IEEE Transactions on Quantum Engineering

Received 1 February 2022; revised 17 April 2022; accepted 3 May 2022; date of publication 23 May 2022; date of current version 28 October 2022.

Digital Object Identifier 10.1109/TQE.2022.3176870

Model-Predictive Quantum Control via Hamiltonian Learning

MAISON CLOUÂTRÉ[1,2] (Student Member, IEEE), MOHAMMAD JAVAD KHOJASTEH[1] (Member, IEEE), AND MOE Z. WIN[1,3] (Fellow, IEEE)

[1]Wireless Information and Network Sciences Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA
[2]Department of Electrical and Computer Engineering, Mercer University, Macon, GA 31207 USA
[3]Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Corresponding author: Moe Z. Win (e-mail: moewin@mit.edu).

The fundamental research described in this article was supported in part by the National Science Foundation under Grant CCF-1956211 and in part by the MIT Summer Research Program (MSRP). An earlier version of this article was presented in part at the 36th Annual MSRP Research Forum, Cambridge, MA, USA, Aug. 2021.

ABSTRACT This article proposes an end-to-end framework for the learning-enabled control of closed quantum systems. The proposed learning technique is the first of its kind to utilize a hierarchical design, which layers probing control, quantum state tomography, quantum process tomography, and Hamiltonian learning to identify both the internal and control Hamiltonians. Within this context, a novel quantum process tomography algorithm is presented, which involves optimization on the unitary group, i.e., the space of unitary operators, to ensure physically meaningful predictions. Our scalable Hamiltonian learning algorithms have low memory requirements and tunable computational complexity. Once the Hamiltonians are learned, we formalize data-driven model-predictive quantum control (MPQC). This technique utilizes the learned model to compute quantum control parameters in a closed-loop simulation. Then, the optimized control input is given to a physical quantum system in an open-loop fashion. Simulations show model-predictive quantum control to be more efficient than the current state-of-the-art, quantum optimal control, when sequential quadratic programming (SQP) is used to solve each control problem.

INDEX TERMS Quantum computing, quantum consensus, quantum control, quantum Hamiltonian learning, quantum networks, quantum process tomography.

I. INTRODUCTION

Quantum information science is a rapidly growing field that seeks to utilize quantum dynamical systems to perform sensing, communication, or computation [1]–[6]. The success of many quantum devices, such as those employing quantum bits (qubits), is dependent on the precise control of their dynamics [7]–[17]. In fact, quantum control techniques can be used to realize quantum gates [18]–[20], and quantum gates have numerous applications in quantum information science, including localization, synchronization, communication, and computing [21]–[24].

Quantum optimal control (QOC) has been of significant research interest for the first part of the 21st century [25]–[34]. However, much like classical systems, the dynamics of quantum systems are not always known, which limits the breadth of control strategies that may be used. A promising means for overcoming this uncertainty is via quantum Hamiltonian learning (QHL). Learning quantum dynamics from data is difficult due to the fundamental limitations of observing quantum phenomena via classical measurement. Measurement devices interact with quantum states and can result in probabilistic outcomes. Quantum learning-based system identification strategies must work with incomplete or imperfect information about quantum states.

To overcome the obstacle of incomplete information about quantum states, methods such as quantum state tomography (QST) have been developed to infer quantum states via repeated measurements of identically prepared states [35]–[38]. Quantum process tomography (QPT) has been developed to infer a quantum process, which maps an input state to an output state [6]. A survey on resource consumption and measurement requirements for various QPT algorithms is given in [39].

Both QST and QPT have been used to infer the evolution of closed quantum systems and ultimately learn the system Hamiltonian in [40]. The method in [40] utilizes a

large number of initial states to perform QPT. The work in [40] gives an explicit upper bound on the estimation error for the learned Hamiltonian. Conventional system identification approaches have been used to estimate the Hamiltonian [41]–[44]. In particular, control signals are used in [43] and [44] to infer unknown parts of the Hamiltonian. Alternatively, learning-based approaches have been used to estimate the Hamiltonian [45]–[47].

Inspired by the learning-based control of classical systems [48]–[51], the learning-based control of quantum systems has also been proposed [52]–[59]. These methods work in an iterative fashion to improve the control policy used in each quantum experiment. However, these techniques are model-free and do not come with theoretical guarantees. On the other hand, if a model of the quantum system is established, one may perform model-based control. For several decades, QOC has been the most popular approach in the model-based quantum control literature. In addition to some early works in this area [60]–[62], recent results have reiterated the strength of QOC and expanded upon earlier possibilities [63]–[65]. An influential QOC technique is the so-called gradient ascent pulse engineering (GRAPE) algorithm for computing optimal nuclear magnetic resonance (NMR) pulses [61], which was expanded upon in [66]–[69]. The principle of the GRAPE algorithm is to use gradient-based solvers to compute optimal control parameters for quantum systems. Other highly successful QOC optimization schemes include the chopped random basis (CRAB) [70], [71] and gradient optimization of analytic controls (GOAT) [72] algorithms, which use ansatzes of basis functions to formulate control pulses. The purpose of these algorithms is to identify optimal tuning parameters within the ansatz.

In the classical control literature, layering model-based and model-free control has been shown to produce desirable results [73]. This has been done recently in the so-called C^3 technique for controlling quantum systems [74]. The term C^3 stands for "control, calibrate, and characterize." As the name implies, this work meshes the ideas of controlling and identifying the dynamics of the system. The first phase utilizes a control ansatz based on the derivative removal by adiabatic gate (DRAG) method. The tuning parameters in the DRAG control pulse are then optimized in a closed-loop experimental setting with a physical quantum system in the calibration phase. Finally, the characterization phase uses the data generated in the previous phase to update an ansatz model of the quantum system.

Model-predictive control (MPC) for infinite-dimensional quantum systems has been considered in the case where the model of the system is fully known [75]–[77]. These works only investigate systems with pure states governed by the Schrödinger equation, and their control designs only consider quadratic objectives. Recently, an optimization procedure known as sequential quadratic programming (SQP) has shown great promise to design robust quantum gates in the context of QOC [78]. We note that SQP can also be employed to solve MPC problems. Hence, investigating SQP's use in MPC for quantum systems is of interest.

Most of the methods outlined above are based on the optimal control of closed quantum systems, and additional calibration and refinement may be needed for addressing open quantum systems. Real-world applications of controlling closed quantum systems include, but are not limited to, control of spin systems in NMR [79]–[82], control of molecular systems in physical chemistry [83]–[85], and forming baseline control policies for manipulating superconducting qubits [86]–[91].

Learning and control of quantum dynamics poses challenges not yet solved by the literature. There is a need to learn both internal and control Hamiltonians. The internal Hamiltonian describes the free evolution of the quantum system, and the control Hamiltonian describes the system's interaction with external fields. While successful, the quantum tomography (QT)-based QHL algorithm proposed in [40] requires the practitioner to readily and repeatedly prepare d^2 (where d is the dimension of the system) unique quantum states to identify a system's internal Hamiltonian and does not address learning the effect of control on the system. In terms of control, QOC is often computationally expensive, and closed-form solutions are only known in specific cases with simple models or restrictive assumptions. For instance, time-optimal control of a qubit is known to be of bang–bang type [26]; however, such discontinuous control signals may excite unwanted energy levels in qubit. The work-horse GRAPE algorithm, which has become the standard for computing optimal quantum control pulses, still takes a significant amount of computational time [92]. Moreover, data-driven techniques are desirable when no model of the quantum system is known. While the C^3 method [74] has shown promising results, like [70]–[72], it utilizes a heuristic ansatz for its control policy and can only improve the efficacy of the control by tuning parameters (such as the amplitude and frequency of the actuation) within this policy. In summary, there is a need for *scalable* and *efficient* quantum learning and control algorithms. The fundamental questions related to quantum system identification are: Is it possible to reduce the number of quantum experiments required, computational complexity demanded, and memory storage needed in order to properly identify a quantum process? How can the error of various Hamiltonian identification methods be bounded in terms of the number of quantum experiments performed? A reduction in experimental, computational, and memory complexity will enable the identification of higher dimensional quantum systems—which is paramount to the future of quantum computing. Improved error bounds on this process will allow practitioners to design robust model predictive quantum control (MPQC) strategies. The goal of this article is to provide an entirely data-driven end-to-end solution to the quantum control problem.

In this article, we attempt to bridge the gap between learning and control for quantum systems. By not relying on

IEEE Transactions on Quantum Engineering

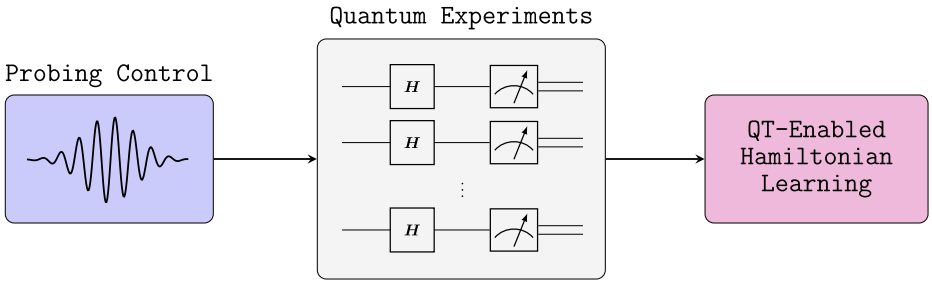

FIGURE 1. QT-enabled QHL utilizes probing control inputs and physical experiments to infer internal and control Hamiltonians. Details of the QT-enabled QHL block are discussed in Section II and Fig. 3.

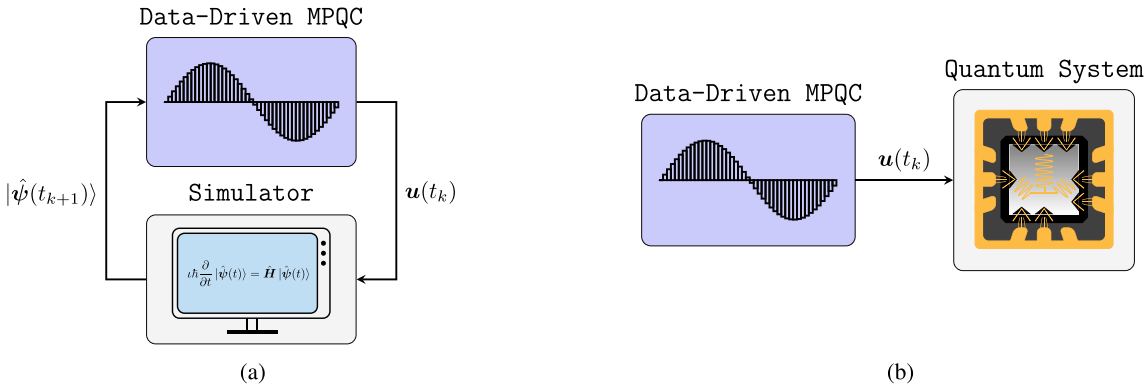

FIGURE 2. Data-driven MPQC utilizes the trained model from the proposed QT-enabled QHL technique to optimize a control sequence offline in simulation. Then, the optimized control sequence is provided to a real-world quantum system in an open-loop fashion. (a) Offline optimization. (b) Online implementation.

manually designed models, which can be brittle, our method will account for the complexity of real-world applications. The contributions of this article are twofold. First, a new QT-based QHL algorithm for identifying both internal and control Hamiltonians for closed quantum systems is proposed. This technique is visually summarized in Fig. 1. Given a physical system under the influence of an external control field, our method utilizes QST and QPT to learn the Hamiltonians of interest. Within this algorithm, a new optimization-based QPT procedure is presented. This optimization procedure is defined over the unitary group, i.e., the set of all unitary operators, which is physically motivated by the dynamics of closed quantum systems. An efficient iterative method for solving this problem is presented, and, in a special case, a closed-form solution and error bound are given. The proposed QT-based QHL algorithm requires less memory than the existing state-of-the-art and, in some cases, greatly reduces the computational complexity.

The second contribution is that, for the first time ever, we formulate data-driven MPQC, which uses the proposed QT-based QHL algorithm to control quantum systems end-to-end with unknown dynamics. MPQC leverages the ideas from MPC for classical systems to produce an optimization-based approach to controlling quantum systems. In our experiments, when an SQP solver is used, MPQC is two orders of magnitude faster at computing control sequences than QOC, yet provides similar control performance. The MPQC formulation is summarized in Fig. 2. First, MPQC generates a control sequence offline in a closed-loop simulation using a model of the quantum system [see Fig. 2(a)]. In this simulation, the state of the quantum system can be fully known, unlike that of a physical quantum system. Once a control sequence has been optimized offline, it may be provided to a physical quantum system in an open-loop fashion so as to preserve the coherence of the physical system [see Fig. 2(b)].

The key contributions of this article are as follows:

1) develop a novel data-driven QHL algorithm to estimate both internal and control Hamiltonians;
2) prove error bounds on the proposed control QHL methods, under appropriate assumptions;
3) formalize an MPQC framework for the learning-enabled open-loop control of quantum systems;
4) illustrate the scalability and efficacy of the proposed data-driven MPQC via numerical experiments.

The rest of this article is organized as follows. Section II describes QPT and our novel approach to QPT. Section III utilizes the proposed QPT method to perform QHL and identify both the internal and control Hamiltonians of a closed quantum system. This is followed by Section IV, which proposes data-driven MPQC. Numerical results are given in Section V. Finally, Section VI concludes this article.

Notations: Sets are denoted by calligraphic font, e.g., $\mathcal{X}$. In particular, $\mathcal{Z}_{m:m+M}$ is defined to be the set of integers from

TABLE 1. Notations and Definitions of Important Quantities

Notation	Definition	Notation	Definition
$\|\boldsymbol{\psi}(t)\rangle$	state vector at time t	$\boldsymbol{\Xi}(t)$	density operator at time t
$\boldsymbol{H}_0$	internal (free) Hamiltonian	$\boldsymbol{H}_\mathrm{c}$	control Hamiltonian
d	Hilbert space dimension	$\mathcal{U}(d)$	unitary group of dimension d
$\boldsymbol{U}(t)$	unitary evolution operator at time t	$\boldsymbol{\Delta}_{\boldsymbol{\phi}}$	diagonal matrix whose n-th element is $e^{\imath\boldsymbol{\phi}_n}$
t_0	the time that a quantum experiment begins	t_f	the time that a quantum experiment ends
$\|\mathring{\boldsymbol{\psi}}_n(t_0)\rangle$	the state of experiment n at time t_0	$\|\mathring{\boldsymbol{\psi}}_n(t_\mathrm{f})\rangle$	the state of experiment n at time t_f
$\mathring{\boldsymbol{\Xi}}_n(t_0)$	the density operator of experiment n at time t_0	$\mathring{\boldsymbol{\Xi}}_n(t_\mathrm{f})$	the density operator of experiment n at time t_f
$\mathring{\boldsymbol{\Psi}}_{1:N_\mathrm{i}}(t_0)$	matrix whose n-th column is $\|\mathring{\boldsymbol{\psi}}_n(t_0)\rangle$ for $n \in \{1,2,\ldots,N_\mathrm{i}\}$	$\mathring{\boldsymbol{\Psi}}_{1:N_\mathrm{i}}(t_\mathrm{f})$	matrix whose n-th column is $\|\mathring{\boldsymbol{\psi}}_n(t_\mathrm{f})\rangle$ for $n \in \{1,2,\ldots,N_\mathrm{i}\}$
$\hat{\boldsymbol{\Xi}}_{1:N_\mathrm{i}}(t_0)$	matrix obtained by concatenating $\{\mathring{\boldsymbol{\Xi}}_n(t_0)\}_{n\in\mathcal{Z}_{N_\mathrm{i}}}$ in a row	$\hat{\boldsymbol{\Xi}}_{1:N_\mathrm{i}}(t_\mathrm{f})$	matrix obtained by concatenating $\{\mathring{\boldsymbol{\Xi}}_n(t_\mathrm{f})\}_{n\in\mathcal{Z}_{N_\mathrm{i}}}$ in a row
N_i	number of input-output state pairs used by QPT	$\boldsymbol{H}_l$	l-th interaction Hamiltonian
N_c	number of control inputs	u_max	maximum magnitude of control input allowed
K_p	prediction horizon	K_c	control horizon
$\|\boldsymbol{\psi}(t_k)\rangle$	state vector at time step k of the optimal control problem	$\boldsymbol{H}(t_k)$	net Hamiltonian at time step k of the optimal control problem
$\boldsymbol{u}(t_k)$	control vector at time step k of the optimal control problem	Δ_t	length of ZOH used in the optimal control problem
$\|\boldsymbol{\psi}_\mathrm{d}\rangle$	desired state	$\|\boldsymbol{\psi}_\mathrm{bad}\rangle$	forbidden state
K	number of control steps	$\hat{A}$	estimate of the quantity A

m up to and including $m+M$. For brevity, set $\mathcal{Z}_M \triangleq \mathcal{Z}_{1:M}$. Random variables are displayed in sans serif upright fonts; their realizations in serif italic fonts. Vectors and matrices are denoted by bold lowercase and uppercase letters, respectively. For example, a random vector and its realization are denoted by x and $\boldsymbol{x}$; a random matrix and its realization are denoted by X and $\boldsymbol{X}$, respectively. In the context of quantum states, complex vectors are denoted using bra-ket notation, i.e., $|\boldsymbol{\psi}\rangle$ and $\langle\boldsymbol{\psi}|$ denote a complex vector and its conjugate transpose, respectively. The d-by-d identity matrix is denoted by $\boldsymbol{I}_d$: the subscript is removed when the dimension of the matrix is clear from the context. The Frobenius norm and spectral norm of $\boldsymbol{X}$ are denoted by $\|\boldsymbol{X}\|_\mathrm{F}$ and $\|\boldsymbol{X}\|$, respectively. The Euclidean and infinity norms of the vector $\boldsymbol{x}$ are denoted $\|\boldsymbol{x}\|$ and $\|\boldsymbol{x}\|_\infty$, respectively. The notation $\mathrm{diag}(A_1, A_2, \ldots, A_n)$ represents a block diagonal matrix with the arguments on its main diagonal. The trace, transpose, conjugate, and conjugate-transpose of the matrix $\boldsymbol{X}$ are denoted by $\mathrm{tr}\{\boldsymbol{X}\}$, $\boldsymbol{X}^T$, $\boldsymbol{X}^*$, and $\boldsymbol{X}^\dagger$, respectively. $[\![\boldsymbol{A}, \boldsymbol{B}]\!]_-$ denotes the commutator of matrices $\boldsymbol{A}$ and $\boldsymbol{B}$. The Kronecker product is denoted by the symbol $\otimes$, and we will denote the Kronecker product of multiple state vectors as $|\boldsymbol{\psi}_1\,\boldsymbol{\psi}_2\cdots\boldsymbol{\psi}_n\rangle \triangleq |\boldsymbol{\psi}_1\rangle \otimes |\boldsymbol{\psi}_2\rangle \otimes \cdots \otimes |\boldsymbol{\psi}_n\rangle$. The notations $\mathfrak{R}\{\cdot\}$ and $\mathfrak{F}\{\cdot\}$ denote the real and imaginary parts of the given argument, respectively. If $f, g : \mathbb{C}^{p\times d} \to \mathbb{R}$ are arbitrary functionals, f is said to grow on the order $\mathcal{O}(g(\boldsymbol{X}))$ if there exists some numerical constant $C \in \mathbb{R}$ such that $\big|f(\boldsymbol{X})\big| \leqslant C\big|g(\boldsymbol{X})\big|$ for all $\boldsymbol{X} \in \mathbb{C}^{p\times d}$. The magnitude of a complex number $z = a + \imath b \in \mathbb{C}$ is denoted by $|z| = \sqrt{a^2+b^2}$, where $\imath = \sqrt{-1}$. Notations and definitions for important quantities used in this article are summarized in Table 1. Important acronyms and optimization problems are also summarized in Tables 2 and 3, respectively.

II. SCALABLE QPT ALGORITHM

In this section, we present a new and scalable QPT algorithm. For ease of exposition, Section II-A considers systems which start from pure states and Section II-B addresses those which start from mixed states. If one's experiments start from both pure and mixed states, the method provided for mixed states will suffice since it is possible to represent any pure quantum state as a mixed state. The proposed QPT formulations rely on solving an optimization problem on the unitary group, which may be achieved via efficient iterative algorithms. Sections II-C gives a method for computing iterative solutions to the QPT problem. Section II-D presents a closed-form solution to the QPT problem in a special case, which provides further insight into the nature of the proposed method.

