FAST HYPOTHESIS FILTERING FOR MULTI-STRUCTURE GEOMETRIC MODEL FITTING

Lokender Tiwari, Saket Anand

IIIT-Delhi
New Delhi, India

ABSTRACT

We propose a fast and efficient two-stage hypothesis filtering technique that can improve performance of clustering based robust multi-model fitting algorithms. Sampling based hypothesis generation is nondeterministic and permits little control over generating poor model hypotheses, often leading to a significant proportion of bad hypotheses. Our novel filtering approach leverages the asymmetry in the distributions of points around the inlier/outlier boundary via the sample skewness computed in the residual space. The output is a set of promising hypotheses which aid multi-model fitting algorithms in improving accuracy as well as running time. We validate our approach on the AdelaideRMF dataset and show favorable results along with comparisons to state-of-the-art.

Index Terms— Hypothesis filtering, Multi-model fitting.

1. INTRODUCTION

Estimating geometric entities (lines, planes, homographies, fundamental matrix etc.) is an integral part of many computer vision applications like 3D reconstruction and object detection and recognition. Following the sampling based hypothesis generation strategy of Random Sample Consensus (RANSAC) [1], several multi-model fitting methods have been proposed like multiRANSAC [2], generalized projection based M-estimator (gpbM) [3], J-linkage [4] and T-linkage [5]. These methods randomly sample a *minimal subset* of points (2 for lines, 3 for planes, etc.), which are either selected

befitting a predefined criterion [2, 1, 3] or are used to cluster the data points [4, 5]. Due to the presence of *outliers* and *pseudo-outliers* in the data, it is crucial to generate *pure* minimal sample subsets, i.e., points arising from the same inlier structure. Performance of clustering based methods like [4, 5] depends significantly on the quality and number of pure hypotheses generated by the random sampling approach. For these methods to work in practical scenarios, an intermediate step capable of efficiently rejecting spurious, irrelevant hypotheses is required, thus necessitating the requirement of a hypothesis filtering algorithm.

Many approaches that aim at generating more promising hypotheses rely on guided or *biased* sampling techniques, primarily driven by using some a priori information available about the data points. Approaches like [7, 8, 9] rely on corresponding matching scores (e.g SIFT [10] matching score) to bias the sampling process. Multi-GS [11] does not assume any a priori information and creates a conditional distribution based on hypothesis preferences (defined in terms of ordered residuals) and is used for sampling minimal subsets, however there is still little control over generating impure minimal subsets.

SCRAMSAC [12] on the other hand employs a spatial consistency filter on the sampled minimal subsets. AKSWH [13] assigns weights to each hypothesis based on the kernel density estimate of its residuals, followed by data driven thresholding for hypothesis rejection. In Dynamic and Hierarchical Filtering (DHF) [14], each data point is associated with the most promising hypothesis based on
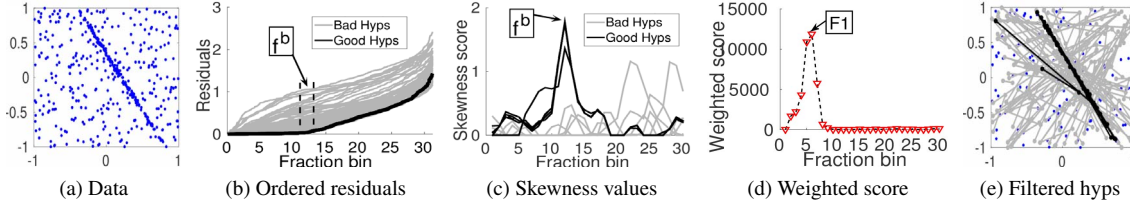
| (a) Data | (b) Ordered residuals | (c) Skewness values | (d) Weighted score | (e) Filtered hyps |

Fig. 1: Fast Hypothesis Filtering (FHF) illustrative example.

its preferences and the hypotheses that are not associated with a data point are filtered. In simultaneous sampling and filtering approaches like DHF [15, 14] or ITKSF [14] that rely solely on preference analysis, it is not straightforward to incorporate prior information so readily taken by guided sampling methods like [7, 8, 9].

