



Assessing the structure of classical molecular optimal control landscapes

Carlee Joe-Wong^{a,b}, Tak-San Ho^b, Herschel Rabitz^{b,*}

^a Department of Electrical and Computer Engineering, Carnegie Mellon University, Silicon Valley, CA, USA

^b Department of Chemistry, Princeton University, Princeton, NJ, USA

ABSTRACT

We consider laser-driven optimal control landscape of a molecule from a classical mechanical perspective. The goal of optimal control in the present work is to steer the molecule from an initial state to a target state, denoted by two distinct points in phase space. Thus, a particular control objective is given as the difference between the final achieved phase space point and the target. The corresponding control landscape is defined as the latter control objective as a functional of the control field. While previous examination of the landscape critical points (i.e., a suboptimal point on the landscape where there is a zero gradient) has shown that the landscape *topology* is generally trap-free, the *structure* of the landscape away from these critical points is not well understood. We explore the landscape structure by examining an underlying metric defined as the ratio R of the gradient-based optimization path length of the control field evolution to the Euclidean distance between a given initial control field and the resultant optimal control field, where the latter field corresponds to a point at the top of the landscape. We analyze the path length-to-distance ratio R analytically for a linear forced harmonic oscillator and numerically for a nonlinear forced Morse oscillator. For the linear forced harmonic oscillator, we find that $R \leq \sqrt{2}$ and reaches its minimum value of 1 (i.e., corresponding to “a straight shot” through control space) in the large target time limit, as well as at special finite target times. The ratio R is similarly small for Morse oscillator simulations when following a steepest-ascent path to the top of the landscape, implying that the landscape is quite smooth and devoid of gnarled features. This conclusion is exemplified for a path discovered with $R \approx 1.0$ where simply following the *initial* gradient direction takes the climb very close to the top of the landscape. These findings are consistent with a variety of previous like simulations examining R in quantum control scenarios.

1. Introduction

Molecular optimal control theory generally considers system dynamics described by quantum mechanics, where it has provided a basis to understand the widely observed success of many control experiments [1,2]. In such studies, an optimally shaped laser field (or other fields appropriate for the applications) is applied to a particular system so as to achieve a desired control objective. For many objectives, one can show that the underlying quantum control landscape, defined as the physical objective as a function of the control field, is generally free of “traps” that could hinder common search algorithms (i.e., especially local gradient-based algorithms) from finding an optimal control field [3–8]. As molecular dynamics is often described well classically [9–11], recent work has extended the control landscape analysis to systems with dynamics described by classical mechanics [12,13]. For appropriate systems, classical mechanics can be more computationally tractable than quantum mechanics, which motivates the further exploration of classical optimal control landscapes.

Under reasonable system assumptions, a careful choice of objective function generally ensures the trap-free nature of classical control landscapes [14]. However, that analysis considered only the landscape *topology* specified by so-called “critical points” and did not explore the

structure of the control landscape lying between the lowest and highest critical points (i.e., the corresponding bottom and top of the landscape). Even if the control field does not encounter any landscape traps as it evolves towards an optimal field, the optimization path length along the excursion depends sensitively on the structure of the landscape and the optimization algorithm employed. Control fields in simulations are frequently optimized with a gradient algorithm, taking a steepest ascent climb of the landscape accompanied by a corresponding path through the space of control fields. Gnarled control landscapes, with many dips and ridges (yet still with no traps), can force the gradient to change direction many times, resulting in a convoluted path between the initial field and final optimal control field. Smooth, rather featureless landscapes, on the other hand, would take the control field along a near straight path from the initial to optimal field. If nearly straight control field trajectories are common, they have computational advantages – instead of re-computing the gradient of the objective at every step along the path, the gradient could be computed infrequently on the way towards the optimal control field. Thus, in this work, we explore the structure of classical optimal control landscapes by examining the nature of optimization paths taken as we proceed from initially poor control fields (i.e., producing dynamical outcomes far from the desired goal) towards optimal fields.

* Corresponding author.

E-mail addresses: cjowong@andrew.cmu.edu (C. Joe-Wong), tsho@princeton.edu (T.-S. Ho), hrabitz@princeton.edu (H. Rabitz).

Previous gradient-based investigations of landscape structures in the quantum mechanical regime of closed finite-level systems have shown that control fields follow remarkably straight paths [15–18], indicating that the landscapes are largely featureless, as well as *almost always* expected to be free of traps [19,20]. These results, although developed under ideal conditions, indicate that optimal control fields may be readily found, which is consistent with the broad laboratory success of quantum optimal control experiments [1,2] as well as simulations, where the gradient or other algorithms are used. Prior to the present work, it was an open question about whether this structural simplicity carried over to classical control landscapes, although they have already been shown to share the common lack of landscape traps under satisfaction of three fundamental assumptions [9,10]. Previous work on optimal control further suggests that classical and quantum molecular dynamical models often give qualitatively similar results [21–24]. However, classical mechanics is fundamentally distinct from quantum mechanics; for instance, the phase space dynamics for classical systems can be severely nonlinear and unstable, unlike the linear Schrödinger equation for quantum state dynamics. These differences further warrant the present investigation into the structure of classical control landscapes.

This paper considers steering the dynamics in the classical phase space from an initial point to a target point, which is analogous to the investigation of state-to-state quantum transition probabilities in [15]. The paper starts with consideration of the linear forced harmonic oscillator, which is studied as a physically relevant model that permits analytically identifying bounds on the length of the control field optimization path. We then conduct numerical simulations for a nonlinear Morse oscillator, showing that even in the presence of nonlinear state dynamics, the control field optimization paths remain remarkably straight. Thus, in this work, we show that classical control landscapes, like their quantum counterparts, are generally devoid of gnarled structure.

We finally remark that classical systems have the potential for chaotic behavior, and numerical instability in evolving the control field can possibly lengthen the path taken towards an optimal field. However, some systems that are chaotic without a control have even been shown to exhibit non-chaotic behavior under control [25–28]. In essence, the potential for chaos could pose a numerical or experimental issue for optimizing control fields within classical dynamics, rather than a fundamental one. Consideration of the chaotic regime, however, is beyond the scope of this paper.

In Section 2, we present the basic formulation for examining the landscape structure in terms of the ratio R of the observed optimization path length to the Euclidean distance between an arbitrarily chosen initial control field and the optimal field discovered by use of a gradient-based landscape climbing algorithm. Section 3, presents analytical bounds on R for a linear forced harmonic oscillator. We then numerically explore the paths taken when controlling a Morse oscillator in Section 4. Section 5 provides concluding remarks.

2. General formulation of landscape structure

We consider the classical dynamics of a molecule consisting of n atoms in the presence of a control field $\epsilon(t)$, $t \in [0, T]$. The evolution of the system in phase space is described by the time-dependent state variable $\mathbf{z}(t) = \begin{bmatrix} \mathbf{q}^T(t) \\ \mathbf{p}^T(t) \end{bmatrix}$, which is governed by Hamilton's equations,
$$\dot{\mathbf{z}}(t) = \begin{bmatrix} \left(\frac{\partial H}{\partial \mathbf{p}}\right)^T \\ -\left(\frac{\partial H}{\partial \mathbf{q}}\right)^T \end{bmatrix} = \begin{bmatrix} \mathbf{M}^{-1}\mathbf{p}(t) \\ -\frac{\partial \mathbf{V}(\mathbf{q}(t))}{\partial \mathbf{q}(t)} + \frac{\partial \mathbf{A}(\mathbf{q}(t))}{\partial \mathbf{q}(t)} \epsilon(t) \end{bmatrix}, \quad (1)$$

where the superscript “T” denotes a vector/matrix transpose, the diagonal square matrix \mathbf{M} is composed of the individual atomic masses [10]. The vectors $\mathbf{q}(t)$ and $\mathbf{p}(t)$ are, respectively, the position and

momentum coordinates of the constituent atoms, and the system Hamiltonian $H(\mathbf{q}(t), \mathbf{p}(t))$ is written as

$$H(\mathbf{q}(t), \mathbf{p}(t), t) = \frac{1}{2} \mathbf{p}^T(t) \mathbf{M}^{-1} \mathbf{p}(t) + \mathbf{V}(\mathbf{q}(t)) - \mathbf{A}(\mathbf{q}(t)) \epsilon(t). \quad (2)$$

We assume that the potential and dipole functions, \mathbf{V} and \mathbf{A} respectively, are twice-differentiable in $\mathbf{q}(t)$.

