Outlier-Robust Multi-Model Fitting on Quantum Annealers

Supplementary Material

This document provides additional details on our Robust Quantum Model Fitting algorithm and its decomposed version (Sec. 7) that could not be included in the main manuscript due to space limitations. It also provides further details and visualizations for the experiments (Sec. 8) presented in the main paper.

7. Algorithmic Details

Decomposed R-QuMF. Real-world size problems are currently intractable on a modern AQC, since the amount of physical qubits required to map logical qubits increases super-linearly. Our decomposed approach, following Farina *et al.* [15], mitigates this issue by decomposing the preference-consensus matrix P (*i.e.*, consensus set of sampled models) into manageable sub-matrices with at most s columns (*i.e.* sub-problem size) that can be confidently sampled on modern quantum hardware using RQuMF. Alg. 1 summarises the decomposed version of our approach that we call De-RQuMF.

The algorithm takes in input a dataset X, a sub-problem size s, and inlier thresholds ϵ. It returns a set of labels l corresponding to a cover of X according to the retrieved models. The parameter s controls how many sampled models are processed in each iteration of the decomposed method. The first step (lines $3-10$) consists of generating a pool of M tentative models via random sampling and of computing their consensus set. The number of hypothesis M is defined as a multiple of the number of input points (line 3). Minimal sample sets are sampled from X using localized sampling (line 5) and used to fit a geometric model (line 6). Hence residuals are computed (line 7). Residuals smaller than the inlier threshold ϵ define the consensus sets of each sampled model, which are stored as columns in the preference matrix (line 8). The process is repeated until M models are sampled (line 4).

The second step (lines $11-25$) involves the decomposition of the preference matrix P to define the logical graph mapping. Specifically, P is partitioned into L sub-block P_j having at most s columns each (line 12). Each sub-problem P_j is converted to a QUBO form (lines $13-21$) with its logical graph and hence solved using RQuMF (line 19). Each solution z represents the selected models, indicated by the corresponding columns of P_j. In the pruning phase (line 22), the process retains only the chosen models, while the rest are eliminated, thereby reducing the dimensionality of P_j. Once the overall number of retained models falls below s, a final execution of the RQuMF method is carried out (line 26) to derive the final solution. The models selected in this last iteration of RQuMF (line 26) undergo label assignment (line 29) where the consensus set corresponding to each model is labelled with the i^{th} index of the model in question. Subsequently, to optimise the coverage of these labels across all data points, a linear assignment problem is tackled (line 33) yielding the final labels l.

Algorithm 1 De-RQuMF Method

1: **Input:** Point set X, problem size s, inlier threshold ϵ
2: **Output:** Predicted labels l
Generate preference consensus matrix
3: $M \leftarrow k \times |X|$, initialize P as a $|X| \times M$ zero matrix
4: **for** $i = 0$ to $M - 1$ **do**
5: Sample points from X
6: fit geometric model
7: compute residuals $R \forall \{(x_i, y_i)\} \in X$
8: Update $P[:, i] \leftarrow [r < \epsilon ? 1 : 0 \; \forall r \in R]$
9: $i \leftarrow i + 1$
10: **end for**
Logical graph mapping
11: **while** $|\text{columns}(P)| > s$ **do**
12: Partition P into L subproblems $\{P_j\}$ of size s
13: **for** $j = 0$ to L **do**
14: z = RQuMF(P_j)
15: **procedure** RQuMF(P_j)
16: $A \leftarrow$ concatenate $[-I; P_j]$
17: $\widetilde{Q} \leftarrow \lambda_2 \times A^T A$
18: $\widetilde{s} \leftarrow$ concatenate $[-\mathbf{1}_N; \lambda_1 \times \mathbf{1}_M]$
19: z = solveQUBO($\widetilde{Q}, \widetilde{s}$)
20: return z
21: **end procedure**
22: retain z from P_j
23: $j \leftarrow j + 1$
24: **end for**
25: **end while**
26: z = RQuMF(P)
Final model selection
27: return z
Generate label assignment
28: **for** each model z_i in z **do**
29: **for** each point x_j in consensus set of z_i **do**
30: Assign label i to x_j
31: **end for**
32: **end for**
33: Solve a linear assignment problem to maximize label coverage across X
34: return predicted labels l

8. Additional Details On Experiments

We provide here additional details about the experiments performed in Sec. 5 of the main paper.

