Outlier-Robust Multi-Model Fitting on Quantum Annealers

Saurabh Pandey[1,2] Luca Magri[3] Federica Arrigoni[3] Vladislav Golyanik[1]
[1]MPI for Informatics, SIC [2]Saarland University [3]Politecnico di Milano

Abstract

Multi-model fitting (MMF) presents a significant challenge in Computer Vision, particularly due to its combinatorial nature. While recent advancements in quantum computing offer promise for addressing NP-hard problems, existing quantum-based approaches for model fitting are either limited to a single model or consider multi-model scenarios within outlier-free datasets. This paper introduces a novel approach, the robust quantum multi-model fitting (R-QuMF) algorithm, designed to handle outliers effectively. Our method leverages the intrinsic capabilities of quantum hardware to tackle combinatorial challenges inherent in MMF tasks, and it does not require prior knowledge of the exact number of models, thereby enhancing its practical applicability. By formulating the problem as a maximum set coverage task for adiabatic quantum computers (AQC), R-QuMF outperforms existing quantum techniques, demonstrating superior performance across various synthetic and real-world 3D datasets. Our findings underscore the potential of quantum computing in addressing the complexities of MMF, especially in real-world scenarios with noisy and outlier-prone data[1].

1. Introduction

Model fitting is a fundamental and challenging problem in computer vision, with applications such as 3D reconstruction, scene layout estimation, motion segmentation, and image stitching. Its objective is to explain input data (e.g., 2D or 3D point sets) using a non-redundant number of parametric models. However, challenges arise with multiple models, whose exact number is typically unknown, and the need for robustness against outliers. These issues are particularly critical in homography and fundamental matrix estimations, where errors can significantly impact downstream tasks. Ultimately, multi-model fitting is an ill-posed problem with a combinatorial nature, where data clustering and model estimation must be solved simultaneously.

Multi-model fitting (MMF) has been actively researched during the last decades, following many different principles (*e.g.*, optimisation and clustering-based approaches, see Sec. 2). One of the latest and newest research directions in the field focuses on adopting quantum computational paradigms, either gate-based or quantum annealing [11, 13, 15]. While gate-based quantum machines are universal in the sense they complement classical computers with a set of additional operations, quantum annealers can be thought of as samplers of a specific type of optimisation problems, *i.e.*, quadratic unconstrained binary optimisation (QUBO) objectives. The latter has recently gained a lot of attention in the community since many computer vision tasks can be rephrased as a QUBO, including matching problems [8, 9, 31, 35, 36], object detection [24], multi-object tracking [45], motion segmentation [3] and neural network weight optimisation [23, 33]. The main motivation for using quantum computers lies in their promise to accelerate the solution search for combinatorial optimisation problems while returning globally optimal solutions with a certain non-zero probability [14]. Advantages in adopting a quantum approach have also been shown for non-combinatorial problems like point-set registration [17, 30].

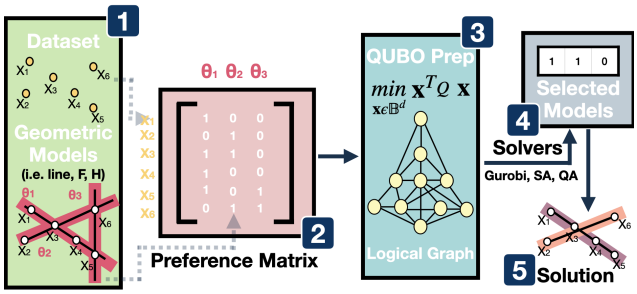

Figure 1. **Overview of our R-QuMF**, a multi-model fitting approach that is robust to outliers and admissible to modern quantum annealers. We first sample models that along with the data define the preference matrix P. Next, a QUBO problem is prepared that can be minimised by quantum annealing (after a minor embedding of the logical problem on quantum hardware) or other solvers. Finally, the best solution is selected. R-QuMF outperforms previous quantum-admissible model fitting approaches.

[1]Project page: https://4dqv.mpi-inf.mpg.de/RQMMF/

Method	Multi-Model	Outlier-Robust
HQF-RF [13]	×	✓
QuMF [15]	✓	×
Ours	✓	✓

Table 1. Comparison of method characteristics.

In the context of model fitting, Chin *et al.* [11] introduced a method based on gate-based quantum hardware, while other works rely on the quantum annealing paradigm [13, 15]. These methods were shown to provide improvements compared to classical methods due to the quantum effects both from the theoretical and practical perspectives; they are compatible with current and upcoming generations of quantum hardware.

Among these approaches, [11, 13] address the case of a single model. The QuMF method of Farina *et al.* [15] for quantum annealers, instead, not only outperforms previous quantum single-model fitting approaches but also supports multiple models and competes with classical state-of-the-art. However, the main drawback of QuMF [15] is that it assumes outlier-free data, which is unrealistic in most practical scenarios. The naive way to extend such an approach to managing outliers is by post-processing, *i.e.*, only the k largest models are selected at the end among all candidate models obtained by random sampling (with k equal to the true number of models explaining the data). Besides requiring the knowledge of k, or equivalent ancillary information, this approach is sub-optimal and prone to inaccuracies, as demonstrated by our experiments on standard datasets.

This paper addresses the challenge of outlier handling in multi-model fitting and tailors for quantum annealers a new multi-model fitting approach, which we call Robust Quantum Multi-Model Fitting (R-QuMF); see Fig. 1. In contrast to previous quantum work [15], it accounts for outliers explicitly in the formulation, resulting in a more general approach while exhibiting superior results on real data. Another advantage of our method is that it does not require any prior information about the optimal number of models explaining the data, which is convenient in practical applications. More details on the differences with respect to previous quantum papers are reported in Tab. 1. To summarise, the contributions of this paper are two-folds:

i) R-QuMF, a new approach for outlier-robust fitting of multiple models;

ii) A formulation compatible with quantum annealers that accounts for outliers explicitly: Our method does not need any post-processing steps or a number of optimal models in advance as input, which are highly advantageous properties in practice.

We also apply to R-QuMF the decomposition principle similar to De-QuMF [15]. It addresses the limitations of current AQC hardware by iteratively decomposing the original large problem into smaller QUBO sub-problems until the final sub-problem selects the solutions among the most promising models. Our approach significantly outperforms previous quantum techniques in the experiments with various multi-model fitting scenarios with different outlier ratios such as geometric model fitting, homography estimation and fundamental matrix estimation. The source code of our method for all solver versions (including D-Wave and demo examples) will be made available.