The goal of this paper is to explore the structure of the classical molecular control landscape defined by a given objective $J[\epsilon(\cdot)] = O(\mathbf{z}(T))$ as a functional of the control field $\epsilon(t)$. While in general $O(\mathbf{z}(T))$ can be any twice-differentiable function, which in itself does not have traps over \mathbf{z} [14], we assume in this work that it has the quadratic form

$$J[\epsilon(\cdot)] = O(\mathbf{z}(T)) = -(\mathbf{z}(T) - \mathbf{z}^{\text{tar}})^T Q(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}). \quad (3)$$

where \mathbf{z}^{tar} is a specified target state and Q is a symmetric, positive-definite matrix. Thus, O is a concave function of $\mathbf{z}(T)$, representing the weighted, negative squared distance between $\mathbf{z}(T)$ and the target \mathbf{z}^{tar} .

To facilitate the landscape structure analysis, we only deal with molecules that are considered to satisfy the following three fundamental assumptions: (1) they are controllable, (2) the gradient of the final state with respect to the control field, $\{\delta \mathbf{z}(T)/\delta \epsilon(t), t \in [0, T]\}$, comprises a full-rank (surjective) collection of $2n$ -dimensional vectors (i.e., linearly independent for any $\mathbf{z}(T)$ generated by an arbitrary $\epsilon(\cdot)$), and (3) the control resources are freely available (i.e., there are no physically relevant constraints on ϵ). Taken together, satisfaction of these three assumptions are sufficient to ensure that the landscape $J[\epsilon]$ contains no traps [9], i.e., a gradient algorithm will evolve the control field so as to successfully maximize J . However, satisfaction of the assumptions does not indicate the *paths* that the field takes to maximize J when starting from an arbitrary initial field. In particular, as pointed out in the introduction, a gnarled landscape with many twists, turns, and ridges could still allow the gradient algorithm to reach the top of the landscape, but a long path would likely be needed to traverse the landscape, instead of the limiting ideal case where the landscape is appropriately smooth to permit taking a nearly straight line path from the initial to the optimal field.

2.1. Path length ratio metric R characterizing the landscape structure

To evaluate the landscape structure, we consider a metric $R(s)$, defined as the ratio of (i) the path length $d_{PL}(s)$ taken by the control field along the gradient climb to (ii) the straight-line Euclidean distance $d_E(s)$ between the initial and optimized fields, i.e.,

$$R(s) = \frac{d_{PL}(s)}{d_E(s)} \geq 1, \quad (4)$$

where the parameter $s (\geq 0)$ is introduced to track (i.e., label) the path of the field traversing the control space and correspondingly its image as a trajectory climbing the landscape. Starting with the same, albeit arbitrary, initial point $\mathbf{z}(s, 0) = \mathbf{z}^{\text{init}}$ $\forall s \geq 0$ in phase space, the initial field $\epsilon(s = 0, t)$, $0 \leq t \leq T$, is generally expected to give a poor (i.e., large negative) value for $O(\mathbf{z}(s = 0, T))$ in Eq. (3), while the final optimal field $\epsilon(s = S, t)$, $0 \leq t \leq T$, yields the maximal value of $O(\mathbf{z}(s = S, T))$ to an acceptable tolerance at the end of the landscape climb where $s = S$. The path length to distance ratio $R(s)$, as a function of s , can be used to assess the features of the landscape encountered in continuously optimizing the control field. Furthermore, since $R(s)$ depends on the initial field, we can evaluate the distribution of R values for a given molecule and its optimization landscape by following the pathway up the landscape using an ensemble of chosen initial fields; note that with the gradient algorithm a specific unique path will be taken, $\epsilon(s = 0, t) \rightarrow \epsilon(s = S, t)$, connecting the specified initial and final phase space points.

Other algorithms not based on the gradient may be used to climb the landscape and explore R , and the nature of the algorithm can

influence R (e.g., a stochastic algorithm is likely to take a rather erratic path up the landscape [15]). Here we use the gradient algorithm as it is (a) most commonly employed in simulations and (b) particularly sensitive to local landscape features. This paper will use the D-MORPH formulation of the gradient algorithm [29]. A gradient algorithm guiding a path up an essentially featureless landscape will likely take a near-straight path, i.e., the gradient will point directly towards the top of the landscape, evolving the field over s in a nearly straight line through control space towards an optimal field. Such a path corresponds to a value of R close to 1. More gnarled landscapes, however, with additional features can force the gradient to repeatedly change direction and increase the length of the path taken by the control field (i.e., to its optimal form $\epsilon(s, t)$, $0 \leq t \leq T$) relative to the Euclidean distance between the initial and final fields, thereby producing a larger value of R .

The gradient-based path length $d_{PL}(s)$ and the associated Euclidean distance $d_E(s)$ can be expediently calculated via the tracking parameter s , where the control field $\epsilon(s, t)$, $t \in [0, T]$, at the value s on the landscape climb corresponds to the yield $J(s) = O(s) \equiv O(\mathbf{z}(s, T))$ in Eq. (3). The generally expected topological absence of traps implies that under the gradient algorithm, the control field $\epsilon(s, t)$ will eventually evolve over s such that the objective yield corresponds to reaching the top of the landscape (i.e., the field will maximize J as $s \rightarrow \infty$), which is in practice stopped at an acceptable value where $s = S$.

Using simple analysis from calculus, the path length traversed from $s = 0 \rightarrow S$ is

$$d_{PL}(S) = \int_0^S ds \sqrt{\frac{1}{T} \int_0^T \left(\frac{\partial \epsilon(s, t)}{\partial s} \right)^2 dt}, \quad (5)$$

while the Euclidean distance $d_E(S)$ between the initial, $\epsilon(0, t)$, and optimal, $\epsilon(S, t)$, control fields is given by

$$d_E(S) = \sqrt{\frac{1}{T} \int_0^T (\epsilon(S, t) - \epsilon(0, t))^2 dt}. \quad (6)$$

From Eqs. (4)–(6) we obtain

$$R(S) = \frac{\int_0^S ds \sqrt{\int_0^T \left(\frac{\partial \epsilon(s, t)}{\partial s} \right)^2 dt}}{\sqrt{\int_0^T (\epsilon(S, t) - \epsilon(0, t))^2 dt}}. \quad (7)$$

Although R can in principle take any value between 1 and ∞ [15], our simulations (cf., Section 4) in this paper show that it generally remains small and below ~ 2 . This result is qualitatively consistent with R determined in a variety of quantum control applications [15–18].

