

Robust Multi-Model Fitting using Density and Preference Analysis

Lokender Tiwari¹, Saket Anand¹, Sushil Mittal²

¹IIT-Delhi, ²Scibler Corporation, Santa Clara

Abstract. Robust multi-model fitting problems are often solved using consensus based or preference based methods, each of which captures largely independent information from the data. However, most existing techniques still adhere to either of these approaches. In this paper, we bring these two paradigms together and present a novel robust method for discovering multiple structures from noisy, outlier corrupted data. Our method adopts a random sampling based hypothesis generation and works on the premise that inliers are densely packed around the structure, while the outliers are sparsely spread out. We leverage consensus maximization by defining the residual density, which is a simple and efficient measure of density in the 1-D residual space. We locate the inlier-outlier boundary by using preference based point correlations together with the disparity in residual density of inliers and outliers. Finally, we employ a simple strategy that uses preference based hypothesis correlation and residual density to identify one hypothesis representing each structure and their corresponding inliers. The strength of the proposed approach is evaluated empirically by comparing with state-of-the-art techniques over synthetic data and the AdelaideRMF dataset.

1 Introduction

Many computer vision applications require estimation of parameters of a mathematical model from a given set of observations. These observations (or features) are typically obtained through a process agnostic to the model being fit and therefore may contain gross outliers, i.e., points that do not belong to any structure. In order to fit the correct model and identify the inlier points, it is imperative for the estimator to be robust to gross outliers and have a high breakdown point. In cases when observations arise from multiple structures, inliers of one structure appear as outliers to the other structures. Such points have come to be called as pseudo-outliers. With multiple structures in the data, the fraction of outliers (both gross and pseudo) can quickly go in excess of 90%. This makes the requirement of a higher breakdown point more stringent.

Traditional robust regression techniques like least median of squares prove to be lacking due to low breakdown points. To achieve higher levels of robustness, random sampling based methods have become popular since RANdom SAMpling Consensus (RANSAC) algorithm [4]. RANSAC and its many variants rely on consensus set maximization for selecting a model from a set of model hypotheses

generated by random sampling. In practice, RANSAC and its variants have been very successful at several geometric fitting problems in computer vision [7, 14, 18]. However, its performance crucially depends on a user-specified inlier threshold measuring the *scale* of the inlier noise.

Apart from RANSAC like *consensus* based methods, another class of robust multi-model fitting algorithms that have emerged use *preference analysis* [13, 1, 3, 8, 9, 12]. These methods rely on rank-ordering the model hypotheses to generate a preference list for each data point. These preference lists define a feature space where an appropriate distance measure like Jaccard distance [13] or Tanimoto distance [8] could be used to cluster the data points. Preference analysis has also been applied successfully to guided sampling for hypothesis generation [3, 15], where preference lists were used to compute similarity between points. In order to recover all the structures in the data, these methods also require a user-defined parameter that is commensurate to the scale of inlier noise [13, 8].

User defined thresholds are not amenable to real-world scenarios where the inlier noise scale may vary over time or may not be known a priori. Moreover, in case of multiple structures, each structure may have different inlier noise distribution, necessitating a data-driven strategy for scale estimation. Inliers yield a small residual value whereas outliers (or pseudo-outliers) can have arbitrarily large residuals. Geometrically, we can say that inliers are densely packed around the regression surface, while outliers are spread sparsely in the ambient space. The density of points around the regression surface is a key factor in distinguishing between inliers and outliers.

Fig. 1 uses an illustrative example to explain the different steps involved in our proposed technique. Fig. 1a shows an example (‘sene’) from the AdelaideRMF [16] planar segmentation dataset. The task is to fit two homographies corresponding to the two planes in the scene and identify the inliers in both

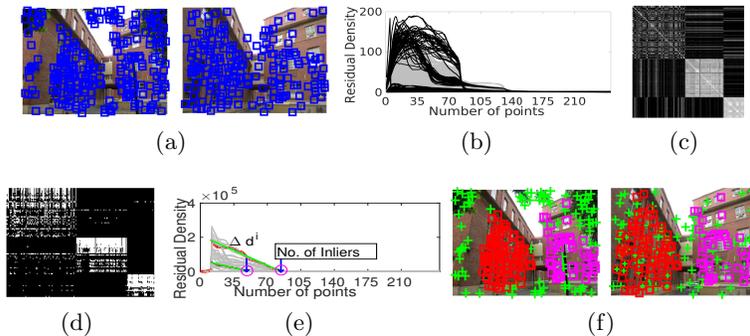


Fig.1: An illustrative example - Density Preference Analysis (DPA), (a)Data, (b)Residual density profile with darker curves indicating selected candidate hypotheses, (c) Density based point correlation, (d)Strong inlier candidates of each hypothesis arrange in the order of structure membership(row wise) and point membership(column wise), (e) Inlier scale (fraction) estimation, (f) Model selection and final segmentation.

cases. We first generate model hypotheses, either using a random sampling or a guided sampling based approach. For each hypothesis, the residual density profile (Section 3.1) is computed. Fig. 1b shows the residual density plotted against the points ordered in increasing order of residuals. The grey curves show the residual density profiles for *all* model hypotheses, while the dark curves show those of selected candidate hypotheses (Section 4.1). We compute the density based point preferences (Section 3.3), which is used to obtain the point correlation matrix as shown in Fig. 1c. For better visualization, the points are ordered by group membership, with the bottom two blocks representing inliers corresponding to the two structures. The point correlation matrix and the selected hypotheses are used to identify strong inlier candidates for each hypothesis as shown in Fig. 1d. The rows correspond to the selected candidate hypotheses, while the columns represent the points. Again, for better visualization, both rows and columns have been reordered using ground truth group membership. For each row, the white pixel indicates a strong inlier candidate. We can see that several rows have strong inliers that comprise of points mostly from the same group. We use these points and the density drop rate to estimate the scale for each model hypothesis (Section 4.2). Fig. 1e shows the number of inliers identified using the scale estimates. In Fig. 1f, we show the final segmentation obtained by applying our model selection algorithm (Section 4.3).