Our hypothesis filtering approach works on the principle that a pure model hypothesis results in a sharp drop in density at the inlier/outlier boundary for that structure. This results in an asymmetric distribution of residuals, which can be measured by the sample skewness computed in the local neighborhood of the inlier/outlier boundary. This intuition is illustrated using a synthetic line fitting example in Fig. 1. The subfigures show the various steps in our algorithm. The dark and grey curves indicate the response of *good* and irrelevant hypotheses respectively. We do not make any assumptions on the distribution of the inlier or the outliers.

The remainder of the paper is organized as follows. Section 2 describes our approach in detail. We show empirical results in Section 3 and conclude in Section 4.

2. PROPOSED APPROACH

Let $\mathcal{X} = \{x_j\}_{j=1}^N$ be a dataset of N data points arising from multiple structures and contaminated with noise and outliers. We generate a set of model hypotheses $\vartheta = \{\theta_i\}_{i=1}^M$ using minimal subsets of $\mathcal{X}$ obtained by a random or guided sampling process. We call a model hypothesis θ_i as *pure* if the corresponding minimal subset is comprised of only inlier points, all belonging to the same structure. The hypotheses generated with mixed minimal subsets comprising points from different structures or outliers are referred to as bad or irrelevant.

The goal of hypothesis filtering is to signifi-

cantly increase the proportion of *pure* hypotheses by rejecting irrelevant hypothesis. Our approach, dubbed Fast Hypothesis Filtering (FHF), is a two-stage process, where the first stage uses a simple and efficient *skewness*-based statistic to filter out a large number of irrelevant hypotheses. The second stage adopts a preference analysis based strategy to select hypotheses that are more likely to be pure.

2.1. Stage-1: Skewness Based Filtering

For each hypothesis $\theta_i \in \vartheta$, we compute the absolute residuals $r_i^j = \phi(\theta_i, x_j)$ for all the N data points and sort them in non-decreasing order. The sorted residual vectors for each hypothesis are stacked along the rows to construct the $M \times N$ matrix $\mathbf{R}$ of *ordered residuals*.

We measure the asymmetry around the inlier/outlier boundary by computing the sample skewness over $\mathbf{R}$. We take a sliding window along each row $\mathbf{R}_{i\cdot}$ of size $\lceil wN \rceil$ points, where $0 < w < 1$ is the fraction of points in a window. The step size for the sliding window is given by $\lceil tN \rceil$, where $\frac{1}{N} < t < w$. We obtain Q overlapping *fraction bins* for the i^{th} row $\mathbf{R}_{i\cdot}$, each denoted by the index set f_i^q, $q = 1, \ldots, Q$ of the constituent $\lceil wN \rceil$ points.

We compute the sample skewness over each bin f_i^q, $i = 1, \ldots, M$, $q = 1, \ldots, Q$ and construct the $M \times Q$ matrix $\mathbf{S}$ with entries as

$$\mathbf{S}_{iq} = \frac{1}{\lceil wN \rceil} \sum_{j \in f_i^q} \left(\frac{\phi(\theta_i, x_j) - \mu}{\sigma} \right)^3 \quad (1)$$

where μ and σ are the sample mean and the sample standard deviation of residuals corresponding to points in f_i^q, respectively. Since the fraction bin containing the inlier/outlier boundary is expected to have a high positive skewness score, for each model

hypothesis, we retain only the *top three* skewness scores and set the rest to zero. This yields an $M \times Q$ matrix $\overline{\mathbf{S}}$, with each row containing only three nonzero values, which are the candidate bins containing the inlier/outlier boundary.

We expect higher skewness values for pure hypotheses and lower values for irrelevant hypotheses. We first determine a threshold τ on the skewness scores, which is computed as $\tau = \min_q \max_i \overline{\mathbf{S}}_{iq}$, where $i = 1, \ldots, M$ and $q = 1, \ldots, Q$. The threshold τ serves as the skewness value corresponding to the lowest degree of asymmetry in the distribution of residuals at the inlier/outlier boundary for a pure hypothesis. Thus we update our skewness matrix as

$$\overline{\mathbf{S}}_{iq} = \begin{cases} \overline{\mathbf{S}}_{iq}, & \overline{\mathbf{S}}_{iq} \geq \tau \\ 0, & \text{otherwise} \end{cases}$$

The nonzero entries in the rows $\overline{\mathbf{S}}_{i\cdot}$ are the candidate fraction bins for *pure* and bad hypotheses alike. However, we expect the pure hypotheses arising from the same structure to have candidate fraction bins correlated in terms of location and skewness value. On the other hand, bad hypotheses will typically have uncorrelated skewness values. We leverage this intuition by considering a weighted sum of the skewness scores.

