

A Quantum Computational Approach to Correspondence Problems on Point Sets

Vladislav Golyanik Christian Theobalt
 Max Planck Institute for Informatics, Saarland Informatics Campus

Abstract

Modern adiabatic quantum computers (AQC) are already used to solve difficult combinatorial optimisation problems in various domains of science. Currently, only a few applications of AQC in computer vision have been demonstrated. We review AQC and derive a new algorithm for correspondence problems on point sets suitable for execution on AQC. Our algorithm has a subquadratic computational complexity of the state preparation. Examples of successful transformation estimation and point set alignment by simulated sampling are shown in the systematic experimental evaluation. Finally, we analyse the differences in the solutions and the corresponding energy values.

1. Introduction

Since their proposal in the early eighties [8, 43, 27], quantum computers have attracted much attention of physicists and computer scientists. Impressive advances both in quantum computing hardware and algorithms have been demonstrated over the last thirty years [40, 30, 61, 58, 42, 19, 25, 49, 65, 48]. Quantum computers are not universally faster than conventional machines, but they can natively execute algorithms relying on quantum parallelism, *i.e.*, the ability to perform operations on exponentially many superimposed memory states simultaneously [59].

To harness the advantages, carefully designed algorithms are required. Nowadays, the motivation to take advantage of quantum effects in computing is also facilitated by the classical computing paradigm approaching its limits, since the quantum effects are becoming non-neglectable while manufacturing and using conventional CPUs. As a result, alternative paradigms such as massively parallel computing devices have been brought into being.

While universal gate quantum computer technology has not yet reached the maturity, modern adiabatic quantum annealers (AQA) are already capable of solving difficult real-world combinatorial optimisation problems [15, 14, 23, 49]. The primary difference of universal gate quantum computing and AQA is that the latter can address objectives formulated as *quadratic unconstrained binary optimisation prob-*




Figure 1: Different 2D point sets — fish [47], qubit, kanji and composer — aligned with our QA approach. For every pair of point sets, the initial misalignment is shown on the left, and the registration is shown on the right. QA is the first transformation estimation and point set alignment method which can be executed on adiabatic quantum computers.

lems (QUBOP) defined as

$$\arg \min_{\mathbf{q} \in \mathbb{B}^n} \mathbf{q}^T \mathbf{P} \mathbf{q}, \quad (1)$$

where \mathbf{q} is a set of n binary variables, and \mathbf{P} is a symmetric matrix of weights between the variables. The operational principle of AQA is grounded on the *adiabatic theorem of quantum mechanics* [17] which states that

if a quantum-mechanical system is in the ground state of a time-dependent Hamiltonian and parameters of this Hamiltonian are changing gradually enough, the system will continue to remain in the ground state during the evolution (see Table 1 for quantum notions). (2)

In their seminal paper, Farhi *et al.* [26] have shown that the adiabatic principle (2) can be used for solving \mathcal{NP} -complete optimisation problems and laid the foundation for adiabatic quantum computing. Several years later, Aharonov *et al.* [2] theoretically showed the equivalence between classical quantum computing and quantum annealing models. As of 2019-2020, general-purpose quantum computers accessible for research purposes and applications contain up to 20 qubits [19]. In contrast, the latest quantum annealers support up to 2^{10} qubits [21]¹. Nevertheless, due

¹the amount of logical qubits which are available on this system is an order of magnitude lower since most qubits are used for the error correction

to design and practical restrictions, quantum algorithms for the gates model such as Shor’s prime number factorisation [61] or Grover’s search algorithms [30] cannot be implemented on current quantum annealers.

Motivation and Contributions. Considering recent successful applications of AQA in several fields of computational science [42, 49, 65], we are motivated to investigate how useful AQA can be for computer vision and which problems can be potentially solved on the new hardware. The vast majority of available materials about quantum annealers are either oriented to physicists or lack technical details and clarity. *Our goal is to fill this gap, introduce the reader into the modern AQA and provide all notions and the background to understand, analyse, simulate and design quantum algorithms for computer vision which can potentially run on modern AQA, as well as interpret the results.*

We consider correspondence problems on point sets which have various applications in computer vision. They consist in finding an optimal rigid transformation between inputs [34, 13, 47, 66]. While transformation estimation assumes known matches, point set alignment is more general and targets, in addition, the recovery of correspondences. We consider two inputs, *i.e.*, a fixed *reference* point set and a *template* undergoing a rigid transformation. *Thus, our goal is to design a quantum approach for point set alignment which can potentially run on AQA and show that it offers advantages compared to the classical counterparts.*

Therefore, we adapt the recent progress in rigid point set alignment and formulate a globally multiply-linked energy functional which does not require any intermediate correspondence updates [29]. In the *gravitational approach* (GA) [29], the optimal alignment is achieved when the gravitational potential energy (GPE) of the system with two interacting particle swarms is locally minimal. Proceeding from GA, we build the weight matrix \mathbf{P} for the associated QUBOP (1) which is unalterably valid in the course of the optimisation. Along with that, we are targeting at a method which is implementable on classical hardware and can solve real-world problems, *cf.* Fig. 1. To summarise, the main **contributions** of this paper are:

- A self-contained and detailed introduction into modern quantum annealers for computer vision problems, including notions from quantum physics and computing (Sec. 2), modern adiabatic quantum annealers (Sec. 3) including D-WAVE (Sec. 3.2), and previous and related works from quantum computing (Sec. 4).
- The first quantum approach (QA) to transformation estimation (Sec. 5) and point set alignment (Sec. 6) which can run on the upcoming quantum annealers (Sec. 6.2).
- Experimental analysis of the proposed method in a simulated environment on several datasets (Sec. 7).

quantum notion	classical counterpart
<i>qubit (states $0\rangle$ and $1\rangle$)</i>	bit (states 0 and 1)
<i>(time-dependent) Hamiltonian eigenstate</i>	energy functional some energy state
<i>ground state</i>	globally optimal energy state
<i>quantum system evolution</i>	optimisation process
<i>quantum annealing [26]</i>	simulated annealing [39]

Table 1: Quantum notions and their counterparts in computer vision.

2. Preliminaries, Definitions and Notations

In this section, we introduce the reader into the basics of quantum computing. See Table 1 for a lookup of notions specific to AQA which have counterparts and interpretation in the classical optimisation theory for computer vision.

Qubit. Quantum computing encompasses tasks which can be performed on quantum-mechanical systems [53]. Quantum *superposition* and *entanglement* are two forms of parallelism evidenced in quantum computers. A *qubit* is a quantum-mechanical equivalent of a classical bit. A qubit $|\phi\rangle$ — written in the *Dirac* notation — can be in the state $|0\rangle$, $|1\rangle$ or an arbitrary *superposition of both states* denoted by $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are the (generally, complex) probability amplitudes satisfying $|\alpha|^2 + |\beta|^2 = 1$. In quantum computing, the state $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ denoted by $|+\rangle$ is often used for initialisation of a qubit register. The state of a qubit remains hidden during the entire computation and reveals when measured. If qubits are *entangled*, measuring one of them influences the measurement outcome of the other one [59]. During the measurement, the qubit’s state irreversibly collapses to one of the basis states $|0\rangle$ or $|1\rangle$. Efficient physical realisation of a qubit demand very low temperatures. Otherwise, thermal fluctuations will destroy it and lead to arbitrary changes of the measured qubit state.

One possible physical implementation of a qubit is an electron which possesses a spin, *i.e.*, its intrinsic magnetic moment [53, 63]. The spin of an electron can be manipulated and brought to the state *spin down*, *spin up*, or a superposition of both. A concrete experimentally realised scheme that uses this property is represented by an atom of phosphorus ^{31}P embedded into a ^{28}Si silicon lattice attached to a transistor [37, 46, 67]. The nucleus of ^{31}P has a positive charge compensated by electrons. The bundle of electrons in the transistor is filled up to the energetic level between the energy of spin-down and spin-up state of ^{31}P . To change a state of a ^{31}P - ^{28}Si qubit, a microwave pulse of the frequency — which is equal to the resonance frequency of the atom — is applied to it. The new state $|\phi\rangle$ depends on the duration of the exposure. A transistor is used to measure a state of the ^{31}P - ^{28}Si qubit. If the extra electron of ^{31}P tunnels into the electron bundle, a positive charge is measured in the transistor indicating the spin-up state (*e.g.*, $|1\rangle$).

