



Exploring the complexity of quantum control optimization trajectories

Journal:	<i>Physical Chemistry Chemical Physics</i>
Manuscript ID:	CP-ART-08-2014-003853.R1
Article Type:	Paper
Date Submitted by the Author:	27-Oct-2014
Complete List of Authors:	Nanduri, Arun; Princeton University, Chemistry Shir, Ofer; Tel-Hai College, School of Computer Science Donovan, Ashley; Princeton University, Chemistry Ho, Tak-San; Princeton University, Chemistry Rabitz, Herschel; Princeton University, Chemistry

SCHOLARONE™
Manuscripts

Exploring the complexity of quantum control optimization trajectories

Arun Nanduri¹, Ofer M. Shir², Ashley Donovan¹, Tak-San Ho¹, and Herschel Rabitz¹

¹*Department of Chemistry, Princeton University, Princeton, NJ 08544*

²*School of Computer Science, Tel-Hai College, Upper Galilee, Israel*

October 24, 2014

Abstract

The control of quantum system dynamics is generally performed by seeking a suitable applied field. The physical objective as a functional of the field forms the quantum control landscape, whose *topology*, under certain conditions, has been shown to contain no critical point suboptimal traps, thereby enabling effective searches for fields that give the global maximum of the objective. This paper addresses the *structure* of the landscape as a complement to topological critical point features. Recent work showed that landscape structure is highly favorable for optimization of state-to-state transition probabilities, in that gradient-based control trajectories to the global maximum value are nearly straight paths. The landscape structure is codified in the metric $R \geq 1.0$, defined as the ratio of the length of the control trajectory to the Euclidean distance between the initial and optimal controls. A value of $R = 1$ would indicate an exactly straight trajectory to the optimal observable value. This paper extends the state-to-state transition probability results to the quantum ensemble and unitary transformation control landscapes. Again, nearly straight trajectories predominate, and we demonstrate that R can take values approaching 1.0 with high precision. However, the interplay of optimization trajectories with critical saddle submanifolds is found to influence landscape structure. A fundamental relationship necessary for perfectly straight gradient-based control trajectories is derived, wherein the gradient on the quantum control landscape must be an eigenfunction of the Hessian. This relation is an indicator of landscape structure and may provide a means to identify physical conditions when control trajectories can achieve perfect linearity. The collective favorable landscape topology

and structure provide a foundation to understand why optimal quantum control can be readily achieved.

1 Introduction

Quantum optimal control theory (OCT) provides a basis to explore applied field tuning of quantum system dynamics [1, 2, 3, 4, 5]. Recent successes include achieving control of electrons in quantum dots through time-varying gate voltages [6], controlling electronic correlations in molecules [7], manipulating the dynamics of quantum many-body systems [8], cooling ultracold atomic gases [9], eliciting quantum revivals [10], and generating adiabatic time evolution [11]. Optimal control experiments (OCE) have found growing success even for complex systems, especially through the use of closed-loop learning algorithms [12] combined with femtosecond laser pulse shaping technology [13]. For example, these advances have enabled vibrational control of population inversion in Bose-Einstein condensates [14], strong-field ionization of silver atoms [15], and coherent energy transfer in light-harvesting complexes [16]. Underlying these achievements is the question of why desirable quantum controls can evidently be found with only modest search effort over the vast space of possible control fields. An important factor in addressing this question is the favorable control landscape *topology*, which results upon satisfaction of certain physical assumptions [17]. The present paper considers a second contributing factor based on the landscape *structure* (i.e., non-topological features) reflected in the behavior of the evolving controls during the optimization process.

The dynamics of a closed quantum system interacting with an applied field is described by the time-dependent Schrödinger equation

$$i\hbar \frac{\partial U(t, 0)}{\partial t} = H(t) U(t, 0), \quad U(0, 0) = 1, \quad (1)$$

where $H(t)$ is the time-dependent Hamiltonian and $U(t, 0)$ is the corresponding evolution operator (propagator) with $0 \leq t \leq T$. Here T is the time at which the objective is optimized. Treating the interaction within the dipole approximation, the time-dependent Hamiltonian is given by

$$H(t) = H_0 - \mu E(t), \quad (2)$$

where H_0 is the field-free Hamiltonian, μ is the dipole moment operator and $E(t)$ is the time-dependent control field.

Through Eq. (1), a control field can be mapped to a value of a cost functional $J[E(t)]$, which may be either the expectation value J_O of an observable O or the cost functional J_W for creating a target unitary transformation W (see Eqs. (3) and (4), respectively). The functional relationship between $J[E(t)]$ and the control field $E(t)$ can be cast in terms of a *quantum control landscape*, where each point $J[E(t)]$ on the landscape corresponds to a particular control field $E(t)$, $0 \leq t \leq T$. Finding successive control fields that eventually lead to an optimal value of the cost functional entails climbing the landscape. We will refer to ‘climbing’ as moving towards the target value of J , regardless of whether the cost functional is to be maximized or minimized. Using a local search algorithm to climb (e.g., the gradient algorithm utilized in this work) results in smooth paths up the landscape $J[E(s, t)]$ that can be parametrized by the continuous variable $s \geq 0$. Every such path up the landscape is associated with a trajectory $E(s, t)$ through *control space*. Thus, starting with an initial control field $E(0, t)$ that produces a low value of the cost functional, the goal is to find the final control field $E(s_{max}, t)$ that yields a high (optimal) value of the cost functional at $s = s_{max}$.

Many recent studies have considered the *topology* of quantum control landscapes [17, 18, 19, 20, 21]. The control landscape topology is specified by the character of the critical points where $\frac{\delta J}{\delta E(t)} = 0$, $\forall t \in [0, T]$. A critical point at a suboptimal value of J where the Hessian $\frac{\delta^2 J}{\delta E(t) \delta E(t')}$ is negative semi-definite would indicate a (second order) trap capable of halting an effort attempting to reach the top of the landscape. The landscape topology has been analyzed for closed quantum systems with N states under the assumptions of: (i) controllability, such that any unitary evolution matrix $U(T, 0)$ can be generated by at least one control field at some sufficiently large time T , (ii) surjectivity, whereby the set of functions $\frac{\delta U_{ij}(T, 0)}{\delta E(t)}$, $\forall i, j$ are linearly independent over $t \in [0, T]$, and (iii) full accessibility, allowing for free access to any control field. Upon satisfaction of these assumptions, the quantum control landscape can be shown to contain only trap-free intermediate value critical saddle submanifolds (when they exist for a particular application) [17, 18, 21] along with the absolute minimum and absolute maximum submanifolds. Thus, the landscape topology is very favorable to performing OCEs or OCT simulations, which is an important factor in explaining the evident ease of locating optimal control fields.

Beyond issues of topology, the *structure* of the quantum control landscape away from the critical points could greatly influence the nature of a control trajectory winding its way to an optimal

field. If landscape structure is highly gnarled, this circumstance could call for complex control trajectories, although the trap free topology assures that the path (e.g., guided by a gradient ascent algorithm) will lead to an optimal solution. The encountered landscape structure upon climbing may be examined by quantifying the linearity of the control trajectories [22]. A predilection towards nearly straight trajectories is indicative of very simple control landscape structural features. Prior experimental work aiming to maximize second harmonic generation found that all of the sampled control trajectories were strikingly close to being straight [26]. Recent theoretical work has also supported the presence of simple landscape structure by examining the linearity of control trajectories in simulations over the state-to-state transition probability landscape [22, 27, 28].

In this work, we extend the state-to-state transition probability quantum control landscape structure analysis to consider two additional landscapes:

$$J_O = \text{Tr}(\rho(T)O) = \text{Tr}\left(U(T,0)\rho(0)U^\dagger(T,0)O\right), \quad (3)$$

and

$$J_W = \|W - U(T,0)\|^2 = 2N - 2\text{Re}\left\{\text{Tr}\left(W^\dagger U(T,0)\right)\right\} \quad (4)$$

where $\|\cdot\|$ denotes the Frobenius norm and N is the number of states in the quantum system. The quantum ensemble control landscape J_O allows for controlling arbitrary observables and consideration of systems which are initially in mixed states $\rho(0)$, and the unitary transformation landscape J_W is relevant to the prospect of creating logic gates for quantum information processing [29]. Both landscapes can contain critical saddle submanifolds [18, 21], which are not present in the simpler state-to-state transition probability landscape. Thus, we aim to assess the influence of these saddle submanifolds on landscape structure. Specifically, we will show that encountering saddle submanifolds at intermediate heights on the landscape can drive control trajectories away from straight paths. This phenomenon is sketched in Figure 1. The trajectory on the right takes a rather straight path up the landscape, while the trajectory on the left is initially attracted to the saddle point, thus altering its direct nature. Nevertheless, we will show in numerical simulations that the trajectories leading to optimal controls are generally still close to being straight, despite the presence of saddles. Lastly, we will derive a mathematical criterion showing that the gradient of the cost functional must be an eigenfunction of the Hessian in order to achieve a perfectly straight gradient-based control trajectory.

The remainder of the paper is organized as follows. The linearity measure R of a control trajectory is defined in Section 2. Section 3 gives the gradient-based search algorithm used in this work to identify control fields that optimize arbitrary observables or achieve unitary transformations, respectively, through J_O or J_W . The presence of landscape saddle submanifolds is examined, and a metric is outlined that gives the distance from any point on a trajectory to a particular saddle submanifold. In Section 4, a derivation is presented specifying the Hessian-gradient eigen-relation for straight control trajectories. Section 5 provides numerical results on the statistical behavior of R over a large ensemble of landscape trajectories for J_O and J_W . We also illustrate satisfaction of the Hessian-gradient eigen-relation for nearly straight control trajectories. Concluding remarks are furnished in Section 6.

2 The control trajectory linearity measure R

The goal of the present work is to explore quantum control landscape structure reflected in the nature of the control trajectories followed during optimizations. Figure 1 indicates that a landscape trajectory has a projected image in the underlying control space. As we use a myopic gradient algorithm to climb the landscape, the degree of gnarled character encountered on the climb is manifested in the trajectory through control space. To this end, we construct a ratio R delineating the linearity of the control trajectory. Specifically, R is the ratio of the path length of the control trajectory between the initial and final fields, given by

$$d_{PL} = \int_0^{s_{max}} \left[\frac{1}{T} \int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt \right]^{\frac{1}{2}} ds, \quad (5)$$

to the Euclidean distance between the initial and final fields in control space, given by

$$d_{EL} = \left[\frac{1}{T} \int_0^T [E(s_{max}, t) - E(0, t)]^2 dt \right]^{\frac{1}{2}}, \quad (6)$$

where s parametrizes the control trajectory. R is expressed as

$$R = \frac{d_{PL}}{d_{EL}} = \frac{\int_0^{s_{max}} \left[\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt \right]^{\frac{1}{2}} ds}{\left[\int_0^T [E(s_{max}, t) - E(0, t)]^2 dt \right]^{\frac{1}{2}}}. \quad (7)$$

The lower bound for the ratio is $R = 1$, which holds for a trajectory that takes a straight path between $E(0, t)$ and $E(s_{max}, t)$. A trajectory with $R = 1$ corresponds to the simplest possible path

a gradient optimization can take to climb the quantum control landscape $J[E(t)]$, and trajectories approaching this limit indicate a lack of encounters with complex structural features. High values of R mean that a control trajectory meanders through control space along a curve whose total length is much greater than the distance between its two endpoints, signifying that the optimization path encountered rugged structural features on the landscape.

