

Fitting Multiple Models via Density Analysis in Tanimoto Space

Luca Magri¹ and Andrea Fusiello²

¹ Dipartimento di Matematica - Università di Milano - Via Saldini, 50 - Milano, Italy

² DIEGM - Università di Udine - Via Delle Scienze, 208 - Udine, Italy

Abstract. This paper deals with the extraction of multiple models from noisy, outlier-contaminated data. We build on the “preference trick” implemented by T-linkage, weakening the prior assumptions on the data: without requiring the tuning of the inlier threshold we develop a new automatic method which takes advantage of the geometric properties of Tanimoto space to bias the sampling toward promising models and exploits a density based analysis in the conceptual space in order to robustly estimate the models. Experimental validation proves that our method compares favourably with T-Linkage on public, real data-sets.

Keywords: Multi-model fitting, Segmentation, Clustering

1 Introduction

The extraction of multiple models from noisy or outlier-contaminated data is an important and challenging problem that emerges in many Computer Vision applications. With respect to single-model estimation in presence of noise and outliers, this problem is even more difficult since it must tolerate both true outliers and *pseudo-outliers* (“outliers to the structure of interest but inliers to a different structure” [15]). Among the wide variety of approaches developed for multi-model fitting, it is possible to identify two mutually orthogonal strategies: *consensus analysis* and *preference analysis*.

The consensus set of a model is defined as the set of data points that are close to the model within a certain threshold. Consensus analysis can be traced back to the popular RANSAC paradigm and its variants [17] and gave rise also to algorithms tailored for the case of multiple structures estimation, e.g. [22]. Consensus-oriented methods generate a pool of putative model hypotheses by random sampling, then retain the models that explain better the data by inspecting their consensus sets. The same idea can be found in the popular Hough transform [20] and its generalization, where multiple models are revealed as peaks in the hypothesis space after it has been adequately quantized. Maximizing the consensus set of models can be also encountered as the foundation of optimization algorithms [9] for geometric fitting.

On the contrary, preference analysis reverses the role of data and models: rather than considering models and examining which points match them, the residuals of individual data points are taken into account [21]. In this way it is possible to exploit the residual information for building a *conceptual space* in which points

are portrayed by the *preferences* they have accorded to provisional models. For example in [16] data points are represented as characteristic functions taking values on the set of hypothesized models, whereas in [2] a data point is represented in a Reproducing Kernel Hilbert Space by the permutation that arranges the models in order of ascending residuals. T-Linkage [12] extends the ideas beyond [16] and is mainly composed by two steps:

1. *Conceptual representation*: given an inlier threshold, points are represented in the m -dimensional unitary cube endowed with the Tanimoto distance, according to the preferences they grant to a set of m random hypotheses. We will refer to this metric space as *Tanimoto space*.
2. *Segmentation*: data are segmented via a tailored version of average linkage.

It is worth to observe that a similar *first-represent-then-clusterize* approach is also adopted by many state-of-the-art algorithms [7, 14, 11] for multiple subspaces estimation that relies on sparse representation of points.

In this paper we elaborate both the conceptual representation and the segmentation step introduced by T-Linkage, conceiving a new automatic method aimed at multi-model extraction resistant to noise and outliers. In particular we present an embedding in Tanimoto space of data points which does not require the tuning of the inlier threshold and In addition we describe how a density based analysis of the Tanimoto space can be used to simultaneously cluster point and identify outliers. The main goal of this work is to gain some insights on the key advantages provided by the use of Tanimoto distance in preference analysis.

2 Method

In this section we explore the geometrical properties of Tanimoto space, introducing a new method for clustering of multi-model data. We motivate and formulate the algorithm using a working example (Fig. 1a) taken from [19]. In this dataset three objects move independently each giving rise to a set of points correspondences in two uncalibrated images: points belonging to the same object are described by a specific fundamental matrix. Outlying correspondences are present.