2. Related Work

Classical Approaches. Multi-model fitting has been addressed since the 1960s, with effective techniques ranging from the Hough transform [43] to more recent approaches based either on clustering or optimising an objective function. Clustering-based methods [2, 6, 10, 18, 21, 26, 28, 29, 32, 37, 38, 40, 46] focus on data segmentation and offer procedural, easy-to-implement solutions that produce promising results in most cases. However, hard clustering of data does not always produce optimal results when models overlap. On the contrary, optimisation methods prioritise the refinement of a precise objective function, offering a quantitative measure to assess the quality of the derived solution. The most common objective functions are typically based on consensus, *i.e.*, they aim at maximising the number of inliers of each model, so optimisation-based methods [4, 5, 20, 27, 39, 47] can be considered as sophisticated extensions of the popular RanSaC paradigm [16] to the case of multiple models.

In this respect, the classical work that is mostly related to our approach is RanSaCov [27], which casts multi-model fitting as a coverage problem. Given a collection of models with their consensus sets, RanSaCov extracts either the minimum number of models that explain all the points (set cover formulation) or, if the number k of the sought models is known in advance, it selects the k models that explain most of the data (maximum coverage formulation). The latter is effective in dealing with data contaminated by outliers that can be recognised as uncovered points.

We borrow from RanSaCov the maximum coverage formulation. However, RanSaCov [27] solves the coverage problems via integer linear programming and branch and bound, hence, it either resorts to approximations or falls back to enumerating all candidate solutions. Instead, our approach exploits quantum effects to optimise the objective directly in the space of qubits, where global optimality is expected with high probability after multiple anneals. In addition, we directly minimize the number of models as done in several traditional MMF frameworks [4, 5, 20], without requiring the knowledge of the true number of models in advance. Contrary to RanSaCov, we do not deal natively with intersecting models, but inliers belonging to multiple models can be identified after the models have been extracted by

inspecting the point-model residuals.

Approaches Compatible with Quantum Hardware. While many classical methods have been developed, the first methods based on quantum computing have only recently attempted to exploit the capabilities of quantum hardware to tackle the combinatorial nature of the problem. The quantum solutions presented so far do not address the multi-model fitting problem under outliers. Moreover, other approaches—both theoretical and practical—start to address closely related problems such as linear regression [12, 34], clustering and segmentation [3, 41] to capitalise on the advantages of quantum computing.

The first attempts to address model fitting with the help of quantum hardware concentrate on single-model fitting [11, 13]. Chin *et al.* and Yang *et al.* [11, 44] introduced a single-model fitting method based on gate-based quantum hardware, while Doan *et al.* [13] rely on quantum annealing. Although both Doan *et al.* [13] and our method are based on linear programming, their approaches diverge significantly. Doan *et al.* [13] employ a hypergraph formalism that relies on multiple QUBOs within an iterative framework, while our method is more streamlined, using a single QUBO. Additionally, while they focus on single-model fitting, our approach extends to multi-model fitting.

Farina *et al.* [15] take this further with their QuMF method for quantum annealers, which considers multiple models and achieves results on par with classical state-of-the-art techniques. However, QuMF is primarily limited by its reliance on outlier-free data. While post-processing can improve outlier robustness, it usually requires prior knowledge of the number k of models and often leads to sub-optimal results, as evidenced by our experiments.

3. Background on Quantum Annealers

Modern quantum annealers (QAs) can sample Quadratic Unconstrained Binary Optimisation (QUBO) problems, which in a general form can be written as

$$\arg\min_{\mathbf{y}\in\mathbb{B}^d} \quad \mathbf{y}^T Q \mathbf{y} + \mathbf{s}^T \mathbf{y} + \sum_i \lambda_i ||A_i \mathbf{y} - \mathbf{b}_i||_2^2, \qquad (1)$$

where $\mathbf{y} \in \mathbb{B}^d$ is a vector of d binary variables, $Q \in \mathbb{R}^{d\times d}$ is a real symmetric matrix, $A_i \in \mathbb{R}^{d\times d}$ are real matrices, $\mathbf{s}, \mathbf{b}_i \in \mathbb{R}^d$, and λ_i are scalar weights. The terms under the ℓ_2-norm are rectifiers expressing soft linear constraints. They preserve the QUBO problem type since Eq. (1) can be written without constraints as follows:

$$\arg\min_{\mathbf{y}\in\mathbb{B}^d} \quad \mathbf{y}^T \widetilde{Q} \mathbf{y} + \tilde{\mathbf{s}}^T \mathbf{y}, \qquad (2)$$

with $\widetilde{Q} = Q + \sum_i A_i^T A_i$ and $\tilde{\mathbf{s}} = \mathbf{s} - 2\sum_i A_i \lambda_i \mathbf{b}_i$. Eq. (2) is the combinatorial QUBO form admissible to modern QAs; the elements of $\widetilde{Q}$ and $\tilde{\mathbf{s}}$ along with the number of binary variables to be optimised have to be provided.

During quantum annealing, $\mathbf{y}_i$ are modelled as qubits weighted by $\tilde{\mathbf{s}}_i$ with the strength of the mutual influence defined by $\widetilde{Q}$. The optimisation takes place in the 2^d-dimensional Hilbert space and involves quantum-mechanical effects of qubit superposition, entanglement and quantum tunnelling.

We will call the connectivity pattern between the binary variables in Eq. (2) the *logical problem graph* (i.e. in which qubits are represented by vertices and edges between the vertices are present if $Q_{i,j}$ entries are non-zero). Since direct interactions between only a small subset of all possible $d(d-1)/2$ pairs of binary variables are enabled by the hardware, mapping the logical problem graph into the physical hardware is necessary for most QUBO problems. This mapping is called *minor embedding*, as the logical problem graph is interpreted as a graph minor of a larger graph, i.e. hardware graph of physical qubits. This means the hardware graph contains qubit chains representing a single logical qubit from the logical problem graph.