Experiments Overview. In Fig. 8 we report a general overview of the experiments conducted on real and synthetic datasets, highlighting the different solvers used (*i.e.*, QA, SA etc) and the different configurations adopted to test scalability and robustness of our proposed methods.

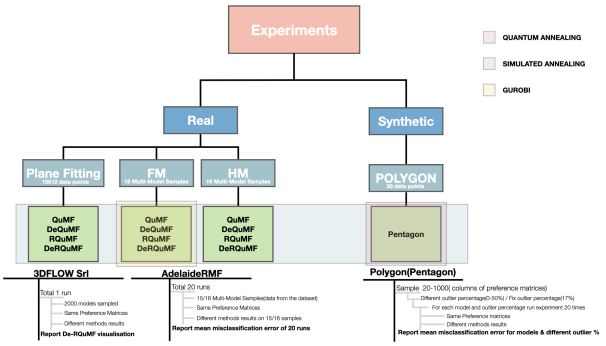

Figure 8. Overview of our experiments. "FM" refers to datasets based on the fundamental matrix model whereas "HM" refers to the homography model.

Outliers Percentage In Real Data Multi-Model Fitting Tasks. The plots in Fig. 9 show the outlier percentage in multiple fundamental matrices and homography fitting problems respectively, as referred in 5.2. Outliers typically correspond to wrong key-point matches that cannot be described by any model. It can be appreciated that in most of the pairs related to fundamental matrices (which are related to motion segmentation in two images) the outlier ratio is greater than 30%, while for homographies (which are related to plane fitting) we have that in 5 pairs out of 16, there are more than 50% of outliers, making the problem particularly challenging. This justifies the higher errors reported in plane fitting with respect to the one attained on motion segmentation.

Logical and physical graphs. In Fig. 10 we report the logical and physical graphs corresponding to sample problems related to the scalability experiment on the synthetic dataset, where we maintain a constant outlier ratio of 17% while expanding the sampled model size from 20 to 140. The left-side images showcase the logical graph representation of the problem, where each node corresponds to a logical qubit and the edges depict the coupling between these qubits. Through minor embedding, these logical qubits are mapped onto physical qubits within the quantum hardware.

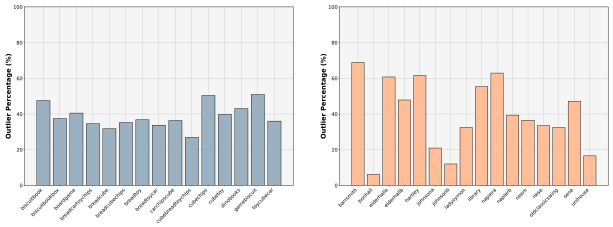

Figure 9. Outlier Percentage of each sequence in AdelaideRMF dataset [42]. Left: 15 image pairs for Fundamental Matrices fitting, Right: 16 image pairs for homographies.

On the right side, the physical representation of these mappings is displayed, with each node representing a physical qubit. The colour inside the node reveals the measured value in its most stable energy state, while the colour of the node's outer ring indicates the direction of bias, specifically pointing out whether the coefficient of the linear component in the optimization is positive or negative.

