On the complexity of quantum control optimization trajectories

Arun Nanduri[1], Ashley Donovan[1], Tak-San Ho[1], and Herschel Rabitz[1]

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

September 8, 2013

Abstract

Quantum control aims to manipulate the dynamics of physical systems by utilizing a shaped external field. The relation between the control field and the physical objective can be realized as a high-dimensional quantum control landscape. Quantum control landscape *topology*, at the critical points, has been thoroughly explored in previous work, and under suitable conditions, no suboptimal extrema that could trap local search algorithms are present. Here, we address the *structure* of the landscape, away from critical points. It has been found that landscape structure is highly favorable for optimization of state-to-state transition probabilities, in that control trajectories take nearly straight, direct paths to global maxima. To explore this issue further, landscape structure has been codified in the metric $R \geq 1$, defined as the ratio of the length of the control trajectory to the Euclidean distance between the initial and optimal controls. We extend the state-to-state transition probability results to the quantum ensemble and unitary transformation control landscapes. The interplay of optimization trajectories with critical submanifolds, possessing saddle point topology, is relevant here, and it is found that these submanifolds complicate landscape structure. A fundamental mathematical relationship of perfectly straight control trajectories is derived, wherein the gradient on the quantum control landscape must be an eigenfunction of the second-order Hessian. This relation can be explored as an indicator of landscape structure and may provide a step towards identifying when control trajectories can achieve perfect linearity. Together with earlier topological results, these findings on landscape structure may illuminate why optimal quantum controls can be readily identified.

1 Introduction

Quantum optimal control theory (OCT) has vividly illustrated the efficacy of using external control fields to tune quantum systems and the ease with which such fields are identified [1, 2, 3, 4]. Optimal control experiments (OCE) has found success as well, thanks to the development of closed-loop learning algorithms [5] made possible through femtosecond laser pulse shaping technology [6]. The

dynamics of a closed quantum system upon interaction with the applied field can be described by the time-dependent Schrödinger equation

$$i\hbar\frac{\partial\left|\psi(t,0)\right\rangle}{\partial t}=H(t)\ \left|\psi(t,0)\right\rangle, \tag{1}$$

where $H(t)$ is the time-dependent Hamiltonian and $|\psi(t,0)\rangle$ is the state's wave function with $0 \leq t \leq T$. Treating the interaction with the dipole approximation, the time-dependent Hamiltonian is given by

$$H(t)=H_0-\mu E(t), \tag{2}$$

where H_0 is the quantum system's Hamiltonian, μ is the dipole moment operator and $E(t)$ is the time-dependent control field.

Through Eq. (1), every control field can be mapped to the value of a cost function $J[E(t)] = \langle\psi(t)|\,O\,|\psi(t)\rangle$, where O is a physical observable. The functional relationship between the cost function $J[E(t)]$ and the control field $E(t)$ can be cast in terms of a *quantum control landscape*, where each point on the landscape corresponds to a certain control field $E(t)$ and the height of the landscape at that point corresponds to the value the cost function takes for that particular control field, $J[E(t)]$. Finding successive control fields that lead to optimal values of the cost function corresponds to climbing this landscape. Using a local search algorithm to climb, such as the gradient ascent algorithm we restrict ourselves to in this work, results in smooth paths up the landscape that can be parametrized by a continuous variable, which we denote s. Every such path up the landscape can be projected onto a trajectory through *control space*. Thus, the initial control field that evaluates to a poor value of the cost function is $E(s=0,t)$ and the final control field that results in an optimal value of the cost function is $E(s_{max},t)$.

Many prior studies have been devoted to the study of the *topology* of these quantum control landscapes. In particular, it has been found that under the assumptions of (i) controllability, which states that any unitary evolution matrix $U(T,0)$ can be generated from the closed quantum system by specifying some control field $E(t)$ via the time-dependent equation $i\hbar\frac{\partial U(t,0)}{\partial t} = H(t)U(t,0), U(0,0)=1$, (ii) constraint-free control, which states that every possible control field is accessible, and (iii) surjectivity, which states that the set of functions $\frac{\delta U_{ij}(T,0)}{\delta E(t)}$ are linearly independent over $t\in[0,T]$, the quantum control landscape contains only trap-free critical submanifolds [7, 8, 9, 10, 11, 12, 13]. This landscape topology is highly favorable to OCT and OCEs, and partly accounts for the success they have had in locating optimal control fields.

However, the *structure* of the quantum control landscape away from the critical points affects the complexity of the resultant control trajectory and plays an equally important role in these successes. Landscape structure may be examined by quantifying the linearity of their control trajectories [14]. Prior experimental work examined the linearity of control trajectories maximizing

second harmonic generation and found that all sampled control trajectories were strikingly close to being straight [15]. Previous theoretical work has also supported the notion of simple landscape structure by examining the linearity of control trajectories for a state-to-state transition probability landscape [14, 16, 17].

In this work, we extend these results to quantum control landscapes induced by two additional cost functions: (i) $J = \mathrm{Tr}\,(\rho(T)O)$ and (ii) $J = \|W - U(T,0)\|$. The first is the value of a general operator O in the presence of an arbitrary density matrix. This cost function allows us to consider systems that are initially in mixed states, thus allowing for control of a broader class of physical observables. The second cost function corresponds to the distance from a desired unitary transformation W, which allows for implementing the quantum analog of logic gates for quantum information processing [18]. We demonstrate that the structure of these landscapes is also relatively simple through performing simulations that show that the control trajectories are close to linear. We also examine the influence of critical submanifolds on landscape structure. Specifically, we find that the presence of critical submanifolds at intermediate heights on the landscape drives control trajectories away from straight paths. This phenomenon is illustrated in Figure 1. The trajectory on the right takes a straightforward path up the landscape, while the trajectory on the left is attracted to the saddle point, thus losing its direct nature. We have also identified a mathematical criterion that the Hessian and the gradient of the cost function must satisfy in order to achieve a perfectly straight control trajectory on any quantum control landscape. In particular, the gradient vector must be an eigenfunction of the Hessian. Furthermore, we show that if 'straight shot' trajectories exist, they would lead to the factorization of the kernels of all higher order derivatives of the gradient function. This leads to a hierarchy of first-order integral equations for straight shot control trajectories.

The rest of this paper is organized as follows. The linearity measure R of control trajectories 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. The existence of saddle points is examined and a metric that gives the distance from a saddle point is also outlined. In section 4, a detailed description and derivation of the relationship between the Hessian and the gradient function for straight control trajectories is presented. Section 5 provides numerical results on the statistical behavior of R and on the effect of saddle points on R for both types of quantum control landscapes. We then illustrate the Hessian-gradient eigenrelation for nearly straight control trajectories. Conclusions are furnished in section 6.

2 The linearity measure R

The goal of the present work is to explore quantum control landscape structure through examining the control trajectories followed during optimizations. To this end, we construct a ratio R measuring 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

$$d_{PL} = \int_0^{s_{max}} \left[\frac{1}{T} \int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt \right]^{\frac{1}{2}} ds, \tag{3}$$

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

$$d_{EL} = \left[\frac{1}{T} \int_0^T \left[E(s_{max},t) - E(0,t) \right]^2 dt \right]^{\frac{1}{2}}, \tag{4}$$

where s parametrizes the control trajectory and T is the target time. R is then given by

$$R = \frac{d_{PL}}{d_{EL}} = \frac{\int_0^{s_{max}} \left[\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt \right]^{\frac{1}{2}} ds}{\left[\int_0^T \left[E(s_{max},t) - E(0,t) \right]^2 dt \right]^{\frac{1}{2}}}. \tag{5}$$

The lower bound for this quantity is $R = 1$, which holds for trajectories that are straight lines between $E(0,t)$ and $E(s_{max},t)$. These correspond to the simplest possible paths an optimization can take to climb the quantum control landscape $J[E(t)]$, and indicate a lack of complicated structural features that could interfere with optimization paths. Higher values of R mean that a control trajectory meanders through control space along a path whose total length is much greater than the distance between its two endpoints, signifying that the optimization path encountered significantly rugged structural components of the landscape. To identify an upper bound on R, d_{PL} must be maximized and d_{EL} must be minimized. An upper bound for the former is given by $\frac{2}{\hbar} s_{max} \|\mu\|$ [14], but there is currently no analytical form for a lower bound on the latter quantity.

In a previous work, we found values of $R-1$ as low as 10^{-4} on a five level state-to-state transition probability landscape [14]. By examining the distances between the endpoints of control trajectories, we also observed that trajectories with low and high R values are distributed identically throughout control space. Furthermore, we found that it was possible to characterize the complexity of low R trajectories through the 'straight shot' climbing procedure. This algorithm calculates the gradient of the control landscape at an initial control field, then proceeds in that direction until a local maximum is encountered. For suitably low R trajectories, substantially optimal control fields may be found in this way. We determined that the threshold for this procedure to yield a final control field with a corresponding landscape height close to the global maximum is $R-1 \leq 10^{-2}$.

3 Negotiating the quantum control landscape

We describe a general quantum system of interest through a field-free Hamiltonian H_0, represented by an $N \times N$ diagonal matrix. Likewise, μ, describing the quantum system's dipole moment, is an $N \times N$ symmetric matrix. Interpreting the value of $J[E(t)]$ as a height on a landscape over the space of possible controls, we see that optimizing an external field to achieve the desired value of J corresponds to climbing the quantum control landscape. Utilizing a gradient-based method to climb corresponds to following the path of steepest ascent up the landscape, and this makes D-MORPH a natural choice for the rule that is used to traverse the landscape and consequently to navigate the underlying control space.

This process can be confounded by the presence of submanifolds of critical points that are located at intermediate heights on the landscape. These points have the potential to halt any local search algorithm from reaching a global maximum, but theoretical analyses have revealed that such 'traps' do not exist, and that these critical submanifolds posses saddle point topology [8, 9, 12, 13]. However, as we shall illustrate, saddle points can still have a deleterious effect on optimizations by distorting the linear nature of climbs to the top of the landscape.

We parameterize the trajectory taken by the controls by a continuous variable $s \geq 0$ so that the external field becomes a function of two variables $E(s,t)$. As s is increased, the change in the cost function can be written using the chain rule as

$$\frac{dJ}{ds} = \int_0^T \frac{\delta J}{\delta E(s,t)} \frac{\partial E(s,t)}{\partial s} \, \mathrm{d}t. \tag{6}$$

If we want to maximize J, then we stipulate that $\frac{dJ}{ds} \geq 0$, so we will set

$$\frac{\partial E(s,t)}{\partial s} = \frac{\delta J}{\delta E(s,t)}. \tag{7}$$

Likewise, when minimizing J, $\frac{dJ}{ds}$ must be less than or equal to zero, so we set $\frac{\partial E(s,t)}{\partial s} = -\frac{\delta J}{\delta E(s,t)}$. This optimization procedure is known as the D-MORPH algorithm, and is the algorithm we make use of in this paper to climb the landscape [19, 20]. The control variables that are adjusted with increasing s are the amplitudes of the control field $E(s,t_i)$ at a set of discretized time points $\{t_i\}$, $i = 1, 2, \ldots$. In order to ensure that every simulation in this work travels the same distance up (or down) the quantum control landscape, all the trajectories for which an R value is calculated begin at the same initial height on the landscape J^I and end at the same final height on the landscape J^F. We accomplish this by using D-MORPH to 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 final height on the landscape J^F is then achieved by $E(s = s_{max}, t)$.

