Multi-Model Estimation in the Presence of Outliers

David F. Fouhey - Advisor: Daniel Scharstein
Department of Computer Science – Middlebury College - Middlebury, VT 05753

Abstract

Abstract: The estimation of models or structures from outlier-contaminated data containing multiple models has a large number of applications in computer vision, the study of the automated understanding of visual data: for instance, geometric figures may be detected from 2D points, and planar surfaces in a scene may be found in pairs of images of the scene using feature matches. This thesis describes a number of contemporary algorithms for multi-model estimation and some of their historical antecedents, as well as an evaluation methodology for the multi-model estimation problem.

Three problem instances are depicted to the right. Left, right: the detection of lines and circles from 2D data points; center: the detection of planar surfaces from feature correspondences.

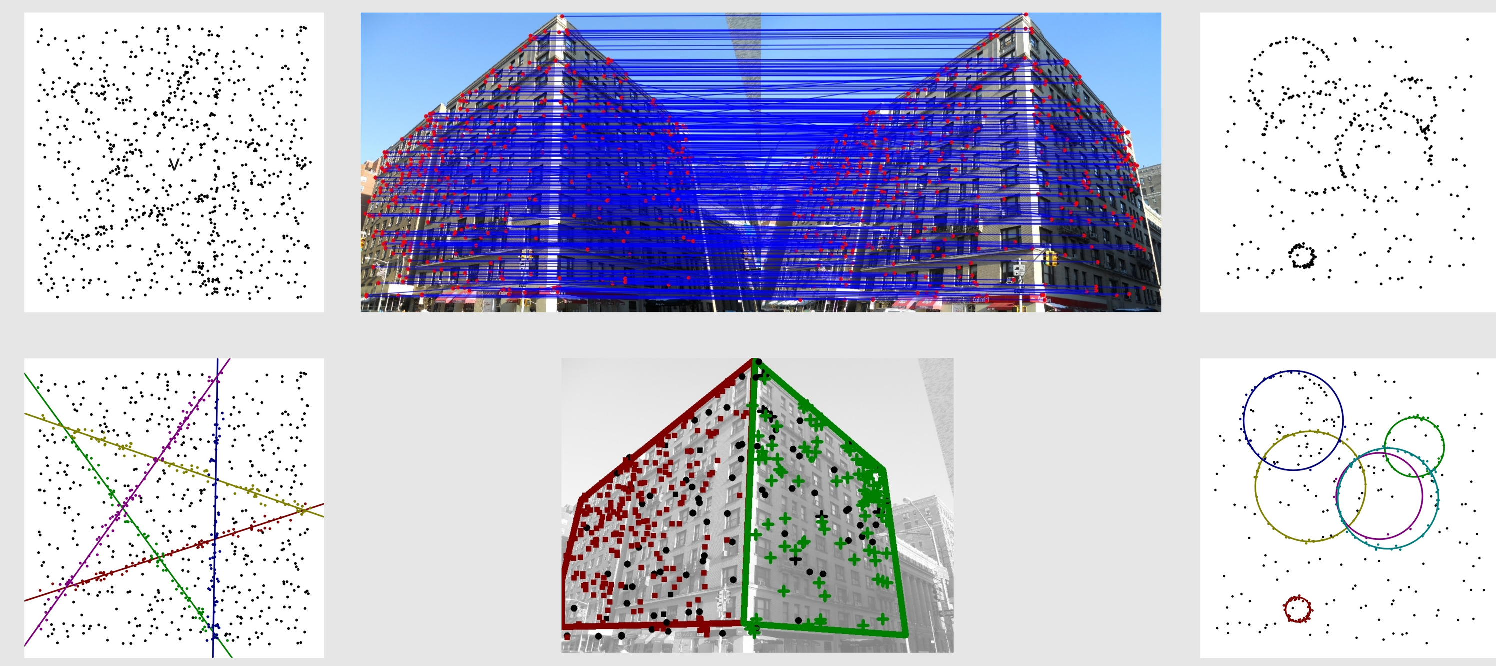

Overview

This thesis describes the development of solutions to the detection of multiple structures from sets of discrete data points. In contrast to techniques such as ordinary or total least-squares, multi-model algorithms do not aim to minimize an objective function over all data points and are not guaranteed to succeed. Instead, they use random sampling to approximate aspects of the model space and heuristic strategies to extract structures from data.

After discussing historical antecedents for outlier-free data (not presented on this poster), we introduce RANSAC [4], a classic approach for the single-model case. We then discuss a number of contemporary approaches for handling the case of multiple models. Finally, we discuss a methodology for the evaluation of multi-model estimation algorithms.