In a recent work, we found values of $R - 1$ as low as $\sim 10^{-4}$ for a five level state-to-state transition probability landscape [22]. By examining the distances between the endpoints of the control trajectories, we also observed that trajectories with low and high R values are distributed almost identically throughout control space. Furthermore, we found that it was possible to characterize the complexity of low R trajectories through a ‘straight shot’ climbing procedure. This procedure operates by calculating the gradient of the objective (e.g., the state-to-state transition probability) at the initial control field, and then proceeds to evolve the field along that specified direction until a local maximum is encountered. For suitably low R trajectories, nearly optimal control fields may be found in this way by merely continuing to march in the initially identified direction. This technique is useful for characterizing the straight nature of a trajectory, but it is still short of a constructive control algorithm without knowledge of a suitable initial field.

3 Negotiating the quantum control landscape

The quantum systems considered here have N -level Hamiltonians H in Eq. (2) consisting of the field-free Hamiltonian H_0 as an $N \times N$ diagonal matrix, and μ as the dipole moment given by an $N \times N$ real symmetric matrix. We will utilize the gradient algorithm to prescribe the path of steepest ascent, thereby forming a natural “rule” to climb the landscape and consequently to navigate the underlying control space. The D-MORPH procedure will be used to implement the gradient search trajectory [30, 31]. The process of climbing the landscape may be confounded by the presence of saddle submanifolds located at certain regions of specific intermediate heights on the landscape. The cartoon in Fig. 1 reflects this issue, and one goal of this work is to examine the degree to which saddles distort control trajectories.

We parameterize a trajectory by a continuous variable $s \geq 0$, such that the control field becomes

$E(s, t)$. As s is increased, the change in the cost functional can be written using the chain rule as

$$\frac{dJ}{ds} = \int_0^T \frac{\delta J}{\delta E(s, t)} \frac{\partial E(s, t)}{\partial s} dt. \quad (8)$$

Maximization of J stipulates that $\frac{dJ}{ds} \geq 0$, which is assured by setting

$$\frac{\partial E(s, t)}{\partial s} = \frac{\delta J}{\delta E(s, t)}. \quad (9)$$

Likewise, when minimizing J , we require $\frac{dJ}{ds} \leq 0$, so in that case we set $\frac{\partial E(s, t)}{\partial s} = -\frac{\delta J}{\delta E(s, t)}$. Minimization applies to J_W , and J_O may be minimized or maximized depending on the nature of the objective operator O . In the present work, the control variables adjusted by the gradient algorithm are the values of the control field $E(s, t_i)$ at a set of discrete time points t_i , $i = 1, 2, \dots$. In order to ensure that each simulation for a particular application traverses the same domain up (or down) the quantum control landscape, every trajectory begins at the same initial height on the landscape J^I and ends at the same final height on the landscape J^F . This criterion is met by using the D-MORPH algorithm [30, 31] to first guide a climb down (up) the landscape to J^I , if an initial trial field yields a value of J that is above (below) J^I . The resultant field at J^I is taken as $E(0, t)$ for the purpose of determining R . The final height on the landscape J^F is then achieved by $E(s = s_{max}, t)$, identified from application of D-MORPH starting with the initial field $E(0, t)$.

Calculation of the gradient $\frac{\delta J}{\delta E(s, t)}$ requires utilization of the relevant cost functional, either J_O or J_W . Sections 3.1 and 3.2 summarize these details, further include the characterization of the critical submanifolds of the landscapes, and introduce a method to determine the distance of a point on the landscape to any particular critical submanifold. In all of the simulations, the assumptions underlying the favorable landscape topology are taken as satisfied [17, 18, 21]; consistent with these assumptions, no landscape trajectories prematurely terminated before reaching J^F in our simulations.

3.1 Quantum ensemble control landscape

The quantum ensemble control landscape $J_O[E(s, t)]$ for optimizing an arbitrary observable with a mixed initial state is given by Eq. (3). The gradient of the objective may be written as [21]

$$\begin{aligned} \frac{\delta J_O}{\delta E(t)} &= \text{Tr} \left\{ O \frac{\delta U(T, 0)}{\delta E(t)} \rho(0) U^\dagger(T, 0) + O U(T, 0) \rho(0) \frac{\delta U^\dagger(T, 0)}{\delta E(t)} \right\} \\ &= \text{Tr} \left\{ O \frac{\delta U(T, 0)}{\delta E(t)} \rho(0) U^\dagger(T, 0) - O U(T, 0) \rho(0) U^\dagger(T, 0) \frac{\delta U(T, 0)}{\delta E(t)} U^\dagger(T, 0) \right\} \\ &= \text{Tr} \left\{ U^\dagger(T, 0) [\rho(T), O] \frac{\delta U(T, 0)}{\delta E(t)} \right\}, \end{aligned} \quad (10)$$

where we have made use of the relation $\frac{\delta}{\delta E(t)} (U^\dagger(T, 0) U(T, 0)) = \frac{\delta U^\dagger(T, 0)}{\delta E(t)} U(T, 0) + U^\dagger(T, 0) \frac{\delta U(T, 0)}{\delta E(t)} = 0$ in the second equality. Climbing the landscape is affected by utilization of Eq. (10) in Eq. (9). This process could possibly be hindered if the trajectory comes near a saddle submanifold where $\frac{\delta J_O}{\delta E(s, t)}$ is small at an intermediate, suboptimal value of J_O . A representative sketch is given in Figure 1, where the optimization path on the left draws near a saddle submanifold, located at the center of the blue landscape surface. In order to assess the impact of encountering saddle manifolds upon the value of R for a control trajectory, we first briefly summarize the nature of the critical submanifolds of the J_O landscape.

At a critical point, the first order variation satisfies $\frac{\delta J_O}{\delta E(t)} = 0$. Thus, examining the last line of Eq. (10), we see that if $\frac{\delta U(T, 0)}{\delta E(t)}$ is of full rank (i.e., the surjectivity assumption (ii) from Section 1 is satisfied), we must have $[\rho(T), O] = 0$ at a critical point [24]. From this relation, it is possible to obtain the form of the critical matrices U_o , and from that information, the distances D to a particular critical submanifold from any point on the landscape; Appendix A contains further details. The distance metric can be used to judge if a control trajectory has strayed close to the saddle submanifolds, and therefore D allows us to assess whether proximity to these manifolds affects the R value of a trajectory.

3.2 Unitary transformation control landscape

In quantum information science, implementing basic logic gate operations necessitates the control of unitary transformations [29]. We therefore consider the optimization of the entire propagator $U(T, 0)$, seeking to attain W , which leads to the unitary transformation cost functional given by Eq. (4). J_W has a minimum value of 0, corresponding to $U = W$, and a maximum of $4N$,

corresponding to $U = -W$, where N is the number of states in the quantum system. Thus, we would like to minimize J_W .

The gradient for the unitary transformation landscape is [2]

$$\begin{aligned} \frac{\delta J_W}{\delta E(t)} &= -\text{Tr} \left\{ W^\dagger \frac{\delta U(T, 0)}{\delta E(t)} + \frac{\delta U^\dagger(T, 0)}{\delta E(t)} W \right\} \\ &= -\text{Tr} \left\{ W^\dagger \frac{\delta U(T, 0)}{\delta E(t)} - U^\dagger(T, 0) \frac{\delta U(T, 0)}{\delta E(t)} U^\dagger(T, 0) W \right\} \\ &= -\text{Tr} \left\{ \left(W^\dagger U(T, 0) - U^\dagger(T, 0) W \right) U^\dagger(T, 0) \frac{\delta U(T, 0)}{\delta E(t)} \right\}. \end{aligned} \quad (11)$$

Optimizations may be carried out on the unitary transformation landscape by combining Eqs. (9) and (11). Trajectories climbing the J_W landscape may pass by saddle submanifolds. In order to assess the effect of these manifolds on R , we briefly review the nature of the critical points of the unitary transformation landscape.

From the last line of Eq. (11), we see that at a critical point, where $\frac{\delta J_W}{\delta E(t)} = 0$, if $\frac{\delta U(T, 0)}{\delta E(t)}$ is of full rank, we must have $W^\dagger U(T, 0) = U^\dagger(T, 0) W$. This implies that $W^\dagger U_o$ is Hermitian. $W^\dagger U_o$ is also a unitary matrix, so at a critical point, the eigenvalues of $W^\dagger U_o$ must be either 1 or -1. As a result, the critical values of J_W in Eq. (4) are $0, 4, \dots, 4N$, implying that there exist $N+1$ critical submanifolds, each of which correspond to a different number of 1's and -1's as the eigenvalues of $W^\dagger U_o$ [18, 19]. A discussion of the distance metric D between any point on the landscape and a particular critical submanifold is contained in Appendix B. As explained in Section 3.1, evaluating D along a landscape climb to reach W can reveal the influence of encountering saddles on the path to an optimal control solution.

4 ‘Straight shot’ eigen-relation

Section 3 laid out how to climb the quantum ensemble and unitary transformation control landscapes using a gradient algorithm, and summarized the topology of the critical points of these landscapes. Here, we present a relation which must be satisfied for ‘straight shot’ control trajectories, utilizing the gradient algorithm for *any* quantum control landscape.

Consider a control trajectory that is perfectly straight, i.e., $R = 1$. Any such ‘straight shot’

trajectory can be written [22]

$$E(s, t) = E(0, t) + \rho(s)\Delta E(t), \quad (12)$$

where $0 \leq \rho(s) \leq 1$ is a monotonically increasing function of s as a result of the gradient algorithm assuring that $dJ/ds \geq 0$ (i.e., without loss of generality we seek to maximize J), and $\Delta E(t) = E(s_{max}, t) - E(0, t)$ depends only on time. Utilizing Eq. (9), the gradient function along this trajectory is

$$\frac{\delta J}{\delta E(s, t)} = \rho'(s) \Delta E(t), \quad (13)$$

where $\rho'(s) = \frac{d\rho(s)}{ds}$.