A simulation that achieves the minimum, $R = 1$, implies that a straight-line path was taken, $d_{PL}(S) = d_E(S)$, through control space to reach the optimal field. It has been shown that $R = 1$, if $\partial \epsilon(s, t)/\partial s$ is a separable function in s and t , i.e.,

$$\frac{\partial \epsilon(s, t)}{\partial s} = \alpha(s)\beta(t) \quad (8)$$

where $\alpha(s)$ and $\beta(t)$ are real scalar functions of s and t , respectively [15]. In particular, such linear “straight line” paths can be expressed as $\epsilon(s, t) = \nu(s)(\epsilon(S, t) - \epsilon(0, t)) + \epsilon(0, t)$, $\frac{\partial \epsilon}{\partial s} = \frac{d\nu(s)}{ds}(\epsilon(S, t) - \epsilon(0, t))$, where $\nu(s) \in [0, 1]$ is a monotonically increasing function with $\nu(0) = 0$ and $\nu(S) = 1$; comparing to Eq. (8), we have $\alpha(s) = d\nu(s)/ds$ and $\beta(t) = (\epsilon(s, t) - \epsilon(\infty, t))$. The separability of s and t dependence in Eq. (8) ensures that the gradient does not shift direction (i.e., it has the same dependence on t) as s changes, allowing the control field to evolve along a straight-line path to its optimal form $\epsilon(S, t)$. The finding of such straight paths would be numerically favorable, as they only require one gradient computation. Even a prevalence toward $R \sim 1$ could allow for infrequent evaluation of the gradient along the landscape climb.

2.2. Gradient-based D-MORPH algorithm

Within the framework of the D-MORPH gradient algorithm [29], the evolution of the form of the control field over time t is governed by a first-order initial-value differential equation in s ,

$$\frac{\partial \epsilon(s, t)}{\partial s} = \frac{\delta J}{\delta \epsilon(s, t)} = \frac{\delta J}{\delta \mathbf{z}(s, T)} \frac{\delta \mathbf{z}(s, T)}{\delta \epsilon(s, t)}, \quad (9)$$

which can be integrated from an initial trial choice $\epsilon(0, t)$ at $s = 0$, following the gradient $\delta J/\delta \epsilon(s, t)$, to produce all intermediate control fields $\epsilon(s, t)$, $t \in [0, T]$ for $s \in [0, S]$. It is readily seen that under the D-MORPH algorithm, the field corresponds to a monotonic climb towards the top of the landscape as s increases,

$$\frac{dJ}{ds} = \int_0^T \frac{\delta J}{\delta \epsilon(s, t)} \frac{\partial \epsilon(s, t)}{\partial s} dt = \int_0^T \left(\frac{\delta J}{\delta \epsilon(s, t)} \right)^2 dt \geq 0. \quad (10)$$

We note that for the numerical simulations in Section 4, Eq. (9) is discretized over time and the field $\epsilon(s, t)$ evolves as $\epsilon(s, t) \rightarrow \epsilon(s + ds, t)$, thereby increasing J . Eq. (9) is integrated with respect to s at each of the time points using a fourth-order Runge-Kutta method and the process is stopped at an acceptable convergence criterion, e.g., $-J \leq 0.01$.

3. Metric bounds for $R(S)$ with a linear forced harmonic oscillator

The linear forced harmonic oscillator has long served as a physically relevant reference case for molecular control, and it is natural to consider as a starting model for assessing landscape structure. Furthermore, we will show that a rigorous upper bound can be found for $R(S)$ in this case, and the later numerical simulation results in Section 4 with a Morse oscillator exhibit very similar qualitative R behavior to that found for the harmonic oscillator.

A one-dimensional linear forced harmonic oscillator is described by the Hamiltonian

$$H = \frac{p(t)^2}{2m} + \frac{kq(t)^2}{2} + aq(t)\epsilon(t), \quad (11)$$

where $aq(t)$ is the dipole function. Consider the cost function

$$J = -\frac{1}{2} \{ \mathbf{z}(T) - \mathbf{z}^{\text{tar}} \}^T \{ \mathbf{z}(T) - \mathbf{z}^{\text{tar}} \}, \quad (12)$$

where $\mathbf{z}(T) = \begin{bmatrix} q(T) \\ p(T) \end{bmatrix}$ is the 2×1 phase space column vector reached at the terminal time $t = T$, starting with some arbitrary phase space column vector $\mathbf{z}(0) = \begin{bmatrix} q(0) \\ p(0) \end{bmatrix} = \begin{bmatrix} q_0 \\ p_0 \end{bmatrix}$ at the initial time $t = 0$, and $\mathbf{z}^{\text{tar}} = \begin{bmatrix} q^{\text{tar}} \\ p^{\text{tar}} \end{bmatrix}$ is the target vector. Here q_0 and p_0 denote the initial position $q(t = 0)$ and momentum $p(t = 0)$, respectively.

For simplicity and without loss of generality, we will assume $m = k = 1$ in the following R metric bound analysis. This can be best realized by adopting the new variables: $t \rightarrow \sqrt{k/m} t$, $q \rightarrow \sqrt{k} q$, $p \rightarrow p/\sqrt{m}$, $a \rightarrow a/\sqrt{k}$. It can then be readily shown that for the linear forced oscillator in the presence of the control field $\epsilon(s, t)$, the terminal time phase point $\mathbf{z}(s, T)$ can be written as

$$\mathbf{z}(s, T) = \exp \left(\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} T \right) \times \left(\begin{bmatrix} q_0 \\ p_0 \end{bmatrix} - a \int_0^T \exp \left(-\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} t \right) \times \begin{bmatrix} 0 \\ 1 \end{bmatrix} \epsilon(s, t) dt \right), \quad (13)$$

where it is understood that the system starts with the same initial phase space point for all $s \geq 0$, i.e., $\mathbf{z}(s, 0) = \begin{bmatrix} q_0 \\ p_0 \end{bmatrix}$. It is then easily seen that Eq. (13) leads to the functional derivative

$$\frac{\delta \mathbf{z}(s, T)}{\delta \epsilon(s, t)} = -a \exp \left(\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} (T - t) \right) \times \begin{bmatrix} 0 \\ 1 \end{bmatrix} = -a \begin{bmatrix} \sin(T - t) \\ \cos(T - t) \end{bmatrix}, \quad (14)$$

which is independent of the control field $\epsilon(s, t)$. By using Eqs. (12) and

(14) in conjunction with the D-MORPH gradient algorithm, Eq. (9), we obtain the relation

$$\begin{aligned} \frac{\partial \epsilon(s, t)}{\partial s} &= \frac{\partial J}{\partial \epsilon(s, t)} \\ &= \frac{\partial J}{\partial \mathbf{z}(s, T)} \frac{\partial \mathbf{z}(s, T)}{\partial \epsilon(s, t)} \\ &= a \times \begin{bmatrix} \sin(T-t) \\ \cos(T-t) \end{bmatrix}^T \times \{\mathbf{z}(s, T) - \mathbf{z}^{\text{tar}}\} \end{aligned} \quad (15)$$

From Eqs. (14) and (15), we then derive the D-MORPH equation for the phase space point $\mathbf{z}(s, T)$ at the final time T as follows:

$$\begin{aligned} \frac{\partial \mathbf{z}(s, T)}{\partial s} &= \int_0^T \frac{\partial \mathbf{z}(s, T)}{\partial \epsilon(s, t)} \frac{\partial \epsilon(s, t)}{\partial s} dt \\ &= -a^2 L(\mathbf{z}(s, T) - \mathbf{z}^{\text{tar}}), \end{aligned} \quad (16)$$

where

$$L = \frac{1}{4} \begin{bmatrix} 2T - \sin(2T) & 1 - \cos(2T) \\ 1 - \cos(2T) & 2T + \sin(2T) \end{bmatrix} \quad (17)$$

is a positive definite matrix for $T > 0$. Since the matrix L is independent of the integration variable s , it follows that the solution of Eq. (16) can be written as

$$\mathbf{z}(s, T) - \mathbf{z}^{\text{tar}} = \exp(-a^2 L s) \times \{\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}\}. \quad (18)$$