4. The Proposed R-QuMF Method

This section presents our robust quantum multi-model fitting (R-QuMF) approach and details of its implementation. The method can be summarised as in Fig. 1: Steps 1 and 2 are described in Secs. 4.1 and 4.2. Steps 3 and 5 correspond to Sec. 4.3, whereas details of Step 4 (implementation/solvers) are in Sec. 4.4.

4.1. Preliminaries, Definitions and Notations

We frame the multi-model fitting problem in the presence of outliers as follows. We are given in input a set of points $X = (x_1, x_2, ..., x_n)$ and a collection of models $\Theta = (\theta_1, \theta_2, ..., \theta_m)$ generated through random sampling akin to those of the RANSAC algorithm [16]. The desidered output is a subset $\{\theta_{i1}, \dots \theta_{ik}\} \subset \Theta$ of *non-redundant* models that explain the data. Having non-redundant models is a form of regularisation to make the ill-posed multi-model fitting problem tractable. Echoing Occam Razors' principle, we favour the interpretation of the data that minimizes the number of models required. Moreover, we assume that: *i*) X can be corrupted by high outlier percentages (up to 50%), and that *ii*) the number k of true models explaining the underlying data is unknown. The above-mentioned assumptions are very desirable in practical applications.

The problem can be equivalently formulated in terms of a preference-consensus matrix defined as

$$P[i,j] = \begin{cases} 1 & \text{if } \mathrm{err}(x_i, \theta_j) < \epsilon, \\ 0 & \text{otherwise,} \end{cases} \qquad (3)$$

where $P \in \mathbb{B}^{n\times m}$ is a $n \times m$ binary matrix; n and m be-

ing the number of points and models, respectively. The i-th data point x_i is assigned to the j-th sampled model θ_j if its residual is below an inlier threshold ϵ. Operator $\text{error}(\cdot,\cdot)$ quantifies the point-to-model distance. Following [27], the rows of P can be interpreted as preference sets, while the columns of P—denoted by $S_1,\dots,S_m$—represent consensus sets (w.r.t. ϵ). MMF reduces to selecting from P the columns that correspond to the sought models $\{\theta_{i1},\dots\theta_{ik}\}$.

4.2. Revisiting Maximum-Set Coverage Objective

In order to gain robustness against outliers, we generalise the QUBO formulation presented in QuMF [15] approach – which, in turn, is based on the Set Cover formulation presented in [27]. Specifically, we revisit the maximum-set coverage (MSC) objective for MMF [27]. The MSC task is to select at most k subsets from Θ such that the coverage of the data points contained in the set X is maximum (or as complete as possible); all the uncovered points are considered outliers. Intuitively, with reference to the matrix P, we want to select k columns that explain most of the points.

Let us introduce n binary variables $y_1,\dots,y_n$ such that: $y_i = 1$ if x_i is covered by one of the θ_j, or, in other terms, it is part of the selected subsets (i.e. the point is an inlier); $y_i = 0$ otherwise (i.e. the point is an outlier). Let us consider additional auxiliary variables $z_1,\dots,z_m$ such that: $z_j = 1$ if model θ_j is selected; $z_j = 0$ otherwise. Using this notation, the MSC problem is formulated as an integer linear programming:

$$\max \sum_{i=1}^{n} y_i \quad s.t. \sum_{j=1}^{m} z_j \leq k, \sum_{j:S_j \ni x_i} z_j \geq y_i \quad \forall x_i \in X \tag{4}$$

where both $y_i \in \{0,1\}$ and $z_j \in \{0,1\}$. In this context, the first constraint imposes that at most k models are selected (with k known in advance); the second constraint ensures that, if $y_i = 1$, then at least one set θ_j containing x_i must be chosen. Recall that n corresponds to the cardinality of the set X and m denotes the number of candidate models. Note that, within this formulation, uncovered points are considered outliers. Therefore, outliers and uncovered points are used interchangeably and no artificial models for outliers need to be introduced.

The combinatorial nature of the problem makes it a suitable choice for designing a QUBO method compatible with a quantum annealer. To accomplish such a task, we make some changes with respect to the MSC objective in (4). First, we replace the inequality in the second constraint with an equality, therefore looking for *disjoint models*. Secondly, instead of demanding the approach to select k models, we opt for directly minimising the number of selected models. This choice is motivated by the fact that, in many practical situations, assuming known k is restrictive. Therefore our final objective, converted into a minimisation problem, is:

$$\min -\sum_{i=1}^{n} y_i + \lambda_1 \sum_{j=1}^{m} z_j \quad s.t. \sum_{j:S_j \ni x_i} z_j = y_i \quad \forall x_i \in X \tag{5}$$

where λ_1 is a regularisation parameter.

4.3. MSC Reformulated as QUBO

The first step to reformulate MSC as QUBO is to vectorise the objective in Eq. (5) and convert the constraint into a matrix form. More precisely, we rewrite it as follows:

$$\min_{\mathbf{y}\in\mathbb{B}^n,\ \mathbf{z}\in\mathbb{B}^m} -\mathbb{1}_n^T\mathbf{y} + \lambda_1(\mathbb{1}_m^T\mathbf{z}) \quad s.t \qquad P\mathbf{z} = \mathbf{y} \tag{6}$$

where $\mathbb{1}$ denotes a vector of ones (whose length is given as a subscript), P is the preference matrix of size $n \times m$, $\mathbf{y}$ and $\mathbf{z}$ are binary vectors collecting the y_i and z_j variables, respectively. To simplify the formulation into a single variable optimisation problem, we can combine the two unknowns into a single variable $\mathbf{w}$:

$$\mathbf{w} = \begin{pmatrix}\mathbf{y}\\ \mathbf{z}\end{pmatrix} \in \{0,1\}^{n+m}, \tag{7}$$

and rewrite the main objective of (6) in terms of variable $\mathbf{w}$:

$$\begin{aligned} -\mathbb{1}_n^T\mathbf{y} + \lambda_1(\mathbb{1}_m^T\mathbf{z}) &= [-\mathbb{1}_n^T,\ \mathbb{0}_m^T]\mathbf{w} + \lambda_1[\mathbb{0}_n^T,\ \mathbb{1}_m^T]\mathbf{w} \\ &= [-\mathbb{1}_n^T,\ \lambda_1\mathbb{1}_m^T]\mathbf{w} \end{aligned} \tag{8}$$

where $\mathbb{0}$ represents a vector of zeros, with its length indicated as a subscript. Furthermore, the constraint in (6) can also be reformulated in terms of the newly introduced variable $\mathbf{w}$ as follows:

$$P\mathbf{z} - \mathbf{y} = \mathbb{0}_n \ \Leftrightarrow\ [-\mathbb{I}_{n\times n},\ P]\mathbf{w} = \mathbb{0}_n \tag{9}$$

where $\mathbb{I}_{n\times n}$ is an identity matrix of size $n \times n$.

