



Universiteit
Leiden
The Netherlands

The travelling salesperson problem and the challenges of near-term quantum advantage

Smith-Miles, K.A.; Hoos, H.H.; Wang, H.; Bäck, T.H.W.; Osborne, T.J.

Citation

Smith-Miles, K. A., Hoos, H. H., Wang, H., Bäck, T. H. W., & Osborne, T. J. (2025). The travelling salesperson problem and the challenges of near-term quantum advantage. *Quantum Science And Technology*, 10(3). doi:10.1088/2058-9565/add61d

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/4260086>

Note: To cite this publication please use the final published version (if applicable).

TOPICAL REVIEW • OPEN ACCESS

The travelling salesperson problem and the challenges of near-term quantum advantage

To cite this article: Kate A Smith-Miles *et al* 2025 *Quantum Sci. Technol.* **10** 033001

View the [article online](#) for updates and enhancements.

You may also like

- [Quantum annealing for industry applications: introduction and review](#)
Sheir Yarkoni, Elena Raponi, Thomas Bäck et al.
- [An integrated programming and development environment for adiabatic quantum optimization](#)
T S Humble, A J McCaskey, R S Bennink et al.
- [An algorithm for solving a 1D seismic data inversion problem on quantum or digital annealers in the presence of a priori information on layer parameters specified by arbitrary functions](#)
Nickolay V Maletin, Anastasiia M Eremenko and Dmitry V Minaev

Quantum Science and Technology



TOPICAL REVIEW

OPEN ACCESS

RECEIVED
24 December 2024

REVISED
22 April 2025

ACCEPTED FOR PUBLICATION
8 May 2025

PUBLISHED
16 May 2025

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



The travelling salesperson problem and the challenges of near-term quantum advantage

Kate A Smith-Miles^{1,*} , Holger H Hoos² , Hao Wang³ , Thomas Bäck³ and Tobias J Osborne⁴

¹ School of Mathematics and Statistics, University of Melbourne, Parkville 3010, Australia

² Department of Computer Science, RWTH Aachen University, 52060 Aachen, Germany

³ Leiden Institute of Advanced Computer Science (LIACS), Universiteit Leiden, 2330 Leiden, The Netherlands

⁴ Institut für Theoretische Physik, Leibniz Universität Hannover, Appelstraße 2, 30167 Hannover, Germany

* Author to whom any correspondence should be addressed.

E-mail: smith-miles@unimelb.edu.au

Keywords: travelling salesperson problem (TSP), quantum computing, quantum annealing, instance space analysis

Abstract

Over the last two decades, the travelling salesperson problem (TSP) has been cited as a benchmark problem to demonstrate the advantage of quantum computers over conventional computers. Its advantage is that it is a well-studied NP-hard optimisation problem that can be easily communicated to highlight the challenges of searching through an exponentially growing number of possible solutions to find the optimal solution. It is therefore a tempting problem to choose to explore quantum advantage. At what point, however, is a call made that quantum advantage is not likely, and efforts should be focused on other problems? This article challenges the continued use of the TSP as a benchmark for quantum optimisation methods—such as quantum annealing and gate-based quantum computing—that require the TSP to be formulated as a quadratic unconstrained binary optimisation (QUBO) problem. We offer explanations for why such quantum approaches are not well suited, nor competitive against state-of-the-art classical methods, for tackling the challenges of the TSP landscape, and we draw parallels with similar observations made almost four decades ago when QUBO-based neural networks proved to be uncompetitive for solving the TSP. After critically reviewing two decades of research effort to solve TSPs using QUBO-based quantum methods, we note a gradual shift in focus: from initial attempts to solve small sized TSPs with general-purpose QUBO-based quantum approaches, to growing evidence that competitiveness is only enhanced where TSP domain knowledge is integrated, via either modified formulations or hybridisation with TSP classical heuristics. We discuss the numerous challenges that must be overcome before QUBO-based quantum optimisers could ever be competitive with classical state-of-the-art TSP solvers. Acknowledging that there may be more promise for non-QUBO-based hybrid approaches, where quantum search accelerates components of conventional algorithms, we offer recommendations for how future studies should be conducted to compare fairly and rigorously any proposed quantum methods against state-of-the-art TSP solvers, or any classical optimisation method, when seeking to establish quantum advantage.

1. Introduction

The well-studied travelling salesperson problem (TSP) is often cited as an example of the potential of quantum advantage over conventional digital computers. In the TSP, given N locations (usually called ‘cities’) to be visited, the objective is to find a minimal cost round-trip visiting each location exactly once and returning to the starting point. Clearly, naïve enumeration of all possible $(N - 1)!$ tours becomes quickly intractable as N grows. Moreover, the TSP is NP-hard, and therefore believed to permit no conventional algorithm that can find provably optimal solutions for arbitrary instances within time polynomial in N . The popular press enjoys explaining the difficulty of solving the TSP by citing, for example, that a conventional computer would take over 1000 years to find the shortest tour through $N = 22$ cities (Wadhwa 2018). Such a

statement, while technically correct if a conventional computer was merely used to sequentially evaluate all possible 10^{19} tours, is misleading since it overlooks decades of research on solving TSPs using mathematical algorithms that significantly prune back the search space and enable rapid solution of very large-scale TSPs. The potential of quantum computing to solve the TSP faster than mere brute force enumeration though is argued by invoking the principles of Grover's algorithm (Grover 1997), which suggests a quantum search that simultaneously considers all possible solutions in a time that scales as $\sqrt{(N-1)!}$ and therefore promises considerable theoretical speed-up over conventional computers for large N , with the current record being $O^*(1.728^N)$ steps (Ambainis *et al* 2019, Severini 2022).

There are many problems therefore with this popular folklore comparison of the seeming inability of conventional computers to solve even small sized TSPs, and citing this as an example of the potential of quantum computers. It is misleading and contributes to unnecessary hype in several ways, including:

- (i) ignoring current state-of-the-art optimisation algorithms (Cook *et al* 2011) which can currently solve TSP instances with many thousands of cities to proven optimality on conventional computers, and even find solutions to problems with millions of cities within minutes, which should be the standard to which quantum advantage over conventional computers is compared;
- (ii) failing to acknowledge that the speed-ups promised by Grover's algorithm will not be effective unless hybridised with classical optimisation methods to reduce the size of the search space—this concept is theoretical only under certain conditions (Moylett *et al* 2017) and has not yet been demonstrated as effective (Ambainis *et al* 2019);
- (iii) not acknowledging that most quantum approaches for solving the TSP have little to do with Grover's algorithm. Instead they use quantum mechanics principles to seek minimum energy states of a constructed Hamiltonian function (Warren 2013a). However, there is no guarantee that this minimum energy state corresponds to a feasible and near-optimal solution of the TSP due to the complex landscape resulting from the required formulation as a quadratic unconstrained binary optimisation (QUBO) problem, which state-of-the-art methods are not constrained to adopt. This distinction makes quantum advantage for such QUBO-based TSP methods harder to reach, as will be discussed in section 3.

Unfortunately, this misleading hype has created an impression of the imminent potential of quantum computing for solving the TSP and other NP-hard optimisation problems, which is not supported by current scientific evidence.

What is the evidence to date about the potential of quantum computers to solve the TSP? Here we must distinguish between the near-term, where so-called noisy intermediate-scale quantum (NISQ) information processing devices are not capable of ideal quantum computations (Preskill 2018), and true scalable fault-tolerant quantum computers. It is worth emphasising that in the far future, when qubits are presumably as plentiful as classical bits, there is no controversy that quantum computers will offer significant quantum advantage for almost all computational problems since they can simulate classical computers, and yet benefit from the additional advantages of the quantum computational model (Nielsen and Chuang 2010).

Due to the current hardware limitations we are squarely in the NISQ era and large-scale problems cannot be solved, so most of the literature has focused on noise-free conventional simulations of quantum algorithms. In order to establish provable near-term *useful* quantum advantage for TSP, four key questions must be addressed:

ABILITY: Can quantum algorithms solve the TSP (i.e. find feasible and at least near-optimal solutions), even if only for a small number of cities, using current quantum hardware, or simulations of quantum approaches?

COMPETITIVENESS: If so, how does the trade-off between execution time (on a physical device) and solution quality compare to state-of-the-art algorithms (rather than brute force enumeration)?

SCALABILITY: How does the execution time and solution quality of simulated quantum approaches scale with increasing problem size compared to scaling of state-of-the-art algorithms?

HARDWARE COMPETITIVENESS: How is solution quality found using noisy quantum circuits expected to degrade compared to simulated performance?

The current literature has focused primarily on exploring the first question on the *ability* of quantum algorithms to solve small sized instances of the TSP, and the answer appears to be largely negative for QUBO-based quantum optimisers. An analysis of recent studies (Osaba *et al* 2022), further augmented in this paper with additional evidence, supports the view that quantum computers, while theoretically able to find feasible and near-optimal solutions to the TSP, are hampered from being able to do so in practice. The

challenges for quantum approaches to constrained optimisation problems like the TSP are comprehensively described in a recent survey paper (Abbas *et al* 2023). In this article we explore further why QUBO-based quantum optimisers are poorly suited to the TSP, drawing analogies from lessons from efforts in the 1980s (Hopfield and Tank 1985, Wilson and Pawley 1988) to solve TSPs using QUBO-based Hopfield neural networks. This ultimately lead to a call to abandon the TSP as a benchmark problem for neural networks (Smith 1996). We argue here that the time has come to recognise that the TSP is also not a useful benchmark for demonstrating imminent useful quantum advantage, failing to meet the criteria outlined in Abbas *et al* (2023) amongst other reasons as will be discussed in following sections.

The scope of this article is limited to the TSP, and we offer no detailed insights into the likely advantage of QUBO-based quantum optimisers or other quantum methods for solving other combinatorial optimisation problems. We note however that several other combinatorial optimisation problems have already been declared as unsuitable benchmarks for NISQ quantum algorithms, including quadratic assignment with constraints, 2-SAT, and linear assignment problems (Jattana *et al* 2020). Results on the unconstrained MAX-CUT problem appear to be slightly more promising (Klemm *et al* 2012, Sathe *et al* 2023), supporting the view that how quantum methods handle complex constraints, as exemplified by the TSP, is critical to their effectiveness (Abbas *et al* 2023).

2. Quantum optimisation of QUBO problems

In this section we briefly review the main strategy that has thus far been investigated for using quantum mechanics principles to solve general combinatorial optimisation problems, formulated as QUBOs, before discussing the effectiveness of this approach for solving the TSP in section 3. Broadly, quantum optimisation methods can be categorised into two main approaches:

- (i) *general-purpose quantum algorithms* to solve any combinatorial optimisation problem, formulated as a QUBO, using either quantum hardware alone or with a conventional computer outer-loop;
- (ii) *specialised hybridisations* coupling both conventional and quantum computers to accelerate the search of classical heuristics.

While the quest for general purpose quantum solvers that can solve QUBOs (including the TSP) has been underway for around two decades, it is only in more recent years that efforts have turned to quantum acceleration of existing classical optimisation methods. This development has partly emerged in recent years from the realisation that quantum advantage is unlikely to be realised by quantum computers working alone with a general-purpose algorithm that lacks combinatorial knowledge about the specific optimisation problem. Hybridisation has been proposed to accelerate classical algorithms based on dynamic programming (Ambainis *et al* 2019) or branch-and-bound approaches (Montanaro 2020) that recursively reduce the viable search space (structured search that exploits combinatorial knowledge), and then Grover's algorithm (unstructured search) can be integrated to accelerate the search for minimal cost solutions within a subset which iterates back into the classical algorithm. This strategy is currently only theoretical, with speed-ups through quantum search and quantum amplitude amplification promised under strict conditions, such as maximum node degrees that limit the theoretical speed-ups to 4-city TSP instances (Dörn 2007, Moylett *et al* 2017). While there are currently no experimental studies to support the applicability, effectiveness, or scalability of the concept, it is a promising direction that we will return to in section 6 when discussing future outlook.

2.1. Mapping QUBOs to hamiltonians for quantum optimisation

For the last two decades, most research efforts have been directed towards the development of general-purpose quantum methods for combinatorial optimisation problems formulated as a QUBO, as shown in equation (1):

$$\min f(\mathbf{x}) = \mathbf{x}^T \mathbf{Q} \mathbf{x} + w_1 (\mathbf{A} \mathbf{x} - \mathbf{b})^2 + w_2 \mathbf{x}^T (\mathbf{1} - \mathbf{x}) \quad (1a)$$

$$= \mathbf{x}^T [\mathbf{Q} + w_1 (\mathbf{A}^T \mathbf{A}) - w_2 \mathbf{I}] \mathbf{x} + \mathbf{x}^T (w_2 \mathbf{1} - 2w_1 \mathbf{A}^T \mathbf{b}) + w_1 \mathbf{b}^T \mathbf{b}. \quad (1b)$$

Here the first term in equation (1a) represents the quadratic objective function (cost) to be minimised, the second term is a penalty for violation of any constraints of the general form $\mathbf{A} \mathbf{x} = \mathbf{b}$, and the third term is a penalty to ensure each decision variable x_i is binary. Any binary solution that satisfies the constraints is called *feasible* and has zero penalty, regardless of the penalty term weights w_1 and w_2 . Any chosen optimisation algorithm—quantum or classical—aims to find a feasible solution with minimal cost (called an *optimal* solution). When the decision variables are binary, the fact that $x_i^2 = x_i$ means that any linear

components of the objective function are captured on the diagonal of the matrix $\mathbf{Q}$, so no additional linear term is required for the cost function. Expanding all three terms shows the QUBO can be expressed as a quadratic term and a linear term, with the final constant term unnecessary to retain since it has no effect on optimal decision variables, only the cost of the solution. Many constrained binary optimisation problems, including the TSP, can be expressed in this QUBO form through construction of suitable penalty functions. We will review such constructions for the TSP in section 3.