Therefore, Eq. (15) can be further expressed as

$$\frac{\partial \epsilon(s, t)}{\partial s} = -a \times \begin{bmatrix} \sin(T-t) \\ \cos(T-t) \end{bmatrix}^T \times \exp(-a^2 L s) \times \{\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}\}, \quad (19)$$

which can be integrated over s to yield the control field

$$\begin{aligned} \epsilon(s, t) &= \epsilon(0, t) + \frac{1}{a} \times \begin{bmatrix} \sin(T-t) \\ \cos(T-t) \end{bmatrix}^T \times L^{-1} \times \{\exp(-a^2 L s) - I\} \\ &\quad \times \{\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}\}, \end{aligned} \quad (20)$$

where I denotes the 2×2 identity matrix and

$$L^{-1} = \frac{1}{T^2 - \sin^2 T} \begin{bmatrix} 2T + \sin(2T) & -1 + \cos(2T) \\ -1 + \cos(2T) & 2T - \sin(2T) \end{bmatrix}, \quad (21)$$

noting that $(T^2 - \sin^2 T) > 0$ for all $T > 0$. Thus, in the limit of $s \rightarrow \infty$, we have from Eq. (18) $\mathbf{z}(s, T) \rightarrow \mathbf{z}^{\text{tar}}$ as L is positive-definite, showing that the final state approaches the target as the control field $\epsilon(s, t)$ evolves towards an optimal control field, i.e.,

$$\lim_{s \rightarrow \infty} \epsilon(s, t) = \epsilon(0, t) - \frac{1}{a} \times \begin{bmatrix} \sin(T-t) \\ \cos(T-t) \end{bmatrix}^T \times L^{-1} \times \{\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}\}, \quad (22)$$

which is in agreement with the exact (inverse) solution of the corresponding classical Hamilton equation (i.e., Eq. (1)) for the linear forced harmonic oscillator [30].

From Eqs. (5), (15), and (18), it can be shown that the path length through the control space is

$$\begin{aligned} d_{PL}(S) &= a \int_0^S \sqrt{[\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]^T \exp(-a^2 L s) L \exp(-a^2 L s) [\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]} ds \\ &\leq \frac{1}{a} (\sqrt{\sigma_1} y_1 + \sqrt{\sigma_2} y_2), \end{aligned} \quad (23)$$

where the vector $\mathbf{y} \equiv \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = UL^{-1}(I - \exp(-a^2 LS))|\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}| \geq 0$, $\sigma_1 = 2/(T + |\sin T|) > 0$, and $\sigma_2 = 2/(T - |\sin T|) > 0$ (for $T > 0$). Here the unitary matrix U diagonalizes L such that $ULU^\dagger = \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix}$. Moreover, we have used the relation $\sqrt{|y_1(s)|^2 + |y_2(s)|^2} \leq |y_1(s)| + |y_2(s)|$ for $y_1, y_2 \in \mathbb{R}$ and $|\mathbf{x}| = \sqrt{x_1^2 + x_2^2}$ for the magnitude of any real number 2×1 column vector $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$. Similarly, using Eqs. (6), (15), and (18),

we find that the Euclidean distance between the initial and final fields, d_E , can be written as

$$\begin{aligned} d_E(S) &= \frac{1}{a} \sqrt{[\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]^T (I - \exp(-a^2 LS)) L^{-1} (I - \exp(-a^2 LS)) [\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]} \\ &= \frac{1}{a} \sqrt{\sigma_1 y_1^2 + \sigma_2 y_2^2}. \end{aligned} \quad (24)$$

By using Eqs. (23) and (24), in conjunction with the chain of inequalities $\sqrt{x_1^2 + x_2^2} \leq x_1 + x_2 \leq \sqrt{2(x_1^2 + x_2^2)}$ for arbitrary numbers $x_1 \geq 0$ and $x_2 \geq 0$, we arrive at an absolute metric upper bound

$$R(S) = \frac{d_{PL}(S)}{d_E(S)} \leq \frac{\sqrt{\sigma_1} y_1 + \sqrt{\sigma_2} y_2}{\sqrt{\sigma_1 y_1^2 + \sigma_2 y_2^2}} \leq \sqrt{2}, \quad (25)$$

indicating that R will be relatively small for the linear forced harmonic oscillator.

Interestingly, at the special times $T = n\pi$ (for the case of $k = m = 1$, or $\omega_0 T = n\pi$ for the more general case, where $\omega_0 \equiv \sqrt{k/m}$ is the characteristic frequency of the harmonic oscillator with a force constant k and mass m), $n = 1, 2, \dots$, the matrix L , Eq. (17), reduces to the simple diagonal form

$$L = \frac{T}{2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (26)$$

and, as a result, Eq. (19) can be written as

$$\frac{\partial \epsilon(s, t)}{\partial s} = -a \exp\left(-\frac{a^2 T}{2} s\right) \times \begin{bmatrix} \sin(T-t) \\ \cos(T-t) \end{bmatrix}^T \{\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}\}, \quad (27)$$

which is a product of two separate scalar functions of the parameter s and time t , respectively, cf., Eq. (8). To this end, it can be shown that the path length $d_{PL}(S)$, Eq. (23), and the corresponding Euclidean distance $d_E(S)$, Eq. (24), can be reduced to the same expression, i.e.,

$$\begin{aligned} d_{PL}(S) &= d_E(S) \\ &= \frac{1}{a} \times \sqrt{\frac{2}{T}} \times [1 - \exp(-a^2 TS/2)] \\ &\quad \times \sqrt{[\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]^T [\mathbf{z}(0, T) - \mathbf{z}^{\text{tar}}]}. \end{aligned} \quad (28)$$

Thus, the metric $R(S)$ shrinks to its lower limit of $R = 1$, implying that the underlying control field follows a straight-line path at the special times $T = n\pi$ (or more generally $\omega_0 T = n\pi$), $n = 1, 2, \dots$. Though the harmonic oscillator is a particular classical system, the findings above suggest that low R values may also arise for other types of systems; this behavior is confirmed in an anharmonic oscillator example in Section 4.

We will verify the analytical results above by numerically calculating R for several simulations with the harmonic oscillator and objective function $J = -(q(T)^2 + p(T)^2)/2$, corresponding to the target $q^{\text{tar}} = 0$ and $p^{\text{tar}} = 0$ in phase space. The parameters in Eq. (11) are set to be $m = 1$, $k = 2$, $a = 0.5$, with a final time of $T = 18\pi$ and initial state $(q(0), p(0)) = (1.69, 1)$. Note that this case corresponds to $k = 2$ and $T \times \sqrt{k/m} = 18\sqrt{2}\pi \neq n\pi$, n an integer; thus, it does not yield $R(S) = 1$. We initialize the control field to be a sum of three sine functions with randomly chosen amplitudes (between 0.0005 and 0.05) and phases (between π and $-\pi$), and frequencies of ω , 2ω , and 3ω , where $\omega = \sqrt{k/m}$ is the natural frequency of the unforced oscillator. Fig. 1(a) shows the distribution of R over 300 simulation runs; we see that the maximum R value is ~ 1.22 , which is below the theoretical limit of $\sqrt{2}$. As another example, Fig. 1(b) shows the distribution for 100 simulation runs with each of three different final states: $(q^{\text{tar}}, p^{\text{tar}}) = (1.565, -0.246)$, $(1.982, 1.537)$, and $(1.989, 0.5)$, utilizing the same initial condition in Fig. 1(a) for all cases. Again, R remains well below $\sqrt{2}$. Each simulation run was stopped when the objective function $J[\epsilon(s, \cdot)] \geq -0.0001$ (i.e., the landscape value J was within four decimal places of its maximum value of 0).