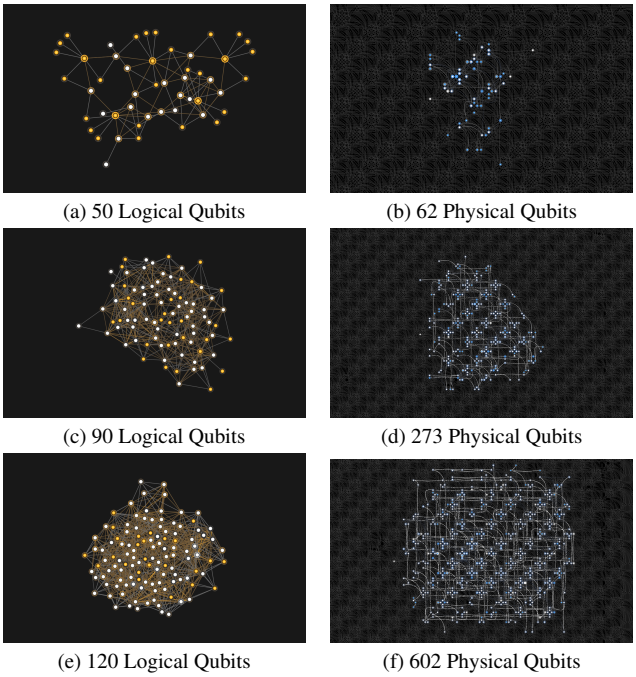

(a) 50 Logical Qubits (b) 62 Physical Qubits

(c) 90 Logical Qubits (d) 273 Physical Qubits

(e) 120 Logical Qubits (f) 602 Physical Qubits

Figure 10. **Left:** These images depict logical graphs for different problem sizes, reported in Fig. Fig. 3 of the main manuscript. **Right:** The corresponding physical qubit embedding of logical graphs(a, c, d) respectively using Pegasus topology on DWave Advantage5.4 as mentioned in Sec. 4.4 of the main paper.

Hyperparameter Tuning. We studied the effect of hyperparameter tuning in our experiments. Specifically, we studied how the objective value defined in Eq. (12) changes with respect to λ_1, λ_2. In Fig. 11 we plot the objective land-

scape for fundamental matrix (FM) fitting problems, using the Tree-structured Parzen Estimator (TPE) discussed in Sec. 4.4. It can be appreciated that the TPE strategy discards suboptimal values of the parameters and concentrates more on the ones that result in lower values (i.e, a cluster of points near the best objective values), contrary to grid search which doesn't consider the previous results for selecting parameters for the future. Additionally, in relation with the Tab. 3 the performance gap between QuMF and our method in the outlier-free setting can be narrowed by adjusting the lambda parameters to suit this specific case. For instance, when optimized lambda values ($\lambda_1 = 4.8$ and $\lambda_2 = 0.6$) are used for the fundamental matrix estimation task in the absence of outliers, the misclassification error for RQuMF decreases to 2.05, while for De-RQuMF, it drops to 6.18. Conversely, these parameters are not ideal for outlier-prone scenarios, where misclassification rates increase from 10.46 to 16.95 for RQuMF and from 12.69 to 15.75 for De-RQuMF. Therefore, we recommend adjusting the lambda parameters based on the specific conditions to achieve optimal results.

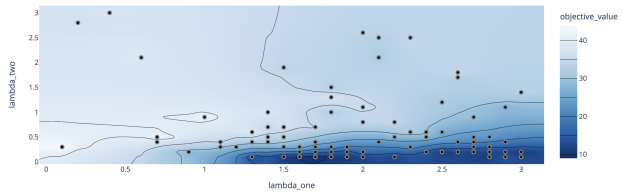

Figure 11. Lambda optimisation contour plot of RQuMF for fundamental matrix data ($\lambda_1 = 1.7\,, \lambda_2 = 0.1$) using TPE.

Number Of Selected Models. Our QUBO formulation does not require knowing in advance the number of models, thus we assess whether the number of estimated models matches the ground truth ones. Results are reported for fundamental matrix fitting problems in different setups, *i.e.*, without outliers in Fig. 12 and with outliers in Fig. 13. We compare the estimated number of models attained by QuMF, DeQuMF, RQuMF, and De-RQuMF. The plots show that our methods mostly select the true number of models, in contrast to previous methods which estimate the right number of models in the outlier-free scenario, but, without any kind of post-processing are prone to over-estimation in the presence of outliers. This is expected as, being based on set-cover rather than on maximum coverage, QuMF and DeQuMF try to maximize the number of inliers at the cost of hallucinating more models in the solution. It is worth noting, that even when coupled with post-processing, *i.e.*, after providing the right number of models, QuMF failed to achieve competitive results compared to our methods, highlighting the fact that these methods, in the presence of outliers, can not segment the data at the first place. See also Tab. 4 from the main paper.