Consensus based methods capture the quality of each model hypothesis based on the consensus set maximization criterion. On the other hand, preference based methods naturally allow measuring the similarity (distance) between a pair of points or a pair of models. By leveraging both these strategies, our technique is able to automatically detect the number of structures in the data, estimate their scale of inlier noise for each structure, perform the point-model assignment of inliers, and identify the gross outliers. We summarize our contributions below:

- We define the residual density profile (Section 3.1), a simple measure of density that effectively captures the disparity between density of inliers and outliers for a given model.
- We define the density based point preferences (Section 3.3) and use it for inlier scale estimation (Section 4.2).
- We develop a simple algorithm that uses density and preference based hypothesis correlations for selecting one model for each structure and identifying their associated inliers (Section 4.3).

The paper is organized as follows. In Section 2 we discuss recent progress in robust multi-model fitting to put our contributions in perspective. We develop the preliminaries in Section 3 and describe our proposed method in Section 4. We show experimental results and comparisons on synthetic and real data in Section 5 and conclude with a discussion of future directions in Section 6.

2 Related Work

A large amount of work has been done in robust model fitting over the last few decades. In this section, we only discuss the work in context of the proposed

method, mainly focusing on consensus based and preference based approaches for discovering *multiple* structures. We also discuss a few recent techniques that address multi-model fitting, albeit they don't strictly fall in these categories.

Consensus based methods: Consensus based methods like RANSAC [4] and its variants work on the premise that valid model hypotheses generate small residuals for inlier points, i.e., inlier points form dense clusters close to zero in the residual space. In [5], the RANSAC method was extended to extract multiple homographies by using a sequential *fit and remove* approach, where the algorithm detects one structure using RANSAC, removes the associated inliers and repeats the process until no more structures are found. This sequential approach has limitations because it can remove inliers erroneously (in case of incorrect scale estimates, or overlapping structures or both), leading to difficulties in recovery of other structures in the following repetitions. The multiRANSAC [18] method avoids the sequential approach by maintaining inlier sets of all structures in parallel. However, it requires the user to input the number of structure along with their respective inlier scales.

Torr et al. [14] proposed MLESAC and MSAC that use a maximum likelihood and M-estimator based criterion respectively. The generalized projection based M-estimator (gpbM) [10] was designed to deal with *heteroscedastic*, i.e., point-dependent noise in the data, which is often encountered in geometric computer vision problems. It estimates the scale automatically and recovers the inlier structure by maximizing the heteroscedastic kernel density estimate in the residual space. Since a single hypothesis is selected based on the maximum density, gpbM is constrained to operate in a sequential *fit-and-remove* manner.

Preference based methods: Contrary to consensus based approaches, preference analysis reverse the roles of data points and model hypotheses. For each data point, residuals are computed with respect to a number of hypotheses. Given a data point, the data preference is defined as the set of hypotheses ordered by their residuals. Intuitively, the data point *prefers* a hypothesis if the corresponding residual is small.

In [19], data preferences are used to estimate the number of modes, followed by model selection and inlier recovery by analyzing histograms of residuals. Toldo and Fusiello [13] pointed out the difficulty of mode detection using residual histogram analysis in [19] and introduced a conceptual space, which represents each data point as a binary vector. A user-specified inlier threshold was used to generate the binary indicator vector identifying the preferred hypotheses. The Jaccard distance was then employed for agglomerative clustering in this conceptual space. In [1, 2], the authors presented a data preferences based Mercer kernel, which following an outlier rejection step, was used for performing Kernel Principal Component Analysis (KPCA) based subspace clustering.

Wang et al [17] introduced the Adaptive Kernel-Scale Weighted Hypotheses (AKSWH) algorithm by combining the Iterative K-th Order Scale Estimator (IKOSE) with a modified version of J-Linkage that clustered hypotheses instead of points. Following a post-processing hypothesis fusion step, the cluster repre-

representatives are chosen as the candidate models, which along with their IKOSE based scale estimates are used to identify the corresponding inliers.

More recently, T-linkage [8], a continuous version of J-linkage based on the Tanimoto distance was proposed. T-linkage also requires the user to adjust a tuning parameter approximating the scale of the inlier noise. Robust Preference Analysis (RPA) [9] used preference analysis along with low-rank approximation and nonnegative matrix facotrization to effectively implement a robust form of spectral clustering. The affinity matrix for RPA is constructed by employing a Cauchy weighting function on the data preference lists. The Random Cluster Model Simulated Annealing (RCMSA)[11] formulated the multi-model fitting problem in a simulated annealing framework. It used data preferences by constructing a weighted graph, which was used to iteratively generate stronger model hypotheses using larger than minimal subsets.

In addition to robust multi-model fitting approaches, preference based methods have also been applied to guided sampling for hypothesis generation, which is an important component of any robust multi-model fitting method. Multi-GS [3], ITKSF [15] and DHF [16, 15] are a few techniques that work well in practice.

3 Notation and Preliminaries

Consider a set of n data points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ in \mathbb{R}^d . Let the set of structures present in the data be denoted by $\mathcal{K} = \{(\boldsymbol{\theta}^{*1}, \mathcal{I}^{*1}), \dots, (\boldsymbol{\theta}^{*k}, \mathcal{I}^{*k})\}$, where $\boldsymbol{\theta}^{*i}$ and \mathcal{I}^{*i} denote the *true* model parameters and the index set of inlier points respectively of the i^{th} structure. We define the *true* inlier fraction of each structure as $\eta^{*i} = \frac{|\mathcal{I}^{*i}|}{n}$, $i = 1, \dots, k$ and $\eta^{*0} = 1 - \sum_{i=1}^k \eta^{*i}$ as the fraction of outliers.

