

Q-Match: Iterative Shape Matching via Quantum Annealing

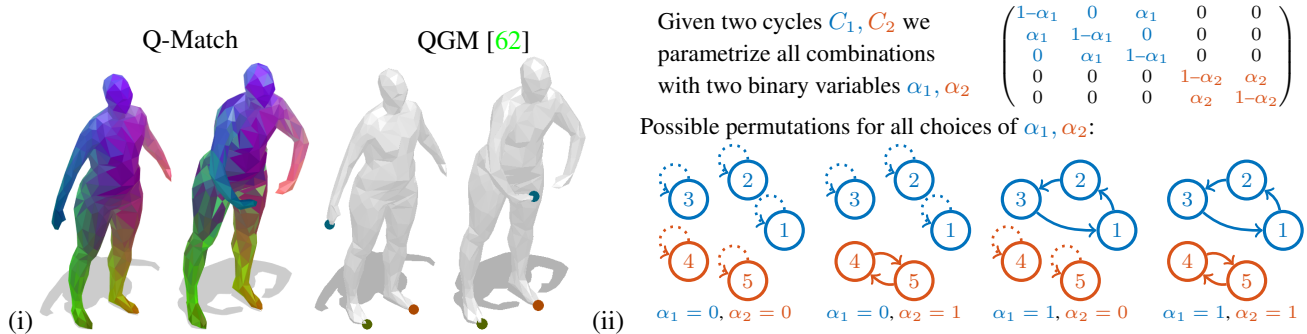
Marcel Seelbach Benkner¹Zorah Löhner¹Vladislav Golyanik²Christof Wunderlich¹Christian Theobalt²Michael Moeller¹¹University of Siegen²MPI for Informatics, SIC

Figure 1. (i) Example correspondences obtained by two quantum approaches, the proposed *Q-Match* (left, 502 points) and QGM [62] (right, three points). Vertices with the same color are mapped to each other. Our iterative method operates directly in the space of permutations and thus can deal with many more vertices compared to QGM. (ii) We represent permutations through a collection of cycles, which allows us to parametrize them through binary variables without the need to enforce any kind of constraints (upper right). *Q-Match* makes the decision whether each cycle C_i should be applied ($\alpha_i = 1$) or not ($\alpha_i = 0$) to progress towards a lower energy solution.

Abstract

Finding shape correspondences can be formulated as an \mathcal{NP} -hard quadratic assignment problem (QAP) that becomes infeasible for shapes with high sampling density. A promising research direction is to tackle such quadratic optimization problems over binary variables with quantum annealing, which allows for some problems a more efficient search in the solution space. Unfortunately, enforcing the linear equality constraints in QAPs via a penalty significantly limits the success probability of such methods on currently available quantum hardware. To address this limitation, this paper proposes *Q-Match*, i.e., a new iterative quantum method for QAPs inspired by the α -expansion algorithm, which allows solving problems of an order of magnitude larger than current quantum methods. It implicitly enforces the QAP constraints by updating the current estimates in a cyclic fashion. Further, *Q-Match* can be applied iteratively, on a subset of well-chosen correspondences, allowing us to scale to real-world problems. Using the latest quantum annealer, the D-Wave Advantage, we evaluate the proposed method on a subset of QAPLIB as well as on isometric shape matching problems from the FAUST dataset.

1. Introduction

Shape correspondence is at the heart of many computer vision and graphics applications, as knowing the relationship between points allows transferring information to new shapes. While this problem has been around for a long time [41], the setting with non-rigidly deformed shapes is still challenging [20]. One reason is that generic matching of n points in two shapes can be formulated as an \mathcal{NP} -hard Quadratic Assignment Problem (QAP)

$$\min_{X \in \mathbb{P}_n} E(X) := \mathbf{x}^T W \mathbf{x}, \quad (1)$$

where $\mathbf{x} = \text{vec}(X)$ is the n^2 -dimensional vector corresponding to a matrix in the set of permutations $\mathbb{P} \subset \{0, 1\}^{n \times n}$ and $W \in \mathbb{R}^{n^2 \times n^2}$ is an energy matrix describing how well certain pairwise properties are preserved between pairs of matches. For some choices of W and assumptions on the input shapes, (1) can become feasible by adding prior knowledge and exploiting geometric structures. For example, rigid shapes can be aligned through rotation and translation only — a problem with only six degrees of freedom. For general cases, QAPs remain challenging nonetheless.

In this work, we propose a new method for shape matching that employs *quantum annealing*, finds high quality solutions of QAPs with high probability, and is suitable for real-world problems. In recent years, quantum computers have made significant progress and machines such as the IBM Q System One (circuit model) and D-Wave Advantage System 1.1 (adiabatic, quantum annealer) became available for research purposes. The circuit-based and adiabatic models are polynomially equivalent to each other in theory [1]. However, the current implementations differ substantially and have different challenges to overcome. One advantage of adiabatic quantum computing is that it is less sensitive to noise and decoherence [16].

Unfortunately, adiabatic quantum computers (AQCer) can a-priori only solve unconstrained problems and thus cannot directly solve QAPs (1) which are constrained to permutation matrices. To overcome this limitation, a recent method [62] adds a penalty term enforcing the solution to be a permutation. However, this reduces the solvable problem sizes to very small $n \leq 4$ and, in practice, the probability of finding the globally optimal solution drops to random guessing. In contrast, we propose a new formulation for AQcing that is guaranteed to result in permutation matrices in problem sizes in the order of magnitude of available (fully connected) qubits. See Fig. 1 for the method overview and Table 1 for a list of acronyms we use. Our method is iterative, and we solve a series of quadratic unconstrained binary optimization (QUBO) problems on D-Wave which were prepared on a classical CPU. In summary, the **contributions** of this paper are as follows:

- Q-Match, *i.e.*, a new scalable quantum approach for searching correspondences between two shapes related by a non-rigid transformation (isometry or near-isometry).
- An iterative problem formulation through a variant of α -expansion which avoids explicit linear constraints for permutations and can be applied to large problem instances.