Fig. 2-(a) visualises a qubit with a so-called Bloch sphere. Every *qubit* can be both in a superposition and en-

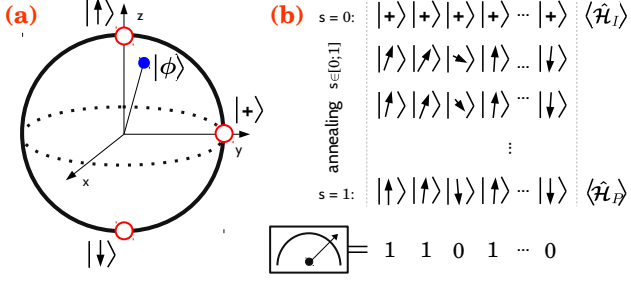


Figure 2: (a): Schematic depiction of a qubit with a Bloch sphere. Spin-up or $|1\rangle$ is located on the north pole, and spin down or $|0\rangle$ is located on the south pole. The state $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ with equal probability amplitudes to measure $|1\rangle$ and $|0\rangle$ values is geodetically equidistant to both poles. A point on the surface of the Bloch sphere corresponds to a valid pure state $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$. (b): Schematic visualisation of adiabatic quantum annealing (AQA). At the beginning, all qubits are initialised in the state $|+\rangle$. After the annealing is finished, the qubit states are measured and returned. After the measurement, the states of variables are classical.

tangled with other qubits. Thus, quantum superposition is the property that calculations are performed on all possible inputs simultaneously which can result in exponential parallelism in the number of qubits. When entangled, states of qubits cannot be described independently from each other.

Schrödinger Equation. In the universal or gates model, changes are expressed by a series of unitary transformations applied to qubits. This is a useful practical simplification, while the evolution of every quantum-mechanical system can be described more precisely by continuous Schrödinger equation, which in common notation reads:

$$-i \frac{d}{dt} |\phi(t)\rangle = \hat{\mathcal{H}}(t) |\phi(t)\rangle. \quad (3)$$

For the simplicity, we denote here by $|\phi\rangle$ a state of n qubits at time t , and $\hat{\mathcal{H}}(t)$ is a Hamiltonian which is, in this case, a $2^n \times 2^n$ Hermitian matrix. Thus, a discrete time evolution of the quantum system is given by a unitary transformation.

Hamiltonian. Hamiltonian $\hat{\mathcal{H}}$ is an energy operator of a system of n qubits. It defines the energy spectrum of a system or, in our case, the space of all possible solutions. The *ground state* of the system is its lowest energy eigenstate. Finding a ground state of a Hamiltonian is equivalent to finding an optimal solution to the problem. The expectation value of Hamiltonian $\langle \hat{\mathcal{H}} \rangle$ provides an instantaneous energy of a given qubit configuration. In correspondence problems, $\langle \hat{\mathcal{H}} \rangle$ is a quantitative characteristic of point set alignment. We denote by $\Delta(\hat{\mathcal{H}})$ the spectral gap of $\hat{\mathcal{H}}$, *i.e.*, the difference between the energies of the ground state and the second lowest eigenstate. The spectral gap influences the annealing rate and is considerable for algorithm design and evaluation in quantum annealing.

Pauli Matrices. An arbitrary Hamiltonian of a n -qubit-system can be expressed by a linear combination of tensor products of Pauli matrices denoted by:

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (4)$$

The Pauli matrices are 2×2 Hermitian and unitary. Together with the identity $\sigma^0 = \mathbf{I}_{2 \times 2}$, they form a basis for $\mathbb{C}^{2 \times 2}$. σ^x flips the probabilities to measure $|0\rangle$ and $|1\rangle$, whereas $\sigma^z|0\rangle = |0\rangle$, and $\sigma^z|1\rangle = -|1\rangle$.

Pseudo-Boolean Functions. A pseudo-boolean function is a real vector-valued function of n boolean variables denoted by \mathbf{x} of the form $\mathbf{F}(\mathbf{x}) : \mathbf{B}^n \rightarrow \mathbb{R}^M$, where M is the number of real-valued outputs.

Quantum Annealing. Quantum annealing is a heuristic combinatorial optimisation method for finding global optima which relies on quantum effects (superposition, entanglement and tunnelling) [11, 36]. In particular, it is used to find a ground state of an Ising Hamiltonian [35, 57], which encodes the target computational problem, see Fig. 2-(b).

Quantum annealing is the quantum counterpart of simulated annealing [45, 39]. Starting from the superposition state $|+\rangle^{\otimes n}$ (this is a shorthand for n qubits in the state $|+\rangle$, *cf.* (9)), the system evolves according to (3) under an external time-dependent magnetic field (a transverse field). When the external field is faded away, the system reaches the ground state of an Ising model [35]. According to (2), if an external magnetic field is changing gradually enough, the system remains near the ground state with high probability throughout the optimisation. Quantum annealing systems taking advantage of (2) are called *adiabatic quantum computers* (AQC). QUBOP is the most common problem form which can be mapped to current realisations of AQC.

3. Modern Adiabatic Quantum Computation

Adiabatic quantum computation is a form of quantum annealing which relies on the adiabatic theorem of quantum mechanics (2) [17]. Starting from a ground state of an initial default Hamiltonian $\hat{\mathcal{H}}_I$, an AQC system adiabatically evolves into the ground state of a problem Hamiltonian $\hat{\mathcal{H}}_P$ which encodes a solution to a problem [26]. In the case of *adiabatic quantum annealing* (AQA), the problem Hamiltonian $\hat{\mathcal{H}}_P$ is given by the Ising model [35]:

$$\hat{\mathcal{H}}_P = \sum_{j \in V} h_j \sigma_j^z + \sum_{(j,k) \in E_P} J_{j,k} \sigma_j^z \otimes \sigma_k^z, \quad (5)$$

with the Kronecker product \otimes , h_j denoting exterior local magnetic fields and $J_{i,j}$ standing for the pairwise connections between the particles. V is a set of particles, and E_P is a set of edges (intra-particle links) of the graph. Eq. (5) is written in a notation common in physics. The first term of (5) on the right side in the explicit notation reads

$$\hat{\mathcal{H}}_P^{j \in V} = \left(\left[\begin{array}{c} [\sigma^z \otimes \mathbf{I} \otimes \dots \otimes \mathbf{I}] \\ [\mathbf{I} \otimes \sigma^z \otimes \dots \otimes \mathbf{I}] \\ \vdots \\ [\mathbf{I} \otimes \mathbf{I} \otimes \dots \otimes \sigma^z] \end{array} \right]^T \right) \begin{bmatrix} h_1 \mathbf{I}_{2^n \times 2^n} \\ h_2 \mathbf{I}_{2^n \times 2^n} \\ \vdots \\ h_n \mathbf{I}_{2^n \times 2^n} \end{bmatrix}, \quad (6)$$

where \mathbf{I} without a subscript is a 2×2 identity matrix. The second term $\hat{\mathcal{H}}_P^{(j,k) \in E_P}$ of (5) can be expressed in a similar manner, involving pairs of σ^z in the tensor product depending on the connectivity of the lattice.

Theoretically, each particle can interact with any other particle from the whole set of qubits. In practice, the couplings are restricted to local neighbourhoods (see Sec. 3.2). Thus, (5) describes a system of N interacting spin-1/2 particles under the influence of distributed magnetic forces, and in the expanded form, $\hat{\mathcal{H}}_P$ is a $2^n \times 2^n$ matrix. Finding a ground state of an Ising model is an \mathcal{NP} -hard problem [7]. In the ground state, the spin configuration of all particles which minimises Ising energy E_{Ising} is given by:

$$E_{\text{Ising}} = \sum_i h_i s_i + \sum_{i,j} J_{i,j} s_i s_j, \quad (7)$$

where $s_i \in \{1, -1\}$ denotes two possible spin measurement outcomes of a spin-1/2 particle.