Calculation of the $\frac{\delta J}{\delta E(s,t)}$ requires knowledge of the cost function that generates the landscape. We will now outline the calculation of the gradient function from the cost functions that result in

optimization of arbitrary observables and the achievement of unitary transformations. We will also characterize the critical submanifolds of these landscapes, and give the definition of a metric that yields the distance to any critical submanifold.

3.1 Quantum ensemble control landscape

3.1.1 Gradient function

The expectation value of an arbitrary observable $\mathbf{O}$ at a target time T for a system that is in a state that can be described by the time-dependent density matrix $\rho(t)$ is given by

$$\langle \mathbf{O} \rangle = \mathrm{Tr}\ (\rho(T)\mathbf{O}) = \mathrm{Tr}\ \Big(U(T,0)\rho(0)U^\dagger(T,0)\mathbf{O}\Big). \tag{8}$$

Our cost function, to be maximized, is $J = \langle \mathbf{O} \rangle$. By manipulating this expression, it can be shown that $\frac{\delta \langle \mathbf{O} \rangle}{\delta E(s,t)}$ is given by [11]

$$\frac{\delta \langle \mathbf{O} \rangle}{\delta E(s,t)} = \frac{2}{\hbar}\, \Im\Big\{ \mathrm{Tr}\, U(T,0)\rho\mu(t)U^\dagger(T,0)\mathbf{O} \Big\}, \tag{9}$$

where $\mu(t) = U(t,0)\mu U^\dagger(t,0)$ is the time-dependent dipole matrix. Having set out how to climb the $\langle \mathbf{O} \rangle$ landscape, we will now give an account of the critical submanifolds of this landscape.

3.1.2 Enumeration of critical submanifolds

Let us first note that we can consider ρ and $\mathbf{O}$ to be in diagonal form without loss of generality [11, 12]. In characterizing the topology of the quantum ensemble control landscape, we will also consider $\langle \mathbf{O} \rangle$ to be a function of the propagator $U(T,0)$ and ignore the propagator's dependence on the time-dependent control field $E(t)$ for the moment. Utilizing this viewpoint is advantageous as it is simpler to consider $\langle \mathbf{O} \rangle$ as a function over the finite dimensional Lie group $U(N)$ than as a functional over the infinite dimensional set of smooth functions on $[0,T]$.

At a critical point $U_0 \in U(N)$ of $\langle \mathbf{O} \rangle$, first order variations in U_0 produce no change in $\langle \mathbf{O} \rangle$. We can expand around U_0 in the variable s by writing $U = U_0 e^{iHs}$, where $H \in \mathfrak{u}(N)$, where $\mathfrak{u}(N)$ is the Lie algebra of $U(N)$ which consists of Hermitian matrices. The condition that U_0 is a critical point is then

$$\frac{d}{ds}\mathrm{Tr}\ \Big(U_0\rho U_0^\dagger \mathbf{O}\Big)\Big|_{s=0} = \mathrm{Tr}\ \Big(iH[O, U_0\rho U_0^\dagger]\Big) = 0. \tag{10}$$

Thus we must have $[O, \rho(T)] = 0$ at a critical point. From this relation, it is possible obtain the form of the critical U matrices. Carrying out this analysis reveals that the number of critical submanifolds is in one-to-one correspondence with the number of ways the eigenvalues of ρ can be permuted such that their overlaps with the eigenvalues of O are different [12]. These permutations,

induced by the critical propagators $U(T)$ which correspond to the distinct critical submanifolds, can be enumerated using contingency tables [12]. These tables possess n rows and m columns and can be constructed as follows. First, write ρ and $\mathbf{O}$ in diagonal form with their eigenvalues in descending order. Then, apply some permutation to the eigenvalues of ρ. The value of the (i,j) entry in the table will be the number of times the eigenvalue λ_i of ρ overlaps with the eigenvalue ϵ_j of $\mathbf{O}$. Furthermore, the rows of the contingency table must add up to the degeneracies of the eigenvalues of $\mathbf{O}$, and the columns of the table must add up to the degeneracies of the eigenvalues of ρ. This fact can be used to easily construct all possible contingency tables. Further study making use of the Hessian shows that on the $\langle \mathbf{O} \rangle$ landscape, all critical submanifolds consist of saddle points, aside from the two submanifolds corresponding to the top and bottom of the landscape [13].

3.1.3 Distance metric for critical submanifolds

The contingency tables can also be used to calculate a distance from a particular critical submanifold. A short explanation of this metric will be given; for a full discussion, consult Sun *et al.* [21].

Consider a point on the landscape $U \in U(N)$ that is also on some critical submanifold. Divide U into blocks of sizes $o_i \times p_j$. Since U can be written as an element of a double coset (c.f. Eq. (**??**)), we have a similar expression for each block

$$U_{ij} = P_i \pi_{ij} Q_j^\dagger \tag{11}$$

where $P_i \in U(o_i)$, $Q_j \in U(p_j)$, and π_{ij} is the block of size $o_i \times p_j$ of the (non-unique) permutation matrix of the critical submanifold in question. This equation states that U_{ij} and π_{ij} have identical singular values. It can be shown that the blocks of permutation matrices π_{ij} have only 0s and 1s as their singular values, and that the number of 1s is equal to c_{ij}, the entry in the ith row and jth column in the contingency table. The same must be true of the block U_{ij} on a critical submanifold. This fact will be used to create a notion of distance to each critical submanifold.

Define a diagonal matrix S_{ij}, which has c_{ij} 1s in the diagonal:

$$S_{ij} = \begin{pmatrix} \mathbb{I}_{c_{ij}} & \\ & 0 \end{pmatrix} \tag{12}$$

To find the distance of an arbitrary point U_0 is from the critical submanifold, write each block U_{ij} in its singular form Σ_{ij}, with the singular values on the diagonal sorted in descending order. Define $V_{ij} = S_{ij} - \Sigma_{ij}$. The distance is then

$$d = \sum_{i=1}^{p} \sum_{j=1}^{o} \mathrm{Tr} \left(V_{ij}^\dagger V_{ij} \right). \tag{13}$$

There is a normalization associated with this metric [22]. Any point U on the landscape can be assigned a vector of the singular values of its blocks U_{jk}. 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 matrices of the critical submanifolds. This 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. We make implicit use of this normalization later in this work.

3.2 Unitary transformation control landscape

3.2.1 Gradient function

The cost function for the unitary transformation landscape is defined as the distance of the propagator $U(T;0)$ from a target unitary transformation W

$$J = \|W - U(T;0)\|^2 = 2N - 2\Re\Big\{ \mathrm{Tr}\,\Big(W^\dagger U(T;0)\Big)\Big\} \tag{14}$$

where $\|\cdot\|$ denotes the Frobenius norm. J has a minimum 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. Note that we would now like to minimize J.

The gradient function for this landscape is

$$\frac{\delta J}{\delta E(t)} = \frac{2}{\hbar}\,\mathrm{Tr}\,\Im\Big\{W^\dagger U \mu(t)\Big\}. \tag{15}$$

Where $\mu(t) = U(t;0)\mu U^\dagger(t;0)$ is the time-dependent dipole matrix. We now proceed to an account of the topology of the unitary transformation control landscape.

3.2.2 Enumeration of critical submanifolds

We mention that we will again consider the controls from the kinematic point of view, for the same reasons as before. To characterize the critical points, expand around U_0 in the variable s as before by writing $U = U_0 e^{iHs}$, where $H \in \mathfrak{u}(N)$, the set of Hermitian $N \times N$ matrices. The condition that U_0 is a critical point is then

$$\frac{d}{ds}\|W - U = U_0 e^{iHs}\|^2\Big|_{s=0} = i\mathrm{Tr}\,\Big(H\,\Big(W^\dagger U - U^\dagger W\Big)\Big) = 0. \tag{16}$$

This is equivalent to requiring $W^\dagger U = U^\dagger W$, or that $W^\dagger U$ be Hermitian. $W^\dagger U$ is also a unitary matrix, so at a critical point, the eigenvalues of $W^\dagger U$ must be either 1 or -1. As a result, the

values J can take on at a critical point are $J_c = 0, 4, \ldots, 4N$, implying that there exist N+1 critical submanifolds, each of which correspond to a certain number of 1s and -1s as the eigenvalues of $W^\dagger U$. A detailed analysis has shown that, as in the case of the quantum ensemble landscape, all critical submanifolds except for the ones corresponding to the top and the bottom of the landscape possess saddle point topology [8, 9].

3.2.3 Distance metric for critical submanifolds

To define a distance to a critical submanifold from an arbitrary point on the kinematic landscape U_0, note that in general the eigenvalues of $W^\dagger U$ take values on the unit circle in the complex plane. Let λ_i stand for the ith eigenvalue of $W^\dagger U_0$, when the eigenvalues are sorted in ascending order by their real parts. To find the distance to the critical submanifold which corresponds to α -1s and $N - \alpha$ 1s in the eigenvalues of $W^\dagger U$, we can 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 [21]. Quantitatively,

$$d = \sum_{i=0}^{\alpha} \left(1 + \Re\{\lambda_i\}\right) + \sum_{\alpha}^{N-1} \left(1 - \Re\{\lambda_i\}\right). \tag{17}$$

The maximum distance to a critical submanifold in this case is $2N$, so we will divide by this quantity and normalize all distances to a maximum of 1.

We have shown how to climb the quantum ensemble and unitary transformation control landscapes using a gradient algorithm, and have examined the topology of the critical points of these landscapes. In the next section, we present a relation that holds for *all* quantum control landscapes which must be satisfied for straight control trajectories which make use of a gradient algorithm.

4 'Straight shot' eigenrelation

One of the goals of this work is to examine the conditions under which control trajectories with $R = 1$ can be achieved. In this regard, we formulate a mathematical criterion that must be satisfied in order for 'straight shot' trajectories to exist in conjunction with the D-MORPH gradient algorithm.

Consider a control trajectory that is perfectly straight, i.e., $R = 1$. This trajectory can be written [14]

$$E(s,t) = E_0(t) + \rho(s)\Delta E(t), \tag{18}$$

where $\rho(s)$ is a monotonically increasing function, as a result of the gradient algorithm, and $\Delta E(t)$ is some function of time. The gradient function along this trajectory is

$$\frac{\delta J}{\delta E(t)}[E(s,\tau)] = \alpha(s)\Delta E(t), \tag{19}$$

where $\alpha(s) = \frac{d\rho(s)}{ds}$, J is the cost function of the quantum control landscape, which, without loss of generality, we will maximize; τ is an auxiliary time variable, and the square brackets denote the argument of the gradient function.