We evaluate our method on a real AQCer, D-Wave Advantage system 1.1, and obtain improved results compared to the previous quantum method [62]. Moreover, Q-Match *performs on par with the classical matching approach* [52]. This paper does not assume prior knowledge of quantum computing from readers. Secs. 3.1 and 3.2 review preliminaries required to understand, implement and apply the proposed Q-Match approach. The source code is available at <https://4dqv.mpi-inf.mpg.de/QMATCH/>.

2. Related Work

This section reviews prior work on the intersection of quantum computing, computer vision and shape matching. For the foundations of AQcing, see Sec. 3.1 and [30, 22].

AQcing for Computer Vision: Fueled by the progress in quantum hardware accessible for a broad research com-

acronym	meaning
AQCer/ing	adiabatic quantum computer/ing
QPU	quantum processing unit
QAP	quadratic assignment problem
QUBO	quadratic unconstrained binary optimization

Table 1. Commonly used acronyms and their meaning.

munity [17, 29], a growing interest is developing to apply quantum computing or its concepts to computer vision. Ways for representing, retrieving and processing images on a quantum computer have been extensively investigated in the theory literature [68, 14, 64, 71]. Methods for image recognition and classification are among the first low-level techniques evaluated on a real AQCer [46, 9, 48, 15, 39]. O’Malley *et al.* [53] learn facial features and reproduce image collections of human faces with the help of AQcing. Cavallaro *et al.* [15] classify multi-spectral images with an ensemble of quantum support vector machines [70]. To account for the limited connectivity of the physical qubit graph of D-Wave 2000Q, they split the training set into multiple disjoint subsets and train the classifier on each of them independently. Li and Ghosh [39] eliminate false positives in multi-object detection on D-Wave 2X using the QUBO formulation from [58]. It provides state-of-the-art accuracy and the implementation with quantum annealing is faster than greedy and tabu search.

Solving image matching problems with a quantum annealer was first theoretically studied in [47]. Practically applicable quantum algorithms for the absolute orientation problem and point set alignment have been introduced in [24]. The work shows that representing rotation matrices using a linear basis allows mapping both problems to QUBOs. A new point set alignment method for gate model quantum computers has been recently derived in [50], inspired by the earlier kernel correlation approach [67]. QGM [62] is the first implementation of graph matching for small problem instances as a single QUBO with a penalty approach, whereas QuantumSync [6] is the first method for permutation synchronization developed for and tested on an AQCer (Advantage system 1.1). In [62], the formulation of permutation matrix constraints with a linear term leads to low probabilities of measuring globally optimal solutions in a single annealing cycle on a modern AQCer. In contrast to [62, 6], we avoid explicit constraints ensuring valid permutations and use a series of QUBOs leading successive improvements in the energy. Consequently, we can align significantly larger graphs and shapes.

Shape Correspondence: Finding point-wise matches between meshes is an actively studied problem in vision and graphics. One common way to establish matches between shapes related by isometry is to compare shape functions or signatures. Multiple methods of this category ana-

lyze spectra of the Laplace-Beltrami operator on the shape surfaces which are invariant under isometric deformations [65, 3, 52, 59, 27]. Functional maps [52, 27] interpret the problem as the alignment of functions in the pre-defined basis (*e.g.*, eigenfunctions of the Laplace-Beltrami operator). Post-processing is required to extract point-wise matches. Among the shape descriptors, wave kernel signature (WKS) [3] is inspired by quantum physics. It relies on solving the Schrödinger equation for the dynamics (dissipation) of quantum-mechanical particles on the shape surface. WKS can resolve fine details and is robust to moderate non-isometric perturbations. Recently, convolution operators were generalized to non-Euclidean structures, and the availability of large-scale shape collections enabled supervised learning of dense shape correspondences [42, 45, 40].

The point-wise correspondence search between two shapes can be formulated as a linear (LAP) or quadratic assignment problem (QAP) over the space of permutations, with the matching costs computed relying on feature descriptors [65, 3, 59]. The solution space of both problems is exponential in the number of points. For LAP, there exists a fast auction algorithm with polynomial runtime [5]. Per construction, LAP formulations for matching are sensitive to surface noise and do not explicitly consider spatial point relations, which often leads to local minima (and, consequently, inaccurate solutions). QAPs [34, 35] add quadratic costs for matching point pairs and regard point neighborhoods; the solutions are spatially smoother. The downside is their \mathcal{NP} -hardness, which makes finding global optima for large inputs unfeasible. Multiple policies to solve QAPs efficiently have been proposed in the literature, such as branch-and-bound [57, 26], spectral analysis [38], alternating direction method of multipliers [36], entropic regularization of the energy landscape [63] and simulated annealing [28]. Convex relaxations are among the most thoroughly investigated policies for QAPs [72, 2, 54, 23, 31, 4]. These methods either have prohibitive worst-case runtime complexity (branch-and-bound), rely on heuristic algorithmic choices (*e.g.*, entropic regularization and simulated annealing), or do not guarantee global optima. In contrast, we address QAPs with a new AQcing metaheuristic to find high quality solutions with a high probability.

Holzschuh *et al.* [28] proposed a non-quantum algorithm closely related to ours. It can be seen a special case of our setup that only applies one 2-cycle per iteration. This means a considerably smaller step size compared to Q-Match and, hence, more iterations until convergence and higher sensitivity to local optima. Several heuristics to gain robustness and speed are further used in [28]. First, the algorithm is based on simulated annealing and accepts worse states with a certain probability to step over local optima. In contrast, our Q-Match explores larger solution spaces simultaneously and converges in fewer iterations. Second, their hierarchical

optimization scheme adds new matches based on geodesic embeddings whereas we can refine any given initialization.

