Image and Vision Computing 67 (2017) 1–15

Contents lists available at ScienceDirect

Image and Vision Computing

journal homepage: www.elsevier.com/locate/imavis

Multiple structure recovery via robust preference analysis☆,☆☆

Luca Magri[a,*], Andrea Fusiello[b,*]

[a] Dipartimento di Informatica, University of Verona, Strada le Grazie, Verona 15-37134, Italy
[b] DPIA, University of Udine, Via Delle Scienze, Udine 208-33100, Italy

ARTICLE INFO

Article history:
Received 13 July 2016
Received in revised form 24 May 2017
Accepted 9 September 2017
Available online 20 September 2017

Keywords:
Multi-model fitting
Model estimation
Spectral clustering
Matrix factorization

ABSTRACT

This paper address the extraction of multiple models from outlier-contaminated data by exploiting preference analysis and low rank approximation. First points are represented in the preference space, then Robust PCA (Principal Component Analysis) and Symmetric NMF (Non negative Matrix Factorization) are used to break the multi-model fitting problem into many single-model problems, which in turn are tackled with an approach inspired to MSAC (M-estimator SAmple Consensus) coupled with a model-specific scale estimate. Experimental validation on public, real data-sets demonstrates that our method compares favorably with the state of the art.

1. Introduction

Geometric multi-model fitting aims at extracting parametric models from unstructured data in order to organize and aggregate visual content in suitable higher-level geometric structures[1]. This ubiquitous task can be encountered in many Computer Vision applications, for example in 3D reconstruction, where it is employed to estimate multiple rigid moving objects in order to initialize multi-body structure from motion (*e.g.*, [1]), or in the processing of 3D point clouds, where planar patches are fitted to produce intermediate geometric interpretations (*e.g.*, [2]).

Several challenges are afoot. First, segmentation and estimation tasks exhibit a chicken-and-egg pattern, for they are two closely entangled aspects of the same problem: data should be segmented based on their geometric proximity to structures whose unknown parameters must be estimated at the very same time. In other words, in order to estimate models one needs to first segment the data, but conversely, in order to segment the data it is necessary to know the structures associated with each data point.

In addition, the presence of multiple structures hinders robust estimation. Not only visual data are typically affected by arbitrarily large measurement errors – and require the adoption of robust estimators – but the multi-modality of the data makes the problem even more demanding, as it is necessary to cope also with pseudo-outliers, a concept introduced by Stewart [3] for describing those measurements that do not match a model of interest because they are inliers of a different structure.

Moreover, the problem is inherently ill-posed, since many different interpretations of the same data are possible. Making the problem tractable requires a regularization strategy that constrains the solution using prior information, usually in the form of one or more parameters, such as the number κ of sought structures. Following the Occam's razor principle – that one should not presume more things than the required minimum – κ should be kept as low as possible, but finding a correct trade-off between data fidelity and model complexity (a.k.a. bias-variance dilemma) is an intricate task, related to the *model selection* problem. Unfortunately estimating this quantity turns to be a thorny problem, and, for this reason, in many scenarios is assumed known.

1.1. Outline

In this article we present an original method henceforth dubbed RPA (robust preference analysis) which attempts to disentangle the

☆ A preliminary version of this work appeared in [53].
☆☆ This paper has been recommended for acceptance by Enrique Dunn.
* Corresponding authors.
E-mail addresses: magri.luca.l@gmail.com (L. Magri), andrea.fusiello@uniud.it (A. Fusiello).
[1] The "structure" pertains to the arrangement and relations among the data, and it is intrinsic to the data itself, whereas the "model" is the mathematical description of the data that an observer fit onto them. Notwithstanding, we will use the two words interchangeably.

https://doi.org/10.1016/j.imavis.2017.09.005
0262-8856/

chicken-and-egg nature of multiple structure recovery reducing it to many single robust model estimation problems. In particular, three main steps can be singled out: First, we employ M-estimator to shift the problem into a conceptual space where data points are depicted by the *preference* they have granted to a pool of tentative structures, instantiated via random sampling. Second, a robust version of spectral clustering is presented: Robust Principal Component Analysis and Symmetric Non negative Matrix Factorization are employed to explicitly model the low rank nature of inlier preferences, in order to produce an accurate soft-segmentation of data. Third, this provisional segmentation is profitably combined with the initial preference representation in a MSAC-like framework to recover the sought structures. A noise scale estimate is computed for each model, with the help of robust statistic.