After incorporating the constraint from Eq. (9) as a penalty term into equation (8) we obtain

$$\min_{\mathbf{w}\in\mathbb{B}^{n+m}} [-\mathbb{1}_n^T,\ \lambda_1\mathbb{1}_m^T]\mathbf{w} + \lambda_2||[-\mathbb{I}_{n\times n},\ P]\mathbf{w}||_2^2. \tag{10}$$

Now, when comparing (10) with (1), the following correspondences emerge:

$$\begin{aligned} Q = 0, \quad \mathbf{s}^T = [-\mathbb{1}_n^T,\ \lambda_1\mathbb{1}_m^T], \\ A = [-\mathbb{I}_{n\times n},\ P], \quad \mathbf{b} = \mathbb{0}_n. \end{aligned} \tag{11}$$

Note that Q equals zero matrix as there are no quadratic terms involved. Finally, we obtain the target QUBO objective admissible on quantum hardware that can be written in the form of Eq. (2):

$$\min_{\mathbf{w}} \ \mathbf{w}^T\widetilde{Q}\mathbf{w} + \widetilde{\mathbf{s}}^T\mathbf{w}, \tag{12}$$

where

$$\mathbf{w} = \begin{pmatrix} \mathbf{y} \\ \mathbf{z} \end{pmatrix} \in \{0,1\}^{n+m}, \\ \widetilde{Q} = \lambda_2 \begin{pmatrix} \mathbb{I}_{n\times n} & -P \\ -P^T & P^T P \end{pmatrix}, \widetilde{\mathbf{s}} = \begin{pmatrix} -\mathbb{1}_n \\ \lambda_1 \mathbb{1}_m \end{pmatrix}. \tag{13}$$

The QUBO formulation in Eq. (12) can be optimised by classical global optimisation algorithms such as simulated annealing, or sampled on a quantum annealer. Note that the number of unknowns scales linearly with the number of points and models. Finally, we point out that the decomposition principle introduced by Farina *et al.* [15] can be applied to our QUBO as well, to manage large-scale applications. The core principle is to iteratively break down the input problem into smaller (tractable) sub-problems and aggregate the respective results. Algorithm A in our supplement provides a summary of our method.

4.3.1. Selection of λ's

The lambda parameters from Eq. (13) are decisive for the performance of our method. To find suitable values, we use Tree-structured Parzen Estimator (TPE) [7] which employs a Bayesian optimisation strategy utilizing a probabilistic model to steer the search process towards hyperparameter configurations that are more likely to improve the performance metric of interest. This model-based approach contrasts sharply with exhaustive grid search, which operates without leveraging prior knowledge or outcomes of evaluations. TPE optimises by constructing and refining a probabilistic model based on past evaluation results, thereby smartly converging to optimal hyperparameters through sequential model fitting and utility-based sampling.

4.4. Implementation Details

Our QUBO objective can be optimized either on CPU using classical solvers, like Simulated Annealing (SA) [22] and Gurobi [19], or on QPU (Quantum Processing Unit) via Quantum Annealing (QA).

For simulated annealing, we use D-Wave's neal package (version 0.6.0), and we fix the number of samples for SA to 100 (the same used in Farina *et al.* [15]), and we adopt this configuration also for the competing methods.

As regard Gurobi, we rely on version 10.0.3 under academic license with a time limit of 120 seconds for both RQuMF and the decomposed version De-RQuMF. The same configuration was used for QuMF and De-QuMF [15].

Experiments with Quantum Anneling are performed on D-Wave quantum annealer Advantage 5.4. We set the number of anneals to $5k$ for the one-shot version and $2.5k$ for the decomposed version (the same values used in Farina *et al.* [15]). In total, we used approximately 16 minutes of QPU time for our experiments. The subproblem size in De-RQuMF is set to 40 (the same as in Farina *et al.* [15]). We adopt the maximum chain length criterion for all conducted experiments to calculate chain length for D-Wave experiments. This involves initially mapping a logical graph onto a physical graph through the process of minor embedding. After determining the final embedding, we calculate the length (l) of the longest chain of qubits. Subsequently, the chain strength parameter is established by adding a small offset to l specifically, an offset of 0.5 inline with the previous works[3, 9]. See further implementation details of R-QuMF in the released source code.

5. Experiments

We assess the effectiveness of the proposed method on both synthetic and real datasets. Specifically, we compare our approach with RanSaCov [27], a classical method, and a quantum apporach, namely the recently introduced QuMF [15], as these methods are the closest competitors (see Sec. 2). To evaluate the performance of the analysed methods we adopt the *misclassification error*, denoted as E_{mis}, which judges the quality of MMF in terms of the segmentations attained. Specifically, the misclassification error counts the number of misclassified points as follows: first, each point is assigned to a label corresponding to the model it belongs to (outliers are assigned to the label 0); then, the map between ground-truth labels and estimated ones that minimises the overall number of misclassified points is found; a point is deemed as correct if one of its labels corresponds to the ground truth. The evaluated multi-model fitting tasks include fitting lines to 2D points (synthetic), fitting planes to 3D points (real) and two-view segmentation on the AdelaideRMF dataset [42] for fitting fundamental matrices or homographies.

5.1. Line Fitting on Synthetic Data

We evaluate our method on line fitting problems by creating a synthetic test-bed that comprises five lines arranged into a pentagon (see Fig. 2 for some visualisations). Each line fitting problem comprises 30 data points, divided into outliers and inliers. Outliers are uniformly distributed, whereas inliers (i.e., points belonging to the lines) are perturbed using Gaussian noise with a standard deviation of 0.01, and they are equally distributed among the lines. Each test is repeated 20 times and the mean misclassification error is reported. The preference matrix is the same for all compared methods and true models have always been sampled in the pool of provisional models Θ (a standard practice).