Once formulated as a QUBO, any optimisation problem can be solved using a variety of approaches including neural networks, meta-heuristics, and other approaches based on statistical mechanics or quantum mechanics. The pathway to methods based on statistical or quantum mechanics involves reformulating the QUBO as an Ising Spin model (Lucas 2014) with a Hamiltonian function that encodes both the objective function and penalty terms for constraint violation. In statistical mechanics, the variables are spin states of particles encoded as $s_i \in \{-1, 1\}$. The Ising Spin Model of ferromagnetism analyses the energy of a spin configuration of N particles, and finds the minimum energy state of a problem Hamiltonian:

$$H(\mathbf{s}) = \mathbf{s}^T \mathbf{J} \mathbf{s} + \mathbf{h}^T \mathbf{s} = \sum_{j=1}^N \sum_{i=j}^N J_{ij} s_i s_j + \sum_{i=1}^N h_i s_i \quad (2)$$

where J is the coupling strength between the particles and $\mathbf{h}$ is an external force acting on each particle. This is clearly equivalent to QUBO shown in equation (1), with a change of variable

$$x_i = (1 - s_i) / 2 \text{ such that } s_i = 1 \leftrightarrow x_i = 0 \text{ and } s_i = -1 \leftrightarrow x_i = 1$$

with comparison between the quadratic and linear terms in equations (1) and (2) enabling the relationships between $(\mathbf{J}, \mathbf{h})$ to be expressed in terms of $(\mathbf{Q}, \mathbf{A}, \mathbf{b}, w_1, w_2)$.

The quantum mechanics equivalent of an Ising spin model is to encode the variables based on their spin in different spatial dimensions, notably using the z -axis as σ_i^z (the Pauli- z operator acting on the i th qubit). Using quantum measurements of qubits, the problem Hamiltonian is expressed using the same Ising model as:

$$H(\sigma^z) = \sum_{j=1}^N \sum_{i=j}^N J_{ij} \sigma_i^z \sigma_j^z + \sum_{i=1}^N h_i \sigma_i^z. \quad (3)$$

Through suitable construction of $\mathbf{J}$ and $\mathbf{h}$, the minimisation of the problem Hamiltonian in equation (3) is equivalent to solving the corresponding QUBO in equation (1).

This mapping of a QUBO onto a quantum mechanics framework has led to the development of two competing methods (quantum annealing (QA)/Adiabatic algorithms and variational quantum algorithms (VQAs)) that use quantum mechanics to seek minimal energy states of the Hamiltonian function, which are intended (but not guaranteed) to be feasible and near-optimal solutions to the encoded combinatorial optimisation problem.

2.2. QA

QA is an instance of the quantum adiabatic algorithm (Farhi *et al* 2001), and is a quantum version of simulated annealing (SA) (Kirkpatrick *et al* 1983) for function minimisation, employing quantum fluctuations rather than thermal fluctuations, and solving Schrödinger's equation to find minimal energy states of a Hamiltonian function.

An initial Hamiltonian H_0 is defined, typically based on the Pauli- x operator acting on the i th qubit, as $H_0(\sigma^x) = \sum_{i=1}^N \sigma_i^x$. A final Hamiltonian function H_{OPT} is defined in the form of equation (3) via a QUBO formulation of the optimisation problem, encoding the cost of a solution and penalties for constraint violation. A feasible and optimal solution to the optimisation problem will correspond to a minimum energy state of H_{OPT} . In order to reach a minimum energy state of H_{OPT} , an N -spin quantum system is annealed over time as:

$$H(t) = A(t) H_0 + B(t) H_{\text{TSP}} \quad (4)$$

where $t \in [0, t_{\text{max}}]$, $A(0) = 1$, $A(t_{\text{max}}) = 0$, $B(0) = 0$, $B(t_{\text{max}}) = 1$. This process slowly transitions from finding minimum energy states of the initial Hamiltonian H_0 to the final Hamiltonian H_{OPT} provided that the annealing time is long enough, and with enough small discretisation steps such that the coefficient functions $A(t)$ and $B(t)$ are sufficiently smooth. In practice, QA uses a finite annealing time t_{max} , and convergence to a globally minimal energy state of H_{OPT} is not guaranteed since most combinatorial optimisation problems

create very complex energy landscapes plagued with local minima. Nor is there any guarantee that the global minimum corresponds to a feasible solution of the optimisation problem, depending on the Hamiltonian construction used and how well it penalises infeasible solutions via tuning w_1 and w_2 . The practical challenges to implementing this elegant theory are well acknowledged (Warren 2013b), as discussed later.

2.3. VQAs

Parameterised quantum circuits are a model for quantum computation constructed as a sequence of quantum gates and measurements of qubit states. When the parameters of these circuits are optimised they can be read off to reveal desired solutions. This is the general framework for VQAs (Cerezo *et al* 2021) which have an outer loop to optimise parameters using classical optimisation algorithms running on a conventional computer, feeding into an inner loop of quantum circuit measurement intended to minimise an energy function. One of the most popular VQA approaches is the quantum approximation optimisation algorithm (QAOA) (Farhi *et al* 2014), which constructs a sequence of circuits to coarsely approximate the time dependent Hamiltonians used in QA. QAOA employs a set of p time-independent operator pairs of the form $e^{-i\beta H_m}$ and $e^{-i\gamma H_{OPT}}$, where H_m is a mixing Hamiltonian that adds quantum amplitude between different solutions to escape local optima. QAOA then involves alternating application of these two operators as the depth of the circuit p is increased incrementally, from an efficiently chosen initial state, with the measurements of the Hamiltonian H_{OPT} on the final state providing a probabilistic distribution of the lowest energy found. Convergence to the minimum energy state of the final Hamiltonian H_{OPT} is theoretically guaranteed as $p \rightarrow \infty$, but $p < 30$ is usually adopted in practice (Pan *et al* 2022, Vizzuso *et al* 2023). However, this also assumes that the parameters β and γ have been optimised for each of the p layers, via a classical optimisation method (often Nelder–Mead Simplex, although other methods are used (Sung *et al* 2020, Fernández-Pendás *et al* 2022)), and thus QAOA is considered a hybrid approach for current quantum computer technology.

These two approaches (QAOA and QA) are discrete and continuous versions of similar ideas, and are considered equivalent to each other in the limit as the QAOA circuit depth $p \rightarrow \infty$ (Mizel *et al* 2007, Koßmann *et al* 2022), assuming that the QAOA problem Hamiltonian is the same as the one used for QA. A review of the literature reveals however, that there are often differences in how this Hamiltonian is formulated that impact whether infeasible solutions are likely to be encoded as minimum energy states (Warren 2021), in addition to the challenges of finding global minima due to the difficulties of parameter tuning (Fernández-Pendás *et al* 2022).

3. Can quantum computing solve the TSP?

While both QA and QAOA offer elegant theoretical solutions to combinatorial optimisation problems in general, in practice they have been far less effective at solving the TSP compared to other problems, particularly in comparison to unconstrained problems like MAX-CUT. The TSP is a particularly challenging constrained optimisation problem, most certainly for methods that are restricted to using a QUBO formulation, which is neither a natural nor competitive TSP formulation compared to those employed by state-of-the-art classical algorithms (Smith 1996, Salehi *et al* 2022). In the following sections we build a mathematical argument for why the TSP is a problem not well suited to solution via QUBO-based approaches, drawing connections with earlier attempts to solve TSPs via QUBO-based neural network approaches.

3.1. TSP formulations

Figure 1 presents two different TSP formulations: edge-based and node-based. The difference in these formulations starts with how the binary decision variables are defined for the set of N cities, with $D_{i,j}$ ($i, j \in \mathcal{N} = \{1, \dots, N\}$) denoting the distance between cities i and j . Consequential differences include the nature of the objective function (linear or quadratic), and how the notoriously difficult ‘sub-tour elimination’ constraints are handled to ensure valid tours. For example, for a 4-city TSP with cities A, B, C, D , a solution with one 2-city sub-tour $A \leftrightarrow B$ and another 2-city sub-tour $C \leftrightarrow D$ is not a valid TSP tour.

The optimal binary matrix $\mathbf{X}^* \in \{0, 1\}^{N \times N}$ is the one that minimises the cost (total distance) function subject to satisfaction of the constraints required for valid tours. While both formulations have the same linear assignment constraints requiring each city to be visited exactly once, the edge formulation requires the addition of challenging sub-tour elimination constraints that are avoided by the node formulation. The trade-off is that the node formulation suffers from a quadratic objective function which is most likely to be non-convex and plagued with many local minima for larger problem sizes.

The concept underlying the edge formulation is the one used by classical state-of-the-art exact solvers (such as CONCORDE (Applegate *et al* 2012)), offering provably optimal solutions) and inexact solvers (such

<i>Edge formulation:</i> $X_{i,j}^e = \begin{cases} 1 & \text{if city } i \text{ is followed by city } j \\ 0 & \text{otherwise} \end{cases}$ minimize $\sum_{i=1}^N \sum_{j=1}^N X_{i,j}^e D_{i,j}$ subject to $\sum_{i=1}^N X_{i,j}^e = 1 \quad \forall j \in \mathcal{N}$ $\sum_{j=1}^N X_{i,j}^e = 1 \quad \forall i \in \mathcal{N}$ $\sum_{i \in S_1} \sum_{j \in S_2} X_{i,j}^e \geq 1 \quad \forall S_1 \subset \mathcal{N}, S_2 = \mathcal{N} \setminus S_1$	<i>Node formulation:</i> $X_{i,j}^n = \begin{cases} 1 & \text{if city } i \text{ is in position } j \\ 0 & \text{otherwise} \end{cases}$ minimize $\sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N X_{i,j}^n D_{i,k} (X_{k,j+1}^n + X_{k,j-1}^n)$ subject to $\sum_{i=1}^N X_{i,j}^n = 1 \quad \forall j \in \mathcal{N}$ $\sum_{j=1}^N X_{i,j}^n = 1 \quad \forall i \in \mathcal{N}$
--	--

Figure 1. Two different formulations for the TSP: Edge-based and Node-based.

as LKH (Helsgaun 2000), offering rapid convergence to at least near-optimal solutions). Quantum approaches cannot use the edge formulation however, due to the inability of the sub-tour elimination constraints to be accommodated within a QUBO formulation. We refer the interested reader to appendix A for details of the mathematical argument, with an illustrative example for a 4-city TSP. It should be noted that the earliest theoretical models for QA did indeed manage to construct a QUBO Hamiltonian using the edge-based TSP formulation (Moser 2003, Martoňák *et al* 2004), but this was only possible due to a clever rewriting of the sub-tour elimination constraint that explicitly prohibits the single type of sub-tour (two 2-city loops) possible in the 4-city TSP they studied. This approach does not generalise to larger TSPs however (Heim *et al* 2017), with an exponential growth in the number of additional constraints required, offering only a necessary but not sufficient condition for a valid tour for $N > 6$ cities (Ruan *et al* 2020). Once larger TSP instances were attempted, it became clear that the TSP must be treated as a proper QUBO problem using the node-based formulation of the Hamiltonian (Warren 2013a, Lucas 2014), which avoids the complications of unintended sub-tours, and makes valid tours theoretically possible as minima of the energy function with appropriate parameter settings.

3.1.1. QUBO-based TSP Hamiltonian for quantum optimisation

The node-based formulation of the TSP is readily mapped to QUBO form as:

$$H_{\text{TSP}} = \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N X_{i,j}^n D_{i,k} (X_{k,j+1}^n + X_{k,j-1}^n) + w_1 \sum_{j=1}^N \left(\sum_{i=1}^N X_{i,j}^n - 1 \right)^2 + w_2 \sum_{i=1}^N \left(\sum_{j=1}^N X_{i,j}^n - 1 \right)^2 \quad (5)$$

where the penalty weights w_1 and w_2 need to be tuned to balance the tradeoff between efforts to minimise constraint violation penalties versus minimising costs that may allow infeasible solutions. In practice, these parameters have been difficult to determine to provide feasible and near-optimal TSP solutions for $N > 6$ (Salehi *et al* 2022, Villar-Rodriguez *et al* 2022).

