Fitting Multiple Heterogeneous Models by Multi-class Cascaded T-linkage

Luca Magri, Andrea Fusiello
Università di Udine - DPIA

No neural network have been harmed in the making of this paper

We address the problem of multi-model fitting in the general context where the sought structures can be described by a mixture of heterogeneous parametric models, i.e., of different type or class.

Also referred to as multi-class/multi-model fitting.

Outline of the idea

We assume nested classes: they form a subsumption hierarchy (*e.g.* lines and circles, homographies and fundamental matrices).

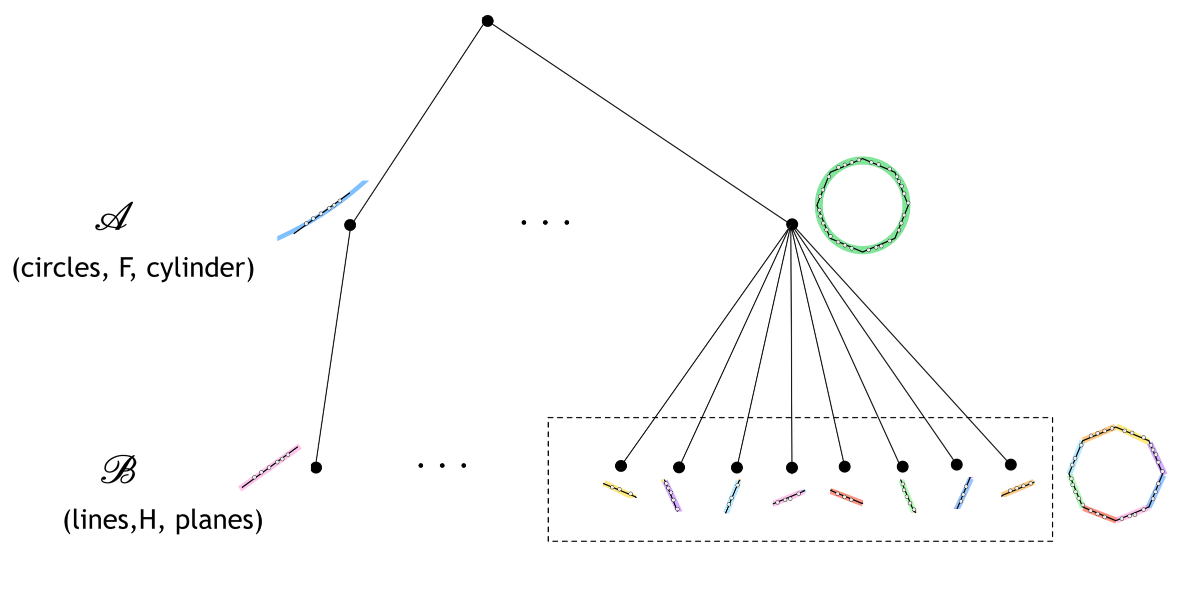

Figure 1: *Hierarchical interpretation of data.*

Let us consider two classes, namely $\mathscr{A}$ and $\mathscr{B}$.

- Let $\mathscr{A}$ be the more "general" class, in the sense that the structures of class $\mathscr{B}$ can be also described in terms of models of type $\mathscr{A}$ without impacting data fidelity.

On the other hand, there are instances in $\mathscr{A}$ that can not be explained by means of a unique model of $\mathscr{B}$ keeping the same level of accuracy.

- Since every structure of type $\mathscr{B}$ can be described by models in $\mathscr{A}$, we may in first instance obtain the segmentation induced by $\mathscr{A}$.

This can be easily obtained with a first run of T-linkage where the pool of tentative hypotheses is instantiated by randomly sampling models of type $\mathscr{A}$.

To recover the structures of type $\mathscr{B}$, we can restrict our search to the refinement of the partition induced by $\mathscr{A}$.