As with most robust multi-model fitting methods, our method begins with a set of model hypotheses, generated by sampling minimal subsets of data. We denote this initial set of model hypotheses by $\vartheta_0 = \{\boldsymbol{\theta}^i | i = 1, \dots, M_0\}$. In order to aid readability and comprehension, we will follow the convention where superscripts will always be indexed over the model hypotheses, while the subscripts will be indexed over the data points. Given a model hypothesis $\boldsymbol{\theta}^i$, the residual for point \mathbf{x}_j is computed using the function $\phi(\boldsymbol{\theta}^i, \mathbf{x}_j) : \mathbb{R}^d \rightarrow \mathbb{R}_+$, where \mathbb{R}_+ denotes the nonnegative real half-line.

Our goal is to recover the correct number of structures k and label the data points in \mathcal{X} as inliers of each of the structures by creating $\mathcal{L} = \{\ell_1, \dots, \ell_n | \ell_i \in \{0, \dots, k\}\}$, where the label 0 identifies the outliers. The inlier points can then be used for parameter estimation using an appropriate estimators such as least squares or M-estimators.

3.1 Residual Density Profile

For a given model hypothesis $\boldsymbol{\theta}^i$, we first compute the residual vector as

$$\mathbf{r}^i = [r_1^i = \phi(\boldsymbol{\theta}^i, \mathbf{x}_1), \dots, r_n^i = \phi(\boldsymbol{\theta}^i, \mathbf{x}_n)] \quad (1)$$

We find a permutation set $\mathbf{u}^i = [u_1^i, u_2^i, \dots, u_n^i]$ for model hypothesis θ^i such that $r_{u_1^i}^i \leq r_{u_2^i}^i \leq \dots \leq r_{u_n^i}^i$ and thus we get a *sorted residual* vector, which is then smoothed using an averaging filter of size $\lceil 0.025n \rceil$. The smoothed ordered residual vector $\boldsymbol{\rho}^i = [r_{u_1^i}^i, r_{u_2^i}^i, \dots, r_{u_n^i}^i]$ is used to define the residual density for θ^i as

$$d_j^i = \frac{j}{\rho_j^i + \varepsilon} = \frac{j}{r_{u_j^i}^i + \varepsilon}, \quad j = 1, \dots, n \quad (2)$$

where ε is of the order of 10^{-4} and is used only to suppress very high densities. The corresponding residual density vector as $\mathbf{d}^i = [d_1^i, d_2^i, \dots, d_n^i]$. This residual density estimate measures the number of points lying in a ball of radius ρ_j^i in the residual space of θ^i .

Using the residual density profile (2) has an advantage over residuals while dealing with outliers. The residual values of outliers are large, however, the magnitude varies with different model hypotheses. However, the residual density values for both gross outliers and pseudo-outliers are small as well as independent of the model hypothesis. This can be seen in the right part of Fig. 1b, where the variance of density values for points in the outlier region is very small compared to that of the inlier region. This property of the residual density profile helps us estimate the inlier scale accurately (Section 4.2).

3.2 Residual based Hypothesis preferences

Similar to [15, 12], we use the residual based hypothesis preferences and define correlation between hypotheses. Given a hypothesis θ^i , its residual based preference list \mathbf{u}^i is the rank ordering of points with smallest residual first. We compute the hypothesis correlation between two hypotheses θ^i and θ^j as

$$h(\theta^i, \theta^j) = \frac{1}{K} |u_{1:K}^i \cap u_{1:K}^j| \quad (3)$$

where $u_{1:K}^i = \{u_1^i, u_2^i, \dots, u_K^i\}$ is the top-K residual based hypothesis preferences for θ^i . We set K to 10% of all the data points in our experiments.

Using (3), a pairwise hypothesis correlation matrix \mathbf{H} can be computed with $\mathbf{H}_{ij} = h(\theta^i, \theta^j)$. However, random sampling based techniques usually lead to a large number of hypotheses. Since the complexity of computing \mathbf{H} is quadratic in the number of hypotheses, it becomes impractical to compute the full $M \times M$ matrix \mathbf{H} . In Section 4.1 we suggest a way of selecting a small subset of promising hypotheses and mitigate this problem of complexity.

3.3 Density based Point Preferences

In general, a point preference list for a given data point is a rank ordering of the model hypotheses based on some criterion. In [3, 15], a smallest residual first criterion was used for ranking hypotheses. Chin et al. [3] made an important empirical observation that inlier points belonging to the same structure have

highly correlated data preferences. As opposed to the traditional residual based point preferences, we use *density* based point preferences, which in turn are used to estimate correlation between data points.

We find a permutation $\mathbf{v}_j = [v_j^1, v_j^2, \dots, v_j^{M_0}]$ for a data point \mathbf{x}_j such that $d_j^{v_j^1} \geq d_j^{v_j^2} \geq \dots \geq d_j^{v_j^{M_0}}$. The permutation \mathbf{v}_j induces a density based point preference list for \mathbf{x}_j , with v_j^1 being its most preferred model hypothesis. We define the $n \times n$ point correlation matrix \mathbf{P} by using the intersection kernel over the top- K density based point preference lists

$$\mathbf{P}_{ij} = \frac{1}{K} |v_i^{1:K} \cap v_j^{1:K}| \quad (4)$$

where $v_i^{1:K} = \{v_i^1, v_i^2, \dots, v_i^K\}$ is the top- K density based preferences for point \mathbf{x}_i and K is a small number of top preferences, set to $\lceil 0.01M_0 \rceil$ in all our experiments. As we use density based preferences, we need a much smaller K than in case of residual based preferences. In Section 4.2, we use a variant of this point correlation matrix \mathbf{P} for identifying potential inliers of a given hypothesis.