3.1. Quantum System Evolution

Solving \mathcal{NP} -hard problems such as QUBOP on a classical computer requires exponential time in the size of the input. The main idea of the AQC is that a QUBOP (1) can be mapped to the Ising model (5) and optimised by allowing the system to evolve according to the adiabatic principle (2). Once annealing is finished, the qubit register will represent the solution to the programmed problem with a high probability [26] (cf. Appendix A on the annealing rate criterion). The initial Hamiltonian of the system is always initialised in the state

$$\hat{\mathcal{H}}_I = - \sum_{j \in V} B_x \sigma_j^x, \quad (8)$$

where $B_x > 0$ stands for a magnetic field pointing in the x direction. The ground state of (8) is a symmetrised superposition with equal normalised probability amplitudes for the states $|0\rangle$ and $|1\rangle$ for all qubits, *i.e.*,

$$[|+\rangle]^{\otimes n} = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2^n}} \right)^{\otimes n}. \quad (9)$$

This initial state (9) is comparably easy to construct by radiating a microwave of the same duration and wavelength to all qubits. In mathematical terms, (9) is obtained by applying a Hadamard transform $H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ to n $|0\rangle$ qubits.

The lowest energy $E_{\text{GS}} = -nB_x$ of (8) is achieved when all qubits in the system point in the anti-parallel direction of the magnetic field, so that $\sigma_j^x |s_j\rangle = |s_j\rangle$. During AQC, the initial Hamiltonian $\hat{\mathcal{H}}_I$ is evolving into the problem Hamiltonian $\hat{\mathcal{H}}_P$, with a high probability of reaching the ground state of $\hat{\mathcal{H}}_P$ [26]. The interpolation between the Hamiltonians can be written as

$$\hat{\mathcal{H}} = [1 - s] \hat{\mathcal{H}}_I + s \hat{\mathcal{H}}_P, \quad (10)$$

with $s \in [0; 1]$ being the time in relative units from the start of annealing at $s = 0$ until reaching the ground state of $\hat{\mathcal{H}}_P$ at $s = 1$. The problem Hamiltonian and the final state of the system depend on the objective function $f(x)$ or the matrix of weights between the qubits \mathbf{P} in (1). After the annealing is accomplished, the state of each qubit is measured, and the result corresponds to the solution of the programmed problem with a high probability. At this stage, the states of all binary variables are classical, and not quantum anymore.

To remain in the ground state during the system evolution, the annealing rate has to be carefully chosen. The condition of adiabaticity (2) is derived from the time-dependent perturbation theory of quantum systems. It is achieved when the average energy pumped into the system per time interval T is smaller than the minimal energy difference between the ground state and the first excited state. This statement was quantified in [4] which generalises the original adiabatic theorem [17] for periodic driving, see Appendix A for further details.

3.2. Quantum Annealer D-WAVE

D-WAVE relies on the adiabatic criterion in its specified form and currently supports up to ≈ 2000 qubits [22]. It reflects the state of the art in physical realisation of quantum processors. It is relatively inexpensive to bring the system in the superposition state, and every computation on D-WAVE starts with the problem-independent $\hat{\mathcal{H}}_I$ (8). Qubits can interact with a restricted number of other qubits, and it is possible to define qubit equality and entanglement constraints [22]. Possible interactions can be seen from the *chimera* graph which schematically depicts the layout of the quantum processor [22, 16]. At the same time, the physically realised connectivity can model QUBOP with arbitrary connectivities through an internal conversion [16]. The drawback is that in the worst case, a quadratic increase in the number of variables is required. A fully connected graph of layers with N qubits would require N^2 qubits for processing. Some QUBOP cannot be mapped to the chimera graph, and some problems can be mapped in multiple ways [55].

4. Previous and Related Work

Universal Quantum Computers. The paradigm of the universal quantum computer originates in the attempts to gain control over individual quantum systems in the early eighties [62, 53]. Later, extending the control to multiple quantum systems has attracted the interest of physicists, promising to facilitate discoveries in quantum physics [53]. By that time, it was noticed that simulating a quantum-mechanical system on a classical computer requires exponential time in the number of simulated elements [43, 27]. “Can you do it² with a new kind of computer – a quantum com-

²to simulate quantum-mechanical effects

puter?” [27] is a famous quote by R. Feynman which has triggered research on quantum computers in the subsequent years. The so-called *no-cloning theorem* [56, 64] belongs to the first discoveries strongly influenced quantum information theory and quantum computations. Nowadays, quantum computers can be used not only to fulfil their primary goal, *i.e.*, to simulate quantum-mechanical systems for different branches of science, but also to solve other computational problems — such as balanced function decision problem [24], quantum Turing machines for complexity analysis [12], prime number factorisation and discrete logarithms [61], database search [30], graph matching [3], data classification [58] and principal component analysis [42] — faster than on classical machines. The related field of quantum communication and quantum key distribution has already found broad practical use nowadays [10, 9, 60].

Classical Methods using Quantum Analogies. Quantum-mechanical effects inspired multiple techniques for conventional computers including variants of genetic and evolutionary algorithms [31, 32], non-rigid mesh analysis [5] and image segmentation [6], among others.

Quantum Annealers in Computer Vision. Only a few theoretical results and applications of AQC to image processing, machine learning and computer vision are known. Neven *et al.* [51] have shown how image recognition can be formulated as QUBOP. Image classification on 12×12 images with AQC was addressed in [52]. The approach of O’Malley *et al.* can learn facial features and reproduce facial image collections [54]. Boyda *et al.* [18] propose an AQC method to detect areas with trees from aerial images. Several methods target classification, dimensionality reduction and training of deep neural networks [50, 38, 1]. Not all theoretical findings of these works are possible to test on the real AQC hardware yet. Nonetheless, we believe that it is essential to explore the theory and highlight the advantages of the upcoming hardware for computer vision tasks.

5. Quantum Transformation Estimation

In this section, we introduce our QA to transformation estimation. The inputs are a reference point set $\{\mathbf{x}_n\} \in \mathbf{X} \in \mathbb{R}^{D \times N}$ and a template point set $\{\mathbf{y}_n\} \in \mathbf{Y} \in \mathbb{R}^{D \times N}$, $n \in \{1, \dots, N\}$. N is the number of points in both point sets and D is the dimensionality of the points. We assume that translation is resolved, the centroids of the point sets coincide, and points are in correspondence.

5.1. Transformation Estimation in 2D

To obtain an advantage in solving transformation estimation on a quantum annealer we should avoid uniform sampling of rotations applied to \mathbf{Y} . Elements of the rotation group are non-commutative, and it is not possible to formulate multiplication of basis rotations as QUBOP. Instead, we propose to represent the transformation matrix as a linear

combination of basis elements. Recall that for any rotation matrix, $\mathbf{R}^{-1} = \mathbf{R}^\top$. Rotation in 2D consists of four elements, *i.e.*, $\mathbf{R} = \begin{pmatrix} r_{1,2} & r_{2,2} \\ r_{2,1} & r_{2,2} \end{pmatrix}$. Additively, we can create a basis for all possible values of \mathbf{R} and encode the influence of the additive elements as binary variables. Consider instead the power series of \mathbf{R} in 2D. Every such matrix has a corresponding skew-symmetric matrix of the form

$$\mathbf{S} = \theta \mathbf{M}, \quad \mathbf{M} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad (11)$$

with a real number θ . According to the Cayley-Hamilton theorem, $\mathbf{S}^2 + \theta^2 \mathbf{I} = 0$ which leads to the following exponential map for \mathbf{R} with power series:

$$\begin{aligned} \mathbf{R} &= \exp(\mathbf{S}) = \\ &= \cos(\theta) \mathbf{I} + \left(\frac{\sin(\theta)}{\theta} \right) \mathbf{S} = \cos(\theta) \mathbf{I} + \sin(\theta) \mathbf{M}. \end{aligned} \quad (12)$$