Robustness to Outliers. First, in order to evaluate the robustness of the methods, we gradually increase the outlier percentage from 0% to 50% while keeping fixed the total number of points (i.e., 30 points). The number of sampled models is set equal to 40 models. Results of this experiment are reported in Top Fig. 3. QuMF and De-QuMF are not designed to deal with outliers and, as expected, they do not

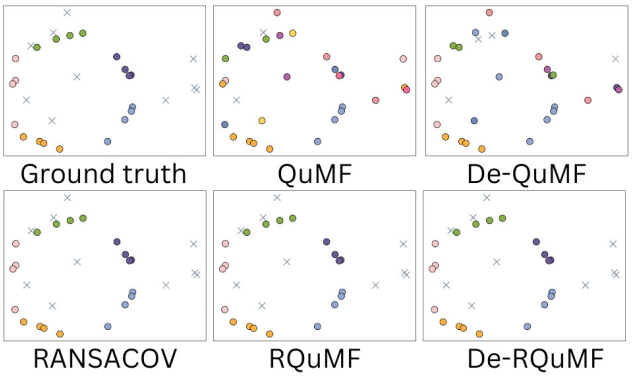

Figure 2. A sample visualization of the synthetic dataset (Ground Truth) and results for various methods for 50 models and 33% outliers (i.e. 10 outliers out of 30 points).

# of Models	QuMF[15]	De-QuMF[15]	RanSaCov[27]	RQuMF	De-RQuMF
20	1.00	1.00	**0**	**0**	**0**
50	16.16	9.50	**0.66**	**0.66**	**0.66**
100	44.83	18.66	**0**	1.33	**0**
500	86.66	36.66	14.66	30.99	**0**
1000	89.50	49.83	15.99	35.99	**3.32**

Table 2. Misclassification Error [%] on synthetic data with varying problem size (i.e. number of sampled models); the outlier percentage is fixed to 17% (i.e. five outliers); data size is fixed to 30 points. SA is used for quantum methods. This experiment is the classical counterpart of the evaluations reported in the bottom Fig. 3. All variants of quantum-enhanced methods in this table (Farina et al. [15] and ours) use simulated annealing.

provide accurate results as the number of outliers increases. This phenomenon can also be observed in Fig. 2, where outliers are erroneously detected as belonging to models supported by very few points. On the contrary, RanSaCov, RQuMF and De-RQuMF, which deal with outliers by design, are able to cope with a higher percentage of outliers without suffering a huge performance degradation.

Scalability. In our second experiment, we focus on how the performance scales with respect to the dimension of the multi-model fitting problem at hand. Thus, having fixed the outlier ratio at 17%, we vary the number m of sampled models between 20 and 1000 (with a step size of 10 until 200 models are reached, and a step size of 100 after that). Results of this experiment are reported in Tab. 2. It can be noted that, when the number m of sampled models increases, the misclassification error achieved by RQuMF worsens, as the problem becomes more difficult to solve. On the contrary, De-RQuMF can handle a higher number of sampled models without impacting the overall misclassification error, thanks to the decomposition principle. It can also be observed that our approach outperforms previous quantum work (QuMF and De-QuMF), similarly to the previous experiment.

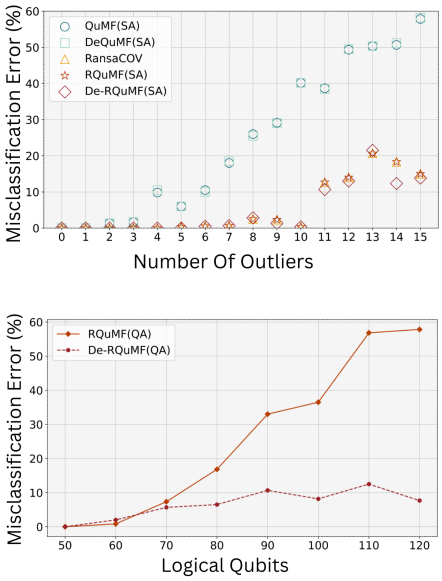

Figure 3. **Top:** Misclassification Error [%] on synthetic data for 40 sampled models with increasing outliers (0-50%); the problem size is fixed to 70 (30 data points + 40 models). SA is used for quantum methods. **Bottom:** Misclassification Error [%] for synthetic data on quantum hardware with increasing problem size; outlier percentage is fixed to 17%. Note that R-QuMF's E_{mis} breaks starting from 120 qubits. Non-robust quantum methods (i.e. QuMF and De-QuMF) are omitted because they fail in this scenario.

Experiments on Quantum Hardware. In order to assess the performance on Quantum Hardware, we also repeat the previous experiment on D-Wave quantum annealer Advantage 5.4. We consider a constant outlier ratio of 17% and problem size varying from 50 to 120 qubits (which is the sum of the number of points and models, namely 30 data points plus 20 to 140 sampled models). Results are reported in Bottom Fig. 3. We observe that RQuMF cannot handle problems starting from 80 models (*i.e.*, 110 qubits)–when its error goes beyond 50%, while De-RQuMF remains robust for the entire range of problem sizes with E_{mis} predominantly staying under 10%. This confirms the advantages of a decomposed and iterative approach for handling high-dimensional problems. The fact that RQuMF can not manage large-scale problems is not surprising since quantum hardware is far from being mature, in agreement with previous work on quantum computer vision [3, 9]. It is worth noticing that results on quantum hardware are worse than the ones with SA reported in Tab. 2, as expected. For example, SA can successfully manage problems with 100 sampled models with 1.33% error whereas QA fails.

To further analyze this aspect, we visualize in Left Fig. 4

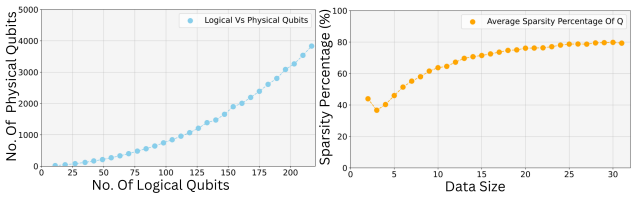

Figure 4. Analysis of R-QuMF runs on our synthetic dataset. **Left:** The number of physical qubits as a function of the number of logical problem qubits for data points varying in the range $[2; 32]$; sampled models are 6 times the data size. **Right:** The sparsity of $\widetilde{Q}$ in % as the function of the input data size.

how the number of physical qubits (that reflects the effective allocation of QPU resources) scales with respect to the increasing problem size (represented by the number of logical qubits) on sample problems. Although $\widetilde{Q}$ in our QUBO is significantly sparse (see Fig. 4-(right)), the number of physical qubits still increases superlinearly with respect to the logical qubits, approaching the maximum size that can typically be handled by an adiabatic quantum computer. These results are in line with previous quantum work [15].