According to Eq. (7), the instantaneous tangent of a control trajectory "points" in the same direction in control space as the gradient of the quantum control landscape at that point. Along a straight control trajectory, then, the gradient vector must point in the same direction at every point on the trajectory; in other words, the gradient functions at two different points on the control trajectory are related by a scale factor. Additionally, for straight trajectories, the difference between any two points on the trajectory is always proportional to the same vector. When using a gradient-based algorithm such as D-MORPH, this vector is the gradient itself. Therefore, in evaluating the gradient function along a straight trajectory, the *argument* of the gradient function changes by an amount proportional to the gradient itself. These two facts form the foundation for the analysis of this section.

To derive the implications, note that we can use Eq. (19) to relate the gradient function at two different points along the straight shot

$$\frac{\delta J}{\delta E(t)}[E(s',\tau)] = \frac{\alpha(s')}{\alpha(s)}\alpha(s)\Delta E(t) = \frac{\alpha(s')}{\alpha(s)}\frac{\delta J}{\delta E(t)}[E(s,\tau)]. \tag{20}$$

Next, let us rewrite equation Eq. (18) to relate the control field at two different points along the trajectory

$$\begin{aligned} E(s,t) &= E_0(t) + \rho(s)\Delta E(t) = E_0(t) + \rho(s')\Delta E(t) + \big(\rho(s) - \rho(s')\big)\,\Delta E(t) \\ &= E(s',t) + \big(\rho(s) - \rho(s')\big)\,\Delta E(t). \end{aligned} \tag{21}$$

If we take the two points to be infinitesimally close, we can write $\rho(s) - \rho(s') = \alpha(s')(s - s')$. Then, combining the two relations above, we have

$$\begin{aligned} \frac{\delta J}{\delta E(t)}[E(s',\tau)] &= \frac{\alpha(s')}{\alpha(s)}\,\frac{\delta J}{\delta E(t)}[E(s,\tau)] \\ &= \frac{\alpha(s')}{\alpha(s)}\,\frac{\delta J}{\delta E(t)}[E(s',\tau) + (s-s')\Delta E(\tau)\alpha(s')] \\ &= \frac{\alpha(s')}{\alpha(s)}\,\frac{\delta J}{\delta E(t)}\left[E(s',\tau) + (s-s')\frac{\delta J}{\delta E(\tau)}[E(s',\tau')]\right]. \end{aligned} \tag{22}$$

We can further expand the argument of the gradient function on the right hand side to lowest order in $s - s'$ obtain

$$\frac{\delta J}{\delta E(t)}[E(s',\tau)] = \frac{\alpha(s')}{\alpha(s)}\left(\frac{\delta J}{\delta E(t)}[E(s',\tau)] + (s-s')\int_0^T \frac{\delta^2 J}{\delta E(t')\delta E(t)}[E(s',\tau)]\,\frac{\delta J}{\delta E(t')}[E(s',\tau)]\,\mathrm{d}t'\right) \tag{23}$$

where $\frac{\delta^2 J}{\delta E(t')\delta E(t)}[E(s',\tau)]$ is the Hessian. Rearranging, we have

$$\int_0^T \frac{\delta^2 J}{\delta E(s',t')\delta E(s',t)}\frac{\delta J}{\delta E(s,' t')}\,\mathrm{d}t' = \frac{\alpha'(s')}{\alpha(s')}\frac{\delta J}{\delta E(s',t)} \tag{24}$$

where we have now dropped the argument of the gradient and the Hessian in brackets and shown the s dependence in the denominator. This equation states that the gradient is an eigenvector of the Hessian along a straight control trajectory.

We wish to stress that the relations examined in this section were derived without making any use of quantum mechanics. The functional relationship between the controls and the height on the landscape has its origin in the quantum control equations set out in Section 3, but the rest of the analysis only makes use of the gradient climbing rule, Eq. (7), and the straight shot condition Eq. (19). Therefore, our analysis applies to all landscape gradient climbs which result in straight trajectories in the underlying space of controls.

5 Results and Illustrations

This section provides numerical illustrations of the details of landscape structure outlined in Sections 2 through 4. Section 5.1 shows that the distribution of R values of control trajectories across the entire landscape takes on very modest values, rarely ever venturing above 2.0. It also contains plots that aim to demonstrate that control trajectories with very low and high values of R are spread evenly throughout control space. The next subsection addresses the effects saddle points have on optimization paths. We find that encountering additional critical submanifolds between the top and bottom of the landscape on an optimization path tends to increase R values. The final subsection provides numerical evidence of the Hessian-gradient eigenrelation discussed in Section 4.1. Dimensionless units are used throughout this section.

5.1 Random sampling of R values

In order to assess the complexity of the quantum ensemble and unitary transformation control landscapes, random control trajectories were constructed and their R values tabulated. Although there is no *a priori* reason to expect values of R at low orders of magnitude since the control space is typically of very high dimensionality, we find that R values are almost always less than 2. We also examine the distribution of distances between control fields that result in trajectories with high and low values of R, and find that they have similar behavior.

5.1.1 Quantum ensemble control landscape

Five different quantum ensemble control landscapes were sampled, and 1000 optimizations of $\langle \mathbf{O} \rangle$ were performed using each landscape. Two different initial density matrices were used in generating

these landscapes:

$$\rho_1 = \begin{smallmatrix} 1&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0 \end{smallmatrix}, \qquad \rho_2 = \begin{smallmatrix} \frac{1}{4}&0&0&0&0&0&0&0\\ 0&\frac{1}{4}&0&0&0&0&0&0\\ 0&0&\frac{1}{4}&0&0&0&0&0\\ 0&0&0&\frac{1}{4}&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0 \end{smallmatrix}. \tag{25}$$

ρ_1 is an almost completely degenerate density matrix. On the other hand, ρ_2 contains two eigenvalues that both exhibit mild degeneracies. The observables used in conjunction with these operators are

$$O_1 = \begin{smallmatrix} 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&\frac{4}{9}&0\\ 0&0&0&0&0&0&0&0&\frac{5}{9} \end{smallmatrix}, \qquad O_2 = \begin{smallmatrix} 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0\\ 0&0&0&0&\frac{4}{17}&0&0&0\\ 0&0&0&0&0&\frac{4}{17}&0&0\\ 0&0&0&0&0&0&\frac{4}{17}&0\\ 0&0&0&0&0&0&0&\frac{5}{17} \end{smallmatrix}. \tag{26}$$

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 different landscapes. Lastly, we also made use of

$$\rho_3 = \begin{smallmatrix} \frac{7}{28}&0&0&0&0&0&0&0\\ 0&\frac{6}{28}&0&0&0&0&0&0\\ 0&0&\frac{5}{28}&0&0&0&0&0\\ 0&0&0&\frac{4}{28}&0&0&0&0\\ 0&0&0&0&\frac{3}{28}&0&0&0\\ 0&0&0&0&0&\frac{2}{28}&0&0\\ 0&0&0&0&0&0&\frac{1}{28}&0\\ 0&0&0&0&0&0&0&0 \end{smallmatrix} \qquad O_3 = \begin{smallmatrix} 0&0&0&0&0&0&0&0\\ 0&\frac{1}{28}&0&0&0&0&0&0\\ 0&0&\frac{2}{28}&0&0&0&0&0\\ 0&0&0&\frac{3}{28}&0&0&0&0\\ 0&0&0&0&\frac{4}{28}&0&0&0\\ 0&0&0&0&0&\frac{5}{28}&0&0\\ 0&0&0&0&0&0&\frac{6}{28}&0\\ 0&0&0&0&0&0&0&\frac{7}{28} \end{smallmatrix} \tag{27}$$

to simulate a completely nondegenerate quantum ensemble control landscape. The optimizations used the Hamiltonian and dipole operators

$$H_0 = \begin{smallmatrix} -10&0&0&0&0&0&0&0\\ 0&-8&0&0&0&0&0&0\\ 0&0&-4&0&0&0&0&0\\ 0&0&0&2&0&0&0&0\\ 0&0&0&0&10&0&0&0\\ 0&0&0&0&0&20&0&0\\ 0&0&0&0&0&0&32&0\\ 0&0&0&0&0&0&0&46 \end{smallmatrix} \qquad \mu = \begin{smallmatrix} 0&\pm1&\pm0.5&\pm0.5^2&\pm0.5^3&\pm0.5^4&\pm0.5^5&\pm0.5^6\\ \pm1&0&\pm1&\pm0.5&\pm0.5^2&\pm0.5^3&\pm0.5^4&\pm0.5^5\\ \pm0.5&\pm1&0&\pm1&\pm0.5&\pm0.5^2&\pm0.5^3&\pm0.5^4\\ \pm0.5^2&\pm0.5&\pm1&0&\pm1&\pm0.5&\pm0.5^2&\pm0.5^3\\ \pm0.5^3&\pm0.5^2&\pm0.5&\pm1&0&\pm1&\pm0.5&\pm0.5^2\\ \pm0.5^4&\pm0.5^3&\pm0.5^2&\pm0.5&\pm1&0&\pm1&\pm0.5\\ \pm0.5^5&\pm0.5^4&\pm0.5^3&\pm0.5^2&\pm0.5&\pm1&0&\pm1\\ \pm0.5^6&\pm0.5^5&\pm0.5^4&\pm0.5^3&\pm0.5^2&\pm0.5&\pm1&0 \end{smallmatrix} \tag{28}$$

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 initial fields of these optimizations were parametrized according to the equation

$$E(t) = \frac{1}{F}\exp[-0.3(t-\frac{T}{2})^2]\sum_{n=1}^{M} a_n \sin(\omega_n t + \phi_\omega), \tag{29}$$

with the target time T being 10. When the eight level landscapes were being sampled, M was set at 60, to ensure that the initial control field had potential resonances with every possible state to state transition. For the four level landscape, $M = 20$. F is a normalization factor, which was picked so that the initial fluence of each field is set to 1. This is to prevent the onset of the strong field regime, where small alterations in a control field can cause massive changes in the dynamics

of the quantum system. Across all the optimizations, this control field was discretized into 1001 time points, and the field amplitude at each of these time points served as the control variables. We define the limits of our optimizations as follows. Let $\langle \mathbf{O} \rangle_{max}$ denote the maximum value that $J = \mathrm{Tr}\,(\rho(t)O)$ can assume for a particular pair of ρ and O, and let $\langle \mathbf{O} \rangle_{min}$ denote the minimum value. When a random field is selected by Eq. (29), it is used to calculate $J = \mathrm{Tr}\,(\rho(t)O)$. If $J < J^I = 0.01\,\langle \mathbf{O} \rangle_{max} + 0.99\,\langle \mathbf{O} \rangle_{min}$, then D-MORPH is used to climb until $J = J^I$; likewise, if $J > J^I$, D-MORPH is used to descend the landscape until $J = J^I$. Then, the resultant control field is optimized using D-MORPH until $J = J^F = 0.99\,\langle \mathbf{O} \rangle_{max} + 0.01\,\langle \mathbf{O} \rangle_{min}$, and it is this latter trajectory, which takes the field from a height of J^I to a height of J^F on the landscape, that is used to calculate R. This is done so that each of the optimizations proceeds over a wide swathe of the landscape and encounters as many structural features as possible. In this regard, we have numerically shown that for a five level state-to-state transition probability landscape, demanding increasing fidelity of J at the beginning and end of the optimizations has a diminutive effect on R [14].