A. LEARNING FROM PURE STATES

We begin by considering a closed quantum system and assume no explicit external control influence. A *pure* state vector, or wave function, $|\boldsymbol{\psi}(t)\rangle \in \mathbb{C}^d$ evolves according to the

IEEE Transactions on Quantum Engineering

TABLE 2. Expansions of Important Acronyms

Acronym	Expansion	Acronym	Expansion
MPQC	model-predictive quantum control	QOC	quantum optimal control
SQP	sequential quadratic programming	GRAPE	GRadient Ascent Pulse Engineering
QST	quantum state tomography	QPT	quantum process tomography
QT	quantum tomography	ZOH	zero-order hold
QHL	quantum Hamiltonian learning	CQHL	control quantum Hamiltonian learning

TABLE 3. Notation for Optimization Problems

Notation	Definition
$\mathscr{P}_1$	QPT for pure states
$\mathscr{P}_2$	Equivalent QPT in terms of the Brockett cost function
$\mathscr{P}_3$	QPT for mixed states
$\mathscr{P}_4$	Problem solved at each time step by MPQC for controlling quantum states
$\mathscr{P}_5$	Problem solved at each time step by MPQC for generating unitary gates
$\mathscr{P}_6$	GRAPE for generating unitary gates

Schrödinger equation

$$\imath\hbar\frac{\partial}{\partial t}\big|\psi(t)\big\rangle = \boldsymbol{H}\big|\psi(t)\big\rangle \tag{1}$$

where $\hbar$ is the reduced Planck constant and $\boldsymbol{H} \in \mathbb{C}^{d\times d}$ is the system Hamiltonian, which is an unknown Hermitian operator. For now, $\boldsymbol{H}$ is a *time-independent* Hamiltonian. In this article, we elect to use atomic units such that $\hbar = 1$. At any time, $|\psi(t)\rangle$ is of unit length such that $\big\||\psi(t)\rangle\big\| = 1$. Given the initial state $|\psi(0)\rangle \in \mathbb{C}^d$, it is possible to calculate the state at any arbitrary time in the future by directly solving the Schrödinger equation:

$$\begin{aligned}\big|\psi(t)\big\rangle &= e^{-\imath\boldsymbol{H}t}\big|\psi(0)\big\rangle \\ &= \boldsymbol{U}(t)\big|\psi(0)\big\rangle\,.\end{aligned} \tag{2}$$

Here, we have defined the operator $\boldsymbol{U}(t) \triangleq e^{-\imath\boldsymbol{H}t}$. The matrix exponential of any skew-Hermitian matrix is unitary. Hence, $\boldsymbol{U}(t)^\dagger\boldsymbol{U}(t) = \boldsymbol{I}_d$. The relationship (2) highlights the importance of closed quantum systems: since the evolution of the state vector is unitary, the system preserves state coherence.

The set of all unitary operators defined on the space $\mathbb{C}^{d\times d}$ is called the *unitary group*, which is denoted

$$\mathcal{U}(d) \triangleq \left\{\boldsymbol{X} \in \mathbb{C}^{d\times d} : \boldsymbol{X}^\dagger\boldsymbol{X} = \boldsymbol{I}_d\right\}.$$

Throughout this article, every problem for learning a quantum system's dynamics will be formalized as an optimization problem defined on the unitary group. Because closed quantum systems evolve unitarily, $\boldsymbol{U}(t) \in \mathcal{U}(d)$ for all t.

We are interested in a system identification problem of the following form. First, let $t_\mathrm{f} \in (0,\infty)$ be some final time where we perform a measurement and end the experiment. The unitary operator $\boldsymbol{U}(t_\mathrm{f}) : \hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_0) \mapsto \hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$ maps each input state to its corresponding output state, where

$$\hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_0) = \Big[\big|\hat{\hat{\psi}}_1(t_0)\big\rangle\ \big|\hat{\hat{\psi}}_2(t_0)\big\rangle\ \cdots\ \big|\hat{\hat{\psi}}_{N_\mathrm{i}}(t_0)\big\rangle\Big] \tag{3}$$

is a matrix of N_i initial states, whose nth column is a state vector $|\hat{\hat{\psi}}_n(t_0)\rangle$, and

$$\hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f}) = \Big[\big|\hat{\hat{\psi}}_1(t_\mathrm{f})\big\rangle\ \big|\hat{\hat{\psi}}_2(t_\mathrm{f})\big\rangle\ \cdots\ \big|\hat{\hat{\psi}}_{N_\mathrm{i}}(t_\mathrm{f})\big\rangle\Big] \tag{4}$$

is a matrix of N_i final states, whose nth column is a state vector $|\hat{\hat{\psi}}_n(t_\mathrm{f})\rangle$. The system identification goal is to learn the mapping $\boldsymbol{U}(t_\mathrm{f})$, which will be accomplished via QPT. As shown in Section III, it is possible to estimate the system Hamiltonian $\boldsymbol{H}$ from $\boldsymbol{U}(t_\mathrm{f})$.

It is assumed that the system is initiated in a known state. However, it is only possible to *estimate* the final state matrix $\hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$, which can be accomplished via QST.[1] The associated *state estimation error* can be written as

$$\boldsymbol{E} = \hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f}) - \hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f}) \tag{5}$$

where $\hat{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$ is the estimate of the final states matrix, with the nth column denoted by $|\hat{\hat{\psi}}_n(t_\mathrm{f})\rangle$. The QPT methods proposed in this article are *agnostic* to the underlying QST algorithms used to estimate the final states.

To identify $\boldsymbol{U}(t_\mathrm{f})$, one needs a measure of distance between two quantum states. An appropriate measure is the *infidelity* [93]. For two pure quantum states $|\psi\rangle$ and $|\phi\rangle$, the infidelity between them is denoted and defined by

$$d_\mathrm{I}(|\psi\rangle, |\phi\rangle) = 1 - |\langle\psi|\phi\rangle|^2\,.$$

This function takes on values in the range of 0 to 1. An infidelity of one corresponds to perfectly orthogonal states, and

[1]Several practical QST algorithms are given in [35]–[38]. Through repeated measurements across multiple experiments, QST estimates probabilities associated with different measurement outcomes. Then, QST uses these estimated probabilities to infer the final state of the quantum system.

an infidelity of zero indicates that $|\psi\rangle$ and $|\phi\rangle$ are identical up to a possible global phase difference. In fact, the selection of the global phase is arbitrary, namely, $e^{\imath\theta}|\psi\rangle$ and $|\psi\rangle$ are functionally equivalent in quantum dynamics, and they represent the same physical state for any global phase $\theta \in \mathbb{R}$. An important interpretation of the infidelity is as follows. The term $|\langle\psi|\phi\rangle|^2$ is the *probability* that $|\psi\rangle$ will be measured to be $|\phi\rangle$. Hence, minimizing the infidelity between two pure quantum states translates into maximizing the probability that one will be observed as the other.

Since $\boldsymbol{U}(t_\mathrm{f}) : |\mathring{\hat{\psi}}(t_0)\rangle \mapsto |\mathring{\hat{\psi}}(t_\mathrm{f})\rangle$, we wish to find an estimate $\hat{\boldsymbol{U}}(t_\mathrm{f})$, which minimizes the distance $d_\mathrm{I}(|\mathring{\hat{\psi}}(t_\mathrm{f})\rangle, \hat{\boldsymbol{U}}(t_\mathrm{f})|\mathring{\hat{\psi}}(t_0)\rangle)$, where $|\mathring{\hat{\psi}}(t_\mathrm{f})\rangle$ is an estimate of the final state obtained by using a QST algorithm. This is the distance between the estimated final state and the one predicted by $\hat{\boldsymbol{U}}(t_\mathrm{f})|\mathring{\hat{\psi}}(t_0)\rangle$. Hereafter, we will omit the explicit dependence of $\hat{\boldsymbol{U}}(t_\mathrm{f})$ on t_f when it is clear from context. Since there are many initial states and final states, we seek to minimize the sum of the infidelities between all the pairs of initial conditions and final conditions. This gives rise to the following optimization problem, which is the basis of our proposed QPT algorithm:

$$\mathscr{P}_1: \quad \underset{\boldsymbol{X}}{\text{minimize}} \quad \sum_{n=1}^{N_\mathrm{i}} d_\mathrm{I}\big(|\mathring{\hat{\psi}}_n(t_\mathrm{f})\rangle, \boldsymbol{X}|\mathring{\psi}_n(t_0)\rangle\big)$$
$$\text{subject to} \quad \boldsymbol{X} \in \mathcal{U}(d)$$

where the objective function sums over all state pairs in (3) and the estimate $\mathring{\hat{\boldsymbol{\Psi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$ of (4). Let $\hat{\boldsymbol{U}}$ denote the optimal solution of $\mathscr{P}_1$. It should be emphasized that Problem 1 is an optimization problem defined over a constraint manifold, which is a distinct feature of our QPT formulation.

Next, Problem 1 is transformed into a simpler problem, which can be solved efficiently. The new problem takes similar form to the well-studied *Brockett cost function* from the controls and optimization literature [94]. The Brockett cost function $f_\mathrm{B} : \mathcal{U}(d) \to \mathbb{R}$ is defined as

$$f_\mathrm{B}(\boldsymbol{X}) = \sum_{n=1}^{N_\mathrm{i}} \mathrm{tr}\left\{\boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n \boldsymbol{X}^\dagger\right\}. \tag{6}$$

where $\boldsymbol{A}_n \in \mathbb{C}^{d\times d}$ and $\boldsymbol{B}_n \in \mathbb{C}^{d\times d}$ are Hermitian positive-semidefinite matrices for all n.

Remark 1: Brockett's original article considered a slightly different functional defined over the real orthogonal group. Since then, the literature has widely referred to functions of the form $f(\boldsymbol{X}) = \mathrm{tr}\{\boldsymbol{N}\boldsymbol{X}\boldsymbol{A}\boldsymbol{X}^\dagger\}$, where $\boldsymbol{N}$ is diagonal and $\boldsymbol{A}$ is real symmetric to be the "Brockett cost function." Our definition, with the summation, is closer to Brockett's definition in [94]. □

Define the density operators $\mathring{\hat{\boldsymbol{\Xi}}}_n(t_0) \triangleq |\mathring{\hat{\psi}}_n(t_0)\rangle\langle\mathring{\hat{\psi}}_n(t_0)|$ and $\mathring{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f}) \triangleq |\mathring{\hat{\psi}}_n(t_\mathrm{f})\rangle\langle\mathring{\hat{\psi}}_n(t_\mathrm{f})|$ associated with the pure states $|\mathring{\hat{\psi}}_n(t_0)\rangle$ and $|\mathring{\hat{\psi}}_n(t_\mathrm{f})\rangle$, respectively.[2] The following proposition relates Problem 1 to the Brockett cost function.

Proposition 1: Solving Problem 1 is equivalent to solving the following optimization problem:

$$\mathscr{P}_2: \quad \underset{\boldsymbol{X}}{\text{maximize}} \quad \sum_{n=1}^{N_\mathrm{i}} \mathrm{tr}\big\{\mathring{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f})\, \boldsymbol{X}\, \mathring{\hat{\boldsymbol{\Xi}}}_n(t_0)\, \boldsymbol{X}^\dagger\big\}$$
$$\text{subject to} \quad \boldsymbol{X} \in \mathcal{U}(d).$$

□

Proof: We prove the case for $N_\mathrm{i} = 1$; the case for $N_\mathrm{i} > 1$ follows directly. First, note that to minimize the objective function in Problem 1, we need only maximize $|\langle\mathring{\hat{\psi}}(t_\mathrm{f})|\boldsymbol{X}|\mathring{\psi}(t_0)\rangle|^2$. Expanding this expression, we obtain

$$\begin{aligned}
&\big|\big\langle\mathring{\hat{\psi}}(t_\mathrm{f})\big|\boldsymbol{X}\big|\mathring{\psi}(t_0)\big\rangle\big|^2 \\
&\quad= \big\langle\mathring{\hat{\psi}}(t_\mathrm{f})\big|\boldsymbol{X}\big|\mathring{\psi}(t_0)\big\rangle\ \big\langle\mathring{\psi}(t_0)\big|\boldsymbol{X}^\dagger\big|\mathring{\hat{\psi}}(t_\mathrm{f})\big\rangle \\
&\quad= \mathrm{tr}\big\{\big\langle\mathring{\hat{\psi}}(t_\mathrm{f})\big|\boldsymbol{X}\big|\mathring{\psi}(t_0)\big\rangle\ \big\langle\mathring{\psi}(t_0)\big|\boldsymbol{X}^\dagger\big|\mathring{\hat{\psi}}(t_\mathrm{f})\big\rangle\big\} \\
&\quad= \mathrm{tr}\big\{\big|\mathring{\hat{\psi}}(t_\mathrm{f})\big\rangle\big\langle\mathring{\hat{\psi}}(t_\mathrm{f})\big|\boldsymbol{X}\big|\mathring{\psi}(t_0)\big\rangle\big\langle\mathring{\psi}(t_0)\big|\boldsymbol{X}^\dagger\big\} \\
&\quad= \mathrm{tr}\big\{\mathring{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})\, \boldsymbol{X}\, \mathring{\hat{\boldsymbol{\Xi}}}(t_0)\, \boldsymbol{X}^\dagger\big\}.
\end{aligned}$$

The second equality holds since the trace of a scalar is equivalent to the scalar itself. The third equality is due to the cyclic property of the trace. The final equality follows from the definition of the density operator. □

In the general case, we propose an iterative gradient-based algorithm for solving Problem 2. This solver and its efficiency will be described in Section II-C. A closed-form solution in the special case where $N_\mathrm{i} = 1$ is given in Section II-D. In the next section, learning from mixed states is addressed.

B. LEARNING FROM MIXED STATES

Recall from Section II-A that the density operator for a pure state $|\psi\rangle$ is denoted and defined by $\boldsymbol{\Xi} = |\psi\rangle\langle\psi|$. In order to properly represent *mixed states*, one should use the density operator $\boldsymbol{\Xi}$ instead of the state vector $|\psi\rangle$. A mixed state is a probabilistic mixture, or classical ensemble, of pure states. In other words, a mixed state is characterized by the density operator

$$\boldsymbol{\Xi} = \sum_i \eta_i |\psi_i\rangle\langle\psi_i|$$

where $\sum_i \eta_i = 1$, $\eta_i \geqslant 0$, and $|\psi_i\rangle$ is a pure state for all i. A mixed state cannot be described by a single state vector alone; however, both the pure and mixed states can be fully characterized by density operators. The density operator $\boldsymbol{\Xi}$ is a Hermitian positive semidefinite matrix satisfying $\mathrm{tr}\{\boldsymbol{\Xi}\} =$

[2] In the quantum physics literature, it is common to use the variable ρ to denote the density operator. However, we elect to use $\boldsymbol{\Xi}$ to emphasize that this is a matrix quantity.

IEEE Transactions on Quantum Engineering

1. Mixed states in a closed quantum system evolve according to the Liouville–von Neumann equation

$$\imath \frac{\partial}{\partial t} \boldsymbol{\Xi} = [\![\boldsymbol{H}, \boldsymbol{\Xi}]\!]_- . \tag{7}$$

The solution to this equation is given by

$$\boldsymbol{\Xi}(t) = \boldsymbol{U}(t)\, \boldsymbol{\Xi}(0)\, \boldsymbol{U}^\dagger(t) \tag{8}$$

where $\boldsymbol{U}(t) = e^{-\imath \boldsymbol{H} t}$ as usual. This solution can also be derived from the fact that each state in the classical ensemble $\boldsymbol{\Xi}$ is governed by (1).

As for the case where the system starts from a pure state, QPT for mixed states starts by fixing some sampling time $t_\mathrm{f} \in (0, \infty)$. A matrix of $N_\mathrm{i} \in \mathbb{N}$ initial states is given

$$\hat{\boldsymbol{\Xi}}_{1:N_\mathrm{i}}(t_0) = \Big[\hat{\boldsymbol{\Xi}}_1(t_0)\ \hat{\boldsymbol{\Xi}}_2(t_0)\ \cdots\ \hat{\boldsymbol{\Xi}}_{N_\mathrm{i}}(t_0) \Big] .$$

For each initial state, quantum experiments are run, and the final states are measured at time t_f. Similar to the case of pure states, an estimate of the final state matrices is denoted by

$$\hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_\mathrm{f}) = \Big[\hat{\hat{\boldsymbol{\Xi}}}_1(t_\mathrm{f})\ \hat{\hat{\boldsymbol{\Xi}}}_2(t_\mathrm{f})\ \cdots\ \hat{\hat{\boldsymbol{\Xi}}}_{N_\mathrm{i}}(t_\mathrm{f}) \Big] .$$

As before, the state estimation error for the final states is given by

$$\boldsymbol{E} = \hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_\mathrm{f}) - \hat{\boldsymbol{\Xi}}_{1:N_\mathrm{i}}(t_\mathrm{f}) . \tag{9}$$

The nth block column of $\hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$ (i.e., the estimate of the nth final state) is simply $\hat{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f})$. Like their pure counterparts, estimates of mixed states are given by QST, and our method remains *agnostic* to what QST method is employed.

Mixed states are defined on the Hilbert space $\mathbb{C}^{d\times d}$ endowed with the canonical Frobenius inner product, which induces the Frobenius norm. This leads to a natural metric, or measure of distance, between two matrices $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{C}^{d\times d}$. Namely, $d_\mathrm{F}(\boldsymbol{X}, \boldsymbol{Y}) = \big\| \boldsymbol{X} - \boldsymbol{Y} \big\|_\mathrm{F}$. In the case of mixed states, we propose the following optimization problem which minimizes the sum of the squared distances between $\hat{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f})$ and $\boldsymbol{X} \hat{\boldsymbol{\Xi}}_n(t_0) \boldsymbol{X}^\dagger$ to produce an estimate of the unitary dynamics:

$$\mathscr{P}_3 : \ \underset{\boldsymbol{X}}{\text{minimize}} \quad \sum_{n=1}^{N_\mathrm{i}} d_\mathrm{F}^2 \big(\hat{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f}), \boldsymbol{X} \hat{\boldsymbol{\Xi}}_n(t_0) \boldsymbol{X}^\dagger \big)$$
$$\text{subject to} \ \ \boldsymbol{X} \in \mathcal{U}(d)$$

where $\hat{\boldsymbol{U}}$ denotes the optimal solution of $\mathscr{P}_3$, which is an optimization problem defined on the unitary group. This problem minimizes the Frobenius norm of the difference between the predicted final states, $\hat{\boldsymbol{U}} \hat{\boldsymbol{\Xi}}_n(t_0) \hat{\boldsymbol{U}}^\dagger$, and the QST estimate of the final states, $\hat{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f})$. Like the proposed QPT formulation for pure states, Problem 3 may be simplified into an easier problem.

Proposition 2: Using the density representation for mixed states, solving Problem 3 is equivalent to solving Problem 2. □

Proof: We prove the case where $N_\mathrm{i} = 1$, which is easily extended to the case for any finite N_i. To begin, we observe that

$$\begin{aligned} &\big\| \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \big\|_\mathrm{F}^2 \\ &\quad = \mathrm{tr}\Big\{ \Big[\boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \Big] \Big[\boldsymbol{X} \hat{\boldsymbol{\Xi}}_k(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \Big]^\dagger \Big\} \\ &\quad = \mathrm{tr}\Big\{ \Big[\boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \Big]^2 \Big\} . \end{aligned}$$

The first equality is the definition of the Frobenius norm. The second equality is due to the fact that $[\boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})]$ is Hermitian. Expanding the argument of the prior expression, we obtain

$$\begin{aligned} &\big[\boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \big]^2 \\ &\quad = \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \\ &\qquad - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})\, \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger + \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})^2 . \end{aligned}$$

The cyclic property of the trace function gives

$$\mathrm{tr}\big\{ \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \big\} = \mathrm{tr}\big\{ \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \big\}$$

and

$$\begin{aligned} &\mathrm{tr}\big\{ \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \big\} \\ &\qquad = \mathrm{tr}\big\{ \boldsymbol{X}^\dagger \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \big\} . \end{aligned}$$

Combining the prior two expressions, the constraint that $\boldsymbol{X}$ is unitary and the linearity property of the trace function, the following is obtained:

$$\begin{aligned} &\big\| \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger - \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \big\|_\mathrm{F}^2 \\ &\quad = \mathrm{tr}\big\{ \hat{\boldsymbol{\Xi}}(t_0)^2 \big\} - 2\mathrm{tr}\big\{ \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \big\} + \mathrm{tr}\big\{ \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})^2 \big\} . \end{aligned}$$

Since $\hat{\boldsymbol{\Xi}}(t_0)$ and $\hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})$ are fixed at the time of optimization and are not functions of the variable $\boldsymbol{X}$, only the term $\mathrm{tr}\{ \hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) \boldsymbol{X} \hat{\boldsymbol{\Xi}}(t_0) \boldsymbol{X}^\dagger \}$ can be optimized. To solve the original optimization problem, we wish to maximize this term subject to the constraint that $\boldsymbol{X}$ is unitary, and the proposition is proved. □

Remark 2: It is a pleasing result that Problem 1 (designed for pure states) and Problem 3 (designed for mixed states) *both* simplify to Problem 2. For the rest of this article, we will refer to "solving Problem 2" unambiguously as it treats both pure and mixed states equally. □

C. GENERAL CASE: $N_\mathrm{i} \geqslant 1$

In the case where $N_\mathrm{i} \geqslant 1$, we provide a computationally efficient algorithm, which iteratively *maximizes* the objective function (i.e., a Brockett cost function) and solves Problem 2. This process is equivalent to performing steepest ascent on the unitary group. Our methodology is inspired by that of [95], which was introduced in the context of *minimizing* the Brockett cost function on the *Stiefel manifold* [96], [97]

when $N_\mathrm{i} = 1$. We now proceed by first introducing preliminaries required to implement the proposed iterative method. For generality, we initially adopt the same general notation as in (6) for the Brockett cost and prove the main result. This is done to emphasize the potential application of said result to other domains and simplify the notation. At the end of this section, the general result is used to solve Problem 2 using the appropriate notation.

