

CCuantuMM: Cycle-Consistent Quantum-Hybrid Matching of Multiple Shapes

Harshil Bhatia^{1,2} Edith Tretschk² Zorah Löhner³ Marcel Seelbach Benkner³
 Michael Moeller³ Christian Theobalt² Vladislav Golyanik²

¹Indian Institute of Technology, Jodhpur ²MPI for Informatics, SIC ³Universität Siegen

Abstract

Jointly matching multiple, non-rigidly deformed 3D shapes is a challenging, \mathcal{NP} -hard problem. A perfect matching is necessarily cycle-consistent: Following the pairwise point correspondences along several shapes must end up at the starting vertex of the original shape. Unfortunately, existing quantum shape-matching methods do not support multiple shapes and even less cycle consistency. This paper addresses the open challenges and introduces the first quantum-hybrid approach for 3D shape multi-matching; in addition, it is also cycle-consistent. Its iterative formulation is admissible to modern adiabatic quantum hardware and scales linearly with the total number of input shapes. Both these characteristics are achieved by reducing the N -shape case to a sequence of three-shape matchings, the derivation of which is our main technical contribution. Thanks to quantum annealing, high-quality solutions with low energy are retrieved for the intermediate \mathcal{NP} -hard objectives. On benchmark datasets, the proposed approach significantly outperforms extensions to multi-shape matching of a previous quantum-hybrid two-shape matching method and is on-par with classical multi-matching methods. Our source code is available at 4dqv.mpi-inf.mpg.de/CCuantuMM/.

1. Introduction

Recently, there has been a growing interest in applying quantum computers in computer vision [3, 20, 32]. Such quantum computer vision methods rely on quantum annealing (QA) that allows to solve \mathcal{NP} -hard quadratic unconstrained binary optimisation problems (QUBOs). While having to formulate a problem as a QUBO is rather inflexible, QA is, in the future, widely expected to solve QUBOs at speeds not achievable with classical hardware. Thus, casting a problem as a QUBO promises to outperform more unrestricted formulations in terms of tractable problem sizes and attainable accuracy through sheer speed.

A recent example for such a problem is shape matching, where the goal is to estimate correspondences between



Figure 1. Our quantum-hybrid method matches all 100 shapes of the FAUST collection [4] with guaranteed cycle consistency (white arrows). Here, we visualise the matchings via texture transfer between all shapes. Our method scales linearly in the number of shapes. See the full figure in the supplement.

two shapes. Accurate shape matching is a core element of many computer vision and graphics applications (*i.e.*, texture transfer and statistical shape modelling). If non-rigid deformations are allowed, even pairwise matching is \mathcal{NP} -hard, leading to a wide area of research that approximates this problem, as a recent survey shows [15]. Matching two shapes is one of the problems that was shown to benefit from quantum hardware: Q-Match [39] iteratively updates a subset of point correspondences using QA. Specifically, its cyclic α -expansion allows to parametrise changes to permutation matrices without relaxations.

The question we ask in this work is: How can we design a *multi-shape* matching algorithm in the style of Q-Match that has the same benefits? As we show in the experiments, where we introduce several naïve multi-shape extensions of Q-Match, this is a highly non-trivial task. Despite tweaking them, our proposed method significantly outperforms them.

If $N > 2$ shapes have to be matched, the computational complexity of naïve exhaustive pairwise matching increases quadratically with N , which does not scale to large N . Furthermore, these pairwise matchings can easily turn out to be inconsistent with each other, thereby violating cycle con-

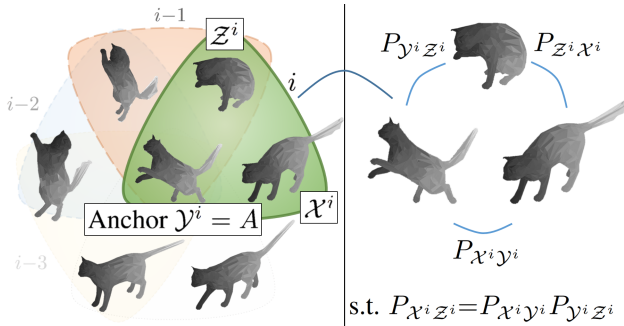


Figure 2. We match N shapes by iteratively matching triplets.

sistency. For example, chaining the matchings $P_{\mathcal{X}\mathcal{Y}}$ from shape \mathcal{X} to \mathcal{Y} and $P_{\mathcal{Y}\mathcal{Z}}$ from \mathcal{Y} to \mathcal{Z} can give very different correspondences between \mathcal{X} and \mathcal{Z} than the direct, pairwise matching $P_{\mathcal{X}\mathcal{Z}}$ of \mathcal{X} and \mathcal{Z} : $P_{\mathcal{X}\mathcal{Z}} \neq P_{\mathcal{X}\mathcal{Y}}P_{\mathcal{Y}\mathcal{Z}}$. (We apply the permutation matrix $P_{\mathcal{X}\mathcal{Y}}$ to the one-hot vertex index vector $x \in \mathcal{X}$ as $x^\top P_{\mathcal{X}\mathcal{Y}} = y \in \mathcal{Y}$.) Thus, how can we achieve cycle consistency by design? A simple solution would be to match a few pairs in the collection to create a spanning tree covering all shapes and infer the remaining correspondences by chaining along the tree. Despite a high accuracy of methods for matching two shapes, this correspondence aggregation policy is prone to error accumulation [36]. A special case of this policy is pairwise matching against a single anchor shape, which also guarantees cycle-consistent solutions by construction [19]. We build on this last option in our method as it avoids error accumulation.

This paper, in contrast to purely classical methods, leverages the advantages of quantum computing for multi-shape matching and introduces a new method for simultaneous alignment of multiple meshes with guaranteed cycle consistency; see Fig. 1. It makes a significant step forward compared to Q-Match and other methods utilising adiabatic quantum computing (AQC), the basis for QA. Our cycle-consistent quantum-hybrid multi-shape matching (CCuantuMM; pronounced “quantum”) approach relies on the computational power of modern quantum hardware. Thus, our main challenge lies in casting our problem in QUBO form, which is necessary for compatibility with AQC. To that end, two design choices are crucial: (1) Our method reduces the N -shapes case to a series of three-shape matchings; see Fig. 2. Thus, CCuantuMM is iterative and hybrid, *i.e.*, it alternates in every iteration between preparing a QUBO problem on the CPU and sampling a QUBO solution on the AQC. (2) It discards negligible higher-order terms, which makes mapping the three-shape objective to quantum hardware possible. In summary, the core technical contributions of this paper are as follows:

- CCuantuMM, *i.e.*, a new quantum-hybrid method for shape multi-matching relying on cyclic α -expansion. CCuantuMM produces cycle-consistent matchings and

scales linearly with the number of shapes N .

- A new formulation of the optimisation objective for the three-shapes case that is mappable to modern QA.
- A new policy in shape multi-matching to address the N -shape case relying on a three-shapes formulation and adaptive choice of an anchor shape.