2.1 Preference Trick

The first step of our method, like T-Linkage, consists in building a conceptual representation of data points $X = \{x_1, \dots, x_n\}$ in the Tanimoto space $\mathcal{T} = ([0, 1]^m, d_{\mathcal{T}})$, given a set of putative models $H = \{h_1, \dots, h_m\}$. We will denote the Tanimoto distance, defined for every $p, q \in [0, 1]^m$ as

$$d_{\mathcal{T}}(p, q) = 1 - \frac{\langle p, q \rangle}{(\|p\|^2 + \|q\|^2 - \langle p, q \rangle)} \quad , \quad (1)$$

whereas $d: X \times H \rightarrow \mathbb{R}$ indicates a suitable distance for computing the residual between a point and a model. This step can be thought as a sort of “*preference trick*”. Echoing the celebrated “kernel trick”, which lifts a non linear problem in an higher dimension space in which it becomes easier, the Tanimoto representation,

shifts the data points from their ambient space to a conceptual one, revealing the multiple structures hidden in the data as groups of neighbouring points. The preference trick can be formalized by defining a vector map

$$\Phi: X \rightarrow \mathcal{T} = ([0, 1]^m, d_{\mathcal{T}}) \quad x \mapsto (\exp(-r_1), \dots, \exp(-r_m)) \quad , \quad (2)$$

where r_j are the standardized residuals of data points with respect to H :

$$r_j = \frac{d(x, h_j)}{\sigma}, \quad \sigma = \text{var} \{d(x_i, h_j), \quad i = 1, \dots, n, \quad j = 1, \dots, m\} \quad . \quad (3)$$

The rationale behind this construction is that $\Phi(x)$ expresses the preference granted by x to the model h_j with a vote in $[0, 1]$ according to its residual. The Tanimoto distance $d_{\mathcal{T}}$ measures the agreement between the preferences of two points. In particular this distance ranges in $[0, 1]$ and achieves its minimum for points sharing the same preferences, whereas $d_{\mathcal{T}} = 1$ when points have orthogonal preferences.

Observe that Φ differs from the embedding Φ_{τ} proposed in T-Linkage in which the inlier threshold τ is needed in order to cut off the preference of points too distant from the models:

$$\Phi_{\tau,j}(x) = \begin{cases} \exp\left(\frac{-d(x, h_j)}{5\tau}\right) & \text{if } d(x, h_j) < \tau \\ 0 & \text{if } d(x, h_j) \geq \tau \end{cases} \quad . \quad (4)$$

In this formulation τ plays a crucial role: it implicitly controls the final number of models because it affects the orthogonality between the conceptual representation of points and T-Linkage connects points until they are orthogonal in \mathcal{T} .

On the contrary our segmentation step does not rely on orthogonality in \mathcal{T} , but exploits other geometric properties of Tanimoto space, such as densities and concentrations. For this reason we do not require a correct estimate of the threshold, which indeed can be a difficult parameter to tune.

Biased Random Sampling in Tanimoto space. The pool of tentative hypotheses is generated instantiating a model on minimal sample sets (MSS), i.e. a subsets of data points having the minimum cardinality (henceforth denoted by k) necessary to estimate the model. Different sampling strategies can be used for generating these 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 [16]. Hence, many strategies have been proposed in order to guide sampling towards promising models both in the case of one model [5, 4] and in the multiple models scenario [3]. With localized sampling [10] neighbouring points are selected with higher probability, thereby reducing the number of hypotheses that have to be generated. However, depending on the application, introducing a local bias can be difficult since different structures may obey different distributions of data in the ambient space (think for example to the case of motion segmentation of moving object with very different shapes).

In order to overcome this difficulty we propose to sample the hypotheses directly in the conceptual space. This can be easily done performing a preliminary

uniform sampling of hypotheses, representing the data in the Tanimoto space according to these putative models and then doing a biased sampling in \mathcal{T} .

In particular if a point x has already been selected, then a point y such that $x \neq y$ has the following probability of being drawn:

$$P(x|y) = \frac{1}{Z} \exp \frac{d_{\mathcal{T}}(\Phi(x), \Phi(y))^2}{\alpha^2} . \quad (5)$$

where Z is a normalization constant and α controls the local bias.

Tanimoto distances are then updated on the fly based on the hypotheses already sampled.