We finally tested the theoretical upper bound of $R \leq \sqrt{2}$ by using a stochastic particle swarm optimization directed to maximize R at the final control field (see Refs. [15,31] for details on the particle swarm

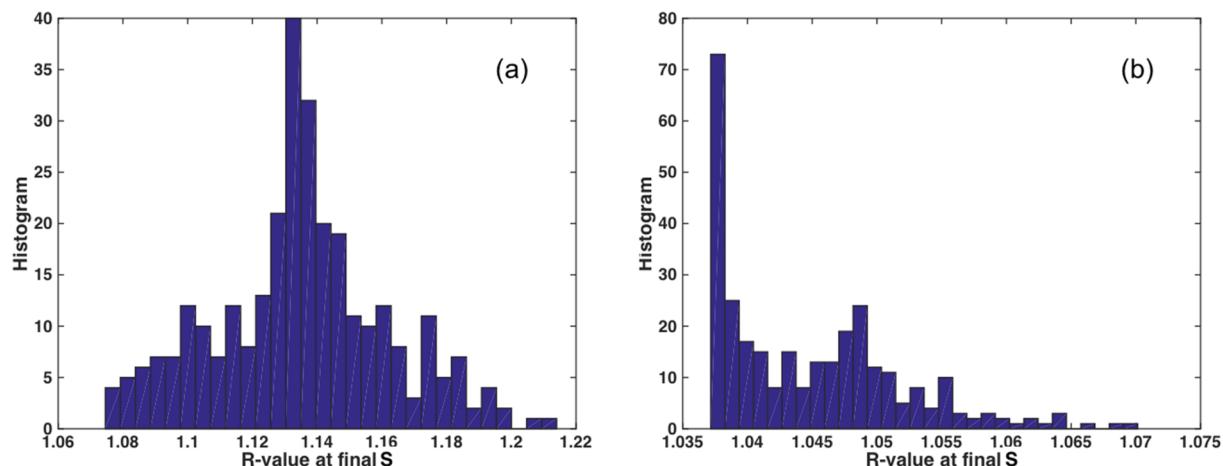


Fig. 1. Histograms of the final R values with the linear forced harmonic oscillator and different target states: (a) Target $(q^{\text{tar}}, p^{\text{tar}}) = (0, 0)$ and (b) the combined distributions for the targets $(q^{\text{tar}}, p^{\text{tar}}) = (1.565, -0.246)$, $(1.982, 1.537)$, and $(1.989, 0.5)$. The initial state in all cases is $(q(0), p(0)) = (1.69, 1)$.

algorithm). The target state is chosen as $q^{\text{tar}} = 0$ and $p^{\text{tar}} = 0$. The initial control field is taken as a sum of 50 sine functions, and the optimization variables to maximize R , given the initial state $(q(0), p(0))$, are the amplitudes and phases of the 50 sine functions. Thus, we are seeking to maximizing R by searching over the indicated space of *initial* fields, with the gradient algorithm then taking a path to the final field that optimized J and permitted a calculation of R for each field. We find that $R \leq 1.34$, which is well below the theoretical bound of $\sqrt{2}$. All of these findings for the linear forced harmonic oscillator will serve as a benchmark reference for the characteristic molecular situation of a nonlinear forced Morse oscillator in Section 4.

4. Simulations for a controlled Morse oscillator

To complement the analytical results for the linear forced harmonic oscillator above, in this section we numerically explore the landscape structure for a Morse oscillator by sampling many different initial control fields. We will see that similar R behavior arises, suggesting that nearly straight, gradient driven control trajectories are common as found in quantum control studies [15–18]. The Hamiltonian of the controlled Morse oscillator is given as

$$H(q(t), p(t), t) = \frac{p(t)^2}{2m} + D_0(1 - \exp(-\alpha(q(t) - r_e)))^2 - Aq(t)\exp(-\xi q(t)^4) \epsilon(t), \quad (29)$$

where we use the parameter values $m = 1732$, $D_0 = 0.2101$, $\alpha = 1.22$, $r_e = 1.75$, $A = 0.4541$, $\xi = 0.0064$. All quantities are given in atomic units (a.u.), and these parameter values approximate the behavior of an HF molecule [32–34] collinear with the applied field. For all the simulations, the final time is $T = 320\pi$, and the initial control field is taken as a sum of three sine functions, with amplitudes randomly drawn from the uniform distribution between 0.0005 and 0.05, and phases drawn from a normal distribution between π and $-\pi$. The frequencies of the three sine functions are, respectively, 1, 2, and 3 times the fundamental frequency of the unforced oscillator. Each control field through a gradient landscape climb has an imposed Gaussian envelope, $\exp[-(t - T/2)^2/20000]$, to assure that the field is essentially zero at the beginning and end of the control interval ($t = 0$ or $t = T$).

4.1. Distributions of the path length ratio metric R

We consider the initial state $(q(0), p(0)) = (1.75, 0)$, which is the equilibrium phase space point of the field-free oscillator; varying the initial state did not yield any significant difference in the resultant distribution of R values. Each of the distinct target phase states are taken to have an energy corresponding to either the 4th, 7th, or 11th

quantum energy level. The successive cases correspond to increasing sampling of the anharmonicity and the nonlinear dipole coupling of the Morse oscillator, cf. Eq. (29), with the objective function

$$J[\epsilon] = -[(q(T) - q^{\text{tar}})^2 + (p(T) - p^{\text{tar}})^2] \quad (30)$$

corresponding to Eq. (3) with $Q = I$. Thus, $-J$ represents the squared Euclidean distance between the final achieved state and the target. The control field optimization was stopped when $J \geq -0.01$, thus ensuring that $q(T)$ and $p(T)$ are within 0.1 a.u. of the target. To calculate R , we derive from Eq. (5) the following equation

$$\frac{d(d_{PL}(s))}{ds} = \sqrt{\int_0^T \frac{1}{T} \left(\frac{\partial \epsilon(s, t)}{\partial s} \right)^2 dt} \quad (31)$$

which is integrated, together with Eq. (9) using a fourth-order Runge–Kutta method, over s as the optimization proceeded to find the total path length of the control field $d_{PL}(s)$ during the optimization. The final value of R at $s = S$ was calculated by dividing $d_{PL}(S)$ by the Euclidean distance $d_E(S)$ between initial and final control fields in Eq. (6).

Fig. 2 shows the combined distribution of the final R values over 100 randomly chosen initial fields for each of the three targets: $(q^{\text{tar}}, p^{\text{tar}})$ is either $(1.565, -0.246)$, $(1.982, 1.537)$, or $(1.989, 0.5)$ having an energy corresponding to the 4th quantum energy level (delimited by the inner and outer classical turning points at $(q_1 - r_e) \approx -0.37$ and $(q_2 - r_e) \approx 0.63$, respectively, relative to the minimum potential position at $r_e = 1.75$) in the anharmonic region of the Morse oscillator and sampling the nonlinear nature of the dipole, cf. Eq. (29). Even the largest R value is quite small, lying below 1.5, which is qualitatively consistent with the analytical findings for the classical harmonic oscillator in Section 3 and qualitatively similar to previous findings for quantum control simulations [15,17]. Several fields are concentrated around $R \approx 1$ in Fig. 2, suggesting that near-straight control fields are quite prevalent. However, there is little obvious distinguishing features between fields that yield low and high R values. For example, Fig. 3 shows two illustrative optimal control fields that respectively correspond to low and high R values. It was found that despite the large differences in the respective R values (a) the fields possess comparable oscillatory structures with some evident character of an approximate sign flip starting at $\epsilon = 0$ and persisting at later times, while (b) the fields are clearly distinct in amplitude with the field at the low R (< 1.01) value having about twice the amplitude of that at the high R (> 1.45) value. We further remark that a detailed examination of all of the collective control fields showed quite varied behavior as well as some similarity in certain cases. Fig. 3 is presented as an example of where fields with a degree of similarity can give significantly different R values.