Interestingly, the original theoretical model of Moser (2003) referred to the edge-based formulation as similar to the Hopfield-Tank approach for using a QUBO-based neural network to solve the TSP (Hopfield and Tank 1985), but this was incorrect since the Hopfield-Tank approach used the node-based formulation to enable valid tours to be found, subject to appropriate balancing of parameters weighting the cost and constraint penalty functions. These early QA papers (Moser 2003, Martoňák *et al* 2004) failed to acknowledge the fundamental difference between the edge-based and node-based formulation, and it would be a decade until the idea of using the node-based formulation as a QUBO representation for the TSP would be re-invented (Warren 2013a, Lucas 2014), with reference again to the neural network approach of Hopfield and Tank. It is a pity that a decade was lost on attempts to use the edge formulation, before it was realised it can never enforce valid tours without adding an exponentially growing number of additional constraints. It is also a pity that it has taken another decade to rediscover what the neural network community already reported (Wilson and Pawley 1988, Smith 1996, 1999) about the challenges of converging to global minima

of the node formulation to strike a balance between optimality and feasibility. We refer the interested reader to appendix B for a brief discussion of efforts to solve the TSP via QUBO-based Hopfield neural networks, and transferable lessons for quantum optimisation.

3.2. Results from two decades of QUBO-based TSP studies

Efforts over the last 20 years to solve the TSP using a general-purpose QUBO-based quantum optimiser (as opposed to quantum acceleration of classical TSP algorithms) are reflected in the outcomes summarised in table 1. We take as a starting point for this summary the systematic literature review (Osaba *et al* 2022) of quantum approaches to TSP, and consider all the papers identified as ‘milestone’ papers. These papers collectively describe much experimentation and discovery, but very little success in consistently finding feasible and near-optimal TSP tours, despite the optimistic conclusion of the review paper. For each paper we review the TSP formulation used (edge or node), and whether experimental simulations or hardware implementations were conducted. We also consider the nature of the TSP instances tested, whether they are synthetic easy instances, randomly generated, or come from the well-studied TSPLIB (a widely used library of benchmark instances for the TSP from real-world and random generators). Where results are available (using a dash symbol where unreported), we report the maximum size N considered, the number and type of instances, and the reported ability of the method to find feasible and optimal (or near-optimal) solutions to the TSP, where optimality gaps have been reported. Many of these important details are unfortunately missing from these milestone papers. We have also augmented the table with additional studies that we consider to be important, that were not included in the published systematic literature review. An annotated bibliography of these papers is provided in appendix C.

The published literature over these two decades traces an interesting sequence of enquiry as studies progressed from theoretical with supporting computational simulations to small-scale hardware implementations as commercial quantum computers became publicly available. In summary, these research efforts show progress has been made towards the following key questions:

How best to formulate the TSP for QUBO-based quantum approaches? It is essential that the minimum energy states of the Hamiltonian correspond to feasible and optimal TSP tours. Early studies focused on $N = 4$ cities enabling tricks to be employed such that the edge-based formulation led to Hamiltonians with this property verified through simulations. Once larger instances were tackled though, it became clear that the node-based formulation is required.

How to find the minimum energy state of the Hamiltonian? The landscape of the TSP node-based QUBO Hamiltonian is complex, with parameters that need to be optimised to balance cost and constraint penalties. By the time the formulation question had been settled, D-WAVE devices for QA and IBM hardware for QAOA were available, and small sized TSPs ($N \leq 8$) could be attempted. Initial results proved that finding feasible and near-optimal solutions was challenging due to the difficulties of tuning parameters, although feasible and optimal solution could be verified as having minimum energy states of the Hamiltonian when evaluated (Srinivasan *et al* 2018). Without modification of the formulation to embed more TSP-specific knowledge (Ruan *et al* 2020, Villar-Rodriguez *et al* 2022), or hybridisation of the algorithm with heuristics to ensure feasibility (Osaba *et al* 2021, Warren 2021), the literature suggests that a dead-end had been reached by about 2020 for the goal of purely QUBO-based quantum methods for solving the TSP.

How do QA and QAOA compare in terms of effectiveness? The results in table 1 suggest that QA can handle slightly larger problem instances with current technology. However, the benchmarking tool QUARK (Finžgar *et al* 2022) shows both approaches are not yet capable of reliably converging to feasible solutions for $N > 4$ (QAOA) or $N > 7$ (QA), with QAOA requiring significantly more run time than QA, with additional parameters to be optimised via classical algorithms.

How can TSP-specific knowledge be integrated to improve performance? Researchers have tried to enhance the competitiveness of the general-purpose Hamiltonian minimisation approach of QA and QAOA by either modifying the distance matrix to penalise unlikely edges (Villar-Rodriguez *et al* 2022), or exploiting knowledge of how to eliminate sub-tours (for small N) to add additional constraints to improve feasibility (Ruan *et al* 2020), improving feasibility and optimality for both QA and QAOA respectively. The ability of classical heuristics such as SA, tabu search of k-opt edge exchange heuristics to quickly restore feasibility of TSP solutions is another way that TSP specific knowledge has been integrated within QA on D-WAVE hybrids with classical computers (Osaba *et al* 2021, Warren 2021). An alternative approach more recently has been to improve QAOA with a ‘warm start’ initialisation based on a classical heuristic solution, enabling the

Table 1. Summary of QUBO-based quantum TSP studies over the last 20 years.

Reference (year)	Formulation	Solution method	Experimental focus	N	Instances	Feasible %	Optimality %
(Moser 2003)	Edge	quantum mechanics	Simulation of spatial energy distribution	6	1 easy	100%	0%
(Martoňák <i>et al</i> 2004)	Edge	QA/ classical	Monte Carlo (MC) simulation with 2-opt heuristic	1002	1 TSPLIB	100%	0.5%
(Chen and Zhang 2006)	Edge ^a (modified)	classical	MC simulation with 3-opt heuristic	512	7 TSPLIB	100%	28%
(Chen <i>et al</i> 2011)	Edge	QA	NMR quantum simulation	4	1 easy	100%	0%
(Suzuki <i>et al</i> 2012)	Edge	QA	—	—	—	—	—
(Heim <i>et al</i> 2017)	Edge ^a (modified)	QA	MC simulation with dynamic constraints	12	100 random	0%	—
(Ruan <i>et al</i> 2020)	Edge	QAOA	Simulation, adding constraints to mixer Hamiltonian	5	1 easy	100%	0%
(Warren 2013a, Warren 2013b)	Node	QA	—	—	—	—	—
(Warren 2018)	Node	QA	Hardware (D-WAVE QAC)	8	1 real world $\times$ 1000 runs	0%	—
(Mehta <i>et al</i> 2019)	Node	QA	Hardware (D-WAVE 2000Q)	5	20 random	0%	—
(Finžgar <i>et al</i> 2022)	Node	QA	Hardware (D-WAVE)	8	1 easy	50%	—
(Finžgar <i>et al</i> 2022)	Node	QAOA	Simulation	5	1 easy	50%	—
(Villar-Rodriguez <i>et al</i> 2022)	Node	QA	Hardware (D-WAVE Advantage 6.1)	7	1 easy $\times$ 2000 runs	16%	16% optimal
(Villar-Rodriguez <i>et al</i> 2022)	Node ^a (modified)	QA	Hardware (D-WAVE Advantage 6.1)	7	1 easy $\times$ 2000 runs	73%	0.8% optimal
(Villar-Rodriguez <i>et al</i> 2022)	Node ^a (modified)	QA	Hardware (D-WAVE Advantage 6.1)	5–14	Solved as smaller N and joined, 10×2000 runs each	100%	10% optimal 0%—4% gap
(Warren 2021)	Node	QA/SA	D-WAVE hybrid	22	4 TSP	100%	16.5% gap
(Osaba <i>et al</i> 2021)	Node	QA/Tabu	D-WAVE hybrid	48	6 TSP asymmetric	100%	27% gap

^aModified formulation to embed TSP-specific knowledge.

quantum algorithm to converge to feasible solutions more rapidly while requiring reduced circuit depth (Chander and Blekos 2025).

How can larger instances be solved with current hardware limitations? With the current limitation to about $N = 8$ cities for QA on D-WAVE, current advances for larger instances have relied thus far on partitioning into smaller sets of cities, with each sub-problem solved by quantum methods, before being stitched together by a classical computer (Villar-Rodriguez *et al* 2022). Of course, there are no guarantees that this merging of sub-tours is optimal.

3.3. Current status of quantum ability to solve TSPs

Intuitively, no general-purpose solver is likely to be competitive with specialised TSP solvers or heuristics that exploit knowledge of the combinatorial structure of a TSP instance. It is therefore not surprising that the results in table 1 show a gradually transition towards enhancing the generic QUBO-based quantum approaches with either TSP-informed modifications to the Hamiltonian, or hybridisation with TSP heuristics.

Returning to the question of the proven *ability* of quantum approaches to solve the TSP, the current state of the art (Villar-Rodriguez *et al* 2022) appears to be limited to 7 cities with current D-WAVE hardware using QA and a modified formulation that adapts the distance matrix, with TSP instances up to 14 cities solvable (feasible and near-optimal) by breaking the cities into two groups, and merging the two sub-tours. With painstaking parameter tuning, a setting can be found for the modified node-based formulation that consistently returns feasible solutions with about 10% being optimal, and all other solutions having an optimality gap under 4%. This is for a clever approach that first modifies the distance matrix to penalise very long distances between cities even further, using domain knowledge that optimal tours are unlikely to include edges that are above the median (Villar-Rodriguez *et al* 2022). Thus, the most promising purely quantum approach so far has come from moving away from a general-purpose QUBO solver, showing significant improvements when TSP knowledge is integrated via a modified formulation.

QA running on D-WAVE architectures appears to be more promising amongst the available quantum devices. Benchmarking of two QAOA implementations in IBM hardware in 2021 found that valid solutions for 5-city TSP instances are not achieved, despite significant running times exceeding 100 s (Mesman *et al* 2021). This paper defines a technology readiness level that describes the TSP as TRL-4, ahead of all other optimisation problems, due to having been implemented in full stack simulation and on a practical application, but stated that it is not yet able to be implemented on available quantum hardware.

In recent years the most successful results have been achieved, not from QA or QAOA alone, but from hybridisation with classical approaches, capable of ensuring feasibility of TSP tours, and biasing the search through the use of ‘warm starts’ (Chander and Blekos 2025). The use of the D-WAVE quantum device for hybridisation with classical heuristics such as SA and tabu search has been reported (Warren 2021), with moderate success solving a $N = 22$ city problem from TSPLIB with a 16.5% optimality gap by decomposing the problem into smaller sub-problems. A hybrid approach that uses D-WAVE to solve smaller sub-problems, in tandem with a classical computer to perform initial clustering and also to maintain a tabu list of previous solutions, has also been reported (Osaba *et al* 2021) for asymmetric TSPLIB instances up to $N = 48$ cities with 27% optimality gap.

Advances in quantum technologies will surely see these results improve over time, but it is already apparent that quantum devices acting alone on generic QUBO formulations of the TSP are unlikely to be competitive with state-of-the-art classical TSP solvers, which do not need to struggle with the challenges of such sensitive parameter tuning of an unnecessarily complex landscape.

4. Quantum competitiveness compared to state-of-the-Art TSP solvers

The ability of quantum algorithms—QUBO-based or other—to reliably solve the TSP has not yet been demonstrated, but if the current challenges were to be overcome so that feasible and near-optimal solutions could reliably be attained for small-sized TSP instances, how should their future competitiveness against conventional computers be demonstrated? Comparisons against brute-force enumeration are not the standard, and it is important to demonstrate advantage against state-of-the-art algorithms running on classical computers, which are dramatically more efficient than naïve enumeration.

4.1. What is state-of-the-art?

Over 40 years ago, methods were proposed for finding provably optimal solutions to the TSP with several hundred cities (Grötschel 1980, Padberg and Hong 1980), using commercial linear programming codes running on a conventional computer of the day, that had less computational power than most of today’s

smart watches. These approaches used powerful cutting plane approaches that exploited knowledge of the polytope structure of the TSP to reduce the search space. Currently, the largest TSP that has been proven optimal has 109 399 cities. The well-known LKH algorithm (Helsgaun 2000), which aims to rapidly find near-optimal solutions with bounded optimality gaps, has provided a solution to a 2079 471-city TSP that is provably within at most 0.00074% of optimality. LKH is based on a sophisticated stochastic local search heuristic that starts with a valid tour and iteratively performs edge swaps within different types of neighbourhoods to reduce tour length towards the global minimum. Under many circumstances, EAX, another inexact TSP solver based on an evolutionary algorithm, has been demonstrated to perform even better (Nagata and Kobayashi 2013) and further performance improvements can be achieved using restarting mechanisms (Dubois-Lacoste *et al* 2015) and automated algorithm selection techniques (Kerschke *et al* 2018).

It should be noted that the proof of optimality or the proof of a quality guarantee, however, are not provided by heuristic methods such as LKH or EAX. Such proofs are the result of a combination of these heuristics with methods that calculate lower bounds for the optimum value. The computational experience in the last 40 years has shown that the cutting plane techniques based on geometric investigations of the travelling salesman polytope combined with branch & bound and special linear programming implementations deliver in practice what is needed to solve these TSP instances of spectacular size. CONCORDE (Applegate *et al* 2012) is an example of a code that combines this primal-dual machinery very efficiently.