From (12) we see that \mathbf{R} is composed of an identity weighted by $\cos(\theta)$ and \mathbf{M} weighted by $\sin(\theta)$. If the basis would resemble additive elements \mathbf{I} and \mathbf{M} of the exponential map, we can stronger constrain the resulting \mathbf{R} . We see that $r_{1,1}$ is entangled with $r_{2,2}$, and $r_{1,2}$ is entangled with $r_{2,1}$. Eventually, we need fewer basis elements, the optimisation will finish faster and the method can be also implemented and tested on a classical computer. Thus, our basis $\mathbf{Q} = \{\mathbf{Q}_k\}$ for \mathbf{R} is a compound of $K = 20$ elements:

$$\begin{aligned} \{\mathbf{Q}_k = \omega \mathbf{C} \in \mathbb{R}^{2 \times 2}, \forall \omega \in \{0.5, 0.2, 0.1, 0.1, 0.05\}, \\ \forall \mathbf{C} \in \{\mathbf{I}, \mathbf{M}, -\mathbf{I}, -\mathbf{M}\}\}. \end{aligned} \quad (13)$$

Since we want to find \mathbf{R} which minimises the distances between the corresponding points $(\mathbf{x}_n, \mathbf{y}_n)$, we multiply each template point with a *negative* sign $-\mathbf{y}_n$ with each basis element \mathbf{Q}_k and stack the result into Φ :

$$\Phi = \begin{bmatrix} \mathbf{x}_1^\top & \mathbf{x}_2^\top & \dots & \mathbf{x}_N^\top \\ -[\mathbf{Q}_1 \mathbf{y}_1]^\top & -[\mathbf{Q}_1 \mathbf{y}_2]^\top & \dots & -[\mathbf{Q}_1 \mathbf{y}_N]^\top \\ -[\mathbf{Q}_2 \mathbf{y}_1]^\top & -[\mathbf{Q}_2 \mathbf{y}_2]^\top & \dots & -[\mathbf{Q}_2 \mathbf{y}_N]^\top \\ \vdots & \vdots & \ddots & \vdots \\ -[\mathbf{Q}_K \mathbf{y}_1]^\top & -[\mathbf{Q}_K \mathbf{y}_2]^\top & \dots & -[\mathbf{Q}_K \mathbf{y}_N]^\top \end{bmatrix}. \quad (14)$$

Next, we set the weight matrix in (1) as

$$\mathbf{P} = \Phi \Phi^\top, \quad (15)$$

and the final QUBOP reads

$$\arg \min_{\mathbf{q} \in \mathbb{B}^{21}} \mathbf{q}^\top \Phi \Phi^\top \mathbf{q}. \quad (16)$$

In total, 21 qubits are required to resolve the transformation on AQC in 2D, with the first qubit of \mathbf{q} being fixed to $|1\rangle$. After solving (16) with quantum annealing and measuring \mathbf{q} , we obtain a classical bitstring $\hat{\mathbf{q}}$. The resulting (perhaps approximate) \mathbf{R} is then obtained by *unembedding* as

$$\mathbf{R} = \sum_{k=1}^K \hat{q}_{k+1} \mathbf{Q}_k. \quad (17)$$

5.2. Transformation Estimation in 3D

In 3D, a skew-symmetric matrix can be represented as

$$\mathbf{S} = \theta \mathbf{M}, \mathbf{M} = \begin{bmatrix} m_{1,1} & m_{1,2} & m_{1,3} \\ m_{2,1} & m_{2,2} & m_{2,3} \\ m_{3,1} & m_{3,2} & m_{3,3} \end{bmatrix} = \begin{bmatrix} 0 & a & b \\ -a & 0 & c \\ -b & -c & 0 \end{bmatrix}, \quad (18)$$

where θ , a , b and c are real numbers, and $a^2 + b^2 + c^2 = 1$. In the 3D case, the Cayley-Hamilton theorem states that $-\mathbf{S}^3 - \theta^2 \mathbf{S} = 0$. The exponential map for \mathbf{R} in 3D with power series reads

$$\mathbf{R} = \exp(\mathbf{S}) = \mathbf{I} + \left(\frac{\sin \theta}{\theta}\right) \mathbf{S} + \left(\frac{1 - \cos \theta}{\theta^2}\right) \mathbf{S}^2 = \mathbf{I} + \sin \theta \mathbf{M} + (1 - \cos \theta) \mathbf{M}^2. \quad (19)$$

Next, \mathbf{M} can be decomposed as follows:

$$\mathbf{M} = a \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{\mathbf{M}_a} + b \underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}}_{\mathbf{M}_b} + c \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}}_{\mathbf{M}_c}. \quad (20)$$

Regarding \mathbf{M} , we see that

- $\{m_{1,2}; m_{2,1}\}$, $\{m_{1,3}; m_{3,1}\}$ and $\{m_{2,3}; m_{3,2}\}$ are mutually dependent or entangled,
- $m_{i,j} \in [-1; 1]$,
- $\mathbf{M} = -\mathbf{M}^T$, $\mathbf{M}_a = -\mathbf{M}_a^T$, $\mathbf{M}_b = -\mathbf{M}_b^T$ and $\mathbf{M}_c = -\mathbf{M}_c^T$, *i.e.*, they are anti-symmetric, and
- $\mathbf{M}^2 = \begin{bmatrix} v_1^- & d & e \\ d & v_2^- & f \\ e & f & v_3^- \end{bmatrix}$ is symmetric negative semi-definite, with $\{v_1^-, v_2^-, v_3^-\} \in \mathbb{R}^-$, and $\{d, e, f\} \in \mathbb{R}$.

The basis for rotation in 3D is comprised of the identity matrix \mathbf{I} , \mathbf{M}_a , \mathbf{M}_b , \mathbf{M}_c as well as the basis for \mathbf{M}^2 :

$$\mathbf{M}_d = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \mathbf{M}_e = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \mathbf{M}_f = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}. \quad (21)$$

Thus, our basis $\mathbf{Q}^{3D} = \{\mathbf{Q}_k^{3D}\}$ for \mathbf{R} in 3D is a compound of $K = 80$ elements:

$$\begin{aligned} \{\mathbf{Q}_k^{3D} = \omega \mathbf{C}^{3D} \in \mathbb{R}^{3 \times 3}, \forall \omega \in \{0.5, 0.2, 0.1, 0.1, 0.05\}, \\ \forall \mathbf{C}^{3D} \in \{\mathbf{I}, -\mathbf{I}, \mathbf{M}_a, -\mathbf{M}_a, \mathbf{M}_b, -\mathbf{M}_b, \mathbf{M}_c, -\mathbf{M}_c, \\ \mathbf{M}_d, -\mathbf{M}_d, \mathbf{M}_e, -\mathbf{M}_e, \mathbf{M}_f, -\mathbf{M}_f\}\}. \end{aligned} \quad (22)$$

The final QUBOP and the unembedding (*i.e.*, decoding the solution to QUBOP) after quantum annealing for the 3D case are obtained similarly to (14)–(17) with $\mathbf{q} \in \mathbf{B}^{81}$ (\mathbf{q}_0 remains fixed to $|1\rangle$) and \mathbf{Q}_k are replaced by \mathbf{Q}_k^{3D} in (17)).

6. Quantum Point Set Registration

In point set registration, the input point sets are of different cardinalities, and correspondences between points are, generally, not known, *i.e.*, $[\mathbf{x}_n] \in \mathbf{X} \in \mathbb{R}^{D \times N}$ and $[\mathbf{y}_m] \in \mathbf{Y} \in \mathbb{R}^{D \times M}$, $m \in \{1, \dots, M\}$. N and M are the numbers of points in the reference and template, respectively, while D is the point dimensionality. The objective of point set alignment is to recover rotation \mathbf{R} ($\mathbf{R}^{-1} = \mathbf{R}^T$, $\det(\mathbf{R}) = 1$) and translation \mathbf{t} aligning \mathbf{Y} to \mathbf{X} . We assume that the translation is resolved in the pre-processing step by bringing the point set centroids into coincidence.

