that can describe the dynamic scenes more accurately, ILP-RansaCov enjoys comparable performance and still achieves the best overall results.

7.3 Two-views segmentation.

In this experiment we additionally compare ILP-RansaCov against RCMSA [18] on the Adelaide Robust Model Fitting Data Set, or AdelaideRMF in short, which consists of 38 image pairs, 19 related by multiple homographies (H) and 19 by multiple fundamental matrices (F), with outliers. The task involves segmenting different planes/moving objects by fitting homographies/fundamental matrices to subsets of corresponding points.

All the methods are given the inlier threshold computed from the available ground truth.

Some failure examples are reported in Figure 10. The left image is an example of under-segmentation, where a unique fundamental matrix explains both the cube and the toy (red points). In the middle image ILP-RansaCov fails in detecting one planar structure (second wall of the building from the left). In the right image the campanile (on the very right) is over-segmented, and this “consumes” one of the available k models, thereby preventing the nearby wall to be detected.

From the data reported in Figure 11 and Table 4, the reader can appreciate that the ME achieved by ILP-RansaCov is consistently lower than RCMSA and in the same range of RPA.

In order to evaluate the relative importance of multiple membership with respect to the optimization method we have rephrased Multi-Ransac in the framework of maximal coverage: the strategy is similar to Greedy-RansaCov, the difference is that, after a set is picked, the subsequent ones are searched among those having maximal Jaccard distance with the currently covered elements, thereby maximizing disjointedness. Even if a point can be assigned to multiple model, experiments demonstrated that the performances are inferior to ILP solutions (ME is: 2.97 for VP, 4.58 for video sequences, 26.85 for homographies and 17.01 for fundamental matrices), confirming the crucial role of adopting a global optimization technique.

To complete the picture Table 5 reports the time spent by ILP-RansaCov in each experiment on real data.




Fig. 10: A sample of the worst ILP-RansaCov results on AdelaideRMF (**two-views segmentation**). Point membership is color-coded, black crosses are points classified as outliers.




Fig. 11: Results on AdelaideRMF. (a) is the cumulative distributions of the errors per sequence; (b) shows the area above the curve (the smaller the better).

		J-Lnkg	T-Lnkg	RPA	RCMSA	Grdy-RansaCov	ILP-RansaCov
F	Mean	16.43	9.37	5.49	12.37	17.08	6.04
	Med	14.29	7.80	4.57	9.87	21.65	4.27
H	Mean	25.50	24.66	17.20	28.30	26.85	12.91
	Med	24.48	24.53	17.78	29.40	28.77	12.34

Table 4: Misclassification error (ME %) for motion segmentation (F) and plane segmentation (H) on AdelaideRMF.

A comparison with other methods would have been meaningless for they are not fully coded in MATLAB as ours.

While it is understood that ILP is a NP-hard problem, execution time can be improved by fine tuned heuristics (*e.g.* [47]) and careful implementation in compiled language.

	YorkUrbanDB	Hopkins155	Adelaide (F)	Adelaide (H)
mean	8.09	41.19	52.24	146.34
median	1.14	11.56	48.79	51.12

Table 5: Execution time [s] of ILP-RansaCov on real data.

8 Qualitative experiments on plane fitting.

Multiple structure recovery can be fruitfully employed in the contest of automatic architectural and urban modelling from images. In this scenario usually a reconstruction technique produces arbitrarily dense but unstructured points clouds. Fitting multiple geometric primitives to these point clouds is a first step in organizing it in a higher informative semantic level. In this section we asses the performance of ILP-RansaCov on this application, in particular we consider the problem of fitting plane to 3D-point clouds produced by structure from motion. The data are the same used in [38] and they have been retrieved from the Oxford Colleges building reconstruction dataset [48] (Merton College 3) and from [49] (Pozzoveggiani and Piazza Dante). Since ground truth is not available, the assessment can be only qualitative; however the reader will appreciate in Figure 12 that the main constituting planes are correctly recovered.

9 Conclusions

We formulated multi-model fitting in terms of SET COVER and MAXIMUM COVERAGE problems, yielding a simple and easy to implement method that generalizes Ransac to the case of multiple structures in a neat and principled manner.