- The solution is to perform T-linkage separately on every $\mathscr{A}$ structure, and then to adopt a model selection criterion to compare a model of type $\mathscr{A}$ with one or several models of type $\mathscr{B}$.

The pool of tentative models of type $\mathscr{B}$ needed to perform T-linkage are generated enforcing geometric compatibility with the type $\mathscr{A}$ models (*e.g.* tangency).

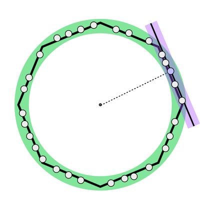

Figure 2: *Tentative lines in circular structures are instantiated by sampling individual points and enforcing the direction of tangency with the circumference.*

The algorithm

Algorithm: *Multi-class Cascaded T-Linkage (MCT)*
1. Extract models of class $\mathscr{A}$ with T-linkage
2. Reject outliers
3. For each $\mathscr{A}$ model
   (a) Extract sub-model(s) of class $\mathscr{B}$ with T-linkage, sampling models $\mathscr{B}$ compatible with $\mathscr{A}$
   (b) Model selection $\mathscr{A}$ vs sub-model(s) $\mathscr{B}$
4. Extract models of class $\mathscr{B}$ from the outliers with T-linkage
5. Reject outliers
6. Run a greeedy set-cover

T-linkage (1, 3(a), 4). A single class multi-model fitting algorithm based on a *first-represent-then-clusterise* scheme: at first, the input data are represented by the "preferences" they grant to a pool of provisional model hypotheses obtained by random sampling, then a greedy bottom-up agglomerative clustering is performed to yield a partition of the data merging points with similar preferences according to the Tanimoto distance [3].

Model selection (3(b)). Let us call $L = -\frac{1}{2}\sum\left(\frac{e_i}{\sigma}\right)^2$ where e_i/σ are the normalised residuals, and let n be the number of data points, d the dimension of model manifold, k the number of model parameters and r the dimension of the measurement space (where data points belong); let $p = dn + k$ the total number of parameters to be estimated.

$$\text{GRIC} = -2L + \lambda_1 dn + \lambda_2 k$$

When $\ell \geq 2$ type $\mathscr{B}$ models are compared against one type $\mathscr{A}$ model, the score is computed substituting k with ℓk to account for the actual number of parameters. In our experiments we set $\lambda_1{=}1$ and adjusted λ_2.

Outlier rejection (2,5) Following the route of the a-contrario approaches, outlying models are pruned using the statistical validation technique described in [4]. In the most simple terms, if there are n inliers at a distance ϵ from a model, assuming a local uniform distribution of the residuals, we expect to have on average κ times more elements at a distance $\kappa\epsilon$. If this does not hqold, the model has a very low probability of occurring by chance and can be retained as inliers.

Greedy set cover (6). Re-assign the inlier points to the models produced by the previous step (each point is assigned to the models for which it is an inlier) and then solve a set cover problem in a greedy fashion, selecting the models in decreasing cardinality order, until all the inlier points have been covered by at least one model. The result is not a partition of the points but a cover.

Comments

- MCT sequentially extract simpler nested models starting form the more general ones. In this way, all the intra-class model selection problems are implicitly solved by T-linkage, whereas the inter-class model selection issues is cast into an explicit comparison of GRIC scores among models explaining the same data.

- This means that MCT does not have to compare all the possible combinations of type $\mathscr{B}$ models against the whole type $\mathscr{A}$ models

- In principle, T-linkage could be used to extract a segmentation from a multi-class soup of heterogeneous models, but in this way one cannot recover the models underlying the clustering, as the class of the models are lost during the clustering fusion, in much the same way as in [5].

- The model selection framework is simple and in principle is agnostic about the multi-model fitting technique adopted. In practice it can be generalised to other preference based multi-model fitting algorithm which can benefit from model-constrained sampling.

Results

The inlier threshold ε and λ_2 have been tuned per-problem (*i.e.*, they are constant across istances of the same problem).