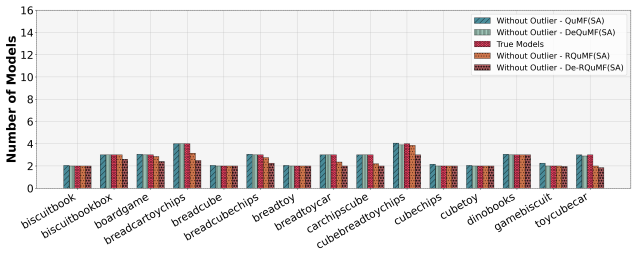

Figure 12. Average number of models selected by different methods for 15 multimodel sequences from AdelaideRMF [42] dataset for fundamental matrix fitting in the absence of outliers. The middle bar in the grouped bars represents the ground truth number of models, the left two bars represent QuMF, DeQuMF respectively and the right bars represent our proposed methods.

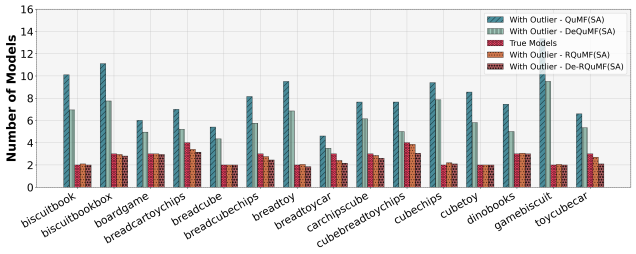

Figure 13. Average number of models selected by different methods for 15 multimodel sequences from AdelaideRMF [42] dataset for fundamental matrix fitting in the absence of outliers. The middle bar in the grouped bars represents the ground truth number of models, the left two bars represent QuMF, DeQuMF respectively and the right bars represent our proposed methods.

Execution Times. Tab. 5 presents the average execution time per sample of different methods on the AdelaideRMF [42] dataset. It can be noticed that DeQuMF by far outperforms in terms of execution time among all the methods confirming the advantages of the decomposed approach. Our formulation is less efficient as the dimension of the Q matrix has to encode also the number of points, while in the previous approach, Q scales with the number of sampled models. For reference, *biscuitbook* of the fundamental matrix fitting dataset has 341 data points. In one of the runs, the corresponding Q matrix dimension and node count in the logical graph are $(2046, 2046), 1793282$ for QuMF, $(40, 40), 744$ for DeQuMF, $(2387, 2387), 1890341$ for RQuMF, $(381, 381), 3059$ for DeRQuMF respectively. Thus De-RQuMF has to solve a problem of almost $\times 5$ bigger than DeQuMF (in terms of the logical graph). It's crucial to emphasise that the extended execution time of our method significantly enhances the reliability of the results,

contrary to the DeQuMF approach, which exhibits limitations as detailed in Tab. 5 of the main paper, our proposed method delivers reliable performance without failure. Additionally, we omit the reporting of methods execution time on quantum hardware, as the anneal time remains constant at $20\mu s$, independent of problem size. Note that, with the advent of stable Adiabatic Quantum Computers (AQC), our approach is not only expected to become significantly faster but also stay reliable.

Method	**FM**	**HM**
QuMF(SA) [15]	45.88	81.11
DeQuMF(SA) [15]	**4.40**	**4.93**
RQuMF(SA)(ours)	51.61	130.10
DeRQuMF(SA)(ours)	21.48	77.99

Table 5. Execution time (in seconds) of methods on different real datasets using Apple Silicon M1 machine with 8GB RAM.

Qualitative Results. We report the best and worst results for our RQuMF method for a few sequences of the Adelaide dataset in Fig. 14 and 15. Notably, our RQuMF method consistently surpasses the previous QuMF method, its non-decomposed counterpart, in performance across all scenarios, including both best and worst cases. Moreover, it excels beyond all other methods in three out of the four instances as depicted in Fig. 14 and 15. Additional visualizations are given in Fig. 16 and 17, confirming previous considerations.