The multi-model fitting problem is formulated in terms of optimization of a global cost function, thereby eluding the greediness of techniques such as Sequential/-Multi-Ransac and J-linkage, but at the same time avoiding the difficult trade-off between data and complexity terms of other formulations, by resorting to consensus maximization. In both cases, we tackle the problem of intersecting models at the root, by replacing partitions with coverages.

ILP-RansaCov is modular with respect to the ILP solver and to the sampling strategy. There are few pa-

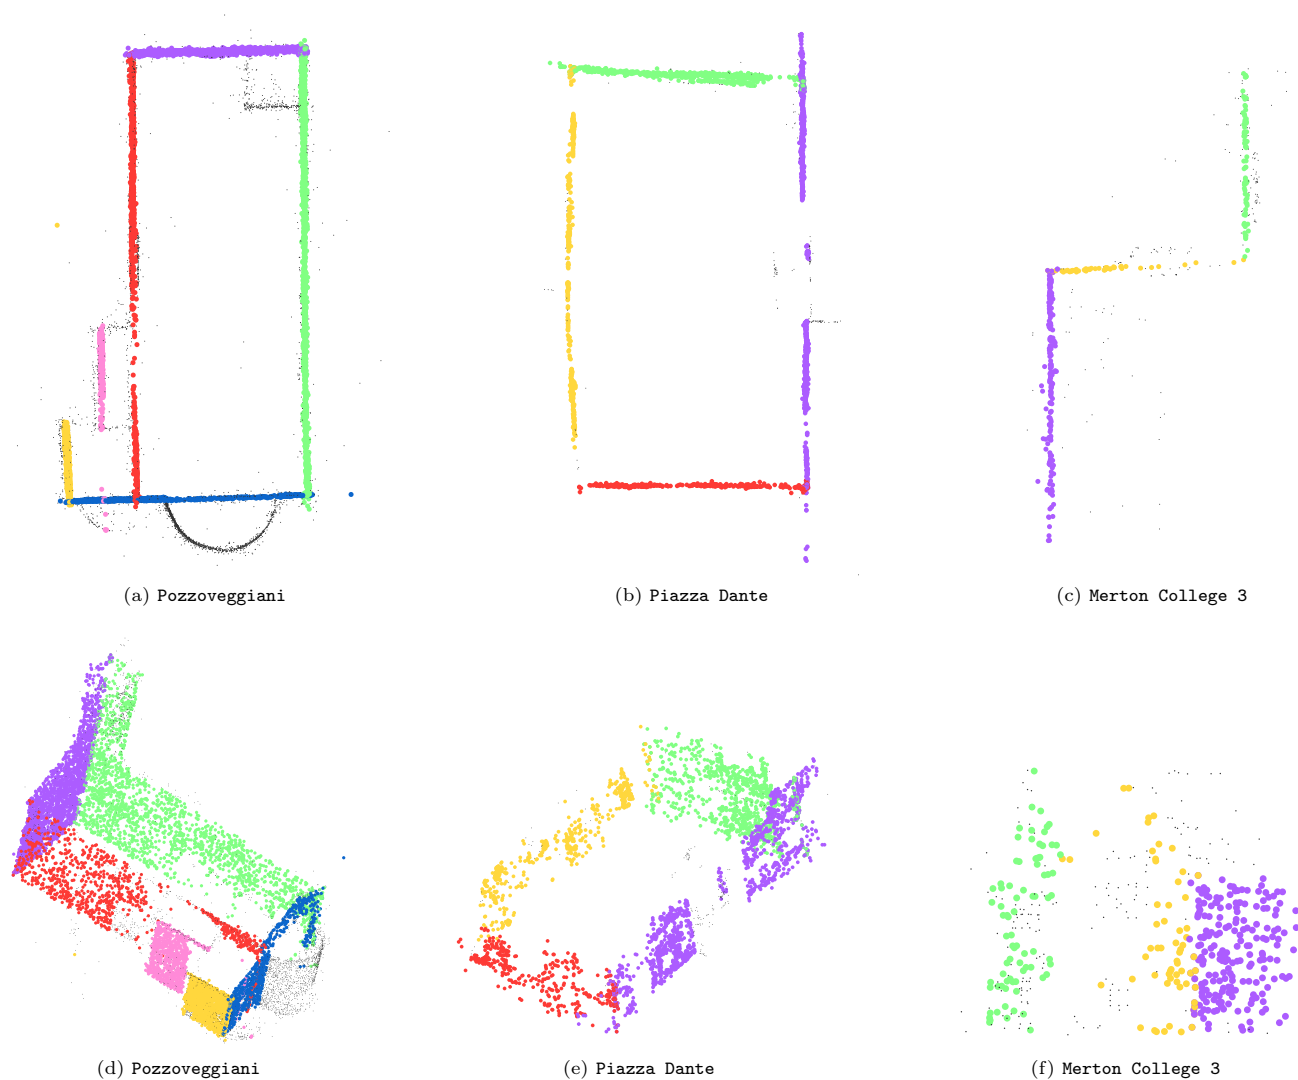


Fig. 12: Qualitative assessment of ILP-RansaCov for fitting planes to 3D point clouds. The first row show a plan view of the fitting results of ILP-RansaCov, whereas the second row shows the same cloud of points from a casual viewpoint. Plane membership is colour-coded.

rameters that require tuning – namely the inlier threshold and the number of desired model – and these have the added virtue of being intuitive in their meaning and function.

In summary, we expect that this paper will offer practitioners a manageable tool for addressing a difficult and ubiquitous problem, and will provide the community a reference baseline for further advancements.

References

1. L. Magri, A. Fusiello, Multiple models fitting as a set coverage problem, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 3318–3326. doi:10.1109/CVPR.2016.361.
2. M. A. Fischler, R. C. Bolles, Random Sample Consensus: a paradigm model fitting with applications to image analysis and automated cartography, *Communications of the ACM* 24 (6) (1981) 381–395.
3. P. H. S. Torr, A. Zisserman, MLESAC: A new robust estimator with application to estimating image geometry, *Computer Vision and Image Understanding* 78 (2000) 2000.
4. K. Lebeda, J. Matas, O. Chum, Fixing the locally optimized RANSAC—full experimental evaluation, in: *British Machine Vision Conference*, 2012, pp. 1–11.
5. J. Choi, G. G. Medioni, StaRSaC: Stable random sample consensus for parameter estimation., in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2009, pp. 675–682.
6. R. Raguram, O. Chum, M. Pollefeys, J. Matas, J. Frahm, USAC: a universal framework for random sample consensus, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8) (2013) 2022–2038.