5.2. Motion Segmentation on Real Data

We consider the AdelaideRMF dataset [42], which encompasses two distinct types of multi-model fitting tasks: fitting fundamental matrices (15 image pairs with at least two moving objects) and homographies (16 image pairs with at least two planes). Our evaluation specifically targets the multi-model sequences associated with both types of fitting problems and we do not take into account single-model fitting. The outlier percentages for these data are depicted in the supplementary material, with nearly all sequences exhibiting an outlier rate exceeding 30%, and some reaching as high as 68% (for more detail about the outlier distribution see Fig. 9 from supplementary material): this presents a substantial challenge for accurate model fitting. The number of sampled models for each instance is six times the number of points. As before, the preference matrices used as input remain consistent across all evaluated methods. We conducted each experiment 20 times, reporting the average E_{mis}.

Unlike the synthetic experiments, where computational resources are less constrained, for real data we do not report results obtained using QA due to limited QPU time. In addition to SA, we also consider the Gurobi solver in order to enrich the evaluation. In addition to the original outlier-contaminated sequences, we also consider the same sequences where outliers have been removed, to study the impact of outliers and diverse behaviour of the considered methods. We also analyse the efficacy of post-processing, especially related to QuMF. Specifically, we examine whether 1) selecting the top k models identified by a method and 2) designating all points not accounted for by these k models as outliers, would yield robustness. Note that this post-processing does not make sense for RanSaCov, for it enforces hard constraints in its formulation and returns at most k models. Aggregated results for fitting fundamental matrices are given in Tab. 3, whereas those for homographies are reported in Tab. 4. See also Fig. 5 and 6 for sample qualitative results. (More visualizations can be referenced from the supplementary material in Fig. 16 and 17)

Similar conclusions can be drawn for fitting fundamental matrices and homographies. It is not surprising that QuMF outperforms our approach in the outlier-free scenario (as seen in the first row of the aforementioned tables), given that QuMF is specifically designed for such settings. Additionally, the lambda parameters in RQuMF have been fine-tuned for scenarios involving outlier contamination.

In the case with outliers (second row of the tables), however, QuMF is significantly worse than RQuMF, therefore showing that explicitly modeling outliers is indispensable for achieving robustness. Using the post-processing largely improves the performance of QuMF, which, however, is still not comparable to RQuMF. Note also that the post-processing assumes the knowledge of the number of true models k which is typically unavailable in practice. Our method is not influenced by post-processing, thereby showing that it selects the right number of models in most cases, and it is better than RanSaCov in many scenarios. Concerning quantum methods, there are no significative differences between using Gurobi or SA solvers. The fact that the decomposed version of our approach does not improve upon using the full QUBO, could be due to the task difficulty in terms of outlier corruption compared to the simplified scenarios with synthetic data. From Fig. 5 it becomes apparent that QuMF can segment the entire model accurately only with the aid of the post-processing phase. Without post-processing, QuMF struggles to segment the entire model accurately, often choosing multiple models to explain what essentially constitutes a single true model (see Fig. 6).

5.3. Plane Fitting on 3D Point Clouds

Finally, we illustrate the versatility and practicality of our approach in a 3D plane fitting scenario. We consider a 3D point cloud obtained through image-based 3D reconstruction [1]. The dataset comprises 10812 points; we sample 2000 models from those, focusing exclusively on planar structures, with an inlier threshold set to 0.5 and use the SA solver. Fig. 7 provides a visual example of a fitting performed using our decomposed method, De-RQuMF. The results demonstrate that our method identifies distinct planes within the point cloud. As expected, the fitting accuracy for cylindrical sections of the building is lower as our method supports plane sampling exclusively per design.

Outliers Settings		QuMF(SA)[15]	De-QuMF(SA)[15]	RanSaCov[27]	RQuMF(SA)	De-RQuMF(SA)	QuMF(G)[15]	De-QuMF(G)[15]	RQuMF(G)	De-RQuMF(G)
No Outliers	Mean	3.61	**0.84**	9.79	6.46	11.47	3.14	1.29	12.56	11.96
	Median	2.68	**0.28**	7.97	2.41	10.55	1.87	0.93	11.35	10.55
With Outliers	Mean	40.37	26.19	**7.22**	10.46	12.69	45.81	26.21	13.14	12.84
	Median	39.82	26.94	**5.76**	8.33	11.18	46.58	27.35	10.96	11.33
With Outliers + Post Processing	Mean	19.76	**8.89**	NA	9.70	12.48	25.61	8.94	12.67	12.59
	Median	19.67	**6.55**	NA	8.02	11.09	27.92	7.12	10.75	11.33

Table 3. Misclassification Error [%] for the 15 multi-model fundamental matrix sequences from AdelaideRMF [42] using SA or Gurobi (for quantum methods). Results for RanSaCov without and with outliers are taken from [15] and [25], respectively.

Outliers Settings		QuMF[15]	De-QuMF[15]	RanSaCov[27]	RQuMF	De-RQuMF
No Outliers	Mean	54.10	**13.94**	-	20.77	19.11
	Median	55.44	**16.20**	-	25.11	24.96
With Outliers	Mean	86.23	49.32	14.72	17.01	**14.33**
	Median	86.25	46.35	**14.38**	16.72	15.76
With Outliers + Post Processing	Mean	51.22	26.57	-	16.75	**14.21**
	Median	50.76	20.46	-	16.63	**16.05**

Table 4. Misclassification Error [%] for the 16 multi-model homography sequences from AdelaideRMF [42] using SA (for quantum methods). Results for RanSaCov without outliers are not available whereas those with outliers are taken from [25]. De-QuMF with post-processing fails on at least one sequence in 16 out of 20 trials.