Our experiments show that CCuantuMM significantly outperforms several variants of the previous quantum-hybrid method Q-Match [39]. It is even competitive with several non-learning-based classical state-of-the-art shape methods [19, 33] and can match more shapes than them. In a broader sense, this paper demonstrates the very high potential of applying (currently available and future) quantum hardware in computer vision.

2. Related Work

Quantum Computer Vision (QCV). Several algorithms for computer vision relying on quantum hardware were proposed over the last three years for such problems as shape matching [20, 32, 39], object tracking [29, 44], fundamental matrix estimation, point triangulation [17] and motion segmentation [1], among others. The majority of them address various types of alignment problems, *i.e.*, transformation estimation [20, 32], point set [20, 34] and mesh alignment [39], graph matching [3, 38] and permutation synchronisation [3].

Only one of them, QSync [3], can operate on more than two inputs and ensure cycle consistency for the underlying matchings. In contrast to QSync, we can align inputs with substantially larger (by two orders of magnitude) shapes in the number of vertices. Furthermore, we address a different problem, *i.e.*, mesh alignment, for which an algorithm for two-mesh alignment with the help of AQC exists, namely Q-Match [39], as we discuss in the introduction.

To maintain the valid structure of permutation matrices, Quantum Graph Matching, QGM [38] and Q-Sync [3] impose linear constraints. However, this requires that the corresponding penalty parameter is carefully chosen. If the parameter is chosen too big and the linear constraints are enforced too strongly, this severely limits QGM and Q-Sync’s ability to handle large sets of vertices. On the other hand, if the linear constraints are enforced too weakly, there is no guarantee to obtain valid permutations as solutions. As discussed in the introduction, our approach follows Q-Match to ensure valid permutation matrices by construction.

Multi-Shape Matching. We focus this section on multi-shape and non-learning methods as CCuantuMM falls in this category. As our approach is not learning-based, it trivially generalises to unknown object categories without a need for training data. For a general survey of recent advances in shape matching, see Sahillioglu [37].

Matching shape pairs is a classical problem in geometry processing [33]. When more than two shapes of the

same class exist, stronger geometric cues can be leveraged to improve results by matching all of them simultaneously. Unfortunately, the already very high problem complexity increases even further the more shapes are used. Hence, existing multi-shape matching methods limit the total number of shapes and their resolution [11, 19], work in spectral space [23], or relax the permutation constraints [25]. Early multi-matching methods computed pair-wise matchings and subsequently used permutation synchronisation to establish cycle consistency [30, 35, 40]. Still, permutation synchronisation requires the eigendecomposition of a matrix with quadratically increasing dimensions [35].

HiPPI [2] is a computationally efficient method that takes geometric relations into account while generalising permutation synchronization but is still limited in resolution. Instead of looking at permutations directly, ZoomOut [33] reduces the dimensionality of the problem by projecting it onto the spectral decomposition. This idea has been extended to take cycle consistency within the spectral space into account [22], which does not guarantee a point-wise consistent matching. To circumvent this issue, IsoMuSh [19] jointly optimises point and functional correspondences. The method detangles the optimisation into smaller subproblems by using a so-called *universe shape* that all shapes are mapped to instead of each other, as Cao and Bernard do [10]. Using a *universe* is similar to requiring a template shape, as many learning-based approaches do [16, 21, 41]: Both synchronise all correspondences by matching them through a unified space. This is similar to the concept of anchor shape we use but inherently less flexible because the universe size or template have to be given *a priori*. Our anchor is chosen from the given collection as part of the method. Although using an anchor slightly improves our results, we note that our method does not necessarily require one for operation. Hence, a random shape could be picked instead in each iteration without an increase in complexity if using an anchor is not feasible or does not represent the shape collection well.

3. Background

3.1. Adiabatic Quantum Computing (AQC)

AQC is a model of computation that leverages quantum effects to obtain high-quality solutions to the \mathcal{NP} -hard problem class of Quadratic Unconstrained Binary Optimisation (QUBO) problems: $\min_{x \in \{0,1\}^k} x^T Q x$, for $k \in \mathbb{N}$ and a QUBO weight matrix $Q \in \mathbb{R}^{k \times k}$. Each entry of x corresponds to its own logical qubit, the quantum equivalent of a classical bit. The diagonal of Q consists of linear terms, while the off-diagonals are inter-qubit *coupling weights*. A QUBO can be classically tackled with simulated annealing (SA) [42] or a variety of other discrete optimisation techniques [18, 28], which, for large k , typically

yield only approximate solutions as QUBOs are in general \mathcal{NP} -hard. AQC holds the potential to systematically outperform classical approaches such as SA, see [14, 27] for an example. AQC exploits the adiabatic theorem of quantum mechanics [5]: If, when starting from an equal superposition state of the qubits (where all solutions $\{0, 1\}^k$ have the same probability of being measured) and imposing external influences corresponding to the QUBO matrix on the qubits sufficiently slowly (called *annealing*), they will end up in a quantum state that, when measured, yields a minimizer x of the QUBO. Not all physical qubits on a real quantum processing unit (QPU) can be connected (coupled) with each other. Thus, a *minor embedding* of the logical-qubit graph (defined by non-zero entries of the QUBO matrix) into the physical-qubit graph (defined by the hardware) is required [9]. This can lead to a chain of multiple physical qubits representing a single logical qubit. For details of quantum annealing on D-Wave machines, we recommend [31].

3.2. Shape Matching

The problem of finding a matching for non-rigidly deformed shapes having n vertices can be formulated as an \mathcal{NP} -hard Quadratic Assignment Problem (QAP) [8, 25]:

$$\min_{P \in \mathbb{P}_n} p^T W p, \quad (1)$$

where $p = \text{vec}(P) \in \{0, 1\}^{n^2}$ is a flattened permutation matrix, and $W \in \mathbb{R}^{n^2 \times n^2}$ is an energy matrix describing how well certain pairwise properties are conserved between two pairs of matches. If two shapes \mathcal{X}, \mathcal{Y} are discretised with n vertices each, W is often chosen as [25]:

$$W_{x_1 \cdot n + y_1, x_2 \cdot n + y_2} = \|d_{\mathcal{X}}^g(x_1, x_2) - d_{\mathcal{Y}}^g(y_1, y_2)\|, \quad (2)$$

where x_1, x_2 are vertices on \mathcal{X} ; y_1, y_2 are vertices on \mathcal{Y} ; and $d_{\mathcal{I}}^g(\cdot, \cdot)$ is the geodesic distance on the shape \mathcal{I} . Therefore, $W_{x_1 \cdot n + y_1, x_2 \cdot n + y_2}$ represents how well the geodesic distance is preserved between corresponding pairs of vertices on the two shapes. Instead of pure geodesics, Gaussian-filtered geodesics are also a popular choice for W [43]:

$$g_{\mathcal{X}}(x_1, x_2) = \frac{1}{\rho \sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{d_{\mathcal{X}}^g(x_1, x_2)}{\rho}\right)^2\right). \quad (3)$$

$g_{\mathcal{I}}$ can be used to directly replace $d_{\mathcal{I}}^g$ in (2). A small value of ρ focuses the energy on a local neighbourhood around the vertex, while a large value increases the receptive field. Using Gaussian kernels in W places more emphasis on local geometry whereas geodesics have higher values far away from the source vertex. Thus, geodesics work well for global alignment and Gaussians for local fine-tuning.