7. M. Zuliani, C. S. Kenney, B. S. Manjunath, The multi-RANSAC algorithm and its application to detect planar homographies, in: Proceedings of the IEEE International Conference on Image Processing, 2005, pp. III-153-6.
8. R. Toldo, A. Fusiello, Robust multiple structures estimation with J-Linkage, in: Proceedings of the European Conference on Computer Vision, 2008, pp. 537-547.
9. L. Xu, E. Oja, P. Kultanen, A new curve detection method: randomized Hough transform (RHT), *Pattern Recognition Letters* 11 (5) (1990) 331-338.
10. R. Subbarao, P. Meer, Nonlinear mean shift for clustering over analytic manifolds, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2006, pp. 1168-1175.
11. D. Comaniciu, P. Meer, Mean shift: A robust approach toward feature space analysis, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (5) (2002) 603-619.
12. W. Zhang, J. Kosecká, Nonparametric estimation of multiple structures with outliers, in: Workshop on Dynamic Vision, European Conference on Computer Vision 2006, Vol. 4358 of Lecture Notes in Computer Science, Springer, 2006, pp. 60-74.
13. R. Toldo, A. Fusiello, Image-consistent patches from unstructured points with J-Linkage, *Image and Vision Computing* 31 (10) (2013) 756-770.
14. R. Duin, E. Pekalska, P. Paclik, D. Tax, The dissimilarity representation, a basis for domain based pattern recognition?, in: Pattern representation and the future of pattern recognition, ICPR 2004 Workshop Proceedings, 2004, pp. 43-56.
15. L. Magri, A. Fusiello, T-Linkage: A continuous relaxation of J-Linkage for multi-model fitting, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 3954-3961.
16. T. Chin, H. Wang, D. Suter, Robust fitting of multiple structures: The statistical learning approach, in: Proceedings of the International Conference on Computer Vision, 2009, pp. 413-420.
17. L. Magri, A. Fusiello, Robust multiple model fitting with preference analysis and low-rank approximation, in: Proceedings of the British Machine Vision Conference, BMVA Press, 2015, pp. 20.1-20.12.
18. T.-T. Pham, T.-J. Chin, J. Yu, D. Suter, The random cluster model for robust geometric fitting, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36 (8) (2014) 1658-1671.
19. S. Agarwal, J. Lim, L. Zelnik-manor, P. Perona, D. Kriegman, S. Belongie, Beyond pairwise clustering, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2005, pp. 838-845.
20. V. M. Govindu, A Tensor Decomposition for Geometric Grouping and Segmentation, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 1 (2005) 1150-1157.
21. S. Jain, V. M. Govindu, Efficient higher-order clustering on the grassmann manifold, in: Proceedings of the International Conference on Computer Vision, 2013, pp. 3511-3518.
22. R. Zass, A. Shashua, A unifying approach to hard and probabilistic clustering, in: Proceedings of the International Conference on Computer Vision, Vol. 1, 2005, pp. 294-301.
23. H. Isack, Y. Boykov, Energy-based geometric multi-model fitting, *International Journal of Computer Vision* 97 (2) (2012) 123-147.
24. N. Gillis, F. Glineur, Nonnegative factorization and the maximum edge biclique problem, arXiv preprint arXiv:0810.4225.
25. D. H. Wolpert, W. G. Macready, No free lunch theorems for optimization, *Transaction on evolutionary computation* 1 (1) (1997) 67-82.
26. J. Kleinberg, An impossibility theorem for clustering, *Advances in neural information processing systems* (2003) 463-470.
27. T.-J. Chin, J. Yu, D. Suter, Accelerated hypothesis generation for multistructure data via preference analysis, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2012) 533-546.
28. Y. Kanazawa, H. Kawakami, Detection of planar regions with uncalibrated stereo using distributions of feature points, in: *British Machine Vision Conference*, 2004, pp. 247 - 256.
29. R. M. Karp, *Reducibility among combinatorial problems*, Springer, 1972.
30. U. Feige, A threshold of $\ln(n)$ for approximating set cover, *Journal of the ACM* 45 (4) (1998) 634-652.
31. V. V. Vazirani, *Approximation algorithms*, Springer Science & Business Media, 2013.
32. P. H. Torr, D. W. Murray, Stochastic motion clustering, in: *Proceedings of the European Conference on Computer Vision*, Springer, 1994, pp. 328-337.
33. H. Li, Two-view motion segmentation from linear programming relaxation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1-8.
34. K. Schindler, D. Suter, Two-view multibody structure-and-motion with outliers through model selection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28 (6) (2006) 983-995.
35. N. Lazic, I. E. Givoni, B. J. Frey, P. Aarabi, FLoSS: Facility location for subspace segmentation., in: *Proceedings of the International Conference on Computer Vision*, 2009, pp. 825-832.
36. J. Yu, T. Chin, D. Suter, A global optimization approach to robust multi-model fitting, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 2041-2048.
37. A. Delong, O. Veksler, Y. Boykov, Fast fusion moves for multi-model estimation, in: *Proceedings of the European Conference on Computer Vision*, 2012, pp. 370-384.
38. T. T. Pham, T.-J. Chin, K. Schindler, D. Suter, Interacting geometric priors for robust multimodel fitting, *IEEE Transactions on Image Processing* 23 (10) (2014) 4601-4610.
39. <http://www.diegm.uniud.it/fusiello/demo/jlk/>.
40. <http://www.diegm.uniud.it/fusiello/demo/rpa/>.
41. M. Soltanolkotabi, E. Elhamifar, E. J. Candès, Robust subspace clustering, *Annals of Statistics* 42 (2) (2014) 669-699.
42. <http://cs.adelaide.edu.au/~trung>.
43. <http://www.vision.jhu.edu/code/>.
44. P. Denis, J. H. Elder, F. J. Estrada, Efficient edge-based methods for estimating manhattan frames in urban imagery, in: *Proceedings of the European Conference on Computer Vision*, 2008, pp. 197-210.
45. R. Tron, R. Vidal, A benchmark for the comparison of 3-d motion segmentation algorithms, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2007, pp. 1-8.
46. Y. Sugaya, Y. Matsushita, K. Kanatani, Removing mistracking of multibody motion video database Hopkins155, in: *British Machine Vision Conference*, 2013.
47. S. Skiena, *The Algorithm Design Manual*, Springer London, 2009.
48. <http://www.robots.ox.ac.uk/vgg/data/data-mview.html>.
49. <http://www.diegm.uniud.it/fusiello/demo/samantha>.