Point set alignment can be alternately solved on AQC by finding some point matches and estimating the transformation with the given correspondences in the ICP fashion [13]. This would result in a sequence of QUBOP of the form (16). To express alignment as a single QUBOP, we have to find an energy functional which is correspondence-free and which, when minimised in one shot on AQC, would result in an optimal alignment. The desired form of the energy functional has been recently shown in the literature [29].

6.1. Particle Dynamics Based Alignment

Barnes-Hut Rigid Gravitational Approach (BHRGA) [29] is a recent point set alignment method with a single energy functional which remains unchanged during the entire optimisation. BHRGA is a globally multiply-linked approach, *i.e.*, all \mathbf{y}_m interact with all \mathbf{x}_n . In [29], point sets are aligned by minimising the mutual *gravitational potential energy* (GPE) \mathbf{E} of the corresponding system of particles in the force field induced by \mathbf{X} :

$$\mathbf{E}(\mathbf{R}, \mathbf{t}) = \sum_m \sum_n \mu_{\mathbf{y}_m} \mu_{\mathbf{x}_n} \|\mathbf{R} \mathbf{y}_m + \mathbf{t} - \mathbf{x}_n\|_2, \quad (23)$$

where $\mu_{\mathbf{y}_m}$ and $\mu_{\mathbf{x}_n}$ denote masses of \mathbf{y}_m and \mathbf{x}_n , respectively. With no imposed boundary conditions, particles are initialised with unit masses. In [29], (23) is optimised with the Levenberg-Marquardt algorithm [41, 44], and the optimum is achieved when the system's GPE is locally minimal. Without acceleration by a 2^D -tree, the method has quadratic complexity and (23) involves all possible interactions between the template and reference points.

We can now derive a QUBOP in the similar fashion as in Sec. 5 for the transformation estimation. Note, however, that the bases (13) and (22) allow for affine transformations and scaling. Thus, implicitly, we would optimise

$$\mathbf{E}(\mathbf{R}, \mathbf{t}, s) = \sum_m \sum_n \mu_{\mathbf{y}_m} \mu_{\mathbf{x}_n} \|\mathbf{R} \mathbf{y}_m s + \mathbf{t} - \mathbf{x}_n\|_2, \quad (24)$$

where the scalar s is the scaling of the template. As proven in [28], *allowing for scale in globally multiply-linked point set alignment results in the shrinkage of the template to a single point with a very high probability*. To remedy the

problem, either prior correspondences can be used, or point interactions can be restricted to local vicinities [28]. In our QA, we opt for the second solution which allows to use the rotational bases (13) and (22) elaborated in Sec. 5. Eventually, the $\Phi \in \mathbb{R}^{(K+1) \times (D)(L(1)+L(2)+\dots+L(N))}$ matrix encoding point interactions for point set alignment reads

$$\Phi = [\Phi_1 \Phi_2 \dots \Phi_N], \quad (25)$$

with Φ_n , $n \in \{1, \dots, N\}$, of the form

$$\begin{bmatrix} \mathbf{x}_n^\top & \mathbf{x}_n^\top & \dots & \mathbf{x}_n^\top \\ -[\mathbf{Q}_1 \mathbf{y}_1^n]^\top & -[\mathbf{Q}_1 \mathbf{y}_2^n]^\top & \dots & -[\mathbf{Q}_1 \mathbf{y}_{L(n)}^n]^\top \\ -[\mathbf{Q}_2 \mathbf{y}_1^n]^\top & -[\mathbf{Q}_2 \mathbf{y}_2^n]^\top & \dots & -[\mathbf{Q}_2 \mathbf{y}_{L(n)}^n]^\top \\ \vdots & \vdots & \ddots & \vdots \\ -[\mathbf{Q}_K \mathbf{y}_1^n]^\top & -[\mathbf{Q}_K \mathbf{y}_2^n]^\top & \dots & -[\mathbf{Q}_K \mathbf{y}_{L(n)}^n]^\top \end{bmatrix}, \quad (26)$$

with \mathbf{Q}_k being as in (13) or (22) for the 2D and 3D case, respectively. Φ_1 , Φ_2 and Φ_N encode point interactions between every \mathbf{x}_n and corresponding $L(n) \ll M$ points of the template denoted by superscripted $\{\mathbf{y}_1^n, \mathbf{y}_2^n, \dots, \mathbf{y}_{L(n)}^n\}$. Note that the latter build N subsets of $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M\}$ of different cardinalities $L(n)$, \bar{L} on average. If different \mathbf{x}_n interact with the same \mathbf{y}_m , the corresponding subcolumns $[[\mathbf{Q}_1 \mathbf{y}_m]^\top [\mathbf{Q}_2 \mathbf{y}_m]^\top \dots [\mathbf{Q}_K \mathbf{y}_m]^\top]^\top$ of Φ_n can be computed only once and reused. The final QUBOP for point set alignment with Φ_n as in (26) reads

$$\arg \min_{\mathbf{q} \in \mathbb{B}^{K+1}} \mathbf{q}^\top \Phi \Phi^\top \mathbf{q}. \quad (27)$$

In total, $K + 1 = 21$ and $K + 1 = 81$ qubits are required to align point sets on AQC in the 2D and 3D case, respectively. Both transformation estimation and point set alignment need the same number of qubits in the same dimensions, and the difference lies in the complexity to construct \mathbf{P} (see Sec. 6.2). Note that if the same template has to be aligned to multiple references, the corresponding Φ can be obtained by reusing $[[\mathbf{Q}_1 \mathbf{y}_m]^\top [\mathbf{Q}_2 \mathbf{y}_m]^\top \dots [\mathbf{Q}_K \mathbf{y}_m]^\top]^\top$ (which has to be computed only once). The first qubit of \mathbf{q} has to be fixed to $|1\rangle$, since the first element of every column contains a reference point which has to be active during the entire optimisation. The unembedding is performed similarly to the case of transformation estimation, see Sec. 5.

6.2. Complexity to Prepare $\mathbf{P} = \Phi \Phi^\top$

To prepare Φ , $\mathcal{O}(KDN\xi)$ and $\mathcal{O}(KDN\bar{L}\xi)$ operations are required for the transformation estimation and point set alignment, respectively. ξ denotes the number of operations for multiplying \mathbf{y}_m with one element of the additive basis \mathbf{Q}_k . To obtain the final \mathbf{P} , we need to transpose Φ and multiply Φ with Φ^\top which, in the worst case, takes $\mathcal{O}(K^2DN)$ operations for the transformation estimation and $\mathcal{O}(K^2DN\bar{L})$ operations for the point set alignment. There are also slightly faster algorithms for matrix multiplication compared to the naïve way [20].

	TE	K				
		10	20	30	40	50
e_{2D}	0.023	0.026	0.041	0.078	0.17	0.3
σ_{2D}	0.012	0.013	0.012	0.012	0.012	0.013
$e_{\mathbf{R}}$	0.058	0.062	0.083	0.22	0.47	0.764
$\sigma_{\mathbf{R}}$	0.041	0.044	0.041	0.036	0.031	0.03

Table 2: The accuracy of QA under random initial misalignments, for the transformation estimation (" TE ") and point set alignment ($K > 1$).

7. Experimental Evaluation

The current generation of D-WAVE annealers does not support the precision of weights in \mathbf{P} necessary for our method [22]³. It is foreseeable that future generations will enable a higher accuracy for couplings. We thus implement and test QA with an AQC sampler on a conventional computer (Intel *i7-6700K* CPU with 32GB RAM). All quantitative tests are performed with 21 binary variables corresponding to the size of the \mathbf{Q} basis in 2D.