Single Model Estimation

To estimate a model in the presence of outliers, Random Sample Consensus (RANSAC) [4] employs the strategies of random sampling and validating models by their consensus set size. The consensus set of a model is defined as the set of data points with error with respect to that model below some inlier threshold ε. Informally, it is the set of points that fit a model. RANSAC aims to find a model that maximizes this rather than minimize an error function over all data points.

To accomplish this, RANSAC randomly draws minimum sample sets (MSS), or sets with the minimum number of points required to estimate a model (e.g., two points for a line), and fits models to the points.

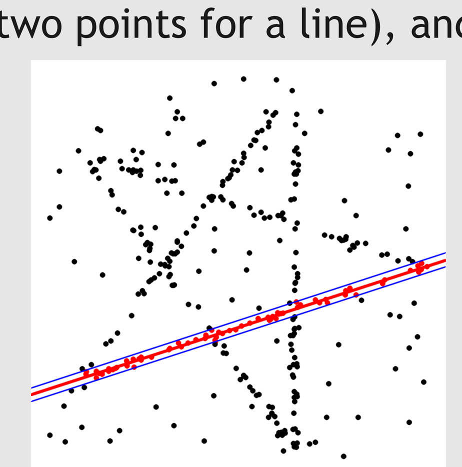

Left: the points in the consensus set of the red line are colored red; the region included in the consensus set is bounded with blue lines.

Right: four randomly drawn minimum sample sets. The size of the consensus set is indicated for each.

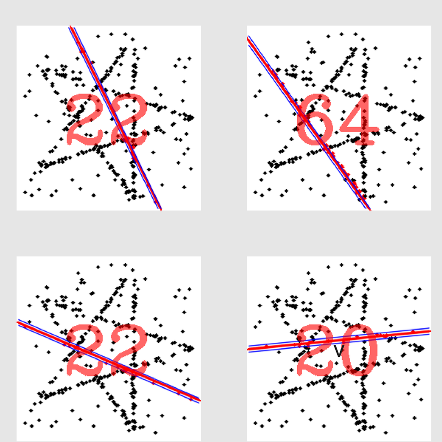

RANSAC is not guaranteed to succeed, since its sampling strategy might not draw an MSS that captures a model present. An approach that increases the chance of drawing a coherent MSS ascribed to Kanazawa et al. [6] is often used: one point is selected uniformly and subsequent points are selected nearby. Further, RANSAC requires a-priori knowledge of the scale of the noise and a threshold on consensus set size to distinguish valid models.

RANSAC as originally proposed is incapable of estimating multiple models, but a simple extension named Sequential RANSAC, depicted below, permits it: Subject to a stopping condition RANSAC is repeatedly performed; after each iteration, the consensus set of the detected model is removed from the data.

A depiction of Sequential RANSAC: stopping conditions aim to detect the absence of valid models in the data, as in the far right image.

Multi-Model Estimation

Starting in 2005, a number of approaches, aiming to perform better than Sequential RANSAC, have been proposed:

- **MultiRANSAC** [12] iteratively updates a collection of consensus sets by repeatedly drawing a collection of MSSs.
- **Residual Histogram Analysis** (RHA) [11] examines the histogram of the residuals of each point with respect to a collection of models fitted to MSSs. RHA requires no a-priori knowledge of the data set.
- **J-linkage** [9] defines the preference set of a point as the set of models that a point fits well enough (i.e., the consensus set with the role of models and points reversed), and the preference set of a set of points as the set of models that every point fits well enough. Beginning with each data point in its own cluster, clusters with minimum Jaccard distance between their preference sets are merged, until no two clusters prefer a model in common.