According to Eq. (9), the instantaneous tangent of a control trajectory, $\frac{\partial E(s, t)}{\partial s}$, “points” in the same direction in control space as the gradient of the quantum control landscape, $\frac{\delta J}{\delta E(s, t)}$, at location s . Furthermore, along a straight control trajectory, the gradient function must point in the *same* direction at *every* location on the trajectory; thus, the gradient functions $\frac{\delta J}{\delta E(s, t)}$ and $\frac{\delta J}{\delta E(s', t)}$ at two different locations s and s' on a straight control trajectory are related by a scale factor. Additionally, for straight trajectories, the difference between any two fields on the trajectory, $E(s, t)$ and $E(s', t)$ is always proportional to the same function $\Delta E(t)$. When using a gradient-based climbing algorithm, the function $\Delta E(t)$ is proportional to the gradient itself, as shown by Eq. (13). These facts allow us to conclude that an eigen-relation exists between the Hessian and the gradient on a straight control trajectory [32, 33].

To make these considerations more specific, we expand the gradient of the cost functional $J[E(s + ds, t)]$ to first order in ds along an arbitrary control trajectory at location s

$$\frac{\delta J}{\delta E(s + ds, t)} = \frac{\delta J}{\delta E(s, t)} + ds \int_0^T \frac{\delta^2 J}{\delta E(s, t') \delta E(s, t)} \frac{\partial E(s, t')}{\partial s} dt'. \quad (14)$$

Here, $\frac{\delta^2 J}{\delta E(s, t') \delta E(s, t)} = \mathcal{H}(s; t, t')$ is the Hessian. Now assume that $R = 1$ for the control trajectory $E(s, t)$. If we bring $\frac{\delta J}{\delta E(s, t)}$ onto the left hand side, divide by ds , and take the limit, we have

$$\lim_{ds \rightarrow 0} \frac{1}{ds} \left(\frac{\delta J}{\delta E(s + ds, t)} - \frac{\delta J}{\delta E(s, t)} \right) = \frac{d}{ds} \frac{\delta J}{\delta E(s, t)} = \rho''(s) \Delta E(t). \quad (15)$$

where Eq. (13) was used for the second equality. Using Eq. (13), we can rewrite the result as $\rho''(s) \Delta E(t) = \frac{\rho''(s)}{\rho'(s)} \frac{\delta J}{\delta E(s, t)}$. On the right hand side of Eq. (14), if we use Eq. (9), we obtain

$$\frac{\rho''(s)}{\rho'(s)} \frac{\delta J}{\delta E(s, t)} = \int_0^T \frac{\delta^2 J}{\delta E(s, t') \delta E(s, t)} \frac{\delta J}{\delta E(s, t)} dt', \quad (16)$$

or more compactly, with the time dependence implicitly understood,

$$\mathcal{H}(s) \cdot g(s) = \frac{\rho''(s)}{\rho'(s)} g(s), \quad (17)$$

where $g(s)$ is the gradient of the cost functional, $\frac{\delta J}{\delta E(s,t)}$, and the scalar product $\mathcal{H}(s) \cdot g(s)$ denotes the integral operation in Eq. (16). As the action of the Hessian on the gradient merely rescales the gradient, this equation states that the gradient is an eigenfunction of the Hessian along a straight control trajectory, with an associated eigenvalue $\frac{\rho''(s)}{\rho'(s)}$. We will illustrate this relationship in Section 5.3. In addition, one may show that repeating the above procedure by expanding the Hessian about the point s and making use of Eq. (12) yields a higher order eigen-relation, stating that the gradient is an eigenfunction of the third-order Hessian. Iterating this process yields a hierarchy of eigen-relations between the gradient and all higher-order derivatives of J ; however, all of these relationships contain no more information than Eq. (17), since the entire hierarchy holds upon satisfaction of Eqs. (9) and (12).

Near a critical point, the eigen-relation is always true, even for trajectories with $R > 1$. To see this relation, expand the gradient of $J[E(s_0 + ds, t)]$ to first order about the location s_0 , which is infinitesimally close to a critical point at $E(s_c, t)$

$$\frac{\delta J}{\delta E(s_c, t)} \equiv \frac{\delta J}{\delta E(s_0 + ds, t)} = \frac{\delta J}{\delta E(s_0, t)} + ds \int_0^T \frac{\delta^2 J}{\delta E(s_0, t') \delta E(s_0, t)} \frac{\partial E(s_0, t')}{\partial s_0} dt' \quad (18)$$

where $ds = s_c - s_0$. Utilizing the gradient rule in Eq. (9) inside the integral and noting that $\frac{\delta J}{\delta E(s_c, t)} = 0$ at a critical point, we have

$$\frac{\delta J}{\delta E(s_0, t)} = -ds \int_0^T \frac{\delta^2 J}{\delta E(s_0, t') \delta E(s_0, t)} \frac{\delta J}{\delta E(s_0, t')} dt'. \quad (19)$$

Equation (19) shows that the gradient function is always an eigenfunction of the Hessian in the immediate neighborhood of a critical point, whether or not the global trajectory is a straight shot. When $R = 1$, however, this behavior extends to the entire control trajectory. From the definition of ds , one can see that if $s_c < s_0$, so that the gradient trajectory is leaving the critical manifold, the eigenvalue in Eq. (19) is positive, and likewise if $s_c > s_0$, so that the trajectory is arriving at the critical manifold, the eigenvalue is negative. This circumstance will be shown in Section 5.3 for the case of a trajectory moving away from the critical point at the bottom of the landscape where $ds < 0$ and then transitioning to the neighborhood of the top of the landscape where $ds > 0$. Thus, along a straight shot from the bottom to the top of the landscape, the Hessian eigenvalue

associated with the gradient will change sign (i.e., at least once, and more often if saddle critical points are closely approached) at some point along the trajectory. In addition, the eigen-relation is nominally an independent consideration at each value of s . However, the diffeomorphic nature of the gradient climb over s in Eq. (9) implies that a smooth eigen-relation should exist, regardless of which eigenvector (corresponding to the gradient function) initiates the climb. This behavior is affirmed in the simulations in Section 5.3.

We note that the relations examined in this section were derived without making use of quantum mechanics, as the analysis leading to the gradient being an eigenfunction of the Hessian only utilizes the gradient climbing rule, Eq. (9), and the straight shot condition Eq. (12). The interest in this paper is directed towards applications where $J[E(s, t)]$ arises in a quantum control context.

5 Illustrations

This section provides numerical illustrations of the landscape structure analysis tools outlined in Sections 2-4. Section 5.1 shows that for randomly chosen control trajectories, the distribution of R takes on very modest values, which did not exceed 2.0 in the present examples. The results also demonstrate that control trajectories with either very low or very high values of R are spread in a like fashion throughout control space. Section 5.2 addresses the effect of saddle submanifolds upon optimization paths. Encountering saddle critical submanifolds on the way from the bottom to the top along an optimization path tends to modestly increase the associated R value. Finally, Section 5.3 provides numerical evidence for the validity of the Hessian-gradient eigenrelation in Section 4. Dimensionless units are used throughout the simulations.

5.1 Random sampling of control trajectories

In order to assess the structural complexity of the quantum ensemble and unitary transformation control landscapes, we recorded the R values of landscape trajectories constructed using randomly chosen initial control fields.

5.1.1 Quantum ensemble control landscape

Here we consider several eight-level quantum systems to illustrate the landscape structural findings. As the nature of a landscape's saddle manifolds is determined by the number and degree

of degeneracies of the initial density matrix and observable matrix [21, 23, 25], we will consider combinations of matrices with varying degrees of degeneracy so as to examine a broad class of landscapes. Increasing numbers of nondegenerate eigenvalues generally correspond to larger numbers of saddle submanifolds. In particular, we are interested in the effects of such submanifolds upon landscape structure. Accordingly, five different quantum ensemble control landscapes are sampled using the ρ and O matrices presented in Table 1, and 1000 optimizations of J_O are performed over each landscape. The Hamiltonian in Table 1 is utilized together with the dipole matrix

$$\mu = \begin{pmatrix} 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 & \pm 0.5^3 & \pm 0.5^4 & \pm 0.5^5 & \pm 0.5^6 \\ \pm 1 & 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 & \pm 0.5^3 & \pm 0.5^4 & \pm 0.5^5 \\ \pm 0.5 & \pm 1 & 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 & \pm 0.5^3 & \pm 0.5^4 \\ \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 & \pm 0.5^3 \\ \pm 0.5^3 & \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 \\ \pm 0.5^4 & \pm 0.5^3 & \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 & \pm 1 & \pm 0.5 \\ \pm 0.5^5 & \pm 0.5^4 & \pm 0.5^3 & \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 & \pm 1 \\ \pm 0.5^6 & \pm 0.5^5 & \pm 0.5^4 & \pm 0.5^3 & \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 \end{pmatrix} \quad (20)$$

where the signs on the dipole matrix elements were chosen randomly under the constraint of μ remaining symmetric. Once these signs were chosen, they remained fixed for the 1000 optimizations over each individual landscape.

The density matrix ρ_1 , which corresponds to the case of a pure initial state, represents an almost completely degenerate density matrix. In contrast, ρ_2 contains two eigenvalues which are both four-fold degenerate. The observables used in conjunction with these density matrices O_1 and O_2 . As these observable matrices also exhibit varying levels of degeneracy, the pairs ρ_1 and O_1 , ρ_1 and O_2 , ρ_2 and O_1 , ρ_2 and O_2 are used to generate four distinct landscapes. Lastly, we also utilized ρ_3 and O_3 as an example of matrices with fully nondegenerate spectra to generate a fifth landscape.