4 Proposed Method: Density Preference Analysis

We will now develop our complete algorithm Density Preference Analysis¹ (DPA) for recovering all inlier structures present in the data using the building blocks described in the previous section.

4.1 Candidate Hypotheses Selection

We use the density based point preferences described in Section 3.3 to select promising hypotheses from ϑ_0 . The residual density \mathbf{d}_j^i given by (2) roughly measures the likelihood of \mathbf{x}_j being an inlier for the model hypothesis θ^i . If a data point \mathbf{x}_j has a hypothesis θ^i in their top-5 density based preference list, we say the point votes for θ^i . Since the points voting for θ^i are likely to be inliers, a selected hypothesis better represents the structure if more inliers vote for it. Recall that the set \mathbf{v}_j is a permutation of the indices of ϑ_0 , i.e., $\{1, \dots, M_0\}$. We create an index set by taking the union of hypothesis indices in the top-5 preferences of all points, i.e., $\bigcup_{j=1}^n v_j^{1:5}$. We then reject spurious hypotheses from this set by eliminating the ones having fewer than two votes. We refer to this reduced set of hypothesis indices as $\mathcal{I}_{\vartheta_s}$ and the corresponding set of selected hypotheses as $\vartheta_s \subseteq \vartheta_0$ having cardinality $M_s \leq M_0$. The set ϑ_s contains hypotheses that appear in the top-5 preferences of at least two points and thus is expected to contain all the hypotheses that represent a structure well.

We emphasize that this step applies a conservative rule for rejecting poorly generated hypotheses and is primarily used for a computational advantage of

¹ Matlab code available at <https://www.iiitd.edu.in/~anands/files/code/dpa.zip>

processing fewer hypotheses in the following steps. The goal is to ensure retaining all the good hypotheses while rejecting many of the bad ones. We verified empirically that values between top-2 to top-10 preferences and 2-5 minimum votes do not affect the good hypothesis selection and thus the final accuracy significantly, and only affect the computational time for the next steps.

4.2 Inlier Scale Estimation

The residual density values drop significantly across the inlier-outlier boundary. We make use of this disparity in density to estimate the scale of inlier noise for a given hypothesis.

Hypothesis Refinement: We construct \mathcal{J}^i , $i \in \mathcal{I}_{\vartheta_s}$ as the potential inlier set for θ^i . \mathcal{J}^i is initialized with all the points voting for θ^i . Note that \mathcal{J}^i will always contain indices pointing to the original set \mathcal{X} of data points. We normalize the point correlation matrix \mathbf{P} (4), such that each columns has unit ℓ_1 -norm. We compute the correlation score of all points with the points voting for θ^i as

$$\bar{\mathbf{p}}^i = \prod_{j \in \mathcal{J}^i} \mathbf{P}_{.j} \quad (5)$$

Points with a large \bar{p}_j^i are highly correlated with the voting points in \mathcal{J}^i and are strong contenders for inliers of θ^i . We define a threshold as $\tau^i = \min_{j \in \mathcal{J}^i} \bar{p}_j^i$ and impose the following criterion to include potential inliers in the set \mathcal{J}^i

$$\mathcal{J}^i = \{j \mid \bar{p}_j^i \geq \tau^i, j \in [n]\} \quad , \text{ where } [n] = \{1, \dots, n\}. \quad (6)$$

We create a similar set of potential outliers defined as the set of points that do not appear as strong inliers for any of the selected hypotheses

$$\mathcal{O} = [n] \setminus \left\{ \bigcup_{i=1}^{M_s} \mathcal{J}^i \right\} \quad (7)$$

The set \mathcal{O} contains all the points that do not belong to any of the structures, and thus are gross outliers. However, since \mathcal{J}^i only contain a subset of inliers, \mathcal{O} may contain some inliers as well. We eliminate the hypotheses for which the largest residual of potential outliers is smaller than the largest residual of the strong inliers, i.e., reject θ^i if $\max_{j \in \mathcal{O}} \rho_j^i \leq \max_{j \in \mathcal{J}^i} \rho_j^i$. This is a conservative step to reject poor hypotheses that have unreliable inliers.

We use the strong inlier points in \mathcal{J}^i to obtain a refined estimate of θ^i using least squares and recompute the corresponding residual density profile (Section 3.1). This step has two positive effects: pure hypotheses that are representative of the inlier structure result in better residual density profiles with increased disparity between the densities of inliers and outliers. On the contrary, impure hypotheses generate density profiles with a lower density drop across the inlier-outlier boundary as will be seen in Section 4.3.

Density Drop Rate: Let $\mathcal{J}_{\uparrow 10}^i \subset \mathcal{J}^i$ contain 10% of points in \mathcal{J}^i that have the largest residuals ρ_j^i . We expect these points to lie closer to the inlier-outlier boundary. Using these potential inlier points, we define the *density drop rate* as

$$\Delta d^i = \frac{\frac{1}{|\mathcal{J}_{\uparrow 10}^i|} \sum_{j \in \mathcal{J}_{\uparrow 10}^i} (d_{j_{max}}^i - d_j^i)}{\frac{1}{|\mathcal{J}_{\uparrow 10}^i|} \sum_{j \in \mathcal{J}_{\uparrow 10}^i} (\rho_j^i - \rho_{j_{max}}^i)} \quad (8)$$

where $j_{max} = \arg \max_j d_j^i$ is the sorted residual index corresponding to maximum residual density. We use the density $d_{k_{max}}^i$, where $k_{max} = \arg \max_j \rho_j^i, j \in \mathcal{J}_{\uparrow 10}^i$ and the density drop rate Δd^i in (8) to linearly extrapolate and estimate the scale of inlier noise for θ^i as

$$\hat{\rho}^i = \frac{(d_{k_{max}}^i - d_{\mathcal{O}_{max}}^i)}{\Delta d^i} + \rho_{k_{max}}^i \quad (9)$$

where $d_{\mathcal{O}_{max}}^i = \max_{j \in \mathcal{O}} d_j^i$ is the maximum density of outlier points for θ^i . In the unlikely event when $\mathcal{O} = \{\emptyset\}$, we set $d_{\mathcal{O}_{max}}^i = 0$. The residual index $\hat{j}^i = \arg \min_{j \in [n]} |\rho_j^i - \hat{\rho}^i|$ can be used to compute the inlier fraction as $\hat{\eta}^i = \hat{j}^i/n$. The inlier scale estimation using linear extrapolation of the density works very well in practice as we observed in our experiments.