With the growing ability of such methods to solve large-scale TSP instances, the benchmarks have also been growing. The World TSP instance is a 1904 711-city instance derived from a database of all populated places in Earth, and the best currently known tour length has been found in 2021 by the LKH algorithm. In terms of computational speed-ups, heuristics are available that provide near-optimal (within 1% of best known) solutions in under 1 h of computation, using a divide-and-conquer strategy that clusters the cities into smaller TSP instances which are then solved using LKH before being joined with greedy heuristics (Drori *et al* 2020). This approach has also been used to tackle the Galaxy TSP instance of 1.69 billion stars, taking about 100 d of computation to solve the smaller-scale TSP instances in parallel, and then around 100 more hours to merge the tours. In a push for more practical solutions to large-scale TSP instances, the Santa Clause Challenge was launched in 2020 (Mariescu-Istodor and Fränti 2021), where competitors were asked to find solutions to a TSP through 1.4 million households in Finland within 1 h of computation time. LKH was the most effective strategy, and with parallel processing would have found a similar solution in around 2 min, instead of the 1 hour competition limit.

Using advanced statistical techniques, the empirical scaling of the running time of state-of-the-art conventional TSP solvers has been characterised. For a widely studied class of TSP instances, the median running time required by CONCORDE for finding optimal solutions (in CPU seconds) has been demonstrated to scale according to a function of the form $a \cdot b^{\sqrt{N}}$, with $a = 0.21$ and $b = 1.24194$ (Hoos and Stützle 2014), and even better empirical scaling results have been obtained for improved versions of LKH and EAX (Mu *et al* 2018). These results further illustrate the relative ease with which state-of-the-art conventional TSP solvers find optimal solutions to instances with thousands of cities.

The TSP webpage www.math.uwaterloo.ca/tsp/ provides very broad and high-quality information about the TSP, its history, applications, and the largest problems solved. We refer the interested reader there for a quick and good introduction to the many aspects of the TSP and state-of-the-art methods. Moving beyond classical algorithms running on conventional computers, the state-of-the-art also includes significant speed-ups reported from other (non-quantum) hardware implementations, such as the FPGA-based 2-opt TSP heuristic that achieves around 600% speed-up on CONCORDE solutions up to several hundreds of cities (Mavroidis *et al* 2007). Meanwhile a different direction has been pursued that promises ‘AI chips’ based on SA of the node-based TSP Hamiltonian. Someya *et al* (2016) simulated an $N = 13$ city TSP and found feasible solutions in each of 10 runs. No optimality gaps were reported, but tour length appears to be around 30% above optimal based on the plotted graphs. By 2020, the idea had been implemented (Iimura *et al* 2020) in a CMOS hardware chip to solve a 21-city TSP in 128 ms, and while the energy was reported to have reduced by 19% from its initial state, how this translates into an optimality gap was not reported. Interestingly, despite this paper forming the basis for the press release about ‘quantum chips’ (Bush 2020), the term ‘quantum’ is not used once in the paper, which is clearly using statistical mechanics rather than quantum mechanics, providing another example of media hype surrounding quantum computing.

When considering the question of likely quantum advantage for TSP in the future, it is therefore important that any purely quantum or quantum-classical hybrids should also be compared to the state-of-the-art in classical solvers and non-quantum hardware devices to accurately assess their competitiveness.

4.2. Assessing competitiveness

Clearly current quantum technology is not competitive with state-of-the-art TSP solvers in terms of scale, but even for very small-sized TSP instances the competitiveness has not been demonstrated. The review of existing literature on quantum algorithms for solving the TSP, shown in table 1, has confirmed an unhelpful tendency to omit comparisons and relevant metrics in terms of reported optimality gaps to known optimal solutions, or against heuristics adopting the same formulations. Similarly, there is often a failure to report the robustness of the method to find feasible solutions in terms of repeatability (where the algorithm is stochastic), sensitivity to parameter tuning, as well as the robustness of the approach to instance difficulty.

Many studies have focused on reporting results for solving one very easy instance with 4 or 6 cities, usually as vertices of a convex polygon (Gonzalez-Bermejo *et al* 2022), which has an energy landscape very different from real-world TSP instances as well as from widely studied families of randomly generated TSP instances, as found, e.g., in TSPLIB. Assessing the strengths and weaknesses of any optimisation method should be done in an objective manner that seeks to explore how performance depends on instance characteristics (see for example the methodology known as Instance Space Analysis (Smith-Miles and Muñoz 2023)). To do this effectively it is important to challenge the algorithm with multiple test instances with a variety of characteristics known to create difficulty for TSP algorithms (Smith-Miles and van Hemert 2011, Smith-Miles and Lopes 2012). This is not a practice seen to date in the testing of quantum methods for TSP, although Instance Space Analysis has recently been used to evaluate the performance of QAOA on the MAX-CUT problem, showing how performance depends on circuit design considerations and the interplay with instance characteristics (Katial *et al* 2024).

Acknowledging that the current practice in computational evaluation of quantum methods for optimisation must be improved, some useful benchmarking tools are emerging to facilitate more useful comparisons (McGeoch 2019, Mesman *et al* 2021, Finžgar *et al* 2022), and guidelines for metrics to support benchmarking studies have been proposed (Abbas *et al* 2023). QUARK (Finžgar *et al* 2022) is a tool that enables comparison of QAOA and QA approaches against SA and greedy heuristics operating also on the QUBO formulation of TSP. One of the TSPLIB instances (dsj1000) is included in the toolkit, but reduced to a $N = 14$ city instance by removing nodes in a manner that is not described. The quantum approaches are found to be uncompetitive compared to the greedy and SA methods, even acknowledging that SA is not ideally suited to a QUBO formulation. This is a fair comparison of capability on one instance though, comparing QA to its classical counterpart using the same formulation. The failure of the quantum approaches to find valid solutions as problem size increases to around $N = 5$ for QAOA and $N = 8$ for QA, even after extensive hyperparameter search, is explored and represented with useful graphical tools.

Other papers have compared quantum implementations of optimisation problems against classical solvers like CPLEX (McGeoch and Wang 2013), but notably not on the TSP. It is a significant improvement to benchmarking practice however to see these papers reporting optimality gaps and percentage feasible solutions, rather than the common practice of reporting minimum energy found, which provides no indication of whether that solution is a feasible and near-optimal solution to the TSP. Indeed it has been confirmed that there is often no alignment between the positions of the (local) minima in the energy landscape and the (local) maxima of the landscape of the TSP success probability (Willsch *et al* 2020), unless the energy landscape is biased through placement of the cities to have favourable structure such as the weak-strong cluster landscapes with tall and narrow energy barriers separating local minima (Mandra *et al* 2016).

To support fairer comparisons against state-of-the-art methods running on conventional computers, it is critical that quantum studies ensure:

- (i) fairness of parameter tuning—i.e., same level of sophistication and computational resources for optimising the values of parameters that impact performance, ideally with the sensitivity of performance to parameter settings (robustness) also reported;
- (ii) all methods are given equal computational power (wall clock time or TSP tour length evaluations);
- (iii) all methods report statistical analysis (at the very least standard metrics of location and dispersion, such as means and standard deviations) over multiple independent runs of any randomised algorithms (within this same computational budget) of the optimality gap and percentage feasible TSP tours found;
- (iv) performance is evaluated on challenging instances with more complex and realistic landscapes, that are more informative than easy (polygon or random) instances.

Furthermore, it is worth noting that classical methods tend not to adopt the QUBO (node-based formulation) of the TSP, since it is inefficient. To understand the effectiveness of quantum search, it would be ideal to compare quantum versus classical on the same QUBO formulation. If quantum methods cannot compete with classical methods using the QUBO formulation, then it is unlikely that they will be competitive

against state-of-the-art solvers operating on their preferred formulations. It was this same point that led to an argument to abandon the TSP as a neural network benchmark in the 1990s (Smith 1996), following the realisation that Hopfield neural networks were unlikely to ever be competitive against state-of-the-art methods for solving the TSP, for a variety of reasons relating to the challenges created by the QUBO formulation (see appendix B).

5. Challenges to quantum advantage for TSP

This article has described current progress in solving the TSP using QUBO-based quantum methods, and showed that there is currently little evidence for the optimism that is often communicated in research papers (Osaba *et al* 2022). While other optimisation problems may show more promising results than TSP, we reiterate the concern expressed in the Introduction that the TSP is commonly used, especially in the popular press, as a prime example of the power of quantum computing versus conventional computers. We believe this is misleading given current evidence, and there are numerous reasons that the potential is unlikely to be realised for TSP in the foreseeable future. The TSP is seemingly not well suited to demonstrate the competitiveness of purely quantum solutions to optimisation problems for a variety of reasons as discussed below.

5.1. Lack of combinatorial knowledge

The current approach adopted by QA and QAOA treats the challenge of solving all combinatorial optimisations in the same manner: QUBO mapping to a Hamiltonian for energy minimisation. It does not make any special use of combinatorial knowledge of the TSP to modify how it searches the space of possible feasible solutions, unlike state-of-the-art methods. Such quantum approaches are aiming to be general-purpose solvers for any combinatorial optimisation problem that can be modelled as a QUBO, and the price they pay for generality is a lack of competitiveness with methods that exploit combinatorial knowledge. It is certainly conceivable that quantum approaches could be developed to exploit structures for certain types of instances, rather than expecting a general purpose approach to offer advantage for all instances of a problem. While this is a direction worth further investigation, it is not yet obvious how such structures could be exploited within a quantum framework (Abbas *et al* 2023). Perhaps more promising are efforts to hybridise Grover-style quantum search within classical heuristics or solvers where possible (Ambainis *et al* 2019, Montanaro 2020). However, this is not an outcome that supports the competitive quest to demonstrate that quantum computers will eventually solve the TSP faster and/or better than conventional computers, except in the case where qubits are so cheap and plentiful that a quantum simulation of the classical algorithm is considered acceptable.

5.2. Ineffective formulation

The competitiveness of quantum approaches for solving the TSP has been limited by the restriction to the QUBO formulation. Recently there have been more compact QUBO formulations proposed (Gonzalez-Bermejo *et al* 2022), reducing the number of variables and hence qubits required, but the challenges of a highly complex landscape and sensitivity to parameters remain, and the modified formulations have so far only been tested on small and very easy TSP instances. Likewise, investigations into the impact of different mixer Hamiltonians in QAOA have begun, exploring robustness to expected noise in real quantum circuits (Qian *et al* 2023). Acknowledging that the landscape of the QUBO formulation creates a tendency for QA to converge to a local minima of the energy landscape, modifications have been proposed to allow noisy QA to escape local minima (Dong and Huang 2020, Rivera-Dean *et al* 2021), similar to the hill-climbing modifications proposed to address the same issue in Hopfield neural networks (Smith *et al* 1998).

5.3. Sensitivity to parameters

There are only a few computational studies thus far that have rigorously explored the robustness of parameter settings to enable quantum algorithms to consistently find feasible solutions to TSP instances (Salehi *et al* 2022, Villar-Rodriguez *et al* 2022). These studies have shown that there are some ranges of parameters that can generate feasible (not necessarily near-optimal) tours for a small number of cities, but this ability to find suitable parameters from extensive grid search of parameters rapidly vanishes as the number of cities increases above $N = 5$. For QAOA the sensitivity to parameters is even more critical (Willsch *et al* 2020), since the challenge of finding optimal parameters within each layer p using classical optimisation methods has proven to be very difficult, with barren plateaus often encountered. Experiments have shown no single classical optimiser can consistently find the optimal parameter settings (Fernández-Pendás *et al* 2022,

Blekos *et al* 2023) that are required theoretically for QAOA to converge to feasible and optimal TSP solutions as the circuit depth is increased, making the approach currently impractical (Herrman *et al* 2021).

5.4. Instance properties

Since most research to date has focused on very small TSP instances, it is natural that all such instances tend to be trivially easy to solve by visual inspection, since they lack the kinds of structures that can only be created with a larger number of cities. Many studies have only focused on very small and synthetic instances, often selecting cities from vertices of polygons (Chen *et al* 2011), which tend to create more convex landscapes. Other studies (Finžgar *et al* 2022) have taken more realistic TSP instances from TSPLIB and removed cities (not stating how this selection was done), to attempt a more realistic small sized instance. A few studies (Akshay *et al* 2020, 2021, Herrman *et al* 2021) have explored the impact of instance structure however, and this is a critically important direction. While QAOA seems largely ineffective in general, there have been some interesting results showing the characteristics of instances where QAOA is guaranteed to find the globally optimal solution in a single layer circuit, while QA is deceived into convergence to poor quality local minima (Mandra *et al* 2016, Mandra and Katzgraber 2018, Streif and Leib 2019). Specifically, energy landscapes that create hard instances for QA tend to feature a large number of minima separated by wide barriers, with a single global minimum with a significantly lower energy than all other local minima which are considered not near-optimal. For QA, such instances cause challenges to find the global minimum, whereas the interference-based approach of QAOA provides all states with a destructive interference, except the target state which alone benefits from a constructive amplitude boost from interference (Streif and Leib 2019).

