
Quantum Theory and Application of Contextual Optimal Transport

Nicola Mariella¹ Albert Akhriev¹ Francesco Tacchino² Christa Zoufal² Juan Carlos Gonzalez-Espitia^{2,3}
Benedek Harsanyi^{4,5} Eugene Koskin^{1,6} Ivano Tavernelli² Stefan Woerner² Marianna Rapsomaniki⁴
Sergiy Zhuk¹ Jannis Born⁴

Abstract

Optimal Transport (OT) has fueled machine learning (ML) across many domains. When paired data measurements (μ, ν) are coupled to covariates, a challenging conditional distribution learning setting arises. Existing approaches for learning a *global* transport map parameterized through a potentially unseen context utilize Neural OT and largely rely on Brenier’s theorem. Here, we propose a first-of-its-kind quantum computing formulation for amortized optimization of contextualized transportation plans. We exploit a direct link between doubly stochastic matrices and unitary operators thus unravelling a natural connection between OT and quantum computation. We verify our method (QontOT) on synthetic and real data by predicting variations in cell type distributions conditioned on drug dosage. Importantly we conduct a 24-qubit hardware experiment on a task challenging for classical computers and report a performance that cannot be matched with our classical neural OT approach. In sum, this is a first step toward learning to predict contextualized transportation plans through quantum computing.

1. Introduction

Optimal transport (OT) (Villani, 2008) provides a mathematical framework for finding transportation plans that minimize the cost of moving resources from a source to a target distribution. The cost is defined as a distance or a dissimilarity measure between the source and target points,

¹IBM Quantum, IBM Research Europe - Dublin ²IBM Quantum, IBM Research Europe - Zurich, Switzerland ³Politecnico di Milano, Milan, Italy ⁴IBM Research, IBM Research Europe - Zurich, Switzerland ⁵École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland ⁶University College Dublin, Ireland. Correspondence to: <nicola.mariella@ibm.com>, <jab@zurich.ibm.com>.

and the OT plan aims to minimize this cost while satisfying certain constraints. OT theory has found applications across several fields, including biology where it gained popularity in single-cell analysis, an area of research rich in problems of mapping cellular distributions across distinct states, timepoints, or spatial contexts (Klein et al., 2023). Notable biological tasks are reconstructing cell evolution trajectories (Schiebinger et al., 2019), predicting responses to therapeutic interventions (Bunne et al., 2023; 2022) and aligning datasets across different omic modalities (Cao et al., 2022). From the OT perspective, source and target distributions are measurements of biomolecules of single cells.

In many OT applications, data measurements μ_i (initial state) and ν_i (final state) are coupled to a context \mathbf{p}_i that induces μ_i to develop into ν_i . One thus might aspire to *learn* a global transport map T parameterized through \mathbf{p}_i and thus facilitate the *prediction* of target states $\hat{\nu}_j$ from source states μ_j , even for an unseen context \mathbf{p}_j (cf. Figure 1A). This work is largely based on Brenier’s theorem (1987) which postulates the existence of a unique OT map T given by the gradient of a convex function, i.e., $T = \nabla f_\theta$. Makkuva et al. (2020) showed that OT maps between two distributions can be learned through neural OT solvers by using a minimax optimization where f_θ is an input convex neural network (ICNN, Amos et al. (2017)). A notable example of such a neural OT approach is CondOT (Bunne et al., 2022) which estimates transport maps conditioned on a context variable and learned from quasi-probability distributions (μ_i, ν_i) , each linked to a context variable \mathbf{p}_j ¹. Limitations of such approaches include the dependence on squared Euclidean cost induced by Brenier’s theorem (Peyré et al., 2019) or the unstable training due to the min-max formulation in the dual objective as well as the architectural constraints induced by the partial ICNN. The Monge Gap (Uscidda & Cuturi, 2023), an architecturally agnostic regularizer to estimate OT maps with any cost C , overcomes these challenges; however, unlike CondOT it cannot generalize to new context. This is resolved in our concurrent work on the *Conditional Monge Gap* (Harsanyi et al., 2024).

¹We use "contextual" rather than "conditional" to differentiate from OT on conditional probabilities (Tabak et al., 2021)

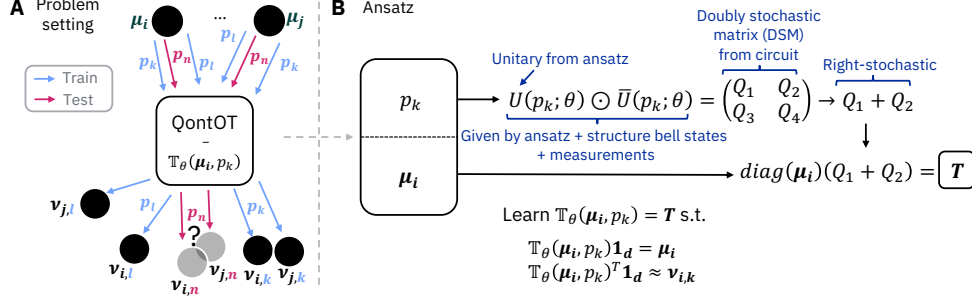


Figure 1: **A)** Contextual OT is a conditional distribution learning problem. **B)** Our proposed ansatz receives the context (p_k) and the initial distribution μ_i and produces a DSM that can be rescaled to a transport plan T with marginal distributions μ_i and $\hat{\nu}_{i,k}$.

On a separate realm, quantum computing (QC) offers a new paradigm with the potential to become practically useful in ML (Havlíček et al., 2019; Liu et al., 2021a; Harrow et al., 2009; Huang et al., 2021) and fuel applications in, e.g., life sciences (Basu et al., 2023) or high-energy physics (Di Meglio et al., 2023). A general hope of quantum ML lies in the identification of structures that can be produced more easily with quantum operators rather than classical principles. Here, we propose a quantum contextual OT approach inspired by a natural link between OT and unitary operators, a fundamental concept in QC. This link is formed thanks to doubly stochastic matrices (DSMs). DSMs are real, nonnegative square matrices with row and column sums of 1. They are a core structure behind the rise of OT in ML, for details we refer the reader to Appendix A. While ensuring the row *or* column constraint is trivial (left- or right-stochastic matrices can be obtained with a simple softmax), producing DSMs parametrically is challenging with classical ML, typically iterative, non-parametric approximations like the Sinkhorn algorithm (1967) are used. Thanks to the link of DSMs and unitary operators, we can turn the analytical problem of computing OT plans into a parameterizable approach to *estimate* them. In contrast to existing neural OT methods like CellOT (Bunne et al., 2023), our quantum formulation does not depend on Brenier’s theorem (i.e., it is cost-agnostic) and unlike CellOT and the Monge Gap it estimates transportation plans explicitly which is more interpretable, e.g. the map topology can be studied.

Our key contributions can be summarized as follows:

1. We are first to bridge QC with OT and ML. As shown in Figure 1B, we devise an ansatz that performs amortized optimization of contextual OT plans (or contextual DSMs) $Q \in \Omega_d$ given a context p_{new} (Section 3.1). Given an unseen distribution μ_{new} , we extract a structure that can be rescaled to a transport map with a desired initial marginal distribution (Section 3.2).
2. We identify and leverage a previously unreported algebraic link between unitary operators and DSMs which connects OT and QC. We prove that the constraints for a

DSM can be obtained with quantum, yielding a "quantum inductive bias for DSMs". This is notable because it is currently unknown whether a similarly natural classical approach to parametrically produce DSMs exists.

3. We report a promising result on the relaxed assignment problem (i.e., contextual prediction of DSMs) obtained with 24-qubits on real hardware (IBM Quantum Platform) that outperforms a classical neural OT approach.

The remaining paper begins with the contextual OT problem and proceeds with the quantum theory and the details of the ansatz (i.e., a parametric quantum circuit to approximate a quantum state) for encoding DSMs and transport plans. As a proof-of-concept, we first verify our method on synthetic and real drug perturbation data (with drug dosage as context). We then turn to a more constrained task, the contextual relaxed assignment problem which emphasizes better the strengths of our approach compared to classical neural OT.

2. Preliminaries

2.1. Notation

The sets of non-negative and positive reals are defined respectively as \mathbb{R}_+ and \mathbb{R}_{++} . The n -dimensional vector of ones is denoted by $\mathbf{1}_n$, the $n \times n$ identity matrix as \mathbb{I}_n and the $n \times n$ matrix of all ones as $J_n := \mathbf{1}_n \mathbf{1}_n^\top$. The set (with group structure) of unitary matrices of order n is denoted $\mathfrak{U}(n)$. Given a set X we denote by $\text{conv}(X)$ the convex hull of X , that is the minimal and unique convex set containing X . We define the linear operator $\text{diag}(\cdot)$ as the mapping from a vector $\mathbf{v} \in \mathbb{C}^n$ to a diagonal matrix having as diagonal the same vector. The trace operator is denoted $\text{Tr}(\cdot)$. The notation $(\cdot)^\dagger$ is used to denote the complex conjugate transpose of the matrix or vector argument. Let A, B be matrices of the same size, then we denote by $A \odot B$ their *Hadamard product*, which is defined as the entry-wise multiplication $(A \odot B)_{i,j} = A_{i,j} B_{i,j}$. We use the *Bra-ket notation* to denote the quantum states and their dual. An integer symbol i in the Ket $|i\rangle$, refers to the i -th basis vector in the computational basis. Also, we make use of the

following shortcut $|ij\rangle := |i\rangle \otimes |j\rangle$.

2.2. Linear Algebra of Doubly Stochastic Matrices

Let Δ_n denote the *probability simplex* in $n - 1$ dimensions,

$$\Delta_n = \{ \mathbf{v} \in \mathbb{R}_+^n \mid \mathbf{1}_n^\top \mathbf{v} = 1 \}. \quad (1)$$

Fundamental to this work is the *Birkhoff polytope* $\Omega_n := \mathcal{N}(\mathbf{1}_n, \mathbf{1}_n)$ (Brualdi, 2006), defined as the (convex) set of $n \times n$ *doubly stochastic matrices* (DSM). A DSM $Q \in \Omega_n$ is a real, non-negative matrix with row/column sums of 1

$$Q\mathbf{1}_n = \mathbf{1}_n, \quad Q^\top \mathbf{1}_n = \mathbf{1}_n, \quad Q_{i,j} \geq 0. \quad (2)$$

Permutation matrices are special cases of DSMs where the entries belong to $\{0, 1\}$. Moreover, permutations are the only DSM that are unitary, so for any permutation P we have that $PP^\top = P^\top P = \mathbb{I}_n$. A convex polytope is defined as the convex hull of a finite set of objects called extremes, in the case of the Birkhoff polytope, the extremes are the permutation matrices, this result is known as the *Birkhoff theorem*. Every DSM $Q \in \Omega_n$ can be decomposed as a convex combination of permutation matrices, that is

$$Q = \sum_{i=1}^N \lambda_i P_i, \quad (3)$$

for some probability vector $\boldsymbol{\lambda} \in \Delta_N$, $n \times n$ permutation matrices $\{P_i\}$, and the number of extreme points $N \leq n^2$. We note that the decomposition is not unique. Given a positive integer n , the number of permutation matrices is $n!$, which is the number of extreme points of Ω_n . However, the Birkhoff theorem bounds the number of permutations N required to represent an arbitrary DSM by $n \cdot n$.

Another fundamental structure is the set \mathfrak{S}_n (a subset of the Birkhoff polytope Ω_n) of $n \times n$ *unistochastic matrices*. Given any $n \times n$ unitary matrix U , the matrix obtained by substituting each element of U with its absolute value squared, is unistochastic. In other words, let $U \in \mathfrak{U}(n)$, then $U \odot \bar{U}$ is doubly stochastic, where $\bar{U} = (U^\dagger)^\top$. The latter result is an implication of unitarity. The set of unistochastic matrices is a non-convex proper subset of the Birkhoff polytope, however the $n \times n$ permutations matrices² all belong to such set, hence its convex hull corresponds to the Birkhoff polytope, that is $\text{conv}(\mathfrak{S}_n) = \Omega_n$. The constraints required for an arbitrary DSM to be unistochastic are still unknown³.

Let $|b_n\rangle$ denote the quantum state acting on $2n$ qubits⁴,

²A permutation matrix P fulfills the unitarity constraints $P^\top P = PP^\top = \mathbb{I}_n$, so $P \in \mathfrak{U}(n)$ (unitary). Also $P \odot \bar{P} = P$, then $P \in \mathfrak{S}_n$ (unistochastic).

³Unistochastic matrices of order 3 cover $\approx 75\%$ of the Birkhoff polytope (Dunkl & Życzkowski, 2009).