The initial fields of these optimizations were parametrized as

$$E(t) = \frac{1}{F} \exp[-0.3(t - \frac{T}{2})^2] \sum_{n=1}^M a_n \sin(\omega_n t + \phi_n), \quad (21)$$

with the target time T being 10. The amplitudes a_n and phases ϕ_n were chosen randomly from the uniform distributions $[0, 1]$ and $[0, 2\pi]$, respectively. The frequencies are $\omega_n = n$ and $M = 60$ to ensure that the initial control field was nearly resonant with every possible transition in the

Table 1: Quantum ensemble landscape matrices

Matrix	Diagonal entries
ρ_1	(1, 0, 0, 0, 0, 0, 0, 0)
ρ_2	($\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0, 0, 0$)
ρ_3	($\frac{7}{28}, \frac{6}{28}, \frac{5}{28}, \frac{4}{28}, \frac{3}{28}, \frac{2}{28}, \frac{1}{28}, 0$)
O_1	(0, 0, 0, 0, 0, 0, $\frac{4}{9}, \frac{5}{9}$)
O_2	(0, 0, 0, 0, $\frac{4}{17}, \frac{4}{17}, \frac{4}{17}, \frac{5}{17}$)
O_3	(0, $\frac{1}{28}, \frac{2}{28}, \frac{3}{28}, \frac{4}{28}, \frac{5}{28}, \frac{6}{28}, \frac{7}{28}$)
H_0	(-10, -8, -4, 2, 10, 20, 32, 46)

Hamiltonian in Table 1. F is a normalization factor, picked so that the initial fluence of each field is 1. This normalization was chosen to avoid operating with strong initial fields. After the initial choice of the control field, the time interval $[0, T]$ was discretized into 1001 evenly spaced time points, and the value of the field at each of these time points served as the control variables. We define the limits of the optimizations on the landscape as follows. We denote J_O^{\max} as the maximum value that $J_O = \text{Tr}(\rho(T)O)$ can attain for a particular pair of $\rho(0)$ and O , and J_O^{\min} denotes the minimum value. When a random field drawn from Eq. (21) is used to calculate $J_O = \text{Tr}(\rho(T)O)$, if $J_O < J_O^I = J_O^{\min} + 0.01[J_O^{\max} - J_O^{\min}]$, then D-MORPH is used to climb until $J_O = J_O^I$. Likewise, if $J_O > J_O^I$, D-MORPH is used to descend the landscape until $J_O = J_O^I$. Then, the resultant control field is optimized using D-MORPH until $J_O = J_O^F = J_O^{\max} - 0.01[J_O^{\max} - J_O^{\min}]$. This latter trajectory, which takes the field from a height of J_O^I to a height of J_O^F on the landscape, is used to calculate R . The large window from $J_O^I \rightarrow J_O^F$ is chosen so that each of the optimizations proceeds over a wide swathe of the landscape and thereby encounters as many structural features as possible. In this regard, we have numerically shown that for the state-to-state transition probability landscape, pushing the fidelity of J closer to either 0 or 1 at the beginning and end of the optimizations, respectively, has a diminutive effect on R [22].

Histograms of the resulting R values from 1000 optimizations over the five different landscapes are presented in Figure 2 for the cases (a) ρ_1 and O_1 , (b) ρ_1 and O_2 , (c) ρ_2 and O_1 , (d) ρ_2 and O_2 , and (e) ρ_3 and O_3 . The distributions are skewed towards the right, indicating that the R values

tend to accumulate near 1.0. The smallest values are found in the histogram of Figure 2(a) which has a mean of $R = 1.22$, followed by Figures 2(b) and 2(c), which have mean values of $R = 1.31$ and $R = 1.33$, respectively. The highest values of R are found in the histograms of Figs. 2(d) and 2(e), which have mean values of $R = 1.65$ and $R = 1.53$, respectively. The distinctions between the encountered structural features are modest over all of the cases examined. In an absolute sense, all of the trajectories produced $R < 2.0$ indicating only very mild landscape structural features (considering that there is no *a priori* reason for R to be close to 1). Prior experience with state-to-state transition probability landscapes suggests that ‘straight shot’ trajectories are reserved for cases with $R < 1.05$, and a search is generally needed to find a proper $E(0, t)$ to initiate such a trajectory [22].

A natural question is whether control trajectories that possess low or high R values are concentrated in some particular region(s) of control space. As the initial controls in the statistical sampling considered above were randomly chosen, this assessment directly bears on the investigation of the range of R values encountered. To address this question, Euclidean distances were calculated between (1) all pairs of initial fields, (2) all pairs of final fields, and (3) all initial and final fields from the collection of control trajectories for the landscape associated with ρ_1 and O_1 . These distances were calculated using

$$d_{EL}^{ij} = \left[\frac{1}{T} \int_0^T [E_i(t) - E_j(t)]^2 dt \right]^{\frac{1}{2}} \quad (22)$$

for control fields E_i and E_j , and the results are displayed in Figure 3. Importantly, in computing these distances, we first used only the 250 trajectories with the lowest values of R out of the 1000 trajectories that were sampled in generating Figure 2(a). Then, separately, we used the 250 trajectories with the highest values of R . In particular, trajectories with $R < 1.173$ were used to generate the histograms in Figure 3(a), and likewise with $R > 1.256$ for the histograms in Figure 3(b); in both cases the histograms show all possible pairs of distances, with $250 \times 249/2 \sim 3.1 \times 10^4$ pairs. The distribution of distances between initial fields, in blue, is centered to the left and narrower than the distribution corresponding to the distances between optimal fields, in red. This behavior implies that the optimal fields are scattered more widely throughout control space than the randomly picked initial fields. The histogram in green of the distances between all initial-final field pairs is centered between the other two distributions. The most significant feature in Figure 3 is that the distributions are very similar, whether generated from trajectories

with low values of R (Figure 3(a)) or from trajectories with high values of R (Figure 3(b)). This result suggests that trajectories with low and high values of R are distributed in the same way across control space. Thus, straighter, more direct trajectories and those less so are likely to be equidistant from a trajectory originating from a random initial control field. Similar results have been observed in other cases with different ρ and O matrices.

5.1.2 Unitary transformation control landscape

Here we consider a four-level quantum system for the target unitary transformation

$$W = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (23)$$

with the Hamiltonian and dipole matrices being

$$H_0 = \begin{pmatrix} -10 & 0 & 0 & 0 \\ 0 & -7 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 8 \end{pmatrix} \quad \mu = \begin{pmatrix} 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 \\ \pm 1 & 0 & \pm 1 & \pm 0.5 \\ \pm 0.5 & \pm 1 & 0 & \pm 1 \\ \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 \end{pmatrix} \quad (24)$$

where the random signs on the dipole matrix were again chosen so that it remained symmetric. A total of 1000 initial fields were randomly chosen according to Eq. (21), with $T = 10$ and M set to 20. The initial fields were adjusted as explained in Section 5.1.1 to arrive at a value of $J_W^I = 0.99 \times 16 = 15.84$, and the final achieved value after optimization was $J_W^F = 0.01 \times 16 = 0.16$. Thus, in this case we climb down the landscape as a minimization process. Figure 4 is a histogram of the R values collected in this fashion. The plot is skewed right with a mean of $R = 1.41$, indicating that the unitary transformation control landscape contains minimally complex structural features. The quantum ensemble landscapes considered earlier were for eight level systems, while this unitary transformation landscape is for a four level system. Together with previous work which observed that typical R values for generating unitary transformations scales linearly with the number of states N [27], these findings suggest that the unitary transformation landscape appears to be structurally similar to the ensemble control landscape with fully nondegenerate ρ and O matrices.

Figure 5 displays the histograms resulting from calculating the pairwise distances between the

control fields using Eq. (22), as explained in Section 5.1.1. The two sets of control trajectories, each of size 250, were used to produce the histograms in Figure 5(a) and 5(b) which satisfy $R < 1.34$ and $R > 1.48$, respectively. Significantly, the histograms computed with the lowest values of R and with the highest values of R are almost identical, indicating that trajectories with varying values of R for generating unitary transformations are distributed in the same way across control space, as found in Figure 3 for the ensemble control landscape. This result was also found for a five level state-to-state transition probability landscape [22], suggesting that similar distributions of fields for high and low R trajectories is a general property of quantum control landscapes.

5.2 Interactions of optimization trajectories with saddle submanifolds

Although the saddle submanifolds of the quantum ensemble and unitary transformation landscapes are not able to prevent reaching the global optimum, they may still significantly influence the optimization trajectories. In order to examine the extent to which saddle submanifolds affect the climb up the landscape, the metrics set out in Section 3.1.3 and 3.2.3 (and Appendices A and B, respectively) were used to calculate the distance from each point along a landscape climb trajectory to all critical submanifolds. As examples, the quantum ensemble control landscape was generated by the initial density matrix ρ_2 and observable operator O_1 from Table 1, while the unitary control landscape was generated with W in Eq. (23).

5.2.1 Quantum ensemble control landscape

The values of the critical submanifolds of the system on the quantum ensemble control landscape [21] of ρ_2 and O_1 are given by $J_O^{\min} = 0$, $J_O^1 = \frac{1}{9}$, $J_O^2 = \frac{5}{36}$, and $J_O^{\max} = \frac{1}{4}$. Here, J_O^1 and J_O^2 correspond to saddles, although a trajectory passing through these values may not actually come close to the associated submanifold of saddle controls. The normalized distances D to these four critical submanifolds, using Eq. (28), are computed during the optimizations with the lowest and highest R values out of the 1000 runs performed, which were $R = 1.17$ and $R = 1.59$, respectively. The results are shown in Figures 6(a) and 6(b). The magnitude of the gradient for each trajectory, scaled down by a factor of 60, is also plotted to show the rate of climbing, as an optimization typically slows down near a critical submanifold. In Figure 6(a), the run with a small R value does not closely approach any of the saddle submanifolds. This observation is corroborated by examining the magnitude of

the gradient, which falls only slightly in the middle of the climb. The gradient is naturally small near the beginning and end of the climb, where the trajectory leaves (approaches) the critical submanifolds corresponding to the bottom (top) of the landscape. In contrast, Figure 6(b) shows that the run with a high R value hovers in the vicinity of the two saddle submanifolds for extended periods during the climb, indicated by the distance to the first and second saddle manifolds, D_1 and D_2 , respectively, becoming very small during portions of the run. Concurrently, the magnitude of the gradient falls sharply in those two areas, recovering only when D_2 starts to rise, thus confirming that the run passed closely by these two saddle submanifolds. Comparison of Figures 6(a) and 6(b) suggests that control trajectories that are ‘diverted’ to pass closely by saddle submanifolds can have higher R values.

5.2.2 Unitary transformation control landscape

Figures 6(c) and 6(d) show the distances to the critical submanifolds while traversing the unitary transformation control landscape using Eq. (29); the magnitude of the gradient (scaled down by a factor of 20) is also shown. D_1 , D_2 , and D_3 denote the normalized distances to the saddles at landscape heights of $J_W^1 = 12$, $J_W^2 = 8$, and $J_W^3 = 4$, respectively, while D_{\min} and D_{\max} denote the normalized distances to the bottom and top critical submanifolds at $J_W^{\min} = 16$ and $J_W^{\max} = 0$, respectively. Figure 6(c) corresponds to a run with a low R value of 1.15, and Figure 6(d) corresponds to a run with a high R value of 1.73. The distances to the saddle submanifolds D_1 and D_3 assume lower values during the high R run (as compared to the low R run), and correspondingly the magnitude of the gradient also dips when the distances to D_1 and D_3 are at a minimum. Neither of the trajectories approach the second saddle submanifold, as indicated by the behavior of D_2 on the trajectory. These plots also support the notion that saddle submanifolds are capable of attracting nearby paths of steepest descent away from proceeding in a direct manner to a global minimum.