We illustrate the effectiveness of this sampling strategy on the *biscuitbookbox* sequence. In Fig. 1 we compare our biased sampling in Tanimoto space with respect to uniform sampling, localized sampling, and Multi-GS a method proposed in [3]. All these methods can be lead back to the conditional sampling scheme presented here, substituting $d_{\mathcal{T}}$ in (5) with an appropriate distance function: $d_U \equiv 1$ for uniform sampling, $d_L = \|\cdot\|$ for localized sampling and the intersection kernel d_{GS} . We run these methods with different values of α ; in particular we set $\alpha = q_w$ as the w -th quantile of all these distances, varying $w \in [0.1, 1]$. Our biased sampling provides results comparable with localized sampling for more values of α (Fig. 1b) and produces many pure MSS per model (Fig. 1c).




Fig. 1: Comparison of guided sampling methods on *biscuitbookbox* sequence. Model membership is colour coded; black crosses are outliers.

2.2 Density based analysis for model extraction

In this section we will see how to cluster the high dimensional cloud of points $Y = \Phi(X)$ originated by the preference trick in order to extract the models that explain the data. The desired segmentation is obtained by exploiting the geometric properties of the discrete set of points Y induced by the topology defined in \mathcal{T} . In particular the notion of *density-connected component* will serve as the geometric counterpart of the statistical notion of cluster. An analysis of points density is performed using OPTICS [1], producing a *reachability plot* from which density-connected components can be identified.

Geometric insight. We recall that the Tanimoto distance measures the agreement between the preferences of points: data sharing the same preferences, i.e. belonging to the same models, are close according to $d_{\mathcal{T}}$. As can be appreciated from Fig. 2a, points belonging to the same model are clustered in high density region whereas outliers occupy region with low density.

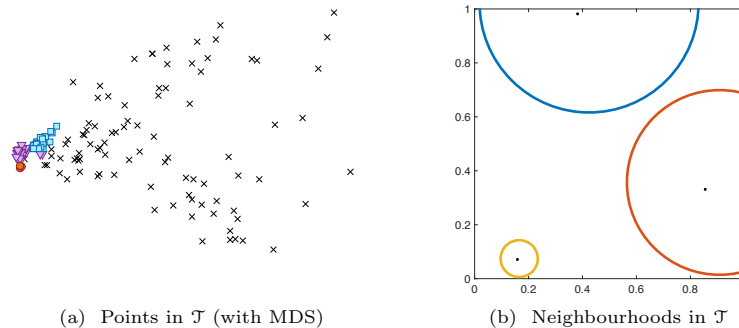


Fig. 2: Insights on the geometry of Tanimoto space. (a) the *biscuitbookbox* data in Tanimoto space are projected using Multi-Dimensional Scaling for visualization purposes. Outliers (x) are recognized as the most separated points. (b) Tanimoto neighbourhoods with the same radius in $[0, 1]^2$ have smaller Euclidean diameter if the centre lies near the origin.

Some insight into this property of Tanimoto space can be reached considering a system of neighbourhoods: Fixed some $\epsilon \in (0, 1)$ and some $y \in \mathcal{T}$ the Tanimoto ball of radius ϵ and centre y is denoted by $N_\epsilon(y)$.

As illustrated in Fig. 2b, the Euclidean diameter of N_ϵ changes accordingly to the position of the centre y . In particular this quantity tends to be smaller for points lying near the origin of \mathcal{T} , that corresponds to the region of \mathcal{T} prevalently occupied by outlying points. In fact outliers grant their preferences to very few sampled hypotheses, they have small Euclidean norm and consequently tend to lie near the origin. Hence the probability that two outliers live in the same ball of radius ϵ is significant lower than the probability that two inliers (with higher Euclidean norm) are contained in a ball with the same radius. For this reason outliers can be recognized as the most separated points in \mathcal{T} .