⁴We consider Bell's states (Nielsen & Chuang, 2011) which are defined on a bipartite system and maximize the Von Neumann entanglement entropy.

consisting of n maximally entangled states on 2 qubits, so

$$\begin{aligned} |b_n\rangle &= \frac{1}{\sqrt{2^n}} \sum_{i_1, \dots, i_n=0}^1 (|i_1\rangle |i_2\rangle \cdots |i_n\rangle) \otimes (|i_1\rangle |i_2\rangle \cdots) \\ &= \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle \otimes |i\rangle. \end{aligned} \quad (4)$$

We denote with vec_r the *row-major vectorization operator*. Given a $n \times n$ matrix M in \mathbb{C} , the latter operator is defined by the rule

$$\text{vec}_r(M) := \sum_{i=0}^{n-1} M |i\rangle \otimes |i\rangle. \quad (5)$$

Moreover, we will be using the well-known identities linking vectorization to the Kronecker product⁵

$$\text{vec}_r(M) = (M \otimes \mathbb{I}_n) \text{vec}_r(\mathbb{I}_n) \quad (6a)$$

$$\text{vec}_r(M^\top) = (\mathbb{I}_n \otimes M) \text{vec}_r(\mathbb{I}_n). \quad (6b)$$

We note that the state $|b_n\rangle$ in Eq. (4), corresponds to the vectorization of the identity operator up to a scalar multiple, that is $\frac{1}{\sqrt{2^n}} \text{vec}_r(\mathbb{I}_2^{\otimes n}) = |b_n\rangle$. The following lemma establishes a relation between unitary operators and their vectorization.

Lemma 2.1. *Let $\{U_i\}$ be a set of unitary operators $U_i \in \mathfrak{U}(n)$ such that $\text{Tr}(U_i U_j^\dagger) = n\delta_{ij}$, that is the unitaries are orthogonal w.r.t. the Frobenius inner product. Then the set $\{\text{vec}_r(U_i)\} \subset \mathbb{C}^{n^2}$ consists of orthogonal vectors, that is*

$$\text{vec}_r(U_j)^\dagger \text{vec}_r(U_i) = \text{Tr}(U_i U_j^\dagger) = n\delta_{ij}. \quad (7)$$

Proof in Appendix C.3.

2.3. The Canonical OT Problem

In the Kantorovich relaxation of the Monge problem (Peyré et al., 2019), $C \in \mathbb{R}^{n \times m}$ is a non-negative matrix representing the cost of mass displacement from entity i to j (so C is called *cost matrix* hereafter). Let $\boldsymbol{\mu}, \boldsymbol{\nu}$ be strictly positive real vectors (i.e., $\boldsymbol{\mu} \in \mathbb{R}_{++}^n$ and $\boldsymbol{\nu} \in \mathbb{R}_{++}^m$), representing the quasi-probability discrete distributions⁶ (also referred as states) for the source and destination entities. The discrete (regularized) Kantorovich's OT problem is defined as

$$\min_{Q \in \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\nu})} \text{Tr}(QC^\top) + \gamma h(Q), \quad (8a)$$

$$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\nu}) = \{ Q \in \mathbb{R}_+^{n \times m} \mid Q\mathbf{1}_m = \boldsymbol{\mu}, Q^\top \mathbf{1}_n = \boldsymbol{\nu} \} \quad (8b)$$

⁵In the context of quantum information theory the identity in Eq. (6a) is known as the Choi-Jamiołkowski correspondence.

⁶We say that a d -dimensional vector \mathbf{v} is a quasi-probability discrete distribution when it is non-negative and non-zero. Then $\mathbf{v} / (\mathbf{1}_d^\top \mathbf{v})$ is a probability distribution.

where h is a regularization function (Cuturi, 2013) with trade-off $\gamma \geq 0$. The set $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\nu})$ is the *transportation polytope* (Brualdi, 2006) whose elements are *transportation maps*. Given a map $T \in \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\nu})$, $T_{i,j}$ represents the mass moved from source i to destination j (cf. Figure A1). As noted, the elements of the *Birkhoff polytope*, a special case of the transportation polytope, are the DSMs which are the solutions to the (relaxed) *assignment problem*, where quasi-probability distributions are uniform $\boldsymbol{\mu} = \boldsymbol{\nu} = \mathbf{1}_n$.

2.4. The Contextual OT Problem

Neural OT is concerned with *learning* the optimal transport between distributions from samples (Makkuva et al., 2020; Korotin et al., 2023), s.t. $\hat{\nu}_i$ can be estimated from unseen μ_i . Contextual OT generalizes this scenario: if (ν_i, μ_i) are not observed in isolation but linked to a context \mathbf{p}_i , a conditional distribution learning task arises (Bunne et al., 2022; Nguyen et al., 2024; Harsanyi et al., 2024).

Formally, let $\mathcal{K}_d := \Delta_d \cap \mathbb{R}_{++}^d$ denote the subset of the probability simplex with vectors presenting non-zero components⁷. We consider a dataset of contextualised measures each represented by a tuple $(\mathbf{p}_i, (\boldsymbol{\mu}_i, \boldsymbol{\nu}_i)) \in \mathcal{X} \times \mathcal{K}_d^2$, where the vector $\mathbf{p}_i \in \mathcal{X} \subseteq \mathbb{R}^s$ defines the context. The initial and final states $\boldsymbol{\mu}_i$ and $\boldsymbol{\nu}_i$ are from the same set \mathcal{K}_d . The cost matrix C is not required to be constant across all samples, and can be interpreted as a materialization of the perturbation. At inference time, we are given an unseen perturbation \mathbf{p}_{new} and initial state $\boldsymbol{\mu}_{\text{new}}$, and aim to predict a transportation map $T^* \in \mathcal{N}(\boldsymbol{\mu}_{\text{new}}, \boldsymbol{\nu}^*)$, s.t. marginalization yields the final states $\boldsymbol{\nu}^*$. At training time, we use classical OT solvers to obtain a map for each sample, providing a list of tuples (\mathbf{p}_i, T_i) where $T_i \in \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\nu}_i)$ solves the i -th OT problem.

3. Quantum Formulation

Our quantum formulation leverages the following fundamental concept. Let $\overline{(\cdot)}$ denote the complex conjugate of the argument (i.e. in the case of a matrix argument, the transpose of the adjoint), and \odot the Hadamard product between matrices. If U is a unitary matrix, then $U \odot \overline{U} \in \Omega_n$ is a DSM. Hence we can represent (with some approximation) the solution of the assignment problem using unitary operators. This principle produces DSM independently of the construction of the unitary U , which offers great freedom in the choice of the ansatz for U supporting both variational and possibly kernel-based learning. Furthermore, such natural link between transportation maps and unitary operators may lead to quantum models enjoying better expressivity compared to classical counterparts (Bowles et al., 2023; Abbas et al., 2021; Liu et al., 2021b; Anschuetz & Gao, 2024). For example, the multitask model by Bowles et al. (2023)

⁷To avoid degeneracy, Remark 2.1 (Peyré et al., 2019).

corresponds to our DSM prediction which intrinsically possesses the required linear bias (i.e., all row and column sums equal 1).

3.1. Quantum Circuit for the Birkhoff Polytope

In this section, we assume that the part of the circuit acting on the first n qubits has a dimension comparable to the input $d = 2^n$, where d is the number of entities for the discrete distributions considered in Section 2.4. Let U_p be a parametric unitary operator acting on the bipartite Hilbert space $(\mathbb{C}^2)^{\otimes m} \otimes (\mathbb{C}^2)^{\otimes n}$, and $m \in \mathbb{N}$ such that the classical simulation of a circuit on $m + n$ qubits is intractable (with $m \geq n$) in general⁸. The operator U_p depends on the input vector $\mathbf{p} \in \mathcal{X}$ (perturbation) as well as on the learning parameters $\boldsymbol{\theta}$. To prove the construction, we consider the *Operator-Schmidt decomposition* (Appendix B) of U_p (on $m + n$ qubits) determined by the quantum-mechanical subsystems A_1, B_1 , consisting of respectively m and n qubits.

$$U_p(\mathbf{p}, \boldsymbol{\theta}) = \sum_{i=1}^{d^2} \lambda_i V_i(\mathbf{p}, \boldsymbol{\theta}) \otimes W_i(\mathbf{p}, \boldsymbol{\theta}) \quad (9)$$

with $\{V_i\}$ and $\{W_i\}$ being sets of unitary operators orthogonal w.r.t. the Frobenius inner product $\text{Tr}(V_i V_j^\dagger) = 2^m \delta_{ij}$. The same follows similarly for the set $\{W_i\}$. As a consequence of the SVD⁹, we have $\lambda_i \geq 0$, with unitarity of U_p implying $\sum_i \lambda_i^2 = 1$. Notably, the matrix U_p depends on the input and the parameters vectors, then the components of the Operator-Schmidt decomposition, namely λ_i, V_i and W_i , are functions of $(\mathbf{p}, \boldsymbol{\theta})$. Moreover, to assure the consistency of the formulation we impose the *Schmidt rank* of U_p (i.e. the number of strictly positive λ_i) to be greater than one¹⁰. Using the unitary U_p (omitting the dependency from \mathbf{p} and $\boldsymbol{\theta}$ for clarity) and the states $|b_n\rangle$ and $|b_m\rangle$ (defined in Eq. (4)) we obtain the following state (on $2m + 2n$ qubits)

$$\begin{aligned} |\varphi\rangle &= (\mathbb{I}_2^{\otimes m} \otimes U_p \otimes \mathbb{I}_2^{\otimes n}) \cdot (|b_m\rangle \otimes |b_n\rangle) \\ &\stackrel{(9)}{=} \sum_k \lambda_k (\mathbb{I}_2^{\otimes m} \otimes V_k) |b_m\rangle \otimes (W_k \otimes \mathbb{I}_2^{\otimes n}) |b_n\rangle \\ &\stackrel{(6a)}{=} \sum_k \lambda_k \frac{\text{vec}_r(V_k^\top)}{\sqrt{2^m}} \otimes \frac{\text{vec}_r(W_k)}{\sqrt{2^n}}, \end{aligned} \quad (10)$$

where $\text{vec}_r(\cdot)$ is the vectorization operator defined in Section 2.2. We note that the last equality is obtained using the

⁸We note that DSMs reside in a classical memory so they have a reduced space resources complexity. If we were not using m auxiliary qubits for producing the DSM, then the circuit would require a logarithmic number of qubits w.r.t. the data size, so the unitary would have a reduced space complexity. Consequently, the resulting circuit would be tractable to classical simulation.

⁹The Operator-Schmidt decomposition is obtained through SVD.

¹⁰We impose the Schmidt rank for U_p w.r.t. the split on $m + n$ qubits, to be > 1 . Otherwise the partial trace, introduced on the m qubits, makes the part of the unitary on m qubits uninformative.

identities in Eq. (6a) and Eq. (6b). Now, we partition the Hilbert space on which $|\varphi\rangle$ lays into two subsystems. The first, A_2 , consists of the first $2m$ qubits (auxiliary qubits). The second, B_2 , takes the last $2n$ qubits (data qubits)¹¹. We obtain the mixed state ρ by applying the partial trace over the system A_2 , to the pure state $|\varphi\rangle\langle\varphi|$, that is

$$\rho = \text{Tr}_{A_2} (|\varphi\rangle\langle\varphi|) \quad (11a)$$

$$\begin{aligned} &= \frac{1}{2^{m+n}} \sum_{i,j} \lambda_i \lambda_j \text{Tr} \left(V_i V_j^\dagger \right) \text{vec}_R(W_i) \text{vec}_R(W_j)^\dagger \\ &\stackrel{(7)}{=} \frac{1}{2^n} \sum_i \lambda_i^2 \text{vec}_R(W_i) \text{vec}_R(W_i)^\dagger. \end{aligned} \quad (11b)$$

Recall that by the Operator-Schmidt decomposition and unitarity of U_p we have that $\sum_i \lambda_i^2 = 1$. Given that the action of the unitary $U_p(\mathbf{p}, \boldsymbol{\theta})$ is generally not classically efficiently simulable, the state ρ has the potential to represent correlations that cannot be captured with classical models. Moreover, here we can appreciate the role of the auxiliary m qubits, that is enlarging the function space as a result of the convex combination of density matrices in Eq. (11b). Indeed we note that if $m = 0$, then the number of terms in Eq. (11b) reduces to 1. The recovery of the DSM is completed with the projective measurements explained in the lemma that follows (see resulting circuit in Figure 2a).

Lemma 3.1. *Let*

$$p(i, j) := 2^n \text{Tr}(\rho |ij\rangle\langle ij|) \quad (12a)$$


for $i, j \in [0 \dots d-1]$ and ρ as defined in Eq. (11a). Let $\{\mathbf{e}_i\}$ be the set of canonical basis vectors (with index i starting from 0) for the vector space \mathbb{R}^{2^n} . Then the matrix

$$Q = \sum_{i,j=0}^{2^n-1} p(i, j) \mathbf{e}_i \mathbf{e}_j^\top \quad (12b)$$

is doubly stochastic.

Proof in Appendix C.3. In the latter result, the rank 1 matrix $\mathbf{e}_i \mathbf{e}_j^\top$ corresponds to the $2^n \times 2^n$ matrix with 1 in position i, j and zeros elsewhere (i.e. the canonical basis for $2^n \times 2^n$ matrices in \mathbb{R}). In other words, given the density matrix ρ prepared as in Eq. (11a), the expectations w.r.t. the observables $|ij\rangle\langle ij|$ provide the corresponding (i, j) entry of the resulting matrix. In practice, fixed an observable for entry (i, j) , we obtain a single bit of information for each execution of the circuit, so the resulting matrix is guaranteed to fulfil the constraints of doubly stochasticity when the number of shots approaches infinity. However, the convexity of the Birkhoff polytope offers great advantage in terms of restoring the DSM on a circuit, more details are given in Appendix C.1. Our circuit is shown in Figure 2a.