5.3 Hessian-gradient eigen-relation

The results of Section 4 identify that the gradient of the cost functional should be an eigenfunction of the Hessian along a straight trajectory. Here, we provide evidence to verify the Hessian-gradient eigen-relation for nearly straight shots found in simulations. In order to illustrate this phenomenon

as clearly as possible, we consider a simple three level quantum system with Hamiltonian and dipole matrix

$$H_0 = \begin{pmatrix} -10 & 0 & 0 \\ 0 & -5 & 0 \\ 0 & 0 & 5 \end{pmatrix} \quad \mu = \begin{pmatrix} 0 & -1 & -0.5 \\ -1 & 0 & 1 \\ -0.5 & 1 & 0 \end{pmatrix}, \quad (25)$$

and we use the cost functional of Eq. (3), with

$$\rho = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad O = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (26)$$

corresponding to maximizing the pure state transition probability $P_{1 \rightarrow 3}(T)$ with $T = 10$. This landscape contains no saddle submanifolds, and therefore the control trajectories tend to have smaller values of R . We used the stochastic Particle Swarm Optimization (PSO) algorithm [34, 35] to locate an initial control $E(0, t)$ taking the form in Eq. (21) with $M = 20$ that produced a nearly straight trajectory with $R = 1.002$. In order to assess whether a lower bound for R (other than 1.0) might exist, a Derandomized Evolution Strategy (ES) [36, 37] was used in a more extensive search for a straight path. Using this algorithm together with an n -point Gauss-Legendre quadrature rule (with $n = 1000$) to accurately evaluate the integrals in Eqs. (5) and (6) allowed reaching an even lower value of $R = 1.00003$. As this result was at the limit of our numerical accuracy, we cannot rule out the prospect of finding a still lower value of R . Below, the eigen-relation is tested with the former conservative near straight trajectory of $R = 1.002$.

We seek to affirm Eq. (17), which states that the gradient function aligns with some eigenfunction of the Hessian. As in Eq. (17), we simplify notation with the Hessian and gradient given by $\mathcal{H}(s)$ and $g(s)$, respectively, and a ‘dot’ \cdot denotes integration over time $t \in [0, T]$. Thus, we calculated $\hat{g}(s) \cdot \mathcal{G}(s)$, where $\mathcal{G}(s) = \mathcal{H}(s) \cdot \hat{g}(s)$ is the function resulting from the action of the Hessian on the unit gradient function, $\hat{g}(s) = \frac{\delta P_{1 \rightarrow 3}}{\delta E(s, t)} / \left(\frac{\delta P_{1 \rightarrow 3}}{\delta E(s, t)} \cdot \frac{\delta P_{1 \rightarrow 3}}{\delta E(s, t)} \right)^{\frac{1}{2}}$. In Figure 7, $\hat{g} \cdot \mathcal{G}(s) = \hat{g} \cdot \mathcal{H} \cdot \hat{g}$ is plotted versus s as the thick black circles, along with every eigenvalue of the Hessian shown as the smaller green circles. Intriguingly, at every point on the climb, the black circle overlaps almost exactly with one of the green circles, suggesting that the gradient function can be identified with some eigenfunction of the Hessian. The sign of $\hat{g}(s) \cdot \mathcal{G}(s)$ reveals that \mathcal{G} and g are parallel at the beginning of the climb and antiparallel at the end, which is consistent with the assessment discussed

below Eq. (19)

We will now illustrate the validity of Eq. (17), rewritten in the notation here as $\hat{g} \cdot \mathcal{H} \cdot \hat{g} \equiv \hat{g} \cdot \mathcal{G} = \rho''(s)/\rho'(s)$. This assessment is performed by replotting $\hat{g}(s) \cdot \mathcal{G}(s)$ versus s in Figure 8 along with $\rho''(s)/\rho'(s)$, calculated for this control trajectory. According to Eq. (13), for a $R = 1$ straight control trajectory, $\frac{\delta P_{1 \rightarrow 3}}{\delta E(s,t)} = \rho'(s) \Delta E(t)$. Since the trajectory considered here is not exactly straight with $R = 1.002$, we determined the average value of $\rho'(s)$ by setting $\rho'(s) = \frac{1}{T} \int_0^T \frac{\delta P_{1 \rightarrow 3}}{\delta E(s,t)} \frac{1}{\Delta E(t)} dt$, where $\Delta E(t) = E(s_{max}, t) - E(0, t)$ and $\rho''(s)$ was then obtained by numerically differentiating $\rho'(s)$. Examining Figure 8 reveals that $\hat{g} \cdot \mathcal{G} \simeq \rho''(s)/\rho'(s)$ to good accuracy, providing evidence for the validity of the eigen-relation derived in Section 4.

6 Conclusions

The exploration of landscape structure first performed for the state-to-state transition probability [22] is extended in this paper to more general quantum ensemble and unitary transformation control landscapes. These collective results point towards the existence of simple quantum control landscape structure, which was first encountered experimentally [26] as well as hinted at by preliminary numerical studies [27, 28].

A special feature of the landscapes considered in this work is the presence of critical submanifolds that possess saddle topology. These submanifolds do not prevent optimization trajectories from reaching the top of the landscape, but they can have an impact on the search effort, reflected in the identified R value. We find that the presence of saddle submanifolds on the landscape has a small but discernible effect on the ease of optimization reflected in the distortion of the control trajectories. Importantly, the appearance of saddle submanifolds did not cause R to take on values in excess of 2.0 for the systems examined here, indicating that these quantum control landscapes, in addition to state-to-state transition probability landscapes [22], possess markedly simple structure.

A condition for the existence of control trajectories with $R = 1.0$, which would indicate exceedingly simple landscape structure, was also derived. It was found that in order for the gradient algorithm to yield a straight trajectory, the gradient function must not only be separable in s and t , but must also be an eigenfunction of the Hessian. This situation also implies that all higher-order derivatives of the cost functional possess a common eigenfunction. The eigen-relation may form a basis to identify the conditions under which exactly straight control trajectories derived from the

gradient algorithm exist. Additional simulations exploring landscape structure on a wider variety of landscapes would also allow for a better understanding of the conditions that lead to low values of R . One direction for further investigation is whether achieving control of unitary transformations, which requires control over every matrix element in the propagator, leads to higher values of R than even for the fully nondegenerate ensemble control landscapes. A simple assessment of this matter was performed for a two level system using an ES algorithm [36, 37] to search for straight trajectories. For two-level systems, in the cases of maximizing $P_{1 \rightarrow 2}$ we find $R - 1 \leq 10^{-6}$, while for unitary transformation control $R = 1.0016$ was achieved.

Quantum control landscapes are expected to nominally be very complex, as they are defined on extremely high dimensional spaces and represent the generally nonlinear relationship, encoded through Schrödinger's equation, between the unitary propagator and a dynamical set of controls. Thus, a natural expectation is that trajectories traveling over these landscapes would require considerable effort to reach an optimal outcome. However, the topology of quantum control landscapes is very simple when certain reasonable physical assumptions are met, providing one factor to explain the relative ease of performing optimizations [4, 17, 38, 39]. These topological results are complemented by the finding in this work of simple landscape structure. This study, along with Ref. [22], supports the notion that simple structure is a universal and robust property of quantum control landscapes. Collectively, the twofold simplicity of landscape features, encapsulated by topology and structure, provides a foundation for more clearly explaining the widespread efficiency of finding optimal controls over quantum phenomena.

Appendix A Critical submanifolds of the quantum ensemble control landscape

This Appendix gives a brief overview of the critical submanifolds of the quantum ensemble control landscape, along with a metric which can be used to determine the distance to all critical submanifolds from any given point in control space. A complete characterization of the critical submanifolds is given in Ref. [23].

The analysis of Section 3.1 reveals that at a critical point, $[O, \rho(T)] = 0$. This condition implies that the number of critical submanifolds is in one-to-one correspondence with the number of ways the eigenvalues of $\rho = \rho(0)$ can be permuted such that their overlaps with the eigenvalues of O are distinct. These permutations are induced by critical propagators $U(T, 0)$, each of which belongs to a distinct critical submanifold, and can be enumerated using contingency tables [23]. These tables possess n rows and m columns and can be constructed as follows. First, ρ and O are expressed in diagonal form with their eigenvalues in descending order. Then, a particular permutation is applied to the eigenvalues of ρ ; such a permutation matrix U_o is associated with the said critical submanifold. The number of times the eigenvalue λ_i of ρ overlaps with the eigenvalue ϵ_j of O , after the permutation is applied, is the value c_{ij} of the (i, j) entry in the contingency table. The rows of the contingency table must add up to the sum of the degeneracies of the eigenvalues of O , and the columns of the table must add up to the sum of the degeneracies of the eigenvalues of ρ . This fact can be used to easily construct all possible contingency tables and therefore reveal every critical submanifold along with its associated height on the landscape.

The contingency tables can also be used to calculate a distance to a particular critical submanifold from any point on the landscape. The distance is expressed in terms of the unitary transformation $U[E(t)]$ corresponding to a particular control field $E(t)$. A short exposition of this metric is given here; a full discussion is given in Ref. [40]. Consider a point on the landscape generated by U_o belonging to some critical submanifold; i.e, U_o is a permutation matrix. The matrix U_o may be divided into blocks $\{U_o^{ij}\}$, starting in the upper left corner, of sizes $\{o_i \times p_j\}$, where o_i is the degeneracy of the i th eigenvalue of O , and p_j is the degeneracy of the j th eigenvalue of $\rho(0)$, with the eigenvalues arranged in descending order. It can be shown that the singular value decomposition of each block contains only 0s and 1s, and that for the ij th block of size $o_i \times p_j$ the number of 1s is equal to c_{ij} , which is the entry in the i th row and j th column in the contingency

table. This fact can be used to specify the distance to each critical submanifold. For each block and its associated entry c_{ij} , define an $N \times N$ diagonal matrix S_{ij} , which has c_{ij} values of 1 starting in the upper left corner in the diagonal:

$$S_{ij} = \begin{pmatrix} \mathbb{I}_{c_{ij}} & \\ & 0 \end{pmatrix}. \quad (27)$$

To find the distance of an arbitrary unitary matrix U from the critical submanifold U_o , we divide U into blocks $\{U^{ij}\}$ in the manner described above. Each block U^{ij} of U is written in its singular form Σ_{ij} , with the singular values on the diagonal sorted in descending order. For each block in U , we define $V_{ij} = S_{ij} - \Sigma_{ij}$. The distance from U to the critical submanifold of U_o is then

$$D = \sum_{i=1}^p \sum_{j=1}^o \text{Tr} \left(V_{ij}^\dagger V_{ij} \right), \quad (28)$$

where p and o are the number of distinct eigenvalues of $\rho(0)$ and O , respectively. It is convenient to normalize this metric [41]. To do so, the unitary matrix U can be assigned a vector consisting of the singular values of all of its blocks $\{U^{ij}\}$. Since the squares of these singular values sum to N , each vector can be thought of as lying in the spherical polygon spanned by the vectors corresponding to the U_o matrices of the critical submanifolds. This formulation implies that the point which is the greatest distance from a critical submanifold is another vertex, which itself corresponds to another critical submanifold. By calculating the pairwise distances between all the critical submanifolds, we can find the maximum distance any point can be from a particular critical submanifold, and then normalize the metric by dividing all distances to that critical submanifold by the maximum distance, thereby giving $0 \leq D \leq 1$. This normalization is used in the results reported in Figures 6(a) and 6(b).