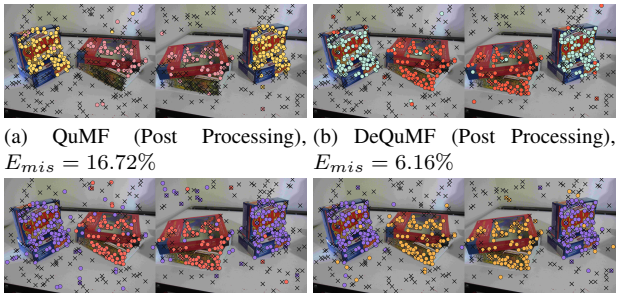

(a) QuMF (Post Processing), $E_{mis} = 16.72\%$ (b) DeQuMF (Post Processing), $E_{mis} = 6.16\%$

(c) RQuMF, $E_{mis} = 5.87\%$ (d) De-RQuMF, $E_{mis} = 5.57\%$

Figure 5. Sample results of fundamental matrix fitting on *biscuitbook* using SA. Our method performs as well as QuMF and De-QuMF which use the information about the number of ground-truth models.

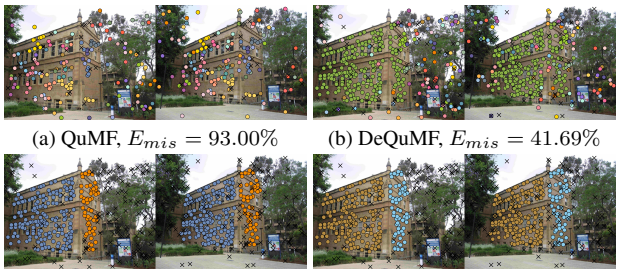

(a) QuMF, $E_{mis} = 93.00\%$ (b) DeQuMF, $E_{mis} = 41.69\%$

(c) RQuMF, $E_{mis} = 2.9\%$ (d) De-RQuMF, $E_{mis} = 0.53\%$

Figure 6. Sample result of homography fitting on *oldclassicswing* (32% outliers) using SA. In the absence of the true number of models both QuMF and De-QuMF fail. Both our proposed methods achieve a near-perfect score.

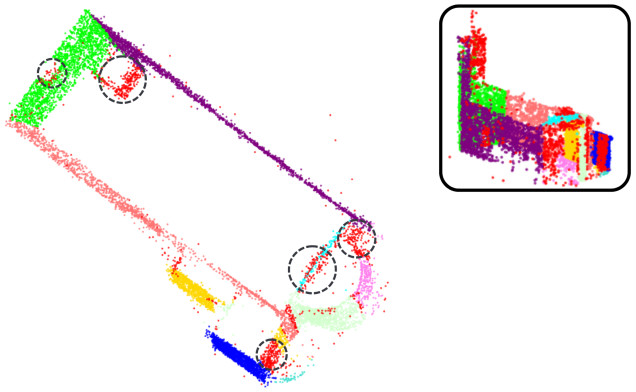

Figure 7. A plane fitting example using the De-RQuMF(SA) method. The encircled dark red points are uncovered and treated as outliers. The inset view on the top-right shows the same result from a different virtual camera perspective.

6. Conclusion

Based on experimental evidence, we conclude that explicitly accounting for outliers in the model significantly lowers the misclassification error across a wide variety of scenarios, compared to all competing quantum-admissible methods. With respect to QuMF [15], the price to pay in RQuMF for outlier robustness is an increased dimensionality of the Q matrix, which is, however, a minor factor in our iterative version and not a major limitation. More importantly, RQuMF does not assume the true number of models explaining the data, which is highly advantageous in practice.

Although the attained results are promising, one of the limitations of our approach is that performance is unpredictable when outliers exceed 50% of the data. While RQuMF already works on real quantum hardware for small problems, we believe its usefulness will increase as the quantum hardware is improving. Managing heterogeneous models (e.g. both planar and cylindrical models) is a promising direction for future extensions.

Acknowledgements. This paper is partially supported by the PNRR-PE-AI FAIR project funded by the NextGeneration EU program and by Geopride (ID: 2022245ZYB CUP: D53D23008370001) funded by PRIN 2022.

References

[1] 3dflow srl. Samantha. https://www.3dflow.net/, 2024. online; accessed on the 01.08.2024. 7

[2] Sameer Agarwal, Jongwoo Lim, Lihi Zelnik-manor, Pietro Perona, David Kriegman, and Serge Belongie. Beyond pairwise clustering. In *CVPR*, pages 838–845, 2005. 2

[3] Federica Arrigoni, Willi Menapace, Marcel Seelbach Benkner, Elisa Ricci, and Vladislav Golyanik. Quantum motion segmentation. In *ECCV*, pages 506–523. Springer, 2022. 1, 3, 5, 6

[4] Daniel Barath and Jiri Matas. Multi-class model fitting by energy minimization and mode-seeking. In *ECCV*, pages 229–245. Springer, 2018. 2

[5] Daniel Barath and Jiri Matas. Progressive-x: Efficient, anytime, multi-model fitting algorithm. In *CVPR*, pages 3780–3788, 2019. 2

[6] Daniel Barath, Denys Rozumny, Ivan Eichhardt, Levente Hajder, and Jiri Matas. Finding geometric models by clustering in the consensus space. In *CVPR*, pages 5414–5424, 2023. 2

[7] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *NIPS*, 2011. 5

[8] Harshil Bhatia, Edith Tretschk, Zorah Lähner, Marcel Benkner, Michael Möller, Christian Theobalt, and Vladislav Golyanik. Ccuantumm: Cycle-consistent quantum-hybrid matching of multiple shapes. In *CVPR*, 2023. 1

[9] Tolga Birdal, Vladislav Golyanik, Christian Theobalt, and Leonidas Guibas. Quantum permutation synchronization. In *CVPR*, 2021. 1, 5, 6

[10] Tat-Jun Chin, Hanzi Wang, and D. Suter. Robust fitting of multiple structures: The statistical learning approach. In *ICCV*, pages 413–420, 2009. 2

[11] Tat-Jun Chin, David Suter, Shin-Fang Ch'ng, and James Q. Quach. Quantum robust fitting. In *ACCV*, 2020. 1, 2, 3