3. Background

3.1. Adiabatic Quantum Annealing (AQcing)

The seminal work of Kadowaki and Nishimori [30] introduced quantum annealing. The authors argued that a quantum computer could be built based on the principle of finding the ground state of the Ising model under quantum fluctuations (causing state transitions) induced by a transverse magnetic field. The theoretical foundation of such a computational machine is grounded on the adiabatic theorem [8]. In follow-up work, Farhi *et al.* [22] tested the quantum annealing algorithm simulated on a classical computer on hard instances of an \mathcal{NP} -complete problem. Although it is *not* believed that quantum annealing can solve \mathcal{NP} -complete problems in polynomial time, for rugged energy landscapes with high and narrow spikes quantum annealing can be faster than simulated annealing [18, 19] and for some instances D-Wave scales better than the classical path integral Monte Carlo method [33].

AQcing can solve *quadratic unconstrained binary optimization* (QUBO) problems which read

$$\min_{s \in \{-1, 1\}^n} s^\top J s + b^\top s, \quad (2)$$

where s is a binary vector of unknowns, J is a symmetric matrix of inter-qubit couplings, and b contains qubit biases. To change the binary variable to $x \in \{0, 1\}^n$ one simply needs to insert $s_i = 2x_i - 1$ everywhere in (2).

Quantum computers operate with *qubits*, *i.e.*, quantum mechanical systems which can be modeled with normalized vectors $|\phi\rangle$ in a Hilbert space \mathbb{H} . One writes $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$ with $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$, where $|0\rangle, |1\rangle$ denote an orthonormal basis. In contrast to classical bits, which are either in state 0 or 1, quantum mechanical systems can be in a superposition $\alpha|0\rangle + \beta|1\rangle$, where after measurement in the so-called computational basis ($\{|0\rangle, |1\rangle\}$) one obtains $|0\rangle$ with probability $|\alpha|^2$ and $|1\rangle$ with probability $|\beta|^2$. A composite system of, *e.g.*, two qubits can be described with vectors from the tensor product of the individual Hilbert spaces $|\psi\rangle \in \mathbb{H} \otimes \mathbb{H}$. If two qubits can be described independently from each other, we have $|\phi\rangle \in \mathbb{H}$ and $|\eta\rangle \in \mathbb{H}$, and the two-qubit state $|\psi\rangle$ is described by the product state $|\psi\rangle = |\phi\rangle \otimes |\eta\rangle$. Since the state space of n qubits is 2^n dimensional, a classical computer would have problems even storing all the coefficients for $n = 50$ [49]. A second ingredient required for quantum computing, in addition to the superposition of qubit states, is the interference between different possible computational paths in the 2^n -dimensional state space caused by interactions between qubits. Thus, during the course of a quantum

algorithm *entangled* (non-classically correlated) states such as $\frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle)$ are created that cannot be decomposed in the above form. Entanglement between quantum states is often considered a necessary resource for quantum computing [55]. At the end of a quantum algorithm the coefficients α and β for each qubit are determined in a suitable measurement leaving the (adiabatic or circuit-model) quantum computer in a classical (non-entangled) state.

Another central notion in quantum mechanics are linear *Hamilton operators*, acting on elements of \mathbb{H} , which are used to describe the static and dynamic properties of a quantum system with the Schrödinger equation. The eigenstates of a (time-independent) Hamiltonian are the stationary energy states of a quantum system. The eigenvalues give the corresponding possible values of the system’s energy. The key idea for solving problems like (2) is to prepare a quantum system in a known state of lowest energy of a Hamiltonian H_I , most commonly a product state of the form,

$$|\psi(t=0)\rangle = \bigotimes_{i=1}^n \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle), \quad (3)$$

with “ \bigotimes ” denoting the Kronecker product over n subsystems, and each qubit is prepared in a superposition. Next, one constructs a *problem Hamiltonian* H_P in such a way that the lowest energy state of H_P is a solution of (2). Finally, one slowly switches from the initial Hamiltonian H_I to the final Hamiltonian H_P with, e.g., a linear combination

$$H(t) = \left(1 - \frac{t}{\tau}\right) H_I + \frac{t}{\tau} H_P. \quad (4)$$

If the transition in (4) is sufficiently slow (i.e., if τ is sufficiently large), and if the state of lowest energy (ground state) of $H(t)$ remains unique, the adiabatic theorem [8] guarantees that the quantum system will remain in the ground state for all t . Thus, the solution to (2) can simply be *measured* after the time evolution starting in an eigenstate of H_I (that is typically easy to prepare) and ending in an eigenstate of H_P . The theoretical requirements for adiabatic evolution can be made precise in terms of the size of the *spectral gap* of $H(t)$, i.e., the difference between the smallest and second smallest eigenvalues of $H(t)$, during the time evolution. Building a system to actually prepare a quantum system and evolve it in an adiabatic way in practice, however, remains highly challenging. A spectral gap that is too small will lead to excitation of higher energy states during the annealing process, thus preventing the system from ending up in the desired ground state of H_P . Next, we discuss some details regarding such practical realizations.

3.2. Algorithm Design and Programming the D-Wave Quantum Annealer

Every AQCing algorithm includes six steps, i.e., QUBO preparation, minor embedding, quantum annealing (sam-

pling), unembedding, bitstring selection (in the case of multiple annealings) and solution interpretation.

QUBO preparation: Note that many computer vision problems are naturally formulated in forms other than a QUBO. A lot of research thus focuses on the QUBO preparation, i.e., how to formulate a target problem as a QUBO mappable to an AQCer (Sec. 4 is devoted entirely to the QUBO preparation for Q-Match). Note that this step is performed entirely on the CPU. In QUBOs, each binary variable is interpreted as one *logical* qubit in the quantum context. Thus, the biases b and couplings between the qubits J define a graph of logical problem qubits.

Minor embedding is the mapping of logical qubits to the grid of physical hardware qubits of the AQCer. Due to the limited connectivity between the physical qubits, each logical qubit often requires several physical qubits in the minor embedding. Physical qubits realizing a single logical qubit—and having as similar quantum states during annealing as possible—build a *chain*. To find a minor embedding for two arbitrary graphs is an \mathcal{NP} -hard optimization problem on its own [13]. However the problem becomes easier since the one graph is fixed by the hardware and heuristic methods such as [13] can be used. The core criteria for an optimal minor embedding is jointly minimizing the chain length and the total number of required physical qubits.