We report two error metrics, *i.e.*, the alignment error e_{2D} and the transformation discrepancy $e_{\mathbf{R}}$, together with their standard deviations denoted by σ_{2D} and $\sigma_{\mathbf{R}}$, respectively. The alignment error $e_{2D} = \frac{\|\mathbf{R}\mathbf{Y} - \mathbf{X}\|_{\mathcal{HS}}}{\|\mathbf{X}\|_{\mathcal{HS}}}$ ($\|\cdot\|_{\mathcal{HS}}$ denotes the Hilbert-Schmidt norm) measures how accurately the aligned shape coincides with the reference and requires ground truth correspondences. The transformation discrepancy is defined as $e_{\mathbf{R}} = \|\mathbf{I} - \mathbf{R}\mathbf{R}^\top\|_{\mathcal{HS}}$, where \mathbf{R} is the recovered rotation. It measures how closely the recovered transformation resembles a valid rigid transformation. The usage of two complementary metrics is necessary because a low $e_{\mathbf{R}}$ does not automatically imply an accurate registration. On the other hand, a low e_{2D} does not quantify how rigid the recovered transformation is.

Datasets and Proof of Concept. We use four 2D datasets, *i.e.*, *fish* [47], *qubit*, *kanji* and *composer* with cardinalities varying from 91 (*fish*) to 7676 (*composer*), see Fig. 1 for qualitative registration results. For point sets with up to a few thousand points, the simulation time $\tau_{\mathbf{P}} < 1$ sec. For $\sim 7.7k$, $\tau_{\mathbf{P}}$ grows to 20.178 sec (by a factor of $\sim 10^4$). Simulation with $n = 30$ takes already ~ 2.5 days. More binary variables allow for more elements in the basis \mathbf{Q} resulting in more accurate alignment. Note that even with 80 qubits, *i.e.*, for problems with $n = 80$, annealing on AQC takes around 100 *ms*. A simulation with $n = 80$ is not possible even on a conventional supercomputer in a reasonable time.

Initial Misalignment and Point Linking. We test how accurately our method recovers the transformation under the random angle of initial misalignment θ and the different size of the point linking region. We generate 500 random transformations in the range $\theta \in [0; 2\pi]$ of the *fish* dataset and resolve them with QA, for each $K \in \{1, 10, 20, 30\}$. The results are summarised in Table 2. We see that e_{2D} corre-

³the current generation natively supports 9-bit floating-point numbers

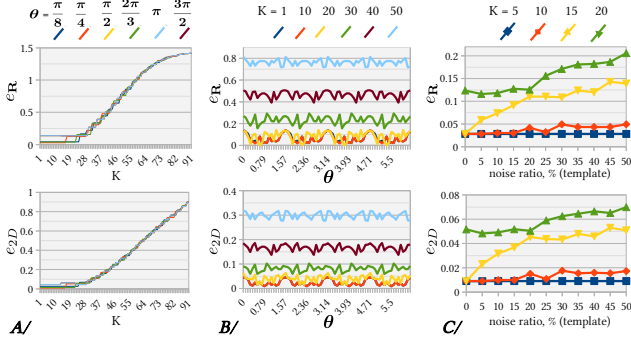


Figure 3: The metrics as the functions of *A*/: the size of the point interaction region parametrised by K ; *B*/: the angle of initial misalignment θ ; *C*/: the template noise ratio.

lates with $e_{\mathbf{R}}$ for all tested K . For $K = 30$ — which corresponds to one third of the template points — both metrics are still comparably low. We also study how the choice of the point interaction region or K affects the accuracy of the transformation recovery and plot e_{2D} and $e_{\mathbf{R}}$ as the functions of K for several angles of initial misalignment θ in Fig. 3-A. Interacting points are determined with the K nearest neighbour rule for each \mathbf{x}_n . Recall that according to the singularity theorem [28], the globally multiply-linked alignment (here, $K = 91$) results in a shrinkage of the template to a single point, which is observed experimentally.

Next, we systematically vary the angle of initial misalignment θ in the range $[0; 2\pi]$ with the angular step $\frac{\pi}{36}$ and report e_{2D} and $e_{\mathbf{R}}$ as the functions of θ , for $K \in \{1, 10, 20, 30, 40, 50\}$. This test reveals the differences in the transformations caused by θ , which arise due to the composition and the expressiveness of the chosen basis M , see Fig. 3-B. QA is almost agnostic to θ , which is a desirable property of every point set alignment method.

Sensitivity to Noise. We systematically add uniformly distributed noise to the template and test the robustness of the proposed QA to outliers in the data, since real data often contains outliers. The highest template noise ratio amounts to 50%. Each metric for every noise ratio and every K is averaged over 50 runs, see Fig. 3-C. $\sigma_{\mathbf{R}}$ and σ_{2D} do not exceed 0.057 and 0.03, respectively. We observe both the increasing alignment error and the discrepancy in the obtained transformations with the increasing noise level. For small K , nonetheless, even large noise ratios seem not to influence the metrics significantly.

Spectral Gap Analysis. Spectral gap $\Delta(\hat{\mathcal{H}})$ is the difference between the energy of the ground state and the second-lowest eigenstate. Each problem has an intrinsic and unique $\Delta(\hat{\mathcal{H}})$. Even though a rigorous analysis of the spectral gap is out of the scope of this paper, we make several qualitative observations about the energy landscape of QA, the difference in the energy values and the corresponding registrations for one exemplary problem. In Fig. 4, we plot the se-

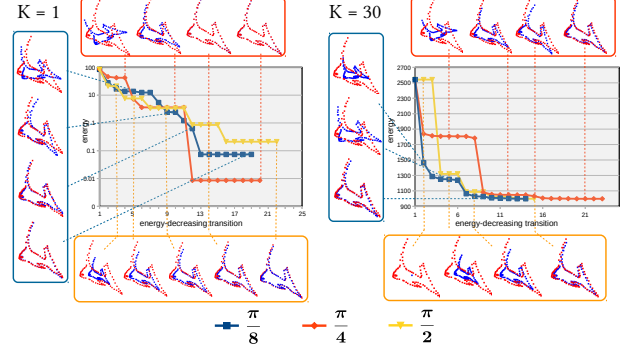


Figure 4: The sequences of energy-decreasing transitions and the corresponding energy values observed in our sampler, for transformation estimation ($K = 1$) and point set alignment with $K = 30$ interactions per \mathbf{x}_n . Besides the graphs, we visualise alignment results for selected energy values and the angle of initial misalignment $\theta \in \{\frac{\pi}{8}, \frac{\pi}{4}, \frac{\pi}{2}\}$.

quences of energy-decreasing transitions together with the energy values in the experiment with *fish*, for three θ values. We notice that some solutions have very small differences in the energies and are qualitatively indistinguishable from each other. This is accounted for by the choice of the additive basis, *i.e.*, that the same alignment can be encoded in different ways. In contrast, we see significant differences in the energy values of the qualitatively different solutions (orders of magnitudes larger in the analysed experiment).

We conclude that even though $\Delta(\hat{\mathcal{H}})$ is small, the alignments corresponding to several few lowest eigenstates are qualitatively similar. This suggests that our selection of the basis leads to problems with sufficient spectral gaps.

8. Conclusions

This paper introduces AQC for the computer vision community and shows that fundamental low-level problems can be brought to a representation suitable for solving on AQC. In simulations on a classical computer and in a wide range of scenarios, our QA is shown to successfully recover 2D transformations which are close approximations of globally optimal transformations. With the chosen basis of 20 elements, the solutions result in low transformation discrepancy and alignment errors. Observations on how to avoid singularities as well as the noise sensitivity and spectral gap analysis complement the experimental section.

In future work, our technique can be extended to affine transformations and other related computer vision problems. We hope to see more research on computer vision methods with quantum hardware in the next decades.

Acknowledgements. This work was supported by the ERC Consolidator Grant 770784. VG is grateful to Polina Matveeva for many enlightening discussions on the physical foundations of adiabatic quantum computing. The authors thank Bertram Taetz and Hanno Ackermann for reviewing an earlier version of this paper.