Line and conic fitting In Tab. 1 MCT is contrasted with PEARL [2] in terms of misclassication error (ME). Both methods were given the same preference matrix as input and results come from a single run; the parameters of both have been tuned to achieve the best results and kept fixed in all instances.

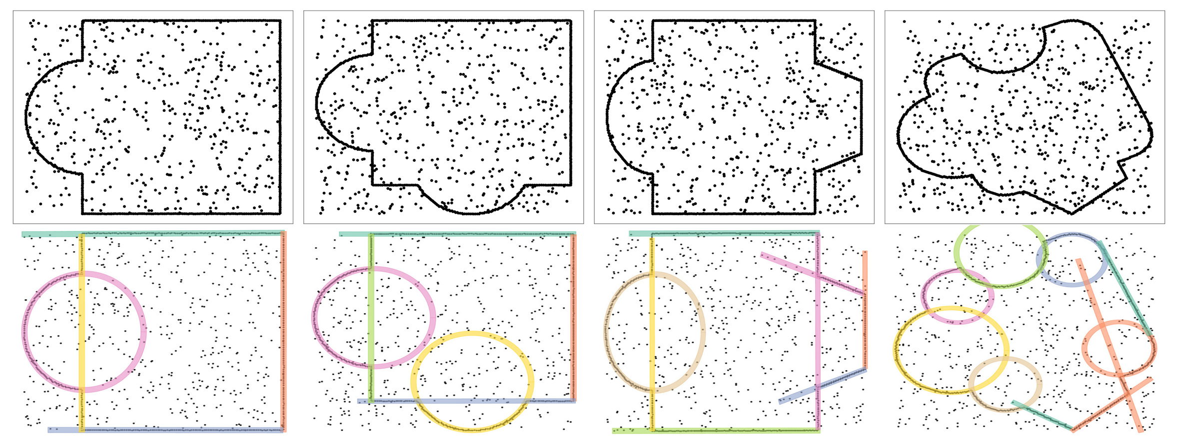

Figure 3: ***Line and conic fitting*** *on "blueprint-like" synthetic data. Top row: input data. Bottom row: detection of multiple lines and circles by MCT.*

PEARL [2]	8.05	8.33	17.38	19.21
MCT	**5.23**	**7.12**	**5.38**	**6.23**

Table 1: *ME (%) for* ***line and conic fitting*** *instances of Fig. 3.*

Plane and cylinder fitting Results on plane and cylinder fitting are presented in Fig. 4. The first two point clouds come from the Aim@Shape repository.

The benefits of the the final refinement with set cover are shown in the bridge example, where portions of the pier have been fitted with cylinders which however ended up to be redundant because they were already covered by the plane that fitted the rest of the facade of the bridge. In the other two cases the refinement brought no visible improvement.