While there has not been much research thus far on how properties of the TSP instance impact solvability of quantum approaches (Smith-Miles 2025), there has been more investigation on related optimisation problems such as MAX-CUT, 2-SAT and 3-SAT, where QAOA performance has been shown to be impacted by instance characteristics, such as density of constraints (Akshay *et al* 2020, 2021, Katial *et al* 2024). Compared to these other optimisation problems, the TSP has a very different landscape structure (Reidys and Stadler 2002), with far less neutrality and more ruggedness, meaning that small moves create big differences (Klemm *et al* 2012). This insight may hold the key to understanding the lack of suitability of QUBO-based quantum methods for TSP, and why unconstrained problems like MAX-CUT, which naturally have a less complex energy landscape, may be more suitable than TSP.

Related to this insight into how the instance properties may affect quantum algorithm performance, there is an emerging direction in applying machine learning methods and ideas from the fields of meta-learning (Wilson *et al* 2021) and algorithm selection (Deshpande and Melnikov 2022). Only recently has research in quantum methods for MAX-CUT started to look beyond the commonly studied 3-regular graphs to explore how different graph structures affect quantum algorithm performance (Moussa *et al* 2020, Herrman *et al* 2021) or how optimal QAOA parameters depend on instance properties (Boulebnane and Montanaro 2021, Moussa *et al* 2022). Similar calls for different instances with a variety of structures and complexities have recently been made for other problems like the maximum independent set (MIS) problem (Andrist *et al* 2023), since it is acknowledged that much of the quantum literature for optimisation has been focused on solving very easy instances. Much more could be done to study TSP instance properties such as skewness of distances (Qian *et al* 2023), and understand their impact on quantum performance, as has been done for classical heuristics (Smith-Miles and van Hemert 2011). The foundational work applying Instance Space Analysis (Smith-Miles and Muñoz 2023) to evaluate how test instance properties impacts quantum algorithm performance has already been done on MAX-CUT (Katial *et al* 2024), and can be readily extended to other combinatorial optimisation problems such as the TSP (Smith-Miles 2025).

5.5. Noise in hardware implementation

The potential of VQAs such as VQE and QAOA to solve optimisation problems is notably compromised when implemented on actual quantum devices due to physical noise processes. Recent research (Resch and Karpuzcu 2021, Wang *et al* 2021, Scriva *et al* 2024) underscores the detrimental effects of the two primary quantum noise sources, namely measurement shot noise and hardware noise (decoherence), on these quantum heuristics. Decoherence not only impacts the current scalability of implementations of quantum algorithms, but significant challenges their ability to outperform classical optimisers. Notably, both theoretical analyses and realistic simulations reveal that VQE and QAOA, despite their promise for the noisy-intermediate scale quantum devices, suffer from severe limitations. For example, the scaling of QAOA, as demonstrated in recent studies (Stilck França and Garcia-Patron 2021, De Palma *et al* 2023, Scriva *et al* 2024), indicates problematically long run times for large problem sizes unless sophisticated parameter initialisation and hybrid optimisation strategies are employed. Furthermore, measurement noise additionally critically hinders the training processes of these algorithms, rendering them practically untrainable in certain contexts due to noise-induced barren plateaus (Xue *et al* 2021). Although there are heuristic strategies

emerging to mitigate some of these issues (Zhou *et al* 2020), the persistent impact of real-world noise continues to seriously impair their practical application. Despite increasingly sophisticated noise models and experimental assessments showing that classical optimisers like Adam and AMSGrad perform better under shot noise conditions (Pellow-Jarman *et al* 2024), the overall advantage of quantum over classical approaches remains an open question. Therefore, while the field of quantum optimisation continues to evolve rapidly, the challenges of measurement shot noise and decoherence suggest that these quantum heuristics, in their current form, may be outperformed by classical methods for instances of practical interest.

6. Conclusions and future outlook for quantum optimisation

The TSP has been cited in the popular press as an example of a hard problem that quantum computers may be able to solve faster, and at larger scale, than conventional computers. However, by focusing on the exponential size of the search space and the corresponding effort for complete enumeration of all possible tours, an impression can be created that misrepresents the state-of-the-art of classical optimisation approaches to solve the TSP on conventional computers. Acknowledging this biased comparison, coupled with the lack of evidence that the popular quantum approaches of QA and QAOA can reliably find feasible solutions to even small-scale TSP instances, means that the TSP should not continue to be used as a benchmark problem that can in good faith be cited for the potential of quantum advantage.

In terms of the question of whether there is evidence that purely quantum approaches (as opposed to hybrid quantum-classical approaches) can solve even small-sized TSPs, this article has reviewed two decades of literature to conclude there is little cause for optimism. We have discussed the major reason that QA and QAOA have struggled to solve even small sized TSPs: their requirement to formulate the TSP in an inefficient and unnatural QUBO form, and the challenging landscape this creates for quantum dynamics to search. Drawing parallels to the near identical debate in the 1980s and 1990s for neural network capability for solving the TSP, we have argued here, just as this article's first author argued then (Smith 1996), that the TSP is not well suited to solution by any method from a QUBO framework. We do not merely highlight this history to point out that research developments in QUBO-based quantum methods would have been faster, with less re-discovery required, if there had been more familiarity with the older literature reviewed in appendix B. We also share this history since it provides a longer timeline of related evidence to support our less optimistic interpretation of the potential for quantum advantage of QUBO-based TSPs approaches. The difficulty of tuning parameters of complex TSP landscape arising from inefficient QUBO formulations was a challenge never overcome in the neural network community, and is the reason we do not currently have the neural chips that were originally promised for solving large scale TSPs via field-programmable gate arrays (Smith 1999). While this history seems to have been forgotten, is highly relevant and important to acknowledge when assessing likelihood of quantum advantage.

Despite the elegance of the theory, in practice, QUBO-based quantum approaches such as QA and QAOA suffer from too many issues relating to the inefficient formulation, the extreme sensitivity to parameters that must be optimised in the formulation, algorithms and hardware, as well as the noise of real circuits requiring the additional overhead and complexity of error correction. Clearly many challenges need to be overcome before QUBO-based quantum optimisers—in simulation or hardware—will be able to solve TSP instances large enough to be of practical interest, and they are likely to remain uncompetitive with state-of-the-art conventional algorithms for the foreseeable future.

It is of course possible that other formulations, tailored with more domain knowledge of the TSP's combinatorial structure, could increase the effectiveness of the QA or QAOA approaches (Gonzalez-Bermejo *et al* 2022, Villar-Rodriguez *et al* 2022) and eventually free them from their QUBO limitations; or that some of the recent variants of QA that have been proposed could offer some enhanced ability to search such complex landscapes more reliably (Crosson and Lidar 2021). More promising though is to recognise that the TSP is never tackled well by general purpose solvers compared to specialised methods that exploit its combinatorial structure.

Of course, the literature is also filled with reports of QA and QAOA implementations on other combinatorial optimisation problems, and perhaps some of these may hold more promise, depending on the complexity of their constraints and how well a QUBO formulation offers an efficient encoding of the search space. For unconstrained problems like MAX-CUT, or those with constraints that can be naturally embedded within a QUBO formulation such as the MIS problem (Brady and Hadfield 2024, Wu *et al* 2024) there is certainly more potential for quantum advantage. It is noteworthy that many studies address multiple optimisation problems, but stay well clear of the TSP. Given that the capability of QA and QAOA to solve even small-sized TSP instances has not yet been robustly demonstrated, it is somewhat surprising to note the growing literature proposing to use these methods to solve even more complex variations of the TSP, such as

vehicle routing problems and the TSP with time windows (Salehi *et al* 2022), albeit in classical-quantum hybrids requiring only very small TSP sub-problems to be solved on quantum devices.

The future of QUBO-based quantum approaches for optimisation in general certainly seems to be continuing the exploration of how problem- and domain-specific insights can be embedded into the method, moving away from expectations of a competitive general-purpose solver. QUBO-based quantum optimisers will also undoubtedly improve in performance and scale via increased hybridisation with classical computers: to decompose into smaller sub-problems, initialise with warm start solutions, and restore feasibility when the quantum method is struggling to balance competing terms in the Hamiltonian energy function. Even for problems that are better suited to a QUBO formulation, such as MAX-CUT and MIS (Wu *et al* 2024), there is, however, a growing consensus that a quantum advantage is unlikely to be achievable for QA and QAOA until several hundred quantum computers with several hundred qubits, each, are available (Farhi *et al* 2012, Guerreschi and Matsuura 2019) with suitable error correction (Stilck França and Garcia-Patron 2021).

More promising, however, for future quantum advantage for TSP and other combinatorial optimisation problems in general, seem to be recent developments in specialised classical-quantum hybrids, as briefly reviewed at the start of section 2. Such a hybrid approach involves best-in-class classical optimisation approaches enhanced by replacing critical sub-routines with functionally equivalent quantum sub-routines. This opens the door to a completely new *data-driven* approach to the assessment of, and design of, quantum algorithms called *hybrid benchmarking*. Here classical solvers are executed on accepted benchmark libraries and the wall-clock runtime recorded. Also, all the instance-specific arguments to critical sub-routines are recorded. These critical sub-routines are then replaced with quantum counterparts. By deriving instance-specific upper and lower bounds to the number of quantum gates required by the quantum sub-routine, one can estimate the quantum resources required to outperform the equivalent classical sub-routine on real benchmark instances. Critically, this assessment can be carried out on large problem instances without requiring access to quantum hardware. This methodology, first introduced by Cade *et al* (2023), has now been applied to a handful of situations, with sometimes sobering results. For example, a recent application to the simplex algorithm (Ammann *et al* 2023) indicates that quantum sub-routines are never likely to provide competitive speedups for benchmark instances unless quantum gate times are infeasibly small. However, there is also hope: a recent study of the 0–1 knapsack problem (Wilkening *et al* 2023) indicates that there are indeed families of hard instances, involving as few as 600 variables, where Grover-type sub-routines may offer useful quantum-accelerated solutions in the intermediate term.

Of course, the focus of quantum advantage for optimisation has thus far been on the speed-ups that quantum computers could offer over conventional computers, but it should be noted that there are other advantages of quantum computing that could be considered attractive in the future, such as potential reduction in energy requirements (Jaschke and Montangero 2023), even if the computational time required to find acceptable solutions is not significantly reduced.

In this article we make no comment on the suitability of quantum approaches to solve many other problems where quantum advantage is more reasonably expected from Grover search, such as machine learning (Huang *et al* 2022), quantum sensing (Pirandola *et al* 2018), or factorisation of large composite numbers (Jiang *et al* 2018). Our focus has been to draw attention to the fact that after two decades of research trying to solve small-sized TSP instances using QA and QAOA, we can understand well why success has not been forthcoming, and this article has highlighted how lessons learned from studying the history of related ideas in neural networks perhaps may have reached this same conclusion much earlier.

If the goal is to solve large TSP instances fast, then the best strategy may be to abandon efforts based on QUBO-based quantum optimisers, and focus efforts on acceleration of TSP heuristics or exact methods through classical-quantum hybridisation. If the goal is to continue to learn more about the capabilities of quantum computers, then the TSP has been an important application to drive fundamental research over the last two decades. While acknowledging the value of such motivating problems, this paper has argued that it is now time to abandon the quest for QUBO-based quantum advantage for the TSP, and turn attention and resources to more promising quantum methods and applications.

Data availability statement

No new data were created or analysed in this study.

Acknowledgment

The authors are grateful to the organisers of the Dagstuhl seminar ‘Challenges in Benchmarking Optimization Heuristics’ (23251) held in June 2023 for providing a stimulating environment for discussions leading to this paper. Further helpful discussions have helped shape the paper, and the authors would like to

thank in particular the two anonymous referees, and the following colleagues: Martin Grötschel, Tony Guttman, Peter Taylor, Ron Sandland, Carleton Coffrin, Michelle Simmons, Charles Hill, Lennart Binkowski, Sandor Fekete, Andreea Lefterovici, Rene Schwonnek, Soeren Wilkening, and Timo Ziegler. Funding for this research has been supported by the Australian Research Council Grant IC200100009, ARC Training Centre in Optimisation Technologies, Integrated Methodologies and Applications (OPTIMA), German Federal Ministry of Education and Research (BMBF), project QuBRA, German Federal Ministry for Economic Affairs and Climate Action (BMWK), project ProvideQ, and The Humboldt Foundation, Alexander von Humboldt Professorship in AI.

Appendix A. Limitations of the Edge-based formulation of the TSP as a QUBO

While it may appear at first glance that the edge-based TSP formulation presented in figure 1 could be recast as a QUBO by expressing the constraints as penalty terms, the choice of encoding of the decision variable (indicating only which cities are connected, and not indicating the exact sequence as per the node-based formulation) necessitates the addition sub-tour elimination constraints which are not readily converted into a penalty term in QUBO form, except for the special case of $N = 4$ cities. To illustrate the challenges, consider the TSP instance of 4 cities (A, B, C and D) with distances between them given by the distance matrix:

$$D = \begin{pmatrix} 0 & 3 & 4 & 5 \\ 3 & 0 & 5 & 4 \\ 4 & 5 & 0 & 3 \\ 5 & 4 & 3 & 0 \end{pmatrix}.$$

Figure A1 provides some solutions to this 4-city TSP including the single optimal solution with minimal tour length of 14; two sub-optimal tours with greater tour lengths; an invalid solution due to presence of sub-tours; and an invalid solution due to a city being omitted from the tour.

