

Exploring quasi-linear control trajectories in quantum optimal control

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

September 27, 2012

Abstract

Quantum control seeks to manipulate chemical phenomena through the maximization of a target observable. It achieves this by casting the relationship between the observable and the control variables in terms of a functional landscape, which is ‘climbed’ as a result of the optimization process. Previous works have characterized the controllability and topology of the quantum control landscape, and have shown that they are simple in nature. In the current work we examine the structure of the quantum control landscape, encoded in the number R which measures how far an optimization path on the landscape deviates from a straight line. Through numerical simulations we find that R takes values close to 1 for many paths on the landscape. We also present analytic results, including bounds on R and the condition under which R must equal 1. A stochastic algorithm is employed to minimize R , and a new, less expensive algorithm for optimizing control fields that can be used in the laboratory when R is low is presented.

1 Introduction

A number of recent events, most notably the advances in femtosecond laser pulse shaping technology [1] and the improvement of closed loop learning algorithms in the laboratory [2], have allowed quantum control to find success in both simulations and in the laboratory. Quantum control has been used to selectively break interatomic bonds in molecules [3], control electron flow in aromatic molecules and electron states in semiconductors [4],[5], and manipulate energy inside proteins [6].

In a typical quantum control experiment, a quantum system of interest interacts with an applied electromagnetic field, whose control variables (e.g., phases and/or amplitudes) are manipulated in

an effort to achieve high fidelity of a specified physical observable. The dynamics of a quantum system interacting with an external electromagnetic field are described by the time-dependent Schrödinger equation,

$$i\hbar \frac{\partial U(t; t_0)}{\partial t} = H(t)U(t; t_0) \quad U(t_0; t_0) = \mathbb{1}, \quad (1)$$

where $H(t)$ is the system's Hamiltonian and $U(t; t_0)$ describes the unitary evolution propagator. Using the dipole approximation, the Hamiltonian can be written as

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

where H_0 is a field-free Hamiltonian, μ describes allowed dipole transitions and $E(t)$ is an external electromagnetic field that operates over a finite time interval $[0, T]$. The search for a field $E(t)$ that produces a target physical observable value can be viewed as occurring on an underlying *quantum control landscape*, which is defined as the observable as a functional of the field's control variables. It has been proven that when these field controls are unconstrained and can assume any numerical value, the control landscape is free of suboptimal extrema [7, 8, 9, 10]; this explains in part why excellent and even 'optimal' control can be attained for a variety of physical observables, as it implies that good control can be achieved using a local search algorithm initiating from any field. Optimizing the control variables to achieve a high target observable fidelity then corresponds to taking 'trajectories' on the quantum control landscape.

As a quantum control optimization proceeds, a curvilinear trajectory is taken in *control space*, which is the high-dimensional space of the control variables themselves (i.e., independent of the observable). Using a local search algorithm (e.g., gradient ascent) implies that these control trajectories are smoothly varying and can be parametrized by a diffeomorphic variable s ; the initial control field can then be represented by $E(s = 0, t)$ and the final control field that yields an acceptable degree of control by $E(s = s_{max}, t)$. Figure 1 illustrates the relationship between the quantum control landscape, control space, and the optimization trajectory, where a trajectory on the control landscape is projected onto control space (in two dimensions; generally the landscape and control space are of high dimensionality).

In recent experimental work [11], gradient-based optimizations were performed with the goal of maximizing second harmonic generation where the spectral phase acted as the control variable. To quantify the linearity of the trajectory followed by the controls during optimization, the ratio of the control trajectory's total path length to the Euclidean length was computed, and the results

