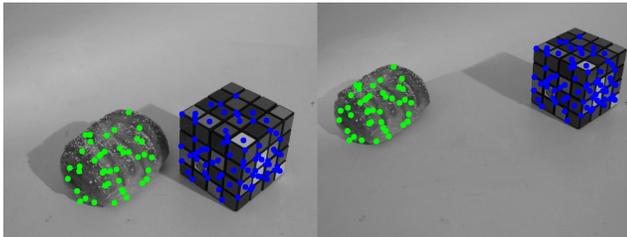


Quantum Multi-Model Fitting

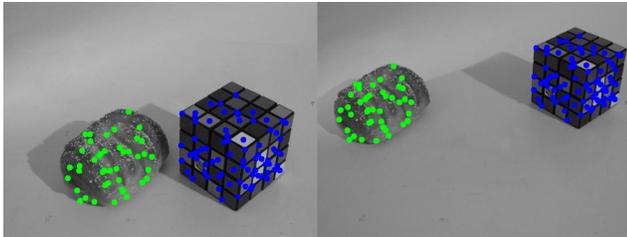
Supplementary Material

Matteo Farina¹ Luca Magri² Willi Menapace¹
 Elisa Ricci^{1,3} Vladislav Golyanik⁴ Federica Arrigoni²

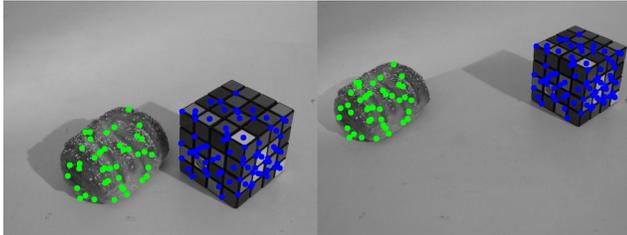
¹University of Trento ²Politecnico di Milano ³Fondazione Bruno Kessler ⁴MPI for Informatics, SIC



(a) DEQUMF outcome, misclassification error = 1.7%.



(b) DEQUMF (SA) outcome, misclassification error = 0%.



(c) Ground-truth segmentation.

Figure 7. A sample of the *best* case for our approach on the *bread-cube* sequence of the *AdelaideRMF* dataset [7]. Point membership is color-coded and ground-truth segmentation is also reported.

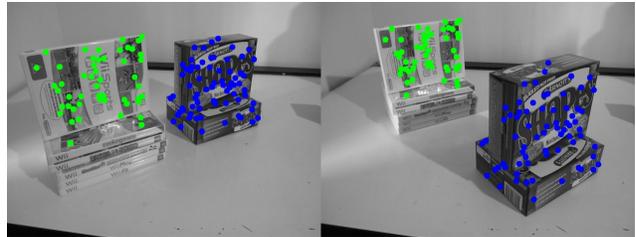
This supplementary document provides the reader with further details on our experiments.

A. Parameters and Implementation Details

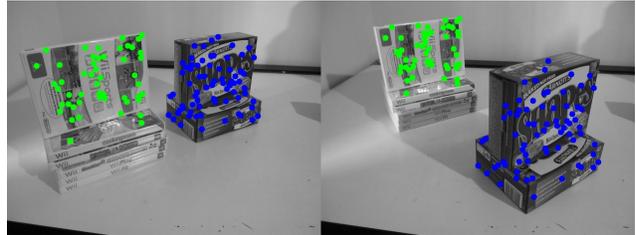
This section highlights the parameter settings of the experiments presented in Sec. 5 of the main paper.



(a) DEQUMF outcome, misclassification error = 47.9%.



(b) DEQUMF (SA) outcome, misclassification error = 1.7%.



(c) Ground-truth segmentation.

Figure 8. A sample of the *worst* case for our approach on the *gamebiscuit* sequence of the *AdelaideRMF* dataset [7]. Point membership is color-coded and ground-truth segmentation is also reported.

The chain strength parameter. Following previous work (e.g., [1]), we employ the *maximum chain length criterion* for *all* our experiments: given a logical graph (i.e., a QUBO problem to be solved), we first minor-embed it to the physical one. Once the final embedding is found, the length l of the longest qubit chain is computed. The chain strength parameter is then set to l plus a small offset, i.e., 0.5 in our experiments.



(a) DEQUMF outcome, misclassification error = 0.2%.



(b) DEQUMF (SA) outcome, misclassification error = 0%.



(c) Ground-truth segmentation.

Figure 9. A sample case for our approach on the *cars2B* sequence of the *Hopkins* benchmark [6]. Point membership is colour-coded and ground-truth segmentation is also reported.

Number of anneals. We set the number of anneals as follows: algorithms solving QUBOs built with the entire preference matrix (“one-sweep” methods QUMF and QUMF (SA)) are executed with 5000 and 100 anneals, respectively. Algorithms employing our iterative decomposition principle are configured differently:

1. DEQUMF executes 2500 anneals for each subproblem, with subproblems containing 20 models each;
2. DEQUMF (SA) runs with 100 anneals for each subproblem, with subproblems having 40 models.