Using the scale estimates, we update the sets $\mathcal{J}^i, i \in \mathcal{I}_{\vartheta_s}$ to include all inliers

$$\mathcal{J}^i = \{j \mid r_j^i \leq \hat{\rho}^i, j \in [n]\} \quad \forall i \in \mathcal{I}_{\vartheta_s} \quad (10)$$

4.3 Model Selection and Labeling

We now have a tentative point-model hypothesis association in $\mathcal{J}^i, i \in \mathcal{I}_{\vartheta_s}$. From the candidate model hypotheses in $\mathcal{I}_{\vartheta_s}$, we need to identify *exactly* one model hypothesis for each structure. We further use the intuition that selected models should be able to explain *all* the inlier points. We describe a simple process (Algorithm 1) that uses the density drop rate and hypothesis preferences based correlation to identify a model corresponding to each structure, associate their inliers with the respective models and label the remaining points as gross outliers. While the process is greedy, we empirically show that it works well in practice.

We use the set of candidate models $\mathcal{I}_{\vartheta_s}$, the corresponding inlier sets $\mathcal{J}^i, i \in \mathcal{I}_{\vartheta_s}$ and the preference based hypothesis correlation matrix \mathbf{H} . We initialize the set of all inliers \mathcal{J} and the set of final models \mathcal{I}_{ϑ} as empty sets. The output of the algorithm is a set of models, each corresponding to a structure in the data. When a candidate model hypothesis has an associated inlier set \mathcal{J}^{i_t} with a sufficiently small overlap ($\leq \tau_o$) with \mathcal{J} , we refer to it as a novel hypothesis. We use τ_o , which is typically set as a small percentage of points indicating the tolerable overlap between structures. The second threshold τ_h determines the smallest acceptable correlation between hypotheses. Both these thresholds are intuitive, easy to set and gracefully affect the performance accuracy. We identify a model hypothesis $\theta^{i_t} \in \vartheta_s$ with the largest density drop rate (line 7). This

Algorithm 1: Model Selection

```

1 Input: variables -  $\mathcal{I}_{\vartheta_s}$ ,  $\mathcal{J}^i$  ( $i \in \mathcal{I}_{\vartheta_s}$ ),  $\mathbf{H}$ 
2     thresholds -  $\tau_o$ ,  $\tau_h$ 
3 Output:  $\mathcal{I}_{\vartheta}$  - index set of selected models
4 begin
5     Initialize:  $\mathcal{I}_{\vartheta} = \mathcal{J} = \{\emptyset\}$ 
6     repeat
7          $i_t \leftarrow \arg \max_{i \in \mathcal{I}_{\vartheta_s}} \Delta d^i$ 
8         if  $|\mathcal{J} \cap \mathcal{J}^{i_t}| \leq \tau_o$  then
9              $\mathcal{I}_{\vartheta_t} = i_t \cup \{i \mid |\mathcal{J} \cap \mathcal{J}^i| \leq \tau_o, \mathbf{H}_{i_t i} \geq \tau_h, \forall i \in \mathcal{I}_{\vartheta_s}\}$ 
10             $\hat{i}_t \leftarrow \arg \max_{i \in \mathcal{I}_{\vartheta_t}} \hat{\eta}^i$ 
11             $\mathcal{I}_{\vartheta} = \mathcal{I}_{\vartheta} \cup \{\hat{i}_t\}$ 
12             $\mathcal{J} = \mathcal{J} \cup \mathcal{J}^{\hat{i}_t}$ 
13             $\mathcal{I}_{\vartheta_s} = \mathcal{I}_{\vartheta_s} \setminus \mathcal{I}_{\vartheta_t}$ 
14        else
15             $\mathcal{I}_{\vartheta_s} = \mathcal{I}_{\vartheta_s} \setminus \{i_t\}$ 
16    until  $\mathcal{I}_{\vartheta_s} = \{\emptyset\}$ 
17 return  $\mathcal{I}_{\vartheta}$ 

```

aligns well with our intuition that a sharp density drop provides the strongest evidence of an inlier structure. Furthermore, if θ^{i_t} is novel, we construct an index set $\mathcal{I}_{\vartheta_t}$ of model hypotheses that are novel *and* highly correlated with θ^{i_t} (line 9), else we ignore it by eliminating the corresponding entry from $\mathcal{I}_{\vartheta_s}$ (line 15). Candidate models that are not novel are likely to arise from hypotheses that merge two distinct structures. From the subset $\mathcal{I}_{\vartheta_t}$, we select the model hypothesis with the largest fraction of inliers and use it as the representative for the corresponding structure (lines 10-11). We update the set of all inliers \mathcal{J} and the set of candidate models $\mathcal{I}_{\vartheta_s}$ as in (lines 12-13). We repeat the process until all candidate models have been explored and $\mathcal{I}_{\vartheta_s}$ is empty and return the final set of models as \mathcal{I}_{ϑ} . We post-process the set of models, \mathcal{I}_{ϑ} , obtained from Algorithm 1 by ignoring the models that have a very small fraction of inliers $\frac{|\mathcal{J}^i|}{n} \leq 0.05$. Now the number of structures recovered is the cardinality of \mathcal{I}_{ϑ} . Finally, due to the overlap tolerance τ_o , we have a small number of points that may be in multiple inlier sets. We reassign such points to models for which the residual density is the highest. This post-processing step results in a unique point-model assignment and all the unassigned points are labeled as outliers. Using the unique point-model assignment in \mathcal{J}^i , $i \in \mathcal{I}_{\vartheta}$, we can generate the label set \mathcal{L} . Our overall algorithm is summarized in Algorithm 2.