With this perspective as guide, we tailor the definition of density-connected component [8] to Tanimoto space:

Definition 1. Given $p, q \in \mathcal{T}$, the cardinality k of MSS and $\epsilon \in (0, 1)$

- p is said a core point if $|N_\epsilon(p)| > k$;
- p is directly density-reachable from q with respect to ϵ if $p \in N_\epsilon(q)$ and q is a core point;
- p is density reachable from q with respect to ϵ if there is a chain of points p_1, \dots, p_ℓ s.t. $p_1 = p, p_\ell = q$ and p_{i+1} is directly density reachable from p_i ;

- p is density-connected to point q with respect to ϵ if there is a point o such that both p and q are density reachable from o .
- a density-connected component is a maximal set of density-connected points.

Density-connectivity is an equivalence relation hence all the points reachable from core points can be factorized into maximal density-connected components yielding the desired segmentation. A crucial advantage of this definition is that it deals directly with outliers which can be recognized as points not connected to any core point. In topological words, outliers can be identified as isolated points of Y , whereas inliers are either internal or boundary points of a density-connected component. Moreover density-connected components may have arbitrary shape. Note that, by definition, a density-connected component must contain at least $k + 1$ points; this is coherent with the fact that at least $k + 1$ points are needed to instantiate a non-trivial model (k points always define a model), and gives a very natural choice of this parameter which in [8] is user defined.

Ordering points to identify the clustering structure. In Tanimoto space clusters of inliers could have relatively varying density, though higher than outliers. For this reason we do not fix a global value of ϵ for finding the density-connected components. Instead we adopt the multi-scale approach offered by OPTICS (Ordering Points to Identify the Clustering Structure) [1]. OPTICS is a density-based technique which frame the geometry of the data in a reachability plot thanks to the notion of reachability distance.

Definition 2. Given the cardinality k of MSS,

- if p is a core point, the core-distance of p refers to the distance between p and its k -nearest neighbour.
- if p is a core point, the reachability-distance of a point p with respect to a point q is the maximum between the core distance of p and the distance $d_{\mathcal{T}}(p, q)$.

After the data have been ordered so that consecutive points have minimum reachability distance, OPTICS produces a special kind of dendrogram, called reachability plot, which consists of the reachability values on the y -axis of all the ordered points on the x -axis. The valleys of this plot represent the density-connected regions: the deeper the valley, the denser the cluster.

Figure 3, where the *biscuitbookbox* reachability plot is shown, illustrates this. Outliers have high reachability values, on the contrary genuine clusters appears as low reachability valley and hence are density-connected components in \mathcal{T} .

The points ordering originated by OPTICS resembles the classical single Linkage clustering, in which at each iteration the closest points are linked together, however the use of reachability distance, which exploits local density information, mitigates the so-called chain effect.

Flooding. The final step of our method is aimed at automatically find the valleys in the reachability plot produced by OPTICS in order to robustly extract clusters and hence models. For this purpose we adapt the Meyer flooding algorithm [13],




Fig. 3: Output of OPTICS. (a) The reachability plot of *biscuitbookbox* shows the reachability distance of ordered points (model membership is colour coded according to the ground truth). (b) Significant local minima are marked according to the θ at which they are found.

originally developed for Watershed segmentation, to deal with the 1D reachability plot:

1. Local minima of the reachability plot are found.
2. Those minima that are at least θ below their adjacent maxima are retained as significant.
3. The reachability plot is flooded starting from the basin of significant local minima.
4. The flooding of a basin stops when different water sources meet.
5. Points flooded with water coming from the same source are clustered together, whereas points untouched by the water are labelled as outliers.

Once data points are segmented, models are fitted via least square regression on points belonging to the same cluster, and outliers are reassigned to their nearest model – if they have distance smaller than the furthest inlier. Finally, models are re-estimated according to this new segmentation. Optionally, a final outlier rejection, as the one proposed in [12], could be performed.

The threshold θ controls the significance of critical points in the reachability plot, allowing us to dichotomize between significant minima and spurious ones due to noise. It is found empirically that $\theta = 0.05$ provides good local minima in several applicative scenarios, consequently the generality of the preference trick is not affected by the use of this fixed threshold.