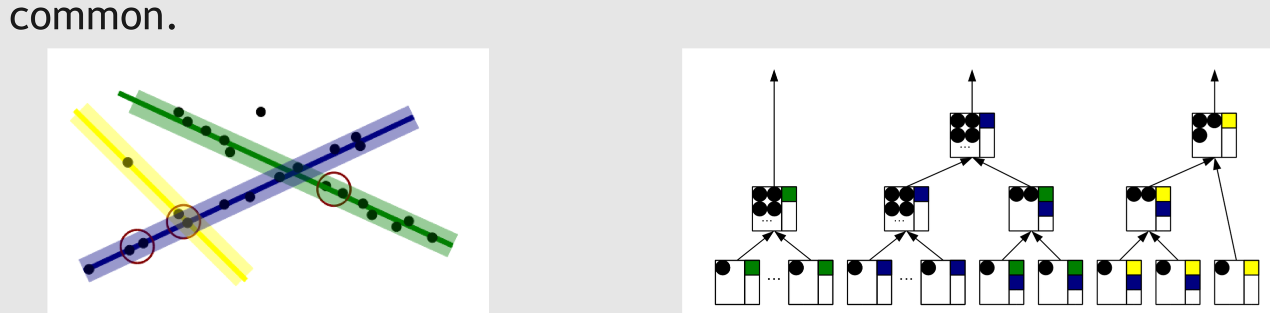

A depiction of J-linkage. Left: MSSs are drawn and models (Green, Blue, Yellow) are fitted. Right: clusters of points are merged. Each cluster is represented by a box: the left compartment contains the clusters' points; the right, its preference set.

- **Merging J-linkage** [5] aims to correct J-linkage's tendency to fragment models, which was observed in [5,8]. It continues merging clusters with a distance defined by the average error under the least-squares fit, until the minimum average error exceeds ε.
- **Kernel Fitting** [1] clusters points by computing the Ordered Residual Kernel between pairs of data points, inducing a feature space that represents data points saliently. Linear algebra techniques are used to remove outliers and cluster data points to find models. Kernel fitting requires no a-priori knowledge of the data set.

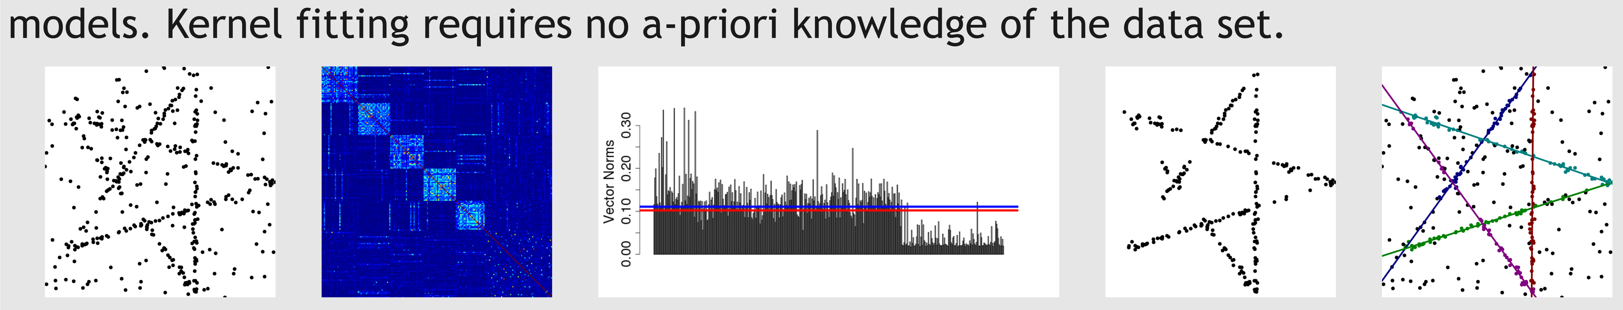

Some steps of Kernel Fitting: (a) the input data; (b) A visualization of the kernel matrix; the kernel function evaluates to high values for data points belonging to the same structure; (c) outlier detection is accomplished by thresholding the norm of the data points projected onto a subspace; (d) The result of outlier removal; (e) the final result, detected by clustering in a salient subspace.

Evaluation of Multi-Model Algorithms

Given the range of approaches for multi-model estimations, one might be curious which to use on a particular task. One considers the a-priori knowledge required; the efficiency; and the accuracy of an algorithm in selecting it. The first two are simple to characterize; we developed a novel quantitative approach to evaluate the latter.

We evaluated algorithms on two geometric figure fitting tasks (finding lines and circles in 2D data points) and a plane fitting task. We use standard synthetic data for figure-fitting, and real correspondence data for plane-fitting.