It is clear from the edge-based encoding of the invalid solutions that the constraints ensuring one ‘1’ per row and column are insufficient to avoid sub-tours, and additional constraints are required. In the edge-based TSP formulation, these sub-tour elimination constraints are enforced by:

$$\sum_{i \in S_1} \sum_{j \in S_2} X_{i,j}^e \geq 1 \quad \forall S_1 \subset \mathcal{N}, S_2 = \mathcal{N} \setminus S_1 \quad (\text{A1})$$

but this constraint is complex and not linear or quadratic in the decision variables as required for a QUBO penalty term. For $N = 4$ cities however, a special case exists since there is only one type of sub-tour to avoid: two 2-city loops, as shown in the invalid (sub-tours) solution in figure A1. To explicitly prevent such 2-city loops, a set of N additional constraints can be introduced, replacing the more general equation (A1), by exploiting the fact that for 4-cities it is sufficient to ensure that each city must be connected to exactly two other cities (not including itself):

$$\sum_{j \neq i} X_{i,j}^e = 2 \quad \forall i. \quad (\text{A2})$$

It is clear that 3-city loop sub-tours of the kind possible for $N > 4$ are still permitted with these constraints, and so it is not an approach that scales. However, this set of $N = 4$ linear constraints can be recast in QUBO form as:

$$\sum_{i=1}^N \left(\sum_{j \neq i} X_{i,j}^e - 2 \right)^2 \quad (\text{A3})$$

enabling the problem Hamiltonian to be constructed for $N = 4$ cities. The earliest studies ((Moser 2003, Martoňák *et al* 2004) not only adopted equation (A2), but employed a further trick, again uniquely possible for $N = 4$, to utilize these constraints to avoid the penalty term completely. Note that for 4 cities, the symmetrical decision matrix can be specified by the 6 upper triangle elements only:

$(X_{1,2}^e, X_{1,3}^e, X_{1,4}^e, X_{2,3}^e, X_{2,4}^e, X_{3,4}^e)$. We only need to treat two of these as independent variables to express the remaining four as dependent variables due to the presence of 4 constraints in equation (A2). For example, if we select $X_{1,2}$ and $X_{1,3}$ as the independent variables we can set the value of $X_{1,4}^e = (2 - X_{1,2}^e - X_{1,3}^e)$ in the objective function and penalty terms. The same process of solving simultaneous equations can be adopted for the remaining 3 dependent variables, giving $X_{2,3}^e = (2 - X_{1,2}^e - X_{1,3}^e)$, $X_{2,4}^e = X_{1,3}^e$, and $X_{3,4}^e = X_{1,2}^e$.




SOLUTION (TSP TOUR)					
Validity	Valid	Valid	Valid	Invalid (sub-tours)	Invalid (omitted city)
Optimality	Optimal	Sub-optimal	Sub-optimal	-	-
Tour length	14	16	18	-	-
Edge-based encoding	$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$
Node-based encoding	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}$	-	-

Figure A1. Four-city TSP example, showing the encodings of possible tours and their optimality and validity.

The edge-based QUBO for the 4-city TSP can now be expressed in terms of only two independent variables. The cost (tour length) of valid solutions (with zero penalty terms) is given as:

$$20 - 4X_{1,2}^e - 2X_{1,3}^e \quad (\text{A4})$$

which can be verified by checking the tour lengths of the optimal tour ($X_{1,2}^e = X_{1,3}^e = 1$) and sub-optimal tours ($X_{1,2}^e = 1, X_{1,3}^e = 0$) and ($X_{1,2}^e = 0, X_{1,3}^e = 1$) in figure A1.

Appendix B. Lessons from the history of Hopfield-Tank neural networks for solving the TSP

In 1982 John Hopfield (Hopfield 1982) proposed a neural network model with fixed weights and neuronal state updating given by a coupled set of differential equations. Inspired by neurobiology, but readily adaptable to circuit implementation, the network dynamics could recover memories stored in the weights based on initial neuronal states capturing partial memories. This work was recently recognized with the award of the 2024 Nobel Prize in Physics. Two years after this seminal work, Hopfield showed that while the neural network states are updated within the interior of a hypercube, with vertices representing firing or non-firing states, the network converges to stable states which are at the binary vertices and correspond to the global minima of a quadratic Lyapunov function (Hopfield 1984). The following year, he collaborated with David Tank (Hopfield and Tank 1985) to show how the weights could be fixed in a manner that enabled this energy function to correspond to a TSP cost function with constraint violation treated as additional penalty function terms. This Hopfield-Tank neural network model represented the TSP as a node-based QUBO formulation, with neurons representing binary decision variables with $X_{ij} = 1$ if city i is in position j , and 0 otherwise. As the neurons update according to a set of coupled differential equations to simulate neuron activation, the quadratic energy function is provably minimised, resulting in a theoretically minimal energy state which should ideally be a feasible and optimal solution to the TSP. It was a brilliant idea that sparked much excitement about the power of neural network models for solving NP hard problems. Initially the plan was to demonstrate ability to converge to optimal solutions using simulations on conventional computers, with the argument that future hardware implementations on FPGA chips would help enormous speed-ups to be realised. Hopfield and Tank reported computational simulation results showing that a TSP with 10 randomly placed cities could find feasible solutions in around 80% of random starts, and that around 50% of those were a shortest tour. They stated, ‘In our simulations, an appropriate general size of the parameters was easily found, and an anecdotal exploration of parameter values was used to find a good (but not optimized) operating point’. Increasing to $N = 30$ cities however, the choice of parameters was found to impact the results greatly, and the authors acknowledged that they had not found a suitable setting for all

parameters. When feasible tours were found, the optimality gap appeared to be around 50% longer tours. Nevertheless, it was an elegant idea and the impressive results for $N = 10$ was enough for some hype to commence. However, it was short lived.

Three years later in 1988, Wilson and Pawley (Wilson and Pawley 1988) published a paper reporting a failure to reproduce the successful $N = 10$ city results of Hopfield and Tank. They showed that feasible solutions were difficult to encounter, with convergence to local minima corresponding to infeasible solutions more commonly achieved in 85% of 100 random starts. For the 15% of trials that resulted in a valid tour, the tour length was only slightly better than a randomly generated tour on average, and certainly not consistent with the Hopfield and Tank results with the same reported parameter settings. An exploration of alternative parameter settings demonstrated extreme sensitivity to initial conditions and various parameters, with the search for optimal settings to achieve feasible solutions becoming more challenging as problem size increased. Wilson and Pawley provided some insights into the causes for the inability of the Hopfield-Tank method to solve the TSP including i) sensitivity to parameters that cannot easily be optimised, and ii) $2N$ -fold symmetry and degeneracy that causes a competition amongst degenerate tours with sub-optimal sub-tours being merged together. They attempted several modifications to try to enable consistently good TSP tours to be found for $N = 10$, but concluded that ‘Hopfield and Tank were very fortunate in the limited number of TSP simulations they attempted. Even at the value $N = 10$ it transpires that their basic method is unreliable and does not offer much scope for improvement.’ (Wilson and Pawley 1988).

Various modifications have since been made to improve convergence and escape from local minima (Smith *et al* 1998, Smith 1999), but comparisons to more standard heuristics running on conventional computers showed no real advantage when requiring such heuristics to adopt the same QUBO formulation for fair comparison, and acknowledging that the heuristics would perform even more effectively when permitted to use their more native formulations. In 1996, an argument (Smith 1996) was made that the TSP was not a useful benchmark for demonstrating the power of neural networks compared to conventional computers, since the quadratic formulation that neural networks are required to adopt makes them barely competitive with heuristics using the same quadratic formulation, and naturally makes them extremely uncompetitive with other heuristics and exact solvers that have freedom to operate on a linear formulation with subtour elimination constraints. Comparison of new technology should always be made against current state-of-the-art, and for TSP such methods do not use a quadratic formulation.

These same issues have appeared again with quantum approaches to solving TSP, and the argument for abandoning the TSP as a benchmark are similar. The insights offered by Wilson and Pawley (Wilson and Pawley 1988) about the degeneracy and impact of parameters on the landscape are highly relevant to the difficulties seen by quantum methods operating on this same QUBO formulation of the TSP. This seminal work of Hopfield-Tank and Wilson and Pawley unfortunately seems to have been overlooked in the two decades of effort to explore quantum advantage for TSP.

Appendix C. Annotated bibliography of QUBO-based quantum approaches to solve the TSP

The first theoretical proposal before quantum computers existed was by Moser (2003), whereby a TSP solution was encoded as the minimum energy state of a Hamiltonian function defined using the edge-based TSP formulation (see figure 1). Operating on a rectangular grid with equidistantly spaced mesh points, some of which are TSP cities, the edges between cities are represented as Gaussian-shaped narrow potential valleys. Using the finite difference method, the solution to the time-independent 2D Schrödinger equation is provided across the whole grid, and the intensity distribution of the one particle wave function is inspected for some negative total energy E . The energy is expected to be minimised at all grid points comprising the minimal cost tour. The method was illustrated on a numerical example of a 6-city TSP on a 100×80 grid, but no implementation details were provided. Indeed, the paper acknowledged the parameter dependency of the approach in practice, and that feasibility or minimal tours may not be attainable without suitable parameter tuning, with likely dependence on city locations acknowledged.

The following year a QA algorithm was simulated on a conventional computer (Martoňák *et al* 2004) and compared to its thermal relative, SA. The Hamiltonian was constructed for the first time using an edge formulation, but with a clever variable dependency trick that was used to ensure valid tours, unfortunately valid only for $N < 6$ cities. The theory proposed to exploit knowledge of TSP heuristics with 2-opt edge exchanges to create a Hamiltonian with a kinetic energy term that generates 2-opt like quantum fluctuations. The theory was not tested with true simulation of the QA strategy since the actual Schrödinger annealing evolution of the quantum Hamiltonian was deemed intractable. Instead, a quantum Monte Carlo (QMC) search was used to simulate time, and the possibility of infeasible solutions (permitted by minimal states of the Hamiltonian for $N \geq 6$) was avoided by allowing 2-opt exchanges within the QMC search. A simulation

on a single 1002-city instance from the TSPLIB benchmark library showed that the ‘QA’ outperformed SA, but not the classical Lin–Kernighan 2-opt heuristic for the same number of evaluations.

The next milestone paper identified in the review by Osaba *et al* (2022) is the 2006 study by Chen and Zhang (Chen and Zhang 2006), which proposed an alternative formulation of the energy landscape for TSP, exploiting combinatorial knowledge of TSP landscapes by studying distributions of nearest neighbours of local optima. Rather than constructing a Hamiltonian for QA however, it used Metropolis MC simulation with 3-opt exchanges to show that this new energy function performs better than SA operating on the cost (distance) function. A simulation with 532 cities showed that SA converged more quickly to the optimal solution with the new energy function. While this is not really a true QA approach, it was a study that showed the value of incorporating TSP-related landscape knowledge into the energy function, rather than just using distances between cities.

It was not until 2011 that we start to see the appearance of more general-purpose approaches, QA and QAOA, operating on a Hamiltonian/Ising formulation based on the edge formulation, with various attempts to handle the sub-tour elimination constraints. Adopting the same trick to embed the valid tour constraint in variable dependencies (Martoňák *et al* 2004) (valid only for $N < 6$), Chen *et al* (2011) ran a QA simulation using an NMR quantum simulator on a single easy instance with 4 cities. They showed that the algorithm can converge within about 15 ms to the optimal solution with probability 0.94, but the instance selected was a very easy layout (vertices of a polygon) with one obvious solution.

The edge formulation continued to be used until 2013, with Suzuki *et al* (2012) including a flawed formulation in the QA chapter of a book on Ising models. The proposed Hamiltonian does not prevent sub-tours from being generated, but there were no simulations conducted to demonstrate its limitations. In 2017, Heim *et al* (2017) finally put to rest the idea that the edge formulation could form the basis of a valid quantum approach. Using the edge formulation they used the exact solver CPLEX on a conventional computer to solve 100 randomly generated $N = 12$ TSP instances, and showed that 75% of them contained sub-tours making them invalid solutions. They then added additional constraints to CPLEX to avoid sub-tours and analysed the exponential growth of such additional constraints, declaring it impractical for QA to adopt such an approach. Instead, they proposed the idea of only adding necessary constraints once a sub-tour has been created to force a new solution that removes the sub-tour. Testing this idea on CPLEX showed only two iterations were required to find feasible and optimal solutions in 95% of the 100 random $N = 12$ TSP instances tested. However, this strategy failed for Monte Carlo simulation of QA and SA with success probability decreasing exponentially with more iterations. They also considered adding slack variables to eliminate sub-tours, but acknowledged this was not scalable or practical for quantum methods. Based on their experimentation with the edge formulation of TSP they concluded that ‘analog QA devices are unlikely to be of interest as TSP solvers in the near future’ primarily due to the QUBO requirement. They were more optimistic about future digital QA devices however with perhaps more flexibility in how inequality constraints can be handled to accommodate the sub-tour elimination constraints true to the edge-based TSP formulation.