5 Experimental Results

In this section we evaluate our proposed approach, DPA, and compare with three recent robust multi-model fitting approaches RPA [9], T-Linkage [8] and RCMSA

Algorithm 2: Density Preference Analysis(DPA)

```

1 Input: Data -  $\mathcal{X}$  , Model type -{synthetic, homography, fundamental}
2 Output:  $\mathcal{L}$  - Labeled data points
3 begin
4    $\vartheta_0 \leftarrow$  generate model hypotheses using DHF [15].
5    $\mathbf{d}^i \leftarrow$  for each  $\boldsymbol{\theta}^i$  compute residual density vector using (2).
6    $\mathbf{P} \leftarrow$  compute density based point correlation using (4).
7    $\vartheta_s \leftarrow$  select candidate hypotheses (Section 4.1).
8    $\mathcal{J}^i \leftarrow$  for each  $\boldsymbol{\theta}^i \in \vartheta_s$  identify strong inliers using (6).
9    $\mathcal{O} \leftarrow$  identify potential outliers using (7).
10   $\hat{\eta}^i \leftarrow$  for each  $\boldsymbol{\theta}^i \in \vartheta_s$  estimate inlier scale and update  $\mathcal{J}^i$  using (8-10)
11   $\{\mathcal{I}_\vartheta\} \leftarrow$  model selection (Algorithm 1).
12   $\mathcal{L} \leftarrow$  post-processing and point-model assignments (Section 4.3)
13 return  $\mathcal{L}$ 

```

[11] using three different datasets. We used publicly available implementations provided by the respective authors and set the parameters as suggested in the corresponding publications.

We report the classification accuracy (CA) averaged over 10 runs on all our experiments. For our synthetic dataset, we use points along three concentric circles, with each structure corrupted with a different scale of noise. For the real experiments, we use the AdeleideRMF [16] data set, which consists of 19 image pairs each for homography based planar segmentation and fundamental matrix based motion segmentation. Each image pair is provided with SIFT[6] point matches corrupted with outliers and their ground truth labeling.

For DPA, we set the hypothesis correlation threshold τ_h to 0.6 for planar segmentation and 0.75 in case of motion segmentation and the synthetic experiments. We fixed the value of the overlap threshold $\tau_o = \lceil 0.025n \rceil$ for planar segmentation and concentric circles. It was set to $(\lceil 0.1n \rceil)$ for fundamental matrix based motion segmentation. A higher overlap threshold is justified as the epipolar constraint imposed by the fundamental matrix on each rigid object is a weaker constraint than the one from homography. We use DHF [15] for hypothesis generation for all experiments.

We emphasize that contrary to the competing methods, our method, DPA, does not require an estimate of the inlier noise scale. We follow [9] and provide the standard deviation of *all* inlier residuals to each of the competing methods (T-Linkage, RPA, and RCMSA). The parameters used are $\sigma_n = 0.0098$ for synthetic data, $\sigma_n = 0.013$ for planar segmentation and $\sigma_n = 0.005$ for motion segmentation. As we will show in the synthetic experiments, this dependency on the scale of noise is detrimental when different structures have significantly different noise. This may also be the reason of relatively poor performance of the competing methods on the planar segmentation experiments.

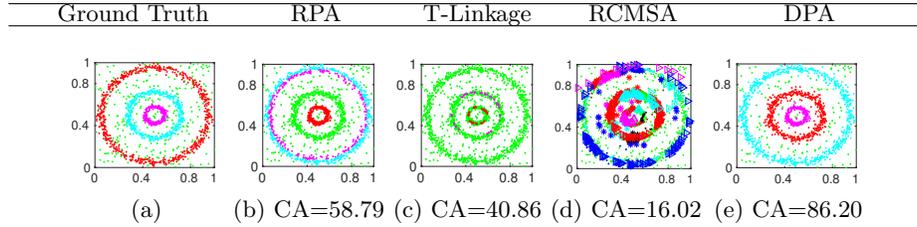


Fig. 2: **Three Concentric Circles**: (CA=Classification Accuracy(%)). Point membership is color coded. Outliers in green.

5.1 Synthetic Example: Three Concentric Circles

In our synthetic experiment, we generate three concentric circles centered at $[0.5, 0.5]$ (radius - 0.08, 0.22, 0.45), each having a different number of inlier points (250, 450, 650) corrupted with Gaussian noise of different standard deviation (0.01, 0.015, 0.018). We added uniformly distributed 300 gross outliers in the square region defined by (0,0) and (1,1). The data and the segmentation results of all the methods are shown in Fig. 2. This seemingly simple experiment resulted in poor performances by each of the methods RPA, T-Linkage and RCMSA, with the average CA as 58.79%, 40.86% and 16.02%. DPA managed to achieve an average CA of 86.20%. In the sample result shown in Fig. 2 both RPA and T-Linkage fail to find one circle, while oversegment the others. RCMSA results in a large oversegmentation and recovers 45 distinct structures, while RPA and T-Linkage recover two structures on average. DPA is able to recover all the three structures in each of the 10 runs. The poor performance of the competing methods is seemingly due to the variation in noise scale of different structures.