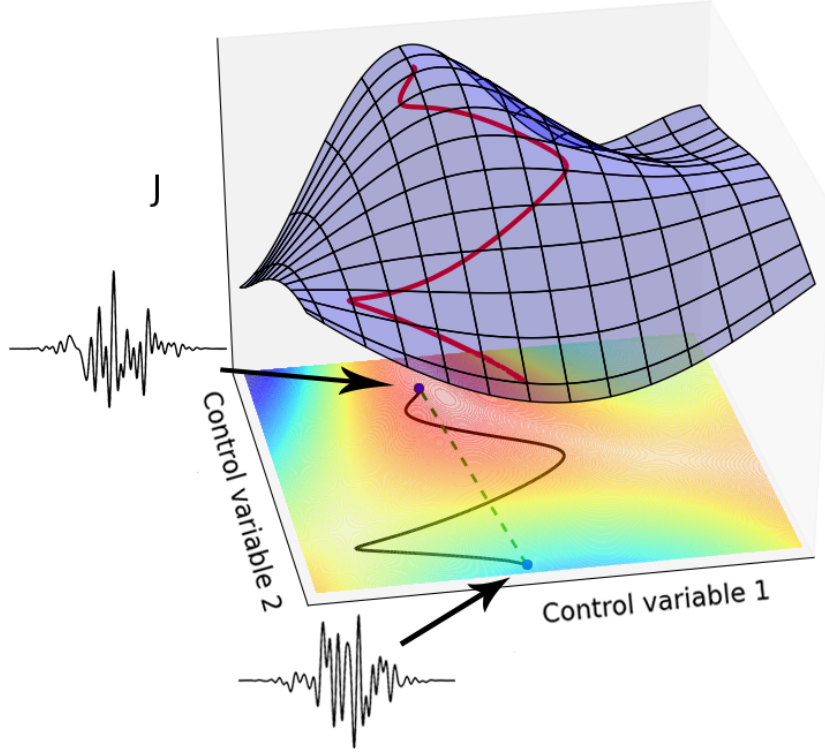


Figure 1: An example of a quantum control landscape as it relates to control space. The two horizontal axes represent two of the many control variables. The vertical axis represents the value of the target observable and the ‘height’ on the landscape, which is shown as the blue surface above the x-y plane. An optimization path, beginning at the ‘bottom’ of the landscape, is displayed in red; it takes the path of steepest ascent up the landscape, and reaches the optimal value of the target observable at the ‘top’ of the landscape. The path’s projection into control space is shown in black; the controls start at some point $E(s = 0, t)$ and end at $E(s = s_{max}, t)$. The black curve is nonlinear, as shown by a comparison to the straight green dashed line.

are reproduced in figure 2. Specifically, the path length between the initial and final control fields may be defined by

$$d_{PL} = M \int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s, t)}{\partial s} \right)^2 dt} ds \quad (3)$$

and the Euclidean distance by

$$d_{EL} = M \sqrt{\int_0^T (E(s_{max}, t) - E(0, t))^2 dt} \quad (4)$$

where $M = \sqrt{1/\tau}$ for τ points in the interval $[0, T]$. The ratio R is then defined by

$$R = \frac{d_{PL}}{d_{EL}} \quad (5)$$

The closer R is to unity, the ‘straighter’ or more linear the trajectory is. If R assumes higher values, then the controls must follow a more tortuous route through control space during optimization. From figure 2, the observed values of this ratio taken from experimental data are all of the same order of magnitude, and many are close to one. Other works [12, 13] have examined this ratio and also found near-unity behavior. This implies that, using a local search algorithm, a surprisingly straight path may be taken by the controls to achieve high observable fidelity. In general, both the quantum control landscape and control space consist of a large number of controls and are presumably highly complex in structure. However, the findings that gradient-based search techniques can yield near-straight control trajectories may imply that the structure of the landscape and/or control space is actually simpler than previously envisioned. The goal of the present work, which is motivated by the experimental findings of $R \approx 1$ [11], is to explore the nature of quasi-linear control trajectories. The observable under consideration in this work is a state-to-state transition probability, and a gradient-based algorithm is used to identify control fields that yield high transition probability.

The remainder of this work is structured as follows. After presenting the local search algorithms utilized to identify optimal control fields in section 2, section 3 provides an in-depth mathematical analysis of the ratio R and discusses practical upper and lower bounds. The form of a control field that would yield an identically ‘straight’ trajectory in control space where $R = 1$ is described. Section 4 introduces an additional algorithm that can be employed to find control fields that further linearize the control trajectory. Section 5 provides numerical illustrations that shows the distribution of R values after optimal control is achieved and discusses how these ratios can be further decreased by using the linearization algorithm from section 4. Section 6 provides concluding remarks.




Figure 2: Reproduction of results from Roslund *et al* [11]. The scatter plot shows the total length of an optimization trajectory, d_{PL} versus the straight line distance between the endpoints of the optimization trajectory, d_{EL} for 4804 experiments. The inset is a histogram of the ratio $R = d_{PL}/d_{EL}$.

2 Traversing the quantum control landscape

The quantum systems used in the present work are of finite dimension N , that is, H_0 and μ (c.f., Eq. (2)) are treated as $N \times N$ diagonal and complex Hermitian matrices, respectively. The observable under consideration is a state-to-state transition probability, denoted $P_{i \rightarrow f}$ and defined as

$$P_{i \rightarrow f} = |\langle f | U(T; 0) | i \rangle|^2, \quad (6)$$

where $P_{i \rightarrow f}$ is a function of the unitary propagator U (c.f., Eq. (1)) at time T , and the states $|i\rangle$, $|f\rangle$ are treated as eigenstates of H_0 .