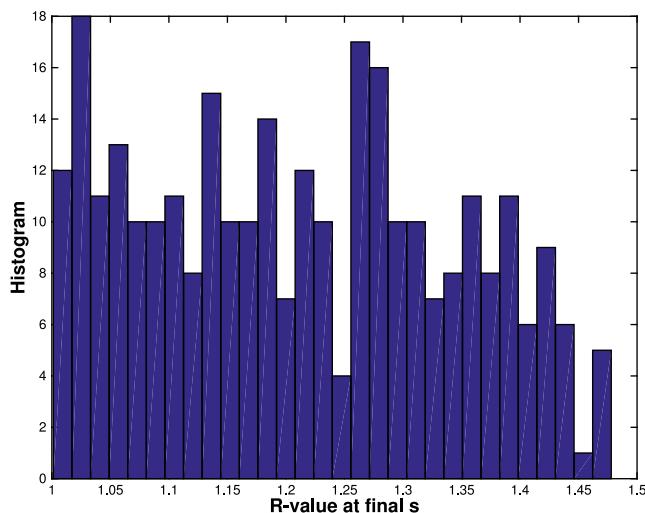


Fig. 2. Histogram of the combined final R values over 100 randomly chosen initial fields for each of three target states $(q^{\text{tar}}, p^{\text{tar}}) = (1.565, -0.246), (1.982, 1.537), (1.989, 0.5)$ for the anharmonic oscillator with phase space targets corresponding to the 4th quantum energy level. Each simulation was stopped when $J \geq -0.01$. The averaged final R value is $\langle R \rangle \approx 1.21$.

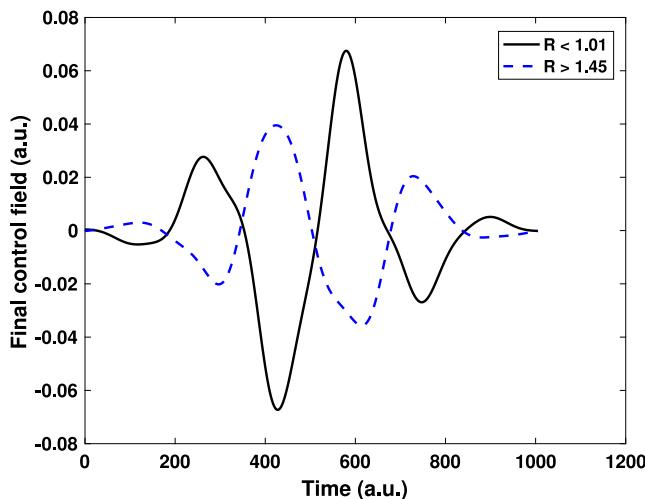


Fig. 3. Two randomly chosen optimal fields that gave either a low (solid curve) or a high (dashed curve) R value in Fig. 2 for the anharmonic oscillators corresponding to target phase space points at the 4th quantum energy level.

Some portion of the simulations in Fig. 2 started near the top of the landscape, which in the extreme case yielded a very small R_{\min} value of 1.001 as a result. In order to eliminate any potential such bias we now consider starting with initial fields that produce points in phase space quite far from the top of the landscape. Thus, we first “normalize” the fields to ensure that the initial value of J is sufficiently far from the final accepted value of $J \geq -0.01$; after choosing a random control field, we evolve it to an interim control field ϵ_i for which $J[\epsilon_i] \in [-25, -9]$. By following this procedure, the achieved state under an interim field ϵ_i is within a ring (i.e., of inner and outer radii, respectively, 9 and 25) around the target point $(q^{\text{tar}}, p^{\text{tar}})$. We then use a set of such interim fields $\{\epsilon_i\}$ to initiate the subsequent optimizations. Fig. 4 shows the histogram of the resultant R values over 120 runs with 40 random choices of ϵ_i to respectively reach each of the three target values $(1.565, -0.246), (1.982, 1.537)$, and $(1.989, 0.5)$ (i.e., the same targets utilized previously) in Fig. 2. While the R values remain below $R = 2$, we see that they have shifted away from $R = 1$, compared to the R -values shown in Fig. 2. Thus, starting very far from the target reduced

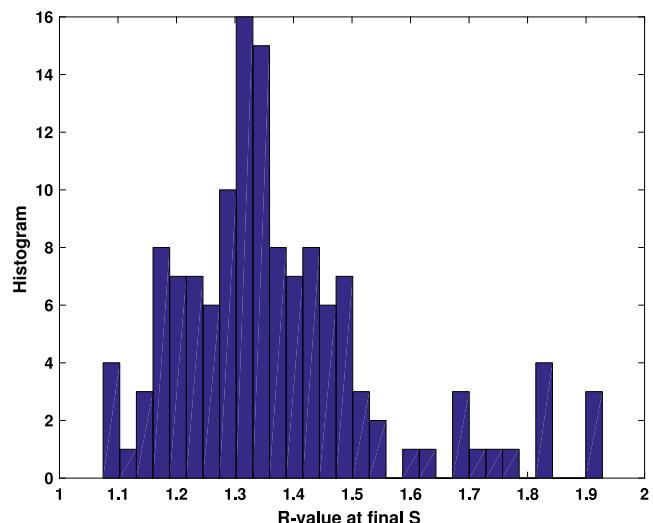


Fig. 4. Histogram of the resultant final R values reaching each of three target states $(q^{\text{tar}}, p^{\text{tar}}) = (1.565, -0.246), (1.982, 1.537), (1.989, 0.5)$, with initial control fields producing $J \in [-25, -9]$ values lying far from the targets at $J > -0.010$ for the anharmonic oscillator with the target phase space points corresponding to the 4th quantum energy level. The average value is $\langle R \rangle \approx 1.36$. A comparison with Fig. 2 shows the evident difference that very small R values (i.e., $R \lesssim 1.1$) are not present above, but the average values of $\langle R \rangle$ are very similar and no R value exceeds 2 in either case. Thus, all control paths are not strongly distorted, regardless of whether the initial field produces a J value close to, or far from, the target at $J = 0$.

the number of essentially straight shots across control space. Notwithstanding this change, Figs. 2 and 4 still share the common characteristic of dramatically small R values (i.e., the average values are $\langle R \rangle \approx 1.21$ in Fig. 2 and $\langle R \rangle \approx 1.36$ in Fig. 4), regardless of the initial field and the associated J value. Thus, all of the results show that subtle changes in $\langle R \rangle$ can arise depending on the imposed circumstances, but there is remarkable general similarity in the final R distribution.