3.3. Cyclic α -Expansion (CAE)

CCuantuMM represents matchings as permutation matrices. In order to update them, we build on Seelbach *et al.*'s CAE algorithm [39] (similar to a fusion move [24]), which we describe here. A permutation matrix P is called an r -cycle, if there exist r disjoint indices i_1, \dots, i_r such that $P_{i_j i_{(j+1)\%r}} = 1$ for all $j \in \{1, \dots, r\}$ and $P_{l,l} = 1$ for all $l \notin \{i_1, \dots, i_r\}$, in which case $P = (i_1 i_2 \dots i_r)$ is a common notation. Two cycles, *i.e.* two permutation matrices, are *disjoint* if these indices are pairwise disjoint. We know that disjoint cycles commute, which allows us to represent any permutation P as $P = \left(\prod_{i=1}^k c_i\right) \left(\prod_{i=1}^l \tilde{c}_i\right)$, where $\{c_i\}_i$ and $\{\tilde{c}_i\}_i$ each are sets of disjoint 2-cycles.

Given a set $\{c_i\}_{i=1}^k$ of k disjoint 2-cycles, an *update*, or modification, of P can therefore be parameterised as: $P(\alpha) = \prod_{i=1}^k c_i^{\alpha_i} P$, where $\alpha \in \{0, 1\}^k$ is a binary decision vector determining the update. (Note that α_i in $c_i^{\alpha_i}$ is an exponent, not an index.) Crucially, to make this parameterisation compatible with QUBOs, we need to make it linear in α . To this end, CAE uses the following equality:

$$P(\alpha) = P + \sum_{i=1}^k \alpha_i (c_i - I)P. \quad (4)$$

4. Our CCuantuMM Method

Previous adiabatic quantum computing methods [38, 39] can only match two shapes. We present a method for matching N shapes. To ensure cycle consistency on N shapes, it is sufficient that all triplets of shapes are matched cycle consistently [22]. CCuantuMM iteratively solves three-shape problems, which preserve cycle consistency by construction and fit on existing quantum annealers with limited resources. We introduce our formulation for matching three shapes in Sec. 4.1 and then extend it to N shapes in Sec. 4.2.

4.1. Matching Three Shapes

Consider the problem of matching three non-rigidly deformed shapes $\mathcal{S} = \{\mathcal{X}, \mathcal{Y}, \mathcal{Z}\}$ of n vertices each, while preserving cycle consistency. We formulate this as an energy minimisation with respect to $\mathcal{P} = \{P_{\mathcal{I}\mathcal{J}} \in \mathbb{P}_n | \mathcal{I}, \mathcal{J} \in \mathcal{S}\}$, the set of permutations between all pairs in \mathcal{S} :

$$\begin{aligned} \min_{\mathcal{P}} \quad & \sum_{\mathcal{I}, \mathcal{J} \in \mathcal{S}; \mathcal{I} \neq \mathcal{J}} \text{vec}(P_{\mathcal{I}\mathcal{J}})^\top W_{\mathcal{I}\mathcal{J}} \text{vec}(P_{\mathcal{I}\mathcal{J}}), \\ \text{s.t.} \quad & P_{\mathcal{X}\mathcal{Z}} = P_{\mathcal{X}\mathcal{Y}} P_{\mathcal{Y}\mathcal{Z}}, \end{aligned} \quad (5)$$

where $W_{\mathcal{I}\mathcal{J}} \in \mathbb{R}^{n^2 \times n^2}$ is the energy matrix describing how well certain pairwise properties are conserved between shapes \mathcal{I} and \mathcal{J} (see Sec. 3.2), and $P_{\mathcal{X}\mathcal{Z}} = P_{\mathcal{X}\mathcal{Y}} P_{\mathcal{Y}\mathcal{Z}}$ enforces cyclic consistency. An overview of the algorithm for three shapes is shown in Alg. 1.

4.1.1 QUBO Derivation

To perform optimisation on the quantum annealer, we need to transform (5) into a QUBO problem. We adapt the CAE formulation from [39] (see Sec. 3.3) and iteratively update the permutations to decrease the value of (5). Given a set $C = \{c_i\}_{i=1}^k$ of k disjoint 2-cycles and binary decision variables α , we can parameterise our permutation matrices as $P_{\mathcal{I}\mathcal{J}}(\alpha) = P_{\mathcal{I}\mathcal{J}} + \sum_{i=1}^k \alpha_i (c_i - I) P_{\mathcal{I}\mathcal{J}}$. However, CAE alone is not sufficient to transform (5) into a QUBO as cyclic consistency is still missing. A simple solution would be to encourage cyclic consistency as a quadratic soft penalty, but then there are no guarantees on the solution. Instead, we enforce cyclic consistency by construction:

$$\begin{aligned} P_{\mathcal{X}\mathcal{Z}}(\alpha, \beta) = \quad & (P_{\mathcal{X}\mathcal{Y}} + \sum_{i=1}^k \alpha_i (c_i - I) P_{\mathcal{X}\mathcal{Y}}) \cdot (P_{\mathcal{Y}\mathcal{Z}} + \sum_{j=1}^k \beta_j (\tilde{c}_j - I) P_{\mathcal{Y}\mathcal{Z}}), \end{aligned} \quad (6)$$

where $\{c_i\}_i$ ($\{\tilde{c}_j\}_j$) are cycles and α (β) are decision variables for the updates to $P_{\mathcal{X}\mathcal{Y}}$ ($P_{\mathcal{Y}\mathcal{Z}}$). For brevity, we write $C_i = (c_i - I) P_{\mathcal{X}\mathcal{Y}}$ and $\tilde{C}_j = (\tilde{c}_j - I) P_{\mathcal{Y}\mathcal{Z}}$. We explain how we construct the cycles in Secs. 4.1.2-4.1.3. Thus, we iteratively solve (5) via a sequence of problems of the form:

$$\min_{\alpha, \beta \in \{0, 1\}^k} E_{\mathcal{X}\mathcal{Y}}(P_{\mathcal{X}\mathcal{Y}}(\alpha)) + E_{\mathcal{Y}\mathcal{Z}}(P_{\mathcal{Y}\mathcal{Z}}(\beta)) + E_{\mathcal{X}\mathcal{Z}}(P_{\mathcal{X}\mathcal{Z}}(\alpha, \beta)), \quad (7)$$

where $E_{\mathcal{I}\mathcal{J}}(P, Q) = \text{vec}(P)^\top W_{\mathcal{I}\mathcal{J}} \text{vec}(Q)$ and $E_{\mathcal{I}\mathcal{J}}(P) = E_{\mathcal{I}\mathcal{J}}(P, P)$. While the first two terms are in QUBO form, the third term contains cubic and bi-quadratic terms (see the supplement for details) which are not compatible with current quantum annealer architectures.