¹¹We note that the systems A_2 and B_2 contain respectively the systems A_1 and B_1 , defined for the unitary in Eq. (9).



(a) Section 3.1: DSM-encoding circuit (b) Section 3.2: Embedded transport map

Figure 2: Circuit structures for the transportation map prediction. The registers $\{i_k\} \cup \{j_k\}$ represent the bits for the index (i, j) related to the entry $Q_{i,j}$ of the resulting DSM. The registry $\{a_k\}$ refers to the $2m$ auxiliary qubits as per Section 3.1. Regarding Figure 2a, we remark that the registry i has been added for construction reasons, however in practice it can be removed and substituted with a classical uniform sampling (using the computational basis states on n qubits) over the registry j . Consequently, the number of required qubits for DSM-encoding can be reduced to $2m + n$. In Figure 2b, we have applied that trick to embed transportation maps, with $|i_1\rangle \in \{|0\rangle, |1\rangle\}$, as per Section 3.2.

3.2. Embedding of Transportation Maps

Since in our applications, the initial distribution $\boldsymbol{\mu}$ is user-provided at inference time, the problem is twofold; (1) embedding the transport map into a DSM to fit the representation presented in Section 3.1, and (2) predicting maps which can be rescaled to an arbitrary initial distribution. Given a data set $\{(\mathbf{p}_i, T_i)\}$ (i.e., tuples of contexts and transport maps), we assume $T_i \in \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\nu}_i)$ $\boldsymbol{\mu}_i, \boldsymbol{\nu}_i \in \mathbb{R}_{++}^d$, i.e., the margins of the transport maps are strictly positive¹². Let $\mathbf{v} \in \mathbb{R}^d$, and $D_{\mathbf{v}} = \text{diag}(\mathbf{v})$ the $d \times d$ diagonal matrix having the elements of the vector \mathbf{v} as diagonal elements. Now, given T_i as defined above, we define $\hat{T}_i := D_{\boldsymbol{\mu}_i}^{-1} T_i$, and observe that $\hat{T}_i \mathbf{1}_d = D_{\boldsymbol{\mu}_i}^{-1} T_i \mathbf{1}_d = D_{\boldsymbol{\mu}_i}^{-1} \boldsymbol{\mu}_i = \mathbf{1}_d$, that is $\hat{T}_i \in \mathcal{N}(\mathbf{1}_d, \boldsymbol{\nu}'_i)$ is a right stochastic matrix¹³, with $\boldsymbol{\nu}'_i = \hat{T}_i^\top \mathbf{1}_d$. At inference, when given a perturbation \mathbf{p}_i , the model predicts a right stochastic matrix $\hat{T} \in \mathcal{N}(\mathbf{1}_d, \boldsymbol{\nu}')$ for some $\boldsymbol{\nu}' \in \mathbb{R}_{++}^d$. The latter, alongside the user-provided initial distribution $\boldsymbol{\mu}$, determines the final predicted map $T = D_{\boldsymbol{\mu}} \hat{T}$, s.t. $T \mathbf{1}_d = D_{\boldsymbol{\mu}} \hat{T} \mathbf{1}_d = D_{\boldsymbol{\mu}} \mathbf{1}_d = \boldsymbol{\mu}$. In other words, we learn the transportation pattern in a margin-independent fashion and rescale to the required margin at inference time. Note that, when the context is $\mathbf{0}$ (null perturbation) then $\hat{T}_i = \mathbb{I}_d$. Given some $\boldsymbol{\mu}$ we have $D_{\boldsymbol{\mu}} \hat{T}_i = D_{\boldsymbol{\mu}}$, hence $D_{\boldsymbol{\mu}} \mathbf{1}_d = D_{\boldsymbol{\mu}}^\top \mathbf{1}_d = \boldsymbol{\mu} = \boldsymbol{\nu}$, ergo the initial and final distributions agree (consistent with the notion of null perturbation), inducing a stationarity inductive bias.

To confirm generality: any transport map $T \in \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\nu})$ with $\boldsymbol{\mu}, \boldsymbol{\nu} \in \mathbb{R}_{++}^d$, can be decomposed as $T = DR$ where $R = \text{diag}(\boldsymbol{\mu})^{-1} T$ is right stochastic and $D = \text{diag}(\boldsymbol{\mu})$ a

¹²Justification on strict positivity of margins in Section 2.4.

¹³ $Q \in \mathbb{R}_{++}^{n \times n}$ is right stochastic iff $Q \mathbf{1}_n = \mathbf{1}_n$. All DSMs are right stochastic but the converse is not true.

positive diagonal¹⁴. Conversely, the product DR of a right stochastic R and a positive diagonal D (both of order d) is a transport map with $\boldsymbol{\mu} = D\mathbf{1}_d$ and $\boldsymbol{\nu} = (DR)^\top \mathbf{1}_d$.

Aggregation scheme. To complete the structure, we now expand on the link between the $d \times d$ right stochastic matrix \widehat{T} and a $2d \times 2d$ DSM Q . This step is necessary since the formulation in Section 3.1 produces only a DSM. First, consider the DSM block decomposition

$$Q = \begin{pmatrix} Q_1 & Q_2 \\ Q_3 & Q_4 \end{pmatrix} \in \Omega_{2d}, \quad (13)$$

with $Q_i \in \mathbb{R}_+^{d \times d}$. Now, note that $(Q_1 \quad Q_2) \mathbf{1}_{2d} = \mathbf{1}_d$ implies $(Q_1 + Q_2) \mathbf{1}_d = \mathbf{1}_d$. We embed the right stochastic matrix \widehat{T}_i into the sum $Q_1 + Q_2$ of the top quadrants of a DSM $Q \in \Omega_{2d}$. Since this structure does not consider the submatrices Q_3 and Q_4 , in Appendix C.2 we describe a custom designed ansatz that takes into account such invariant. Moreover, as depicted in Figure 2b, we can obtain the sum $Q_1 + Q_2$ directly from the state preparation and measurements. Specifically, by initialising the registry j_2 to $|0\rangle$ we obtain the top half of the matrix (w.r.t. rows) and by tracing out the same registry we mimic the sum of the top two quadrants¹⁵. We call this "atop" aggregation. To obtain the number of required qubits, let m be the number of qubits (as per Section 3.1) that makes the function space achievable by the ansatz hard to be computed classically. Also, let $d = 2^n$ and consider $d \times d$ transportation maps, then using the reduction introduced in Figure 2b, the circuit requires $2(n + m + 1)$ qubits¹⁶ where $m \geq n + 1$. In practice, we set $m := n + 1$ unless indicated otherwise.

3.3. Training Objective

Let $f : \mathcal{X} \rightarrow \Omega_d^{(r)}$ be a function from the set of perturbations \mathcal{X} to the set of $d \times d$ row-stochastic matrices, and let \mathcal{F} be the function space of such functions related to our model. Then, given the training set $\{(\mathbf{p}_i, T_i)\}$ we define our learning problem via the loss

$$\mathcal{L}_T = \min_{f \in \mathcal{F}} \sum_i \|D_{\boldsymbol{\mu}_i} f(\mathbf{p}_i) - T_i\|_F^2, \quad (14)$$

where $\boldsymbol{\mu}_i = T_i \mathbf{1}_d$ is the initial distribution for the i -th training sample and $\|\cdot\|_F$ is the Frobenius matrix norm. An optimal function in \mathcal{F} that minimizes the loss \mathcal{L}_T is denoted f^* . At inference time, given the initial (quasi-)distribution $\boldsymbol{\mu} \in \mathbb{R}_{++}^d$ and the perturbation $\mathbf{p} \in \mathcal{X}$, the predicted transportation map is obtained as $T = D_{\boldsymbol{\mu}} f^*(\mathbf{p})$. Alternatively,

¹⁴Indeed $R\mathbf{1}_d = \text{diag}(\boldsymbol{\mu})^{-1} T \mathbf{1}_d = \text{diag}(\boldsymbol{\mu})^{-1} \boldsymbol{\mu} = \mathbf{1}_d$ (right stochastic), so $DR = \text{diag}(\boldsymbol{\mu}) \text{diag}(\boldsymbol{\mu})^{-1} T = T$.

¹⁵see *principle of implicit measurement* (Nielsen & Chuang, 2011).

¹⁶note that removing registry i and using classical sampling of matrix rows could reduce to $(n + 1) + 2m$ qubits

the predicted target distribution $\boldsymbol{\nu} = (D_{\boldsymbol{\mu}} f^*(\mathbf{p}))^\top \mathbf{1}_d$ can be directly optimized via:

$$\mathcal{L}_M = \min_{f \in \mathcal{F}} \sum_i \left\| (D_{\boldsymbol{\mu}_i} f(\mathbf{p}_i))^\top \mathbf{1}_d - \boldsymbol{\nu}_i \right\|_2^2. \quad (15)$$

This can be interpreted as a weakly-supervised learning of transportation maps and is considered for comparison purposes. The ansatz parameters are obtained via gradient-free optimization with COBYLA (Powell, 1994).

Evaluation. Accuracy of transportation plan prediction is measured twofold. First, the relative Frobenius norm

$$F(\bar{T}_i, T_i) = \frac{\|\bar{T}_i - T_i\|_F}{\|\bar{T}_i\|_F} \quad (16)$$

where $\bar{T}_i = D_{\boldsymbol{\mu}_i} f^*(\mathbf{p}_i)$. Secondly, we report the sum of the absolute errors (SAE). Accuracy of the predicted marginals $\bar{\boldsymbol{\nu}}$ is measured through L_2 norm and R^2 .

3.4. Multidimensional OT

This subsection shows how to estimate the bare minimum of necessary quantum resources, including even the case of discrete multidimensional OT (Solomon, 2018), reflecting that many OT applications utilize multivariate rather than univariate measures (as assumed above). Let the source data have K covariates, i.e. $x_i = (x_i^1, \dots, x_i^K)$. Assume that each covariate is defined on a discrete sample space \mathbb{X}_k with cardinality d_k^μ and d_k^ν ($k \in [1, K]$) for source and target data, and define the probability space $\mathcal{P}_k = (\mathbb{X}_k, \mathcal{F}_k, \boldsymbol{\mu}_k)$, where $\mathcal{F} = \sigma(\mathbb{X}_k)$ is the σ -algebra generated by \mathbb{X}_k , and $\boldsymbol{\mu}_k$ a measure on $(\mathbb{X}_k, \mathcal{F}_k)$. Then, the multi-dimensional source measure space is written as $\mathcal{P} = \left(\bigotimes_{k=1}^K \mathbb{X}_k, \bigotimes_{k=1}^K \mathcal{F}_k, \bigotimes_{k=1}^K \boldsymbol{\mu}_k \right)$. Analogously, the target measure space is $\tilde{\mathcal{P}} = \left(\bigotimes_{l=1}^L \mathbb{Y}_l, \bigotimes_{l=1}^L \mathcal{G}_l, \bigotimes_{l=1}^L \boldsymbol{\nu}_l \right)$. With $K = L = 1$, $d^\mu = P$, $d^\nu = R$, we recover the case discussed in Section 2.3, i.e., $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are vectors, and the transportation plan $T \in \mathbb{R}^{P \times R}$. In general, $\boldsymbol{\mu} \in \bigotimes_{k=1}^K \mathbb{R}_{++}^{d_k^\mu}$ and $\boldsymbol{\nu} \in \bigotimes_{l=1}^L \mathbb{R}_{++}^{d_l^\nu}$, i.e., the size is governed by the state space cardinality. Since the source and target distributions are represented by K and L rank tensors, the cost function will be a $(K + L)$ -tensor. Assuming identical source and target spaces with K covariates, and d states per covariate (i.e., $K = L$ and $d_k^\mu = d_k^\nu = d \forall k \in [0, K]$) the OT plan spans $\mathcal{O}(d^K)$ rows/columns. Notably, the discrete N-dimensional OT problem is $\#P$ -hard (Taşkesen et al., 2023). Computing explicit OT plans thus quickly becomes demanding.

Now consider the application by Bunne et al. (2022) on predicting single-cell perturbation responses among 200 cells with 50 gene-based features. In that case, $K = L = 1$ and $d^\mu = d^\nu = 50 \cdot 200$, hence our ansatz would require

at least 56 qubits (cf. Section 3.2 bottom). We prove in Appendix C.1 that the minimal number of required shots N_0 in the right stochastic matrix scales in $\mathcal{O}(n \log n)$ with the number of rows/columns in the OT plan (see Eq. (25)). In this case $N_0 > 161\text{k}$. With less shots the likelihood of empty rows in the matrix is high. Furthermore to obtain a satisfactory sampling error for each entry we need $\mathcal{O}(d^2/\varepsilon^2)$ shots; for a precision $\varepsilon = 0.01$ this is more than a trillion shots. Even abandoning single cell resolution and setting $d^\mu = d^\nu = 50$ still requires 28 qubits, $N_0 > 541$ and $> 25\text{M}$ shots to obtain low error.