The unitary group is an example of a *Lie group*, which is also a differentiable manifold [96]. The tangent space of $\mathcal{U}(d)$ at any point $\boldsymbol{X} \in \mathcal{U}(d)$ is denoted and defined as

$$T_{\boldsymbol{X}} \triangleq \left\{ \boldsymbol{X}\boldsymbol{A} : \ \boldsymbol{A} \in \mathbb{C}^{d\times d},\ \boldsymbol{A} + \boldsymbol{A}^\dagger = \boldsymbol{0} \right\}.$$

The projection of an arbitrary matrix $\boldsymbol{Z} \in \mathbb{C}^{d\times d}$ onto the unitary group is denoted and defined by

$$\mathcal{P}_\mathcal{U}(\boldsymbol{Z}) \triangleq \operatorname*{argmin}_{\boldsymbol{X}\in\mathcal{U}(d)} \left\| \boldsymbol{Z} - \boldsymbol{X} \right\|_\mathrm{F}^2.$$

In fact, assuming that the singular value decomposition of $\boldsymbol{Z}$ is $\boldsymbol{Z} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\dagger$, then $\mathcal{P}_\mathcal{U}(\boldsymbol{Z}) = \boldsymbol{U}\boldsymbol{V}^\dagger$ [98]. For any $\boldsymbol{X}$ in either $\mathcal{U}(d)$, the canonical inner product on the tangent space $T_{\boldsymbol{X}}$ is defined as follows:

$$\langle \boldsymbol{Z}_1, \boldsymbol{Z}_2 \rangle_T = \frac{1}{2}\, \Re\{ \mathrm{tr}\{\, \boldsymbol{Z}_2^\dagger\, \boldsymbol{Z}_1 \} \}$$

where $\boldsymbol{Z}_1, \boldsymbol{Z}_2 \in T_{\boldsymbol{X}}$. A local parameterization $h : \Gamma \to \mathcal{U}(d)$ around any point $\boldsymbol{X} \in \mathcal{U}(d)$ maps an open subset Γ of the tangent space $T_{\boldsymbol{X}}$ to the manifold $\mathcal{U}(d)$. Under such parameterization, the tangent space should be sufficient to locally characterize points on the manifold. In rigorous terms, this translates to: for any point $\boldsymbol{Y} \in \mathcal{U}(d)$ sufficiently close to $\boldsymbol{X}$, there must exist some $\boldsymbol{Z} \in T_{\boldsymbol{X}}$ such that $\boldsymbol{Y} = h(\boldsymbol{Z})$. This article uses the parameterization $h(\boldsymbol{Z}) \triangleq \mathcal{P}_\mathcal{U}(\boldsymbol{X} + \boldsymbol{Z})$. Associated with the local parameterization is a local cost function $g : T_{\boldsymbol{X}} \to \mathbb{R}$. In particular, we set $g(\boldsymbol{Z}) \triangleq f \circ h(\boldsymbol{Z}) = f\big(\mathcal{P}_\mathcal{U}(\boldsymbol{X} + \boldsymbol{Z})\big)$. Steepest ascent on the unitary group selects the steepest ascent direction to maximize the local cost $g(\boldsymbol{Z})$. The following proposition provides an analytic expression for this direction in the general case and is a direct corollary of [95, Th. 14], which computes the direction of steepest *descent* for the local cost g. Noting that the direction of steepest *ascent* for the cost g is the same as the direction of steepest *descent* for $-g$ proves this proposition.

Proposition 3: Given an arbitrary cost function $f : \mathcal{U}(d) \to \mathbb{R}$ and the local cost function $g(\boldsymbol{Z}) = f\big(\mathcal{P}_\mathcal{U}(\boldsymbol{X} + \boldsymbol{Z})\big)$ about any point $\boldsymbol{X} \in \mathcal{U}(d)$, the direction of steepest ascent of the function g about the point $\boldsymbol{Z} = \boldsymbol{0}$ under the inner product $\langle \cdot, \cdot \rangle_T$ is given by

$$\boldsymbol{G} = \boldsymbol{D}_X - \boldsymbol{X}\boldsymbol{D}_X^\dagger \boldsymbol{X}$$

where $\boldsymbol{D}_X$ is any matrix on $\mathbb{C}^{d\times d}$ satisfying

$$f(\boldsymbol{X} + \boldsymbol{Z}) = f(\boldsymbol{X}) + \Re\big\{\mathrm{tr}\big\{\boldsymbol{Z}^\dagger \boldsymbol{D}_X\big\}\big\} + \mathcal{O}\big(\|\boldsymbol{Z}\|_\mathrm{F}^2\big)$$

for all $\boldsymbol{Z} \in T_{\boldsymbol{X}}(d)$. □

Using Proposition 3, we present a closed-form expression for the steepest ascent direction of the local cost function when the global cost is given by (6). This will be used in our iterative algorithm for solving Problem 2. The proof of the following proposition is presented in Appendix A.

Algorithm 1: Steepest Ascent for Maximizing (6) on the Unitary Group:

1: **Input:** $\epsilon > 0$
2: **Initialize:** Set $\boldsymbol{X} = \boldsymbol{I}_d$, $j = 1$, and $\gamma \triangleq 1$.
3: **while** $\langle \boldsymbol{G}, \boldsymbol{G} \rangle_T \geqslant \epsilon$ **and** $j \leqslant C$ **do**
4: Compute the steepest ascent direction:
$$\boldsymbol{G} = 2\sum_{n=1}^{N_\mathrm{i}} \boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n - \boldsymbol{X}\boldsymbol{A}_n \boldsymbol{X}^\dagger \boldsymbol{B}_n \boldsymbol{X}$$
5: **while** $f_\mathrm{B}(\mathcal{P}_\mathcal{U}(\boldsymbol{X} + 2\gamma\boldsymbol{G})) - f(\boldsymbol{X}) \geqslant \gamma \langle \boldsymbol{G}, \boldsymbol{G} \rangle_T$ **do**
6: Set $\gamma = 2\gamma$
7: **while** $f_\mathrm{B}(\mathcal{P}_\mathcal{U}(\boldsymbol{X} + \gamma\boldsymbol{G})) - f(\boldsymbol{X}) < \frac{1}{2}\gamma \langle \boldsymbol{G}, \boldsymbol{G} \rangle_T$ **do**
8: Set $\gamma = \gamma/2$
9: Set $\boldsymbol{X} = \mathcal{P}_\mathcal{U}(\boldsymbol{X} + \gamma\boldsymbol{G})$
10: Set $j = j + 1$
11: **Return:** $\boldsymbol{X}$

Proposition 4: Let $f_\mathrm{B} : \mathcal{U}(d) \to \mathbb{R}$ be the Brockett cost function (6). Moreover, for any $\boldsymbol{X} \in \mathcal{U}(d)$, define the local cost $g : T_{\boldsymbol{X}} \to \mathbb{R}$ to be $g(\boldsymbol{Z}) \triangleq f_\mathrm{B}\big(\mathcal{P}_\mathcal{U}\big(\boldsymbol{X} + \boldsymbol{Z}\big)\big)$. Then, the direction of steepest ascent in the sense of Proposition 3 is

$$\boldsymbol{G} = 2\sum_{n=1}^{N_\mathrm{i}} \boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n - \boldsymbol{X}\boldsymbol{A}_n \boldsymbol{X}^\dagger \boldsymbol{B}_n \boldsymbol{X}. \tag{10}$$

□

Exploiting this result, an iterative method for maximizing (6) may be formalized. The strategy is fully characterized by Algorithm 1. The program starts with the initial guess $\boldsymbol{X} = \boldsymbol{I}_d$ and a step size $\gamma \triangleq 1$. At each iteration, the direction of steepest ascent is computed according to (10). The algorithm is stopped if one of two criteria are met: 1) if, for some (small) user-defined threshold $\epsilon > 0$, the inequality $\langle \boldsymbol{G}, \boldsymbol{G} \rangle_T < \epsilon$ holds, i.e., the direction of steepest ascent is sufficiently small. This is the case when $\boldsymbol{X}$ approaches a stationary point or 2) if the number of iterations exceeds some maximum threshold $C \in \mathbb{N}$. Steps 5–8 of the algorithm adaptively update the step size according to the Armijo criteria, which ensures that the algorithm converges to a stationary point [95]. In step 9, the variable $\boldsymbol{X}$ is updated according to the direction of steepest ascent and the step size. As mentioned at the beginning of this section, Algorithm 1 be used to solve Problem 2 by simply replacing $\boldsymbol{A}_n \triangleq \hat{\boldsymbol{\Xi}}_n(t_0)$ and $\boldsymbol{B}_n \triangleq \hat{\hat{\boldsymbol{\Xi}}}_n(t_\mathrm{f})$.

Recall that N_i is the number of state pairs used in QPT and d is the dimension of the quantum system. Ignoring the complexity of computing the Armijo step length, which may be replaced with a constant step size, the worst-case computational complexity of our optimization procedure is

$\mathcal{O}(CN_\mathrm{i}d^3)$. This is due to the matrix multiplication and addition performed in Step 4 and the maximum number of iterations C. Tuning the hyperparameters N_i and C gives the practitioner control over the computational complexity of our approach at the potential cost of accuracy. The current state-of-the-art QPT method for closed systems (see[40]) has computational complexity $\mathcal{O}(d^{12})$ (or $\mathcal{O}(d^6)$ in a special case).

In terms of memory, our method has space complexity $\mathcal{O}(N_\mathrm{i}d^2)$, which is determined by the N_i matrices of dimension d^2 that must be stored. The state-of-the-art (see[40]) has space complexity of $\mathcal{O}(d^8)$ (or $\mathcal{O}(d^4)$ in a special case). The number N_i of input–output state pairs used in our QPT method may be chosen by the designer. However, one can expect to produce better results with more samples. In our numerical experiments, we find that $N_\mathrm{i} = d$ is a good choice.

Remark 3: The choice of N_i is influenced by both the theoretical and practical limitations. In [99], the use of "unitarily informationally complete (UIC)" quantum states for QPT was studied. It was shown that as few as d pure states or two mixed states are needed to uniquely identify a unitary process. In practice, one may wish to use more states to overcome noise and error in QST estimates. We have not studied the use of UIC quantum states in our QPT approach; however, this is an interesting research direction. □

Remark 4: The computational complexity studied above is that of the QPT algorithm alone. In practical scenarios, one should also consider the complexity of QST since QPT is dependent on state estimates. Popular QST algorithms proposed in [36], [100], and [101] have complexity $\mathcal{O}(d^4)$ and $\mathcal{O}(d^3)$. Using such algorithms, the computational burden resides with QPT. □

D. NOTE ON THE SPECIAL CASE WHEN $N_\mathrm{i} = 1$

In the special case where QPT is performed using a single pair of initial and final states, i.e., $N_\mathrm{i} = 1$, we provide a closed-form solution and error bounds on the recovered unitary operator $\hat{\boldsymbol{U}}(t_\mathrm{f})$. To begin, note that the spectral theorem states that the eigendecomposition of the Hermitian matrix $\hat{\boldsymbol{\Xi}}(t_0)$ may be taken to be

$$\hat{\boldsymbol{\Xi}}(t_0) = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^\dagger \tag{11}$$

where $\boldsymbol{\Lambda} \in \mathbb{R}^{d\times d}$ is the diagonal matrix of real eigenvalues and the columns of $\boldsymbol{V} \in \mathbb{C}^{d\times d}$ consist of orthonormal eigenvectors. Similarly, $\hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})$ is taken to be

$$\hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f}) = \hat{\boldsymbol{Q}}\hat{\boldsymbol{\Lambda}}\hat{\boldsymbol{Q}}^\dagger \tag{12}$$

where $\hat{\boldsymbol{\Lambda}}, \boldsymbol{Q} \in \mathbb{C}^{d\times d}$ are matrices of eigenvalues and orthonormal eigenvectors, respectively. It is always assumed that eigenvalues are placed in *nondecreasing order*. With this setup, we are able to provide a closed-form solution for calculating the maximizer $\hat{\boldsymbol{U}}$ for Problem 2. The proof of the following proposition is postponed to Appendix B.

Proposition 5: When $N_\mathrm{i} = 1$, $\hat{\boldsymbol{U}} = \hat{\boldsymbol{Q}}\boldsymbol{V}^\dagger$ is a solution to Problems 1, 2, and 3. □

Now that we understand the closed-form solution to Problem 2 in this special case, it is possible to bound the error incurred by our approach. For an arbitrary vector of phases $\boldsymbol{\phi} \triangleq [\phi_1\ \phi_2\ \cdots\ \phi_d]^T \in \mathbb{R}^d$, define a diagonal unitary matrix $\boldsymbol{\Delta}_{\boldsymbol{\phi}}$ as

$$\boldsymbol{\Delta}_{\boldsymbol{\phi}} \triangleq \mathrm{diag}\left(e^{\imath\phi_1}, e^{\imath\phi_2}, \ldots, e^{\imath\phi_d}\right). \tag{13}$$

The following proposition gives an upper bound on the unitary dynamics recovered in terms of the amount of error in the QST process to obtain the estimate $\hat{\hat{\boldsymbol{\Xi}}}(t_\mathrm{f})$. The proof of the following proposition is postponed to Appendix C.

Proposition 6: Let $\hat{\boldsymbol{\Xi}}(t_0)$ have simple eigenvalues. Then, let $\hat{\boldsymbol{U}}$ be calculated according to Proposition 5 and let $\boldsymbol{E}$ be the state estimation error defined in (9). Then, the following error bound holds:

$$\min_{\boldsymbol{\phi}} \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{U}} - \boldsymbol{U}\right\|_\mathrm{F}^2 \leqslant \max_{i\neq j} \frac{8d}{|\lambda_i - \lambda_j|^2}\|\boldsymbol{E}\|^2$$

where λ_i denotes the ith eigenvalue of $\hat{\boldsymbol{\Xi}}(t_0)$. □

Remark 5: A matrix is said to have *simple* eigenvalues if none of its eigenvalues are repeated. For $\hat{\boldsymbol{\Xi}}(t_0)$ to possess this property, it must necessarily be a mixed state. Interestingly, mixed quantum states with this property occur naturally in certain quantum systems, e.g., see the model in [74, Sec. IV.A.3]. Moreover, since $\hat{\boldsymbol{\Xi}}(t_0)$ is chosen by the designer, it may be designed so that its eigenvalues tighten the bound stated in the previous proposition. □

Proposition 6 shows that from a *single* initial condition, it is possible to identify $\boldsymbol{U}$ up to a diagonal unitary matrix $\boldsymbol{\Delta}_{\boldsymbol{\phi}}$. This is because the solution to Problem 2 given by Proposition 5 remains a solution when multiplied by any diagonal unitary matrix. In the identification setting studied in [43], which uses population measurements (which are important in various applications including physical chemistry [27]), it is only possible to identify unitary dynamics up to a diagonal unitary matrix. However, in general, it may be possible to uniquely identify a quantum system's dynamics up to a global phase using other measurements, such as positive operator valued measurements (POVMs) [102]. Nonetheless, since the closed-form solution presented in Proposition 5 is nonunique, more than one set of experimental states is necessary to identify the quantum system's dynamics using our QPT formulation. The iterative solution given for the general case of $N_\mathrm{i} \geqslant 1$ can achieve this in a scalable and efficient manner. Section V, which is devoted to numerical experimentation, gives some examples of the number of states, N_i, needed for obtaining "good" results.

III. HAMILTONIAN LEARNING

This section will discuss how to use the results of the proposed QPT algorithm (see Section II) to estimate both internal (see Section III-A) and control (see Section III-B) Hamiltonians of a closed quantum system.

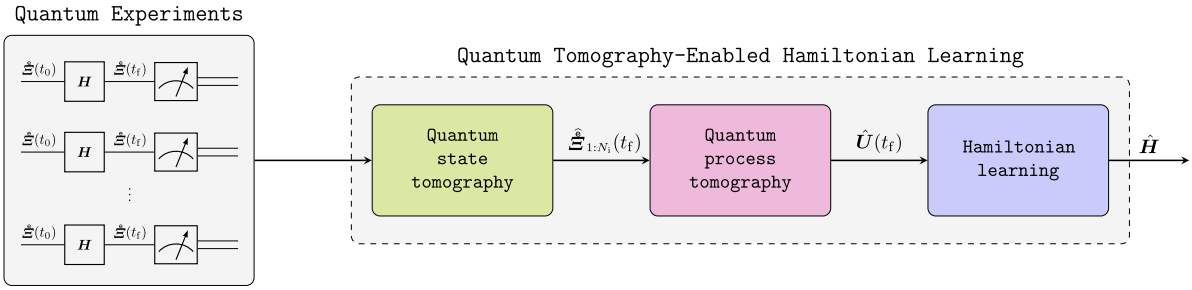

FIGURE 3. Summary of QT-enabled QHL. State estimates from QST drive QPT, which estimates a unitary operator that maps the initial state of the quantum system to those measured at the termination of the experiment. QHL uses QPT's result to learn an estimate of the system Hamiltonian.

A. RECOVERING THE INTERNAL HAMILTONIAN

Once we have estimated $\hat{\boldsymbol{U}}(t_\mathrm{f})$ using QPT, we seek to recover an estimate $\hat{\boldsymbol{H}}$ of the internal Hamiltonian. All the unitary matrices are diagonalizable by a unitary transformation. Suppose that the eigendecomposition of $\hat{\boldsymbol{U}}(t_\mathrm{f})$ is

$$\hat{\boldsymbol{U}}(t_\mathrm{f}) = \hat{\boldsymbol{V}}\hat{\boldsymbol{\Theta}}\hat{\boldsymbol{V}}^\dagger .$$

Here, $\hat{\boldsymbol{\Theta}}, \hat{\boldsymbol{V}} \in \mathbb{C}^{d\times d}$ are the diagonal matrix of eigenvalues and the matrix of orthonormal eigenvectors, respectively. Then, according to the relationship $\hat{\boldsymbol{U}}(t) = e^{-\imath\hat{\boldsymbol{H}}t}$, it follows that

$$\begin{aligned}\hat{\boldsymbol{H}} &= \frac{\imath}{t_\mathrm{f}} \log\left(\hat{\boldsymbol{U}}(t_\mathrm{f})\right) \\ &= \frac{\imath}{t_\mathrm{f}} \hat{\boldsymbol{V}} \log(\hat{\boldsymbol{\Theta}}) \hat{\boldsymbol{V}}^\dagger \end{aligned} \tag{14}$$

where $\log(\hat{\boldsymbol{\Theta}})$ is just the usual logarithm applied entrywise to the diagonal matrix $\hat{\boldsymbol{\Theta}}$. In order for $\hat{\boldsymbol{H}}$ to be a unique estimate, we must limit the duration of the experiment. This well-known result is in the spirit of the Nyquist–Shannon theorem for digitally sampling classical signals and is due to the fact that the logarithm of a complex variable is nonunique.