Quantum annealing is initiated, once the minor embedding is complete. It corresponds to the evaluation of the system with respect to a time-dependent Hamiltonian such as (4) during which the quantum system ideally remains in the ground state during the entire annealing, see Sec. 3.1. In practice, due to such factors as 1) the expected qubits lifetime (e.g., influenced by interaction with the environment, which cannot be entirely prevented), and 2) a small spectral gap of $H(t)$, multiple anneals are needed, each of which leads to the global optimum with some probability <1 . After each annealing, an **unembedding** algorithm assigns measured values to the logical qubits.

The final solution over multiple annealings is chosen in the **bitstring selection** step based on the occurrence frequency or the resulting energies. Finally, the solution is **interpreted** in the context of the original problem and returned to the user. In this regard, the interpretation involves transformation of the measured bitstring to the initial data modality (e.g., permutation matrices).

D-Wave programming is done in Python, using the *Leap 2* programming interface and remote access tools [17]. The D-Wave machine uses superconducting flux qubits [43]. The qubits and the couplers are realized with loops of niobium, which




Figure 2. Superconducting flux qubit.

have to be cooled to temperatures below 17mK. The current, that flows clockwise or anti-clockwise through the niobium loop can be modeled as a qubit, see Fig. 2. Thus, it is possible to have a superposition between a current flowing clockwise and anticlockwise as long as no measurement is performed. The specified couplings J and biases b in (2) determine the time evolution of the magnetic fluxes during quantum annealing.

3.3. α -Expansion

The α -expansion algorithm was introduced in [10] to solve labeling optimization problems (e.g., for semantic segmentation or disparity estimation). A notable extension is [37], where the method is generalized to allow for parallel computation amongst other things. In each iteration, a vector of binary variables α that correspond to two proposed values for the labels is computed so that the decision between the proposals improves the energy. [10] finds the optimal expansion through a graph cut such that each pixel in an image is assigned to the first proposed label if $\alpha_i = 1$ or the second proposed label if $\alpha = 0$.

Apart from this application, α -expansion works particularly well in case the subproblem, i.e., the optimization for α , is submodular. Then the subproblem can be solved efficiently using graph cut techniques.

3.4. Cyclic Permutations

We define the set of permutation matrices as $\mathbb{P}_n = \{X \in \{0, 1\}^{n \times n} \mid \sum_i X_{ij} = 1, \sum_j X_{ij} = 1 \forall i, j\}$. A permutation matrix X is called a k -cycle, if there exist k many disjoint indices i_1, \dots, i_k such that $X_{i_1 i_k} = 1$ and $X_{i_{j+1} i_j} = 1$ for all $j = 1, \dots, k-1$, and $X_{ll} = 1$ for all $l \notin \{i_1, \dots, i_k\}$, in which case $X = (i_1 \ i_2 \ \dots \ i_k)$ is a common notation. Two permutations X and X' are called *disjoint* if $X_{ii} \neq 1$ implies $X'_{ii} = 1$ and $X'_{ii} \neq 1$ implies $X_{ii} = 1$. It is easy to show that disjoint permutations commute. Furthermore, it holds that any X can be written as the product of 2-cycles c_i , i.e., $X = \prod_{i=0}^N c_i$. Finally, we use the notation $X^1 = X$ and $X^0 = I$ where I is the identity for any matrix X .

4. Method

The main difficulty of previous quantum computing methods such as [62] for solving (1) is handling the linear equality constraints arising from the optimization over permutation matrices. Instead of enforcing them via a penalty, the core idea of Q-Match is to minimize (1) iteratively: We propose to use a variant of the α -expansion algorithm that chooses k -many disjoint candidate cycles c_i and uses quantum computing to solve the (still \mathcal{NP} -hard but smaller and unconstrained) subproblem of deciding whether to apply c_i or not (identified with $\alpha_i = 1$ and $\alpha_i = 0$, respectively). In the following, we detail the idea of the cyclic α -expansion as well as our proposed Q-Match algorithm.

4.1. Cyclic α -Expansion

Let $C = \{c_1, \dots, c_m\}$ be a set of m disjoint cycles. We consider the following optimization problem for an initial permutation P_0 :

$$\arg \min_{\{P \in \mathbb{P}_n \mid \exists \alpha \in \{0,1\}^m: P = (\prod_i c_i^{\alpha_i}) P_0\}} E(P), \quad (5)$$

where E is defined in (1) and α is a binary vector parametrizing P . As noted in Sec. 3.4, the order in the product of all $c_i^{\alpha_i}$ does not matter as disjoint permutations commute. (5) would be equivalent to (1) if C was not fixed with disjoint cycles but could also be optimized.

Problem complexity: The optimization problem in (5) makes a binary decision for each cycle whether it should be applied. In this section, we show that (5) is an optimization problem with complexity dependent on the number of cycles and not n . We can convert the multiplication of cycles from (5) into the following linear combination

$$P(\alpha) = P_0 + \sum_{i=1}^m \alpha_i (c_i - I) P_0, \quad (6)$$

where I is the identity. See the supplement for a proof. Next, we use this representation to show that (5) is a problem in the form of (2) with size m . Let P, Q be two permutations and $E(Q, R) = \text{vec}(Q)^T W \text{vec}(R)$ which is linear in each component. We write $C_i = (c_i - I) P_0$ to simplify the notation. Thus, to solve (5), we need to minimize

$$\begin{aligned} & E(P_0 + \sum_i \alpha_i C_i, P_0 + \sum_j \alpha_j C_j) \\ &= E(P_0, P_0 + \sum_j \alpha_j C_j) + \sum_i \alpha_i E(C_i, P_0 + \sum_j \alpha_j C_j) \\ &= E(P_0, P_0) + \sum_j \alpha_j E(P_0, C_j) + \sum_i \alpha_i E(C_i, P_0) \\ &+ \sum_i \sum_j \alpha_j \alpha_i E(C_i, C_j) = E(P_0, P_0) + \alpha^T \tilde{W} \alpha, \end{aligned}$$

with

$$\tilde{W}_{ij} = \begin{cases} E(C_i, C_j) & \text{if } i \neq j, \\ E(C_i, C_i) + E(C_i, P_0) + E(P_0, C_j) & \text{otherwise.} \end{cases} \quad (7)$$

As $E(P_0, P_0)$ is constant w.r.t. α , we are left with

$$\min_{\alpha \in \{0,1\}^m} \alpha^T \tilde{W} \alpha, \quad (8)$$

which can be solved directly by AQCs. The interpretation of the optimal α is a binary indicator for whether to apply the corresponding cycle or not. With the basic methodology of (8) being established, we proceed iteratively: Given a current estimate P_{i-1} minimizing (1), we choose a set of disjoint cycles $C = \{c_1, c_2, \dots, c_m\}$, optimize (8), and set $P_i = \left(\prod_j c_j^{\alpha_j}\right) P_{i-1}$.