5.2 Planar Segmentation

We perform planar segmentation by fitting multiple homographies to the point matches. We use isotropic scaling of the point matches such that they have zero mean and have unit average distance from the origin. We use Sampson error as our residual function $\phi(\theta, \mathbf{x})$. We report the quantitative results in Table 1 and see the average and median performance of DPA is superior than the other methods. We also show sample results on two examples in Fig. 3, where the other three competing methods perform poorly. The scale of noise in the sample sequences, *library* and *oldclassicwing*, are too high and too low respectively as compared to the input parameter σ_n [9]. The superior performance of DPA in these cases is largely due to its residual density based scale estimation.

5.3 Motion Segmentation

We perform motion segmentation on the AdelaideRMF images by fitting multiple fundamental matrices to the point matches. As preprocessing, we use isotropic

Table 1: Planar Segmentation. **Outlier%**- Percentage of outliers, **#S** - Number of true structures, Evaluation metric - Classification Accuracy (CA%) and corresponding running time (T) in seconds.

	Outlier%	#S	RPA	T	T-Link	T	RCMSA	T	DPA	T
barrsmith	68.87	2	62.90	201.67	57.93	62.71	84.81	2.88	97.68	35.79
bonhall	06.17	6	52.88	3712.26	60.41	1447.47	81.67	9.25	77.98	52.75
bonython	73.73	1	84.34	105.91	64.34	48.28	87.27	3.71	96.57	30.06
elderhalla	60.74	2	99.07	176.65	69.53	57.67	75.23	3.13	96.17	37.58
elderhallb	47.84	3	81.96	269.08	57.80	83.63	71.45	3.08	85.90	39.02
hartley	61.56	2	81.38	290.07	71.63	136.58	77.38	3.75	96.91	43.45
johnsona	20.91	4	91.10	535.49	57.80	179.43	83.00	3.94	87.08	63.94
johnsonb	12.01	7	66.84	1676.11	70.72	540.97	79.41	6.18	74.36	75.83
ladysymon	32.48	2	79.16	185.15	77.72	62.01	75.27	2.78	90.46	8.33
library	55.34	2	63.53	163.32	82.51	50.80	77.02	2.89	95.21	6.08
napiera	62.91	2	73.25	269.69	81.26	102.22	70.66	3.39	80.56	3.36
napierb	39.51	3	75.14	271.21	76.76	73.62	74.29	2.96	83.63	11.05
neem	36.51	3	78.51	247.19	53.03	63.91	71.87	3.23	80.21	7.79
nese	33.46	2	99.21	204.25	53.70	70.53	77.56	2.96	97.40	16.25
oldclassswing	32.45	2	76.73	369.72	73.77	262.34	92.45	3.59	96.33	163.14
physics	45.28	1	100.0	40.12	68.49	26.54	54.53	15.85	98.40	11.84
sene	47.20	2	99.44	200.35	84.32	81.33	71.68	5.88	99.76	7.76
unihouse	16.55	5	88.00	9232.87	71.86	5890.28	97.02	16.37	93.17	85.83
unionhouse	76.50	1	76.14	226.33	77.29	116.83	90.06	5.70	98.34	16.22
mean			80.50	967.23	68.99	492.48	78.55	5.31	90.90	37.68
std			13.51	2174.70	9.92	1347.1	9.54	4.14	8.10	39.19
median			79.16	247.19	70.72	81.33	77.38	3.39	95.21	30.06

Table 2: Motion Segmentation. Notations are same as in table 1

	Outlier%	#S	RPA	T	T-Link	T	RCMSA	T	DPA	T
biscuit	57.16	1	98.36	38.75	83.09	19.17	95.15	4.50	82.12	14.42
biscuitbook	47.51	2	96.42	49.42	97.77	20.37	92.52	5.21	97.24	123.66
biscuitbookbox	37.21	3	95.83	45.70	88.80	11.69	83.71	3.36	95.14	46.78
boardgame	42.48	1	87.53	46.13	83.73	13.53	78.46	3.79	83.69	29.52
book	21.48	1	97.54	16.27	82.57	6.31	94.01	12.62	90.16	94.86
breadcartoychips	35.20	4	91.73	45.31	80.51	9.75	78.82	3.24	91.56	15.53
breadcube	32.19	2	95.95	29.98	85.62	10.33	87.27	3.61	94.09	67.46
breadcubechips	35.22	3	95.57	36.91	82.00	9.66	83.17	3.38	94.61	24.64
breadtoy	37.41	2	97.15	44.39	96.81	14.61	78.37	3.77	90.59	15.45
breadtoycar	34.15	3	92.17	30.47	84.70	4.97	83.07	2.60	88.67	15.60
carchipscube	36.59	3	94.30	27.69	88.00	4.89	78.85	2.70	86.30	14.13
cube	69.49	1	97.15	33.65	46.29	15.86	87.98	5.79	96.89	9.81
cubebreadtoychips	28.03	4	93.21	67.63	80.18	19.12	81.62	3.83	87.28	43.75
cubechips	51.62	2	96.48	39.59	95.14	13.99	90.32	5.82	92.92	75.64
cubetoy	41.42	2	96.31	31.60	78.80	10.74	89.64	5.45	93.61	86.55
dinobooks	44.54	3	84.78	64.21	78.56	22.90	72.28	5.76	84.17	73.19
game	73.48	1	95.97	21.93	77.6	9.32	90.77	4.88	97.47	83.83
gamebiscuit	51.54	2	96.95	49.04	70.61	19.05	85.40	4.55	90.95	74.25
toycubecar	36.36	3	91.70	27.21	70.70	7.17	83.45	2.94	84.65	47.46
mean			94.47	39.25	81.65	12.81	84.99	4.62	90.64	50.34
std			3.55	13.21	11.36	5.44	6.16	2.20	4.86	34.03
median			95.95	38.75	82.57	11.69	83.71	3.83	90.95	46.78

scaling of the point matches such that they have zero mean and have unit average distance from the origin. We use the epipolar constraint to compute the residuals. We report the quantitative results in Table 2 and see the average and median performance of DPA is inferior only to RPA. We also show sample results on two examples in Fig. 3. We emphasize that despite not knowing the scale of inlier noise, DPA performs competitively resulting in average CA of excess of 90%.