As a mitigation strategy, we cluster cells in our experiments into d clusters and compute $\mu, \nu \in \mathbb{R}^d$, i.e., we set $K = L = 1$ and d to 8 (or 16) which requires 16 (or 20) qubits and 640k (2.56M) shots. Note that prior art on neural OT for single-cell data (Bunne et al., 2022; 2023) optimizes over the push-forwarded measure so the OT plans can not be directly accessed, unlike in our quantum method.

4. Experimental Setup




Figure 3: Application overview. A population of cells treated with varying drug dosages, resulting in (X_i, Y_i, p_i) where X_i (Y_i) represent scRNA-seq measurements before (after) a drug administered with dosage $p_i \in [0, 1]$. We cluster the measurements to identify cell types and compute for each batch the distribution of cell types before and after perturbation, i.e., μ and ν . A classical OT solver computes the ground truth OT plan T_i based on μ_i, ν_i (not shown). Given our initial cluster distribution before perturbation μ_i and the dosage p_i our ansatz predicts a transport plan \tilde{T}_i .

We applied our method on predicting changes in the composition of a cell population due to drug perturbations and tested it on synthetic and real data as shown in Figure 3. Starting from a population of heterogeneous cells, each living in a high-dimensional state space we know that administering a drug has a direct effect on the composition of the cell population, by eliminating certain cell types or pushing some other cell types to proliferate. We denote as μ, ν_i the cell type distribution of a cell population *before* and *after* the drug perturbation with a context variable p_i , i.e., the drug dosage $\in [0, 1]$. We measure performance on unseen dosages for different data splitting strategies.

4.1. Cell Type Assignment via Clustering

We represent each cell through a single label, obtained by clustering from the original \mathbb{R}^l space into d clusters (i.e., cell

types). We compute $\mu, \nu \in \mathbb{R}^d$ as distributions over clusters for the population of cells before and after perturbation. We cluster with k -Means and $k = d = 8$ or 16 to adhere with the circuit requirements, i.e., d must be a power of 2 (in general, if $\log_2(k) \notin \mathbb{N}$, we pad and set $d := 2^{\lceil \log_2(k) \rceil}$ and fix the transport plan to be diagonal for the padded entries). We then solve Eq. (8) and compute the OT map between μ and ν with the Sinkhorn solver (Cuturi, 2013), which may be sped up through approximate solvers (Haviv et al., 2024). Repeating this procedure for all dosages yields a dataset $\{T_i, p_i\}_{i=1}^N$ of transport plan-perturbation tuples, processed as described in Section 3.3. The cost $C \in \mathbb{R}^{d \times d}$ is the Euclidean or cosine distance between centroids.

5. Experimental Results

This section verifies that QontOT (Quantum Contextual Optimal Transport) can learn to predict transportation maps *contextualized* through a perturbation variable.

5.1. Synthetic Data

Leveraging the established sc-RNA-seq generator Splatter (Zappia et al., 2017), we devised a perturbation data generator that allows to control the number of generated cells, genes, cell types, perturbation functions and more. We experiment with different perturbation functions to up-/down-regulate gene expression linearly or nonlinearly, different distance metrics, number of clusters and data splits. Even though the perturbation functions are simple and only affect expression of a subset of cells and genes, the induced changes in cell type distribution are significant, locally continuous and nonlinear (cf. Figure A3). Details on the synthetic data generator and the used datasets are in Appendix D.1 and D.2.

We compare QontOT to two baselines, *Average* and *Identity*. *Average* always predicts the same transportation plan, obtained by solving the regularized OT problem (Eq. (8)) on *all* training samples at once, disregarding the context. *Identity* always predicts the identity OT plan, s.t., $\nu = \mu$. The results in Table 1 show that QontOT outperforms both baselines in all cases by a wide margin. The two flavors of QontOT, \mathcal{L}_T and \mathcal{L}_M both have respective advantages; \mathcal{L}_T models are explicitly trained on the transport plan. They provide solutions with lower cost but instead the \mathcal{L}_M models only optimize the marginal distribution and give typically better results in L_2 and R^2 . Unlike related work (Bunne et al., 2022), our method supports various costs like Euclidean or cosine distances of centroids, not just squared Euclidean. The exemplary real and predicted transportation plans in Figure A4A show that QontOT learns context-dependent shifts in cell type frequencies, by capturing the change in the distribution of cluster labels induced by the perturbation. Predicting the effect of stronger perturbations (higher dosages) is more

Dist.	Method	\mathcal{L}	OT Plan		Marginals	
			SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
L_2	Identity	–	1.50	1.41	0.69	0.28
L_2	Average	–	1.07	0.79	0.52	0.27
L_2	QontOT	\mathcal{L}_T	0.97	0.70	0.45	0.56
L_2	QontOT	\mathcal{L}_M	0.97	0.79	0.41	0.55
Cos.	Identity	–	1.67	1.50	0.69	0.29
Cos.	Average	–	1.11	0.82	0.52	0.27
Cos.	QontOT	\mathcal{L}_T	0.97	0.71	0.44	0.59
Cos.	QontOT	\mathcal{L}_M	1.10	0.86	0.40	0.59

(a) Recovering effect of linear perturbation for different distances.

Dist.	Method	\mathcal{L}	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
L_2	Identity	–	1.22	1.19	0.50	0.45
L_2	Average	–	0.97	0.72	0.41	0.42
L_2	QontOT	\mathcal{L}_T	0.86	0.62	0.34	0.47
L_2	QontOT	\mathcal{L}_M	0.97	0.77	0.32	0.48

(b) More realistic and challenging scenario of nonlinear perturbations.

Table 1: Transportation plan prediction. Performance in predicting transportation plans for unseen dosages and linear (a) and nonlinear (b) perturbations; comparing QontOT to two baselines. Different distance metrics were used to derive the cost matrix from the k -means centroids and both linear and non-linear perturbation effects were recovered. SAE denotes sum of absolute errors and Frobenius is the relative Frobenius norm. Means across three simulations are shown, full results in Table A2.

challenging (cf. Figure A4B). This is expected because in the control condition ($p_i = 0$), the cell type distribution remains identical, subject only to stochastic effects in data generation and batch assembly.

Circuit ablations. Next we sought to assess the robustness of QontOT to different configurations: the choice of ansatz, the optimizer, the number of layers and auxiliary qubits and the aggregation scheme to obtain a right stochastic matrix. This time we used synthetic data with four cell groups (rather than one) which simulates more complex tissue and yields richer OT plans. Overall, QontOT is robust to small alterations in circuit structure (cf. Table A4). Even though many settings use a slightly larger computational budget, none of them improve consistently across metrics over the base configuration, validating the imposed inductive biases, e.g., our ansatz type and the "atop" aggregation. For example, replacing our gradient-free optimizer (COBYLA) with another one (Nevergrad (Bennet et al., 2021)) yields identical results. Extending the number of layers in the ansatz or adding more auxiliary qubits was generally found beneficial but also does not always improve performance (more results in Table A3). Moreover, we find that in an OOD (out-of-distribution) setting where we kept the 10% highest dosages out, the test error increases only mildly compared to the training error (cf. Figure A5).

Note that the embedding for transport maps with given initial distribution μ proposed in Section 3.2 can be adapted to classical neural networks. The resulting quantum-inspired

algorithm, which we call NeuCOT, is described in Appendix E and performs explicit optimization of transport plans without relying on Brenier. Unlike QontOT it cannot be applied to the contextual (relaxed) assignment problem of predicting DSMs. While QontOT uses gradient-free optimization, the quantum-inspired approach can be trained conventionally with backpropagation and thus outperforms QontOT on the aforementioned datasets. Gradient-based, quasi-Newton optimization through BFGS substantially improved QontOT’s performance in simulation but it is currently not amenable to quantum hardware. Encouragingly even with gradient-free optimization there are cases where QontOT yields identical or slightly better performance e.g., if the Hamiltonian of the system is known (cf. Table A1).

5.2. SciPlex Data

To facilitate comparison with prior art, we compared QontOT to CellOT (Bunne et al., 2023) and CondOT (Bunne et al., 2022) on two drugs from the SciPlex dataset (Srivatsan et al., 2020) each administered in four dosages. For each of the dosages and the control condition, 20% of cells were randomly held out for validation. Table 2 indicates that

Method	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
Identity	1.10	1.04	0.18	0.47
QontOT- \mathcal{L}_T	0.78	0.61	0.17	0.49
QontOT- \mathcal{L}_M	0.92	0.68	0.16	0.57
CellOT	0.46	0.41	0.17	0.52
CellOT-homo	0.68	0.60	0.29	0.37
CondOT	0.45	0.40	0.18	0.56

(a) Mocetinostat

Method	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
Identity	1.10	1.01	0.28	0.43
QontOT- \mathcal{L}_T	0.82	0.60	0.17	0.49
QontOT- \mathcal{L}_M	0.97	0.82	0.24	0.47
CellOT	0.44	0.40	0.18	0.56
CellOT-homo	0.93	0.80	0.45	0.27
CondOT	0.70	0.55	0.32	0.49

(b) Pracinostat

Table 2: SciPlex comparison. Means across three runs are shown.

CellOT largely yields the best results, note however, that it is an unconditional model and five models were trained (one per condition) inducing an unfair comparison. When aggregating data across conditions (CellOT-homo), performance drops below the level of QontOT- \mathcal{L}_T . CondOT is an evolution of CellOT that leverages partial ICNNs (PICNNs) and can be parameterized by dosage. This yielded overall the best results on transportation plan metrics but on the marginal metrics L_2 and R^2 , QontOT is on par or even superior. Notably, in many applications, such marginal metrics are of higher importance. They can be directly optimized

in QontOT’s \mathcal{L}_M optimization mode which depicts a form of weakly supervised amortized optimization of transport plans that also overcomes the dependence on conventional OT solvers to compute training samples (cf. Section 2.4). The duality of QontOT’s optimization mode is visualized and explained further in Figure A6.

5.3. Contextual Relaxed Assignment Problem

Since our ansatz naturally emits a DSM, it can be applied to the contextual relaxed assignment problem directly. In this task we predict DSMs rather than generic transport plans. In this case, initial and final distributions are fixed, thus we do not need the structures described in Section 3.2. We hypothesized that this task lies closer to the heart of our ansatz (e.g. it does not ignore Q_3 and Q_4) and thus decided to challenge QontOT in a stress test on quantum hardware.

We applied QontOT on synthetic data to contextually predict DSMs and compared it to the classical NeuCOT approach. We used 24 qubits (6 data, 18 auxiliary), a circuit with depth 50 and ~ 70 ECR gates and a dataset of fourty 8×8 DSMs (randomly split into train and test with 20% test data) obtained by sampling from our circuit with random parameters $\in \mathcal{U}(-0.8\pi, 0.8\pi)$. For efficiency, the circuit was split (cf. Appendix G.1) and the smallest physical circuit layout was picked after 50,000 transpilation with the standard Qiskit transpiler. Parameters were optimized for 235 steps over 13 days on a 127-qubit device (IBM Sherbrooke) available through the IBM Quantum Platform. Since the algorithm requires state sampling, no error mitigation was performed and we collected only 8192 shots per iteration. Further details about the setup can be found in Appendix G.




Figure 4: In a 24-qubit hardware experiment, the performance of QontOT surpasses a 18-qubit simulation and various classical neural OT models trained with different hyperparameter settings (cf. Appendix E). The NeuCOT models of size XS, S, M and L optimize respectively 3k, 8k, 81k and 5M parameters with ADAM compared to 124 gradient-free optimized parameters in our ansatz.

The result shows that the objective function was converged well (cf. Figure A8). A few sudden spikes were observed during optimization; those are due to the device recalibration

which was beyond our control. The final result in Figure 4 demonstrates that QontOT achieves a better Frobenius Distance than any of the 15 tested NeuCOT configuration, each with different hyperparameters. Interestingly, despite the device noise, the lack of error mitigation and the limited number of collected shots, the hardware run with a larger qubit budget (24) yielded better results than the 18 qubit simulation. Further performance metrics confirm this result.

6. Discussion

Here, we introduced QontOT, a principled approach to represent transportation maps on quantum computers. We proposed an ansatz for learning to predict OT maps conditioned on a context variable, without requiring access to the cost. Our empirical results on synthetic and real data show that our method learns to predict contextualized transport maps which represent distributional shifts in cell type assignments. While our method does not always match performance of the best classical models on this task, it constitutes, to the best of our knowledge, the first approach to bridge QC, OT and ML. Notably, our approach does not impose constraints on the dimensionality of the context variable(s), thus more complex perturbations such as continuous drug representations, combinatorial genetic perturbations or other covariates could be employed. However, given that the dosage-induced shifts in cluster assignments are also driven by the initial cell states (not only the dosage), future work could devise an ansatz fully parametric for μ_i , potentially through (unbalanced) co-optimal transport (Titouan et al., 2020; Tran et al., 2023). Concurrently, we introduced a new classical neural OT baseline (NeuCOT) which may be further improved through a combination of neural and tensor networks (Wang et al., 2023) or a hybrid quantum-classical tensor network (Schuhmacher et al., 2024).