Algorithm 1: Preference based pruning

1   **Input:**$\{\mathcal{I}, \vartheta_{\mathrm{I}}, k, \beta, \alpha\}$
2   **Output:**$\{\vartheta_{\mathrm{II}}\}$ (set of hypotheses after Stage-II)
3   **begin**
4   **for** $j \leftarrow 1$ *to* $|\vartheta_I|$ **do**
5       $\vartheta_{temp} \leftarrow \{\emptyset\}$, $c \leftarrow 0$
6   **for** $l \leftarrow 1$ *to* $|\vartheta_I|; l \neq j$ **do**
7         $olap = \left(\left| \mathcal{I}_k^{\theta_j^{\mathrm{I}}} \cap \mathcal{I}_k^{\theta_l^{\mathrm{I}}} \right| \right) / k$
8   **if** $olap \geq \beta$ **then**
9           $c \leftarrow c + 1$
10           $\vartheta_{temp} \leftarrow \{\vartheta_{temp} \cup \theta_l^{\mathrm{I}}\}$
11   **if** $c \geq \alpha$ **then**
12         $\vartheta_{\mathrm{II}} \leftarrow \{\vartheta_{\mathrm{II}} \cup \vartheta_{temp} \cup \theta_j^{\mathrm{I}}\}$
13         $\vartheta_{\mathrm{I}} \leftarrow \{\vartheta_{\mathrm{I}} \setminus \vartheta_{\mathrm{II}}\}$
14   **return** ϑ_{II}

Let $\mathcal{H}_q, q = 1, \ldots, Q$ be the index set of hypotheses at the q^{th} fraction bin that yield a skewness value at least τ. We compute the weighted score as

$$\mathbf{c}_q = |\mathcal{H}_q| \sum_{i \in \mathcal{H}_q} \overline{\mathbf{S}}_{iq} \qquad (2)$$

and chose $\widehat{q} = \arg\max_q \mathbf{c}_q$ as the largest fraction bin index for an inlier structure. A large structure (with high inlier fraction) will generate a large number of pure hypotheses, thus leading to positive skewness scores at the inlier/outlier boundary. The multiplicative factor $|\mathcal{H}_q|$ biases $\mathbf{c}_q$ to peak at the largest bin index. Empirical evidence in our experiments support this intuition.

We retain all hypotheses from ϑ that yield a large positive skewness value ($\geq \tau$) at a fraction smaller than $\widehat{q}$. Our filtered hypothesis set after Stage-I is

$$\vartheta_{\mathrm{I}} = \{\theta_i | \ \overline{\mathbf{S}}_{iq} \geq \tau, \ \theta_i \in \vartheta, q = 1, \ldots, \widehat{q}\} \qquad (3)$$

2.2. Stage-2: Preference Based Pruning

The filtered set of hypotheses $\vartheta_{\mathrm{I}} = \{\theta_j^{\mathrm{I}}\}$, $j = 1, \ldots, |\vartheta_{\mathrm{I}}|$, may still contain bad hypotheses. However, this set is considerably smaller than ϑ and we can use preference analysis efficiently to eliminate irrelevant hypotheses.

The preference based pruning is summarized in Algrotihm 1. We create a set $\mathcal{X}^{\vartheta_{\mathrm{I}}} \subset \mathcal{X}$, comprising all minimal subsets corresponding to hypotheses in ϑ_{I}. We expect many of the points in $\mathcal{X}^{\vartheta_{\mathrm{I}}}$ to be inliers as they are minimal subsets which generated several pure hypotheses.