A search for an optimal control field, or one that yields $P_{i \rightarrow f} = 1.0$ to specified numerical precision, entails following trajectories in control space and on the control landscape, where in the latter case the path begins near the ‘bottom’ of the landscape (an arbitrary field will generally yield $P_{i \rightarrow f} \ll 1$) and end at the ‘top’ of the landscape. Increasing $P_{i \rightarrow f}$ may be viewed as

‘climbing’ the associated transition probability control landscape. To ensure smooth movement during this landscape climb, a diffeomorphic parameter s is introduced, which ensures that a continuous trajectory is followed by the controls and thus permits a meaningful analysis of the trajectory’s linearity. Monotonically increasing $P_{i \rightarrow f}$ then requires that

$$\frac{dP_{i \rightarrow f}}{ds} = \int_0^T \frac{\delta P_{i \rightarrow f}}{\delta E(s, t)} \frac{\partial E(s, t)}{\partial s} dt \geq 0 \quad (7)$$

over all s , which can be satisfied by choosing

$$\frac{\partial E(s, t)}{\partial s} = \gamma(s) \frac{\delta P_{i \rightarrow f}}{\delta E(s, t)} \quad (8)$$

for $\gamma(s) \geq 0$ for all s . This choice for $\frac{\partial E(s, t)}{\partial s}$ is an implementation of the previously developed first-order DMORPH (diffeomorphic modulation under observable preserving homotopy) algorithm [14, 15]. In this work, the control variables are treated as the field amplitude at each discretized time point. Changes in $E(t)$ will continue until $dP_{i \rightarrow f}/ds = 0$, which occurs at a landscape critical point; for unconstrained controls, the only critical points exist at the top and bottom of the landscape, where $\delta P_{i \rightarrow f}/\delta E(s, t) = 0 \forall t$ [10]. For clarity, the gradient $\delta P_{i \rightarrow f}/\delta E(t)$ can be written as [10]

$$\frac{\delta P_{i \rightarrow f}}{\delta E(t)} = -\frac{2}{\hbar} \Im \left\{ \left\langle i \left| U^\dagger(T; 0) \right| f \right\rangle \left\langle f \left| U(T; 0) U^\dagger(t; 0) \mu U(t; 0) \right| i \right\rangle \right\}. \quad (9)$$

An analytical gradient-based search is one of many techniques that can be employed to find optimal control fields. In the laboratory, gradient-based optimization methods may be used [11] but are generally expensive, and evolutionary (stochastic) algorithms are often used instead [3, 16].

In lieu of computing the gradient at each obtained control field, as is done using DMORPH, an alternative optimization method is to compute the gradient only once at the initial control field and then follow a ‘straight shot’ in control space by forcing changes in $E(t)$ to assume the linear form

$$E(u, t) = G_0 u + E(0, t), \quad (10)$$

where u indexes the optimization step, $E(0, t)$ is the initial control field, and $G_0 = \delta P_{i \rightarrow f}/\delta E(0, t)$. Unlike the DMORPH method, which terminates at $dP_{i \rightarrow f}/ds = 0$, there is no guarantee of complete maximization of $P_{i \rightarrow f}$ using this algorithm. Instead, the ‘optimal’ field is considered to correspond to the first encountered maximum of the computed $P_{i \rightarrow f}$. Section 5 will present numerical examples to illustrate the benefits of this technique and will compare the linearity of corresponding control trajectories.

3 A quantitative characterization of R

The linearity measure of a control space trajectory considered in this work is the ratio R of path length to Euclidean distance, defined as

$$R = \frac{\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} ds}{\sqrt{\int_0^T [E(s_{max}, t) - E(0, t)]^2 dt}}. \quad (11)$$

Nominally, the minimum value for R is unity, which implies a straight control trajectory. As a means to understand the structure of control fields that would yield $R = 1$, the following analysis formally derives this lower bound. The Euclidean distance can be rewritten as

$$d_{EL} = \sqrt{\int_0^T \left(\int_0^{s_{max}} \frac{\partial E(s,t)}{\partial s} ds \right)^2 dt} \quad (12)$$

$$= \sqrt{\int_0^{s_{max}} \int_0^{s_{max}} \int_0^T \frac{\partial E(s,t)}{\partial s} \frac{\partial E(s',t)}{\partial s'} dt ds ds'}. \quad (13)$$

Using the Cauchy-Schwartz inequality,

$$\int_0^T \frac{\partial E(s,t)}{\partial s} \frac{\partial E(s',t)}{\partial s'} dt \leq \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} \sqrt{\int_0^T \left(\frac{\partial E(s',t)}{\partial s'} \right)^2 dt}. \quad (14)$$

This implies that

$$d_{EL} \leq \sqrt{\int_0^{s_{max}} \int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} \sqrt{\int_0^T \left(\frac{\partial E(s',t)}{\partial s'} \right)^2 dt} ds ds'} \quad (15)$$

$$= \sqrt{\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} \int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s',t)}{\partial s'} \right)^2 dt} ds ds'} \quad (16)$$

$$= \sqrt{\left(\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} \right)^2} \quad (17)$$

$$= \int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} ds. \quad (18)$$

Thus,

$$R \geq \frac{\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} ds}{\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s,t)}{\partial s} \right)^2 dt} ds} = 1. \quad (19)$$