Proposition 7 ([40] and [103]): To uniquely identify $\boldsymbol{H}$ from the unitary operator $\boldsymbol{U}(t_\mathrm{f})$, the duration t_f of the experiment should satisfy

$$0 < t_\mathrm{f} < \frac{\pi}{\mu_\mathrm{max} - \mu_\mathrm{min}} \tag{15}$$

where μ_max and μ_min are the maximum and minimum eigenvalues of $\boldsymbol{H}$, respectively. □

While, in general, the spectrum of $\boldsymbol{H}$ is unknown, choosing a sufficiently short sampling time will ensure that this condition is satisfied. If such a sampling time is employed and the system Hamiltonian is estimated using the approach presented above, it was shown in [40] that the following bound holds:

$$\left\|\boldsymbol{H} - \hat{\boldsymbol{H}}\right\|_\mathrm{F} \leqslant \frac{\pi}{2t_\mathrm{f}} \left\|\boldsymbol{U} - \hat{\boldsymbol{U}}\right\|_\mathrm{F} . \tag{16}$$

Hence, bounding the error in the estimate of $\boldsymbol{U}$ bounds the error in the estimate of $\boldsymbol{H}$.

With these results in hand, the time has come to present the proposed QT-enabled QHL algorithm for identifying an internal Hamiltonian. This technique utilizes both QST and QPT, and a visual depiction is given in Fig. 3. The process is also summarized in Algorithm 2. A set of prepared states are given to the quantum system and are measured at time t_f. Through repeated measurement across multiple experiments, QST estimates the state of the quantum system at the termination of each experiment and returns a matrix of estimates $\hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$. QPT uses these state estimates to solve Problem 2 and infer a unitary operator $\hat{\boldsymbol{U}}(t_\mathrm{f})$ that best maps the prepared input states to the experimentally measured output states. Using the inferred unitary map, QHL is used to reconstruct an estimate $\hat{\boldsymbol{H}}$ of the system Hamiltonian.

Algorithm 2: QT-Enabled QHL:

1: **Input:** $\hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_0)$
2: Perform QST to estimate the final states: $\hat{\hat{\boldsymbol{\Xi}}}_{1:N_\mathrm{i}}(t_\mathrm{f})$
3: Use Algorithm 1 to solve Problem 2 and produce $\hat{\boldsymbol{U}}(t_\mathrm{f})$
4: Estimate $\hat{\boldsymbol{H}}$ according to (14)
5: **Return:** $\hat{\boldsymbol{H}}$

Before concluding this section, an inequality relating the error in QST to the error in QHL is in order. In particular, we consider the case where $N_\mathrm{i} = 1$, i.e., a single pair of states is used in QPT. For any given $\boldsymbol{\phi}$, denote $\hat{\boldsymbol{U}}_{\boldsymbol{\phi}} \triangleq \boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{U}}$ and let $\hat{\boldsymbol{H}}_{\boldsymbol{\phi}}$ be the Hamiltonian inferred from $\hat{\boldsymbol{U}}_{\boldsymbol{\phi}}$ according to (14). Then, Propositions 6 and 7 can be combined with (16) to produce the following result.

Proposition 8: Under the assumptions of Proposition 6 and when the sampling time t_f satisfies (15), the following error bound holds:

$$\min_{\boldsymbol{\phi}} \left\|\boldsymbol{H} - \hat{\boldsymbol{H}}_{\boldsymbol{\phi}}\right\|_\mathrm{F} \leqslant \max_{i\neq j} \frac{4\pi d}{t_\mathrm{f}\,|\lambda_i - \lambda_j|^2} \|\boldsymbol{E}\|^2.$$

□

B. LEARNING THE CONTROL HAMILTONIANS

In the case of a controlled quantum system, we consider a time-varying Hamiltonian, which may be decomposed as

$$\boldsymbol{H}(t) = \boldsymbol{H}_0 + \boldsymbol{H}_\mathrm{c}(t) \tag{17}$$

Algorithm 3: QT-Enabled CQHL:

Input: $\{c_l : l = 1, 2, \ldots, N_c\}$
Initialize: $u_l = 0$ for all $l \in \{1, 2, \ldots, N_c\}$
Learn $\hat{\boldsymbol{H}}_0$ using Algorithm 2
for $l = 1, 2, \ldots, N_c$ **do**
  Set $u_l = c_l$
  Use Algorithm 2 to learn $\hat{\boldsymbol{H}} = [\hat{\boldsymbol{H}}_0 + c_l \hat{\boldsymbol{H}}_l]$
  Estimate $\hat{\boldsymbol{H}}_l = c_l^{-1}[\hat{\boldsymbol{H}} - \hat{\boldsymbol{H}}_0]$
  Set $u_l = 0$
Return: $\{\hat{\boldsymbol{H}}_0, \hat{\boldsymbol{H}}_1, \ldots, \hat{\boldsymbol{H}}_{N_c}\}$

where $\boldsymbol{H}_0$ is the *internal* Hamiltonian that appears in (1) and $\boldsymbol{H}_c(t)$ is the *control* Hamiltonian. Furthermore, we assume that the control Hamiltonian has the following structure:

$$\boldsymbol{H}_c(t) = \sum_{l=1}^{N_c} u_l(t) \boldsymbol{H}_l$$

where $N_c \in \mathbb{N}$, $\boldsymbol{H}_l$ is a Hermitian matrix for each $l \in \{1, 2, \ldots, N_c\}$, and $u_l(t) \in \mathbb{R}$ is a control input.[3] Each $\boldsymbol{H}_l$ is known as an *interaction Hamiltonian*, and it describes the effect of the lth control field on the quantum system. We emphasize that each $\boldsymbol{H}_l$ is time independent and that the only temporal dependence is in the real control input $u_l(t)$. The Schrödinger equation becomes

$$\imath \frac{\partial}{\partial t} |\psi(t)\rangle = \underbrace{\left[\boldsymbol{H}_0 + \sum_{l=1}^{N_c} u_l(t) \boldsymbol{H}_l\right]}_{\boldsymbol{H}(t)} |\psi(t)\rangle \tag{18}$$

and the von Neumann equation becomes

$$\imath \frac{\partial}{\partial t} \boldsymbol{\Xi} = \left[\!\!\left[\boldsymbol{H}_0 + \sum_{l=1}^{N_c} u_l(t) \boldsymbol{H}_l \,,\, \boldsymbol{\Xi} \right]\!\!\right]_- . \tag{19}$$

During the learning phase, i.e., for $t \in [0, t_f]$, it is assumed that the control $u_l(t)$ can be placed under a zero-order hold (ZOH), which removes the *time dependence* in (18) and (19) and allows Algorithm 2 to be applied. This is outlined in Algorithm 3, which we call QT-enabled CQHL. The premise is to learn the Hamiltonians one by one via selectively turning on the control $u_l(t)$ and placing it under a ZOH for the duration of the sampling time. The value of $u_l(t)$ under the ZOH hold is denoted $c_l \in \mathbb{R}$. This parameter may be chosen by the designer and is called a *probing control input*. First, the internal Hamiltonian is learned according to Algorithm 1, while all the control inputs are set to zero. Then, the first control input is set to c_1 (the other inputs remain zero) and the composite (internal and control) Hamiltonian is estimated by repeating Algorithm 1. To recover the control Hamiltonian, the estimated internal Hamiltonian from step one is subtracted from the composite Hamiltonian learned in the second step. This process is repeated for all $l \in \{1, 2, \ldots, N_c\}$. At each step l, the sampling time t_f can be arbitrarily chosen (and possibly different across steps) as long as it satisfies Proposition 7 for the composite Hamiltonian, $\boldsymbol{H} = \boldsymbol{H}_0 + c_l \boldsymbol{H}_l$. The following proposition gives an error bound on the recovery of the control Hamiltonians when Algorithm 3 is used.

Proposition 9: Let $\hat{\boldsymbol{H}}_0$ and $\hat{\boldsymbol{H}}_l$, for $l = 1, 2, \ldots, N_c$, be estimates of the internal and lth control Hamiltonian produced by Algorithm 3. Furthermore, suppose that $\hat{\boldsymbol{U}}_0$ is the unitary operator produced by QPT for the internal dynamics and $\hat{\boldsymbol{U}}$ is the unitary operator produced by QPT for the dynamics governed by the composite Hamiltonian $\boldsymbol{H} = \boldsymbol{H}_0 + c_l \boldsymbol{H}_l$. Finally, let the sampling time t_f used by QPT satisfy Proposition 7 for both the internal and composite Hamiltonians. Then, the following error bound on the lth control Hamiltonian holds:

$$\left\|\boldsymbol{H}_l - \hat{\boldsymbol{H}}_l\right\|_F \leqslant \frac{\pi}{2t_f |c_k|} \left(\left\|\boldsymbol{U} - \hat{\boldsymbol{U}}\right\|_F + \left\|\boldsymbol{U}_0 - \hat{\boldsymbol{U}}_0\right\|_F \right). \tag{20}$$

□

Remark 6: Consider the special case where QPT uses a single pair of states to estimate both $\hat{\boldsymbol{U}}$ and $\hat{\boldsymbol{U}}_0$. A corollary to Propositions 8 and 9 could be given, which bounds the error in our estimates of the control Hamiltonians in terms of the error in QST. □

IV. DATA-DRIVEN MPQC

Once the dynamics (Hamiltonians) of the closed quantum system are learned, one may turn attention to controlling the system's behavior. MPQC is inspired by MPC for classical systems. Traditionally, MPC works by computing optimal controls over a short prediction horizon based on a model of the system's dynamics [104]. It then implements only the first control in this optimal sequence. After executing this action, it measures the state of the system, recomputes an optimal control sequence based on its new knowledge of the state, executes the first control action in this sequence, and repeats. Unlike classical systems, state measurements of quantum systems are not always available as measurement may cause decoherence of their states. Hence, MPQC performs MPC in a closed-loop computer simulation of the quantum system where its state can be fully known [see Fig. 2(a)]. Once a control sequence is computed, it can be delivered to a physical quantum system in an open-loop manner so as to preserve the coherence of the quantum state [see Fig. 2(b)]. Here, for the first time, we propose the data-driven MPQC framework, where the model of the quantum system is inferred by the QHL algorithm proposed in the previous section.

This article is concerned with controlling a closed quantum system whose dynamics are of the form (18) or (19). All the systems studied in the rest of this article are assumed to be *controllable* [28], which ensures that any control problem considered is well posed. For ease of exposition, the theory of MPQC will be presented in terms of pure states with comments on the related formulation for mixed states when necessary.

[3] It should be noted that the control input $u_l(t)$ and the unitary propagator $\boldsymbol{U}(t)$ are separate entities despite the notational similarity.

TABLE 4. Sample Cost Functions for MPQC

	Name:	For pure states:	For mixed states:
g_1:	Occupation of forbidden state	$\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \left\lvert \langle\boldsymbol{\psi}_\mathrm{bad}\vert\boldsymbol{\psi}(t_s)\rangle \right\rvert^2$	$\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} -\left\lVert\boldsymbol{\Xi}_\mathrm{bad}-\boldsymbol{\Xi}(t_s)\right\rVert_\mathrm{F}^2$
g_2:	Distance to target state	$1-\left\lvert \langle\boldsymbol{\psi}_\mathrm{d}\vert\boldsymbol{\psi}(t_{k+K_\mathrm{p}})\rangle \right\rvert^2$	$\left\lVert\boldsymbol{\Xi}_\mathrm{d}-\boldsymbol{\Xi}(t_{k+K_\mathrm{p}})\right\rVert_\mathrm{F}^2$
g_3:	α-weighted average distance to target state	$1-\frac{1}{\sum_s \alpha_s}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \alpha_s\left\lvert \langle\boldsymbol{\psi}_\mathrm{d}\vert\boldsymbol{\psi}(t_s)\rangle \right\rvert^2$	$\frac{1}{\sum_s \alpha_s}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \alpha_s\left\lVert\boldsymbol{\Xi}_\mathrm{d}-\boldsymbol{\Xi}(t_s)\right\rVert_\mathrm{F}^2$
g_4:	β-minimum norm controller	$\left(1-\left\lvert \langle\boldsymbol{\psi}_\mathrm{d}\vert\boldsymbol{\psi}(t_{k+K_\mathrm{p}})\rangle \right\rvert^2\right)$ $+\frac{\beta}{K_\mathrm{p}u_\mathrm{max}}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \left\lVert\boldsymbol{u}(t_s)\right\rVert^2$	$\left\lVert\boldsymbol{\Xi}_\mathrm{d}-\boldsymbol{\Xi}(t_{k+K_\mathrm{p}})\right\rVert_\mathrm{F}^2$ $+\frac{\beta}{K_\mathrm{p}u_\mathrm{max}}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \left\lVert\boldsymbol{u}(t_s)\right\rVert^2$
g_5:	(α,β)-minimum norm controller	$1-\frac{1}{\sum_s \alpha_s}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \alpha_s\left\lvert \langle\boldsymbol{\psi}_\mathrm{d}\vert\boldsymbol{\psi}(t_s)\rangle \right\rvert^2$ $+\frac{\beta}{K_\mathrm{p}u_\mathrm{max}}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \left\lVert\boldsymbol{u}(t_s)\right\rVert^2$	$\frac{1}{\sum_s \alpha_s}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \alpha_s\left\lVert\boldsymbol{\Xi}_\mathrm{d}-\boldsymbol{\Xi}(t_s)\right\rVert_\mathrm{F}^2$ $+\frac{\beta}{K_\mathrm{p}u_\mathrm{max}}\sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} \left\lVert\boldsymbol{u}(t_s)\right\rVert^2$

The goal of the controller is to drive $|\boldsymbol{\psi}(t)\rangle$ to a desired state, $|\boldsymbol{\psi}_d\rangle$, via an appropriate control input while minimizing a cost function. MPQC works by placing the control vector $\boldsymbol{u}(t)$ under a ZOH on each time interval of length $\Delta_t > 0$. If $t_k = k\Delta_t$, the ZOH is realized as $\boldsymbol{u}(t) \equiv \boldsymbol{u}(t_k)$ for all $t \in [t_k, t_k + \Delta_t)$. Let $|\boldsymbol{\psi}(t_k)\rangle$ denote the state of the system at time step $k \in \mathbb{N}$ such that $|\boldsymbol{\psi}(t_0)\rangle$ is the initial state at time $t_0 = 0$. The lth element of $\boldsymbol{u}(t_k)$ is denoted by $u_l(t_k)$. Finally, $\boldsymbol{H}(t_k)$ is the system Hamiltonian at time step k, i.e.,

$$\boldsymbol{H}(t_k) = \boldsymbol{H}_0 + \sum_{l=1}^{N_\mathrm{c}} u_l(t_k)\boldsymbol{H}_l\,.$$

The discrete-time dynamics that advance state $|\boldsymbol{\psi}(t_k)\rangle$ to $|\boldsymbol{\psi}(t_{k+1})\rangle$ are given by the solution to the Schrödinger equation over the time interval $[t_k, t_k + \Delta_t)$:

$$\big|\boldsymbol{\psi}(t_{k+1})\big\rangle = e^{-\imath\boldsymbol{H}(t_k)\Delta_t}\big|\boldsymbol{\psi}(t_k)\big\rangle\,. \tag{21}$$

Define an arbitrary cost as $g(|\boldsymbol{\psi}(t_k)\rangle, \boldsymbol{u}(t_k), t_k)$ where $g : \mathbb{C}^d \times \mathbb{R}^{N_\mathrm{c}} \times \mathbb{R} \to \mathbb{R}$. Examples of such a cost are given in Table 4 and discussed later. In many cases, the cost is chosen independently of time. Let $K_\mathrm{p}, K_\mathrm{c} \in \mathbb{N}$ denote the prediction horizon and control horizon, respectively, for the controller such that $K_\mathrm{c} \leqslant K_\mathrm{p}$. During closed-loop quantum simulation, at each time step t_k, MPQC attempts to minimize the total cost over the prediction horizon K_p by manipulating the control $\boldsymbol{u}(t_k)$ over the control horizon K_c. The optimization problem performed at time step t_k is given in the following:

$$\begin{aligned}\mathscr{P}_4:\quad \underset{\boldsymbol{u}(t_s):\, s\in\mathcal{Z}_{k:k+K_\mathrm{c}}}{\text{minimize}} \quad & \sum_{s\in\mathcal{Z}_{k:k+K_\mathrm{p}}} g\left(\big|\boldsymbol{\psi}(t_s)\big\rangle, \boldsymbol{u}(t_s), t_s\right)\\ \text{subject to} \quad & \big|\boldsymbol{\psi}(t_{s+1})\big\rangle = e^{-\imath\boldsymbol{H}(t_s)\Delta_t}\big|\boldsymbol{\psi}(t_s)\big\rangle\\ & \left\lVert\boldsymbol{u}(t_s)\right\rVert_\infty \leqslant u_\mathrm{max}, \quad \forall s \in \mathcal{Z}_{k:k+K_\mathrm{p}}\end{aligned}$$

where $u_\mathrm{max} > 0$ is the maximum control magnitude allowed. If $K_\mathrm{c} < K_\mathrm{p}$, then for all $s > k + K_\mathrm{c}$, MPQC sets $\boldsymbol{u}(t_s) \equiv \boldsymbol{u}(t_{k+K_\mathrm{c}})$. That is, MPQC will only optimize controls over the horizon K_c—after that, allowing the control to remain constant—while considering the cost incurred over the longer horizon K_p. After the optimal sequence $\{\boldsymbol{u}(t_k), \boldsymbol{u}(t_{k+1}), \ldots, \boldsymbol{u}(t_{k+K_\mathrm{c}})\}$ is computed, only the control action $\boldsymbol{u}(t_k)$ is recorded and given to the simulator, which then returns the next state according to (21). This is repeated until a control sequence of the desired length is constructed or the cost g becomes smaller than a designed threshold. At the termination of this simulation, the sequence of recorded control actions is given to the physical quantum system in an open-loop manner.

The choice of cost function g determines the types of control generated by MPQC, and Table 4 outlines several different choices. In all cases, a summation over s in Table 4 is the sum over all $s \in \mathcal{Z}_{k:k+K_\mathrm{p}}$. Cost g_1 penalizes the system's state for occupying a forbidden state $|\boldsymbol{\psi}_\mathrm{bad}\rangle$. Costs g_2 and g_3 both reward the system's state for tending toward the desired state $|\boldsymbol{\psi}_\mathrm{d}\rangle$; however, g_3 promotes quick convergence by considering the distance of the state $|\boldsymbol{\psi}(t_s)\rangle$ from the desired state over the entire prediction interval. Cost g_2 is the infidelity between the desired state and the system's state at the end of the prediction horizon. In g_3, $\{\alpha_s\}_{s=k}^{k+K_\mathrm{p}}$ is a set of positive weighting parameters that can be designed to put emphasis on different parts of the trajectory. If one would like to consider the expense of control, g_4 and g_5 are a good starting place. A controller utilizing g_4 is called a β-minimum norm controller. Function g_4 adds a penalty for the control expended over the prediction horizon K_p to the state cost in g_2. The term $\beta/(K_\mathrm{p}u_\mathrm{max})$ is a scaling parameter, which ensures that the control cost is the same order of magnitude as the state cost. Here, $\beta \in (0, 1]$ may be chosen by the designer.

IEEE Transactions on Quantum Engineering

An (α, β)-minimum norm controller is given by optimizing with respect to g_5, which combines the α-weighted average state cost with the β-weighted control cost. This promotes quick convergence to the desired state while considering the cost of actuation.

Solving the classical MPC problem has been a subject of research interest for many years [105]. We would like to utilize these methods to solve Problem 4; however, this is not immediately feasible since the state $|\psi(t_k)\rangle$ is a complex variable, and most of these methods are designed to handle real variables. To overcome this hurdle, we make use of the isomorphism

$$|\psi\rangle \mapsto \begin{bmatrix} \Re\{|\psi\rangle\} \\ \Im\{|\psi\rangle\} \end{bmatrix}$$

between $\mathbb{C}^d$ and $\mathbb{R}^{2d}$. The MPQC formulation (Problem 4) can then be cast, for the purpose of optimization, into one involving only real variables. In this article, we make use of SQP [106], which has been used to solve nonlinear programs for classical MPC [107] and recently to design robust quantum gates [78]. SQP solvers are part of many available numerical optimization packages, including the MPC toolbox in MATLAB [108].