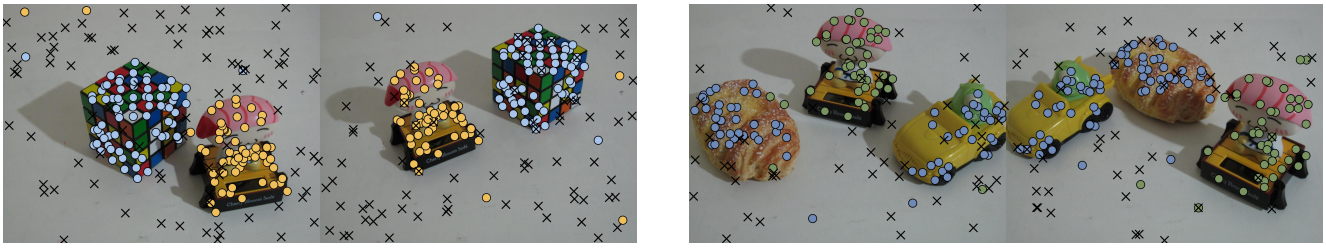

Figure 14. **Left:** A sample of the *best-case* for RQuMF on the *cubetoy* (average $E_{mis} = 3.73\%$) sequence of the AdelaideRMF [42] dataset for fundamental matrix (for the same sample average E_{mis} for De-RQuMF, QuMF and DeQuMF is $4.13\%, 42.95\%, 23.71\%$ respectively). **Right:** A sample of the *worst-case* for RQuMF on the *breadtoycar* (average $E_{mis} = 21.02\%$) sequence of the AdelaideRMF [42] dataset for fundamental matrix (for the same sample average E_{mis} for De-RQuMF, QuMF and DeQuMF is $26.23\%, 35.81\%, 21.95\%$ respectively).

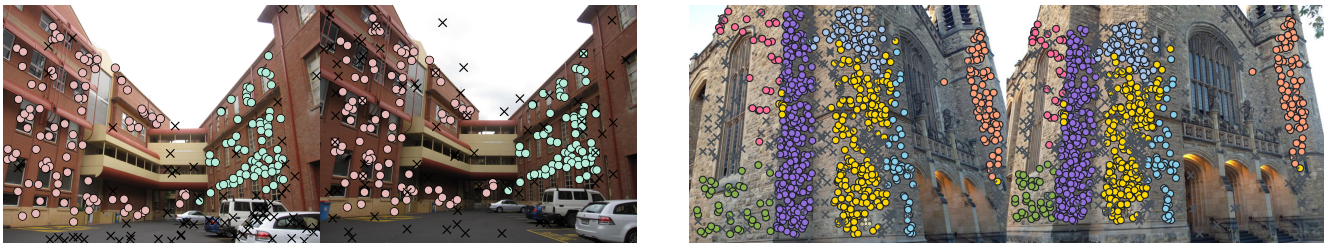

Figure 15. **Left:** A sample of the *best-case* for RQuMF on the *nese* (average $E_{mis} = 1.92\%$) sequence of the AdelaideRMF [42] dataset for homography matrix (for the same sample average E_{mis} for De-RQuMF, QuMF and DeQuMF is $2.14\%, 77.70\%, 28.29\%$ respectively). **Right:** A sample of the *worst-case* for RQuMF on the *bonhall*(average $E_{mis} = 41.60\%$) sequence of the AdelaideRMF [42] dataset for homography matrix (for the same sample average E_{mis} for De-RQuMF, QuMF and DeQuMF is $23.13\%, 74.20\%, 25.63\%$ respectively).