A. Appendix

In this additional section, we provide details on the selection of the annealing rate, analyse the structure of \mathbf{P} and formalise the unembedding, *i.e.*, the conversion of the solution to QUBOP (16) to the solution of the original alignment problem on point sets. We preserve the notations referring to the sections and equations from the main matter. The equations and the figure introduced in this supplement are equipped with Roman numerals.

Annealing Rate. Suppose $E_n(s)$ is the ground state of instantaneous Hamiltonian, $E_n(0)$ is the initial state (ground state) of the system and $E_m(s)$ is any other excited state of the instantaneous Hamiltonian. Let $s = \frac{t}{T} \in [0; 1]$, where T is the overall time of interpolation and t is physical time. Then, according to [4], T has to be chosen so that

$$T \gg \frac{|\langle E_m(s) | dH/ds | E_n(s) \rangle|}{E_{nm}(s)^2}, \quad \forall m \neq n, \quad (\text{i})$$

where dH/ds is the rate of change of Hamiltonian with respect to s and E_{nm} is the difference in the corresponding instantaneous energies.

Analysis of \mathbf{P} . Fig. 1 visualises several exemplary weight matrices \mathbf{P} from the experiments with clean and noisy data (see Sec. 7). There are several observations. First, $\mathbf{P} = \Phi\Phi^\top$ is symmetric upon algorithm design. We also see that the columns of Φ can be arbitrarily reshuffled as long as the correspondences are preserved⁴. Second, \mathbf{P} contains regularly arranged zero submatrices, due to our choice of the basis. As soon as a row of Φ induced by $q\mathbf{C}_I$, where $\mathbf{C}_I \in \{\mathbf{I}, -\mathbf{I}\}$, is multiplied by a column of Φ^\top induced by $q\mathbf{C}_M$, where $\mathbf{C}_M \in \{\mathbf{M}, -\mathbf{M}\}$, and vice versa, we obtain a zero entry in \mathbf{P} . The reason is that

$$\begin{cases} [\mathbf{I} \sum_i \mathbf{y}_i]^\top & [\mathbf{M} \sum_j \mathbf{y}_j] = 0 \\ [-\mathbf{I} \sum_i \mathbf{y}_i]^\top & [\mathbf{M} \sum_j \mathbf{y}_j] = 0 \\ [\mathbf{I} \sum_i \mathbf{y}_i]^\top & [-\mathbf{M} \sum_j \mathbf{y}_j] = 0 \\ [-\mathbf{I} \sum_i \mathbf{y}_i]^\top & [-\mathbf{M} \sum_j \mathbf{y}_j] = 0 \end{cases}, \quad (\text{ii})$$

if $\sum_i \mathbf{y}_i = \sum_j \mathbf{y}_j$, which holds in our case since each row of Φ except the first row includes all points of \mathbf{Y} multiplied by a single basis element \mathbf{Q}_k (see Fig. 1-(top left) for \mathbf{C} pairs resulting in zero matrices). Third, the structure of \mathbf{P} reflects that its diagonal entries encode biases, and non-diagonal elements represent couplings between the qubits.

With the increasing K , the span of the absolute energy values increases, due to the higher number of point interactions. As expected, \mathbf{P} depends on data and the angle of initial misalignment between the point sets. For all possible inputs and initial conditions — point sets of different cardinalities, K and θ — the structure of \mathbf{P} is the same for

⁴a reshuffling of rows requires changing the order of elements in \mathbf{Q}

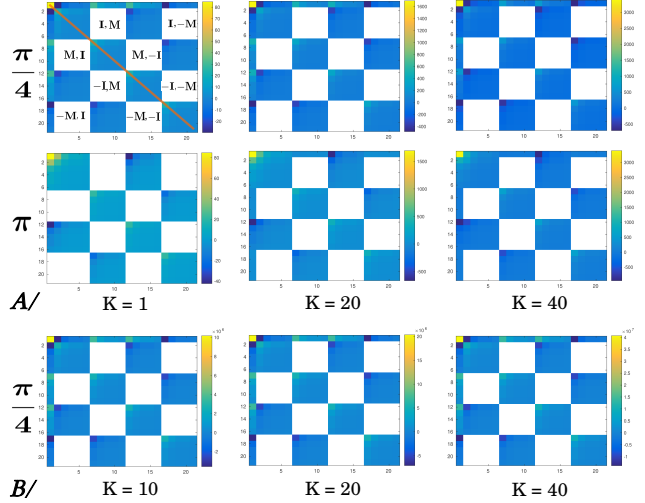


Figure 1: Exemplary visualisations of the weight matrix $\mathbf{P} = \Phi\Phi^\top$ in the experiment with clean (A/) and noisy data with 35% of outliers in the template (B/), for $K \in \{1, 10, 20, 40\}$ and $\theta \in \{\frac{\pi}{4}, \pi\}$. The colour scheme and the range of energy values are given to the right of each \mathbf{P} . White colour stands for zero entries. The diagonal values in \mathbf{P} represent biases (marked in orange on the top left), and non-zero elements represent couplings between the qubits. In the visualisation on the top left, we list the pairs of $\mathbf{C} \in \{\mathbf{I}, \mathbf{M}, -\mathbf{I}, -\mathbf{M}\}$ eventually leading to zero matrices.

the chosen basis. From \mathbf{P} , we also recognise that the considered alignment problem is not purely combinatorial and requires high-precision weights $J_{j,k}$ in (5).

Unembedding. Unembedding is the decoding of the solution to QUBOP (16) to the solution of the original alignment problem. Upon the design, our QA method assembles the entries of the transformation matrix in the additive basis \mathbf{Q}_k (see Secs. (5.1)–(6.1)). Suppose $\hat{\mathbf{q}}$ is the measurement result of \mathbf{q} , *i.e.*, it is a classical bitstring with $K + 1$ elements. Recall that \mathbf{q}_1 is reserved for reference points and does not contribute to the assembly of the transformation. Once $\hat{\mathbf{q}}$ is measured and returned, we obtain the corresponding transformation \mathbf{R} by summing up \mathbf{Q}_k multiplied by $\hat{\mathbf{q}}_{k+1}$:

$$\mathbf{R} = \sum_k \hat{\mathbf{q}}_{k+1} \mathbf{Q}_k. \quad (\text{iii})$$

The obtained \mathbf{R} is an affine transformation. If the solution has to represent a valid rotation matrix \mathbf{R}_r , \mathbf{R} can be projected to the rotation group by solving the *closest orthogonal approximation problem with constraints*:

$$\begin{aligned} \min \|\mathbf{R}_r - \mathbf{R}\|_{\mathcal{HS}}^2, \\ \text{s. t. } \mathbf{R}_r^{-1} = \mathbf{R}_r^\top \text{ and } \det(\mathbf{R}_r) = 1. \end{aligned} \quad (\text{iv})$$

For a solution to (iv) by singular value decomposition, see [33].