In Sec. 5.3, when evaluating DEQUMF and HQC-RF [3], to provide a fair comparison we fix the amount of available resources per-routine. In this setup, we configured all the optimization routines (SAs/QAs) with 100 anneals, regardless of the algorithm.

Synthetic experiments. To collect statistics, we execute each CPU-based algorithm ten times and report average results. Due to limited QPU availability in our subscription, DEQUMF is run 5 times. In Figs. 3-5 (main paper), the 95% confidence intervals are displayed around data points.

		<i>AdelaideRMF</i> [7]	<i>Traffic2</i> [6]	<i>Traffic3</i> [6]	<i>York</i> [2]
n	<i>mean</i>	160	241	332	119
	<i>min</i>	105	41	123	25
	<i>max</i>	239	511	548	627
m	<i>mean</i>	960	1446	1994	1188
	<i>min</i>	630	246	738	250
	<i>max</i>	1434	3066	3288	6270
k	<i>min</i>	2	2	3	2
	<i>max</i>	4	2	3	3
σ	<i>const</i>	6	6	6	10

Table 4. Details of each real dataset used in the experiments of Sec. 5 in the main paper: n refers to the number of points, k to the number of ground-truth structures, m is the number of models computed as $m = \sigma n$.

To provide a comparison on an equal basis, all the methods have been tested on the *same* preference matrices. In addition, for this test we used a modified version of RANSACOV [4], which enforces disjointedness among the sets of the retrieved solution, as done in the (soft) constraint of QUMF.

Real experiments. In the context of multi-model experiments of Sec. 5 of the main paper (Tables 1-2), for a fair comparison, the inlier threshold ϵ has been tuned per sequence as in [4]. Preference matrices are instantiated with a fixed points-to-models ratio, *i.e.*, the number of models m is always $m = \sigma n$. In our experiments, we select $\sigma = 6$ for *AdelaideRMF* [7] and the *Hopkins* [6] benchmarks, $\sigma = 10$ for the *York* dataset [2]. Recall that the number of models is equal to the number of logical qubits, which essentially determines the problem complexity. Additional insights on the characteristics of the involved datasets are provided in Tab. 4. Following synthetic experiments, CPU-based algorithms are run ten times and averaged results are reported; DEQUMF, when available, is executed once.

Runtimes for real experiments. Runtimes of the experiments executed on real world datasets are reported in Tab. 5 to provide completeness to our experimental evaluation. However, these numbers may report a distorted comparison because of:

1. *implementation differences*, as RANSACOV [4] is MATLAB-based, whilst HQC-RF [3] and all of our methods are Python-based;
2. *hardware differences*, highlighted in the caption of Tab. 5;
3. *overheads* related to AQCs; specifically, we are including *network communication times* since we must access a shared AQC we don’t have in loco, and *resource allocation times* due to other users that would not be present with an in-house machine.

	<i>AdelaideRMF-S</i>	<i>AdelaideRMF-M</i>	<i>Traffic2</i>	<i>Traffic3</i>	<i>York</i>
[†] RANSACov [4]	-	0.48	1.46	2.14	0.24
[†] DEQUMF (SA) (ours)	0.99	0.77	2.54	3.65	1.93
[‡] QUMF (SA) (ours)	23.92	10.72	53.05	76.46	102.38
[‡] DEQUMF (ours)	89.62	116.71	-	376.09	-
[‡] HQC-RF [3]	2349.27	-	-	-	-

Table 5. Mean runtimes [s]. *AdelaideRMF-S* is single-model, with *M* is multi-model. [†] run on an Intel-i7-8575U; [‡] on an Intel-i9-7900X. Please zoom in.

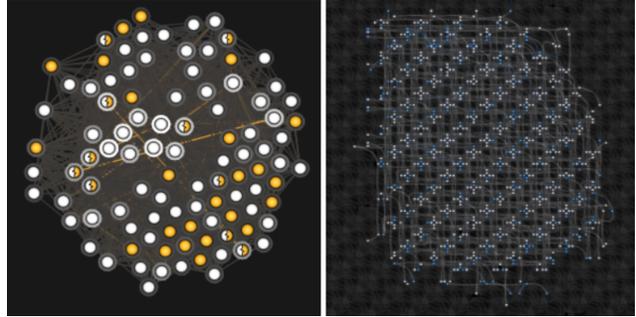
These technicalities currently make the comparison *unfair*. Additionally, we point out that the annealing time ($20\mu s$) is independent on the problem size. Hence, in anticipation of stable AQCs, we believe our quantum approach is a prominent research direction in terms of runtime, too.