Histograms of the resulting R values from 1000 optimizations over five different landscapes are presented in Figure 2. Starting at the top-left, the histograms show the results of optimizations over the landscape generated by (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 right, indicating that the R values tend to accumulate near 1.0. The smallest values out of the eight level systems are found in the histogram of subfigure (a) which has a mean of $R = 1.2167$. This is followed by the histograms of Figures 2(b) and 2(c), which have means of $R = 1.3053$ and $R = 1.3295$ respectively. The landscapes sampled in these two figures correspond to ρ and O matrices with lower numbers of degeneracies. The highest values of R are found in the histograms of Figures 2(d) and 2(e), which have mean values of $R = 1.6516$ and $R = 1.6455$, respectively. In Figure 2(d), ρ_2 and O_2, the combination with the least amount of degeneracies, was used to generate the landscape sampled, and in Figure 2(e) a fully nondegenerate landscape was sampled. These statistics illustrate that the lack of degeneracies lead to more structurally complex landscapes, although evidently only up to a point, evidenced by the similar resulting histograms in Figures 2(d) and 2(e) between a landscape with four degenerate values in both ρ and O and a landscape with fully nondegenerate ρ and O. Also pictured in Figure 3 is a histogram of the R values of 1000 optimizations over a unitary transformation landscape generated by the transformation of Eq. (31). This plot is skewed right and centered near 1.0, with a mean of $R = 1.4142$ indicating that the unitary transformation control landscape also contains minimally complex structural features.

A natural question to consider is whether or not control trajectories that possess low or high R values are concentrated in some region(s) of control space. To address this issue, pairwise Euclidean

distances were calculated between (1) all pairs of initial fields, (2) all pairs of final fields, and (3) all initial and final fields of each control trajectory for the landscape generated by ρ_1 and O_1. These distances were calculated using

$$d_{EL}^{ij} = \left[\frac{1}{T}\int_0^T [E_i(t) - E_j(t)]^2 dt\right]^{\frac{1}{2}} \tag{30}$$

for control fields E_i and E_j, and the results are displayed in Figure 4. Importantly, in computing these distances, we used only the 250 trajectories with the lowest values of R and, separately, the 250 trajectories with the highest values of R out of the 1000 trajectories that were sampled. This means that $R < 1.1725$ for the control trajectories used to generate the histograms in Figure 4(a), and $R > 1.2558$ for the histograms in 4(b). The distribution of distances between initial fields, in blue, is centered to the left of and is narrower than the distribution corresponding to the distances between all optimal fields, in red. This implies that the optimal fields are scattered more widely throughout control space than the randomly picked initial fields. The curve in green is a histogram of the distribution of distances between all initial-final field pairs, and is centered in between the other two curves. The most important feature to discern from this figure, however, is that the distributions generated from the trajectories with low values of R and the distributions generated from the trajectories with high values of R appear to be very similar. This suggests that trajectories with low and high values of R are distributed in the same way across control space, with no deviating behavior becoming apparent as low or high values of R are approached.

5.1.2 Unitary transformation control landscape

1000 optimizations were also randomly carried out over a unitary transformation control landscape. The target unitary transformation was

$$W = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{31}$$

and the Hamiltonian and dipole operators utilized were

$$H_0 = \begin{pmatrix} -10 & 0 & 0 & 0 \\ 0 & -7 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 8 \end{pmatrix} \qquad \mu = \begin{pmatrix} 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 \\ \pm 1 & 0 & \pm 1 & \pm 0.5 \\ \pm 0.5 & \pm 1 & 0 & \pm 1 \\ \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 \end{pmatrix} \tag{32}$$

where the signs on the dipole matrix were again chosen so that it remained symmetric. The initial fields were chosen according to Eq. (29), with $T = 10$ and $M = 20$. The initial height on this

landscape was normalized to $J^I = 0.99 \times 16 = 15.84$, and the final height was normalized to $J^F = 0.01 \times 16 = 0.16$. Figure 2 contains a histogram of the 1000 R values collected in this way, with a mean value of $R = 1.4142$. Although this is smaller than the values contained in Figures 2(a)-(e), the R values for unitary transformation landscapes increases linearly with the size of the quantum system [16]. The quantum ensemble landscapes considered earlier were for 8 level systems, while this unitary transformation landscape is only a four level system, suggesting that the unitary transformation landscape is structurally more complex than any quantum ensemble landscape for a system of the same size, even one with a minimal number of degeneracies. This is likely because achieving a unitary transformation entails specifying every matrix element of the propagator, while optimizing the expectation value of arbitrary observables only places restrictions on the absolute value of some elements of the propagator.

Figure 5 displays the histograms resulting from calculating the pairwise distances using the same methodology of Section 5.1.1. Again, the pairwise distances between (1) all initial fields, (2) all final fields, and (3) all initial-final field pairs were computed using Eq. (30). The two sets of control trajectories used to produce the histograms in Figure 5(a) and 5(b) satisfy $R < 1.3423$ and $R > 1.7496$, respectively. Significantly, the histograms computed using the 250 control trajectories with the lowest values of R and the 250 trajectories with the highest values of R are almost identical, indicating that trajectories with varying values of R are distributed uniformly across the unitary transformation control landscape as well. This result was also found for a five level state-to-state transition probability landscape [14], suggesting that the even distribution of high and low R trajectories is a general property of quantum control landscapes.

5.2 Interactions with saddle points

Although the saddle points of the quantum ensemble and unitary transformation landscapes are not able to prevent optimizations from reaching a global optimum, they may still significantly influence the trajectory of the optimizations. In order to examine the extent to which saddle points affect the climb up the landscape, the metrics set out in Section 3.1.3 and 3.2.3 were used to calculate the distance to all critical submanifolds on the landscape during optimizations on the unitary transformation control landscape generated by the transformation of Eq. (31), and also on the quantum ensemble control landscape generated by the initial density matrix ρ_2 and observable operator O_1 from Eqs (25) and (26), respectively. For both of the systems outlined above, 1000 optimizations were performed. The initial fields were parametrized using Eq. (29) with $M = 20$. We will first discuss the quantum ensemble control landscape case.

5.2.1 Quantum ensemble control landscape

The contingency tables of the critical submanifolds of the system are

$$C_{min} = \begin{bmatrix} 2 & 4 \\ 1 & 0 \\ 1 & 0 \end{bmatrix}, \qquad C_1 = \begin{bmatrix} 3 & 3 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad C_2 = \begin{bmatrix} 3 & 3 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad C_{max} = \begin{bmatrix} 4 & 2 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}, \tag{33}$$

with corresponding heights on the quantum ensemble control landscape of $J_{min} = 0$, $J_1 = \frac{1}{9}$, $J_2 = \frac{5}{36}$, and $J_{max} = \frac{1}{4}$. This information was used to plot the normalized distances to these four critical submanifolds, using Eq. (13), during the optimizations with the lowest R value and highest R values out of the 1000 optimizations performed, which were 1.1724 and 1.5902, respectively. The magnitude of the gradient vector during the run is also plotted to illustrate the speed of the climb, as optimizations typically slow down near a critical submanifold.

The results are shown in Figure 6. In panel (a), the run with a small R value appears not to closely approach any saddle points. This is confirmed by examining the magnitude of the gradient, which falls slightly only slightly in the middle of the climb and also near the beginning and end of the climb, where the trajectory approaches the critical submanifolds corresponding to the bottom and top of the landscape. In contrast, panel (b) shows that the run with a high R value hovers in the vicinity of two saddle points for an extended period of time, indicated by the distance to the first and second saddle point manifolds D_1 and D_2 becoming very small in portions of the run. Concurrently, the magnitude of the gradient falls drastically in these two areas, and recovers only when D_2 starts to rise, confirming that the run indeed passed closely by a saddle point. The combination of these two plots suggests that control trajectories with higher R values pass more closely by saddle points.

5.2.2 Unitary transformation control landscape

Figure 6 also illustrates the distance to critical submanifolds on the unitary transformation control landscape along with the magnitude of the gradient, plotted using Eq. (17). In panel (c) the run with the lowest R value of 1.1533 out of the 1000 runs performed is shown, and in panel (d) a run with a high R value of 1.7300 is displayed. The distance to the first saddle point D_1 assumes lower values during the high R run as compared to the low R run, and correspondingly the magnitude of the gradient dips when the system approaches the critical submanifold. These plots support the notion that optimization trajectories generated by a gradient-based search algorithm which pass closely by saddle point critical submanifolds tend to have higher R values than trajectories that stay clear of saddle points. One way of interpreting this result is that saddle points tend to attract the paths of steepest ascent away from proceeding to a global maximum in a straight line. Instead

they veer towards the closest critical submanifold, and as result take longer to reach an optimal control due to the drop in the magnitude of the gradient near critical submanifolds.

5.3 Hessian-gradient eigenrelation

The analytical results of Section 4 that apply to straight shot trajectories when using the gradient rule are striking in that they postulate a novel and unusual connection between the gradient function and the Hessian along the trajectory. In this section, we provide numerical evidence to verify the Hessian-gradient eigenrelation for straight shots. In order to obtain a control trajectory with an R value very close to 1.0, we consider a very simple three level quantum system with Hamiltonian and dipole matrix

$$H_0 = \begin{pmatrix} -10 & 0 & 0 \\ 0 & -5 & 0 \\ 0 & 0 & 5 \end{pmatrix} \qquad \mu = \begin{pmatrix} 0 & -1 & -0.5 \\ -1 & 0 & 1 \\ -0.5 & 1 & 0 \end{pmatrix} \tag{34}$$

where the signs of the dipole matrix elements were chosen randomly, and we use the cost function of Eq. (8), with

$$\rho = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \qquad O = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{35}$$

This landscape was chosen because empirically, the it contains the least amount of structural features that would compromise straight shots (c.f. Section 5.1). We then used a stochastic algorithm, the Particle Swarm Optimization (PSO) algorithm [24], to locate a control trajectory for which $R = 1.0026$. More information about the PSO algorithm is presented in Appendix B.