We also found that similar results shown in the histograms given in Fig. 5, obtained from simulations for target phase space states having relatively much higher energies corresponding to the 7th quantum energy level [i.e., delimited by the inner and outer classical turning points $(q_1 - r_e) \approx -0.45$ and $(q_2 - r_e) \approx 1.05$] and the 11th quantum energy level [i.e., delimited by the inner and outer classical turning points $(q_1 - r_e) \approx -0.5$ and $(q_2 - r_e) \approx 1.57$] in the very anharmonic region and the strong nonlinear dipole domain of the Morse oscillator, Eq. (29). In each case, either corresponding to the 7th or 11th quantum energy level, 100 simulations was performed starting with random initial fields with the stopping criterion $J \leq 0.01$. The simulations sampled a range of target q^{tar} and p^{tar} values (i.e., p^{tar} ranged from 0.04 to 4 at intervals of 0.04 with q^{tar} taking whatever values needed to get to the corresponding energy level.) Fig. 5 shows that in both cases the resultant R values are all below 2 (i.e., between 1.0067 and 1.6125 with an average value $\langle R \rangle \approx 1.2455$ for the 7th level and between 1.0314 and 1.8934 with an average value for the 11th level). These findings demonstrate that the ubiquitous small R -value behavior persists also for the strongly anharmonic region of the Morse oscillator while also sampling the very nonlinear domain of dipole moment. The latter characteristics are quite distinct from the linear forced harmonic oscillator where the analytical results in Eq. (25) show that $R \leq \sqrt{2} \approx 1.4142$, but the distribution of R in all the simulations was found to be over a small range (i.e., less than 2).

4.2. Extrema of the Path Length Ratio Metric R

Although physically the lower bound of R is 1.0, its practically attainable lower limit in any particular physical system calls for

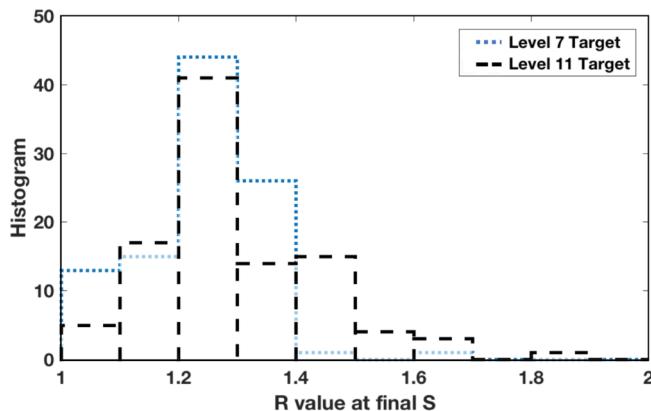


Fig. 5. Histogram of the combined final R values over 100 randomly chosen initial fields for target phase states having higher energies corresponding to the 7th quantum energy level (dotted lines) and the 11th quantum energy level (the dashed lines) in the very anharmonic region and the strong nonlinear dipole region of the Morse oscillator. Each simulation was stopped when $J \geq -0.01$. The averaged final R values, respectively, are $\langle R \rangle \approx 1.2455$ for the 7th level and $\langle R \rangle \approx 1.3003$ for the 11th level. All the collective results in Figs. 1, 2, 4, and 5 under quite varied physical situations show characteristically small R values..

numerical assessment. Furthermore, the highest value of R upon a gradient climb of the landscape is an important indication of the degree of confounding structural features that may be encountered. Specifically, given the generally low R values for all cases found in the preceding section, we now attempt to *maximize* and *minimize* R for the case of the 4th quantum energy level of the Morse oscillator to assess its extreme values using a stochastic particle swarm optimization (PSO) method [15,31]. Since the gradient algorithm is used for evolving the field, the value of R for a given simulation run is uniquely determined by the initial control field. To this end, we use the PSO algorithm to search over initial fields that either minimize or maximize the R value.

To facilitate the PSO algorithm, we take the optimization variables (i.e., those searched over to seek either a minimum or maximum value of R) to be the respective amplitude and phase parameters for the initial field, which were arbitrarily drawn from probability distributions in the simulations as in the previous section. Since the initial field was the sum of three sine functions, there are then 6 optimization variables, i.e., 3 amplitudes and 3 phases. Throughout the particle swarm optimizations, we took the initial point to be $(q(0), p(0)) = (1.7, -0.1)$ and the target state to be $(q^{\text{tar}}, p^{\text{tar}}) = (1.8205, 0.5)$. Moreover, the particle swarm was taken to have 100 initial members with each member initially chosen so that its cost satisfied $J \in [-25, -9]$; however, subsequent fields found by the particle swarm algorithm could take any initial cost, as dictated by the algorithm. Specifically, the following three steps were implemented to assess the minimum and maximum attainable values of R : (1) Select random amplitude and phase parameters for the initial field $\epsilon(0, \cdot)$, (2) compute the final $R(S)$ after a gradient optimization to reach the optimal control field where $J \geq -0.01$, and (3) evolve, using the PSO algorithm, the initial control field's amplitude and phase parameters so as to optimize the final R value.

The maximum value of $R_{\max} = 1.9$ was found essentially coinciding with the value already found in Fig. 4, while the minimum was found to be $R_{\min} = 1.0001$ (despite starting from a relatively large initial cost J , e.g., the field considered for illustrating the straight line trajectory in Fig. 6 corresponds to such a case), which is nearly three order of magnitude smaller than the lowest R_{\min} values ≈ 1.075 shown in Fig. 4. These findings are particularly important since they indicate that (1) the control landscape structure is generally smooth and (2) some regions of control landscape, followed by a gradient climb path, are especially featureless and devoid of gnarled structure. Nonetheless, the

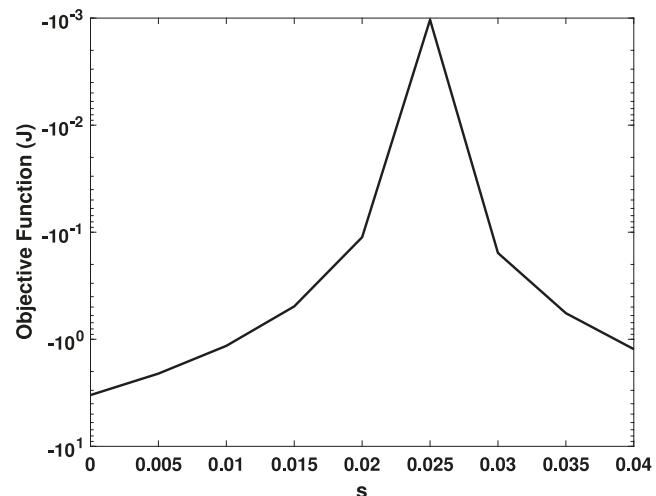


Fig. 6. Achieved J as we just follow the initial gradient $\delta J / \delta \epsilon(s = 0, t)$, Eq. (32), for a field corresponding to the minimal value $R = 1.0001$ obtained in the PSO simulations of Section 4.2 for the case corresponding to the quantum mechanical 4th energy level. The figure shows that marching in the initially identified gradient direction proceeds, from a relatively large initial cost J , directly to the top of the landscape at $s \approx 0.025$ producing $J = -10^{-3}$.

two extreme R generating fields are distinct (not shown), but no dramatic features are present to indicate that either one would give an extreme R value.

4.3. Straight paths through control space

To access the physical nature of a straight line trajectory, cf. Eq. (8), we consider the converse of considering a small R value found *post facto* (i.e., for the case corresponding to the 4th quantum energy level) from a gradient-based optimization of a control field, as illustrated in the previous sections. The relatively low values of R found in all of the simulations presented in both Figs. 2, 4 and 5 indicate that near-straight controls are prevalent in landscape climbs. The extreme limit of this behavior is to just follow the *initial* gradient to assess how far up the landscape the climb reaches. Thus the field $\epsilon(s, t)$, $0 \leq s \leq S$, is given by

$$\epsilon(s, t) = \left(\frac{\delta J}{\delta \epsilon(s = 0, t)} \right) s + \epsilon(0, t) \quad (32)$$

Naturally, the success of this procedure is expected to rely on the appropriate R value being sufficiently small for the field trajectory $\epsilon(0, t) \rightarrow \epsilon(S, t)$. Fig. 6 shows the evolution of J with Eq. (32) using the prior determined control field giving the minimal value $R = 1.0001$, obtained in our simulations by combining D-MORPH and PSO algorithms, see Section 4.2. The objective (i.e., the value of J) initially rises to reach a maximum value of $J = -0.001$ at $\sim s = 0.025$ as shown in Fig. 6. This result demonstrates the existence of an initial field whose cost function gradient points directly to a high value optimal field; many such cases were found in the analogous quantum mechanical studies [15–18]. However, no procedure exists at the present time to *a priori* identify such ideally behaving fields. Notwithstanding, the distribution of R values for all of the D-MORPH gradient optimizations in Figs. 2 and 4 are indicative of nearly straight line trajectories.