Remark 7: Many nonlinear programming routines, including SQP, require gradients of the cost and constraint functions with respect to states and controls. These may either be computed analytically or numerically approximated. The literature on QOC provides an exhaustive review of how these may be computed, and the reader is referred to [61], [67], and [109] for a few examples of both the analytic and approximate methods. □

Suppose one wishes to generate a control sequence of length $K \in \mathbb{N}$. If the prediction horizon and control horizon are taken to be this entire length, i.e., $K = K_\mathrm{p} = K_\mathrm{c}$, then MPQC reduces to QOC. Hence, MPQC can be considered as a generalization to QOC that allows the practitioner to simplify the optimization problem performed by breaking it into smaller pieces. This simplification is intrinsically dependent on the underlying optimization routine used. We use SQP to solve the optimal control problem. Suppose, for simplicity, that $K_\mathrm{c} = K_\mathrm{p}$ in MPQC, i.e., the control horizon and prediction horizon are the same length. The number of decision variables used by the SQP program that solves Problem 4 is given by $K_\mathrm{p}(d + N_\mathrm{c})$. At each iteration of the optimization routine, SQP solves a quadratic program, which has complexity that is cubic in the number of decision variables. That is $\mathcal{O}\big(K_\mathrm{p}^3(d + N_\mathrm{c})^3\big)$. Hence, for large values of K and $K_\mathrm{p} < K$, one can expect a speedup in the optimization problem performed at each iteration.

Remark 8: We note that there are several different notions of "QOC." The one we refer to in the prior paragraph is similar to that in which GRAPE is applied: assuming ZOH on the control input, optimize a given cost function g over the control horizon K. The differences are twofold. First, GRAPE attempts to solve this problem via gradient ascent, which is different than SQP. Second, the cost in GRAPE is a terminal cost—only the cost of the final state at step K is considered. QOC, more generally defined, could incur a cost at each time step as is done in Problem 4. □

Remark 9: The implementation of MPQC in physical scenarios is highly dependent on the application. In the case of superconducting qubits, high-fidelity arbitrary waveform generators (AWGs) have enabled the physical implementation of control signals generated by QOC algorithms such as GRAPE [68]. We, therefore, expect that MPQC could also be physically implemented by such methods. In general, we expect MPQC would perform well in any physical scenario in which QOC has been applied previously. □

So far, MPQC has been formalized. Data-driven MPQC is the same procedure; however, the Hamiltonians learned in Section II are used as the simulation's model. Once a control sequence is generated using the learned model, it can be delivered to the physical system in an open-loop fashion. This method is validated in Section V on numerical experiments, and the effect of error in the learned Hamiltonians on the performance of the proposed data-driven MPQC is analyzed.

A. GENERATING UNITARY GATES

In the previous section, MPQC was used to govern the evolution of the state of a closed quantum system. MPQC can also be used to generate control sequences for implementing arbitrary unitary gates. Take the Schrödinger operator equation

$$\imath \frac{\partial}{\partial t} \boldsymbol{U}(t) = \underbrace{\left[\boldsymbol{H}_0 + \sum_{l=1}^{N_\mathrm{c}} u_l(t) \boldsymbol{H}_l\right]}_{\boldsymbol{H}(t)} \boldsymbol{U}(t) \tag{22}$$

which describes the evolution of the propagator $\boldsymbol{U}(t)$. This allows one to ignore the initial state $|\psi(0)\rangle$ or density operator $\boldsymbol{\Xi}(0)$ and only consider what its net evolution is over time. At time zero, $\boldsymbol{U}(0) = \boldsymbol{I}$. The objective of the control may be to generate any number of unitary gates, such as an $\boldsymbol{X}$-gate, $\boldsymbol{Y}$-gate, $\boldsymbol{Z}$-gate, or Hadamard gate. In this scenario, the MPQC formulation becomes

$$\begin{aligned} \mathscr{P}_5: \quad \underset{\boldsymbol{u}(t_s):\, s \in \mathcal{Z}_{k:k+K_\mathrm{c}}}{\text{minimize}} \quad & \sum_{s \in \mathcal{Z}_{k:k+K_\mathrm{p}}} g\left(\boldsymbol{U}(t_s), \boldsymbol{u}(t_s), t_s\right) \\ \text{subject to} \quad & \boldsymbol{U}(t_{s+1}) = e^{-\imath \boldsymbol{H}(t_s)\Delta t}\, \boldsymbol{U}(t_s) \\ & \left\|\boldsymbol{u}(t_s)\right\|_\infty \leqslant u_\mathrm{max}, \quad \forall s \in \mathcal{Z}_{k:k+K_\mathrm{p}} \end{aligned}$$

where $\boldsymbol{U}(t_s)$ is the unitary operator at time step s. A commonly used cost function in this control scenario is the "infidelity" between the unitary $\boldsymbol{U}(t_s)$ and a desired unitary $\boldsymbol{U}_\mathrm{d}$ as follows:

$$g_6\left(\boldsymbol{U}(t_s)\right) \triangleq 1 - \frac{1}{d^2}\left|\mathrm{tr}\left\{\boldsymbol{U}(t_s)\, \boldsymbol{U}_\mathrm{d}^\dagger\right\}\right|^2. \tag{23}$$

This infidelity is also referred to as the "gate error."

V. CASE STUDIES

In this section, various simulation scenarios are presented to validate both the proposed QT-based QHL algorithm and data-driven MPQC. In the first experiment, the accuracy of the QHL algorithm is assessed for estimating both the internal and control Hamiltonians. Then, data-driven MPQC is tested, and its performance is compared to QOC when the *same* underlying optimization routine is used. In this section, the percent error metric

$$\varepsilon(\boldsymbol{H}, \hat{\boldsymbol{H}}) \triangleq \frac{\left\|\boldsymbol{H} - \hat{\boldsymbol{H}}\right\|_{\mathrm{F}}}{\left\|\boldsymbol{H}\right\|_{\mathrm{F}}} 100\% \tag{24}$$

is used to evaluate the recoverability properties of the proposed QHL algorithm. All the experiments were performed on a basic laptop with 8-GB RAM and an Intel® Core™ i5-7300 U CPU @ 2.60 GHz.

A. EXPERIMENT 1

In this experiment, the proposed QHL algorithm is used to infer the dynamics of a four-qubit network. For the remainder of Section V, let $\boldsymbol{I}$ denote the 2-by-2 identity matrix, $\boldsymbol{X}$ denote the Pauli-x operator, $\boldsymbol{Y}$ denote the Pauli-y operator, and $\boldsymbol{Z}$ denote the Pauli-z operator. The internal Hamiltonian of the network is

$$\begin{aligned}\boldsymbol{H}_0 = \omega_1\ \boldsymbol{X}^{(1)} + \omega_2\ \boldsymbol{X}^{(2)} + \omega_3\ \boldsymbol{X}^{(3)} + \omega_4\ \boldsymbol{X}^{(4)} \\ + \omega_{12}\,\boldsymbol{X}^{(12)} + \omega_{13}\,\boldsymbol{X}^{(13)} + \omega_{14}\,\boldsymbol{X}^{(14)} \\ + \omega_{23}\,\boldsymbol{X}^{(23)} + \omega_{24}\,\boldsymbol{X}^{(24)} + \omega_{34}\,\boldsymbol{X}^{(34)}.\end{aligned}$$

Here, $\boldsymbol{X}^{(1)} = \boldsymbol{X} \otimes \boldsymbol{I} \otimes \boldsymbol{I} \otimes \boldsymbol{I}$ represents the evolution of qubit 1 and ω_1 is the frequency of this evolution. The term $\boldsymbol{X}^{(12)} \triangleq \boldsymbol{X} \otimes \boldsymbol{X} \otimes \boldsymbol{I} \otimes \boldsymbol{I}$ represents the coupling between qubits 1 and 2 in the network and ω_{12} denotes the strength of this coupling. In general, $\boldsymbol{X}^{(i)}$ represents the independent evolution of the ith qubit and is the matrix defined by a tensor product of four matrices, which has $\boldsymbol{X}$ at the ith element of the tensor product and $\boldsymbol{I}$ at all other locations. The coupling between the ith and jth qubits is specified by $\boldsymbol{X}^{(ij)}$, which is the matrix defined as the tensor product of four matrices, which has $\boldsymbol{X}$ at both the ith and jth locations in the product and $\boldsymbol{I}$ at the other locations. The frequency of the ith qubit is denoted ω_i and the coupling frequencies between the ith and jth qubits are denoted ω_{ij}. In the following experiments, $\omega_1 = 0.1$ GHz, $\omega_2 = 0.025$ GHz, $\omega_3 = 0.075$ GHz, and $\omega_4 = 0.13$ GHz; moreover, all qubit coupling frequencies are $\omega_{ij} = 0.01$ GHz. Each qubit is controlled by two external fields. The control Hamiltonian is

$$\boldsymbol{H}_{\mathrm{c}}(t) = \sum_{l=1}^{4} u_l(t)\,\boldsymbol{X}^{(l)} + \sum_{l=5}^{8} u_l(t)\,\boldsymbol{Z}^{(l-4)}$$

where $\boldsymbol{Z}^{(l)}$ is defined analogously to $\boldsymbol{X}^{(l)}$. In this experiment, $d = 16$ is the dimension of the quantum system; however, only $N_{\mathrm{i}} = 8$ pairs of initial and final states are used to infer the system's Hamiltonians. Random initial states were generated using quantum entanglement theory laboratory (QETLAB)'s random density matrix function [110], which generates a random density matrix uniformly according to the Hilbert–Schmidt measure. The final states used in this experiment were generated by evolving the initial states forward in time according to (19). In addition, the state estimation error (9) is assumed to be a random (traceless) noise. The sampling time of the final states is chosen as $t_{\mathrm{f}} = 1$ns. The maximum number of iterations in Algorithm 1, which is called by Algorithms 2 and 3, is chosen to be $C = 15$.

One hundred numerical experiments were performed with various noise intensities. The results using Algorithm 2 to infer the internal Hamiltonian are depicted in Fig. 4(a). The recovery error of our proposed QHL method is sublinear to the level of noise in the experimental data. Next in each experiment, Algorithm 3 was used to infer the interaction Hamiltonians. For inferring each control Hamiltonian, $\boldsymbol{H}_l$, the corresponding control input was placed under a ZOH with value $c_l = 1$ GHz. Fig. 4(b) presents the error in recovering the internal and control Hamiltonians over the 100 experiments. The recovery error for the control Hamiltonians is usually higher than that of the internal Hamiltonian, which is expected due to the design of Algorithm 3. In both figures, the y-axis denotes the percent error. In all the numerical experiments, the percent error in recovering any Hamiltonian was less than 1%.

B. EXPERIMENT 2

Now, the efficacy of data-driven MPQC will be demonstrated using the same four-qubit network studied in the previous experiment. The Hamiltonians inferred in the previous experiment, which correspond to the *worst-case error*, were given to MPQC. The goal of the control is to drive the qubit network from the state $|0011\rangle$ to the consensus state $|++++\rangle$, where $|+\rangle \triangleq \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. This task is known as achieving "consensus" in a quantum network or driving the network to a "consensus state," i.e., all nodes (qubits) in the network share the same state [111]. An (α, β)-minimum norm controller was used with $\{\alpha_s\}_{s=k}^{k+K_{\mathrm{p}}} = \{1, 2, \ldots, K_{\mathrm{p}}\}$ and $\beta = 0.3$ to generate a control input offline in simulation. The control is limited with $u_{\max} = 1$GHz, the prediction horizon is $K_{\mathrm{p}} = 4$, the control horizon is $K_{\mathrm{c}} = 4$, and the discretization time is $\Delta_t = 0.05$ns. The optimized control input, which was calculated using the learned model, was given to the ground-truth model and the outcome recorded. As depicted in Fig. 5, data-driven MPQC quickly drove the qubit network to a consensus state.

C. EXPERIMENT 3

In this insightful experiment, the performance of two different MPQC controllers is analyzed in the case of a simple qubit, where the Hamiltonians are perfectly known and their efficacy is compared to QOC. The internal Hamiltonian of the system is $\boldsymbol{H}_0 = \omega\,\boldsymbol{Z}$, the qubit frequency is $\omega = 5$ GHz, and the interaction Hamiltonian is $\boldsymbol{H}_1 = \boldsymbol{Y}$. The control

IEEE Transactions on Quantum Engineering

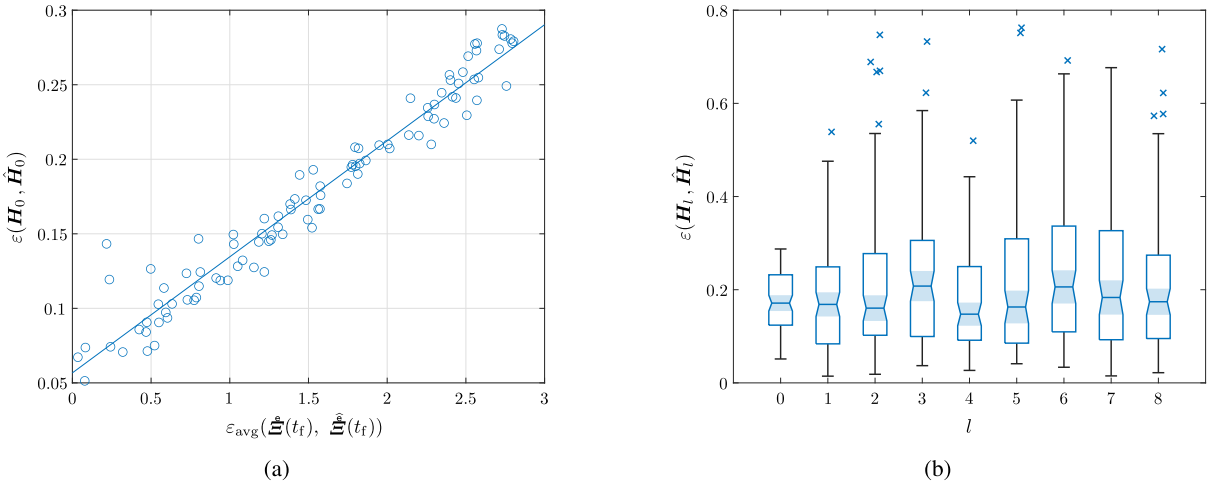

FIGURE 4. One hundred numerical experiments were performed for various noise intensities. (a) plots the percent error in recovering the internal Hamiltonian as a function of average percent noise in the experimental data (the N_f final states) used to infer the Hamiltonian. The circles represent the outcome of each of 100 experiments and the lines represent the best linear approximation of the outcomes. (b) is a box chart which summarizes the recovery error for each of the nine Hamiltonians across the 100 numerical experiments performed. The "×" marks are outlying data points, which are defined as being beyond 1.5 times the interquartile range away from their corresponding box.

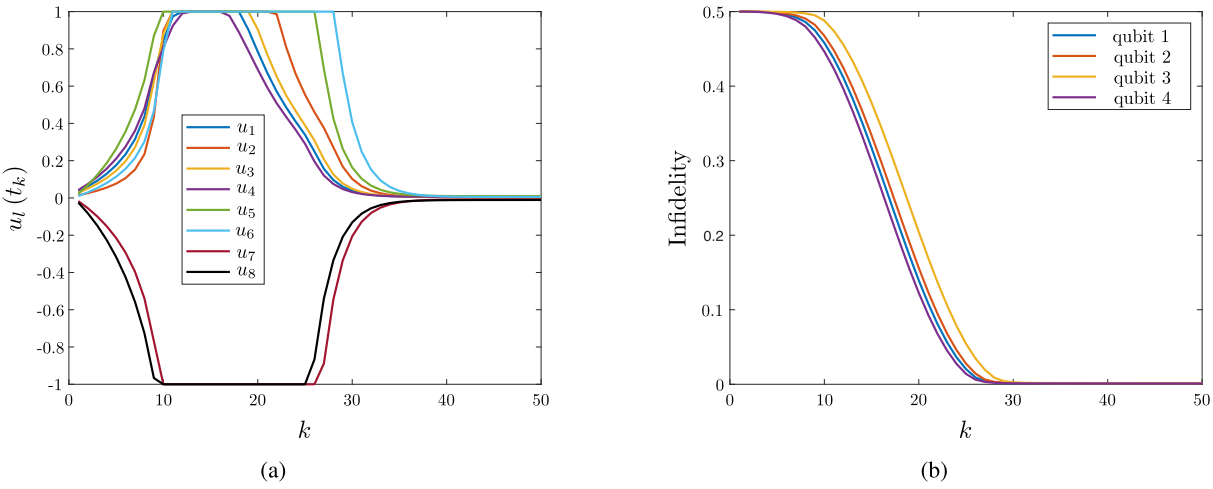

FIGURE 5. Results of Experiment 2. (a) Eight control inputs to the qubit network generated by data-driven MPQC. Then, these controls were given to the ground-truth model of the system. (b) Infidelity between the states of the four ground-truth qubits and their target states at each time step of the experiment.

is limited with $u_{\max} = 1\,\text{GHz}$ and the discretization time is $\Delta_t = 0.01\,\text{ns}$. The system is initiated in the excited state $|1\rangle$ and driven to the ground state $|0\rangle$.

First, a β-minimum norm controller with $\beta = 0.005$ is considered; second is an (α, β)-minimum norm controller with $\{\alpha_s\}_{s=k}^{k+K_\text{p}} = \{1, 2, \ldots, K_\text{p}\}$ and $\beta = 0.005$. A prediction horizon of $K_\text{p} = 5$ and a control horizon of $K_\text{c} = 2$ are used for both of these controllers. If we allow K_p and K_c to be the entire length of the control sequence, then MPQC is equivalent to the QOC problem on that interval. Using this approach, the third controller is an (α, β)-minimum norm quantum optimal controller. Fig. 6 depicts the results of this experiment. Fig. 6(a) plots the control input generated by each controller, and Fig. 6(b) depicts the evolution of the qubit's state along the Bloch sphere in response to each control input. Moreover, Fig. 7 shows the infidelity between the qubit's state and the desired state at each time step for the three scenarios.

For this simple experiment, it took the basic laptop discussed at the beginning of Section V roughly 6 s to compute each MPQC control sequence and over 31 min to solve the QOC problem despite using the same underlying optimization procedure. QOC achieved complete population transfer in the fewest time steps; however, the (α, β)-minimum norm MPQC achieved quite similar performance. The β-minimum norm MPQC took longer to achieve population transfer, which was expected. While

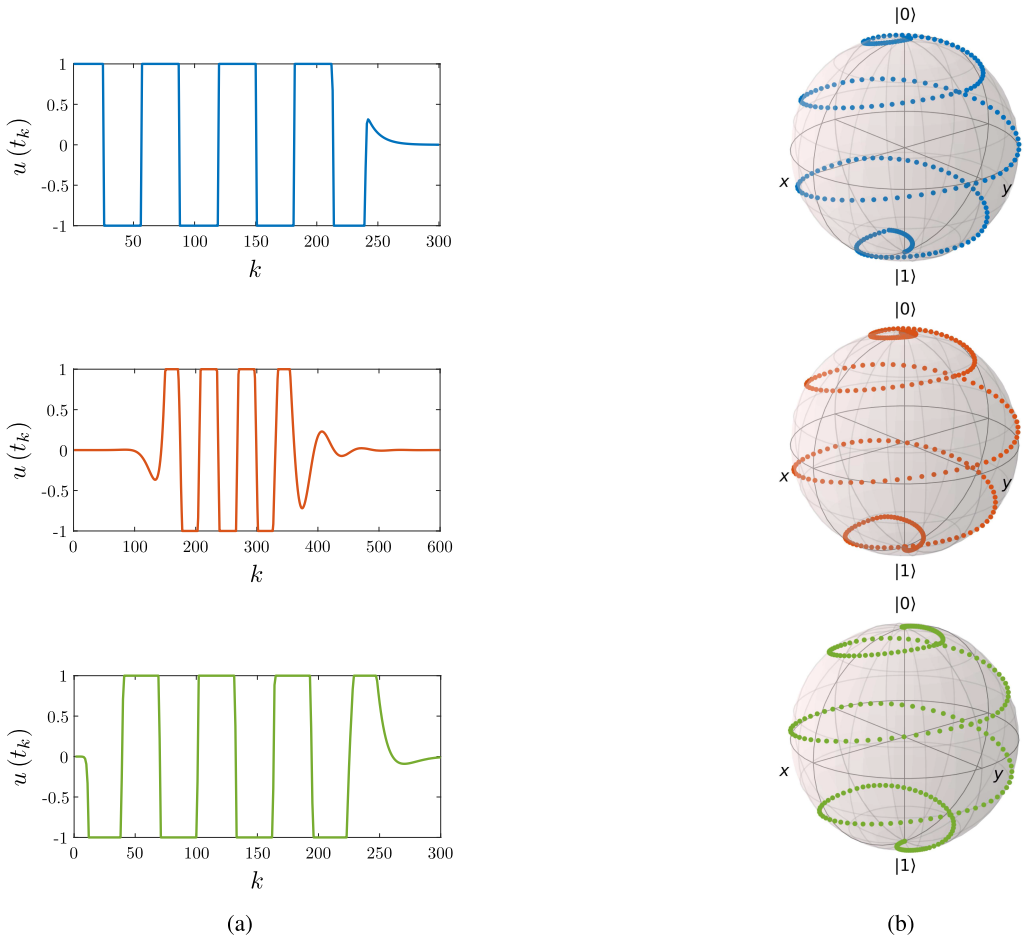

FIGURE 6. (a) Control input selected by (α, β)-minimum norm QOC in blue, the β-minimum norm MPQC in orange, and the (α, β)-minimum norm MPQC in green. (b) Evolution of the qubit's state on the Bloch sphere in response to each corresponding control input.