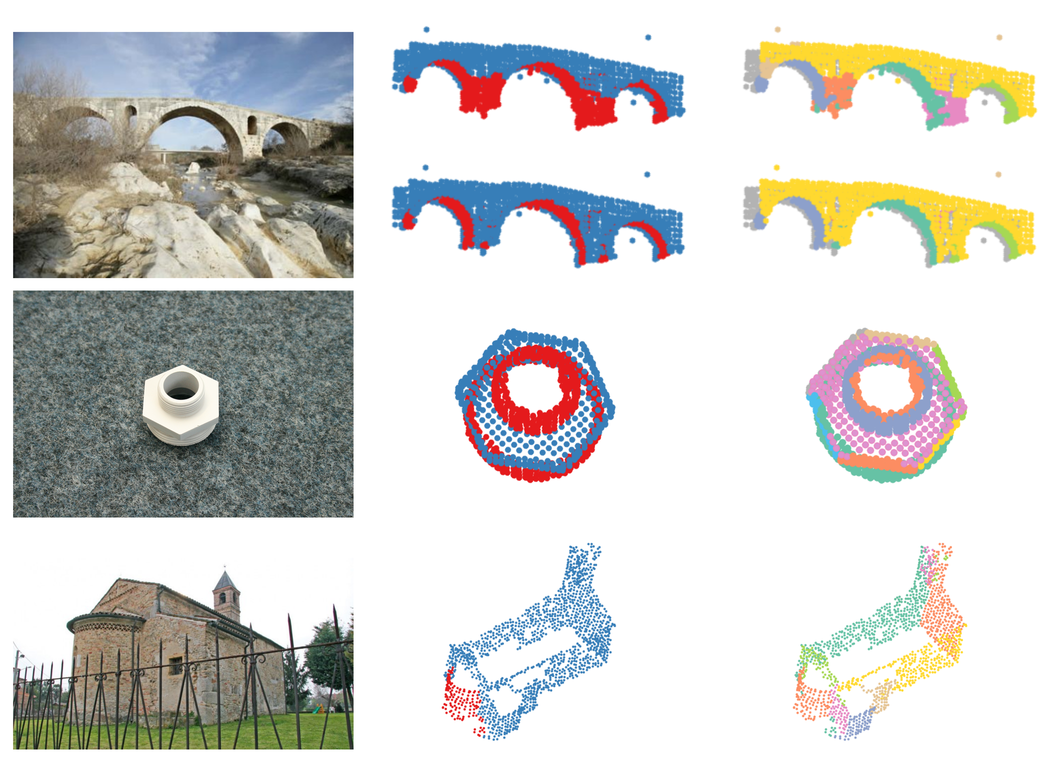

Figure 4: ***Plane and cylinder fitting*** *by MCT on 3D point clouds. Class/model assignment is color coded. For the bridge the class/model assignment is show before (top) and after (bottom) the final refinement with set cover.*

Homography and fundamental matrix fitting We used the "`cube*`" subset of the AdelaideRMF motion dataset, consisting of 8 image pairs depicting a cube and other objects. We manually labelled correspondences according to the face of the cube they belong to, in order to have a true multi-class problem, where homographies (cube faces) and fundamental matrices (other objects) coexist, and the homographies are subsumed by a fundamental model.

	Multi-H	Multi-X	T-lnkg	MCT
mean	14.35	9.72	6.60	**6.13**
median	9.56	**2.49**	4.68	4.93

Table 2: *ME (%) on AdelaideRMF for* ***multiple homography fitting****. The first two columns are copied from [1].*