Algorithm 1: Q-Match

Input: Shapes M, N
Output: Correspondences P

- 1 Initialize P (Sec. 4.2.3)
- 2 **repeat**
- 3 Calculate the k worst vertices I_M, I_N (10)
- 4 Construct sub-matrix of worst matches W_s (S. 4.2.2)
- 5 **repeat**
- 6 Choose random set of new 2-cycles
- 7 Calculate \tilde{W} (Sec. 4.2.2 and (7))
- 8 Solve the QUBO (8) with AQCing
- 9 Update permutation according to (6)
- 10 **until** Every 2-cycle occurred in one set
- 11 Apply permutation to worst vertices in P
- 12 **until** P does not change anymore
- 13 **return** P

4.2. Q-Match

Cyclic α -expansion can optimize any QAP up to a problem size that fits within the AQC, *e.g.*, by iterating over random sets of cycles. To solve larger problems of (isometric) shape correspondence between two 3D shapes M and N , we propose *Q-Match*. If both shapes are discretized with the same mesh of n vertices, W has the following form:

$$W_{i \cdot n + k, j \cdot n + l} = |d_M^g(i, j) - d_N^g(k, l)| \quad (9)$$

where i, j are vertices on M , k, l are vertices on N , and $d^g(a, b)$ is the geodesic distance between two points on the same shape. Therefore, for two matches $(i, k), (j, l)$ $W_{i \cdot n + k, j \cdot n + l}$ describes how well the geodesic distance is preserved between the elements of this pair. The optimal solution is the correspondence for which the distance between all pairs of points is preserved the best, for isometries the optimal energy of (1) is zero. Since the geodesic distances have geometric meaning, this propagates into the QAP, and Q-Match leverages this to choose better sets of cycles. In Q-Match, the set of tested cycles is based on the contribution of each vertex to (1) (Sec. 4.2.1), we can solve subproblems iteratively (Sec. 4.2.2), and the initial permutation is descriptor-based (Sec. 4.2.3). An overview of the entire Q-Match algorithm can be found in Alg. 1.

4.2.1 Choosing Cycles

One option to select C is by randomly drawing disjoint cycles, which might, however, take many iterations to reach the optimum. Instead we make use of the fact that the isometric shape matching problem creates QAPs in which the entries are highly correlated.

Worst Vertices: For a point with index v in M and a permutation P , we can quantify the influence of v on (1) via

$$I(v) = \sum_{w \in M} W_{v \cdot n + P(v), w \cdot n + P(w)}, \quad (10)$$

where $P(v)$ denotes the index v is mapped to. For vertices on N we proceed equivalently. If $I(v)$ is high, this is an indicator for $P(v)$ being inconsistent with the majority of other matches in P . Therefore, we collect the m vertices with highest values for I on each shape in the two sets I_M and I_N . The potential improvement of (1) is maximized by this choice. See Sec. 5.3 for an analysis of the size of m .

2-Cycles: Subsequently, we choose a set of k random but disjoint 2-cycles from I_M, I_N . While the limitation to 2-cycles might sound restrictive, the following lemma (proven in the appendix and in many abstract algebra textbooks [60]) highlights the expressive power of disjoint 2-cycles:

Lemma 4.1. *Every permutation P can be written as $P = QR$, where Q and R are products of disjoint 2-cycles.*

4.2.2 Subproblems

3D meshes normally consist of more than thousand vertices, this means $W \in \mathbb{R}^{n^2 \times n^2}$ cannot be precomputed and stored in memory. Therefore, in each iteration all needed entries of W to compute \tilde{W} for (8) have to be computed from scratch. As can be seen in (7) several evaluations of the cost function E are needed for \tilde{W} which makes this operation expensive.

To be more efficient, we split the computation into the calculation of W_s , a $k^2 \times k^2$ reduction of W based on the k worst vertices in I_M, I_N . This is the most expensive operation in our algorithm (see Fig. 6). For a set of 2-cycles on I_M, I_N , computing \tilde{W} is now efficient. To reduce the number of times we have to compute W_s , instead of just choosing one set of 2-cycles on each I_M, I_N , we evaluate several sets. In practice, we sample sets of 2-cycles on I_M, I_N until every possible 2-cycle has been chosen at least once. A detailed description how to compute W_s and \tilde{W}_s and a runtime analysis can be found in the supplementary material.

4.2.3 Initialization

The only thing left is choosing P_0 for the first iteration. Instead of a random initialization, we calculate a descriptor-based similarity $S \in \mathbb{R}^{N \times N}$ between all vertices of M and N , equivalent to the strategy of [69]. S_{mn} contains the inner product of the normalized HKS [65] and SHOT [66] descriptors of vertices $m \in M, n \in N$, and we solve a linear assignment problem on S to get the first permutation.




Figure 3. Experiments on how often Q-Match retrieves the optimum for (1). (Left) Fraction of random instances with increasing problem size for which our iterative method leads to the optimum with brute-force calculation. (Right) Success rate of Q-Match and QGM [62] on 20 random problems. We count (QGM) or lower bound (ours) the percentage of anneals that reach the optimum.