Appendix B Critical submanifolds of the unitary transformation control landscape

Here we outline a distance metric for the unitary transformation control landscape; full details are given in Ref. [40]. Recall, from Section 3.2 that for a unitary transformation U_o corresponding to a critical point, the eigenvalues of $W^\dagger U_o$ are ± 1 . The distance from an arbitrary unitary matrix U to a submanifold containing the critical point U_o can be specified by noting that the eigenvalues of

$W^\dagger U$ take values on the unit circle in the complex plane. The i th eigenvalue of $W^\dagger U$ is λ_i , and the eigenvalues are sorted in ascending order by their real parts. To find the distance to the critical submanifold which corresponds to α eigenvalues of -1 and $N - \alpha$ eigenvalues of 1 for $W^\dagger U_o$, we compare the real part of the smaller α eigenvalues of $W^\dagger U$ with -1, and the real part of the larger $N - \alpha$ eigenvalues with 1 [40]. The distance between U and the critical submanifold containing U_o then is

$$D = \left[\sum_{i=0}^{\alpha} (1 + \operatorname{Re}\{\lambda_i\}) + \sum_{\alpha}^{N-1} (1 - \operatorname{Re}\{\lambda_i\}) \right] / 2N. \quad (29)$$

The maximum distance to a critical submanifold in this case is $2N$, and the denominator of Eq. (29) assures that $0 \leq D \leq 1$.

Acknowledgments

We thank Greg Riviello and Julia Yan for helpful discussions. A.N. acknowledges support from the Program in Plasma Science and Technology at Princeton University. We also acknowledge support from ARO Grant No. W911NF-13-1-0237, ARO-MURI Grant No. W911NF-11-1-2068, DOE Grant No. DE-FG02-02ER15344, and NSF Grant No. CHE-1058644.

References

- [1] G. G. Balint-Kurti, S. Zou, and A. Brown, *Advances in Chemical Physics*, page 43. Wiley, 2008.
- [2] T.-S. Ho, J. Dominy, and H. Rabitz, *Phys. Rev. A* **79**, 013422 (2009).
- [3] D. Dong and I. R. Peterson, *IET Control Theory Appl.* **4**, 2651 (2010).
- [4] K. W. Moore and H. Rabitz, *Phys. Rev. A* **84**, 012109 (2011).
- [5] K. W. M. Tibbetts, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T.-S. Ho, R. B. Wu, and H. Rabitz, *Phys. Rev. A* **86**, 062309 (2012).
- [6] T. Blasi, M. F. Borunda, E. Räsänen, and E. J. Heller, *Phys. Rev. B* **87** 241303 (2013).
- [7] M. Nest, M. Ludwig, I. Ulusoy, T. Klamroth, and P. Saalfrank, *J. Chem. Phys.* **138**, 164108 (2013).

- [8] P. Doria, T. Calarco, and S. Montangero, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [9] A. Rahmani, T. Kitagawa, E. Demler, and C. Chamon, *Phys. Rev. A* **87**, 043607 (2013).
- [10] E. Räsänen and E. J. Heller, *Eur. Phys. J. B* **86**, 1 (2013).
- [11] C. Brif, M. D. Grace, M. Sarovar, and K. C. Young, arXiv:1310:3443v1.
- [12] R. S. Judson and H. Rabitz, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [13] A. Weiner, *Rev. Sci. Instrumen.* **71**, 1929 (2000).
- [14] R. Bücker, T. Berrada, S. van Frank, J. Schaff, T. Schumm, J. Schmiedmayer, G. Jäger, J. Grond, and U. Hohenester, *J. Phys. B* **46**, 104012 (2013).
- [15] N. X. Truong, P. Hulse, S. Göde, A. Przystawik, T. Döppner, Th. Fennel, Th. Bornath, J. Tiggesbäumler, M. Schlanges, G. Gerber, and K. H. Meiwes-Broer, *Phys. Rev. A* **81**, 013201 (2010).
- [16] R. Hildner, D. Brinks, J. B. Nieder, R. J. Cogdell, and N. F. van Hulst, *Science* **340**, 1448 (2013).
- [17] R. Chakrabarti and H. Rabitz, *Int. Rev. Phys. Chem.* **26** 671 (2007).
- [18] H. Rabitz, M. Hsieh, and C. Rosenthal, *Phys. Rev. A* **72**, 052337 (2005).
- [19] M. Hsieh and H. Rabitz, *Phys. Rev. A* **77**, 042306 (2006).
- [20] K. W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz, *Phys. Rev. A* **83**, 012326 (2011).
- [21] T.-S. Ho and H. Rabitz, *J. Photochem. Photobiol. A* **180**, 226 (2006).
- [22] A. Nanduri, A. Donovan, T.-S. Ho, and H. Rabitz, *Phys. Rev. A* **88**, 033425 (2013).
- [23] R. Wu, H. Rabitz, and M. Hsieh, *J. Phys. A*, **41**, 015006 (2008).
- [24] R.-B. Wu, R. Long, J. Dominy, T.-S. Ho, and H. Rabitz, *Phys. Rev. A* **86**, 013405 (2012).
- [25] M. Hsieh, R. Wu, and H. Rabitz, *J. Chem. Phys.* **130**, 104109 (2009).
- [26] J. Roslund and H. Rabitz, *Phys. Rev. A* **80**, 013408 (2009).

- [27] K. W. Moore, M. Hsieh, and H. Rabitz, *J. Chem Phys.* **128**, 154117 (2008).
- [28] A. Donovan, V. Beltrani, and H. Rabitz, *Phys. Chem. Chem. Phys.* **13**, 7348 (2011).
- [29] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2004.
- [30] A. Rothman, T.-S. Ho, and H. Rabitz, *Phys. Rev. A* **72**, 023416 (2005).
- [31] A. Rothman, T.-S. Ho, and H. Rabitz, *J. Chem. Phys.* **123**, 134104 (2005).
- [32] R. Magnanini, *Appl. Anal.* **41**, 171 (1991).
- [33] D. Drucker and S. A. Williams, *Am. Math. Mon.* **110**, 869 (2003).
- [34] J. Kennedy and R. Eberhart, in *Proceedings of the IEEE International Conference on Neural Networks*, Perth, WA, 1995 (IEEE, Piscataway, NJ, 1995), Vol. 4, p. 1942.
- [35] J. Kennedy and R. Eberhart, *Swarm Intelligence* (Morgan Kaufman, San Francisco, CA, 2001).
- [36] N. Hansen and A. Ostermeier, *Evol. Comput.* **9**, 159 (2001).
- [37] R. Ros and N. Hansen, *Lect. Notes Comput. Sc.* **5199**, 296 (2008).
- [38] H. Rabitz, M. Hsieh, and C. Rosenthal, *Science* **303**, 1998 (2004).
- [39] H. Rabitz, T.-S. Ho, M. Hsieh, R. Kosut, and M. Demiralp, *Phys. Rev. A* **74**, 012721 (2006).
- [40] Q. Sun, R. Wu, G. Riviello, and H. Rabitz (to be published).
- [41] G. Riviello and H. Rabitz, (to be published).

Figures

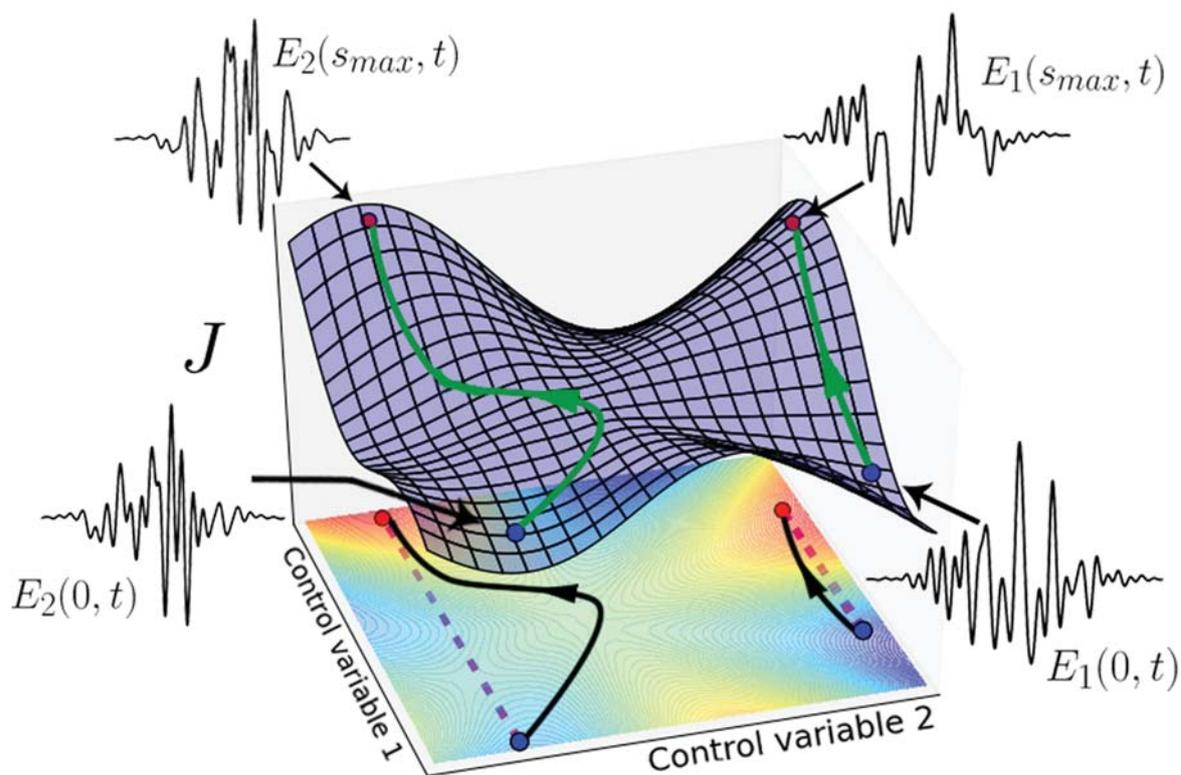


Figure 1: Illustrative sketch of a quantum control landscape (blue surface) and its underlying control space (colored contour map), also displayed as the projection onto the plane spanned by two control variables (in practice many more are used). The vertical axis represents a cost functional J subject to maximization. The two optimization paths shown, in green, differ in the nature of their R values. The path on the right, proceeding from $E_1(0, t)$ to $E_1(s_{max}, t)$, climbs in a direct manner to the top of the control landscape. Consequently, its projection into control space, in black, has a path length very close to the Euclidean distance along the straight line between the projection's endpoints, shown as a magenta dashed line. Thus, this path's R value is low. On the other hand, the optimization path on the left, from $E_2(0, t)$ to $E_2(s_{max}, t)$, is attracted away from proceeding directly to a maximum by the saddle point at the center of the landscape. As a result, its projection into control space is much longer than the straight line between its endpoints, and so this path possesses a higher R value. The paths do not strictly terminate at the top and bottom of the landscape, as optimizations generally begin and end at values of the cost functional which approximate J^{min} and J^{max} , respectively, to acceptable precision.