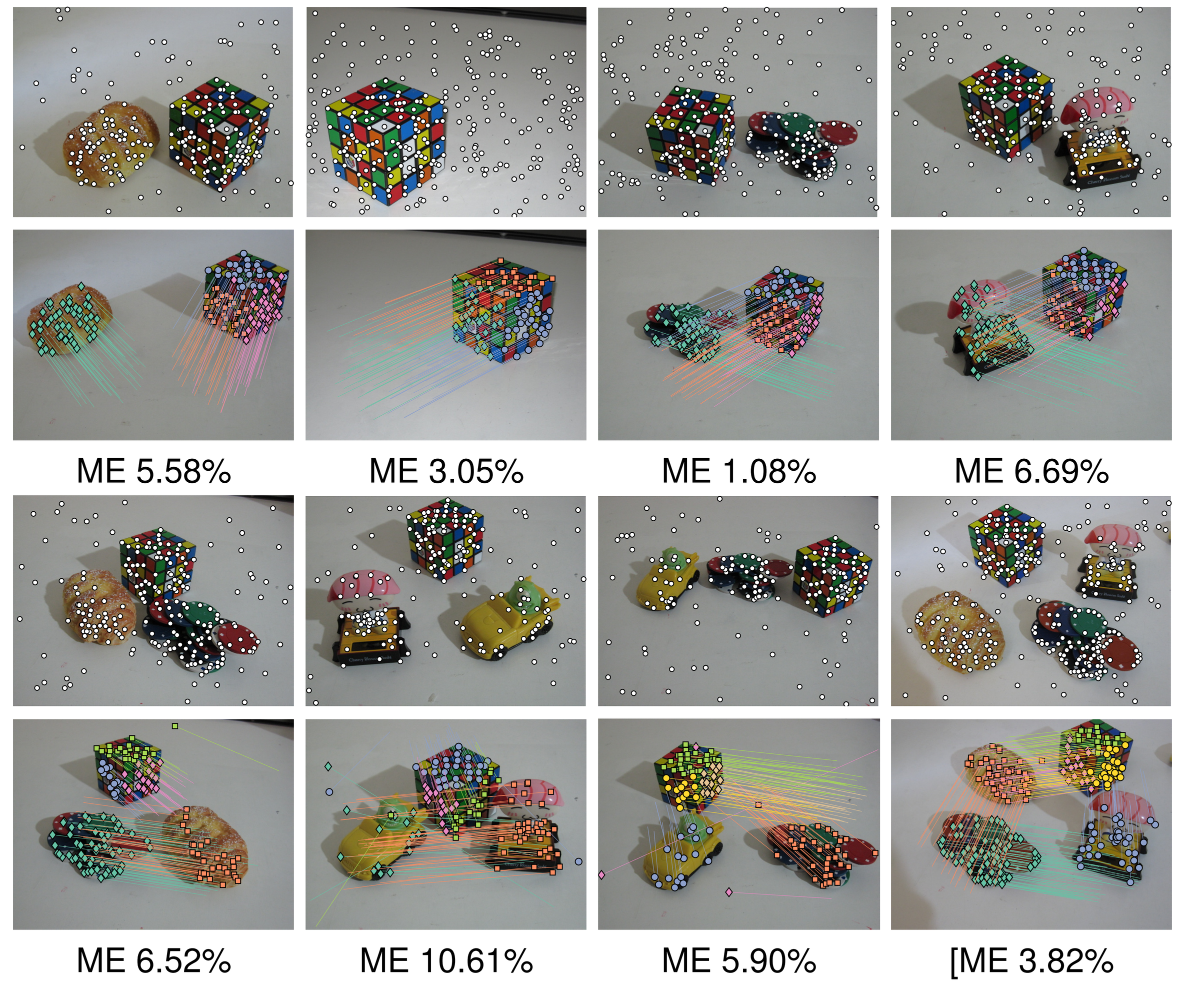

Figure 5: ***Homography and fundamental matrix fitting*** *by MCT on* `cube*` *image pairs. Odd rows: first input image with outlier-contaminated points overlaid. Even rows: F and H models superimposed on the second image. The mean run time per image pair is 21s in Matlab on 2.6 GHz i7 machine.*

Following the same protocol used in [1], we applied MCT to a single class problem, namely **multiple homography fitting** on the AdelaideRMF homography dataset, consisting of 19 image pairs with ground truth point correspondences assigned to planes (homographies). In order to bias MCT toward homographies we set $\lambda_1 = 500$ and $\lambda_2 = 0$ in this experiment.

In this case MCT works almost like regular T-linkage, except a first clustering is made with F matrices that biases the subsequent sampling of homographies.

The Matlab code of MCT is available on-line (QR-code below)

References

[1] D. Barath and J. Matas. Multi-class model fitting by energy minimization and mode-seeking. In *ECCV*, pages 229–245, 2018.

[2] H. Isack and Y. Boykov. Energy-based geometric multi-model fitting. *IJCV*, 97(2):123–147, 2012.

[3] L. Magri and A. Fusiello. T-Linkage: A continuous relaxation of J-Linkage for multi-model fitting. In *CVPR*, pages 3954–3961, 2014.

[4] M. Tepper and G. Sapiro. Fast L1-NMF for multiple parametric model estimation. *arXiv preprint arXiv:1610.05712*, 2016.

[5] X. Xu, L. F. Cheong, and Z. Li. Motion segmentation by exploiting complementary geometric models. In *CVPR*, pages 2859–2867, 2018.

CVPR 2019