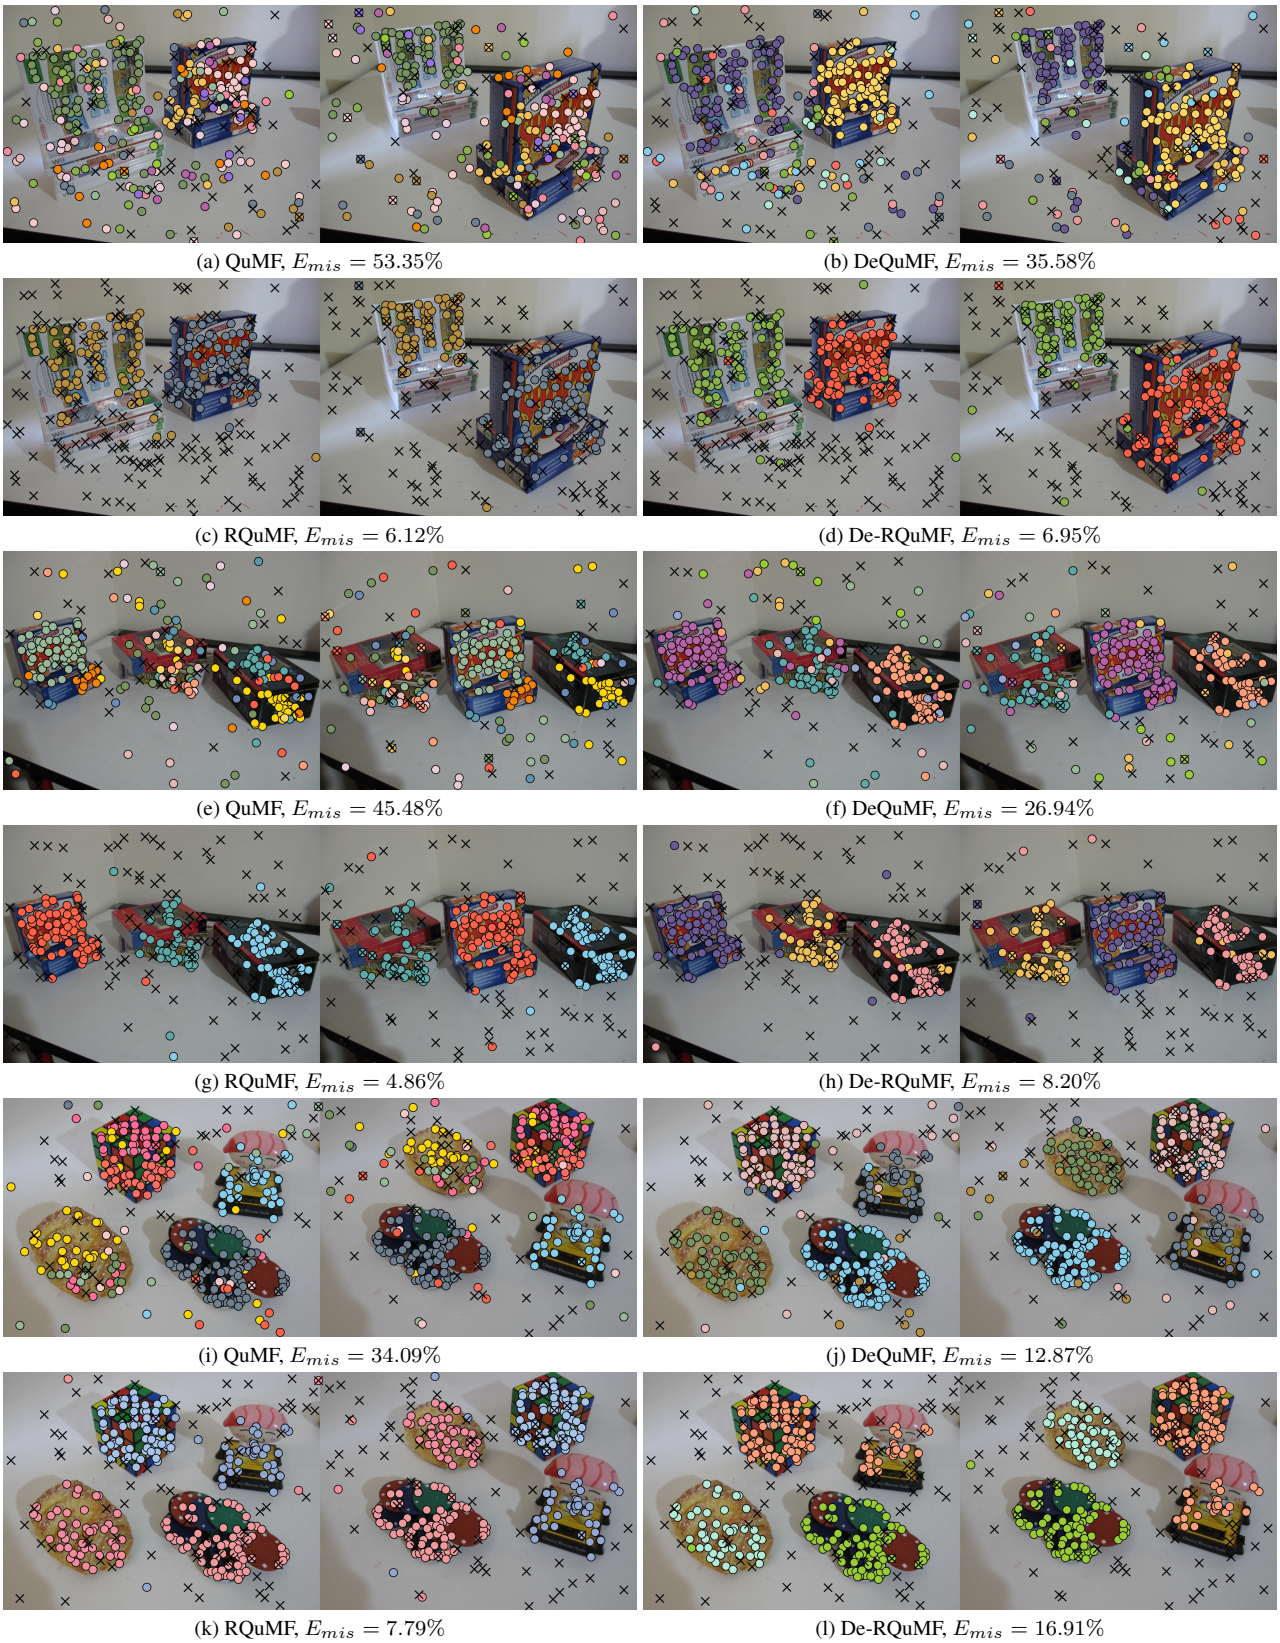

(a) QuMF, $E_{mis} = 53.35\%$

(b) DeQuMF, $E_{mis} = 35.58\%$

(c) RQuMF, $E_{mis} = 6.12\%$

(d) De-RQuMF, $E_{mis} = 6.95\%$

(e) QuMF, $E_{mis} = 45.48\%$

(f) DeQuMF, $E_{mis} = 26.94\%$

(g) RQuMF, $E_{mis} = 4.86\%$

(h) De-RQuMF, $E_{mis} = 8.20\%$

(i) QuMF, $E_{mis} = 34.09\%$

(j) DeQuMF, $E_{mis} = 12.87\%$

(k) RQuMF, $E_{mis} = 7.79\%$

(l) De-RQuMF, $E_{mis} = 16.91\%$

Figure 16. Average E_{mis} on some samples of AdelaideRMF [42] dataset for fundamental matrix fitting in the presence of outliers. one of our proposed methods outperforms the previous methods every time.

(a) QuMF, $E_{mis} = 83.99\%$ (b) DeQuMF, $E_{mis} = 49.88\%$

(c) RQuMF, $E_{mis} = 2.86\%$ (d) De-RQuMF, $E_{mis} = 2.80\%$

(e) QuMF, $E_{mis} = 86.51\%$ (f) DeQuMF, $E_{mis} = 52.59\%$

(g) RQuMF, $E_{mis} = 20.53\%$ (h) De-RQuMF, $E_{mis} = 24.75\%$

(i) QuMF, $E_{mis} = 80.27\%$ (j) DeQuMF, $E_{mis} = 27.40\%$

(k) RQuMF, $E_{mis} = 35.67\%$ (l) De-RQuMF, $E_{mis} = 22.49\%$

Figure 17. Average E_{mis} on some samples of AdelaideRMF [42] dataset for homography fitting in the presence of outliers. One of our proposed methods outperforms the previous methods every time in addition to being reliable.