Eq. (24) states that along a straight shot, the gradient function is an eigenvector of the Hessian matrix. To affirm that the gradient vector aligns with some eigenvector of the Hessian, at each step along the control trajectory, all eigenvectors of the Hessian matrix were found. Then, the dot product of the unit gradient vector with each unit eigenvector of the Hessian was computed. In Figure 7, the square of these values are shown as the smaller, multicolored circles. Each color corresponds to a different eigenvector of the Hessian, although the same color does not correspond to a unique eigenvector for different values of s. The unit vector which results from the Hessian acting on the gradient was also computed, and its dot product with the unit gradient vector was also found for each value of s. The square of these values are displayed as the thicker, wider black circles in Figure 7. $\mathrm{Tr}\left(\rho(T,s)O\right)$ is also displayed for reference. Except for a region in the middle where $0.8 \leq s \leq 1.0$, the black dots remain very close to 1.0, indicating that the Hessian does not change the direction of the gradient vector by acting on it, except for up to a sign. Also, for most of the climb, one of the eigenvectors of the Hessian remains either parallel or antiparallel to the

gradient vector. The exception is again in the region where $0.8 \leq s \leq 1.0$ and at the beginning of the climb. The cause of the discrepancy for $0 \leq s \leq 0.2$ is unknown. However, for the middle of the climb, the sign of the eigenvalue of the gradient vector flips. We have observed this by calculating the dot product of the unit gradient vector with the vector that results when the Hessian acts on the unit gradient vector for all values of s. In Figure 8, this quantity is represented by the larger black circles. When the Hessian does not alter the direction of gradient vector by its action, or when the large black circles in Figure 7 are near 1.0, this quantity is the eigenvalue of the gradient vector. The eigenvalues of the Hessian matrix are plotted as the smaller multicolored circles against s for this control trajectory. It appears to follow a particular eigenvalue of the Hessian throughout the optimization, thus seeming to affirm that the gradient vector is identical to one of the eigenvectors of the Hessian, up to a scale factor.

Further verification of Eq. (24) is provided by Figure 9. There, the quantity displayed in Figure 8 is again plotted against the first factor on the right hand side of Eq. (24), $\frac{\rho''(s')}{\rho'(s')}$. Since the control trajectory used was not an exact straight shot (i.e., it does not have $R = 1.0$), the function $\rho'(s) = \alpha(s)$ obtained from this trajectory is merely a constant t cross-section of the derivative of the trajectory $\frac{\partial E(s,t)}{\partial s}$. Nevertheless, excellent agreement between the two curves is observed.

6 Conclusions

The novel considerations of landscape structure first performed for state-to-state transition probability landscapes [14] have been extended by this paper to more general quantum ensemble and unitary transformation control landscapes. The results considered here also point towards simple landscape structure, which was first encountered experimentally [15] as well as hinted at by numerical studies [16, 17], all of which made use of the D-MORPH or gradient algorithm.

A confounding feature of the new landscapes considered in this work are additional critical submanifolds that possess saddle point topology. These submanifolds do not prevent optimizations from reaching the top of the landscape, but they can have an undesirable effect on their search effort. By quantifying this effort in the ratio R, we have shown that the presence of saddle points on the landscape has a small but non-negligible effect on the ease of optimization by means of distorting the underlying control trajectories. Importantly, the appearance of saddle points has not caused R to take on values in excess of 2.0 for the quantum systems examined here, indicating that these quantum control landscapes continue to subvert the 'curse of dimensionality' and possess unexpectedly simple structure.

A condition for the existence of control trajectories with $R = 1.0$, which would indicate ex-

ceedingly simple landscape structure, was also delineated. It was found that in order to obtain a straight control trajectory, the gradient has to be an eigenfunction of the Hessian, in addition to being separable in s and t. These conditions also happen to imply that all higher order Hessians possess a common eigenfunction, as well as a separable component. This is a strict requirement to place on these functions, and points to the great deal of freedom that a set of controls must be allowed on the control landscape in order to achieve trajectories with $R = 1$. Future work could explore the nature of the Hessian-gradient eigenrelation further, perhaps by using it to identify the conditions under which $R = 1$ control trajectories exist. Additional simulations exploring landscape structure on a wider variety of landscapes would also allow us to better understand the conditions that lead to low values of R. For example, does the fact that achieving unitary transformations requires control over every matrix element in the propagator lead to higher values of R?

Quantum control landscapes are inherently complex objects, being extremely high dimensional spaces and representing a highly nonlinear relationship encoded in Schrödinger's equation between the unitary propagator and a dynamical set of controls. The intuitive expectation is that optimizations taking place over these landscapes would be considerably difficult, but prior work has established that the topology of these landscapes is very simple when certain reasonable physical assumptions are met, paving the way for the success of any potential optimization [3, 7, 25, 26]. This paper has observed and attempted to explain the origin of highly favorable landscape structure as well, which could mean that any potential optimization is efficiently able to be carried out. This twofold simplicity of quantum control landscapes provides a foundation for explaining why optimal controls can be found without a great deal of effort. It is hoped that this paper will provide a stepping stone for advancing studies of control landscape structure.

Appendix A: Hierarchy of eigenrelations

A similar approach to deriving the eigenrelation Eq. (24), which lends itself to extending the relationship to all orders, is now given. Consider expanding the gradient function to first order along an arbitrary control trajectory

$$\frac{\delta J}{\delta E(s+ds,t)} = \frac{\delta J}{\delta E(s,t)} + ds \int_0^T \frac{\delta^2 J}{\delta E(s,t')\delta E(s,t)} \frac{\partial E(s,t')}{\partial s} \,\mathrm{d}t'. \tag{36}$$

Now, assume that $R = 1$ along the trajectory. Invoking Eq. (19) in order to substitute for $\frac{\partial E(s,t')}{\partial s}$ lets us arrive at Eq. (24). However, we shall dividing both sides by $\alpha(s) = \rho'(s)$ and write

$$\beta^{[2]}(s)\Delta E(t) = \int_0^T \frac{\delta^2 J}{\delta E(s,t')\delta E(s,t)} \Delta E(t) \,\mathrm{d}t', \tag{37}$$

where $\beta^{[2]}(s) = \frac{1}{\alpha(s)} \frac{d\alpha(s)}{ds}$. The form of Eq. (37) suggests that the Hessian possesses at least one eigenfunction that only depends on t. This suggests that the Hessian may be written

$$\frac{\delta^2 J}{\delta E(s,t')\delta E(s,t)} = \beta^{[2]}(s) K^{[2]}(t,t') + f_1(s,t,t') \tag{38}$$

where the purely time-dependent symmetric kernel satisfies the integral eigenvalue equation

$$\int_0^T K^{[2]}(t,t')\Delta E(t') \,\mathrm{d}t' = \Delta E(t) \tag{39}$$

and $f_1(s,t,t')$ is an arbitrary function satisfying

$$\int_0^T f_1(s,t,t')\Delta E(t') \,\mathrm{d}t' = 0. \tag{40}$$

Now consider expanding the Hessian to first order along an arbitrary control trajectory. We can write this expression as

$$\frac{\delta^2 J}{\delta E(s+ds,t')\delta E(s+ds,t)} = \frac{\delta^2 J}{\delta E(s,t')\delta E(s,t)} + ds \int_0^T \frac{\delta^3 J}{\delta E(s,t'')\delta E(s,t')\delta E(s,t)} \frac{\partial E(s,t)}{\partial s} \,\mathrm{d}t''. \tag{41}$$

Again, we assume that we are on a trajectory with $R = 1$ and use Eq. (19), and after multiplying by $\Delta E(t)$ and integrating over t' we obtain

$$\beta^{[3]}(s)\Delta E(t) = \int_0^T \int_0^T \frac{\delta^3 J}{\delta E(s,t'')\delta E(s,t')\delta E(s,t)} \Delta E(t'')\Delta E(t') \,\mathrm{d}t'' \,\mathrm{d}t' \tag{42}$$

where $\beta^{[3]}(s) = \frac{1}{\alpha(s)} \frac{d\beta^{[2]}(s)}{ds}$. The above equation similarly suggests that we may write the third-order derivative of J as

$$\frac{\delta^3 J}{\delta E(s,t'')\delta E(s,t')\delta E(s,t)} = \beta^{[3]}(s) K^{[3]}(t,t',t'') + f_2(s,t,t',t'') \tag{43}$$

where the time-dependent symmetric kernel $K^{[3]}(t,t',t'')$ satisfies the equation

$$\int_0^T \int_0^T K^{[3]}(t,t',t'')\Delta E(t'')\Delta E(t') \,\mathrm{d}t'' \,\mathrm{d}t' \tag{44}$$

and the arbitrary function $f_2(s,t,t',t'')$ satisfies

$$\int_0^T f_2(s,t,t',t'')\Delta E(t'')\,\mathrm{d}t'' = 0. \tag{45}$$

Continuing to iterate this process, we obtain the following hierarchy of eigenrelations for straight shot trajectories

$$\begin{aligned}
\beta^{[1]}(s)\Delta E(t_0) &= H^{[1]}(s,t_0)\\
\beta^{[2]}(s)\Delta E(t_0) &= \int_0^T H^{[2]}(s,t_0,t_1)\Delta E(t_1)\,\mathrm{d}t_1\\
\beta^{[3]}(s)\Delta E(t_0) &= \int_0^T\int_0^T H^{[3]}(s,t_0,t_1,t_2)\Delta E(t_2)\Delta E(t_1)\,\mathrm{d}t_2\,\mathrm{d}t_1\\
&\vdots\\
\beta^{[n]}(s)\Delta E(t_0) &= \int_0^T\int_0^T\cdots\int_0^T H^{[n]}(s,t_0,t_1,t_2,\cdots,t_{n-1})\Delta E(t_{n-1})\cdots\Delta E(t_2)\Delta E(t_1)\,\mathrm{d}t_{n-1}\cdots\mathrm{d}t_2\,\mathrm{d}t_1\\
&\vdots
\end{aligned} \tag{46}$$

with

$$\beta^{[n]}(s) = \frac{1}{\alpha(s)}\frac{d\beta^{[n-1]}(s)}{ds}, \qquad \beta^{[1]}(s) = \alpha(s) \tag{47}$$

where the nth order Hessian is defined as

$$H^{[n]}(s,t_0,t_1,\cdots,t_{n-1}) = \frac{\delta^n J}{\delta E(s,t_{n-1})\cdots\delta E(s,t_1)\delta E(s,t_0)}. \tag{48}$$

Each of these higher order Hessians can be expressed as the sum of a separable term and a free function:

$$H^{[n]}(s,t_0,t_1,\cdots,t_{n-1}) = \beta^{[n]}(s)K^{[n]}(t_0,t_1,\cdots,t_{n-1}) + f_{n-1}(s,t_0,t_1,\cdots,t_{n-1}) \tag{49}$$

where the arbitrary function $f_{n-1}(s,t_0,t_1,\cdots,t_{n-1})$ is symmetric in $\{t_i\}_{i=1}^{n-1}$ and satisfies

$$\int_0^T f_{n-1}(s,t_0,t_1,\cdots,t_{n-1})\Delta E(t_{n-1})\,\mathrm{d}t_{n-1}. \tag{50}$$

This implies each of the purely time dependent symmetric kernels satisfies the following hierarchy of eigenrelations

$$\begin{aligned}
\Delta E(t_0) &= \int_0^T K^{[2]}t_1(t_0,t_1)\Delta E(t_1)\,\mathrm{d}t_1\\
&= \int_0^T\int_0^T K^{[3]}(t_0,t_1,t_2)\Delta E(t_2)\Delta E(t_1)\,\mathrm{d}t_2\,\mathrm{d}t_1\\
&\vdots\\
&= \int_0^T\int_0^T\cdots\int_0^T K^{[n]}(t_0,t_1,t_2,\cdots,t_{n-1})\Delta E(t_{n-1})\cdots\Delta E(t_2)\Delta E(t_1)\,\mathrm{d}t_{n-1}\cdots\mathrm{d}t_2\,\mathrm{d}t_1\\
&\vdots
\end{aligned} \tag{51}$$

akin to Eq. (46).