Higher-Order Terms. All of these higher-order terms come from $P_{\mathcal{X}\mathcal{Z}}(\alpha, \beta)$, specifically from the term $H = \sum_i \sum_j \alpha_i \beta_j C_i \tilde{C}_j$. As we only consider 2-cycles, C_i and \tilde{C}_j each have only four non-zero elements. Due to this extreme sparsity, most summands of H become 0.

We could tackle these undesirable terms by decomposing them into quadratic terms by using ancilla variables and adding penalty terms [13]. This gives exact solutions for sufficiently high weights of the penalty terms. However, multiple reasons speak against this: (1) the QUBO matrix is already dense (a clique) under the current formulation (as we will see in (9)) and adding ancilla qubits scales quadratically in k , (2) adding penalties makes the problem harder to solve, and (3) H is sparsely non-zero and in practise we observe no drastic influence on the quality of the solution.

Alternatively, we could assume $H=0$. However, this is unnecessarily strong since (1) H also contributes to quadratic terms ($E(H, \cdot)$), and (2) higher-order terms operating on the same decision variable trivially reduce to

quadratic terms: $\alpha_i \cdot \alpha_i = \alpha_i$ for binary α_i . We thus keep those two types of terms and merely assume all truly cubic and bi-quadratic terms to be zero.

Cycle-Consistent CAE. After eliminating the higher-order terms and ignoring constants from (7), we obtain (with the same colour coding):

$$\begin{aligned}
& \min_{\alpha, \beta} \sum_{i=1}^k \alpha_i \left(F_{\mathcal{X}\mathcal{Y}}(P_{\mathcal{X}\mathcal{Y}}, C_i) + F_{\mathcal{X}\mathcal{Z}}(P_{\mathcal{X}\mathcal{Z}}, C_i P_{\mathcal{Y}\mathcal{Z}}) \right) \\
& + \sum_{j=1}^k \beta_j \left(F_{\mathcal{Y}\mathcal{Z}}(P_{\mathcal{Y}\mathcal{Z}}, \tilde{C}_j) + F_{\mathcal{X}\mathcal{Z}}(P_{\mathcal{X}\mathcal{Z}}, P_{\mathcal{X}\mathcal{Y}} \tilde{C}_j) \right) \\
& + \sum_{i=1}^k \sum_{l=1}^k \alpha_i \alpha_l \left(E_{\mathcal{X}\mathcal{Y}}(C_i, C_l) + E_{\mathcal{X}\mathcal{Z}}(C_i P_{\mathcal{Y}\mathcal{Z}}, C_l P_{\mathcal{Y}\mathcal{Z}}) \right) \\
& + \sum_{j=1}^k \sum_{l=1}^k \beta_j \beta_l \left(E_{\mathcal{Y}\mathcal{Z}}(\tilde{C}_j, \tilde{C}_l) + E_{\mathcal{X}\mathcal{Z}}(P_{\mathcal{X}\mathcal{Y}} \tilde{C}_j, P_{\mathcal{X}\mathcal{Y}} \tilde{C}_l) \right) \\
& + \sum_{i=1}^k \sum_{j=1}^k \alpha_i \beta_j \left(F_{\mathcal{X}\mathcal{Z}}(P_{\mathcal{X}\mathcal{Y}} \tilde{C}_j, C_i P_{\mathcal{Y}\mathcal{Z}}) + F_{\mathcal{X}\mathcal{Z}}(K_{ij}, P_{\mathcal{X}\mathcal{Z}}) \right. \\
& \left. + F_{\mathcal{X}\mathcal{Z}}(K_{ij}, P_{\mathcal{X}\mathcal{Y}} \tilde{C}_j) + F_{\mathcal{X}\mathcal{Z}}(K_{ij}, C_i P_{\mathcal{Y}\mathcal{Z}}) + E_{\mathcal{X}\mathcal{Z}}(K_{ij}, K_{ij}) \right), \tag{8}
\end{aligned}$$

where $P_{\mathcal{X}\mathcal{Z}} = P_{\mathcal{X}\mathcal{Z}}(\mathbf{0}, \mathbf{0}) = P_{\mathcal{X}\mathcal{Y}} P_{\mathcal{Y}\mathcal{Z}}$, and we use the shorthands $F_{\mathcal{I}\mathcal{J}}(A, B) = E_{\mathcal{I}\mathcal{J}}(A, B) + E_{\mathcal{I}\mathcal{J}}(B, A)$ and $K_{ij} = C_i \tilde{C}_j$. Denoting $\alpha_{k+j} = \beta_j$ for an expanded $\alpha \in \{0, 1\}^{2k}$, (8) can be written in the form:

$$\min_{\alpha \in \{0, 1\}^{2k}} \alpha^\top \tilde{W} \alpha. \tag{9}$$

The full formula for \tilde{W} is provided in the supplement. (9) is finally in QUBO form and we can optimise it classically or on real quantum hardware (see Sec. 3.1).

4.1.2 Choosing Vertices

The question of how to choose the sets of cycles $\{c_i\}_i, \{\tilde{c}_j\}_j$ is still open. We first choose a subset of vertices using the ‘‘worst vertices’’ criterion introduced in [39] based on the relative inconsistency $I_{\mathcal{X}\mathcal{Y}}$ of a vertex $x \in \mathcal{X}$ under the current permutation:

$$I_{\mathcal{X}\mathcal{Y}}(x) = \sum_{w \in \mathcal{X}} W_{x \cdot n + x^\top P_{\mathcal{X}\mathcal{Y}}, w \cdot n + w^\top P_{\mathcal{X}\mathcal{Y}}}, \tag{10}$$

where we treat the one-hot vector $x^\top P_{\mathcal{X}\mathcal{Y}}$ as a vertex index on \mathcal{Y} . A high value indicates that x is inconsistent with many other matches under $P_{\mathcal{X}\mathcal{Y}}$ and swapping it will likely improve the matching. We denote the set of the $m=2k$ vertices with the highest $I_{\mathcal{X}\mathcal{Y}}(\cdot)$ as $V_{\mathcal{X}}$. Finally, we follow the permutations to get $V_{\mathcal{Y}} = \{x^\top P_{\mathcal{X}\mathcal{Y}} | x \in V_{\mathcal{X}}\} \subset \mathcal{Y}$.

In practice, we observe a systematic improvement in the matchings when considering all three possibilities (using $I_{\mathcal{X}\mathcal{Y}}, I_{\mathcal{Y}\mathcal{Z}},$ or $I_{\mathcal{X}\mathcal{Z}}$ as the starting point). We thus use three ‘‘sub’’-iterations per iteration, one for each possibility.