References

- [1] Steven H. Adachi and Maxwell P. Henderson. Application of Quantum Annealing to Training of Deep Neural Networks. *arXiv e-prints*, 2015. 5
- [2] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM J. Comput.*, 37(1):166–194, 2007. 1
- [3] Andris Ambainis and Robert Špalek. Quantum algorithms for matching and network flows. In *STACS*, 2006. 5
- [4] Mohammad H. S. Amin. Consistency of the adiabatic theorem. *Physical Review Letters*, 102:220401, 2009. 4, 9
- [5] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *International Conference on Computer Vision (ICCV) Workshops*, pages 1626–1633, 2011. 5
- [6] Caglar Aytekin, Serkan Kiranyaz, and Moncef Gabbouj. Automatic object segmentation by quantum cuts. In *International Conference on Pattern Recognition (ICPR)*, 2014. 5
- [7] Francisco Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253, 1982. 4
- [8] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980. 1
- [9] Charles H. Bennett, François Bessette, Gilles Brassard, Louis Salvail, and John Smolin. Experimental quantum cryptography. *J. Cryptol.*, 5(1):3–28, 1992. 5
- [10] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *International Conference on Computers, Systems, and Signal Processing*, 1984. 5
- [11] Aleta Berk Finnila, Maria Gomez, C. Sebenik, Catherine Stenson, and Jimmie D. Doll. Quantum annealing: A new method for minimizing multidimensional functions. *Chemical Physics Letters*, 219, 1994. 3
- [12] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. In *Symposium on Theory of Computing (STOC)*, 1993. 5
- [13] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 14(2):239–256, 1992. 2, 6
- [14] Zhengbing Bian, Fabian Chudak, Robert Brian Israel, Brad Lackey, William G. Macready, and Aidan Roy. Mapping constrained optimization problems to quantum annealing with application to fault diagnosis. *Frontiers in ICT*, 3:14, 2016. 1
- [15] Zhengbing Bian, Fabian Chudak, William G. MacReady, Lane Clark, and Frank Gaitan. Experimental determination of ramsey numbers with quantum annealing. *Physical Review Letters*, 111(13):130505, 2013. 1
- [16] Tomas Boothby, Andrew D. King, and Aidan Roy. Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing*, 15(1):495–508, 2016. 4
- [17] Max Born and Vladimir Fock. Beweis des adiabatsatzes. *Zeitschrift für Physik*, 51(3):165–180, 1928. 1, 3, 4
- [18] 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(2), 2017. 5
- [19] Patrick J. Coles et al. Quantum Algorithm Implementations for Beginners. *arXiv e-prints*, 2018. 1
- [20] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251 – 280, 1990. 7
- [21] D-Wave Systems. D-Wave Announces D-Wave 2000Q Quantum Computer and First System Order. <https://www.dwavesys.com/press-releases/...>, 2019. online; accessed 15 October 2019. 1
- [22] D-Wave Systems. Technical Description of the D-Wave Quantum Processing Unit. https://docs.dwavesys.com/docs/latest/doc_qpu.html, 2019. online; accessed 5 November 2019. 4, 7
- [23] 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, 2016. 1
- [24] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London Series A*, 439(1907):553–558, 1992. 5
- [25] Simon J. Devitt. Performing quantum computing experiments in the cloud. *Phys. Rev. A*, 94:032329, 2016. 1
- [26] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Laplan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science*, 292(5516):472–475, 2001. 1, 2, 3, 4
- [27] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982. 1, 4, 5
- [28] Vladislav Golyanik and Christian Theobalt. Optimising for scale in globally multiply-linked gravitational point set registration leads to singularities. In *International Conference on 3D Vision (3DV)*, 2019. 6, 7, 8
- [29] Vladislav Golyanik, Christian Theobalt, and Didier Stricker. Accelerated gravitational point set alignment with altered physical laws. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 6
- [30] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996. 1, 2, 5
- [31] Kuk-Hyun Han and Jong-Hwan Kim. Genetic quantum algorithm and its application to combinatorial optimization problem. In *Congress on Evolutionary Computation*, 2000. 5
- [32] Kuk-Hyun Han and Jong-Hwan Kim. Quantum-inspired evolutionary algorithm for a class of combinatorial optimization. *Trans. Evol. Comp.*, 6(6):580–593, 2002. 5
- [33] Nicholas J. Higham. Matrix nearness problems and applications. *Applications of Matrix Theory*, 22, 1989. 9

- [34] Berthold K. P. Horn, Hugh M. Hilden, and Shahriar Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *J. Opt. Soc. Am. A*, 5(7):1127–1135, 1988. 2
- [35] Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, 1925. 3
- [36] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, 1998. 3
- [37] Bruce E. Kane. A silicon-based nuclear spin quantum computer. *Nature*, 393:133137, 1998. 2
- [38] Amir Khoshaman, Walter Vinci, Brandon Denis, Evgeny Andriyash, Hossein Sadeghi, and Mohammad H. Amin. Quantum variational autoencoder. *Quantum Science and Technology*, 4(1), 2018. 5
- [39] Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. 2, 3
- [40] Trevor Lanting et al. Entanglement in a quantum annealing processor. *Phys. Rev. X*, 4:021041, 2014. 1
- [41] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Applied Mathematics*, II(2):164–168, 1944. 6
- [42] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10, 2013. 1, 2, 5
- [43] Yuri Manin. *Computable and Noncomputable*. Sov. Radio., 1980. 1, 4
- [44] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963. 6
- [45] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. 3
- [46] Andrea Morello et al. Single-shot readout of an electron spin in silicon. *Nature*, 467:687691, 2010. 2
- [47] Andriy Myronenko and Xubo Song. Point-set registration: Coherent point drift. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010. 1, 2, 7
- [48] Kae Nemoto, Michael Trupke, Simon J. Devitt, Ashley M. Stephens, Burkhard Scharfenberger, Kathrin Buczak, Tobias Nöbauer, Mark S. Everitt, Jörg Schmiedmayer, and William J. Munro. Photonic architecture for scalable quantum information processing in diamond. *Phys. Rev. X*, 4:031022, 2014. 1
- [49] Florian Neukart, Gabriele Compostella, Christian Seidel, David von Dollen, Sheir Yarkoni, and Bob Parney. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:29, 2017. 1, 2
- [50] 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. 5
- [51] Hartmut Neven, Geordie Rose, and William G. Macready. Image recognition with an adiabatic quantum computer I. Mapping to quadratic unconstrained binary optimization. *arXiv e-prints*, 2008. 5
- [52] Nga T. T. Nguyen and Garrett T. Kenyon. Image classification using quantum inference on the D-Wave 2X. *arXiv e-prints*, 2019. 5
- [53] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2011. 2, 4
- [54] 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, 2018. 5
- [55] Ojas Parekh, Jeremy Wendt, Luke Shulenburg, Andrew Landahl, Jonathan Moussa, and John Aidun. Benchmarking Adiabatic Quantum Optimization for Complex Network Analysis. *arXiv e-prints*, 2016. 4
- [56] James L. Park. The concept of transition in quantum mechanics. *Foundations of Physics*, 1(1):23–33, 1970. 5
- [57] Rudolph Peierls. On ising’s model of ferromagnetism. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32(3):477–481, 1936. 3
- [58] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical Review Letters*, 113, 07 2013. 1, 5
- [59] Eleanor Rieffel and Wolfgang Polak. An introduction to quantum computing for non-physicists. *ACM Computing Surveys (CSUR)*, 32(3):300–335, 2000. 1, 2
- [60] Tobias Schmitt-Manderbach, Henning Weier, Martin Fürst, Rupert Ursin, Felix Tiefenbacher, Thomas Scheidl, Josep Perdigues, Zoran Sodnik, Christian Kurtsiefer, John G. Rarity, Anton Zeilinger, and Harald Weinfurter. Experimental demonstration of free-space decoy-state quantum key distribution over 144 km. In *European Conference on Lasers and Electro-Optics and the International Quantum Electronics Conference*, 2007. 5
- [61] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. 1, 2, 5
- [62] Peter W. Shor. Introduction to Quantum Algorithms. *arXiv e-prints*, 2001. 4
- [63] Janus H. Wesenberg, Arzhang Ardavan, G. Andrew D. Briggs, John J. L. Morton, Robert J. Schoelkopf, David I. Schuster, and Klaus Mølmer. Quantum computing with an electron spin ensemble. *Phys. R. Lett.*, 103:070502, 2009. 2
- [64] William K. Wootters and Wojciech Hubert Zurek. A single quantum cannot be cloned. *Nature*, 299(5886), 1982. 5
- [65] Richard Y. Li, Rosa Di Felice, Remo Rohs, and Daniel Lidar. Quantum annealing versus classical machine learning applied to a simplified computational biology problem. *npj Quantum Information*, 4:14, 02 2018. 1, 2
- [66] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(11):2241–2254, 2016. 2
- [67] Floris A. Zwanenburg, Andrew S. Dzurak, Andrea Morello, Michelle Y. Simmons, Lloyd C. L. Hollenberg, Gerhard Klimeck, Sven Rogge, Susan N. Coppersmith, and Mark A. Eriksson. Silicon quantum electronics. *Rev. Mod. Phys.*, 85:961–1019, 2013. 2