On the more constrained task of contextual prediction of DSMs we report a compelling finding from a noisy quantum computer, obtained without error mitigation and with gradient-free optimization whereas our classical competitor (NeuCOT) performs worse despite using backpropagation and orders of magnitudes more parameters. Since we only used 24 qubits, the physical circuit of that particular DSM prediction experiment may be reproduced by classical calculations. In that sense, our method may be seen as a novel classical algorithm formulated in the language of quantum, however by applying it to larger sizes it could eventually lead to quantum advantage, assuming that the observed benefits are robust to scaling to more qubits. Overall, our experiments suggest that predicting DSMs rather than generic transport maps is the more promising future endeavour. A natural next step is to apply QontOT in scenarios where DSMs have to be estimated, one potential avenue could be Transformers where DSMs were found to emerge naturally (Sander et al., 2022).

Impact Statement

This paper presents work with the goal to advance the field of quantum machine learning (QML). There are many potential societal consequences of advances in quantum computing, first and foremost in cryptanalysis (see [Scholten et al. \(2024\)](#) for a broad overview). Next, advances in QML will have more specific implications, e.g., they could widen disparities between organizations, countries or researchers that have access to, or can afford, leveraging quantum computing and those that cannot. More specific to our work, the promising result on contextual prediction of DSMs has to be seen in light of the limitations of current quantum hardware, e.g., we can only approximate specific states and structures (such as a DSM) since we have constraints in terms of circuit depth, the number of measurements and have to mitigate the device noise.

Acknowledgements

We thank the anonymous reviewers for their thorough feedback which helped to substantially improve the paper.

References

- Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., and Woerner, S. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409, 2021.
- Amos, B., Xu, L., and Kolter, J. Z. Input convex neural networks. In *International Conference on Machine Learning*, pp. 146–155. PMLR, 2017.
- Anschuetz, E. R. and Gao, X. Arbitrary polynomial separations in trainable quantum machine learning. *arXiv preprint arXiv:2402.08606*, 2024.
- Barenco, A., Bennett, C. H., Cleve, R., DiVincenzo, D. P., Margolus, N., Shor, P., Sleator, T., Smolin, J. A., and Weinfurter, H. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, nov 1995. doi: 10.1103/physreva.52.3457. URL <https://doi.org/10.1103%2Fphysreva.52.3457>.
- Basu, S., Born, J., Bose, A., Capponi, S., Chalkia, D., Chan, T. A., Doga, H., Goldsmith, M., Gujarati, T., Guzman-Saenz, A., et al. Towards quantum-enabled cell-centric therapeutics. *arXiv preprint arXiv:2307.05734*, 2023.
- Bengtsson, I. and Życzkowski, K. *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, 2006.
- Bennet, P., Doerr, C., Moreau, A., Rapin, J., Teytaud, F., and Teytaud, O. Nevergrad: black-box optimization platform. *ACM SIGEVOlution*, 14(1):8–15, 2021.
- Blom, G., Holst, L., and Sandell, D. *Problems and Snapshots from the World of Probability*. Springer Science & Business Media, 1993.
- Bowles, J., Wright, V. J., Farkas, M., Killoran, N., and Schuld, M. Contextuality and inductive bias in quantum machine learning. *arXiv preprint arXiv:2302.01365*, 2023.
- Brenier, Y. Décomposition polaire et réarrangement monotone des champs de vecteurs. *CR Acad. Sci. Paris Sér. I Math.*, 305:805–808, 1987.
- Brualdi, R. A. *Combinatorial Matrix Classes*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2006.
- Bunne, C., Krause, A., and Cuturi, M. Supervised training of conditional monge maps. *Advances in Neural Information Processing Systems*, 35:6859–6872, 2022.
- Bunne, C., Stark, S. G., Gut, G., Del Castillo, J. S., Levesque, M., Lehmann, K.-V., Pelkmans, L., Krause, A., and Rättsch, G. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, pp. 1–10, 2023.
- Cao, K., Gong, Q., Hong, Y., and Wan, L. A unified computational framework for single-cell data integration with optimal transport. *Nature Communications*, 13(1):7419, 2022.
- Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Cuturi, M., Meng-Papaxanthos, L., Tian, Y., Bunne, C., Davis, G., and Teboul, O. Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv preprint arXiv:2201.12324*, 2022.
- Di Meglio, A., Jansen, K., Tavernelli, I., Alexandrou, C., Arunachalam, S., Bauer, C. W., Borrás, K., Carrazza, S., Crippa, A., Croft, V., et al. Quantum computing for high-energy physics: State of the art and challenges. summary of the qc4hep working group. *arXiv preprint arXiv:2307.03236*, 2023.
- Dunkl, C. and Życzkowski, K. Volume of the set of unistochastic matrices of order 3 and the mean jarlskog invariant. *Journal of Mathematical Physics*, 50(12), December 2009. ISSN 1089-7658. doi: 10.1063/1.3272543. URL <http://dx.doi.org/10.1063/1.3272543>.
- Genevay, A., Chizat, L., Bach, F., Cuturi, M., and Peyré, G. Sample complexity of sinkhorn divergences. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1574–1583. PMLR, 2019.

- Harrow, A. W., Hassidim, A., and Lloyd, S. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), 2009.
- Harsanyi, B., Rapsomaniki, M., and Born, J. Learning drug perturbations via conditional map estimators. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*, 2024. URL <https://openreview.net/forum?id=FE71RuwmfI>.
- Haviv, D., Kunes, R. Z., Dougherty, T., Burdziak, C., Nawy, T., Gilbert, A., and Pe'er, D. Wasserstein wormhole: Scalable optimal transport distance with transformers. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://arxiv.org/abs/2404.09411>.
- Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., and Gambetta, J. M. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019. doi: 10.1038/s41586-019-0980-2. URL <https://doi.org/10.1038/s41586-019-0980-2>.
- Horn, R. A. and Johnson, C. R. *Matrix Analysis*. Cambridge University Press, USA, 2nd edition, 2012. ISBN 0521548233.
- Huang, H.-Y., Broughton, M., Mohseni, M., Babbush, R., Boixo, S., Neven, H., and McClean, J. R. Power of data in quantum machine learning. *Nature Communications*, 12(1):2631, 2021.
- Khatri, S., LaRose, R., Poremba, A., Cincio, L., Sornborger, A. T., and Coles, P. J. Quantum-assisted quantum compiling. *Quantum*, 3:140, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Klein, D., Palla, G., Lange, M., Klein, M., Piran, Z., Gander, M., Meng-Papaxanthos, L., Sterr, M., Bastidas-Ponce, A., Tarquis-Medina, M., et al. Mapping cells through time and space with moscot. *bioRxiv*, pp. 2023–05, 2023.
- Korotin, A., Selikhanovych, D., and Burnaev, E. Neural optimal transport. In *The Eleventh International Conference on Learning Representations*, 2023.
- Liu, Y., Arunachalam, S., and Temme, K. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 2021a. ISSN 1745-2481.
- Liu, Y., Arunachalam, S., and Temme, K. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 17(9):1013–1017, 2021b.
- Lotfollahi, M., Wolf, F. A., and Theis, F. J. scgen predicts single-cell perturbation responses. *Nature methods*, 16(8):715–721, 2019.
- Madden, L. and Simonetto, A. Best approximate quantum compiling problems. *ACM Transactions on Quantum Computing*, 3(2):1–29, 2022.
- Makkuva, A., Taghvaei, A., Oh, S., and Lee, J. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pp. 6672–6681. PMLR, 2020.
- Nguyen, B., Nguyen, B., Nguyen, H., and Nguyen, V. A. Generative conditional distributions by neural (entropic) optimal transport. In *Forty-first International Conference on Machine Learning*, 2024.
- Nielsen, M. A. and Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, USA, 10th edition, 2011. ISBN 1107002176.
- Peyré, G., Cuturi, M., et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Powell, M. J. *A direct search optimization method that models the objective and constraint functions by linear interpolation*. Springer, 1994.
- Qiskit contributors. Qiskit: An open-source framework for quantum computing, 2023.
- Rousseeuw, P. J. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- Sander, M. E., Ablin, P., Blondel, M., and Peyré, G. Sinkformers: Transformers with doubly stochastic attention. In *International Conference on Artificial Intelligence and Statistics*, pp. 3515–3530. PMLR, 2022.
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- Scholten, T. L., Williams, C. J., Moody, D., Mosca, M., Hurley, W., Zeng, W. J., Troyer, M., Gambetta, J. M., et al. Assessing the benefits and risks of quantum computers. *arXiv preprint arXiv:2401.16317*, 2024.
- Schuhmacher, J., Ballarin, M., Baiardi, A., Magnifico, G., Tacchino, F., Montangelo, S., and Tavernelli, I. Hybrid tree tensor networks for quantum simulation. *arXiv preprint arXiv:2404.05784*, 2024.

- Sinkhorn, R. and Knopp, P. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21:343–348, 1967.
- Solomon, J. Optimal transport on discrete domains. *AMS Short Course on Discrete Differential Geometry*, 2018.
- Srivatsan, S. R., McFaline-Figueroa, J. L., Ramani, V., Saunders, L., Cao, J., Packer, J., Pliner, H. A., Jackson, D. L., Daza, R. M., Christiansen, L., et al. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473):45–51, 2020.
- Tabak, E. G., Trigila, G., and Zhao, W. Data driven conditional optimal transport. *Machine Learning*, 110:3135–3155, 2021.
- Taşkesen, B., Shafieezadeh-Abadeh, S., Kuhn, D., and Natarajan, K. Discrete optimal transport with independent marginals is #p-hard. *SIAM Journal on Optimization*, 33(2):589–614, 2023. doi: 10.1137/22M1482044. URL <https://doi.org/10.1137/22M1482044>.
- Titouan, V., Redko, I., Flamary, R., and Courty, N. Co-optimal transport. *Advances in neural information processing systems*, 33:17559–17570, 2020.
- Tran, Q. H., Janati, H., Courty, N., Flamary, R., Redko, I., Demetci, P., and Singh, R. Unbalanced co-optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 10006–10016, 2023.
- Uscidda, T. and Cuturi, M. The monge gap: A regularizer to learn all transport maps. In *International Conference on Machine Learning*, volume 202, pp. 34709–34733. PMLR, 23–29 Jul 2023.
- Villani, C. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2008. ISBN 9783540710509. URL https://books.google.ie/books?id=hV8o5R7_5tkC.
- Vinh, N. X., Epps, J., and Bailey, J. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *International conference on machine learning*, pp. 1073–1080, 2009.
- Wang, M., Pan, Y., Xu, Z., Yang, X., Li, G., and Cichocki, A. Tensor networks meet neural networks: A survey and future perspectives. *arXiv preprint arXiv:2302.09019*, 2023.
- Zappia, L., Phipson, B., and Oshlack, A. Splatter: simulation of single-cell rna sequencing data. *Genome biology*, 18(1):174, 2017.

Appendix

A. Relation of DSMs and OT

DSMs are a key structure within optimal transport because

- They are at the root of the Kantorovich relaxation of the assignment problem. In the original assignment problem n entities from the source distribution are assigned bijectively to n entities from the target distribution through a mapping represented by a permutation matrix which is indeed invertible and orthogonal. In the Kantorovich relaxation the assignment takes a probabilistic form rather than a deterministic one (permutation matrices). This probabilistic form is represented through DSMs which are convex combinations of permutation matrices. More details can be found in the seminal textbook by [Peyré et al. \(2019\)](#) (see [Section 3.1](#) and especially [Section 3.2](#)).
- They were fundamental to formulate the entropically regularized version of OT ([Cuturi, 2013](#)). Cuturi’s seminal work kicked off the integration of OT into modern ML. The discovery of the Sinkhorn divergence has been enabled through the Sinkhorn rescaling algorithm which converts nonnegative square matrices into DSMs by alternatively rescaling row and column sums ([Cuturi, 2013](#)). In essence, the prohibitively slow computation of OT (or earth-mover) distances — given by a linear program that requires super cubic runtime ($\mathcal{O}(n^3 \log n)$, see [Genevay et al. \(2019\)](#)) — can be accelerated by an entropic regularisation term that converts the LP to Sinkhorn’s matrix rescaling algorithm.
- DSMs can be rescaled to arbitrary transport maps, as demonstrated in [Section 3.2](#) of our paper. Thus they constitute a fundamental building block and can be leveraged not only for constrained applications where transport maps are exactly DSMs (e.g., in our experiment on quantum hardware on the ”contextual relaxed assignment problem”) but even for cases where the transport maps do not follow a specific structure (as shown in the remaining experiments of our paper).
- DSMs can be linked to unitary operators thus connecting OT and quantum computing. This fundamental observation constitutes the foundation of our paper and is critical to address our specific task with quantum.




Figure A1: Transportation maps. The left and top sequences of blobs represent the initial (μ) and final (ν) distributions. The grid blobs denote the mass displaced from row i to column j . The principle of mass preservation manifests as maintaining the total area of initial and final distribution blobs. The left quadrant shows a diagonal transportation (without displacement) so $\mu = \nu$.