Figure 3. We depict the sub-iteration that starts from $I_{\mathcal{X}\mathcal{Y}}$, from which we construct $V_{\mathcal{X}}$, then $C_{\mathcal{X}}^{all}$, and finally $C_{\mathcal{X}} = \{C_{\mathcal{X}}^0, C_{\mathcal{X}}^1, C_{\mathcal{X}}^2\}$. We also build $V_{\mathcal{Y}}$ from $V_{\mathcal{X}}$ and construct $C_{\mathcal{Y}}$ analogously. Matching each element of $C_{\mathcal{X}}$ with one from $C_{\mathcal{Y}}$ (visualised via matching colours) leads to three sub-sub-iterations.




Figure 4. We depict the sub-sub-iteration for $C_{\mathcal{X}}^2 = \{(u_x, v_x), (w_x, t_x)\}$ and $C_{\mathcal{Y}}^1 = \{(u_y, w_y), (v_y, t_y)\}$ from Fig. 3.

4.1.3 Choosing Cycles

Given the worst vertices $V_{\mathcal{X}}$ and $V_{\mathcal{Y}}$ of any sub-iteration, we construct the cycles $\{c_i\}_i, \{\tilde{c}_j\}_j$ from them. Fig. 3 visualises this process. Focusing on $V_{\mathcal{X}}$ for the moment, we want to use all possible 2-cycles $C_{\mathcal{X}}^{all} = \{(uv) | u, v \in V_{\mathcal{X}}, u \neq v\}$ in each sub-iteration. We cannot use all of these cycles at once since they are not disjoint, as CAE requires. Instead, we next construct a set $C_{\mathcal{X}}$ by partitioning $C_{\mathcal{X}}^{all}$ into $m-1$ sets of cycles with each containing $m/2=k$ disjoint cycles. An analogous methodology is used for $C_{\mathcal{Y}}$.

We now have $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$. Since we want to consider each cycle of $C_{\mathcal{X}}^{all}$ and $C_{\mathcal{Y}}^{all}$ once, we need several ‘‘sub-sub’’ iterations. Thus, we next need to pick one set of cycles from each $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$ for each sub-sub-iteration. There are $(m-1)^2$ possible pairs between elements of $C_{\mathcal{X}}$ and $C_{\mathcal{Y}}$. Considering all possible pairs is redundant, does not provide significant performance advantage, and increases the computational complexity quadratically. Hence, we randomly pair each element of $C_{\mathcal{X}}$ with one element of $C_{\mathcal{Y}}$ (without replacement). This leads to $m-1$ sub-sub-iterations, with each one solving (9) with its respective cycles; see Fig. 4.

4.2. Matching N Shapes

In this section, we extend our model to matching a shape collection \mathcal{S} with N elements by iteratively matching three shapes while still guaranteeing cycle consistency. Similar

Algorithm 1 Hybrid Three-Shape Matching

Input: $\mathcal{P}^i, \mathcal{S}$ **Output:** \mathcal{P}^{i+1}

```
1: for  $I \in \{I_{\mathcal{X}\mathcal{Y}}, I_{\mathcal{Y}\mathcal{Z}}, I_{\mathcal{X}\mathcal{Z}}\}$  do ▷ sub-iterations
2:   construct  $V_{\mathcal{X}}, V_{\mathcal{Y}}, V_{\mathcal{Z}}$  (see Sec. 4.1.2)
3:   construct  $\mathcal{C}_{\mathcal{X}}, \mathcal{C}_{\mathcal{Y}}, \mathcal{C}_{\mathcal{Z}}$  (see Sec. 4.1.3)
4:   for  $l=1$  to  $m-1$  do ▷ sub-sub-iterations
5:     compute  $\tilde{W}$ 
6:     optimise QUBO (9) ▷ quantum
7:   end for
8:    $P_{\mathcal{X}\mathcal{Y}} = \prod_{i=1}^k c_i^{\alpha_i} P_{\mathcal{X}\mathcal{Y}}$ 
9:    $P_{\mathcal{Y}\mathcal{Z}} = \prod_{j=1}^k \tilde{c}_j^{\alpha_{m+j}} P_{\mathcal{Y}\mathcal{Z}}$ 
10:   $P_{\mathcal{X}\mathcal{Z}} = P_{\mathcal{X}\mathcal{Y}} \cdot P_{\mathcal{Y}\mathcal{Z}}$ 
11: end for
12: return  $\mathcal{P}^{i+1} = \{P_{\mathcal{X}\mathcal{Y}}, P_{\mathcal{Y}\mathcal{Z}}, P_{\mathcal{X}\mathcal{Z}}\}$ 
```

to the three-shape case (5), this can be formulated as an energy minimisation problem with respect to the set of permutations \mathcal{P} , except \mathcal{S} now has cardinality N :

$$\begin{aligned} \min_{\mathcal{P}} \quad & \sum_{\mathcal{I}, \mathcal{J} \in \mathcal{S}; \mathcal{I} \neq \mathcal{J}} E_{\mathcal{I}\mathcal{J}}(P_{\mathcal{I}\mathcal{J}}), \\ \text{s.t.} \quad & P_{\mathcal{I}\mathcal{K}} = P_{\mathcal{I}\mathcal{J}} P_{\mathcal{J}\mathcal{K}} \quad \forall \mathcal{I}, \mathcal{J}, \mathcal{K} \in \mathcal{S}. \end{aligned} \quad (11)$$

The energy contains summands for each possible pair of shapes. Solving all of them jointly would be computationally expensive and even more complicated than (8). This is the reason most multi-shape matching methods apply relaxations at this point or cannot scale to a large N . However, the cycle-consistency constraints still only span over three shapes; triplets are sufficient for global consistency [22].

We thus iteratively focus on a triplet $\mathcal{X}, \mathcal{Y}, \mathcal{Z} \in \mathcal{S}$ and its set of permutations $\mathcal{P}' = \{P_{\mathcal{X}\mathcal{Y}}, P_{\mathcal{X}\mathcal{Z}}, P_{\mathcal{Y}\mathcal{Z}}\}$. We could then minimise (11) over \mathcal{P}' , leading to a block-coordinate descent optimisation of (11) over \mathcal{P} . This would make the problem tractable on current quantum hardware since it keeps the number of decision variables limited. It would also formally guarantee that our iterative optimisation would never increase the total energy. However, each iteration would be linear in N due to the construction of the QUBO matrix, preventing scaling to large N in practice. We therefore instead restrict (11) to those terms that depend *only* on permutations from \mathcal{P}' . This leads to the same energy as for the three-shape case (5), where the minimisation is now over \mathcal{P}' . Importantly, the computational complexity per triplet becomes independent of N , allowing to scale to large N . While this foregoes the formal guarantee that the *total* energy never increases, we crucially find that it still rarely increases in practice; see the supplement.