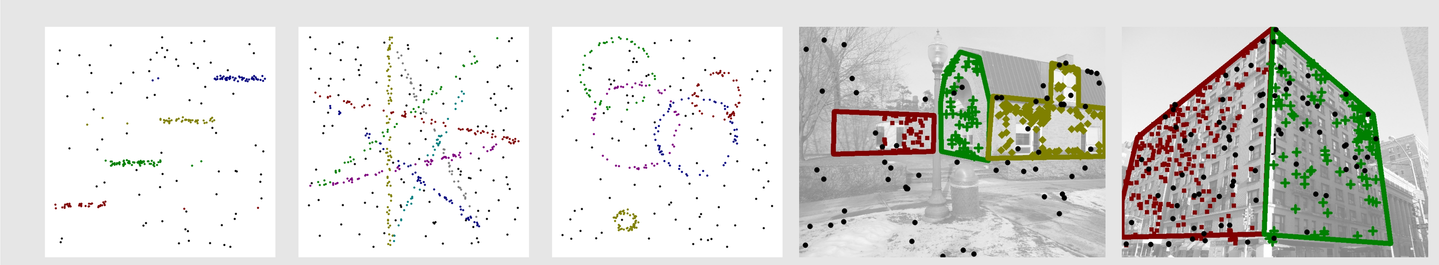

Some of our data sets. Left to right: stairs4, star7, circles5, planes3, planes2

We generate test sets with a range of noise and outlier levels and run each algorithm on each test set 15 times. We then compare the algorithm's estimated models with the ground-truth, and quantify the accuracy with an automatic scoring metrics. Developing these metrics is challenging since a metric must establish correspondences between estimated and ground-truth models and produce accurate results for a wide range of degenerate configurations, some of which are depicted below.

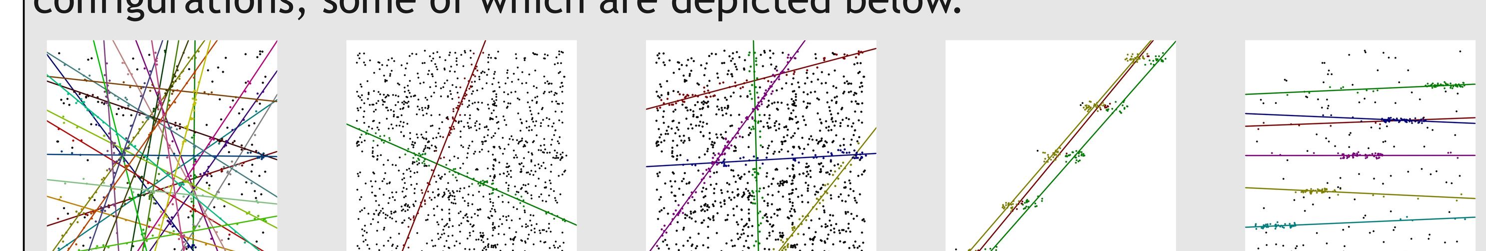

A selection of common degenerate configurations

We both manually analyzed the output and characterized it in the aggregate using automatic scoring metrics. We observed that Sequential RANSAC was fast and often sufficient, but that in some cases, one might need a more complicated algorithm. J-linkage offers increased performance at the cost of increased runtime, and Kernel Fitting offers increased performance and no requirement of a-priori knowledge at the cost of difficult implementation and increased runtime. In the future, we hope to include more algorithms, as well as a motion-segmentation task and real-world geometric figure-fitting.

Selected References

[1] T.-J. Chin, H. Wang, and D. Suter. Robust fitting of multiple structures: The statistical learning approach. In Proc. *ICCV* 2009.
[2] S. Choi, T. Kim, and W. Yu. Performance evaluation of RANSAC family. In Proc. *BMVC* 2009.
[3] A. Delong, A. Osokin, H. Isack, and Y. Boykov. Fast approximate energy minimization with label costs. In Proc. *CVPR* 2010.
[4] M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. In CACM 1981.
[5] D. Fouhey, D. Scharstein, and A. Briggs. Multiple plane detection in image pairs using J-linkage. In Proc. *ICPR* 2010.
[6] Y. Kanazawa and H. Kawakami. Detection of planar regions with uncalibrated stereo using distributions of feature points. In Proc. *BMVC* 2004.
[7] R. Szeliski. Computer Vision: Algorithms and Applications. Springer London, 2010.
[8] J.P. Tardif. Non-iterative approach for fast and accurate vanishing point detection. In Proc. *ICCV* 2009.
[9] R. Toldo and A. Fusiello. Robust multiple structures estimation with J-linkage. In Proc. *ECCV* 2008.
[10] R. Tron and R. Vidal. A benchmark for the comparison of 3-d motion segmentation. In Proc. *CVPR* 2007.
[11] W. Zhang and J. Kosecká. Nonparametric estimation of multiple structures with outliers. In *Dynamical Vision, ICCV 2005 and ECCV 2006 Workshops*, 2006.
[12] M. Zuliani, C.S. Kenney, and B.S. Manjunath. The multiRANSAC algorithm and its application to detect planar homographies. In Proc. ICIP 2005.