Let $\mathcal{I}_k^{\theta_j^{\mathrm{I}}}$ be the index set of points in $\mathcal{X}^{\vartheta_{\mathrm{I}}}$ with the smallest k residuals with respect to $\theta_j^{\mathrm{I}} \in \vartheta_{\mathrm{I}}$. For each hypothesis in ϑ_{I}, we find other hypotheses that are similar, i.e., ones for which the corresponding k-preference sets overlap significantly ($\geq \beta$, lines 7-8). If there are sufficient number of similar hypotheses ($\geq \alpha$), we retain them as potential pure hypotheses in ϑ_{II} (lines 11-12). The retained hypotheses are removed from the input set ϑ_{I} (line 13), and the process is repeated until all hypotheses in ϑ_{I} have been explored.

3. EXPERIMENTAL RESULTS

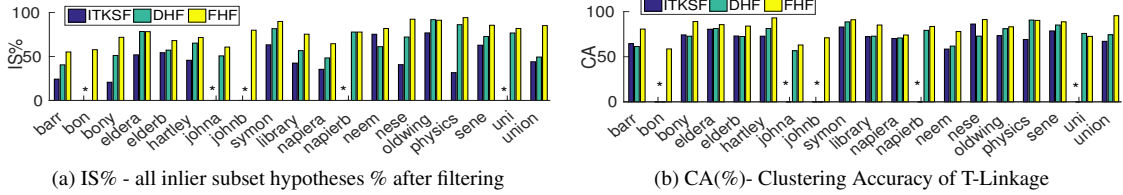We evaluate our algorithm on the AdelaideRMF[15] data set, which contains two sets of 19 image pairs

(a) IS% - all inlier subset hypotheses % after filtering

(b) CA(%)- Clustering Accuracy of T-Linkage

Fig. 2: **Multiple homography fitting.**

(a) IS% - all inlier subset hypotheses % after filtering
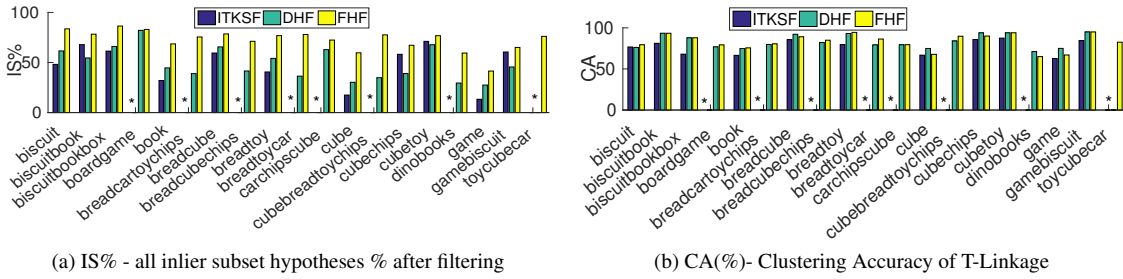
(b) CA(%)- Clustering Accuracy of T-Linkage

Fig. 3: **Multiple fundamental matrix fitting.**

for multiple homography and multiple fundamental matrices each. Each image pair contains SIFT [10] point matches and uniformly distributed outliers. We ran all experiments using MATLAB on a 2.00 GHz Core i7 processor with 4GB RAM *without* any parallel processing. We fix the parameters through all our experiments for FHF. Stage-I: $w = 0.15$ and $t = 0.025$; Stage-II: $\beta = 0.85, k = 0.2 \left| \mathcal{X}^{\vartheta_1} \right|$ and $\alpha = 3$.

We compared our algorithm with two state-of-art hypothesis filtering methods: DHF [14] and ITKSF [14]. Both competing methods use a preallocated time budget for hypothesis generation and filtering. We use multi-GS [11] to generate the set ϑ as input to FHF. Let T be the total time by multi-GS and FHF combined to process an image pair, then we allocate T as the time budget for ITKSF and DHF both. We report results averaged over 10 independent runs in Fig. 2 and Fig. 3 for multiple homographies and multiple fundamental matrices respectively. We say the filtering approach failed if it did not retain any pure hypotheses for at least one structure. The asterisk (*) indicates the algorithms failed in at least three runs out of 10.