B. The Operator-Schmidt decomposition

We start by defining the structure of the *Schmidt decomposition* ([Bengtsson & Zyczkowski, 2006](#)). Let $|\psi\rangle$ denote a bipartite quantum state on the Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$. We assume the dimensions of \mathcal{H} , \mathcal{H}_1 and \mathcal{H}_2 are the positive integers n , k_1 and k_2 , respectively, with $n = k_1 k_2$. Then the Schmidt decomposition of $|\psi\rangle$ w.r.t. the split $\mathcal{H}_1 \otimes \mathcal{H}_2$ is defined as

$$|\psi\rangle := \sum_{i=1}^{\min\{k_1, k_2\}} \lambda_i |a_i\rangle \otimes |b_i\rangle, \quad (17)$$

where $\{|a_i\rangle\}$ and $\{|b_i\rangle\}$ are bases¹⁷ for respectively \mathcal{H}_1 and \mathcal{H}_2 . The coefficients λ_i , called *Schmidt coefficients*, are real non-negative with $\sum_i \lambda_i^2 = 1$. The Schmidt decomposition can be obtained through the singular value decomposition (SVD). However, in our formulation we never compute such decomposition explicitly, instead, its formulation is used for proving the main results. When a state decomposition is characterized by a single non-zero coefficient $\lambda_1 = 1$, we call it a *product state*. This is linked to a fundamental concept in quantum mechanics called *entanglement*, and the product state represents the absence of it.

¹⁷The bases have cardinality respectively k_1 and k_2 , however we note that the decomposition considers subspaces of dimension up to $\min\{k_1, k_2\}$.

We now extend the decomposition to unitary operators. Let U denote a unitary operator acting on \mathbb{C}^n . Consider the following decomposition

$$\text{vec}_r(U) = \sum_{i=1}^{\min\{k_1^2, k_2^2\}} \lambda_i \text{vec}_r(V_i) \otimes \text{vec}_r(W_i) \quad (18)$$

which can be interpreted as the Schmidt decomposition of the n^2 -dimensional vector $\text{vec}_r(U)$. Since the vectorization operator is an isomorphism, by inverting it we obtain

$$U = \sum_{i=1}^{\min\{k_1^2, k_2^2\}} \lambda_i V_i \otimes W_i \quad (19)$$

which is the definition of Operator-Schmidt decomposition for U w.r.t. the split $\mathcal{H}_1 \otimes \mathcal{H}_2$. Finally, we note that Lemma 2.1 shows that the orthogonality of the vectors $\{\text{vec}_r(V_i)\}$ in (18), corresponds to the orthogonality (w.r.t. the Frobenius inner product) of the operators $\{V_i\}$ in (19). Similarly, the same argument follows for $\{\text{vec}_r(W_i)\}$ and $\{W_i\}$.

C. Quantum formulation

C.1. Recovery of finitely-sampled matrices

In Section 3.1 we obtained a DSM from a quantum circuit by considering the measurements asymptotically. The objective of the present section is that of obtaining a method for estimating the DSM from finite measurements. We define a multiset as the tuple $\langle X, c_X \rangle$, where X denotes the underlying set of elements and c_X a function mapping each element $x \in X$ to its cardinality. Given the density matrix ρ prepared as in Eq. (11a), we run a sampling process which produces a (nonempty) multiset $S = \langle \{(i, j)\}, c \rangle$ of pairs $\{(i, j) | i, j \in [0 \dots d - 1]\}$, where i and j correspond respectively to the row and column indices of the DSM being sampled. The pairs (i, j) are counted using a (non-negative) $d \times d$ matrix F whose entries $F_{i,j}$ correspond to the relative frequency of each pair, that is

$$F_{i,j} = \frac{c((i, j))}{\sum_{i', j'} c((i', j'))}. \quad (20)$$

By previous the definition we note that $\mathbf{1}_d^\top F \mathbf{1}_d = 1$ (i.e. F has total mass 1), and asymptotically the matrix dF (i.e. rescaled to have total mass d) approaches the DSM in Eq. (12b). Assuming the matrix F is non-zero (that is we acquired at least one sample), we define the projector onto the Birkhoff polytope as

$$Q^* = \underset{Q \in \Omega_d}{\text{argmin}} \|Q - dF\|_F^2. \quad (21)$$

Then, by the closedness and convexity of Ω_d the solution Q^* always exists and is unique. We note that when $dF \in \Omega_d$, that is the input matrix is already DSM, then the minimizer becomes $Q^* = dF$. In other words, the projection acts as the identity operator when the input matrix belongs to the Birkhoff polytope.

The remaining part of this section is going to focus on estimating the sampling error for the case of right stochastic matrices (see Section 3.2) obtained from the matrix of relative frequencies F . To solve the problem, we employ the Kullback-Leibler (KL) divergence to quantify the informational difference between probability distributions. We minimize the KL divergence, so

$$\min_{Q \in \mathbb{R}_+^{d \times d}} \sum_{i,j=1}^n Q_{i,j} \ln \left(\frac{Q_{i,j}}{dF_{i,j}} \right), \quad (22a)$$

$$\text{s.t. } Q \mathbf{1}_d = \mathbf{1}_d. \quad (22b)$$

Leveraging Lagrange multipliers, denoted as a_i , the objective function becomes

$$\mathcal{L}(Q) = \sum_{i,j=1}^d \left(Q_{i,j} \ln \left(\frac{Q_{i,j}}{dF_{i,j}} \right) + a_i (Q_{i,j} - 1) \right), \quad (23)$$

$$\frac{\partial \mathcal{L}}{\partial Q_{i,j}} = \ln Q_{i,j} - \ln(dF_{i,j}) + 1 + a_i = 0 \implies Q = d \cdot \text{diag}(\tilde{\mathbf{a}})F,$$

with $\tilde{\mathbf{a}} \in \mathbb{R}_{++}^d$. By imposing the constraint $Q\mathbf{1}_d = \mathbf{1}_d$ and assuming that $F\mathbf{1}_d \in \mathbb{R}_{++}^d$, it follows that $\text{diag}(\tilde{\mathbf{a}}) = \frac{1}{d}\text{diag}(F\mathbf{1})^{-1}$, hence we obtain the minimizer

$$Q^* = d \cdot \text{diag}(\tilde{\mathbf{a}})F = \text{diag}(F\mathbf{1})^{-1}F. \quad (24)$$

To obtain the latter we have assumed that $F\mathbf{1}_d \in \mathbb{R}_{++}^d$, that is the row vectors of F are non-zero. To prevent this scenario, it is essential to determine the minimum number of shots required. This count can be derived from the "Coupon Collector's Problem." (Blom et al., 1993). It indicates that to achieve a satisfactory probability (of obtaining F with nonzero rows) p , we need at least

$$N_0 = d \ln \left(\frac{d}{1-p} \right) \quad (25)$$

samples, and $1-p \ll 1$.

We remark that the minimum number of samples N_0 , relates to the requirement regarding the non-zero rows in the matrix F . However, the latter does not cover the minimum number of shots to obtain a given precision ε for each entry of the resulting matrix. Indeed, for each entry of the matrix Q^* we need $\mathcal{O}(1/\varepsilon^2)$ measurements to obtain precision ε , consequently we require $\mathcal{O}(d^2/\varepsilon^2)$ measurements for the entire matrix.

C.2. The checkerboard ansatz

We propose an ansatz construction which is convenient with respect to the structure of the embedding of transportation maps expanded in Section 3.2. Specifically, since in the partitioning of the DSM in Eq. (13), only the top two quadrants contribute to the resulting right stochastic matrix, we aim at devising an ansatz that does not carry additional information in the discarded quadrants. The latter could also be interpreted as making the parametrisation for the ansatz more efficient.

Let σ_i with $i = [1 \dots 3]$ be the Pauli operators commonly denoted with σ_x, σ_y and σ_z , respectively. Also we define $\sigma_0 = \mathbb{I}_2$. The subscript of the σ to determine the Pauli will be indicated interchangeably as symbol or integer index.

For some positive integer k , we define the subset G_k of unitary operators as

$$G_k := \{U \in \mathfrak{U}(2k) \mid U = \mathbb{I}_2 \otimes A + \sigma_x \otimes B\}, \quad (26)$$

where A, B are $k \times k$ matrices, not necessarily unitary. In other words, the operators in $U \in G_k$ have the following block matrix form

$$U = \begin{pmatrix} A & B \\ B & A \end{pmatrix}, \quad (27)$$

which is clearly inherited by the corresponding unistochastic

$$U \odot \bar{U} = \begin{pmatrix} A \odot \bar{A} & B \odot \bar{B} \\ B \odot \bar{B} & A \odot \bar{A} \end{pmatrix} = \begin{pmatrix} Q_1 & Q_2 \\ Q_2 & Q_1 \end{pmatrix}. \quad (28)$$

We now proceed with revealing the group-theoretical structure of the set G_k and also its relation with the tensor product, hence we obtain the construction of the ansatz implementing the unitary in Eq. (27).

The next lemma shows that the set G_k is a subgroup of even degree of the unitary group.

Lemma C.1. *The set G_k is non-empty and endowed with a group structure under operator composition, for all positive integers k .*

Proof. It is immediately verifiable that $\mathbb{I}_{2k} \in G_k$, that is the set G_k is non-empty and it contains the identity element w.r.t. matrix multiplication. Also the composition of operators carries the associativity as required. Finally we verify the closure. Let $U_1, U_2 \in G_k$ such that $U_i = \mathbb{I}_2 \otimes A_i + \sigma_x \otimes B_i$ for $i = 1, 2$, then

$$U_1 U_2 = (\mathbb{I}_2 \otimes A_1 + \sigma_x \otimes B_1) (\mathbb{I}_2 \otimes A_2 + \sigma_x \otimes B_2) \quad (29a)$$

$$= \mathbb{I}_2 \otimes (A_1 A_2 + B_1 B_2) + \sigma_x \otimes (A_1 B_2 + B_1 A_2), \quad (29b)$$

which corresponds to the pattern in (27). Hence $U_1 U_2 \in G_k$. \square

The result that follows shows that the structure is preserved under the tensor product.

Lemma C.2. *Let $U_1 \in G_{k_1}$ and $U_2 \in G_{k_2}$, for some positive integers k_1 and k_2 . Then $U_1 \otimes U_2 \in G_{2k_1k_2}$.*

Proof. Let $U_1 \in G_{k_1}$ and $U_2 \in G_{k_2}$ such that $U_i = \mathbb{I}_2 \otimes A_i + \sigma_x \otimes B_i$ for $i = 1, 2$, then

$$U_1 \otimes U_2 = (\mathbb{I}_2 \otimes A_1 + \sigma_x \otimes B_1) \otimes (\mathbb{I}_2 \otimes A_2 + \sigma_x \otimes B_2) \quad (30a)$$

$$= \mathbb{I}_2 \otimes A + \sigma_x \otimes B, \quad (30b)$$

with $A = A_1 \otimes (\mathbb{I}_2 \otimes A_2 + \sigma_x \otimes B_2)$ and $B = B_1 \otimes (\mathbb{I}_2 \otimes A_2 + \sigma_x \otimes B_2)$. Hence it follows that $U_1 \otimes U_2$ fulfils the pattern in (27) and since A, B are linear maps in $\mathbb{C}^{2k_1k_2}$, then $U_1 \otimes U_2 \in G_{2k_1k_2}$. \square

C.2.1. ANSATZ'S TWO-QUBIT GENERATOR

We obtain a two-qubit circuit $U_g \in G_2$, that by Lemma C.1 and C.2 can be used as a generator for the more general G_{2k} with $k \geq 1$. From the definition in Eq. (26) we obtain the symmetry $U \in G_2 \implies (\sigma_x \otimes \mathbb{I}_2)U(\sigma_x \otimes \mathbb{I}_2) = U$. Using the latter and the general unitary circuit with 2 CNOTs (highlighted) (Barenco et al., 1995), we solve the following circuit equation

$$= \mathbb{I}_2^{\otimes 2}, \quad (31)$$

where A, B, C, D are arbitrary single qubit (special) unitaries and $\alpha, \beta \in \mathbb{R}$. We obtain a solution to the equation. Since the operator $\sigma_z \otimes \mathbb{I}_2$ commutes with the CNOT gate (with the Pauli σ_z acting on CNOT's controlling qubit), we impose on Eq. (31) the conditions

$$B^\dagger \sigma_x B = \sigma_z, \quad (32a)$$

$$R_y(-\beta) \sigma_z R_y(\beta) = \sigma_z, \quad (32b)$$

$$D^\dagger \sigma_z D = \sigma_x. \quad (32c)$$

Then a solution is $B = D = H$ and $\beta = 0$, where H is the Hadamard operator on a single qubit. Hence the generator circuit takes the following form

$$\quad (33)$$

where $A, C \in SU(2)$ and $\alpha \in \mathbb{R}$.

Finally, by using Lemma C.1 and C.2, and the generator block in Eq. (33), we construct the ansatz as exemplified in Figure A2.




Figure A2: An example of depth and connectivity efficient (single) layer for the Checkerboard ansatz. Here the blocks K correspond to the 2-qubits circuit in Eq. (33) and the vectors θ_i are seven dimensional vectors parameterizing gates A, C and R_z of Eq. (33).