5. Experiments

We evaluate accuracy, convergence, and runtime of our method on random QAP instances (Sec. 5.1), FAUST (Sec. 5.2) and QAPLIB (Sec. 5.5). Since quantum computing time is still extremely expensive, we only evaluate on subsets of FAUST and QAPLIB, but our results show how promising this technology can be in the future. All experiments were done on Intel i5 8265U CPU with 16GB RAM using Python 3.8 and D-Wave Advantage system 1.1 accessed by Leap 2. Note that one can use classical QUBO solvers for the subproblems. In the appendix, we present experiments for this with a simulated annealing sampler.

5.1. Comparison to Penalty Based Implementation

We compare Q-Match to the inserted method in [62] for problems of the form (9) with distances $d_{i,j}$ drawn uniformly at random from $[0, 1]$, and W being constructed using a randomly drawn true permutation P . We compare the success probability of both methods on 20 random instances, see Fig. 3-(right). Q-Match finds the optimum in all but 4 cases. For a fair comparison, we restrict the number of anneals per iteration in our method to 10 and use 500 anneals for the inserted method. Additionally, the success probability of the iterative method is bounded by the success probabilities of the individual steps multiplied with each other. In the case that multiple distinct outputs have the lowest energy, both results will be counted as success. For $n \in \{2, 4, 6, 8, 10\}$, we count how often the optimal permutation is reached, if each subproblem is solved to global optimality. For this we generated 20 random instances for $n = 10$ and 100 random instances for the other dimensions. The errors are calculated as a binomial proportion confidence interval that should contain the true probability in 95% of the cases. This experiment was performed first so that all 2-cycles occur only once and a second time where the whole process is carried out twice, see Fig. 3-(left).




Figure 4. Evaluation of cumulative error [32] (left) and convergence (right) on FAUST. (Left) We compare against Simulated Annealing [28] without post-processing and Functional Maps [52]. Dashed lines indicate non-quantum approaches. The results have symmetry-flipped solutions removed, these have a comparable final energy for all methods but are not recognized as correct in the evaluation. (Right) We show convergence of the energy over 30 iterations. The larger the set of worst vertices, the faster our method converges. The dashed grey line shows the optimal energy.




Figure 5. Example correspondences from the FAUST registrations. (i) Our results are overall correct with some outliers. (ii) A failure case. The solution partially swapped front and back. These kinds of solutions are near-isometries, have energies close to the optimum, and often occur in all methods that purely solve (1).

Moreover, we study if the solution with the lowest value returned by the quantum annealer is globally optimal. For 500 runs and coupling matrices up to a size of 13 this was always the case. Therefore, the number of optima found by Q-Match is identical to the probabilities in Fig. 3-(left).

5.2. FAUST

We perform experiments on the registrations of the FAUST dataset [7]. We evaluate on 5 different pairs from one class within FAUST that were downsampled to 502 vertices. It is in theory possible to use our method on the original resolution and more pairs, but QPUs are not as freely available as traditional hardware, and we focus our available computing time on a deeper analysis of our method. We apply Q-Match as described in Alg. 1 with 500 anneals, and compare to two non-quantum methods, Simulated Annealing [28] and Functional Maps [52]. [28] is applied without ZoomOut [44] such that it is close to a version of our algorithm where only one 2-cycle is chosen in each iteration. The functional maps implementation includes a Laplacian commutativity regularizer and retrieves the final point-to-




Figure 6. Influence of the problem size on the runtime. We increase the size of the set of worst vertices (see Sec. 4.2) and measure the runtime development of different parts of our pipeline. While the QPU access time is mostly independent of the problem size, calculating \tilde{W} on the CPU grows drastically.

point correspondence through an LAP [52]. While [62] is the conceptionally closest method to ours, it only produces matches for up to preselected 4 vertices. The quantitative evaluation can be seen in Fig. 4, and some qualitative examples of our method in Fig. 5. While we do not reach the same accuracy as the recent [28], we do outperform the classical but still popular functional maps. Since AQCs are still newly developed and have limited complexity in comparison to traditional computers, we see this as a successful application of quantum annealing for real-world problems.

5.3. Convergence

We evaluate how many iterations Q-Match needs to converge on FAUST. Fig. 4 shows the convergence when choosing 26, 40, 50 worst vertices in each iteration. Since the margin between 40 and 50 is small, we chose to run all further experiments with the 40 worst matches. In each iteration every 2-cycle occurs once in a set C . Some exemplary matches obtained in this way can be seen in Fig. 5.

5.4. Runtime

Fig. 6 shows how the runtime of Q-Match scales with the problem size given by the number of worst vertices. The calculation of the submatrix W_s is the most expensive. Computing the matrix \tilde{W} given W_s takes only ms, while the calculation of W_s is in the order of seconds. We did not optimize or parallelize this operation though. The QPU access time is the time the quantum annealer spends to solve the problem, and is mostly independent of the problem size.

We report results for about 25 minutes of QPU time, which allows us to evaluate 10^4 QUBOs. One FAUST instance using 40 worst vertices and 500 anneals per QUBO requires about 2min of QPU time. For all experiments we used the default annealing path and the default annealing time of $20\mu s$. The chain strength was chosen as 1.0001 times the largest absolute value of couplings or biases.

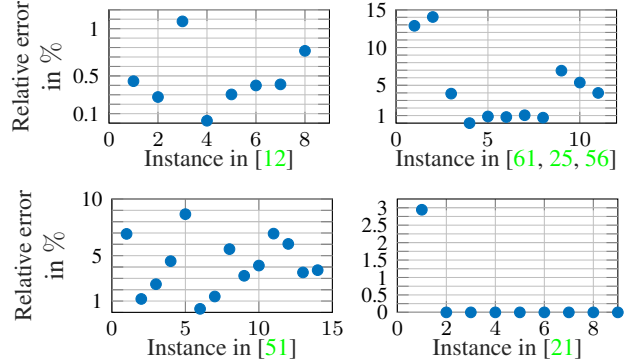


Figure 7. Relative error $\frac{E_{\text{obtained}} - E_{\text{opt}}}{E_{\text{opt}}}$ of our method in percentage for the instances of [12], [61] (1-3), [25] (4-8) and [56] (9-11), [51], and [21] in QAPLIB. The problem sizes range between 12 and 30, of which [51] contains the larger ones where we do less well.