By iterating over different triples $\mathcal{X}^i, \mathcal{Y}^i, \mathcal{Z}^i$, we cover the entire energy term and reduce it iteratively. Specifically, one iteration i of the N -shape algorithm runs Alg. 1 on

$\mathcal{P}' = \{P_{\mathcal{X}^i\mathcal{Y}^i}^i, P_{\mathcal{Y}^i\mathcal{Z}^i}^i, P_{\mathcal{X}^i\mathcal{Z}^i}^i\}$. Here, $\mathcal{X}^i \in \mathcal{S}$ is chosen randomly (we use stratified sampling to pick all shapes equally often), the *anchor* $\mathcal{Y}^i = A \in \mathcal{S}$ is fixed, and $\mathcal{Z}^i = \mathcal{X}^{i-1}$. In practice, we saw slightly better results with this scheme instead of choosing the triplet randomly; see the supplement. We note that we only need to explicitly keep track of permutations into the anchor: $\mathcal{P}^i = \{P_{\mathcal{I}A}^i\}_{\mathcal{I} \in \mathcal{S}, \mathcal{I} \neq A}$. We then get \mathcal{P}^{i+1} from \mathcal{P}^i by replacing $P_{\mathcal{X}^i A}^i$ and $P_{\mathcal{Z}^i A}^i$ with their updated versions from Alg. 1.

Initialisation. We compute an initial set of pairwise permutations \mathcal{P}^{init} using a descriptor-based similarity of the normalised heat-kernel-signatures (HKS) [7] extended by a dimension indicating whether a vertex lies on the left or right side of a shape (a standard practice in the shape-matching literature [19]). Instead of using a random shape as anchor, the results improve when using the following shape:

$$A = \arg \min_{A \in \mathcal{S}} \sum_{\mathcal{I} \in \mathcal{S}; \mathcal{I} \neq A} E_{\mathcal{I}A}(P_{\mathcal{I}A}^{init}), \quad (12)$$

where $P_{\mathcal{I}A}^{init} \in \mathcal{P}^{init}$. We thus have $\mathcal{P}^0 = \{P_{\mathcal{I}A}^{init}\}_{\mathcal{I} \in \mathcal{S}, \mathcal{I} \neq A}$. **Time Complexity.** Our algorithm scales linearly with the number of shapes. Each iteration of Alg. 1 has worst-case time complexity $\mathcal{O}(nk^3)$, as we discuss in the supplement.

Energy Matrix Schedule. In practise, we first use pure geodesics for a coarse matching and then Gaussian-filtered geodesics to fine-tune. Specifically, for a shape collection of three shapes, we use a schedule with $2T$ geodesics iterations followed by $2T$ Gaussian iterations. For each additional shape in the shape collection, we add T iterations to both schedules. We exponentially decrease the variance of the Gaussians every $N-1$ iterations to $\rho(i) = c_2 \exp(-\frac{c_1}{i-T})$ where c_1 and c_2 are chosen such that the variance decreases from 25% to 5% of the shape diameter over the iterations. Thus, *all* shapes undergo one iteration with the same specific variance. We refer to the supplement for more details.

5. Experimental Evaluation

We compare against state-of-the-art multi-matching methods with a focus on quantum methods. We consider classical works for reference. All experiments use Python 3.9 on an Intel Core i7-8565U CPU with 8GB RAM and the D-Wave Advantage System 4.1 (accessed via Leap 2). We will release our code, which is accelerated using Numba.

Hyperparameters. We set $T=11$. We set the number of worst vertices m to 16% of the number of vertices n .

Quantum Comparisons. The closest quantum work, Q-Match [39], matches only two shapes. We consider two adaptations to multi-matching: 1) Q-MatchV2-cc, similar to our CCuantuMM, chooses an anchor and matches the other shapes pairwise to it, implicitly enforcing cycle consistency; and 2) Q-MatchV2-nc matches all pairs of shapes directly, without guaranteed cycle consistency. In both cases,




Figure 5. PCK curves for (left) two three-shape and (right) two ten-shape instances using QA and SA. In each plot, we denote one instance by normal lines and the other one by dotted lines.

we use our faster implementation and adapt our energy matrix schedule, which gives significantly better results.

Classical Comparisons. For reference, we also compare against the classical, non-learning-based multi-matching state of the art: IsoMuSh [19] and the *synchronised* version of ZoomOut [33], which both guarantee vertex-wise cycle consistency across multiple shapes.

Evaluation Metric. We evaluate the correspondences using the Princeton benchmark protocol [26]. Given the ground-truth correspondences $P_{\mathcal{I}\mathcal{J}}^*$ for matching the shape \mathcal{I} to \mathcal{J} , the error of vertex $v \in \mathcal{I}$ under our estimated matching $P_{\mathcal{I}\mathcal{J}}$ is given by the normalised geodesic distance:

$$e_v(P_{\mathcal{I}\mathcal{J}}) = \frac{d_{\mathcal{J}}^g(v^\top P_{\mathcal{I}\mathcal{J}}, v^\top P_{\mathcal{I}\mathcal{J}}^*)}{\text{diam}(\mathcal{J})}, \quad (13)$$

where $\text{diam}(\cdot)$ is the shape diameter. We plot the fraction of errors that is below a threshold in a percentage-of-correct-keypoints (PCK) curve, where the threshold varies along the x-axis. As a summary metric, we also report the area-under-the-curve (AUC) of these PCK curves.

Datasets. The FAUST dataset [4] contains real scans of ten humans in different poses. We use the registration subset with ten poses for each class and downsample to 500 vertices. TOSCA [6] has 76 shapes from eight classes of humans and animals. We downsample to ~ 1000 vertices. SMAL [45] has scans of toy animals in arbitrary poses, namely 41 non-isometric shapes from five classes registered to the same template. (E.g., the felidae (cats) class contains scans of lions, cats, and tigers.) We downsample to 1000 vertices. We use the same number of vertices as IsoMuSh [19], except that they use 1000 vertices for FAUST.

5.1. Experiments on Real Quantum Annealer

We run two three-shapes and two ten-shapes experiments with FAUST on a real QPU. However, since our QUBO matrices are dense, we effectively need to embed a clique on the QPU. (The supplement contains a detailed analysis of the minor embeddings and the solution quality.) Hence, we test a reduced version of our method with 20 worst vertices per shape (40 virtual qubits in total), as more would worsen




Figure 6. Qualitative results on the TOSCA [6] cat class. We colour a source shape and transfer this colouring to target shapes via the matches estimated by our method or IsoMuSh [19].

results significantly on current hardware. To compensate for this change, we use more iterations for the ten-shape experiments. We use 200 anneals per QUBO, the default annealing path, and the default annealing time of $20\mu\text{s}$. As standard chain strength, we choose 1.0001 times the largest absolute value of entries in Q . Each ten-shape experiment takes about 10 minutes of QPU time. In total, our results took about 30 minutes of QPU time for a total of $5.5 \cdot 10^4$ QUBOs. QA under these settings achieves a similar performance as SA under the same settings (Fig. 5). As QPU time is expensive and since we have just shown that SA performs comparably to a QPU in terms of result quality, we perform the remaining experiments with SA under our default settings, on classical hardware. This is common practice [1, 39, 44] since SA is conceptually close to QA. For additional results, including results on the new Zephyr hardware [12], we refer to the supplement.