Nevertheless it is interesting to investigate how θ affects the quality of clustering. In Fig. 3b the sets of significant minima are shown for $\theta = 0.05, 0.005, 0.001$. It is readily seen that as θ decreases the induced segmentation becomes finer, creating a hierarchy of models, which allows multiple levels of details to be captured, as can be appreciated in Fig. 4, where points that at level $\theta = 0.05$ obey a single fundamental matrix are split in clusters described by homographies.

3 Experimental evaluation

This section is devoted to validate our methodology (henceforth called T-OPTICS) in several real multi-model fitting problems. All the code is implemented in Mat-




Fig. 4: Hierarchy of models at different levels of θ . Model membership is colour coded, black crosses are outliers.

lab on a personal computer with 2.6 GHz Intel Core i7 processor. We use the OPTICS implementation of [6] to produce reachability plots and the Matlab function `peakfinder`³ for finding significant minima. A quantitative measure of the attained segmentation is obtained using the misclassification error (ME), that counts the percentage of misclassified points.

Video motion segmentation. Given a video taken by a moving camera, motion segmentation tries to identify moving object across the video. This aim is pursued by fitting subspace of dimension at most 4 to the set of features trajectories in \mathbb{R}^{2F} , where F is the number of frames in the video.

We use three video sequences taken from the Hopkins 155 dataset [18], *1RT2TC* from the chequerboard sequences, *cars5* which belongs to the traffic ones and *people1* which is an example of articulated motion. All the trajectories are outlier free. The *cars5* and *people1* videos, which are properly segmented by T-Linkage, are correctly segmented also by our method. The *1RT2TC* dataset is more challenging for T-Linkage, as confirmed by the experiments conducted in [12]. The tuning of a global inlier threshold τ turns out to be a thorny problem, since the subspace describing the background (red) and the one representing the moving box (blue) are close to each other with respect to their distances to the subspace of the moving glass (as can be also deduced by the reachability plot in Fig. 5d, where a small jump separates the first two valleys and an higher one separates the last object). This difficulty affects also T-linkage, which, indeed, under-segments the data (see

³ www.mathworks.com/matlabcentral/fileexchange/25500-peakfinder

Fig. 8b in [12]). On the contrary, as presented in Fig. 5, our methods is able to distinguish correctly the three moving objects.

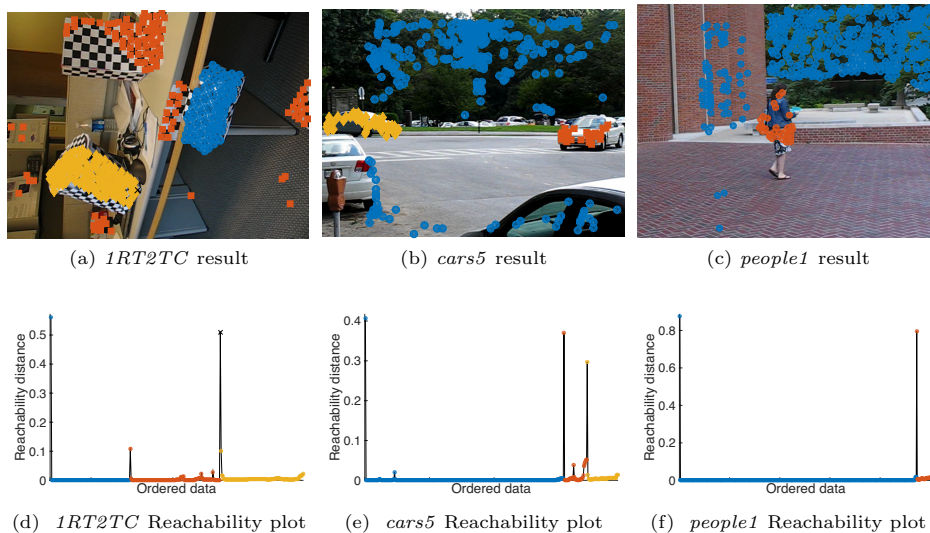


Fig. 5: Video motion segmentation. Top row: segmentation results. Model membership is colour coded, no outliers. Bottom row: reachability plots.