5.5. QAPLIB

QAPLIB [11] provides a collection of QAPs together with the best known solution. This is the most commonly used benchmark for QAPs, with sets of problems from a variety of settings, *e.g.* graphs but also unstructured data. For some instances it is unclear if the best known solution is actually the optimum due to the complexity of the problem.

We ran cyclic α -expansion on some of the subsets ensuring that every 2-cycle occurs once when optimizing, and repeat this three times. We always choose the best solutions from 5000 anneals if the size is > 25 , and 500 anneals otherwise. The relative error we achieved is in Fig. 7. For instances from [12, 21, 25], our solution is almost always within 1% of the best solution. For the $n = 16$ instances in [21], we achieved the optimal solution in 9 out of 10 cases. Other sets contain instances up to size 30 and pose a bigger challenge, our worst result was within 15% of the optimum. We report our exact solutions in the supplementary material.

6. Conclusion

We introduce Q-Match and experimentally show that it can solve correspondence problems of sizes encountered in practical applications. This conclusion is made for the first time in literature for methods utilizing quantum annealing. Refraining from explicit permutation constraints and addressing the problem iteratively allowed us to significantly outperform the previous quantum state of the art, both in terms of the supported problem sizes and the probability to measure a globally optimal solution. We even achieved results comparable to a classical method. We hope that our work inspires further research in quantum annealing and related optimization problems in the vision community.

Acknowledgements. This work was partially supported by the ERC consolidator grant 4DReply (770784).

References

- [1] Dorit Aharonov, Wim Van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008.
- [2] Kurt M. Anstreicher and Nathan W. Brixius. Solving quadratic assignment problems using convex quadratic programming relaxations. *Optimization Methods and Software*, 16(1-4):49–68, 2001.
- [3] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *International Conference on Computer Vision Workshops (ICCV Workshops)*, 2011.
- [4] Florian Bernard, Christian Theobalt, and Michael Moeller. Ds*: Tighter lifting-free convex relaxations for quadratic matching problems. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [5] Dimitri P. Bertsekas. *Network Optimization: Continuous and Discrete Models*. Athena Scientific, 1998.
- [6] Tolga Birdal, Vladislav Golyanik, Christian Theobalt, and Leonidas Guibas. Quantum permutation synchronization. In *Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [7] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [8] Max Born and Vladimir Fock. Beweis des adiabatsatzes. *Zeitschrift für Physik*, 51(3):165–180, 1928.
- [9] Edward Boyda, Saikat Basu, Sangram Ganguly, Andrew Michaelis, Supratik Mukhopadhyay, and Ramakrishna R. Nemani. Deploying a quantum annealing processor to detect tree cover in aerial imagery of california. *PLoS ONE*, 12, 2017.
- [10] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 23(11), 2001.
- [11] Rainer E. Burkard, Stefan E. Karisch, and Franz Rendl. Qaplib - a quadratic assignment problem library. *Journal of Global Optimization*, 1997.
- [12] Rainer E. Burkard and Josef Offermann. Entwurf von schreibmaschinentastaturen mittels quadratischer zuordnungsprobleme. *Zeitschrift für Operations Research*, 1977.
- [13] Jun Cai, William G. Macready, and Aidan Roy. A practical heuristic for finding graph minors. *arXiv e-prints*, 2014.
- [14] Simona Caraiman and Vasile I. Manta. Image processing using quantum computing. In *International Conference on System Theory, Control and Computing (ICSTCC)*, 2012.
- [15] Gabriele Cavallaro, Dennis Willsch, Madita Willsch, Kristel Michielsen, and Morris Riedel. Approaching remote sensing image classification with ensembles of support vector machines on the d-wave quantum annealer. In *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2020.
- [16] Andrew M Childs, Edward Farhi, and John Preskill. Robustness of adiabatic quantum computation. *Physical Review A*, 65(1):012322, 2001.
- [17] D-Wave Systems, Inc. Leap datasheet v10. https://www.dwavesys.com/sites/default/files/Leap_Datasheet_v10_0.pdf, 2021. online; accessed on the 25.02.2021.
- [18] Arnab Das and Bikas K Chakrabarti. *Quantum annealing and related optimization methods*, volume 679. Springer Science & Business Media, 2005.
- [19] 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? *Phys. Rev. X*, 6:031015, Aug 2016.
- [20] Roberto Dyke, Caleb Stride, Yu-Kun Lai, Paul L. Rosin, Mathieu Aubry, Amit Boyarski, Alex M. Bronstein, Michael M. Bronstein, Daniel Cremers, and Matthew Fisher et al. Shrec’19: Shape correspondence with isometric and non-isometric deformations. In *Workshop on 3D Object Retrieval (3DOR)*, 2019.
- [21] Bernhard Eschermann and Hans-Joachim Wunderlich. Optimized synthesis of self-testable finite state machines. In *20th International Symposium on Fault-Tolerant Computing (FTCS 20)*, 1990.
- [22] 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.
- [23] Fajwel Fogel, Rodolphe Jenatton, Francis Bach, and Alexandre D’Aspremont. Convex relaxations for permutation problems. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, 2013.
- [24] Vladislav Golyanik and Christian Theobalt. A quantum computational approach to correspondence problems on point sets. In *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [25] Scott W Hadley, Franz Rendl, and Henry Wolkowicz. A new lower bound via projection for the quadratic assignment problem. *Mathematics of Operations Research*, 1992.
- [26] Peter Hahn, Thomas Grant, and Nat Hall. A branch-and-bound algorithm for the quadratic assignment problem based on the hungarian method. *European Journal of Operational Research*, 108(3):629–640, 1998.
- [27] Oshri Halimi, Or Litany, Emanuele R. Rodolà, Alex M. Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [28] Benjamin Holzschuh, Zorah Lähler, and Daniel Cremers. Simulated annealing for 3d shape correspondence. In *Conference of 3D Vision (3DV)*, 2020.
- [29] Ibm q experience. <https://quantum-computing.ibm.com>, 2021. Accessed on the 26.02.2021.
- [30] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, 1998.
- [31] Itay Kezurer, Shahar Z. Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxation of quadratic matching. In *Eurographics Symposium on Geometry Processing*, 2015.
- [32] Vladimir G. Kim, Yaron Lipman, and Thomas A. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4), 2011.