These relationships show that the Hessians of all orders not only must be separable, but must also have $\Delta E(t)$ as a common eigenfunction. These requirements function as additional restrictions on a straight control trajectory. They signify that the controls must have sufficient freedom to achieve factorization of all higher order Hessians in order to proceed along a trajectory with $R = 1$.

Appendix B: The Particle Swarm Optimization Algorithm

The Particle Swarm Optimization (PSO) algorithm is a stochastic optimization algorithm that we employ to identify control trajectories with minimal R values. It creates a swarm of K particles, or trial control fields, in control space. In selecting the initial set of trial control fields, the PSO algorithm makes use of Eq. (29). The gradient algorithm is then used on each of these trial control fields to bring it to the initial landscape height J^I, as set out in Section 3. D-MORPH is then used to climb the landscape for all K trial fields from this initial height until the final normalized value J^F is reached; the resulting climbs from J^I to J^F are then projected onto control trajectories, and their R values are computed. The trial control field which resulted in the best, or lowest, value of R is then labeled by $E^{best}_{swarm}(t)$, and this point is then used to stochastically update each particle so that a new generation of trial control fields is created. The procedure is repeated for this new swarm, and $E^{best}_{swarm}(t)$ is relabeled. In addition, the best trial control field each particle has encountered in the history of the algorithm's current run is recorded. That information is also used along with the location of $E^{best}_{swarm}(t)$ in updating the set of trial fields. The algorithm then continues until a preset number of generations have been computed, and returns the control trajectory with the lowest value of R found over the course of the algorithm's execution.

There is no guarantee, however, that this algorithm will identify the trajectory possessing the globally minimal value of R, because there exists the possibility of the swarm "collapsing" around local minima of R and failing to explore more remote regions in control space that may contain trajectories with more optimal values of R. This is due to the PSO algorithm's stochastic nature, but also because the nature of the R landscape, that is, the functional relationship between R and the control field $E(t)$, is not known. The application of the PSO algorithm to searching for trajectories with small R values was first studied in [14]; more information regarding the mechanism of the algorithm is available in [24].

Acknowledgments

A. N. thanks Greg Riviello and Julia Yan for useful discussions. We acknowledge support from the Program in Plasma Science and Technology at Princeton University, the ARO, and the DOE.

References

[1] G. G. Balint-Kurti, S. Zou, and A. Brown. *Advances in Chemical Physics*, page 43. Wiley, 2008.

[2] K. W. Moore Tibbetts, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T. S. Ho, R. B. Wu, and H. Rabitz. Exploring the tradeoff between fidelity and time optimal control of quantum unitary transformations. *Phys. Rev. A*, 86:062309, Dec 2012.

[3] K. W. Moore and H. Rabitz. Exploring quantum control landscapes: topology, features, and optimization scaling. *Phys. Rev. A*, 84:012109, 2011.

[4] S. Sharma, H. Singh, and G. G. Balint-Kurti. Genetic algorithm optimization of laser pulses for molecular quantum state excitation. *The Journal of Chemical Physics*, 132(6):064108, 2010.

[5] R. S. Judson and H. Rabitz. Teaching lasers to control molecules. *Phys. Rev. Lett.*, 68:1500–1503, 1992.

[6] A. Weiner. Femtosecond pulse shaping using spatial light modulators. *Rev. Sci. Instrum.*, 71(5), 2000.

[7] R. Chakrabarti and H. Rabitz. Quantum control landscapes. *Int. Rev. Phys. Chem.*, 26(4), 2007.

[8] H. Rabitz, M. Hsieh, and C. Rosenthal. Landscape for optimal control of quantum-mechanical unitary transformations. *Phys. Rev. A*, 72:052337, Nov 2005.

[9] M. Hsieh and H. Rabitz. Optimal control landscape for the generation of unitary transformations. *Phys. Rev. A*, 77:042306, Apr 2008.

[10] K. W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz. Search complexity and resource scaling for the quantum optimal control of unitary transformations. *Phys. Rev. A*, 83:012326, Jan 2011.

[11] T. S. Ho and H. Rabitz. Why do effective quantum controls appear easy to find? *Journal of Photochemistry and Photobiology A: Chemistry*, 180(3):226 – 240, 2006. ¡ce:title¿Coherent Control of Photochemical and Photobiological Processes¡/ce:title¿.

[12] R. Wu, H. Rabitz, and M. Hsieh. Characterization of the critical submanifolds in quantum ensemble control landscapes. *Journal of Physics A: Mathematical and Theoretical*, 41(1):015006, 2008.

[13] M. Hsieh, R. Wu, and H. Rabitz. Topology of the quantum control landscape for observables. *The Journal of Chemical Physics*, 130(10):104109, 2009.

[14] A. Nanduri, A. Donovan, T. S. Ho, and H. Rabitz. Exploring Quantum Control Landscape Structure. *ArXiv e-prints*, May 2013.

[15] J. Roslund and H. Rabitz. Experimental quantum control landscapes: Inherent monotonicity and artificial structure. *Phys. Rev. A*, 80:013408, 2009.

[16] K. W. Moore, M. Hsieh, and H. Rabitz. On the relationship between quantum control landscape structure and optimization complexity. *J. Chem. Phys.*, 128(15), 2008.

[17] A. Donovan, V. Beltrani, and H. Rabitz. Quantum control by means of hamiltonian structure manipulation. *Phys. Chem. Chem. Phys.*, 13:7348–7362, 2011.

[18] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information (Cambridge Series on Information and the Natural Sciences)*. Cambridge University Press, 1 edition, 2004.

[19] A. Rothman, T. Ho, and H. Rabitz. Observable-preserving control of quantum dynamics over a family of related systems. *Phys. Rev. A*, 72(2), 2005.

[20] A. Rothman, T. Ho, and H. Rabitz. Quantum observable homotopy tracking control. *J. Chem. Phys.*, 123(13), 2005.

[21] Q. Sun, R. Wu, G. Riviello, and H. Rabitz. Searching for saddle points over the quantum control landscapes. In preparation.

[22] G. Riviello. Private communication.

[23] J. Dominy and H. Rabitz. Volume fractions of the kinematic 'near-critical' sets of the quantum ensemble control landscape. *Journal of Physics A: Mathematical and Theoretical*, 44(25):255302, 2011.

[24] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proc. IEEE Symposium on Swarm Intelligence*, pages 1942–1948, 1995.

[25] H. Rabitz, M. Hsieh, and C. Rosenthal. Quantum optimally controlled transition landscapes. *Science*, 303(5666):1998–2001, 2004.

[26] H. Rabitz, T. S. Ho, M. Hsieh, R. Kosut, and M. Demiralp. Topology of optimally controlled quantum mechanical transition probability landscapes. *Phys. Rev. A*, 74:012721, 2006.

Figures

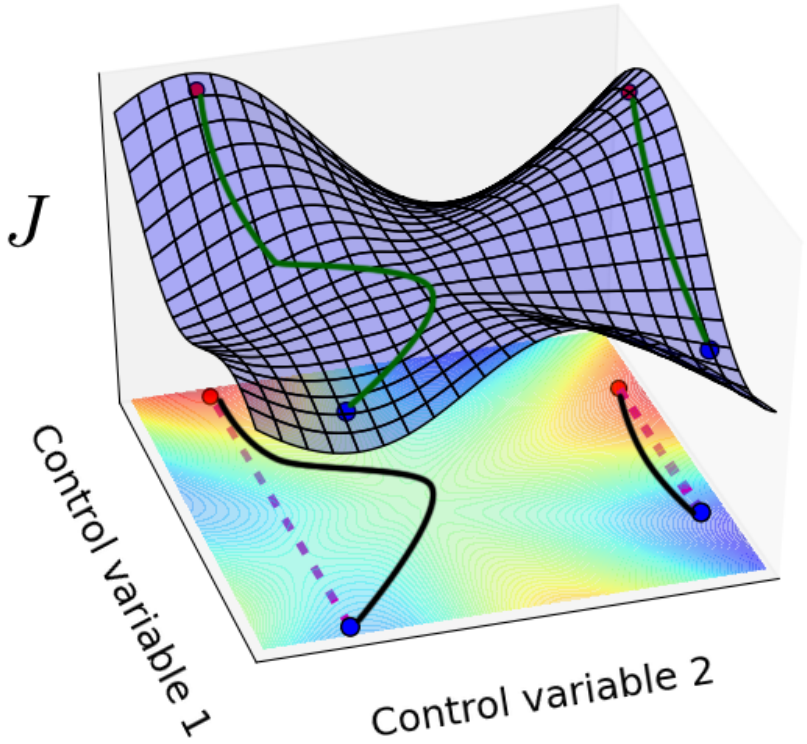

Figure 1: Illustrative sketch of a quantum control landscape (blue surface) and its underlying control space (colored contour map), displayed as the projection onto the plane spanned by two control variables (in practice many more are used). The vertical axis represents a cost function that we wish to maximize. The two optimization paths shown, in green, differ in the nature of their R values. The path on the right climbs in a straightforward manner to the top of the control landscape. Consequently, its projection into control space, in black, has a length very close to the length of the straight line between the projection's endpoints, shown as a magenta dashed line. This path's R value is low. On the other hand, the optimization path on the right 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 high R value.