Comparison with quantum state of the art. Concerning Tab. 3 of the main paper, we fix the inlier threshold as $\epsilon = 0.045$ for all the sequences for both our method and HQC-RF [3]. Apart from ϵ (that essentially defines a bound on the distance of inlier points to models), when evaluating HQC-RF, we use the source code provided by the authors and the default parameter settings therein.

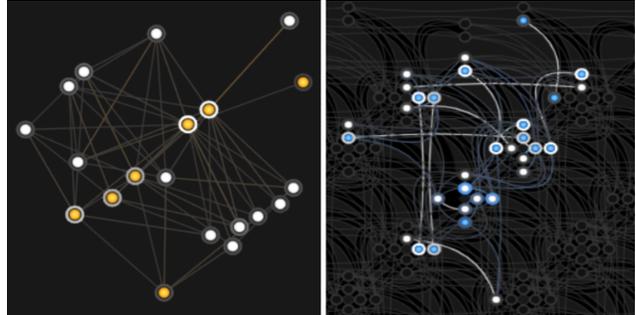
Additional insights on QUMF. Recall that the quadratic term in our QUBO formulation is determined by the product $P^T P$, with P being the preference matrix. Hence, from an interpretative perspective, the quadratic term between models i and j (which we denote by q_{ij} from now on) is equal to the number of points the two have in common. By leveraging this interpretation, a useful observation can be made for a very peculiar situation, namely the case where, for a fixed i , we have $q_{ij} = 0 \forall j \neq i$. This means that model i does not share any point with the remaining models in P : In such a situation, we can conclude that model i is an essential component of the final solution \mathbf{z} , as it covers at least one point that is not explained by any other model. In other terms, $z_i = 1$ must be true at the global optimum. This simple observation thus allows to prune the search space and proceed with embedding a *reduced* (i.e., smaller) QUBO. On the contrary, if the full logical graph was considered, then the minor embedding procedure would avoid mapping model i to the QPU; this would result in unmapped logical qubits, which is undesirable in practice.

B. Qualitative Results

In this section, we provide qualitative results to visually interpret the behaviour of the proposed approaches. Results are given in Figs. 7, 8 and 9. In particular, a failure case is reported in Fig. 8a where DEQUMF exhibits pseudo-random behaviour. Apart from that, in general, both variants of our approach—DEQUMF and DEQUMF (SA)—



(a) Visualization of a full QUBO from *Star5* [5]. The 100 nodes composing the logical graph (left) are mapped to 869 physical qubits (right).



(b) Visualization of a subproblem extracted from *Star5* [5]. The 20 nodes composing the logical graph (left) are mapped to 32 physical qubits (right).

Figure 10. Visual representation of logical and physical graphs when working with the *Star5* dataset [5]. Sub-figure (a) refers to a preference matrix with $n = 250$ points and $m = 100$ models while sub-figure (b) considers a small portion of such preference matrix, corresponding to a sub-problem with $m = 20$ models (as tackled by our decomposed approach).

return accurate results, confirming the outcome of the quantitative analysis reported in the main paper.

C. Minor Embeddings

Finally, we visually analyze the concepts of logical and physical graphs. Specifically, Fig. 10 reports both the logical graph and its mapping to the QPU (obtained as a result of a minor embedding procedure) on exemplary problems with 20 and 100 logical qubits, respectively. Fig. 10 clearly shows that the problem with 100 logical qubits requires a significantly larger amount of physical qubits, as already observed in the main paper.

References

- [1] Federica Arrigoni, Willi Menapace, Marcel Seelbach Benkner, Elisa Ricci, and Vladislav Golyanik. Quantum motion segmentation. In *Proceedings of the European Conference on Computer Vision*, 2022. 1
- [2] Patrick Denis, James H. Elder, and Francisco J. Estrada. Efficient edge-based methods for estimating manhattan frames in

- urban imagery. In *Proceedings of the European Conference on Computer Vision*, pages 197–210, 2008. [2](#)
- [3] Anh-Dzung Doan, Michele Sasdelli, David Suter, and Tat-Jun Chin. A hybrid quantum-classical algorithm for robust fitting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 417–427, June 2022. [2](#), [3](#)
- [4] Luca Magri and Andrea Fusiello. Multiple models fitting as a set coverage problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3318–3326, June 2016. [2](#), [3](#)
- [5] Roberto Toldo and Andrea Fusiello. Robust multiple structures estimation with j-linkage. In *Computer Vision–ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part I 10*, pages 537–547. Springer, 2008. [3](#)
- [6] Roberto Tron and René Vidal. A benchmark for the comparison of 3-d motion segmentation algorithms. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [2](#)
- [7] Hoi Sim Wong, Tat-Jun Chin, Jin Yu, and David Suter. Dynamic and hierarchical multi-structure geometric model fitting. In *2011 International Conference on Computer Vision*, pages 1044–1051. IEEE, 2011. [1](#), [2](#)