- [33] 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.
- [34] Tjalling C. Koopmans and Martin Beckmann. Assignment Problems and the Location of Economic Activities. *Econometrica*, 25(1):53, Jan. 1957.
- [35] Eugene L. Lawler. The quadratic assignment problem. *Management science*, 9(4):586–599, 1963.
- [36] D. Khuê Lê-Huu and Nikos Paragios. Alternating direction graph matching. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [37] Victor Lempitsky, Carsten Rother, Stefan Roth, and Andrew Blake. Fusion moves for markov random field optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 32(8):1392–1405, 2010.
- [38] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *International Conference on Computer Vision (ICCV)*, 2005.
- [39] Junde Li and Swaroop Ghosh. Quantum-soft qubo suppression for accurate object detection. In *European Conference on Computer Vision (ECCV)*, 2020.
- [40] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *International Conference on Computer Vision (ICCV)*, 2017.
- [41] Sven Loncaric. A survey on shape analysis techniques. *Pattern Recognition*, 31(8), 1998.
- [42] Jonathan Masci, Davide Boscaini, Michael M. Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *International Conference on Computer Vision Workshop (ICCVW)*, 2015.
- [43] Catherine C. McGeoch. Adiabatic quantum computation and quantum annealing: Theory and practice. *Synthesis Lectures on Quantum Computing*, 5(2):1–93, 2014.
- [44] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 2019.
- [45] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [46] Hartmut Neven, Vasil S. Denchev, Geordie Rose, and William G. Macready. Qboost: Large scale classifier training with adiabatic quantum optimization. In *Asian Conference on Machine Learning (ACML)*, 2012.
- [47] Hartmut Neven, Geordie Rose, and William G Macready. Image recognition with an adiabatic quantum computer i. mapping to quadratic unconstrained binary optimization. *arXiv preprint arXiv:0804.4457*, 2008.
- [48] Nga T. T. Nguyen and Garrett Kenyon. Image classification using quantum inference on the d-wave 2x. In *International Conference on Rebooting Computing (ICRC)*, 2018.
- [49] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.
- [50] Mohammadreza Noormandipour and Hanchen Wang. A parameterised quantum circuit approach to point set matching. *arXiv e-prints*, 2021.
- [51] Christopher E. Nugent, Thomas E. Vollmann, and John Ruml. An experimental comparison of techniques for the assignment of facilities to locations. *Operations Research*, 1968.
- [52] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *Trans. Graphics*, 31(4):30, 2012.
- [53] Daniel O’Malley, Velimir V. Vesselinov, Boian S. Alexandrov, and Ludmil B. Alexandrov. Nonnegative/binary matrix factorization with a d-wave quantum annealer. *PLOS ONE*, 13(12), 12 2018.
- [54] Janez Povh and Franz Rendl. Copositive and semidefinite relaxations of the quadratic assignment problem. *Discrete Optimization*, 6(3):231–241, 2009.
- [55] Robert Raussendorf and Hans J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86:5188–5191, May 2001.
- [56] Catherine Roucairol. *Du séquentiel au parallèle: la recherche arborescente et son application à la programmation quadratique en variables 0 et 1*. PhD thesis, Université Pierre et Marie Curie, 1987.
- [57] Catherine Roucairol. A parallel branch and bound algorithm for the quadratic assignment problem. *Discrete Applied Mathematics*, 18(2):211–225, 1987.
- [58] Sitapa Rujikietgumjorn and Robert T. Collins. Optimized pedestrian detection for multiple and occluded people. In *Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [59] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Comput. Vis. Image Underst. (CVIU)*, 125:251–264, 2014.
- [60] William Raymond Scott. *Group theory*. Courier Corporation, 2012.
- [61] Michael Scriabin and Roger C. Vergin. Comparison of computer algorithms and visual based methods for plant layout. *Management Science*, 1975.
- [62] Marcel Seelbach Benkner, Vladislav Golyanik, Christian Theobalt, and Michael Moeller. Adiabatic quantum graph matching with permutation matrix constraints. In *Conference of 3D Vision (3DV)*, 2020.
- [63] Justin Solomon, Gabriel Peyré, Vladimir G. Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Trans. Graph.*, 35(4), 2016.
- [64] Bo Sun, Abdullah Iliyasa, Fei Yan, Fangyan Dong, and Kaoru Hirota. An rgb multi-channel representation for images on quantum computers. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 17:404–417, 2013.
- [65] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature-based on heat diffusion. *Computer Graphics Forum*, 2009.

- [66] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *European Conference on Computer Vision (ECCV)*, 2010.
- [67] Yanghai Tsin and Takeo Kanade. A correlation-based approach to robust point set registration. In *European Conference on Computer Vision (ECCV)*, 2004.
- [68] Salvador E. Venegas-Andraca and Sougato Bose. Storing, processing, and retrieving an image using quantum mechanics. In *Quantum Information and Computation*. International Society for Optics and Photonics, 2003.
- [69] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodola, Alex Bronstein, Michael Bronstein, Ron Kimmel, and Daniel Cremers. Efficient deformable shape correspondence via kernel matching. In *Conference on 3D Vision (3DV)*, 2017.
- [70] Dennis Willsch, Madita Willsch, Hans De Raedt, and Kristel Michielsen. Support vector machines on the d-wave quantum annealer. *Computer Physics Communications*, 248:107006, 2020.
- [71] Fei Yan, Abdullah M. Iiyasu, and Salvador E. Venegas-Andraca. A survey of quantum image representations. *Quantum Information Processing*, 15(1), 2016.
- [72] Qing Zhao, Stefan Karisch, Franz Rendl, and Henry Wolkowicz. Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2, 10 1997.