Two views segmentation. In these experiments we are given a set of correspondences across two uncalibrated images. When the scene is dynamic, as in *biscuitbookbox* and in *breadcubechips*, we want to identify the moving objects by fitting fundamental matrices. If the scene is static our purpose is to identify the points belonging to 3D planar surfaces by estimating multiple homographies. The latter is the case of *elderhalla* and *sene* datasets. All these datasets are contaminated by the presence of gross outliers. Segmentation results are presented in Fig. 6. In motion segmentation T-OPTICS succeeds in extracting the correct fundamental matrices and outliers are correctly detected after the flooding step. As plane segmentation is concerned the two dominant planes are correctly identified both in *sene* and in *elderhalla*. In the latter two additional models are detected composed by outliers which happen to lies in homographic configurations.

For each tested sequence the corresponding ME is reported in Tab. 1, where T-OPTICS is compared to T-Linkage, showing that in some cases T-OPTICS increases the performances of T-Linkage without requiring the tuning of the inlier threshold. By collating the time elapsed (in seconds) for the clustering steps of these two algorithms it is evident that T-OPTICS attains the best trade-off between accuracy and computational time. In fact, in T-Linkage every time two clusters are merged, Tanimoto distances need to be updated, whereas in OPTICS distances in \mathcal{J} are computed only once, considerably reducing the computational load.

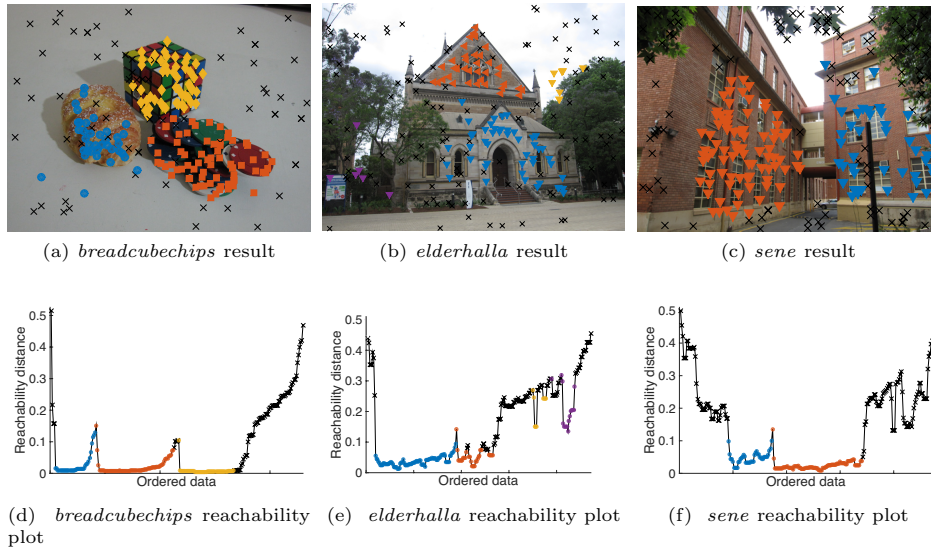


Fig. 6: Two views segmentation. Top row: segmentation results. Model membership is colour coded, black crosses are outliers. Bottom row: reachability plots.

Table 1: Segmentation result: comparison between T-Linkage and T-OPTICS

	T-Linkage		T-OPTICS	
	ME (%)	Time (s)	ME (%)	Time (s)
<i>cars5</i>	0	11.50	0	1.09
<i>1RT2TC</i>	31.39	3.76	0.32	0.51
<i>people1</i>	0	23.92	0	2.35
<i>biscuitbookbox</i>	1.54	3.17	2.70	0.37
<i>breadcubechips</i>	0.86	2.20	3.09	0.25
<i>elderhalla</i>	7.51	1.65	5.14	0.15
<i>sene</i>	0.40	2.43	2.12	0.25

4 Conclusion

We have presented a novel approach to multi model fitting based on density analysis in Tanimoto Space. We leverage the property of Tanimoto space at different levels: at first to guide sampling towards promising models, then to recognize clusters as density-connected components in \mathcal{T} and finally in order to robustly extract models from the reachability plot of the data. With respect to T-Linkage, T-OPTICS produces reliable segmentations in less time.