5.2. Comparison to Quantum and Classical SoTA

	Ours	Q-MatchV2-cc	Q-MatchV2-nc	IsoMuSh	ZoomOut	HKS
FAUST	0.989	0.886	0.879	0.974	0.886	0.746
TOSCA	0.967	0.932	0.940	0.952	0.864	0.742
SMAL	0.866	0.771	0.813	0.926	0.851	0.544

Table 1. AUC averaged over all classes of each dataset. For reference, we also include classical methods on the right.

FAUST. We outperform both quantum and classical prior work, as Fig. 7a and Tab. 1 show. Because we downsample FAUST more, IsoMuSh’s results are better in our experiments than what Gao *et al.* [19] report.

Matching 100 Shapes. Next, we demonstrate that, unlike IsoMuSh and ZoomOut, our approach can scale to matching all 100 shapes of FAUST. Fig. 1 contains qualitative results. Tab. 2 compares the runtime of our method (using SA) to others. Only ours and Q-MatchV2-cc scale well to 100 shapes while ZoomOut and IsoMuSh cannot.

TOSCA. Fig. 7b and Tab. 1 show that our method achieves state-of-the-art results. While IsoMuSh’s PCK curve starts higher (better), the AUC in Tab. 1 suggests that our method performs better overall. Fig. 6 has qualitative examples.

SMAL. Our CCuantuMM outperforms the quantum baselines, both in terms of PCK (Fig. 7c) and AUC (Tab. 1).

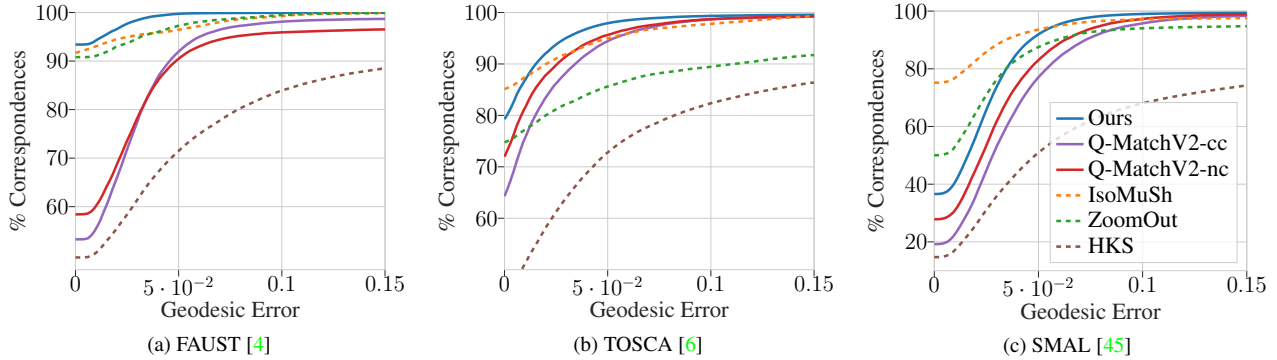


Figure 7. Quantitative results on all three datasets. For each dataset, we match all shapes within a class and then plot the average PCK curve across classes. We plot classical methods with dashed lines as they are only for reference. HKS is our initialisation (see Sec. 4.2).

# Shapes	Ours	Q-MatchV2-cc	Q-MatchV2-nc	IsoMuSh	ZoomOut
10	97	16	81	(4+)0.3	4
100	1137	175	~8000 [†]	OOM	OOM

Table 2. Runtime (in *min*) for FAUST. IsoMuSh uses ZoomOut for initialisation. “OOM” (out of memory): memory requirements are infeasible. “[†]” denotes an estimate.




Figure 8. We ablate (left) the usage of Gaussian kernels, and (right) the large-scale multi-shape setting. Gaussian kernels improve the results greatly. Matching more shapes improves results.

At the same time, it achieves performance on par with ZoomOut and below IsoMuSh. SMAL is considered the most difficult of the three datasets due to the challenging non-isometric deformations of its shapes. All methods thus show worse performance compared to FAUST and TOSCA.

5.3. Ablation Studies

We perform an ablation study on FAUST to analyse how different components of our method affect the quality of the matchings. We refer to the supplement for more ablations.

Gaussian Energy Schedule. Our schedule, which starts with geodesics and afterwards uses Gaussians, provides a significant performance gain over using only geodesics, under the same number of iterations, see Fig 8. That is because Gaussians better correct local errors in our approach.

Does Using More Shapes Improve Results? We analyse what effect increasing the number of shapes N has on the matchings’ quality. We first randomly select three shapes and run our method on them, to obtain the baseline. Next,

we run our method again and again from scratch, each time adding one more shape to the previously used shapes. This isolates the effect of using more shapes from all other factors. In Fig. 8, we plot the PCK curves for the three selected shapes. We repeat this experiment for several randomly sampled instances. Our results show that including more shapes improves the matchings noticeably overall.

5.4. Discussion and Limitations

Our method and all considered methods are based on intrinsic properties like geodesic distances. Thus, without left-right labels for initialisation, they would produce partial flips for inter-class instances in FAUST and intra-class instances in TOSCA and SMAL. For a large worst-vertices set, contemporary quantum hardware leads to embeddings (see Sec. 3.1) with long chains, which are unstable, degrading the result quality. Finally, while our method is currently slower in practice than SA, it would immediately benefit from the widely expected quantum advantage in the future.

6. Conclusion

The proposed method achieves our main goal: improving mesh alignment w.r.t. the quantum state of the art. Furthermore, it is even highly competitive among classical state-of-the-art methods. This suggests that the proposed approach can be used as a reference for comparisons and extensions of classical mesh-alignment works in the future. (For such cases, classical SA is a viable alternative when access to quantum computers is lacking.) Our results show that ignoring certain higher-order terms still allows for high-quality matchings, which is promising for future quantum approaches that could use similar approximations. Finally, unlike classical work, we designed our method within the constraints of contemporary quantum hardware. We found that iteratively considering shape triplets is highly effective, perhaps even for classical methods.

Acknowledgements. This work was partially supported by the ERC Consolidator Grant 4DReply (770784). ZL is funded by the Ministry of Culture and Science of the State of NRW.