Meanwhile, attempts to solve small city TSP instances using QAOA had taken a different approach. In 2020, Ruan *et al* (2020) used the edge formulation coupled with the idea of adding constraints (for $N = 5$ city TSP instances) into the mixer Hamiltonian. Acknowledging it was not a scalable approach, they used an oracle to check feasibility of possible solutions, to prepare an initial state as a superposition over all feasible solutions. They claimed this to be a more efficient approach than QA, with less qubits required and guarantees of feasibility for $N \leq 6$. One easy instance was simulated and the optimal solution was found. This was based on the idea of Hen and Spedalieri (2016) of encoding constraints in the mixer Hamiltonian, but they never tested it on TSP, only less constrained problems such as graph colouring and 3-SAT. However, it clearly has limited potential for reasonable sized TSP instances.

After a decade of efforts focused on the edge-based TSP formulation, by 2013 the feasibility issue was finally tackled by remembering that the Hopfield–Tank model of the TSP actually used the node-based formulation, and not the edge-based formulation as suggested by Martoňák *et al* (2004). Warren (2013a) proposed the TSP Hamiltonian for QA using the node-based formulation—which was also included around the same time in a list of Ising formulations of various combinatorial optimisation problems by Lucas (2014)—acknowledging the benefits to avoiding exponentially growing constraints. No simulations or experiments were conducted yet to explore if the minimum energy states of the Hamiltonian could be found, and if they would correspond to feasible and optimal solutions of the TSP, but at least it was now theoretically possible. We refer the interested reader to the review by Warren (2017) of the history to that point, prior to hardware implementations.

It would take another five years before the first implementations of QA on hardware would start to be reported. In 2018, Warren (Warren 2018) reported the first unsuccessful attempts to solve TSP instances of sizes $N = 6, 7$ and 8 with real city locations, with 9 cities beyond the available qubits on D-WAVE QAC

hardware. He confirmed difficulties in converging to feasible tours in most of the 1000 trial runs, and the trial-and-error nature of the parameter setting process with unpredictable results.

The next D-WAVE 2000Q hardware implementation for QA was reported (Mehta *et al* 2019) in 2019 for 20 TSP instances ranging from size $N = 2$ to 5. For $N = 5$ it failed to produce optimal, or even feasible, solutions with median success probability of 10^{-5} from 20 random instances, taking over 50 s. The classical LKH heuristic was used to declare the optimal solution, however the performance gap was not presented for the 20 runs and the time comparison was not provided against the heuristic. Oddly, the paper paints an optimistic picture of demonstrating feasibility for solving practical robotic applications of the TSP.

Finally, the debate about which method, QA or QAOA, is more effective at solving small-sized TSP instances was progressed by an experimental study utilised a new benchmarking tool known as QUARK (Finžgar *et al* 2022) in 2022. Comparing hardware implementations of QA (on the D-WAVE Advantage 4.1 and D-WAVE 2000 Q6) against a simulation of QAOA they showed both approaches are not yet capable of reliably converging to feasible solutions for $N > 4$ (QAOA) or $N > 7$ (QA). Percentage feasibility is around 50% for 1000 s of run time (QAOA running on GPU) and 100 s for D-WAVE results, with the authors acknowledging the difficulties in tuning parameters such that a balance is found between feasibility and optimality. They conclude by suggesting that better formulations may be needed for the TSP.

Most comprehensive is a 2022 study by Villar-Rodriguez *et al* (2022) that explored the range of performances obtained from a systematic exploration of the effect of Hamiltonian and algorithm parameters on the energy landscape for TSP, and the resulting impacts on feasibility, optimality in terms of tour length, and repeatability from multiple trials. They focused on two different QUBO models based on the node formulations of the TSP, the standard one and a heuristically modified one that penalises edges that are longer than the median from a given city. This is a clever modification that changes the distance matrix rather than adding new constraints with parameters to tune. A single 7-city TSP known as ‘Burma7’ was studied. While the standard Hamiltonian has 5040 feasible tours for this 7-city problem, the heuristically modified Hamiltonian reduces this to only 28 un-penalised tours, with the remainder being penalised in a manner that alters the landscape to make them less attractive for energy minimisation. This is a nice example of using TSP-specific knowledge that very long edges are not likely (although possible) to be included in an optimal tour, and so biases the Hamiltonian with this knowledge. The best feasible tour is retained within this set of 28 feasible tours, and the question is explored whether a QA algorithm can find it. Using the D-WAVE Advantage 6.1 they conduct over 3700 runs across 220 parameter configurations given an annealing time of 400 ms per run. The standard Hamiltonian formulation was four times more likely to find feasible tours when each formulation’s parameters were selected for best performance, but the heuristically modified Hamiltonian was six times more likely for its minimal energy solution to be both feasible and optimal. From 2000 runs, the probability of finding a feasible solution was 73% and 16.2% respectively for the heuristically modified and standard formulations respectively, each with their own optimised parameter settings, but the probability that these feasible solutions were optimal was 0.1% and 16% respectively. The authors suggest that longer annealing time might improve these probabilities further. They then investigate the generalised performance of the heuristically modified Hamiltonian on larger TSP instances from $N = 5$ to $N = 14$, noting that $N > 9$ cannot be handled by the current D-WAVE hardware, and so they decomposed the larger problems into equal sized subproblems, applied the approach to each portion, and then merged the subtours in a minimal cost manner to ensure feasibility. The optimality gaps were in the range of 0%–4% for $N = 5$ to 14; with 10% of feasible solutions found being optimal. This appears to be the most promising QA approach to date for embedding domain knowledge to modify the TSP formulation (distance matrix) before mapping to the QUBO formulation.

ORCID iDs

Kate A Smith-Miles  <https://orcid.org/0000-0003-2718-7680>

Holger H Hoos  <https://orcid.org/0000-0003-0629-0099>

Hao Wang  <https://orcid.org/0000-0002-4933-5181>

Thomas Bäck  <https://orcid.org/0000-0001-6768-1478>

Tobias J Osborne  <https://orcid.org/0000-0002-0688-2345>

References

- Abbas A, Ambainis A, Augustino B, Bärttschi A, Buhrman H, Coffrin C, Cortiana G, Dunjko V, Egger D J and Elmegreen B G 2023 Quantum optimization: potential, challenges, and the path forward (arXiv: 2312.02279)
- Akshay V, Philathong H, Morales M E and Biamonte J D 2020 Reachability deficits in quantum approximate optimization *Phys. Rev. Lett.* **124** 090504

- Akshay V, Philathong H, Zacharov I and Biamonte J 2021 Reachability deficits in quantum approximate optimization of graph problems *Quantum* **5** 532
- Ambainis A, Balodis K, Iraids J, Kokainis M, Prūsis K and Vihrovs J 2019 Quantum speedups for exponential-time dynamic programming algorithms *Proc. 30th Annual ACM-SIAM Symp. on Discrete Algorithms*
- Ammann S, Hess M, Ramacciotti D, Fekete S P, Goedicke P L, Gross D, Lefterovici A, Osborne T J, Perk M and Rotundo A 2023 Realistic runtime analysis for quantum simplex computation (arXiv: [2311.09995](#))
- Andrist R S, Schuetz M J, Minssen P, Yalovetzky R, Chakrabarti S, Herman D, Kumar N, Salton G, Shaydulin R and Sun Y 2023 Hardness of the maximum independent set problem on unit-disk graphs and prospects for quantum speedups (arXiv: [2307.09442](#))
- Applegate D, Bixby R, Chvátal V and Cook W 2012 *The Traveling Salesman Problem, Concorde TSP Solver* (available at: www.math.uwaterloo.ca/tsp/concorde/index.html)
- Blekos K, Brand D, Ceschini A, Chou C-H, Li R-H, Pandya K and Summer A 2023 A review on quantum approximate optimization algorithm and its variants (arXiv: [2306.09198](#))
- Boulebnane S and Montanaro A 2021 Predicting parameters for the quantum approximate optimization algorithm for MAX-CUT from the infinite-size limit (arXiv: [2110.10685](#))
- Brady L T and Hadfield S 2024 Iterative quantum algorithms for maximum independent set *Phys. Rev. A* **110** 052435
- Bush S 2020 CMOS Quantum Chip Solves ‘Travelling Salesman’ Problem for 22 Cities (available at: www.electronicweekly.com/news/research-news/quantum-chip-solves-travelling-salesman-problem-22-cities-2020-01/)
- Cade C, Folkertsma M, Niesen I and Weggemans J 2023 Quantifying Grover speed-ups beyond asymptotic analysis *Quantum* **7** 1133
- Cerezo M, Arrasmith A, Babbush R, Benjamin S C, Endo S, Fujii K, McClean J R, Mitarai K, Yuan X and Cincio L 2021 Variational quantum algorithms *Nat. Rev. Phys.* **3** 625–44
- Chander S and Blekos K 2025 Comparative analysis of quantum approximate optimization algorithms for solving the travelling salesperson problem (<https://doi.org/10.13140/RG.2.2.25958.54080>)
- Chen H, Kong X, Chong B, Qin G, Zhou X, Peng X and Du J 2011 Experimental demonstration of a quantum annealing algorithm for the traveling salesman problem in a nuclear-magnetic-resonance quantum simulator *Phys. Rev. A* **83** 032314
- Chen Y and Zhang P 2006 Optimized annealing of traveling salesman problem from the nth-nearest-neighbor distribution *Physica A* **371** 627–32
- Cook W J, Applegate D L, Bixby R E and Chvatal V 2011 *The Traveling Salesman Problem: A Computational Study* (Princeton university Press)
- Crosson E and Lidar D 2021 Prospects for quantum enhancement with diabatic quantum annealing *Nat. Rev. Phys.* **3** 466–89
- De Palma G, Marvian M, Rouzé C and França D S 2023 Limitations of variational quantum algorithms: a quantum optimal transport approach *PRX Quantum* **4** 010309
- Deshpande A and Melnikov A 2022 Capturing symmetries of quantum optimization algorithms using graph neural networks *Symmetry* **14** 2593
- Dong Y and Huang Z 2020 An improved noise quantum annealing method for TSP *Int. J. Theor. Phys.* **59** 3737–55
- Dörn S 2007 Quantum algorithms for graph traversals and related problems *Proc. CIE*
- Drori I, Kates B J, Sickinger W R, Kharkar A G, Dietrich B, Shporer A and Udell M 2020 GalaxyTSP: a new billion-node benchmark for TSP *Learning Meets Combinatorial Algorithms at NeurIPS 2020*
- Dubois-Lacoste J, Hoos H H and Stützle T 2015 On the empirical scaling behaviour of state-of-the-art local search algorithms for the Euclidean TSP *Proc. 2015 Annual Conf. on Genetic and Evolutionary Computation*
- Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv: [1411.4028](#))
- Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A and Preda D 2001 A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem *Science* **292** 472–5
- Farhi E, Gosset D, Hen I, Sandvik A, Shor P, Young A and Zamponi F 2012 Performance of the quantum adiabatic algorithm on random instances of two optimization problems on regular hypergraphs *Phys. Rev. A* **86** 052334
- Fernández-Pendás M, Combarro E F, Vallecorsa S, Ranilla J and Rúa I F 2022 A study of the performance of classical minimizers in the quantum approximate optimization algorithm *J. Comput. Appl. Math.* **404** 113388
- Finžgar J R, Ross P, Hölscher L, Klepsch J and Luckow A 2022 QUARK: a framework for quantum computing application benchmarking *2022 IEEE Int. Conf. on Quantum Computing and Engineering (QCE)*
- Gonzalez-Bermejo S, Alonso-Linaje G and Atchade-Adelomou P 2022 GPS: a new TSP formulation for its generalizations type QUBO *Mathematics* **10** 416
- Grötschel M 1980 *On the Symmetric Travelling Salesman Problem: Solution of a 120-City Problem* (Springer)
- Grover L K 1997 Quantum mechanics helps in searching for a needle in a haystack *Phys. Rev. Lett.* **79** 325–8
- Guerraeschi G G and Matsuura A Y 2019 QAOA for Max-Cut requires hundreds of qubits for quantum speed-up *Sci. Rep.* **9** 6903
- Heim B, Brown E W, Wecker D and Troyer M 2017 Designing adiabatic quantum optimization: a case study for the traveling salesman problem (arXiv: [1702.06248](#))
- Helsgaun K 2000 An effective implementation of the Lin–Kernighan traveling salesman heuristic *Eur. J. Oper. Res.* **126** 106–30
- Hen I and Spedalieri F M 2016 Quantum annealing for constrained optimization *Phys. Rev. Appl.* **5** 034007
- Herrman R, Ostrowski J, Humble T S and Siopsis G 2021 Lower bounds on circuit depth of the quantum approximate optimization algorithm *Quantum Inf. Process.* **20** 1–17
- Hoos H H and Stützle T 2014 On the empirical scaling of run-time for finding optimal solutions to the travelling salesman problem *Eur. J. Oper. Res.* **238** 87–94
- Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* **79** 2554–8
- Hopfield J J 1984 Neurons with graded response have collective computational properties like those of two-state neurons *Proc. Natl Acad. Sci.* **81** 3088–92
- Hopfield J J and Tank D W 1985 “Neural” computation of decisions in optimization problems *Biol. Cybern.* **52** 141–52
- Huang H-Y, Broughton M, Cotler J, Chen S, Li J, Mohseni M, Neven H, Babbush R, Kueng R and Preskill J 2022 Quantum advantage in learning from experiments *Science* **376** 1182–6
- Imura R, Kitamura S and Kawahara T 2020 Implementation of multi spin-thread architecture to fully-connected annealing processing AI chips *2020 IEEE 63rd Int. Midwest Symp. on Circuits and Systems (MWSCAS)*
- Jaschke D and Montanero S 2023 Is quantum computing green? An estimate for an energy-efficiency quantum advantage *Quantum Sci. Technol.* **8** 025001
- Jattana M S, Jin F, Gonzalez Calaza C D, Michielsen K, Willsch D, Willsch M and Mehta V 2020 Optimization problems for benchmarking the Hybrid Solver Service version 2 and the Advantage QPU Jülich Supercomputing Center (JSC)