C.3. Proofs

Proof of Lemma 2.1. Considering the constraint $\text{Tr} \left(U_i U_j^\dagger \right) = n \delta_{ij}$ we obtain

$$\text{vec}_r(U_j)^\dagger \text{vec}_r(U_i) = \sum_k \left(\langle k | U_j^\dagger \otimes \langle k | \right) \sum_t (U_i |t\rangle \otimes |t\rangle) \quad (34a)$$

$$= \sum_{k,t} \left(\langle k | U_j^\dagger U_i |t\rangle \otimes \langle k |t\rangle \right) \quad (34b)$$

$$= \sum_k \left(\langle k | U_j^\dagger U_i |k\rangle \right) = \text{Tr} \left(U_i U_j^\dagger \right) = n \delta_{ij}. \quad (34c)$$

□

Proof of Lemma 3.1. We expand the function $p : [0 \dots d-1] \times [0 \dots d-1] \rightarrow [0, 1]$ defined in (12a), so

$$p(i, j) := 2^n \text{Tr}(\rho |ij\rangle \langle ij|) \quad (35a)$$

$$= \sum_k \lambda_k^2 \langle ij | \text{vec}_r(W_k) \text{vec}_r(W_k)^\dagger |ij\rangle \quad (35b)$$

$$= \sum_k \lambda_k^2 \langle i | (W_k \odot \overline{W_k}) |j\rangle. \quad (35c)$$

The positivity of the entries of the DSM is clear from the definition of $p(i, j)$. We prove the rows sum constraint for $Q = \sum_{i,j} p(i, j) |i\rangle \langle j|$, that is

$$Q \mathbf{1}_d = \sum_{i,j} p(i, j) |i\rangle \quad (36a)$$

$$= \sum_i |i\rangle \cdot \left(\sum_k \lambda_k^2 \langle i | \sum_j (W_k \odot \overline{W_k}) |j\rangle \right), \quad (36b)$$

where the rightmost sum equals the vector $\mathbf{1}_m$ since $W_k \odot \overline{W_k}$ is unistochastic, also $\sum_k \lambda_k^2 = 1$ (following from (9)), hence $Q \mathbf{1}_d = \mathbf{1}_d$. Similarly the same holds for the columns sum constraint, hence the claim follows. □

D. Experimental details

D.1. Synthetic dosage perturbation data generator

We leverage the established single-cell RNA sequencing generator `Splatter` (Zappia et al., 2017) to form a three-stage generator for drug dosage perturbation datasets:

1. First, `Splatter` samples raw expression counts ($X \in \mathbb{R}^{n_1 \times l}$, with n_1 cells and l genes) from zero-inflated negative binomial distributions (one per gene). Sufficient statistics of all underlying distributions (Poisson, Gamma, Chi-Square) can be controlled.
2. We aim to produce a tuple of (X_i, Y_i, \mathbf{p}_i) where X_i holds unperturbed base states of n_1 cells and $Y_i \in \mathbb{R}^{n_2 \times l}$ holds perturbed states of n_2 cells, resulting from a drug perturbation administered with dosage $\mathbf{p}_i \in [0, 1]$. To derive the perturbed states Y_i , new base states \tilde{Y}_i are sampled with the same configuration used to generate X_i , mimicking that cells are being destroyed during measurement. Subsequently $Y_i = g(\tilde{Y}_i, \mathbf{p}_i)$ where g is the total effect on the cells, governed by a combination of noise terms and the immediate effect $g_p(\cdot)$ of the perturbation. We assume that only 15% of the genes alter their expression upon perturbation. In this case, we apply g_p to the raw cell states, scaled by a response amplitude $\sim \mathcal{U}(0.3, 1)$. Moreover, 10% of the cells are generally unresponsive to the perturbation ($g_p = 0$). We investigate linear and non-linear perturbations, i.e., $g_{p_1}(x) = ax + b$ and a reciprocal root function $g_{p_2}(x) = ax^{-b}$ with $a, b > 0$, respectively. The hyperparameters of the experiments can be found in Appendix D.2.
3. We repeat stage 2 for each dosage by varying smoothly the immediate effect $g_p(\cdot)$ based on \mathbf{p}_i , resulting in a dataset $\{X_i, Y_i, \mathbf{p}_i\}_{i=1}^N$ of N tuples. Responsive genes are fixed across samples. The base states X_i are *identical* across all samples of the dataset, mimicking the common experimental setting where only one control population was measured (Srivatsan et al., 2020).




Figure A3: Exemplary cell type distributions for source and target cell populations. The distribution of cell types in the unperturbed tissue is either entirely static (see left) or varies mildly based on the user-defined noise level (not shown). Instead, the perturbed cells produce highly dissimilar distributions that, however, exhibit some local continuity for similar dosages. The plot was generated with data from linear perturbations and euclidean cost.

D.2. Datasets and hyperparameters

D.2.1. SYNTHETIC DRUG DOSAGE PERTURBATION DATA

Initial experiments For the results shown in Table 1 and Figure A4, we simulate 300 genes and 1000 cells across 50 unique dosages, equidistantly spaced in $[0, 1]$. 15% of the genes respond to the perturbation $g_p(\cdot)$ but 10% of the cells are set as unresponsive. The sinkhorn regularization $\gamma = 0.001$. For the linear case, $f_{p_1}(x) = 3x + 1$ and for the non-linear case $f_{p_2}(x) = 100x^{-0.2}$. For each dosage, four batches of 500 cells each were created, summing to 200 samples which were split randomly with 20% held out dosages for testing. In almost all experiments we set the number of clusters to $k = d = 8$; only in Figure A4 it was set to 16.

Four cell groups. For the ablation studies on circuit structure (cf. Table A4), we simulate populations of 100 genes and 2000 cells, each belonging now to one out of four groups to simulate more complex tissue. We perturb the populations with the function f_{p_2} (as above) and 100 dosages, equidistantly spaced in $[0, 1]$. The initial cell states X_i are resampled for every dosage and only 2% of cells are set as unresponsive. The 100 dosages are split randomly with 20% test data.

D.2.2. SCIPLEX DATA

Selected models were trained on two of the nine compounds from the SciPlex dataset (Srivatsan et al., 2020). High Throughput Screens on three cell lines were conducted with four varying concentrations (10, 100, 1000, and 10000 nM) for each drug. We inherit preprocessing from Bunne et al. (2023) and Lotfollahi et al. (2019) which includes library size normalization, filtering at the cell and gene levels, and \log_{1p} transformation. For mocetinostat and pracinostat we obtained respectively 22,154 and 21,926 cells from which 17,565 and 15,137 were control cells. Data was split per condition (control + four dosages) in a roughly 80/20 ratio. Preprocessing identified 1,000 highly-variable genes, which were compressed with PCA to obtain 50-dimensional latent codes which are clustered with K -Means into 8 clusters.

CellOT & CondOT. CellOT and CondOT are trained with the `ott-jax` package (Cuturi et al., 2022) for 1000 iterations and batches of size 50 on μ_i , i.e., the same 8-dimensional feature vectors (denoting a distribution of cell types over 50 cells) used to train QontOT. We use a cosine decay learning rate scheduler with an initial value of 0.001 and an alpha of 0.01, optimized with ADAM (Kingma & Ba, 2014). CondOT uses the gaussian map initialization proposed in (Bunne et al., 2022). As CellOT and CondOT learn directly a map $f : \mathbb{R}^8 \rightarrow \mathbb{R}^8$ such that $f(\mu_i) = \bar{\nu}_i$, we use the entropically regularized sinkhorn solver (Cuturi, 2013) on (μ_i, ν_i) to obtain the transport maps and compute performance metrics on SAE and relative Frobenius norm.

D.3. Implementation

As mentioned in Section 3, the fact that $U \odot \bar{U} \in \Omega_n$ is a DSM offers great flexibility in the choice of the ansatz. In practice, we implemented two ansätze, *centrosymmetric* and *simple*. Both of them have been trained in Qiskit 0.43.0 (Qiskit contributors, 2023) and all experiments were performed with Qiskit’s sampler class and, unless indicated otherwise (cf. Section 5.3), in statevector simulation.

Centrosymmetric The centrosymmetric ansatz is our default implementation which induces a bias toward properties of centrosymmetric matrices. Specifically, the matrix being modelled can be divided into four quadrants such that the respective diagonals are equal and the respective off-diagonals are also equal.

Simple Instead, the simple ansatz instead is symmetric by construction and has less bias toward a specific class of unitary operators than the centrosymmetric ansatz. This ansatz was first formulated in [Khatri et al. \(2019\)](#) and later refined in [Madden & Simonetto \(2022\)](#). Note that this ansatz implements the identity operator when all parameters are zero.

E. Neural contextual OT (NeuCOT)

E.1. Methodology

In Section 3.2 we have shown that, if we produce a right stochastic matrix $\hat{T} \in \mathcal{N}(\mathbf{1}_d, \nu')$, the latter can be rescaled to a transportation map with the required initial distribution μ . We recall that, at inference time, the right stochastic \hat{T} is predicted by a circuit depending on the perturbation, thus the predicted transportation map results from the rescaling of the rows of \hat{T} , using the elements of the initial distribution μ as coefficients. A complementary case is that of the contextual (relaxed) assignment problem, whose origin is presented in Section 2.3. The latter requires the prediction of a DSM, and we believe that this task is hard for classical machine learning. The former case instead, that is the one requiring a right stochastic matrix, can be shown to be practical for classical ML. Let $M \in \mathbb{R}^{d \times d}$ and $[e^{M_{i,j}}]$ the *Hadamard exponential matrix* ([Horn & Johnson, 2012](#)) of M , that is the matrix resulting from the entry-wise application of the exponential mapping. We extend the notation to a row-rescaled form, so

$$M \mapsto \left[\frac{e^{M_{i,j}}}{\sum_k e^{M_{i,k}}} \right], \quad (37)$$

then

$$\left[\frac{e^{M_{i,j}}}{\sum_k e^{M_{i,k}}} \right] \mathbf{1}_d = \left(\frac{\sum_k e^{M_{1,k}}}{\sum_k e^{M_{1,k}}} \quad \dots \quad \frac{\sum_k e^{M_{d,k}}}{\sum_k e^{M_{d,k}}} \right)^\top = \mathbf{1}_d, \quad (38)$$

that is (37) maps any $M \in \mathbb{R}^{d \times d}$ to a right stochastic matrix. Let the matrix-valued function $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{d \times d}$ represent a neural network parametrised by the vector θ , mapping the space of perturbations \mathcal{X} to a $d \times d$ real matrix. We obtain a right stochastic matrix as a function of the perturbation $\mathbf{p} \in \mathcal{X}$, so

$$\hat{T} = \left[\frac{e^{(f_\theta(\mathbf{p}))_{i,j}}}{\sum_k e^{(f_\theta(\mathbf{p}))_{i,k}}} \right] \in \mathcal{N}(\mathbf{1}_d, \nu'). \quad (39)$$

Hence, the right stochastic matrix can be obtained through a neural network with a softmax activation in the ultimate layer. Finally the procedure continues as outlined in Section 3.2, that is the prediction for the transportation map is given by $\bar{T} = D_\mu \hat{T}$, so $\bar{T} \mathbf{1}_d = \mu$ as required. Optionally, one can make the prediction depending non-linearly on the initial distribution μ by re-defining f_θ as a function of both the perturbation and the distribution μ .

This is a novel, quantum-inspired approach that combines neural and contextual OT through amortized optimization. It can be trained in a fully or weakly supervised setting, either optimizing OT plans directly (i.e., $\mathcal{L}(\bar{T}, T)$) or only the push-forwarded distribution (i.e., $\mathcal{L}(\hat{T} \# \mu, \nu)$). We dub this approach *NeuCOT* for Neural Contextual Optimal Transport. Previous approaches either leverage Brenier’s theorem ([1987](#)) to recast the problem to convex regression (e.g., CellOT and CondOT ([Bunne et al., 2022; 2023](#))) or use regularization ([Uscidda & Cuturi, 2023](#)). In the implementation the optimization occurs over the push-forwarded measure in both cases, so unlike in our method, the OT plans can not be directly accessed.

E.2. Implementation and Result

In practice, we implement this approach with a shallow, dense neural network of two layers (64 and 128 units unless indicated otherwise), a ReLU activation, a dropout (40%, unless indicated otherwise) and use a MSE loss between real and predicted OT plans. We apply an optional residual connection of the context to the last layer of the network. Moreover, we also apply an optional "DSM loss" to penalize deviations of the marginals from uniform ones ($\mathbf{1}_n$). The models have $\sim 25k$ trainable parameters that are optimized with ADAM for 500 epochs with a learning rate of $5e-4$.

The results in [Table A1](#) compare NeuCOT to QontOT and the average baseline for different datasets and splitting strategies.

Dataset	Data split	Method	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
Four cell types	Random	Average	1.23	0.83	0.15	0.56
		QontOT	0.41	0.30	0.13	0.54
		NeuCOT	0.22	0.20	0.10	0.73
Four cell types	Extrapolation	Average	1.21	0.84	0.24	0.09
		QontOT	0.53	0.78	0.20	0.21
		NeuCOT	0.32	0.30	0.17	0.39
Known Hamiltonian 8×8 plans	Random	Average	2.332	0.883	0.538	0.880
		QontOT- \mathcal{L}_T	0.571	0.208	0.085	0.892
		NeuCOT	0.581	0.215	0.089	0.894
Known Hamiltonian 16×16 plans	Random	Average	1.317	0.879	0.055	0.671
		Identity	0.878	1.042	0.078	0.640
		QontOT- \mathcal{L}_T	0.421	0.833	0.028	0.710
		QontOT- \mathcal{L}_M	0.479	0.845	0.029	0.699
		NeuCOT	0.400	0.356	0.028	0.708

Table A1: Comparison of QontOT to NeuCOT. Means across three different runs are shown. All datasets use actual distributions, so $\mu \mathbf{1}_n = \mathbf{1} = \nu \mathbf{1}_n$.