The next section offers a perspective on the literature that is most relevant to our work. Next (Section 2) our robust preference analysis method is detailed. In Section 2.1 we present the preference representation, devoting Section 2.2 to address the problem of sampling model hypotheses. In Section 2.3 we concentrate on the clustering problem and in Section 2.3.1 we explain how Low Rank matrix factorization techniques can be tailored to our preference embedding, paving the way to Section 2.4 where the robust structure-recovery strategy is described. Finally, in Section 3 we explore the performance of our method on public real datasets. Appendix A reviews some ideas, firstly emerged in the context of subspace clustering, that are extended to our general multi-model fitting problem.

1.2. Related work

The analysis of *consensus* and its dual counterpart, the analysis of *preferences*, can be traced as a *fil rouge* linking the vast literature on multi model geometric fitting. The consensus of a model is defined as the set of data points that fit the model within a certain inlier threshold ϵ; likewise, the preference set of a point is the set of models to which that point is closer than ϵ.

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 the focus is on models that have to describe as many points as possible. On the other hand, preference approaches concentrate on the segmentation side of the problem, for they are aimed at finding a proper partition of the data, from which estimation follows. In this section we attempt to retrace the path that, starting from consensus throughout preference analysis, have been followed in the literature to address the challenging issues presented by multiple structures recovery.

1.2.1. Consensus analysis

Consensus-oriented methods start with a pool of provisional model hypotheses, that are randomly instantiated on minimal sample sets (MSS), *i.e.*, samples composed by the minimum number of data points necessary to define a structure. Then the consensus sets of the models are inspected, and the models that better explain the data are kept. This idea is at the core of the well-known RANSAC (Random Sample Consensus) algorithm and derivations, *e.g.*, MSAC (M-estimator Sample Consensus) and MLESAC (Maximum Likelihood Estimation Sample Consensus) [4]. Many ameliorations of RANSAC have been proposed in the literature, *e.g.* [5–8], most of which have been surveyed in [9]. In the case of multiple models, Multi-RANSAC [10] and its modifications [11,12] rely on the same principle; also the usual Hough transform and its randomized version [13] can be regarded as consensus-based algorithms, where models are detected as consensus maxima in a quantized hypothesis space. The approach presented in [14] combines random sampling, scale estimation and Mean Shift to determine the consensus set when models are multiple linear subspaces. More in general, maximizing the consensus of models is the foundation of many optimization-based geometric fitting algorithms [15].

1.2.2. Preference analysis

Preference analysis, introduced by [16], also start with a pool of provisional model hypotheses, but it swaps the role of data and models: rather than considering models and examining which points match them, the preference sets of individual data points are inspected. In particular, [17] build a *conceptual* space in which points are portrayed by the *preferences* they have accorded to provisional models. Within this category, J-Linkage [17] and T-Linkage [18] share the same *first-represent-then-segment* scheme: at first data are represented, respectively, either as characteristic functions or as continuous ones taking values on the hypothesized models, then the preference representations are segmented by greedy bottom-up clustering, exploiting either the Jaccard [19] or the Tanimoto [20] distances in order to measure the agreement between preferences, and using the fact that preferences of inliers from the same structure are correlated. This "preference trick" is a very flexible mechanism that can be applied to a wide varieties of scenarios requiring few assumptions on the desired structures. It is sufficient to have at disposal an error function, aimed at measuring residuals and then the structure recovery problem is shifted in the preference space where it can be addressed using cluster analysis.

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

Multi-model fitting has been successfully cast as higher order clustering problems [22–25], which implicitly adopt a preference based approach. In these works higher order similarity tensors are defined between n-tuple of points as the probability that these points are clustered together. In practice, this measure is approximated exploiting the residual error of the n points with respect to provisional models; this preference information is encoded in a hypergraph or a multi-way order tensor, which are properly reduced to pairwise similarity and fed to spectral clustering-like segmentation algorithms.

For instance, Sparse Grassmann Clustering (SGC) [24], approximates the multi-way similarity tensor as the Gramian matrix defined by the inner product of points in the preference space, hence, following the spirit of spectral clustering, projects the Gramian to its best low rank approximation in a least square sense, using Grouse [26]. At the end, the rows of this approximated matrix, are considered as a new representation of the data points in a proper low dimensional space, and are consequently segmented with k-means.

Alternatively, instead of solving a point-point clustering problem, one can formulate a point-model bi-clustering problem directly on the preference hypergraph [27].

It goes without saying that the state-of-the-art on multi-model fitting can be also described along other dimensions. For example multiple structures recovery can be seen by an optimization perspective as the minimization of a global energy functional composed by two terms: a modeling error which can be interpreted as a likelihood term, and a penalty term encoding model complexity mimicking classical MAP-MRF objectives. A survey of multi-model fitting methods form this point of view can be found in [15].