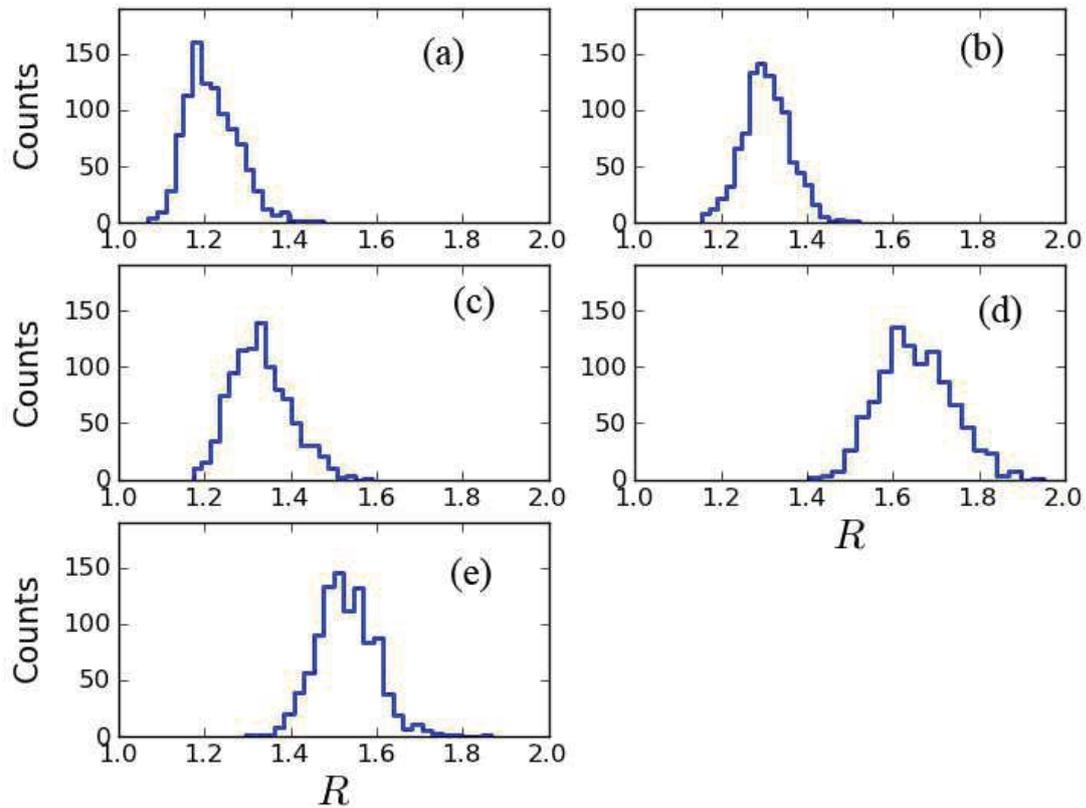


Figure 2: The distribution of R values for five different quantum ensemble control landscapes. The cases are generated using: (a) ρ_1 and O_1 ; (b) ρ_1 and O_2 ; (c) ρ_2 and O_1 ; (d) ρ_2 and O_2 ; and (e) ρ_3 and O_3 . For each of these runs, the Hamiltonian H_0 of Table 1 and dipole μ of Eq. (20) were used. All of the distributions lean towards the left, showing that R tends to accumulate near 1.0. 1000 optimizations were carried out for each landscape.

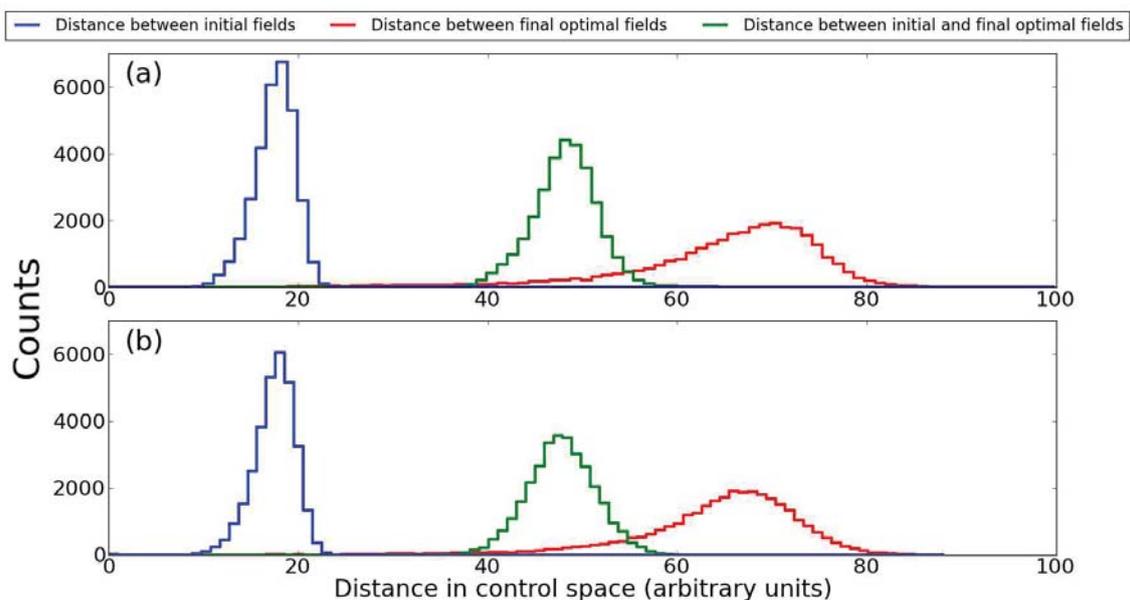


Figure 3: The distributions of pairwise distances between the initial fields, final fields, and initial-final fields for the landscape generated by ρ_1 and O_1 from Table 1, respectively. The distances are calculated with Eq. (22). Plot (a) was produced using the 250 optimizations with the lowest values of R , while plot (b) was produced using the 250 optimizations that had the highest values of R . The similarity between the two plots suggests that low R and high R trajectories are distributed in essentially the same way throughout control space.

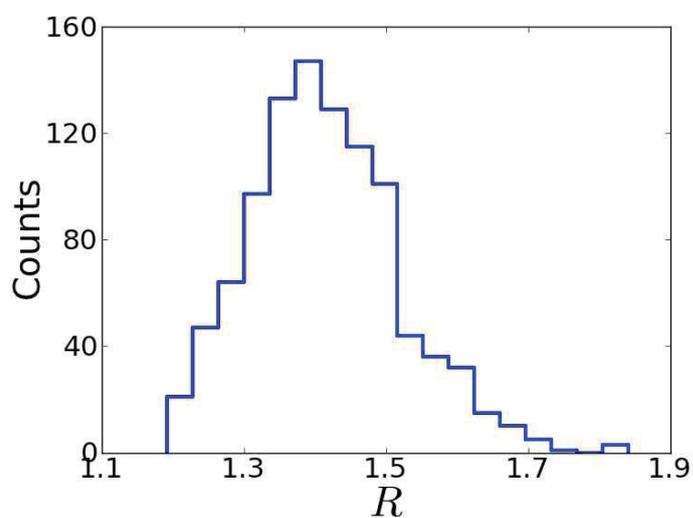


Figure 4: The distribution of R values for a unitary transformation control landscape. 1000 optimizations were carried out, using the target unitary transformation W of Eq. (23) and Hamiltonian H_0 and dipole μ matrices of Eq. (24). Modest values of R , all less than 2.0, characterize this landscape's structure as simple.

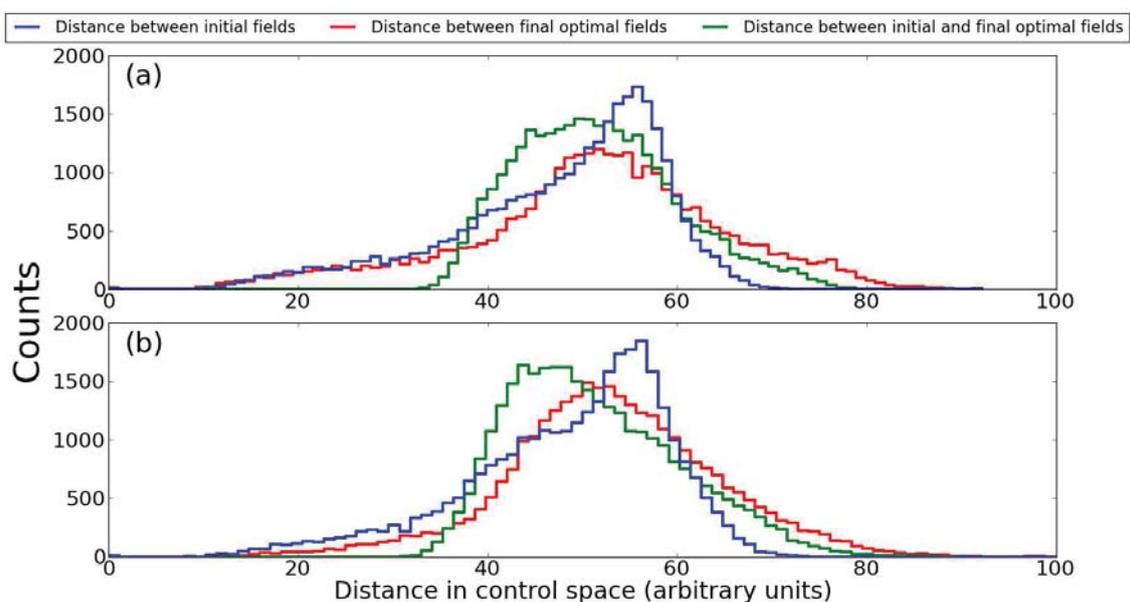


Figure 5: The distributions of pairwise distances between initial fields, final fields, and initial-final fields for the unitary transformation control landscape, calculated using Eq. (22). Plot (a) was produced using the 250 optimizations with the lowest values of R , while plot (b) was produced using the 250 optimizations that had the highest values of R . The similarity between the sets of distributions in (a) and (b) suggests that low R and high R trajectories, as well as all of the associated fields, are distributed in the same way throughout control space.