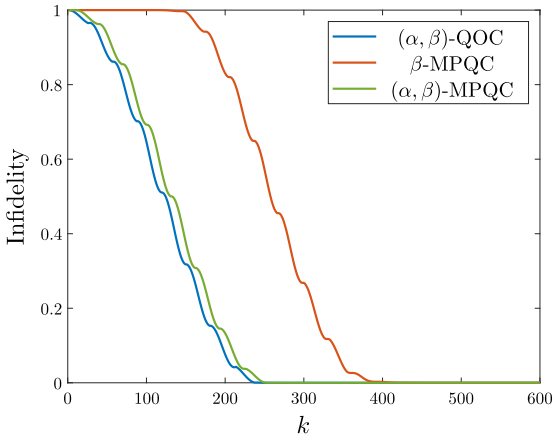

FIGURE 7. Infidelity of the simulated qubit at time step for each of the three control scenarios.

the control sequence generated by this controller is longer than required by the others, it appears to be more smooth. This may be of practical interest in cases where highly discontinuous changes in a control field excite unwanted energy levels in a quantum system.

D. EXPERIMENT 4

For the final experiment, the goal is to generate arbitrary unitary gates for the qubit system studied in Experiment 3. We compare the efficacy of MPQC and GRAPE for this task. Let $K \in \mathbb{N}$ be the total length of the control sequence. GRAPE generates a control signal by solving the following optimization problem:[4]

$$\begin{aligned} \mathscr{P}_6: \quad & \underset{\boldsymbol{u}(t_k):\, k \in \mathcal{Z}_K}{\text{minimize}} \quad g_6\left(\boldsymbol{U}(t_K)\right) \\ & \text{subject to} \quad \boldsymbol{U}(t_{k+1}) = e^{-\imath \boldsymbol{H}(t_k)\Delta_t}\, \boldsymbol{U}(t_k) \\ & \qquad\qquad \left\|\boldsymbol{u}(t_k)\right\|_\infty \leqslant u_{\max}, \quad \forall k \in \mathcal{Z}_K \end{aligned}$$

Notice that GRAPE only considers a terminal cost as opposed to MPQC, which can consider a different cost at each step in the control sequence. For this comparison, Problem 5 uses the gate error, g_6, at each time step as its objective function. The control and prediction horizons are set to

[4]The original GRAPE paper [61] considers several optimization problems. In one scenario, they attempt to maximize the gate fidelity. This is equivalent to minimizing the gate infidelity, which we consider here. The hallmark of GRAPE is the optimization procedure used rather than the cost function.

IEEE Transactions on Quantum Engineering

TABLE 5. Comparison of MPQC and GRAPE for Generating Quantum Gates

	Hadamard		Pauli-$\boldsymbol{X}$		Pauli-$\boldsymbol{Y}$	
	$K = 157$		$K = 476$		$K = 301$	
Controller	MPQC	GRAPE	MPQC	GRAPE	MPQC	GRAPE
Run Time (s)	8.3212	16.3905	19.6201	70.7177	13.2996	51.0152
Fidelity Error	$2.7860 \cdot 10^{-4}$	$9.9985 \cdot 10^{-4}$	$8.7500 \cdot 10^{-6}$	$9.9958 \cdot 10^{-4}$	$2.7540 \cdot 10^{-4}$	$9.9999 \cdot 10^{-4}$

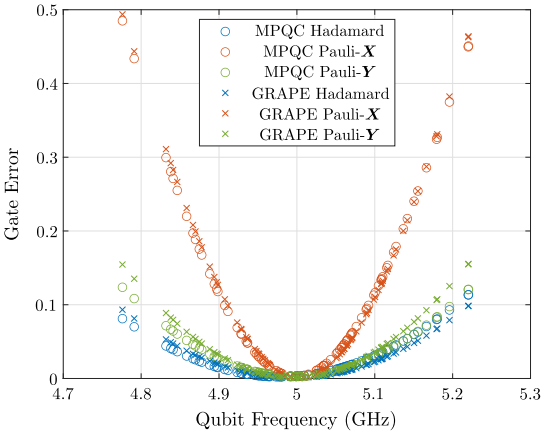

FIGURE 8. Gate error as a function of qubit frequency for the three single-qubit gates studied in Experiment 4. Control signals optimized by MPQC (denoted by ∘'s) and GRAPE (denoted by ×'s) for a nominal 5-GHz qubit were given to a qubit with each new frequency, and the resulting gate error was recorded. In most cases, the control generated by MPQC is more robust to uncertainty in the qubit frequency than that of GRAPE.

$K_\mathrm{c} = 4$ and $K_\mathrm{p} = 12$, respectively. Problem 5 optimizes the variable $\boldsymbol{U}(t_k)$, which is of dimension d^2, at each time step as opposed to Problem 4, which optimized the variable $|\psi(t_k)\rangle$ of dimension d at each time step. Thus, the number of decision variables optimized at each step of SQP in MPQC is greater than that of Experiment 3. Hence, longer computation time can be expected. Our GRAPE implementation follows that of [61], which uses a gradient-based procedure for solving Problem 6. For this comparison, the control is limited to $u_\mathrm{max} = 2\,\mathrm{GHz}$. The iterative optimization procedures of both MPQC and GRAPE were set to terminate after the gate error reached a predefined threshold of 10^{-3}. The qubit dynamics are the same as in Experiment 3. Unlike MPQC, the number of control steps, K, must be defined prior to solving Problem 6 (GRAPE), and there are no established rules for choosing this parameter. However, MPQC is designed to iterate through time, solving Problem 5 at each step, until a desired gate error is achieved. Hence, for this experiment, MPQC was performed first and K was recorded. Then, Problem 6 (GRAPE) was solved using the same value. Table 5 summarizes the results. For all gates studied, we found that MPQC offered a significant speedup over GRAPE and that both the methods achieved low gate errors.

Next, the robustness of the control signals generated by MPQC and GRAPE is studied. Previously, control signals for three single-qubit gates were optimized under the assumption that the qubit's frequency was 5 GHz. These control signals were then provided to qubits whose frequency differs from the 5-GHz value. Fig. 8 plots the resulting gate errors as a function of qubit frequency. In this figure, circles represent data points for MPQC and crosses represent data points for GRAPE. These points are color-coded to indicate which unitary gate they correspond to. We observe that, in almost all cases, MPQC generated gates, which are more robust to uncertainty in the qubit frequency.

VI. CONCLUSION

This article introduced the concept of data-driven MPQC to control general closed quantum systems, where no nominal model is provided. To this end, a novel and efficient approach for QHL was proposed. To learn the internal and control Hamiltonians, a novel QPT algorithm was developed, which involves optimization on the unitary group. The learned Hamiltonians were then used by MPQC to compute quantum control sequences. The proposed QHL algorithm allows for inferring the Hamiltonians of high-dimensional quantum systems due to its low memory requirement. In addition, the MPQC framework is flexible and works for a variety of control costs. In our experiments, when an SQP solver is used, MPQC is observed to be much faster at computing control sequences than the current state-of-the-art, QOC. The success of many quantum technologies depends on the ability to precisely characterize and control quantum systems, and our results offer a promising approach for controlling quantum systems with uncertainty. Our results can serve as guidelines for designing robust and efficient data-driven control policies for the intervention in quantum dynamics such as interactions of atoms and molecules or superconducting qubits.

APPENDIX

A. PROOF OF PROPOSITION 4

Proof: Let $\boldsymbol{X}, \boldsymbol{Z} \in \mathbb{C}^{d\times d}$ be arbitrary matrices. By definition of the cost, we have

$$\begin{aligned} f_\mathrm{B}(\boldsymbol{X}+\boldsymbol{Z}) &= \sum_{n=1}^{N_\mathrm{i}} \mathrm{tr}\{\boldsymbol{B}_n\,[\boldsymbol{X}+\boldsymbol{Z}]\,\boldsymbol{A}_n\,[\boldsymbol{X}+\boldsymbol{Z}]^\dagger\} \\ &= \sum_{n=1}^{N_\mathrm{i}} \mathrm{tr}\left\{\boldsymbol{B}_n\boldsymbol{X}\boldsymbol{A}_n\boldsymbol{X}^\dagger + \boldsymbol{B}_n\boldsymbol{X}\boldsymbol{A}_n\boldsymbol{Z}^\dagger \right. \\ &\qquad \left. + \boldsymbol{B}_n\boldsymbol{Z}\boldsymbol{A}_n\boldsymbol{X}^\dagger + \boldsymbol{B}_n\boldsymbol{Z}\boldsymbol{A}_n\boldsymbol{Z}^\dagger\right\} \end{aligned}$$

$$= f_{\mathrm{B}}(\boldsymbol{X}) + f_{\mathrm{B}}(\boldsymbol{Z}) + \sum_{n=1}^{N_{\mathrm{i}}} \mathrm{tr}\{\boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n \boldsymbol{Z}^\dagger + \boldsymbol{B}_n \boldsymbol{Z} \boldsymbol{A}_n \boldsymbol{X}^\dagger\}. \tag{25}$$

Using the linearity and cyclic properties of the trace, we have

$$\begin{aligned}\mathrm{tr}\{\boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n \boldsymbol{Z}^\dagger + \boldsymbol{B}_n \boldsymbol{Z} \boldsymbol{A}_n \boldsymbol{X}^\dagger\} &= \mathrm{tr}\{\boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n \boldsymbol{Z}^\dagger\} + \mathrm{tr}\{\boldsymbol{Z} \boldsymbol{A}_n \boldsymbol{X}^\dagger \boldsymbol{B}_n\} \\ &= 2\mathfrak{R}\{\mathrm{tr}\{\boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n \boldsymbol{Z}^\dagger\}\} \\ &= 2\mathfrak{R}\{\mathrm{tr}\{\boldsymbol{Z}^\dagger \boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n\}\}. \end{aligned} \tag{26}$$

Above, the second equality holds because, for any complex number $c \in \mathbb{C}$, the relationship $c + c^* = 2\,\mathfrak{R}\{c\}$ holds. Since each $\boldsymbol{A}_n$ and $\boldsymbol{B}_n$ are positive semidefinite, they admit principal square roots $\boldsymbol{P}_n$ and $\boldsymbol{L}_n$, respectively, such that $\boldsymbol{A}_n = \boldsymbol{P}_n^\dagger \boldsymbol{P}_n$ and $\boldsymbol{B}_n = \boldsymbol{L}_n^\dagger \boldsymbol{L}_n$. Thus

$$\begin{aligned} f_{\mathrm{B}}(\boldsymbol{Z}) &= \sum_{n=1}^{N_{\mathrm{i}}} \mathrm{tr}\{\boldsymbol{B}_n \boldsymbol{Z} \boldsymbol{A}_n \boldsymbol{Z}^\dagger\} \\ &= \sum_{n=1}^{N_{\mathrm{i}}} \mathrm{tr}\{\boldsymbol{L}_n \boldsymbol{Z} \boldsymbol{P}_n^\dagger \boldsymbol{P}_n \boldsymbol{Z}^\dagger \boldsymbol{L}_n^\dagger\} \\ &= \sum_{n=1}^{N_{\mathrm{i}}} \|\boldsymbol{L}_n \boldsymbol{Z} \boldsymbol{P}_n^\dagger\|_{\mathrm{F}}^2 \\ &\leqslant \sum_{n=1}^{N_{\mathrm{i}}} \|\boldsymbol{L}_n\|_{\mathrm{F}}^2 \|\boldsymbol{P}_n\|_{\mathrm{F}}^2 \|\boldsymbol{Z}\|_{\mathrm{F}}^2 \end{aligned} \tag{27}$$

which shows that $f_{\mathrm{B}}(\boldsymbol{Z})$ grows on the order $\mathcal{O}(\|\boldsymbol{Z}\|_{\mathrm{F}}^2)$. The definition of the Frobenius norm and its submultiplicative property were used to obtain (27). Combining (25)–(27) gives

$$\begin{aligned} f_{\mathrm{B}}(\boldsymbol{X} + \boldsymbol{Z}) &= f_{\mathrm{B}}(\boldsymbol{X}) + 2\sum_{n=1}^{N_{\mathrm{i}}} \mathfrak{R}\{\mathrm{tr}\{\boldsymbol{Z}^\dagger \boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n\}\} + \mathcal{O}(\|\boldsymbol{Z}\|_{\mathrm{F}}^2) \\ &= 2\,\mathfrak{R}\Big\{\mathrm{tr}\Big\{\boldsymbol{Z}^\dagger \sum_{n=1}^{N_{\mathrm{i}}} \boldsymbol{B}_n \boldsymbol{X} \boldsymbol{A}_n\Big\}\Big\} + f_{\mathrm{B}}(\boldsymbol{X}) + \mathcal{O}(\|\boldsymbol{Z}\|_{\mathrm{F}}^2) \end{aligned}$$

which holds for all $\boldsymbol{X}, \boldsymbol{Z} \in \mathbb{C}^{d\times d}$. Hence, this expression also holds for all $\boldsymbol{Z} \in T_{\boldsymbol{X}}(d) \subset \mathbb{C}^{d\times d}$, and using Proposition 3, the desired result follows. □

B. PROOF OF PROPOSITION 5

Proof: Let $\boldsymbol{Y} \in \mathcal{U}(d)$ and $\boldsymbol{N}, \boldsymbol{A} \in \mathbb{C}^{d\times d}$. When $\boldsymbol{A}$ is Hermitian and $\boldsymbol{N}$ is diagonal, it is well known that the solution to the Brockett optimization problem

$$\hat{\boldsymbol{Y}} = \underset{\boldsymbol{Y}\in\mathcal{U}(d)}{\arg\max}\ \mathrm{tr}\{\boldsymbol{N}\boldsymbol{Y}\boldsymbol{A}\boldsymbol{Y}^\dagger\}$$

is given by some $\hat{\boldsymbol{Y}}$ whose rows are orthonormal eigenvectors of $\boldsymbol{A}$ [95]. Hence, by the spectral theorem, $\hat{\boldsymbol{Y}}\boldsymbol{A}\hat{\boldsymbol{Y}}^\dagger$ is a diagonal matrix whose diagonal entries are the eigenvalues of $\boldsymbol{A}$. Note that $\boldsymbol{N}\hat{\boldsymbol{Y}}\boldsymbol{A}\hat{\boldsymbol{Y}}^\dagger$ is also a diagonal matrix. Hence, the Brockett function is maximized by choosing the rows of $\hat{\boldsymbol{Y}}$ to be an appropriate permutation of the eigenvectors of $\boldsymbol{A}$.

Next, we consider the implications of the previous discussion on Problem 2. Define the orthogonal transformation $\boldsymbol{Y} \triangleq \hat{\boldsymbol{Q}}^\dagger \boldsymbol{X}$, where $\hat{\boldsymbol{Q}}$ is the orthogonal matrix of eigenvectors of $\hat{\tilde{\boldsymbol{\Xi}}}(t_{\mathrm{f}})$, as defined in (12). In Problem 2, the variable $\boldsymbol{X}$ is unitary; hence, the variable $\boldsymbol{Y}$ is unitary as well. Using (12) and the cyclic property of the trace, Problem 2 becomes

$$\hat{\boldsymbol{Y}} = \underset{\boldsymbol{Y}\in\mathcal{U}(d)}{\arg\max}\ \mathrm{tr}\{\hat{\boldsymbol{\Lambda}}\boldsymbol{Y}\tilde{\boldsymbol{\Xi}}(t_0)\boldsymbol{Y}^\dagger\}$$

which is an optimization problem of the form discussed at the beginning of the proof. Once a maximizer $\hat{\boldsymbol{Y}}$ to this problem is obtained, we can recover the unitary maximizer to Problem 2 as $\hat{\boldsymbol{X}} = \hat{\boldsymbol{Q}}\hat{\boldsymbol{Y}}$. Recalling the prior discussion, since $\hat{\boldsymbol{\Lambda}}$ is diagonal and $\tilde{\boldsymbol{\Xi}}(t_0)$ is Hermitian, the solution to this problem is obtained by choosing the rows of $\boldsymbol{Y}$ to be some permutation of the eigenvectors of $\tilde{\boldsymbol{\Xi}}(t_0)$ to maximize $\mathrm{tr}\{\hat{\boldsymbol{\Lambda}}\boldsymbol{Y}\boldsymbol{\Xi}_0\boldsymbol{Y}^\dagger\} = \mathrm{tr}\{\hat{\boldsymbol{\Lambda}}\boldsymbol{\Lambda}\}$. Since the eigenvalues in decompositions (11) and (12) are nonnegative and are in nondecreasing order, $\hat{\boldsymbol{Y}} = \boldsymbol{V}^\dagger$ is a solution to the problem. Therefore, $\hat{\boldsymbol{U}} = \hat{\boldsymbol{X}} = \hat{\boldsymbol{Q}}\boldsymbol{V}^\dagger$ is a solution to Problem 2. In the case of pure states, Problem 2 is equivalent to Problem 1, so $\hat{\boldsymbol{U}} = \hat{\boldsymbol{Q}}\boldsymbol{V}^\dagger$ solves Problem 1 as well. In the case of mixed states, a similar conclusion is drawn for solutions of Problems 2 and 3. □

C. PROOF OF PROPOSITION 6

Prior to presenting the proof, we provide a few remarks. Since $\tilde{\boldsymbol{\Xi}}(t_0)$ and $\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}})$ are related by a unitary transformation, they share the same eigenvalues. Hence, given the assumption that $\tilde{\boldsymbol{\Xi}}(t_0)$ has simple eigenvalues, so does $\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}})$. The eigenspace associated with each simple eigenvalue is of dimension one. Therefore, the orthonormal eigenvectors of each matrix are unique up to a shift by a complex phase. For instance, let $\boldsymbol{q}_k$ represent an arbitrarily computed unit-length eigenvector of $\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}})$ corresponding to the eigenvalue λ_k. Namely, $\boldsymbol{q}_k$ represents the kth column $\boldsymbol{Q}$ from (12). Then, all possible unit-length eigenvectors (corresponding to λ_k) of $\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}})$ can be represented as

$$\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}})\, e^{\imath\phi_k}\boldsymbol{q}_k = \lambda_k\, e^{\imath\phi_k}\boldsymbol{q}_k$$

for an arbitrary phase $\phi_k \in \mathbb{R}$. It should also be noted that, using (8) and (11),

$$\tilde{\boldsymbol{\Xi}}(t_{\mathrm{f}}) = \big[\boldsymbol{U}(t_{\mathrm{f}})\boldsymbol{V}\big]\,\boldsymbol{\Lambda}\,\big[\boldsymbol{U}(t_{\mathrm{f}})\boldsymbol{V}\big]^\dagger \tag{28}$$

IEEE Transactions on Quantum Engineering

is a valid spectral decomposition of the final state; the orthonormal eigenvectors of $\boldsymbol{\Xi}_{\mathrm{f}}$ are the columns of the product $\boldsymbol{U}(t_{\mathrm{f}})\boldsymbol{V}$. Hence, if the error $\boldsymbol{E}$ is zero, the estimate $\hat{\boldsymbol{U}}$ of Proposition 5 exactly recovers $\boldsymbol{U}(t_{\mathrm{f}})$ up to a diagonal phase shift $\boldsymbol{\Delta}_{\boldsymbol{\phi}}$. The proof of Proposition 6 follows below in the general case for any estimation error $\boldsymbol{E}$.