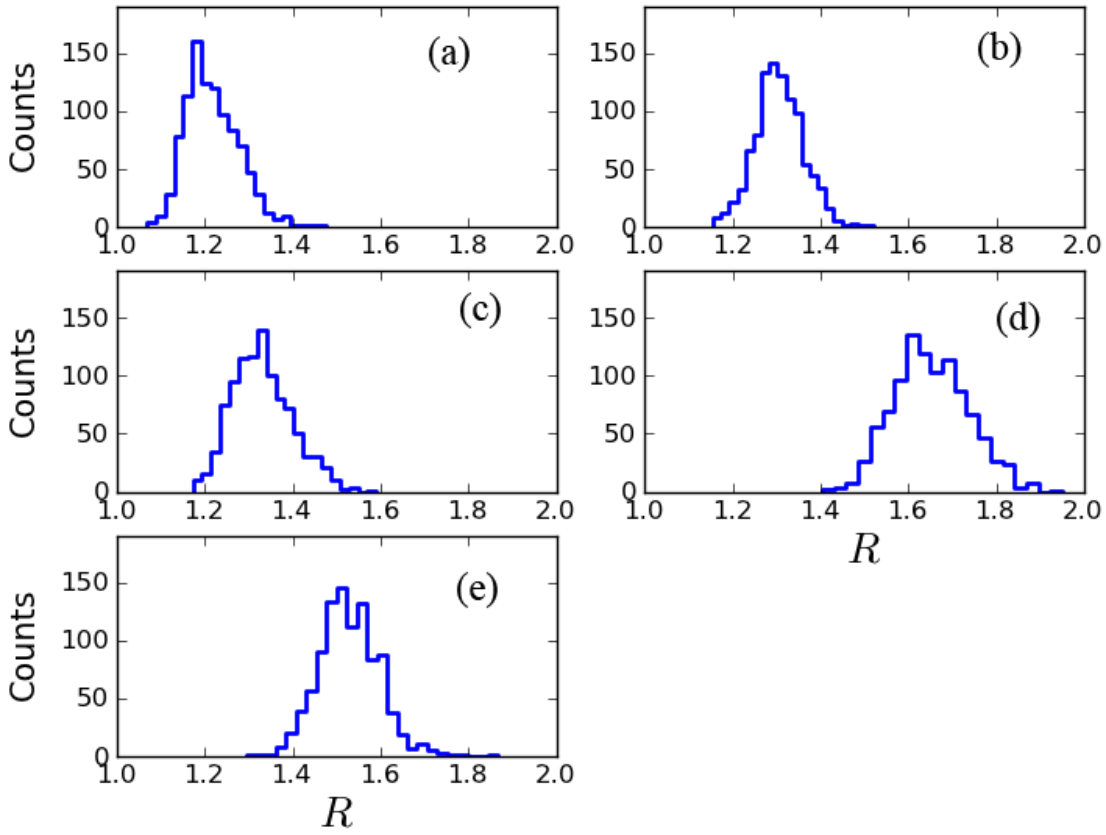

Figure 2: The distribution of R values for five different quantum ensemble control landscapes. 1000 optimizations were carried out for each landscape. (a) was generated using ρ_1 and O_1; (b) was generated using ρ_1 and O_2; (c) was generated using ρ_2 and O_1; (d) was generated using ρ_2 and O_2; and (e) was generated using ρ_3 and O_3. For each of these runs, the Hamiltonian operator of Eq. (28) and the dipole of Eq. (28) were used. All of the plots are centered close to 1.0 and skewed to the right, showing that R values are pushed to 1.0 across landscapes.

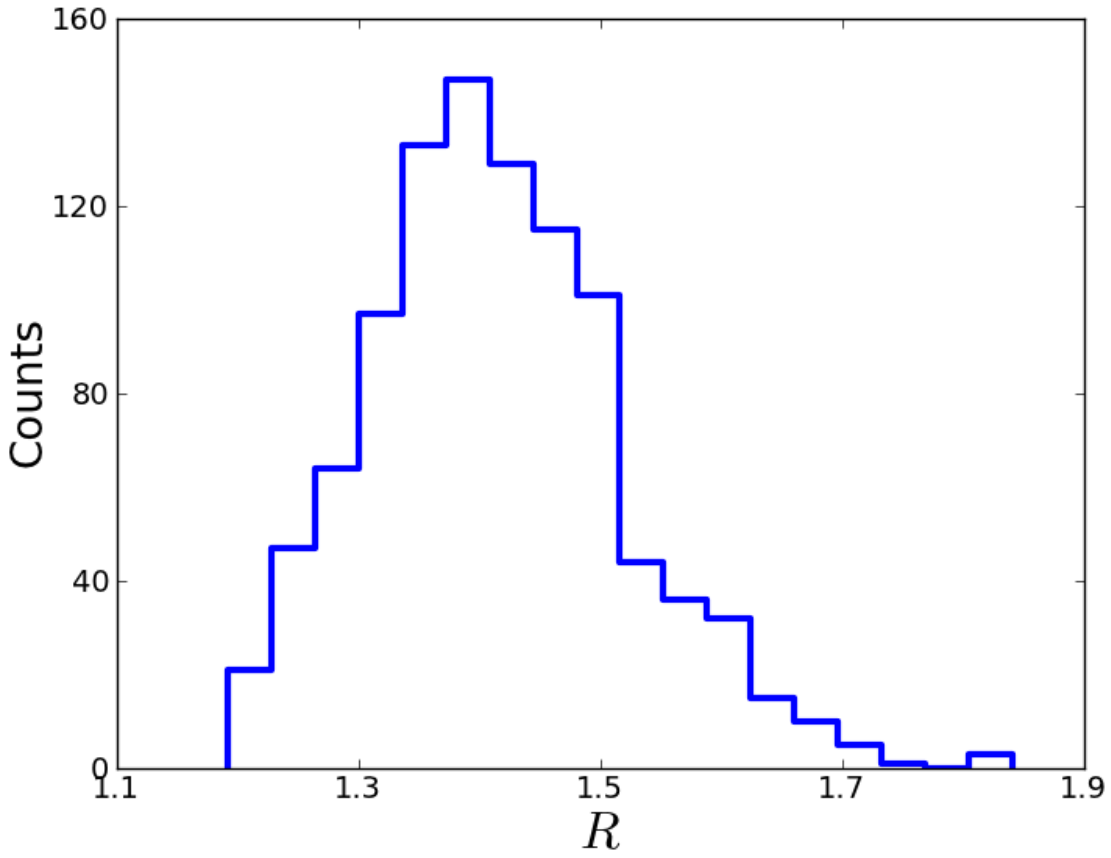

Figure 3: The distribution of R values for a unitary transformation control landscapes. 1000 optimizations were carried out, using the target unitary transformation W of Eq. (31) and Hamiltonian and dipole operators of Eq. (32)and (32). Modest values of R, all still less than 2.0, characterize this landscape's structure as well.

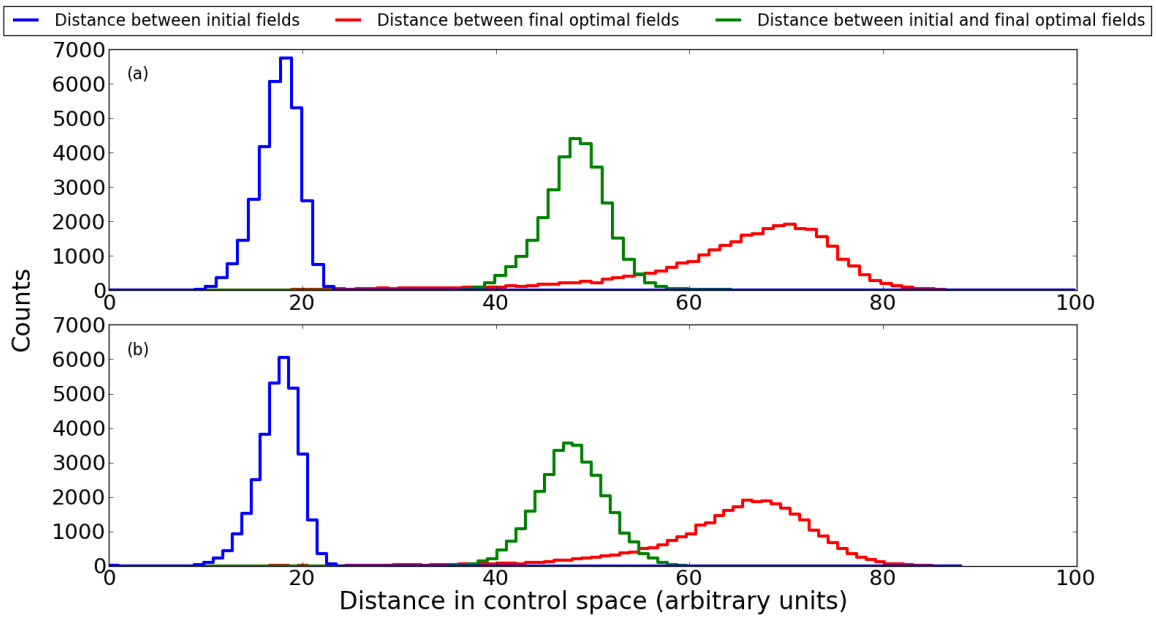

Figure 4: The distributions of pairwise distances between initial fields, final fields, and initial-final field pairs for the landscape generated by ρ_3 and O_3. The top plot was produced using the 250 optimizations with the lowest values of R, while the bottom plot was produced using the 250 optimizations that had the highest values of R. The similarity between the two plots suggests that low R and high R trajectories are distributed in the same way throughout control space.

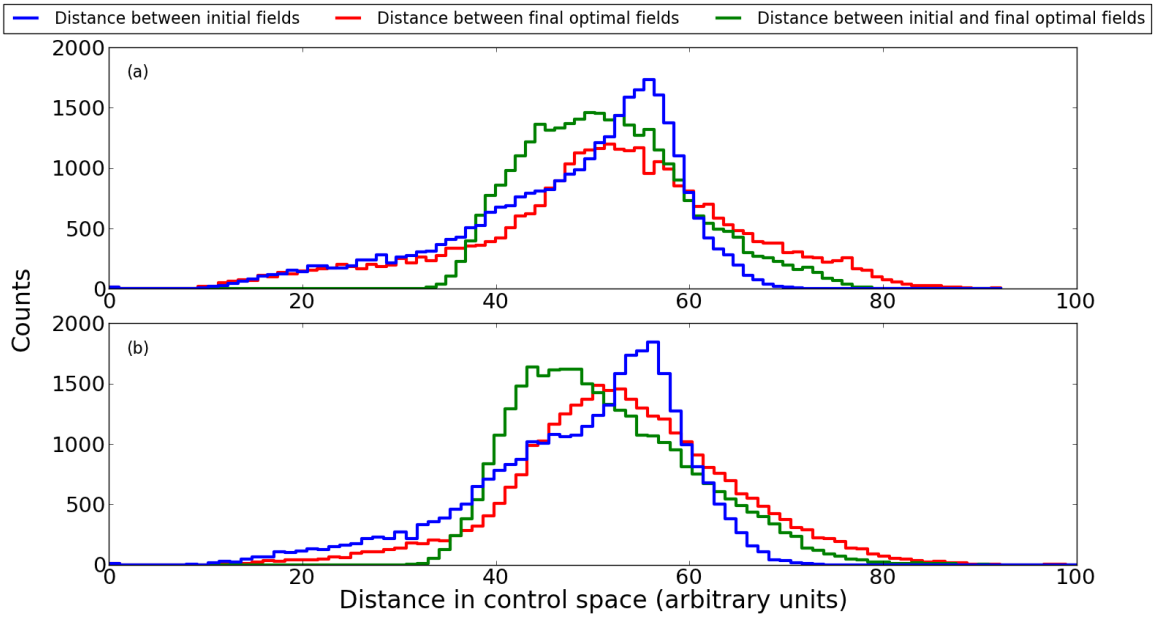

Figure 5: The distributions of pairwise distances between initial fields, final fields, and initial-final field pairs for the unitary transformation control landscape. The top plot was produced using the 250 optimizations with the lowest values of R, while the bottom plot was produced using the 250 optimizations that had the highest values of R. The similarity between the two plots suggests that low R and high R trajectories are distributed in the same way throughout control space.