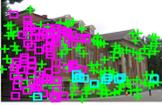
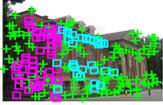
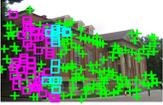
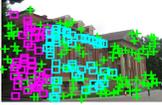
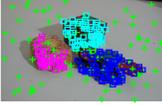
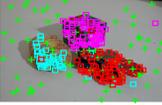
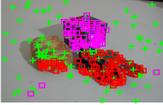
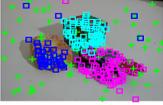
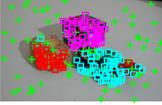
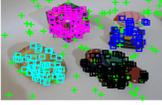
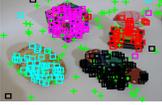
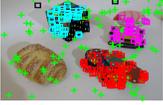
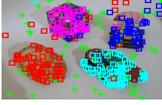
Ground Truth	RPA	T-Linkage	RCMSA	DPA
				
(a)	(b) CA=76.28	(c) CA=83.67	(d) CA=82.33	(e) CA=97.67
				
(f)	(g) CA=81.00	(h) CA=76.20	(i) CA=95.68	(j) CA=96.57
				
(k)	(l) CA=96.52	(m) CA=83.48	(n) CA=90.04	(o) CA=96.09
				
(p)	(q) CA=94.80	(r) CA=89.30	(s) CA=91.04	(t) CA=91.13

Fig. 3: [Sample results] **Planar Segmentation**: Row 1 - *library*, Row 2- *oldclasswing*. **Motion Segmentation**: Row 3- *breadcubechips*, Row 4- *cubebreadtoychips*. (CA=Classification Accuracy(%)). Point membership is color coded. Outliers in green.

6 Conclusion

We proposed a novel algorithm for the problem of robust multi-model fitting that leverages both consensus and preference analysis. To characterize the consensus based approach, we defined a residual density profile, which we further used for computing point preferences. We further leveraged the disparity between the density of potential inliers and outliers to estimate the scale of inlier noise for each model hypothesis. We devised a greedy scheme that uses preference based hypothesis similarity to identify a model hypothesis for representing each structure. Finally, we used a simple post-processing step to eliminate small, spurious structures and to generate a unique point to model assignment.

We showed empirical results using synthetic examples as well as the AdelaideRMF datasets. Without using any information about the scale of the noise, we showed competitive performance with other state-of-the-art multi-model fitting methods that primarily rely on preference analysis. Through our empirical analysis, we have shown that the residual density information judiciously combined with preference analysis can improve the performance in robust model fitting problems while reducing dependence on a priori knowledge about noise in data.

References

1. Chin, T.J., Wang, H., Suter, D.: The Ordered Residual Kernel for Robust Motion Subspace clustering. *NIPS* (2009) 333–341
2. Chin, T.J., Wang, H., Suter, D.: Robust fitting of multiple structures: The statistical learning approach. *ICCV* (2009) 413–420
3. Chin, T.J., Yu, J., Suter, D.: Accelerated Hypothesis Generation for Multistructure Data via Preference Analysis. *IEEE TPAMI* **34** (2012) 625–638
4. Fischler, M.A., Bolles, R.C.: Random Sample Consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24** (1981) 381–395
5. Kanazawa, Y., Kawakami, H.: Detection of planar regions with uncalibrated stereo using distribution of feature points. *BMVC* (2004) 247–256
6. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* **60** (2004) 91–110
7. Lavva, I., Hameiri, E., Shimshoni, I.: Robust methods for geometric primitive recovery and estimation from range images. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **38** (2008) 826–845
8. Magri, L., Fusiello, A.: T-Linkage: A Continuous Relaxation of J-Linkage for Multi-Model Fitting. *CVPR* (2014) 3954–3961
9. Magri, L., Fusiello, A.: Robust Multiple Model Fitting with Preference Analysis and Low-rank Approximation. *BMVC* (2015) 20.1–20.12
10. Mittal, S., Anand, S., Meer, P.: Generalized Projection-Based M-Estimator. *IEEE TPAMI* **34** (2012) 2351–2364
11. Pham, T.T., Chin, T.J., Yu, J., Suter, D.: The Random Cluster Model for Robust Geometric Fitting. *IEEE TPAMI* **36** (2014) 1658–1671
12. Raguram, R., Frahm, J.M.: RECON: Scale-adaptive robust estimation via Residual Consensus. *ICCV* (2011) 1299–1306
13. Toldo, R., Fusiello, A.: Robust multiple structures estimation with J-linkage. *ECCV* (2008) 537–547
14. Torr, P.H.S., Zisserman, A.: MLESAC: A new robust estimator with application to estimating image geometry. *CVIU* **78** (2000) 138–156
15. Wong, H.S., Chin, T.J., Yu, J., Suter, D.: A Simultaneous Sample-and-filter Strategy for Robust Multi structure Model Fitting. *CVIU* **117** (2013) 1755–1769
16. Wong, H.S., Chin, T.J., Yu, J., Suter, D.: Dynamic and Hierarchical Multi-Structure Geometric Model Fitting. *ICCV* (2011) 1044–1051
17. Wang, H., Chin, T.J., Suter, D.: Simultaneously Fitting and Segmenting Multiple-Structure Data with Outliers. *IEEE TPAMI* **34** (2012) 1177–1192
18. Zuliani, M., Kenny, C.S., Manjunath, B.S.: The multiRANSAC algorithm and its application to detect planar homographies. *IEEE ICIP* (2005) 153–156
19. Zhang, K., Kwok, J.T., Tang, M.: Accelerated Convergence using Dynamic Mean Shift. *ECCV* **2** (2006) 257–268