The percentage of pure hypotheses (IS%) retained after filtering is shown in Fig. 2a and Fig. 3a. In most cases, we see that FHF outperforms the other two competing methods for the same time

budget. We also measured the clustering accuracy achieved by T-Linkage [5] when the input was the set of filtered hypotheses from the competing algorithms. For multiple homographies, as seen in Fig. 2b, T-Linkage usually performs better when FHF is used, while in case of multiple fundamental matrices (Fig. 3b) the results for FHF are competitive with the other two methods. Fig. 4 shows sample results from the AdelaideRMF dataset. Due to space limitations, we report the detailed results of these experiments in the supplementary material.
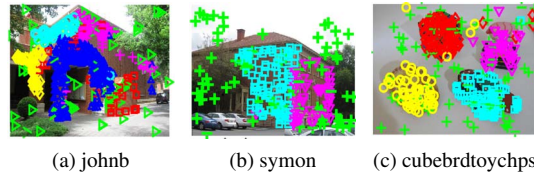

(a) johnb   (b) symon   (c) cubebrdtoychps

Fig. 4: Sample results of T-Linkage with FHF.

4. CONCLUSION

We presented a two-stage hypothesis filtering approach (FHF) that filters out bad hypotheses and improves the performance of clustering based multiple model fitting methods. The algorithm relies on local skewness scores in residual space and leverages preference analysis to eliminate poor hypotheses. We showed competitive performance with state of the art hypothesis filtering techniques on the AdelaideRMF dataset.

5. REFERENCES

[1] M.A. Fischler and R.C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[2] M. Zuliani, C.S. Kenney, and B.S. Manjunath, "The multiRANSAC algorithm and its application to detect planar homographies.," in *ICIP*, 2005, pp. 153–156.

[3] S. Mittal, S. Anand, and P. Meer, "Generalized projection-based M-estimator," *PAMI*, vol. 34, no. 12, pp. 2351–2364, 2012.

[4] R. Toldo and A. Fusiello, "Robust multiple structures estimation with J-Linkage," in *ECCV*, 2008, pp. 537–547.

[5] L. Magri and A. Fusiello, "T-Linkage: a continuous relaxation of J-Linkage for multi-model fitting," in *CVPR*, 2014, pp. 3954–3961.

[6] R. Raguram and J.M. Frahm, "RECON: Scale-adaptive robust estimation via residual consensus," in *ICCV*, 2011, pp. 1299–1306.

[7] B.J. Tordoff and D.W. Murray, "Guided-MLESAC: Faster image transform estimation by using matching priors," *PAMI*, vol. 27, no. 10, pp. 1523–1535, 2005.

[8] O. Chum and J. Matas, "Matching with PROSAC-progressive sample consensus," in *CVPR*, 2005, vol. 1, pp. 220–226.

[9] V. Fragoso, P. Sen, S. Rodriguez, and M. Turk, "EVSAC: accelerating hypotheses generation by modeling matching scores with extreme value theory," in *CVPR*, 2013, pp. 2472–2479.

[10] D.G. Lowe, "Distinctive image features from scale-invariant keypoints," *IJCV*, vol. 60, no. 2, pp. 91–110, 2004.

[11] T.J. Chin, J. Yu, and D. Suter, "Accelerated hypothesis generation for multistructure data via preference analysis," *PAMI*, vol. 34, no. 4, pp. 625–638, 2012.

[12] T. Sattler, B. Leibe, and L. Kobbelt, "SCRAM-SAC: Improving ransac's efficiency with a spatial consistency filter," in *ICCV*, 2009, pp. 2090–2097.

[13] H. Wang, T.J. Chin, and D. Suter, "Simultaneously fitting and segmenting multiple-structure data with outliers," *PAMI*, vol. 34, no. 6, pp. 1177–1192, 2012.

[14] H.S. Wong, T.J. Chin, J. Yu, and D. Suter, "A simultaneous sample-and-filter strategy for robust multi-structure model fitting," *CVIU*, vol. 117, no. 12, pp. 1755–1769, 2013.

[15] H.S Wong, T.J Chin, J. Yu, and D. Suter, "Dynamic and hierarchical multi-structure geometric model fitting," in *ICCV*, 2011, pp. 1044–1051.