References

- [1] Federica Arrigoni, Willi Menapace, Marcel Seelbach Benkner, Elisa Ricci, and Vladislav Golyanik. Quantum motion segmentation. In *Eur. Conf. Comput. Vis.*, 2022. 2, 7
- [2] Florian Bernard, Johan Thunberg, Paul Swoboda, and Christian Theobalt. Hippo: Higher-order projected power iterations for scalable multi-matching. In *Int. Conf. Comput. Vis.*, 2019. 3
- [3] Tolga Birdal, Vladislav Golyanik, Christian Theobalt, and Leonidas Guibas. Quantum permutation synchronization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 1, 2
- [4] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014. 1, 7, 8
- [5] Max Born and Vladimir Fock. Beweis des adiabatsatzes. *Zeitschrift für Physik*, 51(3):165–180, 1928. 3
- [6] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 7, 8
- [7] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2010. 6
- [8] O. Burghard and R. Klein. Efficient lifted relaxations of the quadratic assignment problem. In *Vision, Modeling and Visualization (VMV)*, 2017. 3
- [9] Jun Cai, William G Macready, and Aidan Roy. A practical heuristic for finding graph minors. *arXiv preprint arXiv:1406.2741*, 2014. 3
- [10] Dongliang Cao and Florian Bernard. Unsupervised deep multi-shape matching. In *Eur. Conf. Comput. Vis.*, 2022. 3
- [11] Luca Cosmo, Emanuele Rodola, Andrea Albarelli, Facundo Mémoli, and Daniel Cremers. Consistent partial matching of shape collections via sparse modeling. In *Computer Graphics Forum*, volume 36, pages 209–221. Wiley Online Library, 2017. 3
- [12] D-Wave. Zephyr topology of d-wave quantum processors (d-wave technical report series). https://www.dwavesys.com/media/2uznec4s/14-1056a-a_zephyr_topology_of_d-wave_quantum_processors.pdf, 2021. 7
- [13] Nike Dattani. Quadraticization in discrete optimization and quantum mechanics. *arXiv preprint arXiv:1901.04405*, 2019. 4
- [14] Vasil S Denchev, Sergio Boixo, Sergei V Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the computational value of finite-range tunneling? *Physical Review X*, 6(3):031015, 2016. 3
- [15] Bailin Deng, Yuxin Yao, Roberto M. Dyke, and Juyong Zhang. A survey of non-rigid 3d registration. *Computer Graphics Forum*, 41(2):559–589, 2022. 1
- [16] Theo Deprelle, Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. Learning elementary structures for 3d shape generation and matching. In *Neurips*, 2019. 3
- [17] Anh-Dzung Doan, Michele Sasdelli, David Suter, and Tat-Jun Chin. A hybrid quantum-classical algorithm for robust fitting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2
- [18] Iain Dunning, Swati Gupta, and John Silberholz. What works best when? a systematic evaluation of heuristics for max-cut and qubo. *INFORMS Journal on Computing*, 30(3):608–624, 2018. 3
- [19] Maolin Gao, Zorah Lahner, Johan Thunberg, Daniel Cremers, and Florian Bernard. Isometric multi-shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 2, 3, 6, 7
- [20] Vladislav Golyanik and Christian Theobalt. A quantum computational approach to correspondence problems on point sets. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1, 2
- [21] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. 3d-coded : 3d correspondences by deep deformation. In *Eur. Conf. Comput. Vis.*, 2018. 3
- [22] Qi-Xing Huang and Leonidas Guibas. Consistent shape maps via semidefinite programming. In *Computer graphics forum*, volume 32, pages 177–186, 2013. 3, 4, 6
- [23] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. In *Computer Graphics Forum*, volume 39, pages 265–278, 2020. 3
- [24] Lisa Hutschenreiter, Stefan Haller, Lorenz Feineis, Carsten Rother, Dagmar Kainmüller, and Bogdan Savchynskyy. Fusion moves for graph matching. In *Int. Conf. Comput. Vis.*, 2021. 4
- [25] Itay Kezurer, Shahar Z. Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxation of quadratic matching. In *Symposium on Geometry Processing (SGP)*, 2015. 3
- [26] Vladimir G. Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 2011. 7
- [27] Andrew D King, Jack Raymond, Trevor Lanting, Sergei V Isakov, Masoud Mohseni, Gabriel Poulin-Lamarre, Sara Ejtemaee, William Bernoudy, Isil Ozfidan, Anatoly Yu Smirnov, et al. Scaling advantage over path-integral monte carlo in quantum simulation of geometrically frustrated magnets. *Nature communications*, 12(1):1–6, 2021. 3
- [28] Gary Kochenberger, Jin-Kao Hao, Fred Glover, Mark Lewis, Zhipeng Lü, Haibo Wang, and Yang Wang. The unconstrained binary quadratic programming problem: a survey. *Journal of combinatorial optimization*, 28(1):58–81, 2014. 3
- [29] Junde Li and Swaroop Ghosh. Quantum-soft qubo suppression for accurate object detection. In *Eur. Conf. Comput. Vis.*, 2020. 2
- [30] Eleonora Maset, Federica Arrigoni, and Andrea Fusiello. Practical and efficient multi-view matching. In *Int. Conf. Comput. Vis.*, 2017. 3
- [31] Catherine C. McGeoch. Adiabatic quantum computation and quantum annealing: Theory and practice. *Synthesis Lectures on Quantum Computing*, 5(2):1–93, 2014. 3
- [32] Natacha Kuete Meli, Florian Mannel, and Jan Lellmann. An iterative quantum approach for transformation estimation from point sets. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 1, 2

- [33] Simone Melzi, Jing Ren, Emanuele Rodola, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 2019. [2](#), [3](#), [7](#)
- [34] Mohammadreza Noormandipour and Hanchen Wang. Matching point sets with quantum circuit learning. In *International Conference on Acoustics, Speech and Signal Processing*, 2022. [2](#)
- [35] Deepti Pachauri, Risi Kondor, and Vikas Singh. Solving the multi-way matching problem by permutation synchronization. *Adv. Neural Inform. Process. Syst.*, 2013. [3](#)
- [36] Yusuf Sahillioglu and Yücel Yemez. Multiple shape correspondence by dynamic programming. *Computer Graphics Forum (CGF)*, 33(7), 2014. [2](#)
- [37] Yusuf Sahillioglu. Recent advances in shape correspondence. *The Visual Computer*, 36, 2020. [2](#)
- [38] Marcel Seelbach Benkner, Vladislav Golyanik, Christian Theobalt, and Michael Moeller. Adiabatic quantum graph matching with permutation matrix constraints. In *Int. Conf. 3D Vis. (3DV)*, 2020. [2](#), [4](#)
- [39] Marcel Seelbach Benkner, Zorah Löhner, Vladislav Golyanik, Christof Wunderlich, Christian Theobalt, and Michael Moeller. Q-match: Iterative shape matching via quantum annealing. In *Int. Conf. Comput. Vis.*, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [40] Yanyao Shen, Qixing Huang, Nati Srebro, and Sujay Sanghavi. Normalized spectral map synchronization. *Adv. Neural Inform. Process. Syst.*, 2016. [3](#)
- [41] Ramana Sundararaman, Gautam Pai, and Maks Ovsjanikov. Implicit field supervision for robust non-rigid shape matching. In *Eur. Conf. Comput. Vis.*, 2022. [3](#)
- [42] Peter JM Van Laarhoven and Emile HL Aarts. Simulated annealing. In *Simulated annealing: Theory and applications*, pages 7–15. Springer, 1987. [3](#)
- [43] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodolà, Alex M. Bronstein, Michael M. Bronstein, Ron Kimmel, and Daniel Cremers. Efficient deformable shape correspondence via kernel matching. In *Int. Conf. 3D Vis. (3DV)*, 2017. [3](#)
- [44] Jan-Nico Zaech, Alexander Liniger, Martin Danelljan, Dengxin Dai, and Luc Van Gool. Adiabatic quantum computing for multi object tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. [2](#)
- [45] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. [7](#), [8](#)