5. Conclusion

In this work, we examine the structure of classical optimal control landscapes. The classical molecular trajectories follow Hamiltonian

dynamics and the objective is to steer an initial phase space point to given target point. We assess the structure of the control landscape by examining the path along which the control field evolves towards an optimal field, guided by a steepest ascent gradient algorithm. In particular, we find that the path through control space (i.e., from the initial to optimal field) in many cases is comparable to a straight-line path. For a one-dimensional harmonic oscillator, we show that the length of the control field's path is *at most* a factor of $\sqrt{2}$ longer than the straight-line Euclidean distance between the initial and final fields; numerical investigations were fully consistent with the theoretical upper bound of $\sqrt{2}$ by only coming near to that value. Numerical results for a Morse oscillator further support the near straightness of gradient algorithm generated control paths for phase space target cases corresponding to the 4th, 7th and 11th quantum energy levels. This result indicates that the control landscape is relatively smooth and featureless, and thus gradient algorithms follow efficient, non-convoluted paths when searching for optimal control fields. The same qualitative observation was found prevalent in more extensive quantum mechanical simulations [15–18], leading to the tentative conclusion that the indicated landscape behavior remarked above is universal for closed quantum or classical systems under control.

While our results show that control fields evolve along nearly-straight gradient-based paths for sample classical systems, we do not have an explanation of *why* the landscape is surprisingly devoid of features. However, an explanation was provided in the quantum case [15–18] involving a multiple of interfering pathways that does not appear to have an analog in the classical regime. A more encompassing classical and quantum explanation may exist as a goal of future research to discover. It is possible that more complex classical systems, in particular, in a higher-dimensional phase space, could exhibit longer control paths. For example, it is straightforward to derive the bound $R \leq \sqrt{2n}$ for the n -dimensional harmonic oscillator and it may be possible to generalize these bounds to anharmonic oscillators. However, despite the upper bound, what actually counts is the statistical distribution of R values, and especially whether they are biased to lower values. Finally, an interesting direction for future research would be to take advantage of the collective results to design more efficient search algorithms to discover optimal control fields. We plan to explore these questions in future work.

Acknowledgments

We gratefully acknowledge useful discussions with Arun Nanduri. C.J.-W. acknowledges support from the Princeton Program in Plasma Science and Technology and NDSEG Fellowship. H.R. acknowledges partial support from the ARO (Grant No. W911NF-16-1-0014), while T.-S.H. acknowledges partial support from the DOE (Grant No. DE-FG02-02ER15344).

References

- [1] Q. Sun, I. Pelzner, G. Riviello, R.-B. Wu, H. Rabitz, *Phys. Rev. A* 89 (2014) 033413 .
- [2] C. Brif, R. Chakrabarti, H. Rabitz, *New J. Phys.* 12 (2010) 075008 .
- [3] M. Hsieh, H. Rabitz, *Phys. Rev. A* 77 (2008) 042306 .
- [4] H. Rabitz, T.-S. Ho, M. Hsieh, R. Kosut, M. Demirpal, *Phys. Rev. A* 74 (2006) 12721 .
- [5] H. Rabitz, M. Hsieh, C. Rosenthal, *Science* 303 (2004) 1998 .
- [6] H. Rabitz, M. Hsieh, C. Rosenthal, *J. Chem. Phys.* 124 (2006) 204107 .
- [7] R. Wu, H. Rabitz, M. Hsieh, *J. Phys. A: Math. Theor.* 41 (2008) 015006 .
- [8] T.-S. Ho, H. Rabitz, *J. Photochem. Photobiol. A: Chem.* 180 (2006) 226 .
- [9] C. Joe-Wong, T.-S. Ho, R. Long, H. Rabitz, R. Wu, *J. Chem. Phys.* 138 (2013) 124114 .
- [10] C. Joe-Wong, T.-S. Ho, H. Rabitz, R. Wu, *J. Chem. Phys.* 142 (2015) 154115 .
- [11] M.E. Tuckerman, G.J. Martyna, *J. Phys. Chem. B* 104 (2000) 159 .
- [12] M. Karplus, J.A. McCammon, *Nature Struct. Mol. Biol.* 9 (2002) 646 .
- [13] T.D. Sellwell, D.L. Thompson, *Int. J. Mod. Phys. B* 11 (1997) 1067 .
- [14] C. Joe-Wong, T.-S. Ho, H. Rabitz, *J. Math. Chem.* 1 (2015) .
- [15] A. Nanduri, A. Donovan, T.-S. Ho, H. Rabitz, *Phys. Rev. A* 88 (2013) 033425 .
- [16] A. Nanduri, A. Donovan, T.-S. Ho, H. Rabitz, *Phys. Rev. A* 90 (2014) 013406 .
- [17] A. Nanduri, O.M. Shir, A. Donovan, T.-S. Ho, H. Rabitz, *PCCP* 17 (2015) 334 .
- [18] A. Nanduri, T.-S. Ho, H. Rabitz, *Phys. Rev. A* 93 (2016) 023427 .
- [19] B. Russell, H. Rabitz, R.-B. Wu, *J. Phys. A: Math. Theor.* 50 (2017) 205302 .
- [20] B. Russell, S. Vuglar, H. Rabitz, *J. Phys. A: Math. Theor.* 51 (2018) 335103 .
- [21] E.F. de Lima, M.A.M. de Aguiar, *Phys. Rev. A* 77 (2008) 33406 .
- [22] C.D. Schwiegers, H. Rabitz, *Phys. Rev. A* 44 (1991) 5224 .
- [23] C.D. Schwiegers, H. Rabitz, *Phys. Rev. A* 48 (1993) 2549 .
- [24] R.B. Walker, R.K. Preston, *J. Chem. Phys.* 67 (1977) 2017 .
- [25] J. Botina, H. Rabitz, N. Rahman, *Phys. Rev. A* 51 (1995) 923 .
- [26] Y. Lai, M. Ding, C. Grebogi, *Phys. Rev. E* 47 (1993) 86 .
- [27] C. Cai, Z. Xu, W. Xu, *Automatica* 38 (2002) 1927 .
- [28] E. Ott, C. Grebogi, J. Yorke, *Phys. Rev. Lett.* 64 (1990) 1196 .
- [29] A. Rothman, T.-S. Ho, H. Rabitz, *J. Chem. Phys.* 123 (2005) 134104 .
- [30] S. Shi, H. Rabitz, *J. Chem. Phys.* 92 (1990) 2927 .
- [31] J. Kennedy, Particle swarm optimization, in: C. Sammut, G.I. Webb (Eds.), *Encyclopedia of Machine Learning*, Springer US, Boston, MA, 2010, pp. 760–766 .
- [32] A. Efimov, A. Fradkov, and A. Krivtsov, Proceedings of the European Control Conference Cambridge, UK (2003) pp. 1–4.
- [33] A. Guldberg, G.D. Billing, *Chem. Phys. Lett.* 186 (1991) 229 .
- [34] J.R. Stine, D.W. Noid, *Opt. Commun.* 31 (1979) 161 .