References

1. Ankerst, M., Breunig, M.M., Kriegel, H.P., Sander, J.: Optics: Ordering points to identify the clustering structure. In: ACM Sigmod Record. pp. 49–60 (1999)
2. Chin, T., Wang, H., Suter, D.: Robust fitting of multiple structures: The statistical learning approach. In: Int. Conf. on Computer Vision. pp. 413–420 (2009)
3. Chin, T.J., Yu, J., Suter, D.: Accelerated hypothesis generation for multistructure data via preference analysis. IEEE Trans. Pattern Anal. Mach. Intell. pp. 533–546 (2012)
4. Chum, O., Matas, J.: Randomized ransac with $T_{d,d}$ test. In: Image and Vision Computing. vol. 22, pp. 837–842 (2002)
5. Chum, O., Matas, J.: Matching with PROSAC - progressive sample consensus. In: Int. Conf. on Computer Vision and Pattern Recognition. pp. 220–226 (2005)
6. Daszykowski, M., Walczak, B., Massart, D.L.: Looking for natural patterns in analytical data, 2. Tracing local density with OPTICS. Journal of Chemical Information and Computer Sciences (3), 500–507 (2002)
7. Elhamifar, E., Vidal, R.: Sparse subspace clustering: Algorithm, theory, and applications. IEEE Trans. Pattern Anal. Mach. Intell. 35(11), 2765–2781 (2013)
8. Ester, M., Kriegel, H.P., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Second International Conference on Knowledge Discovery and Data Mining. pp. 226–231 (1996)
9. Isack, H., Boykov, Y.: Energy-based geometric multi-model fitting. International Journal of Computer Vision 97(2), 123–147 (2012)
10. Kanazawa, Y., Kawakami, H.: Detection of planar regions with uncalibrated stereo using distribution of feature points. In: British Machine Vision Conf. pp. 247–256 (2004)
11. Liu, G., Lin, Z., Yan, S., Sun, J., Yu, Y., Ma, Y.: Robust recovery of subspace structures by low-rank representation. IEEE Trans. Pattern Anal. Mach. Intell. pp. 171–184 (2013)
12. Magri, L., Fusiello, A.: T-linkage: A continuous relaxation of J-linkage for multi-model fitting. In: Conf. on Computer Vision and Pattern Recognition (June 2014)
13. Meyer, F., Beucher, S.: Morphological segmentation. Journal of Visual Communication and Image Representation 1(1), 21 – 46 (1990)
14. Soltanolkotabi, M., Elhamifar, E., Candès, E.J.: Robust subspace clustering. Ann. Statist. 42(2), 669–699 (04 2014)
15. Stewart, C.V.: Bias in robust estimation caused by discontinuities and multiple structures. IEEE Trans. Pattern Anal. Mach. Intell. 19(8), 818–833 (1997)
16. Toldo, R., Fusiello, A.: Robust multiple structures estimation with J-linkage. In: European Conf. on Computer Vision. pp. 537–547 (2008)
17. Torr, P.H.S., Zisserman, A.: MLESAC: A new robust estimator with application to estimating image geometry. Comp. Vis. and Image Underst. (1), 138–156 (2000)
18. Tron, R., Vidal, R.: A benchmark for the comparison of 3D motion segmentation algorithms. In: Conf. on Computer Vision and Pattern Recognition (2007)
19. Wong, H.S., Chin, T.J., Yu, J., Suter, D.: Dynamic and hierarchical multi-structure geometric model fitting. In: Int. Conf. on Computer Vision (2011)
20. Xu, L., Oja, E., Kultanen, P.: A new curve detection method: randomized Hough transform (RHT). Pattern Recognition Letters 11(5), 331–338 (1990)
21. Zhang, W., Kosecká, J.: Nonparametric estimation of multiple structures with outliers. In: European Conf. on Computer Vision. vol. 4358, pp. 60–74 (2006)
22. Zuliani, M., Kenney, C.S., Manjunath, B.S.: The multiRANSAC algorithm and its application to detect planar homographies. In: Int. Conf. on Image Processing (2005)