Proof: Suppose, without loss of generality, that $\boldsymbol{U}(t_{\mathrm{f}}) = \boldsymbol{Q}\boldsymbol{V}^{\dagger}$. Then

$$\begin{aligned}\left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{U}} - \boldsymbol{U}\right\|_{\mathrm{F}}^2 &= \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{Q}}\boldsymbol{V}^{\dagger} - \boldsymbol{Q}\boldsymbol{V}^{\dagger}\right\|_{\mathrm{F}}^2 \\ &= \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{Q}} - \boldsymbol{Q}\right\|_{\mathrm{F}}^2 . \end{aligned} \tag{29}$$

The second equality follows from the unitary invariance of the Frobenius norm. It follows that

$$\min_{\boldsymbol{\phi}} \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{U}} - \boldsymbol{U}\right\|_{\mathrm{F}}^2 = \min_{\boldsymbol{\phi}} \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{Q}} - \boldsymbol{Q}\right\|_{\mathrm{F}}^2 . \tag{30}$$

If for each pair of columns $\boldsymbol{q}_k$ and $\hat{\boldsymbol{q}}_k$ of $\boldsymbol{Q}$ and its estimate $\hat{\boldsymbol{Q}}$, respectively, the following inequality holds:

$$\min_{\phi_k} \left\|e^{\imath\phi_k}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 \leqslant \max_{i\neq j} \frac{8}{\left|\lambda_i - \lambda_j\right|^2}\left\|\boldsymbol{E}\right\|^2 \tag{31}$$

then

$$\begin{aligned}\min_{\boldsymbol{\phi}} \left\|\boldsymbol{\Delta}_{\boldsymbol{\phi}}\hat{\boldsymbol{Q}} - \boldsymbol{Q}\right\|_{\mathrm{F}}^2 &= \sum_{k=1}^{d} \min_{\phi_k} \left\|e^{\imath\phi_k}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 \\ &\leqslant \max_{i\neq j} \frac{8d}{\left|\lambda_i - \lambda_j\right|^2}\left\|\boldsymbol{E}\right\|^2 \end{aligned}$$

which via (30) proves the desired result. Hence, it only remains to show that (31) holds.

We will now prove (31) for an arbitrary $k \in \{1, 2, \ldots, d\}$. Therefore, $\boldsymbol{q}_k$ and $\hat{\boldsymbol{q}}_k$ represent the kth columns of $\boldsymbol{Q}$ and its estimate $\hat{\boldsymbol{Q}}$, respectively. For an arbitrary angle $\vartheta \in \mathbb{R}$, note that

$$\begin{aligned}\left\|e^{\imath\vartheta}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 &= \left(e^{-\imath\vartheta}\hat{\boldsymbol{q}}_k^{\dagger} - \boldsymbol{q}_k^{\dagger}\right)\left(e^{\imath\vartheta}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right) \\ &= \left(\hat{\boldsymbol{q}}_k^{\dagger}\hat{\boldsymbol{q}}_k + \boldsymbol{q}_k^{\dagger}\boldsymbol{q}_k\right) - 2\Re\left\{e^{-\imath\vartheta}\hat{\boldsymbol{q}}_k\boldsymbol{q}_k\right\} \\ &= 2\big(1 - \Re\big\{e^{-\imath\vartheta}\hat{\boldsymbol{q}}_k^{\dagger}\boldsymbol{q}_k\big\}\big) . \end{aligned} \tag{32}$$

The last equality above holds because $\boldsymbol{q}_k$ and $\hat{\boldsymbol{q}}_k$ are unitary. Fix the angle ϑ to be ϑ_0 such that $\Re\{e^{-\imath\vartheta_0}\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\} = \left|\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\right|$. Since $\boldsymbol{q}_k$ and $\hat{\boldsymbol{q}}_k$ are unit length, $|\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle| \leqslant 1$ and

$$\begin{aligned}\Re\big\{e^{-\imath\vartheta_0}\hat{\boldsymbol{q}}_k^{\dagger}\boldsymbol{q}_k\big\} &= \left|\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\right| \\ &\geqslant \left|\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\right|^2 \end{aligned}$$

from which (32) provides

$$\left\|e^{\imath\vartheta_0}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 \leqslant 2\Big(1 - |\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle|^2\Big) . \tag{33}$$

Since the columns of $\boldsymbol{Q}$ form a basis for $\mathbb{C}^d$, it is possible to rewrite $\hat{\boldsymbol{q}}_k$ as

$$\hat{\boldsymbol{q}}_k = \sum_{m=1}^{d} \langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\, \boldsymbol{q}_m . \tag{34}$$

Next, using the orthonormality of the eigenvectors, observe that

$$\begin{aligned}\Big\|[\hat{\boldsymbol{\Xi}}(t_{\mathrm{f}}) - \lambda_k\boldsymbol{I}]\hat{\boldsymbol{q}}_k\Big\|^2 &= \Big\|[\hat{\boldsymbol{\Xi}}(t_{\mathrm{f}}) - \lambda_k\boldsymbol{I}]\sum_{m=1}^{d}\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\,\boldsymbol{q}_m\Big\|^2 \\ &= \Big\|\sum_{\substack{m=1\\m\neq k}}^{d}\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\,(\lambda_m - \lambda_k)\,\boldsymbol{q}_m\Big\|^2 \\ &= \sum_{\substack{m=1\\m\neq k}}^{d}\left|\,\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\,(\lambda_m - \lambda_k)\,\right|^2 \\ &= \sum_{\substack{m=1\\m\neq k}}^{d}(\lambda_m - \lambda_k)^2\left|\,\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\,\right|^2 \\ &\geqslant \min_{i\neq j}(\lambda_i - \lambda_j)^2\sum_{\substack{m=1\\m\neq k}}^{d}\left|\,\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_m\rangle\,\right|^2 \\ &= \min_{i\neq j}\left(\lambda_i - \lambda_j\right)^2\Big(1 - \left|\,\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\,\right|^2\Big) . \end{aligned} \tag{35}$$

The first equality follows from (34), the second from the eigenrelation between $\hat{\boldsymbol{\Xi}}(t_{\mathrm{f}})$ and its eigenvector $\boldsymbol{q}_k$, and the third equality comes from applying the definition of the 2-norm and the orthogonality of the columns of $\boldsymbol{Q}$. The final equality follows from the fact that $\hat{\boldsymbol{q}}_k$ is unitary and $\boldsymbol{q}_m$ for $m \in \{1, 2, \ldots, d\}$ forms an orthonormal basis. On the other hand, we have

$$\begin{aligned}\Big\|[\hat{\boldsymbol{\Xi}}(t_{\mathrm{f}}) - \lambda_k\boldsymbol{I}]\hat{\boldsymbol{q}}_k\Big\| &= \Big\|\big[\hat{\tilde{\boldsymbol{\Xi}}}(t_{\mathrm{f}}) - \boldsymbol{E} - \lambda_k\boldsymbol{I}\big]\hat{\boldsymbol{q}}_k\Big\| \\ &\leqslant \Big\|\big[\hat{\tilde{\boldsymbol{\Xi}}}(t_{\mathrm{f}}) - \lambda_k\boldsymbol{I}\big]\hat{\boldsymbol{q}}_k\Big\| + \left\|\boldsymbol{E}\hat{\boldsymbol{q}}_k\right\| \\ &\leqslant \left|\hat{\lambda}_k - \lambda_k\right|\left\|\hat{\boldsymbol{q}}_k\right\| + \left\|\boldsymbol{E}\right\|\left\|\hat{\boldsymbol{q}}_k\right\| \\ &= \left|\hat{\lambda}_k - \lambda_k\right| + \left\|\boldsymbol{E}\right\| \\ &\leqslant 2\left\|\boldsymbol{E}\right\| . \end{aligned} \tag{36}$$

While deriving the previous inequality, we have used the following facts:

1) the definition $\hat{\tilde{\boldsymbol{\Xi}}}(t_{\mathrm{f}}) = \hat{\boldsymbol{\Xi}}(t_{\mathrm{f}}) + \boldsymbol{E}$;
2) the triangle inequality;
3) the absolute homogeneity property of the 2-norm and submultiplicative property of the operator norm;
4) the fact that $\hat{\boldsymbol{q}}_k$ is unit length;
5) for the perturbed Hermitian matrix $\hat{\tilde{\boldsymbol{\Xi}}}(t_{\mathrm{f}}) = \hat{\boldsymbol{\Xi}}(t_{\mathrm{f}}) + \boldsymbol{E}$, Weyl's celebrated eigenvalue perturbation inequality states that $\left|\hat{\lambda}_k - \lambda_k\right| \leqslant \left\|\boldsymbol{E}\right\|$ for all k.

Together, (35) and (36) produce

$$1 - \left|\,\langle\hat{\boldsymbol{q}}_k|\boldsymbol{q}_k\rangle\,\right|^2 \leqslant \max_{i\neq j}\frac{4}{\left|\lambda_i - \lambda_j\right|^2}\left\|\boldsymbol{E}\right\|^2 . \tag{37}$$

Together, (33) and (37) give

$$\left\|e^{\imath\vartheta_0}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 \leqslant \max_{i\neq j} \frac{8}{\left|\lambda_i - \lambda_j\right|^2}\left\|\boldsymbol{E}\right\|^2.$$

Hence

$$\min_{\phi_k}\left\|e^{\imath\phi_k}\hat{\boldsymbol{q}}_k - \boldsymbol{q}_k\right\|^2 \leqslant \max_{i\neq j} \frac{8}{\left|\lambda_i - \lambda_j\right|^2}\left\|\boldsymbol{E}\right\|^2$$

which proves that (31) holds and the proposition is proved. □

D. PROOF OF PROPOSITION 9

Proof: Algorithm 3 estimates $\boldsymbol{H}_k$ according to

$$\hat{\boldsymbol{H}}_k = \frac{1}{c_k}\left[\hat{\boldsymbol{H}} - \hat{\boldsymbol{H}}_0\right].$$

Equation (16) allows the following deductions:

$$\begin{aligned}
\left\|\boldsymbol{H}_k - \hat{\boldsymbol{H}}_k\right\|_{\mathrm{F}} &= \left\|\boldsymbol{H}_k - \frac{1}{c_k}\left[\hat{\boldsymbol{H}} - \hat{\boldsymbol{H}}_0\right]\right\|_{\mathrm{F}} \\
&= \frac{1}{|c_k|}\left\|\hat{\boldsymbol{H}} - \hat{\boldsymbol{H}}_0 - c_k\boldsymbol{H}_k\right\|_{\mathrm{F}} \\
&= \frac{1}{|c_k|}\left\|\hat{\boldsymbol{H}} - \left[\boldsymbol{H}_0 + c_k\boldsymbol{H}_k\right] + \boldsymbol{H}_0 - \hat{\boldsymbol{H}}_0\right\|_{\mathrm{F}} \\
&\leqslant \frac{1}{|c_k|}\Big(\left\|\hat{\boldsymbol{H}} - \left[\boldsymbol{H}_0 + c_k\boldsymbol{H}_k\right]\right\|_{\mathrm{F}} \\
&\qquad + \left\|\boldsymbol{H}_0 - \hat{\boldsymbol{H}}_0\right\|_{\mathrm{F}}\Big) \\
&\leqslant \frac{\pi}{2t_{\mathrm{f}}|c_k|}\left(\left\|\boldsymbol{U} - \hat{\boldsymbol{U}}\right\|_{\mathrm{F}} + \left\|\boldsymbol{U}_0 - \hat{\boldsymbol{U}}_0\right\|_{\mathrm{F}}\right)
\end{aligned}$$

which proves the desired result. □

ACKNOWLEDGMENT

The authors would like to thank S. Marano, C. E. Souza, and F. Zaman for their helpful suggestions and careful reading of this article.

REFERENCES

[1] V. Giovannetti, S. Lloyd, and L. Maccone, "Advances in quantum metrology," *Nature Photon.*, vol. 5, no. 4, pp. 222–229, 2011, doi: 10.1038/nphoton.2011.35.

[2] A. S. Fletcher, P. W. Shor, and M. Z. Win, "Optimum quantum error recovery using semidefinite programming," *Phys. Rev. A*, vol. 75, no. 1, Jan. 2007, Art. no. 012338, doi: 10.1103/PhysRevA.75.012338.

[3] A. S. Cacciapuoti, M. Caleffi, R. Van Meter, and L. Hanzo, "When entanglement meets classical communications: Quantum teleportation for the quantum internet," *IEEE Trans. Commun.*, vol. 68, no. 6, pp. 3808–3833, Jun. 2020, doi: 10.1109/TCOMM.2020.2978071.

[4] A. S. Cacciapuoti, M. Caleffi, F. Tafuri, F. S. Cataliotti, S. Gherardini, and G. Bianchi, "Quantum internet: Networking challenges in distributed quantum computing," *IEEE Netw.*, vol. 34, no. 1, pp. 137–143, Jan./Feb. 2020, doi: 10.1109/MNET.001.1900092.

[5] M. Chiani, A. Conti, and M. Z. Win, "Piggybacking on quantum streams," *Phys. Rev. A*, vol. 102, no. 1, Jul. 2020, Art. no. 012410, doi: 10.1103/PhysRevA.102.012410.

[6] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2010.

[7] M. Hirose and P. Cappellaro, "Coherent feedback control of a single qubit in diamond," *Nature*, vol. 532, no. 7597, pp. 77–80, 2016. doi: 10.1038/nature17404.

[8] D. J. Egger and F. K. Wilhelm, "Optimized controlled-Z gates for two superconducting qubits coupled through a resonator," *Supercond. Sci. Technol.*, vol. 27, no. 1, 2013, Art. no. 014001, doi: 10.1088/0953-2048/27/1/014001.

[9] X. Wang, M. Allegra, K. Jacobs, S. Lloyd, C. Lupo, and M. Mohseni, "Time-optimal quantum control via differential geometry," *Proc. SPIE*, vol. 10118, 2017, Art. no. 101180F, doi: 10.1117/12.2256267.

[10] D. Rossini *et al.*, "Bang-bang control of a qubit coupled to a quantum critical spin bath," *Phys. Rev. A*, vol. 77, no. 5, 2008, Art. no. 052112, doi: 10.1103/PhysRevA.77.052112.

[11] L. Viola and S. Lloyd, "Dynamical suppression of decoherence in two-state quantum systems," *Phys. Rev. A*, vol. 58, no. 4, 1998, Art. no. 2733, doi: 10.1103/PhysRevA.58.2733.

[12] N. Yamamoto, H. I. Nurdin, M. R. James, and I. R. Petersen, "Avoiding entanglement sudden death via measurement feedback control in a quantum network," *Phys. Rev. A*, vol. 78, no. 4, 2008, Art. no. 042339, doi: 10.1103/PhysRevA.78.042339.

[13] H. I. Nurdin, M. R. James, and I. R. Petersen, "Coherent quantum LQG control," *Automatica*, vol. 45, no. 8, pp. 1837–1846, 2009, doi: 10.1016/j.automatica.2009.04.018.

[14] J. Kerckhoff, H. I. Nurdin, D. S. Pavlichin, and H. Mabuchi, "Designing quantum memories with embedded control: Photonic circuits for autonomous quantum error correction," *Phys. Rev. Lett.*, vol. 105, no. 4, 2010, Art. no. 040502, doi: 10.1103/PhysRevLett.105.040502.

[15] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan, "Quantum feedback control and classical control theory," *Phys. Rev. A*, vol. 62, no. 1, 2000, Art. no. 012105, doi: 10.1103/PhysRevA.62.012105.

[16] M. Mirrahimi and R. Van Handel, "Stabilizing feedback controls for quantum systems," *SIAM J. Control Optim.*, vol. 46, no. 2, pp. 445–467, 2007, doi: 10.1137/050644793.

[17] E. A. Jonckheere, A. T. Rezakhani, and F. Ahmad, "Differential topology of adiabatically controlled quantum processes," *Quantum Inf. Process*, vol. 12, no. 3, pp. 1515–1538, 2013, doi: 10.1007/s11128-012-0445-0.

[18] M. Abdelhafez *et al.*, "Universal gates for protected superconducting qubits using optimal control," *Phys. Rev. A*, vol. 101, no. 2, 2020, Art. no. 022321, doi: 10.1103/PhysRevA.101.022321.

[19] S. Y. Huang and H. S. Goan, "Optimal control for fast and high-fidelity quantum gates in coupled superconducting flux qubits," *Phys. Rev. A*, vol. 90, no. 1, 2014, Art. no. 012318, doi: 10.1103/PhysRevA.90.012318.

[20] N. Ghaeminezhad and S. Cong, "Preparation of hadamard gate for open quantum systems by the Lyapunov control method," *IEEE/CAA J. Autom. Sinica*, vol. 5, no. 3, pp. 733–740, May 2018, doi: 10.1109/JAS.2018.7511084.

[21] D. Hocker, Y. Zheng, R. Kosut, T. Brun, and H. Rabitz, "Survey of control performance in quantum information processing," *Quantum Info. Process.*, vol. 15, no. 11, pp. 4361–4390, 2016, doi: 10.1007/s11128-016-1406-9.

[22] S. Guerrini, M. Z. Win, M. Chiani, and A. Conti, "Quantum discrimination of noisy photon-added coherent states," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 2, pp. 469–479, Aug. 2020, doi: 10.1109/JSAIT.2020.3012944.

[23] R. A. Malaney, "Location-dependent communications using quantum entanglement," *Phys. Rev. A*, vol. 81, no. 4, 2010, Art. no. 042319, doi: 10.1103/PhysRevA.81.042319.

[24] R. A. Malaney, "The quantum car," *IEEE Wireless Commun. Lett.*, vol. 5, no. 6, pp. 624–627, Dec. 2016, doi: 10.1109/LWC.2016.2607740.

[25] C. Altafini and F. Ticozzi, "Modeling and control of quantum systems: An introduction," *IEEE Trans. Autom. Control*, vol. 57, no. 8, pp. 1898–1917, Aug. 2012, doi: 10.1109/TAC.2012.2195830.

[26] M. James, "Optimal quantum control theory," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 4, no. 1, pp. 343–367, 2021, doi: 10.1146/annurev-control-061520-010444.

[27] C. Brif, R. Chakrabarti, and H. Rabitz, "Control of quantum phenomena: Past, present and future," *New J. Phys.*, vol. 12, no. 7, 2010, Art. no. 075008, doi: 10.1088/1367-2630/12/7/075008.

[28] D. d'Alessandro, *Introduction to Quantum Control and Dynamics*. London, U.K.: Chapman and Hall/CRC, 2007.

[29] S. J. Glaser *et al.*, "Training Schrödinger's cat: Quantum optimal control," *Eur. Phys. J. D*, vol. 69, no. 12, 2015, Art. no. 279, doi: 10.1140/epjd/e2015-60464-1.

[30] H. Mabuchi and N. Khaneja, "Principles and applications of control in quantum systems," *Int. J. Robust Nonlinear Control*, vol. 15, no. 15, pp. 647–667, 2005, doi: 10.1002/rnc.1016.

[31] D. Dong and I. R. Petersen, "Quantum control theory and applications: A survey," *IET Control Theory Appl.*, vol. 4, no. 12, pp. 2651–2671, 2010, doi: 10.1049/iet-cta.2009.0508.

[32] S. Cong, *Control of Quantum Systems: Theory and Methods*. Hoboken, NJ, USA: Wiley, 2014.

[33] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

[34] M. Clouâtré, M. J. Khojasteh, and M. Z. Win, "Learning-enabled optimal quantum control," *36th Annual MSRP Res. Forum*, Massachusetts Institute of Technology, Cambridge, MA, Aug. 2021.

[35] R. Blume-Kohout, "Optimal, reliable estimation of quantum states," *New J. Phys.*, vol. 12, no. 4, 2010, Art. no. 043034, doi: 10.1088/1367-2630/12/4/043034.

[36] B. Qi, Z. Hou, L. Li, D. Dong, G. Xiang, and G. Guo, "Quantum state tomography via linear regression estimation," *Sci. Rep.*, vol. 3, no. 1, 2013, Art. no. 3496, doi: 10.1038/srep03496.

[37] A. Utreras-Alarcón, M. Rivera-Tapia, S. Niklitschek, and A. Delgado, "Stochastic optimization on complex variables and pure-state quantum tomography," *Sci. Rep.*, vol. 9, no. 1, 2019, Art. no. 16143, doi: 10.1038/s41598-019-52289-0.

[38] L. Zambrano, L. Pereira, S. Niklitschek, and A. Delgado, "Estimation of pure quantum states in high dimension at the limit of quantum accuracy through complex optimization and statistical inference," *Sci. Rep.*, vol. 10, no. 1, 2020, Art. no. 12781, doi: 10.1038/s41598-020-69646-z.

[39] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, "Quantum-process tomography: Resource analysis of different strategies," *Phys. Rev. A*, vol. 77, no. 3, Mar. 2008, Art. no. 032322, doi: 10.1103/PhysRevA.77.032322.