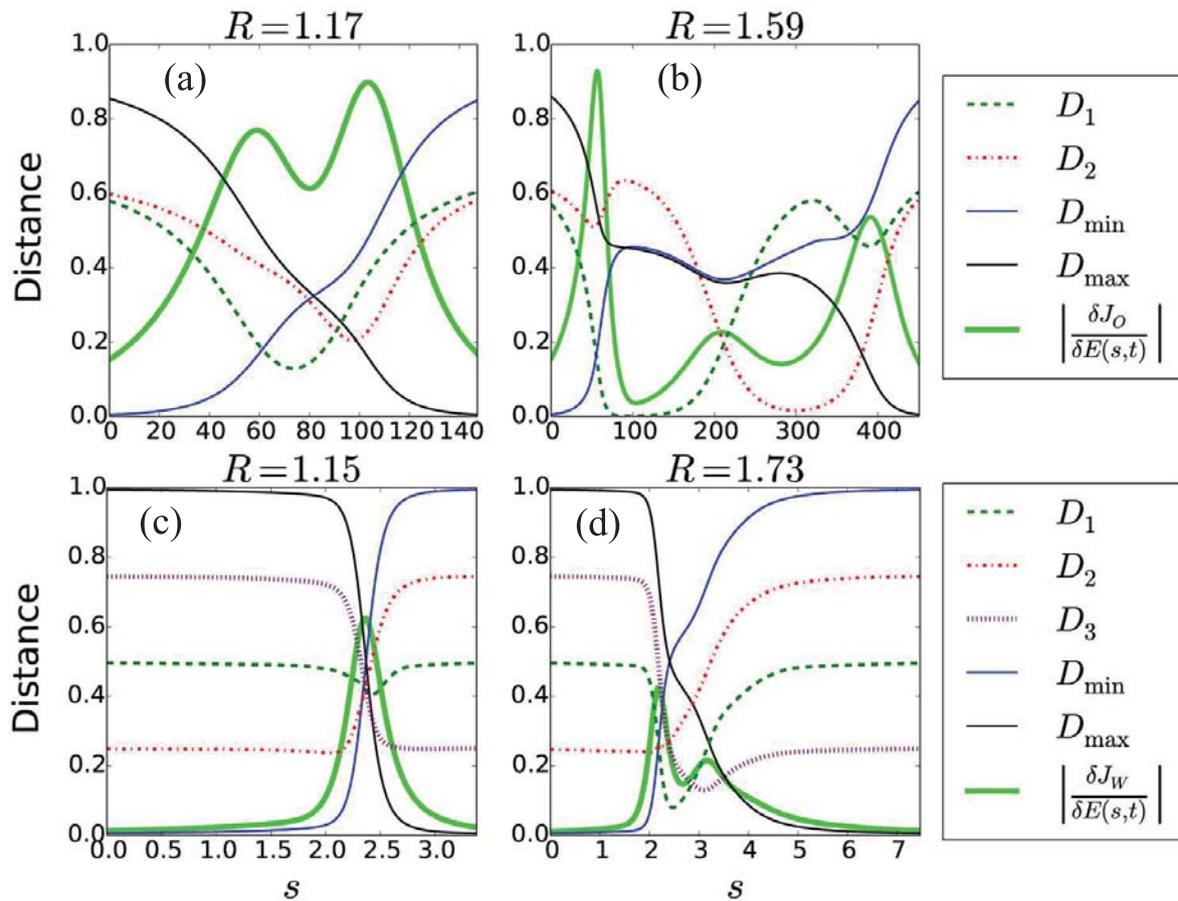


Figure 6: In (a) and (b), the distances D to the critical submanifolds (in arbitrary units) are plotted using Eq. (28) along a low and high R value optimization trajectory, respectively, on a quantum ensemble control landscape. It can be seen in (b) that the distance to saddle submanifolds 1 and 2 become very small along the high R trajectory, indicating that the saddles distorted the trajectory by acting as attractors. The norm of the gradient function along the climb, scaled down by a factor of 60 and shown as the thick green curve, also noticeably dips in (b), further confirming that saddle submanifolds were encountered. For a unitary transformation control landscape in (c) and (d), the distances D to the critical submanifolds, using Eq. (29), are plotted along a low and high R value optimization trajectory. In (c), no saddle submanifolds are encountered, but in (d) the black and green curves (i.e., for saddles 1 and 3, respectively) dip sharply along with the norm of the gradient (scaled down by a factor of 20), signaling attraction towards these saddle submanifolds and resulting in the corresponding high value of R .

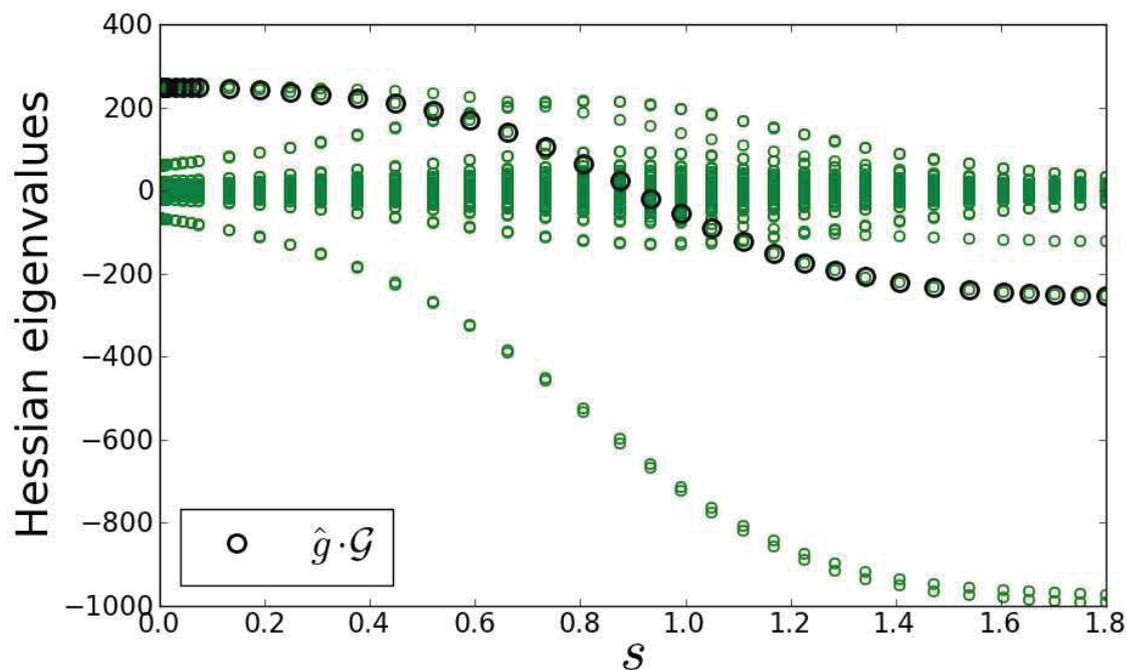


Figure 7: The thick black circles are values of $\hat{g}(s) \cdot \mathcal{G}(s) = \hat{g}(s) \cdot \mathcal{H}(s) \cdot \hat{g}(s)$ along a nearly straight shot to the top of a state-to-state transition probability landscape. Here $\mathcal{H}(s)$ is the Hessian and \hat{g} is the normalized gradient function. The eigenvalues of the Hessian are shown as the smaller green circles, for each point s along the climb of the landscape. The black circles match with a green circle at each point on the climb, supporting the identification of the gradient with an eigenfunction of the Hessian. As the associated eigenvalue of the Hessian changes from positive to negative, the function $\mathcal{G}(s) = \mathcal{H}(s) \cdot \hat{g}$ changes from being parallel to antiparallel to the gradient function.

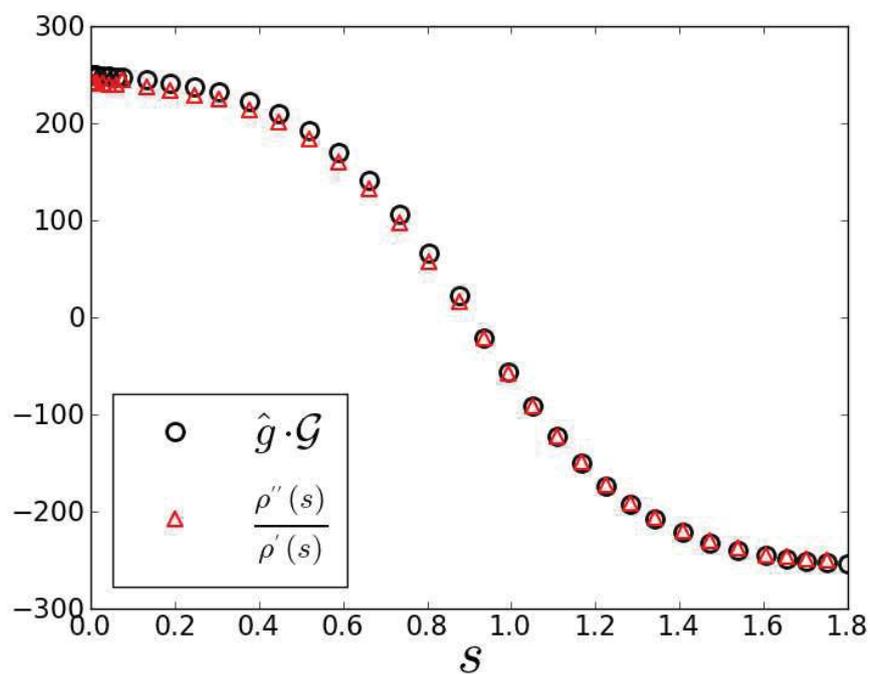


Figure 8: A plot of $\hat{g} \cdot \mathcal{G} = \hat{g} \cdot \mathcal{H} \cdot \hat{g}$, which is the factor by which the Hessian \mathcal{H} scales the normalized gradient \hat{g} . From Eq. (17), along a straight control trajectory we should have $\hat{g} \cdot \mathcal{G} = \frac{\rho''(s)}{\rho'(s)}$. Thus, the plot also shows $\rho''(s)/\rho'(s)$ evaluated along the gradient landscape climb. The good agreement between the two functions signifies the validity of Eq. (17) along a straight control trajectory.