- Jiang S, Britt K A, McCaskey A J, Humble T S and Kais S 2018 Quantum annealing for prime factorization *Sci. Rep.* **8** 17667
- Katyal V, Smith-Miles K, Hill C and Hollenberg L 2024 On the instance dependence of parameter initialization for the quantum approximate optimization algorithm: insights via instance space analysis *INFORMS J. Comput.* **37** 146–71
- Kerschke P, Kotthoff L, Bossek J, Hoos H H and Trautmann H 2018 Leveraging TSP solver complementarity through machine learning *Evol. Comput.* **26** 597–620
- Kirkpatrick S, Gelatt C D Jr and Vecchi M P 1983 Optimization by simulated annealing *Science* **220** 671–80
- Klemm K, Mehta A and Stadler P F 2012 Landscape encodings enhance optimization *PLoS One* **7** e34780
- Koßmann G, Binkowski L, van Luijk L, Ziegler T and Schwonnek R 2022 Deep-circuit QAOA (arXiv: 2210.12406)
- Lucas A 2014 Ising formulations of many NP problems *Front. Phys.* **2** 5
- Mandra S and Katzgraber H G 2018 A deceptive step towards quantum speedup detection *Quantum Sci. Technol.* **3** 04LT01
- Mandra S, Zhu Z, Wang W, Perdomo-Ortiz A and Katzgraber H G 2016 Strengths and weaknesses of weak-strong cluster problems: a detailed overview of state-of-the-art classical heuristics versus quantum approaches *Phys. Rev. A* **94** 022337
- Mariescu-Istodor R and Fränti P 2021 Solving the large-scale TSP problem in 1 h: santa Claus challenge 2020 *Front. Robot. AI* **8** 689908
- Martonák R, Santoro G E and Tosatti E 2004 Quantum annealing of the traveling-salesman problem *Phys. Rev. E* **70** 057701
- Mavroidis I, Papaefstathiou I and Pnevmatikatos D 2007 A fast FPGA-based 2-opt solver for small-scale euclidean traveling salesman problem *15th Annual IEEE Symp. on Field-Programmable Custom Computing Machines (FCCM 2007)*
- McGeoch C C 2019 Principles and guidelines for quantum performance analysis. *Quantum Technology and Optimization Problems: First Int. Workshop, QTOP 2019 Proc. 1, (Munich, Germany, 18 March 2019)*
- McGeoch C C and Wang C 2013 Experimental evaluation of an adiabatic quantum system for combinatorial optimization *Proc. ACM Int. Conf. on Computing Frontiers*
- Mehta A, Muradi M and Woldetsadick S 2019 Quantum annealing based optimization of robotic movement in manufacturing *Quantum Technology and Optimization Problems: First Int. Workshop, QTOP 2019 Proc. 1 (Munich, Germany, 18 March 2019)*
- Mesman K, Al-Ars Z and Möller M 2021 Qpack: quantum approximate optimization algorithms as universal benchmark for quantum computers (arXiv: 2103.17193)
- Mizel A, Lidar D A and Mitchell M 2007 Simple proof of equivalence between adiabatic quantum computation and the circuit model *Phys. Rev. Lett.* **99** 070502
- Montanaro A 2020 Quantum speedup of branch-and-bound algorithms *Phys. Rev. Res.* **2** 013056
- Moser H R 2003 The quantum mechanical solution of the traveling salesman problem *Physica E* **16** 280–5
- Moussa C, Calandra H and Dunjko V 2020 To quantum or not to quantum: towards algorithm selection in near-term quantum optimization *Quantum Sci. Technol.* **5** 044009
- Moussa C, Wang H, Bäck T and Dunjko V 2022 Unsupervised strategies for identifying optimal parameters in quantum approximate optimization algorithm *EPJ Quantum Technol.* **9** 11
- Moylett A E, Linden N and Montanaro A 2017 Quantum speedup of the traveling-salesman problem for bounded-degree graphs *Phys. Rev. A* **95** 032323
- Mu Z, Dubois-Lacoste J, Hoos H H and Stützle T 2018 On the empirical scaling of running time for finding optimal solutions to the TSP *J. Heuristics* **24** 879–98
- Nagata Y and Kobayashi S 2013 A powerful genetic algorithm using edge assembly crossover for the traveling salesman problem *INFORMS J. Comput.* **25** 346–63
- Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information* (Cambridge University Press)
- Osaba E, Villar-Rodríguez E and Oregi I 2022 A systematic literature review of quantum computing for routing problems *IEEE Access* **10** 55805–17
- Osaba E, Villar-Rodríguez E, Oregi I and Moreno-fernandez-de-leceta A 2021 Focusing on the hybrid quantum computing-tabu search algorithm: new results on the asymmetric salesman problem *Proc. Genetic and Evolutionary Computation Conf. Companion*
- Padberg M W and Hong S 1980 *On the Symmetric Travelling Salesman Problem: A Computational Study* (Springer)
- Pan Y, Tong Y and Yang Y 2022 Automatic depth optimization for a quantum approximate optimization algorithm *Phys. Rev. A* **105** 032433
- Pellow-Jarman A, McFarthing S, Sinayskiy I, Park D K, Pillay A and Petruccione F 2024 The effect of classical optimizers and Ansatz depth on QAOA performance in noisy devices *Sci. Rep.* **14** 16011
- Pirandola S, Bardhan B R, Gehring T, Weedbrook C and Lloyd S 2018 Advances in photonic quantum sensing *Nat. Photon.* **12** 724–33
- Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- Qian W, Basili R A, Eshaghian-Wilner M, Khokhar A, Luecke G and Vary J P 2023 Comparative study of variations in quantum approximate optimization algorithms for the traveling salesman problem (arXiv: 2307.07243)
- Reidys C M and Stadler P F 2002 Combinatorial landscapes *SIAM Rev.* **44** 3–54
- Resch S and Karpuzcu U R 2021 Benchmarking quantum computers and the impact of quantum noise *ACM Comput. Surv.* **54** 1–35
- Rivera-Dean J, Huembeli P, Acín A and Bowles J 2021 Avoiding local minima in variational quantum algorithms with neural networks (arXiv: 2104.02955)
- Ruan Y, Marsh S, Xue X, Liu Z and Wang J 2020 The quantum approximate algorithm for solving traveling salesman problem *Comput. Mater. Contin.* **63** 1237–47
- Salehi Ö, Glos A and Miszczak J A 2022 Unconstrained binary models of the travelling salesman problem variants for quantum optimization *Quantum Inf. Process.* **21** 67
- Sathe P, Lubinski T, Coffrin C, Apanavicius J, McGeoch C and Bernal Neira D 2023 Characteristics of optimization applications as quantum performance benchmarks *Bull. Am. Phys. Soc.* **2023** A72–005
- Scriva G, Astrakhansev N, Pilati S and Mazzola G 2024 Challenges of variational quantum optimization with measurement shot noise *Phys. Rev. A* **109** 032408
- Severini S 2022 *Quantum Computing and the Traveling Salesperson Problem: a tale of mystery and imagination (but mostly imagination)* (available at: www.linkedin.com/pulse/quantum-computing-traveling-salesperson-problem-tale-mystery-simone/)
- Smith K A 1999 Neural networks for combinatorial optimization: a review of more than a decade of research *INFORMS J. Comput.* **11** 15–34
- Smith K 1996 An argument for abandoning the travelling salesman problem as a neural-network benchmark *IEEE Trans. Neural Netw.* **7** 1542–4
- Smith K, Palaniswami M and Krishnamoorthy M 1998 Neural techniques for combinatorial optimization with applications *IEEE Trans. Neural Netw.* **9** 1301–18

- Smith-Miles K 2025 Understanding instance hardness for optimisation algorithms: methodologies, open challenges and post-quantum implications *Appl. Math. Modelling* **142** 115965
- Smith-Miles K and Lopes L 2012 Measuring instance difficulty for combinatorial optimization problems *Comput. Oper. Res.* **39** 875–89
- Smith-Miles K and Muñoz M A 2023 Instance space analysis for algorithm testing: methodology and software tools *ACM Comput. Surv.* **55** 1–31
- Smith-Miles K and van Hemert J 2011 Discovering the suitability of optimisation algorithms by learning from evolved instances *Ann. Math. Artif. Intell.* **61** 87–104
- Someya K, Ono R and Kawahara T 2016 Novel Ising model using dimension-control for high-speed solver for Ising machines 2016 14th IEEE Int. New Circuits and Systems Conf. (NEWCAS)
- Srinivasan K, Satyajit S, Behera B K and Panigrahi P K 2018 Efficient quantum algorithm for solving travelling salesman problem: an IBM quantum experience (arXiv: [1805.10928](https://arxiv.org/abs/1805.10928))
- Stilck França D and Garcia-Patron R 2021 Limitations of optimization algorithms on noisy quantum devices *Nat. Phys.* **17** 1221–7
- Streif M and Leib M 2019 Comparison of QAOA with quantum and simulated annealing (arXiv: [1901.01903](https://arxiv.org/abs/1901.01903))
- Sung K J, Yao J, Harrigan M P, Rubin N C, Jiang Z, Lin L, Babbush R and McClean J R 2020 Using models to improve optimizers for variational quantum algorithms *Quantum Sci. Technol.* **5** 044008
- Suzuki S, Inoue J-I and Chakrabarti B K 2012 *Quantum Ising Phases and Transitions in Transverse Ising Models* vol 862 (Springer)
- Villar-Rodriguez E, Osaba E and Oregi I 2022 Analyzing the behaviour of D’WAVE quantum annealer: fine-tuning parameterization and tests with restrictive Hamiltonian formulations 2022 IEEE Symp. Series on Computational Intelligence (SSCI)
- Vizzuso M, Passarelli G, Cantele G and Lucignano P 2023 Convergence of digitized-counterdiabatic QAOA: circuit depth versus free parameters (arXiv: [2307.14079](https://arxiv.org/abs/2307.14079))
- Wadhwa V 2018 Quantum computers may be more of an imminent threat than AI (available at: www.washingtonpost.com/news/innovations/wp/2018/02/05/quantum-computers-may-be-more-of-an-imminent-threat-than-ai/)
- Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L and Coles P J 2021 Noise-induced barren plateaus in variational quantum algorithms *Nat. Commun.* **12** 6961
- Warren R H 2013a Adapting the traveling salesman problem to an adiabatic quantum computer *Quantum Inf. Process.* **12** 1781–5
- Warren R H 2013b Numeric experiments on the commercial quantum computer *Notices of the AMS* **60** 1484
- Warren R H 2017 Small traveling salesman problems *J. Adv. Appl. Math.* **2** 101–7
- Warren R H 2018 Mathematical methods for a quantum annealing computer *J. Adv. Appl. Math.* **3** 82–90
- Warren R H 2021 A benchmark for quantum optimization: the traveling salesman *Quantum Inform. Comput.* **21** 557–62
- Wilkening S, Lefterovici A-I, Binkowski L, Perk M, Fekete S and Osborne T J 2023 A quantum algorithm for the solution of the 0–1 Knapsack problem (arXiv: [2310.06623](https://arxiv.org/abs/2310.06623))
- Willsch M, Willsch D, Jin F, De Raedt H and Michielsen K 2020 Benchmarking the quantum approximate optimization algorithm *Quantum Inf. Process.* **19** 1–24
- Wilson G and Pawley G 1988 On the stability of the travelling salesman problem algorithm of Hopfield and Tank *Biol. Cybern.* **58** 63–70
- Wilson M, Stromswold R, Wudarski F, Hadfield S, Tubman N M and Rieffel E G 2021 Optimizing quantum heuristics with meta-learning *Quantum Mach. Intell.* **3** 1–14
- Wu C, Wang J and Zuo F 2024 From maximum cut to maximum independent set (arXiv: [2408.06758](https://arxiv.org/abs/2408.06758))
- Xue C, Chen Z-Y, Wu Y-C and Guo G-P 2021 Effects of quantum noise on quantum approximate optimization algorithm *Chin. Phys. Lett.* **38** 030302
- Zhou L, Wang S-T, Choi S, Pichler H and Lukin M D 2020 Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices *Phys. Rev. X* **10** 021067