[40] Y. Wang, D. Dong, B. Qi, J. Zhang, I. R. Petersen, and H. Yonezawa, "A quantum Hamiltonian identification algorithm: Computational complexity and error analysis," *IEEE Trans. Autom. Control*, vol. 63, no. 5, pp. 1388–1403, May 2018, doi: 10.1109/TAC.2017.2747507.

[41] Y. Wang, D. Dong, A. Sone, I. R. Petersen, H. Yonezawa, and P. Cappellaro, "Quantum Hamiltonian identifiability via a similarity transformation approach and beyond," *IEEE Trans. Autom. Control*, vol. 65, no. 11, pp. 4632–4647, Nov. 2020, doi: 10.1109/TAC.2020.2973582.

[42] S. Bonnabel, M. Mirrahimi, and P. Rouchon, "Observer-based Hamiltonian identification for quantum systems," *Automatica*, vol. 45, no. 5, pp. 1144–1155, 2009, doi: 10.1016/j.automatica.2008.12.007.

[43] C. Le Bris, M. Mirrahimi, H. Rabitz, and G. Turinici, "Hamiltonian identification for quantum systems: Well-posedness and numerical approaches," *ESIAM—Control Optim. Calculus Variations*, vol. 13, no. 2, pp. 378–395, 2007, doi: 10.1051/cocv:2007013.

[44] Z. Leghtas, G. Turinici, H. Rabitz, and P. Rouchon, "Hamiltonian identification through enhanced observability utilizing quantum control," *IEEE Trans. Autom. Control*, vol. 57, no. 10, pp. 2679–2683, Oct. 2012, doi: 10.1109/TAC.2012.2190209.

[45] N. Wiebe, C. Granade, C. Ferrie, and D. Cory, "Quantum Hamiltonian learning using imperfect quantum resources," *Phys. Rev. A*, vol. 89, no. 4, 2014, Art. no. 042314, doi: 10.1103/PhysRevA.89.042314.

[46] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, "Hamiltonian learning and certification using quantum resources," *Phys. Rev. Lett.*, vol. 112, no. 19, 2014, Art. no. 190501, doi: 10.1103/PhysRevLett.112.190501.

[47] A. A. Gentile *et al.*, "Learning models of quantum systems from experiments," *Nature Photon.*, vol. 17, no. 7, pp. 837–843, 2021, doi: 10.1038/s41567-021-01201-7.

[48] D. P. Bertsekas, *Rollout, Policy Iteration, and Distributed Reinforcement Learning*. Belmont, MA, USA: Athena Sci., 2020.

[49] M. J. Khojasteh, V. Dhiman, M. Franceschetti, and N. Atanasov, "Probabilistic safety constraints for learned high relative degree system dynamics," in *Proc. 2nd Conf. Learn. Dyn. Control*, 2020, pp. 781–792.

[50] R. Cheng, M. J. Khojasteh, A. D. Ames, and J. W. Burdick, "Safe multi-agent interaction through robust control barrier functions with learned uncertainties," in *Proc. IEEE Conf. Decis. Control*, 2020, pp. 777–783, doi: 10.1109/CDC42340.2020.9304395.

[51] S. Tu and B. Recht, "Least-squares temporal difference learning for the linear quadratic regulator," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 5005–5014.

[52] R. B. Wu, H. Ding, D. Dong, and X. Wang, "Learning robust and high-precision quantum controls," *Phys. Rev. A*, vol. 99, no. 4, 2019, Art. no. 042327, doi: 10.1103/PhysRevA.99.042327.

[53] R. B. Wu, B. Chu, D. H. Owens, and H. Rabitz, "Data-driven gradient algorithm for high-precision quantum control," *Phys. Rev. A*, vol. 97, no. 4, 2018, Art. no. 042122, doi: 10.1103/PhysRevA.97.042122.

[54] D. Dong, M. A. Mabrok, I. R. Petersen, B. Qi, C. Chen, and H. Rabitz, "Sampling-based learning control for quantum systems with uncertainties," *IEEE Trans. Control Syst. Technol.*, vol. 23, no. 6, pp. 2155–2166, Nov. 2015, doi: 10.1109/TCST.2015.2404292.

[55] D. Dong, X. Xing, H. Ma, C. Chen, Z. Liu, and H. Rabitz, "Learning-based quantum robust control: Algorithm, applications, and experiments," *IEEE Trans. Cybern.*, vol. 50, no. 8, pp. 3581–3593, Aug. 2020, doi: 10.1109/TCYB.2019.2921424.

[56] C. Wu, B. Qi, C. Chen, and D. Dong, "Robust learning control design for quantum unitary transformations," *IEEE Trans. Cybern.*, vol. 47, no. 12, pp. 4405–4417, Dec. 2017, doi: 10.1109/TCYB.2016.2610979.

[57] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, "Universal quantum control through deep reinforcement learning," *npj Quantum Inf.*, vol. 5, no. 1, 2019, Art. no. 33, doi: 10.1038/s41534-019-0141-3.

[58] D. J. Egger and F. K. Wilhelm, "Adaptive hybrid optimal quantum control for imprecisely characterized systems," *Phys. Rev. Lett.*, vol. 112, no. 24, 2014, Art. no. 240503, doi: 10.1103/PhysRevLett.112.240503.

[59] A. A. Melnikov *et al.*, "Active learning machine learns to create new quantum experiments," *Proc. Nat. Acad. Sci. USA*, vol. 115, no. 6, pp. 1221–1226, 2018, doi: 10.1073/pnas.1714936115.

[60] C. T. Kehlet *et al.*, "Improving solid-state NMR dipolar recoupling by optimal control," *J. Amer. Chem. Soc.*, vol. 126, no. 33, pp. 10202–10203, 2004, doi: 10.1021/ja048786e.

[61] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, "Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms," *J. Magn. Reson.*, vol. 172, no. 2, pp. 296–305, 2005, doi: 10.1016/j.jmr.2004.11.004.

[62] J. S. Li and N. Khaneja, "Ensemble control of bloch equations," *IEEE Trans. Autom. Control*, vol. 54, no. 3, pp. 528–536, Mar. 2009, doi: 10.1109/TAC.2009.2012983.

[63] F. Albertini and D. D'Alessandro, "The K-P problem on tensor products of lie groups and time-optimal control of n quantum bits with a bounded field," *IEEE Trans. Autom. Control*, vol. 63, no. 2, pp. 518–524, Feb. 2018, doi: 10.1109/TAC.2017.2727225.

[64] D. Stefanatos, "Minimum-time transitions between thermal equilibrium states of the quantum parametric oscillator," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 4290–4297, Aug. 2017, doi: 10.1109/TAC.2017.2684083.

[65] Y. Chen, T. T. Georgiou, and A. Tannenbaum, "Matrix optimal mass transport: A quantum mechanical approach," *IEEE Trans. Autom. Control*, vol. 63, no. 8, pp. 2612–2619, Aug. 2018, doi: 10.1109/TAC.2017.2767707.

[66] F. Motzoi, J. M. Gambetta, S. Merkel, and F. Wilhelm, "Optimal control methods for rapidly time-varying Hamiltonians," *Phys. Rev. A*, vol. 84, no. 2, 2011, Art. no. 022307, doi: 10.1103/PhysRevA.84.022307.

[67] P. De Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, "Second order gradient ascent pulse engineering," *J. Magn. Reson.*, vol. 212, no. 2, pp. 412–417, 2011, doi: 10.1016/j.jmr.2011.07.023.

[68] C. Yang *et al.*, "Silicon qubit fidelities approaching incoherent noise limits via pulse engineering," *Nature Electron.*, vol. 2, no. 4, pp. 151–158, 2019, doi: 10.1038/s41928-019-0234-1.

[69] Z. Zong *et al.*, "Optimization of a controlled-Z gate with data-driven gradient-ascent pulse engineering in a superconducting-qubit system," *Phys. Rev. Appl.*, vol. 15, no. 6, 2021, Art. no. 064005, doi: 10.1103/PhysRevApplied.15.064005.

[70] T. Caneva, T. Calarco, and S. Montangero, "Chopped random-basis quantum optimization," *Phys. Rev. A*, vol. 84, no. 2, 2011, Art. no. 022326, doi: 10.1103/PhysRevA.84.022326.

[71] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, "Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape," *Phys. Rev. A*, vol. 92, no. 6, 2015, Art. no. 062343, doi: 10.1103/PhysRevA.92.062343.

[72] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, "Tunable, flexible, and efficient optimization of control pulses for practical qubits," *Phys. Rev. Lett.*, vol. 120, no. 15, 2018, Art. no. 150401, doi: 10.1103/PhysRevLett.120.150401.

[73] V. G. Satorras, Z. Akata, and M. Welling, "Combining generative and discriminative models for hybrid inference," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 20 19, pp. 13820–13830.

[74] N. Wittler *et al.*, "Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits," *Phys. Rev. Appl.*, vol. 15, no. 3, 2021, Art. no. 034080, doi: 10.1103/PhysRevApplied.15.034080.

[75] J.-P. Humaloja and S. Dubljevic, "Linear model predictive control for Schrödinger equation," in *Proc. IEEE Amer. Control Conf.*, 2018, pp. 2569–2574, doi: 10.23919/ACC.2018.8431686.

[76] T. Hashimoto, I. Yoshimoto, and T. Ohtsuka, "Probabilistic constrained model predictive control for Schröinger equation with finite approximation," in *Proc. SICE Annu. Conf.*, 2012, pp. 1613–1618.

[77] T. Hashimoto, "Stability of stochastic model predictive control for Schrödinger equation with finite approximation," *Int. J. Elect. Inf. Eng.*, vol. 11, no. 1, pp. 12–17, 2016, doi: 10.5281/zenodo.1128115.

[78] R. L. Kosut, M. D. Grace, and C. Brif, "Robust control of quantum gates via sequential convex programming," *Phys. Rev. A*, vol. 88, no. 5, 2013, Art. no. 052326, doi: 10.1103/PhysRevA.88.052326.

[79] L. M. Vandersypen and I. L. Chuang, "NMR techniques for quantum control and computation," *Rev. Mod. Phys.*, vol. 76, no. 4, 2005, Art. no. 1037, doi: 10.1103/RevModPhys.76.1037.

[80] N. Khaneja, R. Brockett, and S. J. Glaser, "Time optimal control in spin systems," *Phys. Rev. A*, vol. 63, no. 3, 2001, Art. no. 032308, doi: 10.1103/PhysRevA.63.032308.

[81] J.-S. Li and N. Khaneja, "Control of inhomogeneous quantum ensembles," *Phys. Rev. A*, vol. 73, no. 3, 2006, Art. no. 030302, doi: 10.1103/PhysRevA.73.030302.

[82] M. Nimbalkar *et al.*, "The fantastic four: A plug 'n' play set of optimal control pulses for enhancing NMR spectroscopy," *J. Magn. Reson.*, vol. 228, pp. 16–31, 2013, doi: 10.1016/j.jmr.2012.12.007.

[83] M. Shapiro and P. Brumer, *Quantum Control of Molecular Processes*. Hoboken, NJ, USA: Wiley, 2012.

[84] R. Chakrabarti and H. Rabitz, "Quantum control landscapes," *Int. Rev. Phys. Chem.*, vol. 26, no. 4, pp. 671–735, 2007, doi: 10.1080/01442350701633300.

[85] V. Ramakrishna, M. V. Salapaka, M. Dahleh, H. Rabitz, and A. Peirce, "Controllability of molecular systems," *Phys. Rev. A*, vol. 51, no. 2, pp. 960–966, 1995, doi: 10.1103/PhysRevA.51.960.

[86] M. C. Collodo *et al.*, "Implementation of conditional phase gates based on tunable ZZ interactions," *Phys. Rev. Lett.*, vol. 125, no. 24, 2020, Art. no. 240502, 10.1103/PhysRevLett.125.240502.

[87] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, "Leakage reduction in fast superconducting qubit gates via optimal control," *npj Quantum Inf.*, vol. 7, no. 1, 2021, Art. no. 14, doi: 10.1038/s41534-020-00339-1.

[88] S. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. Murch, and I. Siddiqi, "Mapping the optimal route between two quantum states," *Nature*, vol. 511, no. 7511, pp. 570–573, 2014, doi: 10.1038/nature13559.

[89] J. C. Bardin *et al.*, "Design and characterization of a 28-nm bulk-CMOS cryogenic quantum controller dissipating less than 2 mW at 3 K," *IEEE J. Solid-State Circuits*, vol. 54, no. 11, pp. 3043–3060, Nov. 2019, doi: 10.1109/JSSC.2019.2937234.

[90] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "A quantum engineer's guide to superconducting qubits," *Appl. Phys. Rev.*, vol. 6, no. 2, 2019, Art. no. 021318, doi: 10.1063/1.5089550.

[91] E. Leonard Jr. *et al.*, "Digital coherent control of a superconducting qubit," *Phys. Rev. Appl.*, vol. 11, no. 1, 2019, Art. no. 014009, doi: 10.1103/PhysRevApplied.11.014009.

[92] M. Larocca and D. Wisniacki, "Krylov-subspace approach for the efficient control of quantum many-body dynamics," *Phys. Rev. A*, vol. 103, no. 2, 2021, Art. no. 023107, doi: 10.1103/PhysRevA.103.023107.

[93] R. Jozsa, "Fidelity for mixed quantum states," *J. Modern Opt.*, vol. 41, no. 12, pp. 2315–2323, 1994, doi: 10.1080/09500349414552171.

[94] R. W. Brockett, "Least squares matching problems," *Linear Algebra Appl.*, vol. 122, pp. 761–777, 1989, doi: 10.1016/0024-3795(89)90675-7.

[95] J. H. Manton, "Optimization algorithms exploiting unitary constraints," *IEEE Trans. Signal Process.*, vol. 50, no. 3, pp. 635–650, Mar. 2002, doi: 10.1109/78.984753.

[96] A. W. Knapp, *Lie Groups Beyond an Introduction*, vol. 140. Berlin, Germany: Springer, 1996.

[97] G. E. Bredon, *Topology and Geometry*, vol. 139. Berlin, Germany: Springer, 2013.

[98] A. S. Fletcher, P. W. Shor, and M. Z. Win, "Structured near-optimal channel-adapted quantum error correction," *Phys. Rev. A*, vol. 77, no. 1, Jan. 2008, Art. no. 012320, doi: 10.1103/PhysRevA.77.012320.

[99] C. H. Baldwin, "Efficient and robust methods for quantum tomography," Ph.D. dissertation, Dept. Phys. Astron., Univ. New Mexico, Albuquerque, NM, USA, 2016.

[100] J. A. Smolin, J. M. Gambetta, and G. Smith, "Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise," *Phys. Rev. Lett.*, vol. 108, no. 7, 2012, Art. no. 070502, doi: 10.1103/PhysRevLett.108.070502.

[101] Q. Xu and S. Xu, "Neural network state estimation for full quantum state tomography," 2018, *arXiv:1811.06654*, doi: 10.48550/arXiv.1811.06654.

[102] D. Burgarth and K. Yuasa, "Identifiability of open quantum systems," *Phys. Rev. A*, vol. 89, no. 3, 2014, Art. no. 030302, doi: 10.1103/PhysRevA.89.030302.

[103] J. Zhang and M. Sarovar, "Quantum Hamiltonian identification from measurement time traces," *Phys. Rev. Lett.*, vol. 113, no. 8, 2014, Art. no. 080401, doi: 10.1103/PhysRevLett.113.080401.

[104] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2017.

[105] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Trans. Control Syst. Technol.*, vol. 18, no. 2, pp. 267–278, Mar. 2010, doi: 10.1109/TCST.2009.2017934.

[106] J. Nocedal and S. J. Wright, "Sequential quadratic programming," in *Numerical Optimization*. Berlin, Germany: Springer, 2006, pp. 529–562.

[107] A. L. Dontchev, M. Huang, I. V. Kolmanovsky, and M. M. Nicotra, "Inexact Newton–Kantorovich methods for constrained nonlinear model predictive control," *IEEE Trans. Autom. Control*, vol. 64, no. 9, pp. 3602–3615, Sep. 2019, doi: 10.1109/TAC.2018.2884402.

[108] *Nonlinear MPC*, MathWorks, Natick, MA, USA, 2021. Accessed: Jul. 30, 2021. [Online]. Available: https://www.mathworks.com/help/mpc/ug/nonlinear-mpc.html

[109] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, "Speedup for quantum optimal control from automatic differentiation based on graphics processing units," *Phys. Rev. A*, vol. 95, no. 4, 2017, Art. no. 042318, doi: 10.1103/PhysRevA.95.042318.

[110] N. Johnston, "QETLAB: A MATLAB toolbox for quantum entanglement, version 0.9," Jan. 2022. [Online]. Available: http://qetlab.com

[111] G. Shi, D. Dong, I. R. Petersen, and K. H. Johansson, "Reaching a quantum consensus: Master equations that generate symmetrization and synchronization," *IEEE Trans. Autom. Control*, vol. 61, no. 2, pp. 374–387, Feb. 2016, doi: 10.1109/TAC.2015.2434034.

Maison Clouâtré (Student Member, IEEE) received the B.S.E. degree in electrical engineering and B.S. degree in mathematics from Mercer University, in 2022. He is currently pursuing the Ph.D. degree in aerospace engineering at Texas A&M University, College Station, TX.

Throughout his undergraduate career, Mr. Clouâtré worked in the laboratory of Prof. Makhin Thitsa at Mercer University. During 2019, he worked as a research intern at the Georgia Tech Research Institute and the Vehicle Systems & Control Lab (VSCL) at Texas A&M University. During the summer of 2021, he was a visiting researcher at the Wireless Communication and Network Sciences Laboratory (WINS Lab) as part of the MIT Summer Research Program (MSRP). His research interests include control, optimization, and machine learning and their applications in aerospace engineering and quantum information science.

Mr. Clouâtré is a 2022 Department of Defense (DoD) National Defense Science & Engineering Graduate (NDSEG) Fellow, a 2022 Avilés-Johnson Fellow at Texas A&M, a 2020 Barry Goldwater Scholar, and a 2018 Stamps Scholar. He was also awarded, but declined, a fellowship from the 2022 National Science Foundation (NSF) Graduate Research Fellowship Program (GRFP).

IEEE Transactions on Quantum Engineering

Mohammad Javad Khojasteh (Member, IEEE) received the Ph.D. and M.Sc degrees in electrical and computer engineering from University of California San Diego in 2019 and 2017, respectively. He completed his B.Sc. degree in both electrical engineering and mathematics from Sharif University of Technology, in 2015.

In 2020, he was a Postdoctoral Scholar with the Center for Autonomous Systems and Technology at the California Institute of Technology and a visitor at NASA Jet Propulsion Laboratory, where he worked with Team CoSTAR. Currently, he is a Postdoctoral Associate with the Wireless Communication and Network Sciences Laboratory (WINS Lab) at the Massachusetts Institute of Technology.

Dr. Khojasteh served as a technical program committee member of several IEEE conferences, was co-chair of the IEEE Workshop on Advances in Network Localization and Navigation at the IEEE Globecom 2021, and coorganized an invited session at the 2021 American Control Conference.

Moe Z. Win (Fellow, IEEE) is a Professor at the Massachusetts Institute of Technology (MIT) and the founding director of the Wireless Information and Network Sciences Laboratory. Prior to joining MIT, he was with AT&T Research Laboratories and with the NASA Jet Propulsion Laboratory.

His research encompasses fundamental theories, algorithm design, and network experimentation for a broad range of real-world problems. His current research topics include ultra-wideband systems, network localization and navigation, network interference exploitation, and quantum information science. He has served the IEEE Communications Society as an elected Member-at-Large on the Board of Governors, as elected Chair of the Radio Communications Committee, and as an IEEE Distinguished Lecturer. Over the last two decades, he held various editorial positions for IEEE journals and organized numerous international conferences. Recently, he has served on the SIAM Diversity Advisory Committee.

Dr. Win is an elected Fellow of the AAAS, the EURASIP, the IEEE, and the IET. He was honored with two IEEE Technical Field Awards: the IEEE Kiyo Tomiyasu Award (2011) and the IEEE Eric E. Sumner Award (2006, jointly with R. A. Scholtz). His publications, co-authored with students and colleagues, have received several awards. Other recognitions include the MIT Everett Moore Baker Award (2022), the IEEE Vehicular Technology Society James Evans Avant Garde Award (2022), the IEEE Communications Society Edwin H. Armstrong Achievement Award (2016), the Cristoforo Colombo International Prize for Communications (2013), the Copernicus Fellowship (2011) and the *Laurea Honoris Causa* (2008) from the Università degli Studi di Ferrara, and the U.S. Presidential Early Career Award for Scientists and Engineers (2004). He is an ISI Highly Cited Researcher.