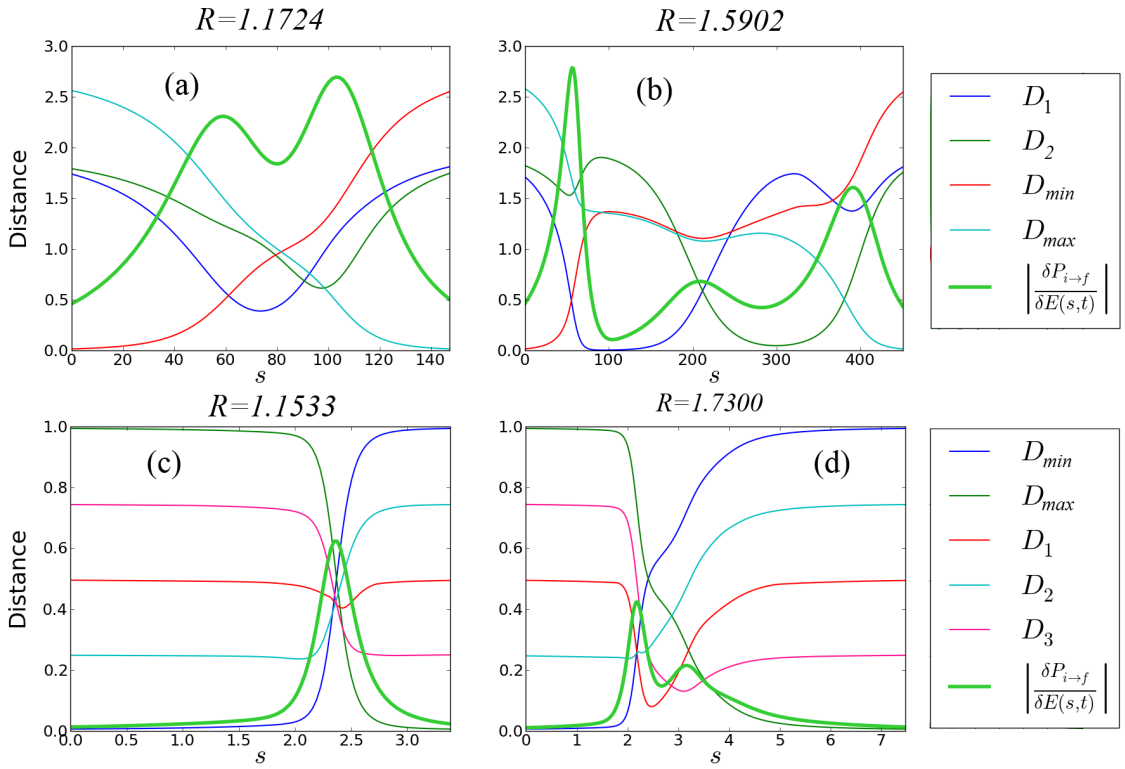

Figure 6: In (a) and (b), the distances to the critical submanifolds along low and high R valued optimizations on a quantum ensemble control landscape, respectively, are plotted using Eq. (13). It can be seen in (b) that the distance to a saddle point becomes exceedingly small along a high R trajectory, indicating that the saddle distorted the trajectory by acting as an attractor. The norm of the gradient function along the climb, shown as the thick bright green curve, also noticeably dips in (b), indicating that a saddle point has been encountered. In (c) and (d), the distances to the critical submanifolds along low and high R valued optimizations on a unitary transformation control landscape are plotted using Eq. (17). In (c), no saddle points are encountered, but in (d) the red curve dips sharply along with the norm of the gradient, signaling the draw of a saddle point and resulting in the corresponding high value of R.

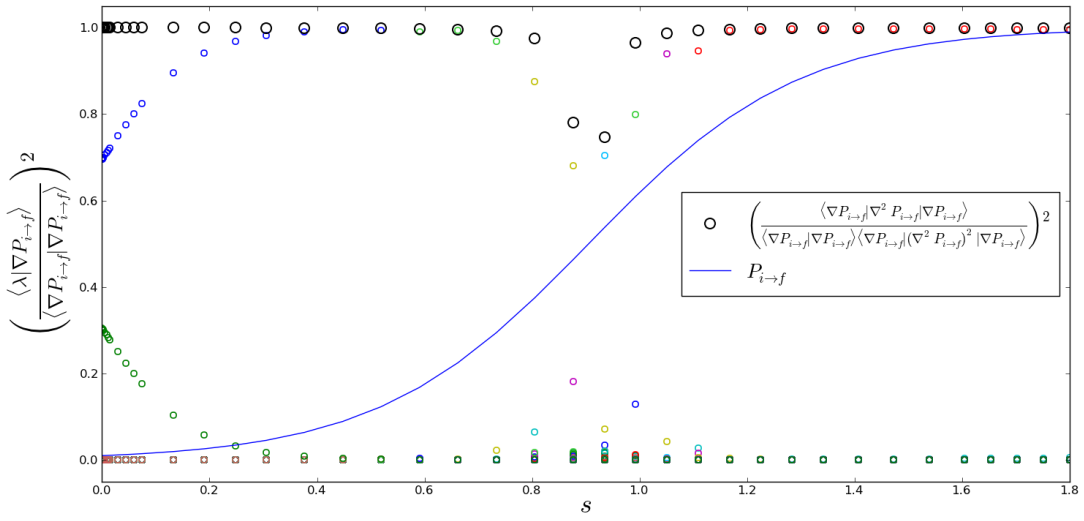

Figure 7: The thick black circles represent how closely aligned the gradient and the vector which results from the Hessian acting on the gradient are. Except for the very middle of the climb, where the orientation of the latter vector flips, they are nearly constantly aligned. The smaller colored circles are the dot product squared of the gradient with every eigenvector of the Hessian, and they indicate along most of the climb, the gradient can be identified with an eigenvector of the Hessian.

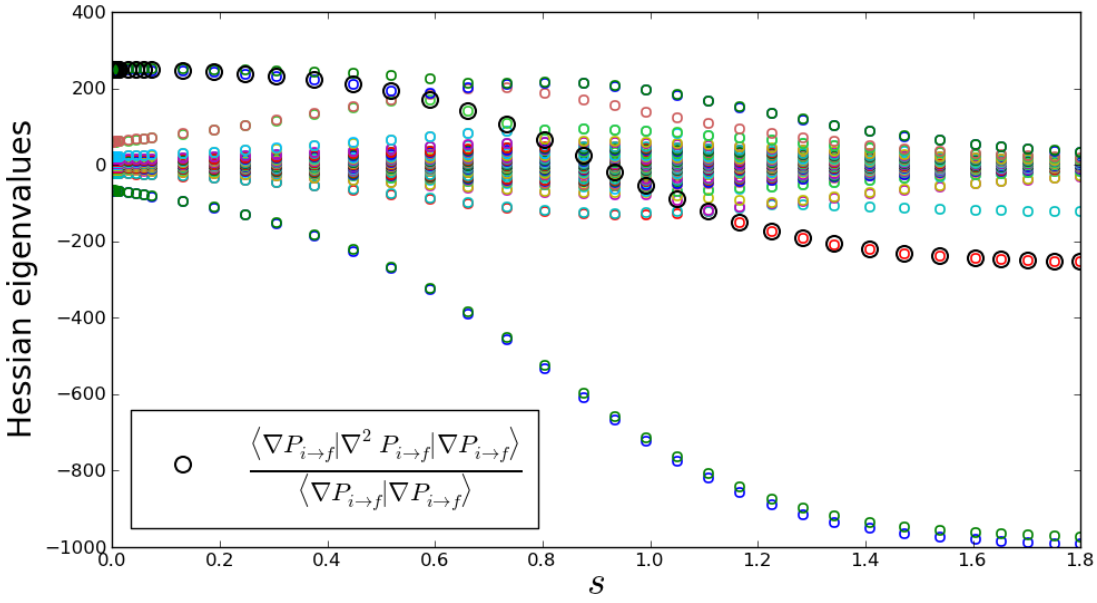

Figure 8: A plot of all the eigenvalues of the Hessian at every iteration of an optimization with $R = 1.0026$. The heavy black dots represent the factor by which the Hessian scales the gradient when it acts on it. Those dots follow an eigenvalue of the Hessian, indicating that the gradient can be identified with an eigenvector of the Hessian.

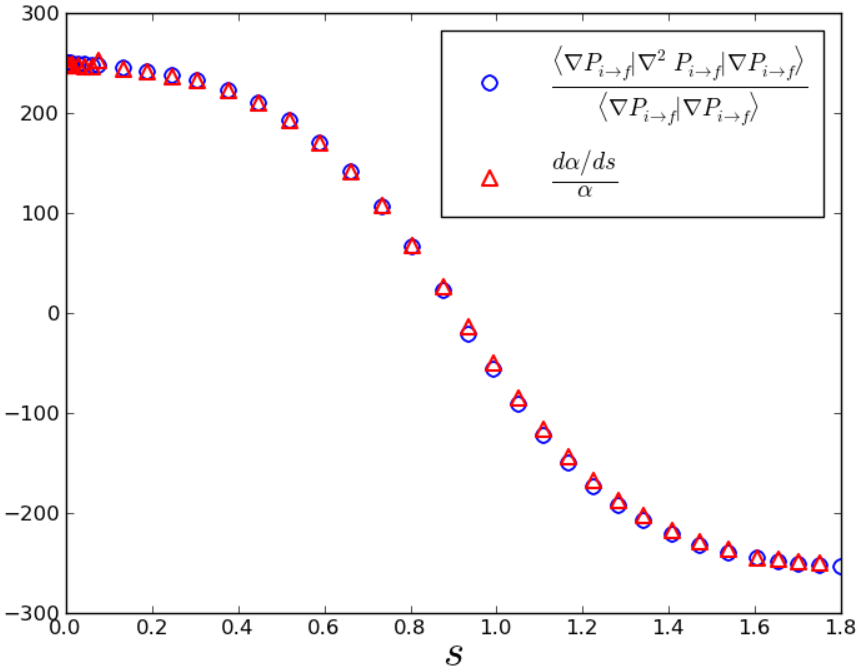

Figure 9: A plot of the factor the Hessian scales the gradient by when it operates on it, in blue, and a plot of $\frac{\alpha'}{\alpha} = \frac{\rho''}{\rho'}$, the eigenvalue belonging to the gradient according to Eq. (24). The agreement between the two values is very good, signifying the validity of Eq. (24).