[12] Prasanna Date and Thomas Potok. Adiabatic quantum linear regression. *Scientific Reports*, 11:21905, 2021. 3

[13] Anh-Dzung Doan, Michele Sasdelli, David Suter, and Tat-Jun Chin. A hybrid quantum-classical algorithm for robust fitting. In *CVPR*, pages 417–427, 2022. 1, 2, 3

[14] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science*, 292(5516):472–475, 2001. 1

[15] Matteo Farina, Luca Magri, Willi Menapace, Elisa Ricci, Vladislav Golyanik, and Federica Arrigoni. Quantum multimodel fitting. In *CVPR*, pages 13640–13649, 2023. 1, 2, 3, 4, 5, 6, 7, 8

[16] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. 2, 3

[17] Vladislav Golyanik and Christian Theobalt. A quantum computational approach to correspondence problems on point sets. In *CVPR*, pages 9182–9191, 2020. 1

[18] Venu Madhav Govindu. A tensor decomposition for geometric grouping and segmentation. In *CVPR*, pages 1150–1157, 2005. 2

[19] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. 5

[20] Hossam Isack and Yuri Boykov. Energy-based geometric multi-model fitting. In *IJCV*, pages 123–147, 2012. 2

[21] Suraj Jain and Venu Madhav Govindu. Efficient higher-order clustering on the grassmann manifold. In *ICCV*, pages 3511–3518, 2013. 2

[22] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598): 671–680, 1983. 5

[23] Maximilian Krahn, Michelle Sasdelli, Fengyi Yang, Vladislav Golyanik, Juho Kannala, Tat-Jun Chin, and Tolga Birdal. Projected Stochastic Gradient Descent with Quantum Annealed Binary Gradients. In *BMVC*, 2024. 1

[24] Junde Li and Swaroop Ghosh. Quantum-soft qubo suppression for accurate object detection. In *ECCV*, 2020. 1

[25] Luca Magri. *Multiple Structure Recovery via Preference Analysis in Conceptual Space*. PhD thesis, University of Milan (Italy), 2015. 8

[26] Luca Magri and Andrea Fusiello. T-Linkage: A continuous relaxation of J-Linkage for multi-model fitting. In *CVPR*, pages 3954–3961, 2014. 2

[27] Luca Magri and Andrea Fusiello. Multiple models fitting as a set coverage problem. In *CVPR*, pages 3318–3326, 2016. 2, 4, 5, 6, 8

[28] Luca Magri and Andrea Fusiello. Multiple structure recovery via robust preference analysis. *Image and Vision Computing*, 67:1–15, 2017. 2

[29] Luca Magri, Filippo Leveni, and Giacomo Boracchi. Multilink: Multi-class structure recovery via agglomerative clustering and model selection. In *CVPR*, pages 1853–1862, 2021. 2

[30] Natacha Kuete Meli, Florian Mannel, and Jan Lellmann. An iterative quantum approach for transformation estimation from point sets. In *CVPR*, pages 529–537, 2022. 1

[31] Natacha Kuete Meli, Vladislav Golyanik, Marcel Seelbach Benkner, and Michael Moeller. Qucoop: A versatile framework for solving composite and binary-parametrised problems on quantum annealers. In *CVPR*, 2025. 1

[32] Pulak Purkait, Tat-Jun Chin, Hanno Ackermann, and David Suter. Clustering with hypergraphs: the case for large hyperedges. In *ECCV*, pages 672–687, 2014. 2

[33] Michele Sasdelli and Tat-Jun Chin. Quantum annealing formulation for binary neural networks. In *DICTA*, pages 1–10, 2021. 1

[34] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. Prediction by linear regression on a quantum computer. *Phys. Rev. A*, 94:022342, 2016. 3

[35] Marcel Seelbach Benkner, Vladislav Golyanik, Christian Theobalt, and Michael Moeller. Adiabatic quantum graph matching with permutation matrix constraints. In *3DV*, pages 583–592, 2020. 1

[36] Marcel Seelbach Benkner, Zorah Lähner, Vladislav Golyanik, Christof Wunderlich, Christian Theobalt, and Michael Moeller. Q-match: Iterative shape matching via quantum annealing. In *ICCV*, pages 7586–7596, 2021. 1

[37] Mariano Tepper and Guillermo Sapiro. Nonnegative matrix underapproximation for robust multiple model fitting. In *CVPR*, pages 2059–2067, 2017. 2

[38] Roberto Toldo and Andrea Fusiello. Robust multiple structures estimation with J-Linkage. In *ECCV*, pages 537–547, 2008. 2

[39] Esther Vincent and Robert Laganiere. Detecting planar homographies in an image pair. In *ISPA*, pages 182–187, 2001. 2

[40] Hanzi Wang, Guobao Xiao, Yan Yan, and David Suter. Searching for representative modes on hypergraphs for robust geometric model fitting. pages 697–711. IEEE, 2018. 2

[41] Dennis Willsch, Willsch Madita, Hans De Raedt, and Kristel Michielsen. Support vector machines on the D-Wave quantum annealer. *Computer physics communications*, 248: 107006 –, 2020. 3

[42] Hoi Sim Wong, Tat-Jun Chin, Jin Yu, and David Suter. Dynamic and hierarchical multi-structure geometric model fitting. In *ICCV*, pages 1044–1051, 2011. 5, 7, 8, 2, 3, 6

[43] Lei Xu, Erkki Oja, and Pekka Kultanen. A new curve detection method: randomized Hough transform (RHT). *Pattern recognition letters*, 11(5):331–338, 1990. 2

[44] Frances Fengyi Yang, Michele Sasdelli, and Tat-Jun Chin. Robust fitting on a gate quantum computer. In *ECCV*, 2024. 3

[45] Jan-Nico Zaech, Alexander Liniger, Martin Danelljan, Dengxin Dai, and Luc Van Gool. Adiabatic quantum computing for multi object tracking. In *CVPR*, pages 8811–8822, 2022. 1

[46] Ron Zass and Amnon Shashua. A unifying approach to hard and probabilistic clustering. In *ICCV*, pages 294–301, 2005. 2

[47] Marco Zuliani, Charles S Kenney, and BS Manjunath. The multiRANSAC algorithm and its application to detect planar homographies. In *ICIP*, pages III–153–6, 2005. 2