F. Extended Results

Perturb.	Dist.	Layers	\mathcal{L}	Transportation plan		Marginals	
				SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
Lin.	L_2	Ident.	–	$1.50_{\pm 0.14}$	$1.41_{\pm 0.12}$	$0.69_{\pm 0.07}$	$0.28_{\pm 0.07}$
Lin.	L_2	Avg.	–	$1.07_{\pm 0.04}$	$0.79_{\pm 0.01}$	$0.52_{\pm 0.04}$	$0.27_{\pm 0.06}$
Lin.	L_2	QontOT	\mathcal{L}_T	$0.97_{\pm 0.08}$	$0.70_{\pm 0.06}$	$0.45_{\pm 0.04}$	$0.56_{\pm 0.03}$
Lin.	L_2	QontOT	\mathcal{L}_M	$0.97_{\pm 0.12}$	$0.79_{\pm 0.15}$	$0.41_{\pm 0.05}$	$0.55_{\pm 0.10}$
Lin.	Cos.	Ident.	–	$1.67_{\pm 0.16}$	$1.50_{\pm 0.13}$	$0.69_{\pm 0.07}$	$0.29_{\pm 0.07}$
Lin.	Cos.	Avg.	–	$1.11_{\pm 0.05}$	$0.82_{\pm 0.02}$	$0.52_{\pm 0.05}$	$0.27_{\pm 0.06}$
Lin.	Cos.	QontOT	\mathcal{L}_T	$0.97_{\pm 0.09}$	$0.71_{\pm 0.07}$	$0.44_{\pm 0.05}$	$0.59_{\pm 0.02}$
Lin.	Cos.	QontOT	\mathcal{L}_M	$1.10_{\pm 0.14}$	$0.86_{\pm 0.10}$	$0.40_{\pm 0.06}$	$0.59_{\pm 0.10}$
NonLin.	L_2	Ident.	–	$1.22_{\pm 0.05}$	$1.19_{\pm 0.05}$	$0.50_{\pm 0.02}$	$0.45_{\pm 0.02}$
NonLin.	L_2	Avg.	–	$0.97_{\pm 0.02}$	$0.72_{\pm 0.01}$	$0.41_{\pm 0.01}$	$0.42_{\pm 0.04}$
NonLin.	L_2	QontOT	\mathcal{L}_T	$0.86_{\pm 0.03}$	$0.62_{\pm 0.02}$	$0.34_{\pm 0.01}$	$0.47_{\pm 0.05}$
NonLin.	L_2	QontOT	\mathcal{L}_M	$0.97_{\pm 0.06}$	$0.77_{\pm 0.05}$	$0.32_{\pm 0.01}$	$0.48_{\pm 0.01}$

Table A2: Transportation plan prediction. Extended results for Table 1, \pm denotes standard deviation across three runs.

Dist.	Layers	\mathcal{L}	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
L_2	6	\mathcal{L}_T	$0.97_{\pm 0.08}$	$0.70_{\pm 0.06}$	$0.45_{\pm 0.04}$	$0.56_{\pm 0.03}$
L_2	12	\mathcal{L}_T	$0.98_{\pm 0.11}$	$0.72_{\pm 0.08}$	$0.45_{\pm 0.05}$	$0.53_{\pm 0.10}$
L_2	6	\mathcal{L}_M	$0.97_{\pm 0.12}$	$0.79_{\pm 0.15}$	$0.41_{\pm 0.05}$	$0.55_{\pm 0.10}$
L_2	12	\mathcal{L}_M	$1.02_{\pm 0.09}$	$0.81_{\pm 0.13}$	$0.42_{\pm 0.04}$	$0.54_{\pm 0.08}$
Cos.	6	\mathcal{L}_T	$0.97_{\pm 0.09}$	$0.71_{\pm 0.07}$	$0.44_{\pm 0.05}$	$0.59_{\pm 0.02}$
Cos.	12	\mathcal{L}_T	$0.99_{\pm 0.10}$	$0.71_{\pm 0.06}$	$0.45_{\pm 0.05}$	$0.59_{\pm 0.17}$
Cos.	6	\mathcal{L}_M	$1.10_{\pm 0.14}$	$0.86_{\pm 0.10}$	$0.40_{\pm 0.06}$	$0.59_{\pm 0.10}$
Cos.	12	\mathcal{L}_M	$1.06_{\pm 0.05}$	$0.83_{\pm 0.12}$	$0.42_{\pm 0.05}$	$0.60_{\pm 0.12}$

Table A3: Ablation study on number of layers in ansatz. Adding more layers in the ansatz and thus more parameters in the circuit does not improve performance. Experiment performed on linear perturbation function (cf. Table 1). For 6 respectively 12 layers in the ansatz, there are 234 respectively 456 circuit parameters to optimize.




Figure A4: Capturing variation in cell type distributions. a) Three predicted transportation plans from the nonlinear perturbation dataset are shown next to their unseen ground truth. b) Frobenius distance of real and predicted transportation across unseen dosages are shown for QondOT and the baseline.




Figure A5: Out-of-distribution scenario. When QontOT is evaluated on dosages outside the training data scenario (redly shaded background), the performance decreases but still remains well above the baseline.

Method	SAE (\downarrow)	Frob. (\downarrow)	L_2 (\downarrow)	R^2 (\uparrow)
Identity	0.65	0.60	0.18	0.46
Average	1.23	0.83	0.15	0.56
\mathcal{L}_M	0.59	0.43	0.13	0.58
\mathcal{L}_T -Nevergrad	0.41	0.30	0.14	0.54
\mathcal{L}_T	0.41	0.30	0.13	0.54
\mathcal{L}_T -AsIs	0.45	0.34	0.14	0.54
\mathcal{L}_T -Simple	0.46	0.32	0.13	0.57
\mathcal{L}_T -Simple-12	0.55	0.40	0.13	0.58
\mathcal{L}_T -Simple-12-Shared	0.41	0.30	0.14	0.57
\mathcal{L}_T -Simple-16-Shared	0.41	0.30	0.14	0.58

Table A4: Ablation studies on circuit structure. \mathcal{L}_T is the base configuration. *Simple* is an circuit constructions alternative to the base type "centrosymmetric" (cf. Appendix D.3 for details). "12" or "16" refers to the number of layers (base is 6) and "Shared" defines whether the ansatz parameters are identical across layers (this is faster to optimize thus allowing deeper circuits). "AsIs" denotes an alternative to the "atop" aggregation to produce a DSM (cf. Section 3.2). Means across three random splits are shown.




Figure A6: Comparison of QontOT optimization modes. **A)** Sankey plots of $\mathcal{L}_{Marginal}$ and $\mathcal{L}_{Transport}$ for a validation sample of pracinostat reveal that \mathcal{L}_M learns transport maps with higher transport cost and unnecessary move of mass (see bucket 4 and 8). **B)** In contrast, the marginal performance of \mathcal{L}_M is typically better, evidenced by the more discriminative cell type identification for the UMaps. The higher AMI (adjusted mutual information (Vinh et al., 2009)) shows that the predicted clusters are more similar to the real ones and the relative silhouette score (Rousseeuw, 1987) shows that the clusters are more consistent for \mathcal{L}_M than for \mathcal{L}_T .

G. Hardware Simulation

G.1. Adhoc circuit depth optimization

The structure of the circuit in Figure 2a can be optimised in terms of circuit depth when the unitary U_p can be factorized as the product of two unitaries. In our case, the latter condition is always possible since the unitary is constructed from elementary gates. We sketch the mechanism and its proof. Assume $U_p = U_p^{(2)}U_p^{(1)}$ (where the circuit depths of the two factors are assumed equal), and assume the state in (10) is generated by $U_p^{(1)}$, so

$$\begin{aligned} |\varphi^{(1)}\rangle &= \left(\mathbb{I}_2^{\otimes m} \otimes U_p^{(1)} \otimes \mathbb{I}_2^{\otimes n} \right) \cdot (|b_m\rangle \otimes |b_n\rangle) \\ &\stackrel{(9)}{=} \sum_k \lambda_k \left(\mathbb{I}_2^{\otimes m} \otimes V_k \right) |b_m\rangle \otimes \left(W_k \otimes \mathbb{I}_2^{\otimes n} \right) |b_n\rangle \\ &\stackrel{(6a)}{=} \sum_k \lambda_k \frac{\text{vec}_r(V_k^\top)}{\sqrt{2^m}} \otimes \frac{\text{vec}_r(W_k)}{\sqrt{2^n}}. \end{aligned} \quad (40)$$

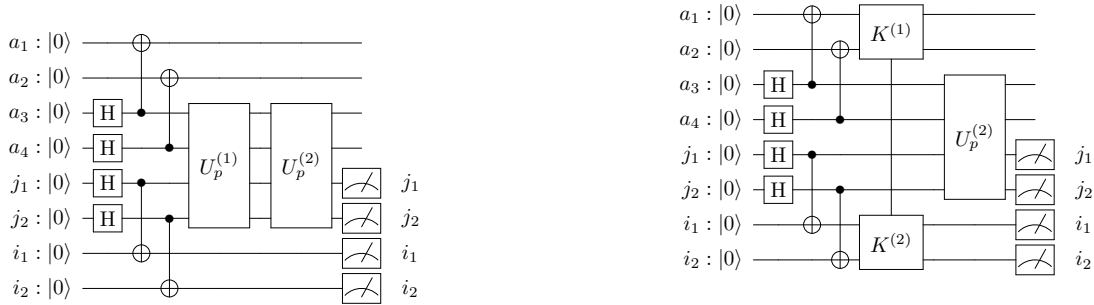
Note that the operators V_k and W_k here are not the same as those in (10), however we keep the same symbols for simplicity. Using the identities for vectorization (6a) and (6b), we rewrite the middle expression as

$$|\varphi^{(1)}\rangle = \sum_k \lambda_k \left(V_k^\top \otimes \mathbb{I}_2^{\otimes m} \right) |b_m\rangle \otimes \left(\mathbb{I}_2^{\otimes n} \otimes W_k^\top \right) |b_n\rangle, \quad (41)$$

Which corresponds to transposing¹⁸ $U_p^{(1)}$ and changing the subsystems on which it acts. We call the new operator K . Subsequently, the operator $U_p^{(2)}$, which now commutes with K , is applied to the state $|\varphi^{(1)}\rangle$ with the resulting halving of the circuit depth (under the assumption stated above). The final state is equivalent to the original one, that is

$$|\varphi\rangle = \left(\mathbb{I}_2^{\otimes m} \otimes U_p^{(2)} \otimes \mathbb{I}_2^{\otimes n} \right) |\varphi^{(1)}\rangle. \quad (42)$$

We illustrate the decomposition and the resulting circuit in Figure A7.



(a) The decomposition of the circuit for $U_p(\mathbf{p}; \boldsymbol{\theta}) = U_p^{(2)}U_p^{(1)}$, where on the right-hand side we omit the parameters for clarity.

(b) The equivalent formulation of $U_p(\mathbf{p}; \boldsymbol{\theta})$ where $U_p^{(1)}$ is substituted by the operator K which commutes with $U_p^{(2)}$. Consequently, the circuit depth of the overall gate U_p is halved, while the total number of qubits remains the same.

Figure A7: Illustration of the adhoc compilation trick that halves the circuit depth for the unitary U_p . On the LHS we have the decomposition on the unitary U_p into two factors, whereas on the RHS we represent pictorially the effect of the commutativity of the factors K and $U_p^{(2)}$.

G.2. Experimental details

During the simulation, ansatz parameters had been optimized for 235 iterations on a 127-qubit device (IBM Sherbrooke) available through the IBM Quantum Platform. In principle, on modern hardware every iteration can be accomplished in seconds. However, due to a very high demand on (shared) quantum computers, the actual execution time can take days.

¹⁸By transposition of a unitary operator we mean the complex conjugate of its adjoint, that is \bar{U}^\dagger .

Result of a quantum computation is not directly observable because unlike classical bits that always take 0 or 1, qubits can exist in a superposition of states. To extract a result, one needs to measure the qubits, and this single act of measurement is called a “shot”. Upon measurement, the superposition state collapses to one of the basis states (vectors of 0s and 1s) with a certain probability.

We did 8192 shots per iteration and every obtained state of 0s and 1s was used to increment a counter of occurrences of a corresponding entry in measured DSM, just as described in the main body of this paper. A typical picture of objective function convergence profile is shown on [Figure A8](#).

One can notice sudden spikes as the curve approaches a plateau. This happens primarily because a quantum hardware should be recalibrated from time to time. To cope with this behaviour we track the best solution (with the smallest objective function value) and report the best one rather than the last solution obtained by an optimizer. In our simulation, the best parameters had been obtained on iteration 193.

No error mitigation was performed. It is still an open problem in quantum computing community how to suppress noise while doing state sampling. There are a number of established techniques for error mitigation when the output is a single scalar value measured on some observable. However, in this study we are interested in measuring probabilities of individual states.




Figure A8: Convergence profile of the objective function while running on IBM Sherbrooke quantum device.