When the inequality in Eq. (19) becomes saturated, the trajectory in control space connecting the initial and optimal control fields constitutes a straight line. Recall that the use of DMORPH as described in section 2 (where $\gamma(s) = 1$) implies that $\partial E(s, t)/\partial s = \delta P_{i \rightarrow f}/\delta E(s, t)$. Using this, the following is posited:

Theorem 1. It is sufficient and necessary that the gradient function is separable, i.e.

$$\frac{\partial E(s, t)}{\partial s} = \frac{\delta P_{i \rightarrow f}}{\delta E(s, t)} = \alpha(s) \times \beta(t) \quad (20)$$

in order to achieve

$$R = 1. \quad (21)$$

Proof. To prove the sufficient condition, Eq. (20) can be substituted into Eq. (12), which may then be used in an expression for R :

$$\begin{aligned} R &= \frac{\int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\partial E(s, t)}{\partial s}\right)^2 dt ds}}{\sqrt{\int_0^T \left(\int_0^{s_{max}} \frac{\partial E(s, t)}{\partial s} ds\right)^2 dt}} \\ &= \frac{\int_0^{s_{max}} \sqrt{\int_0^T \alpha^2(s) \times \beta^2(t) dt ds}}{\sqrt{\int_0^T \left(\int_0^{s_{max}} \alpha(s) \times \beta(t) ds\right)^2 dt}} \\ &= \frac{\int_0^{s_{max}} \alpha(s) ds \sqrt{\int_0^T \beta^2(t) dt}}{\sqrt{\int_0^T \beta^2(t) dt} \int_0^{s_{max}} \alpha(s) ds} = 1. \end{aligned} \quad (22)$$

To prove the necessary condition, assume that $R = 1$. This means the the length of the trajectory in control space is the same as the distance of the straight line between $E(0, t)$ and $E(s_{max}, t)$. This in turn implies that the trajectory is a straight line between those two points, because that curve minimizes the distance between the endpoints.

For a straight line, the gradient, or slope, cannot depend on s . We must have that

$$E(s, t) = A(s) \times \beta(t) + E(0, t). \quad (23)$$

This is simply the infinite dimensional analogue of the slope-intercept form of the equation for a straight line, with $\beta(t)$ the analogue of the slope and $E(0, t)$ the analogue of the y-intercept. The function $A(s)$, the antiderivative of $\alpha(s)$, dictates how quickly the line is traversed. It must be monotonic increasing, which means that $\alpha(s)$ must be nonnegative. Taking the derivative with respect to s , we have equation 20. \square

To obtain an upper bound for R , d_{PL} must be maximized and d_{EL} must be minimized. If the DMORPH algorithm is employed, then the expression for d_{PL} can be written in terms of $\delta P_{i \rightarrow f} / \delta E(s, t)$, where

$$d_{PL} = \int_0^{s_{max}} \sqrt{\int_0^T \left(\frac{\delta P_{i \rightarrow f}}{\delta E(s, t)} \right)^2 dt} ds \quad (24)$$

$$= \int_0^{s_{max}} \left\| \frac{\delta P_{i \rightarrow f}}{\delta E(s, t)} \right\| ds \quad (25)$$

$$\leq \int_0^{s_{max}} \frac{2}{\hbar} \|\mu\| ds \quad (26)$$

$$= \frac{2}{\hbar} s_{max} \|\mu\| \quad (27)$$

where $\|\mu\|$ represents the norm of the transition dipole matrix μ [16]. Thus, the term d_{PL} is bounded from above by the strength of the dipole matrix and the value of s_{max} . In principle, achieving increasingly high values (arbitrarily close to unity) for $P_{i \rightarrow f}$ requires $s_{max} \rightarrow \infty$. In practice, however, s_{max} will be of some finite value that depends on the demanded precision of $P_{i \rightarrow f} = 1.0$. This implies that for suitably high degrees of control, d_{PL} will have a practical upper bound that principally depends on the strength of the dipole, which supports the evidence that R is often found to be surprisingly close to unity. Identifying a lower bound for d_{EL} remains an open challenge that will be addressed in future work [17].

4 Searching for fields to minimize R

In searching for control fields that produce a high state-to-state transition probability, the optimization methods described in section 2 (DMORPH and straight shot) work independently of R ; that is, there is no attempt to extremize R during a typical landscape climb. In order to extremize R itself, we utilized an evolutionary algorithm, namely a particle swarm optimization (PSO), to search for initial control fields that, upon gradient-based $P_{i \rightarrow f}$ optimization, produce low R values. The details of the algorithm are presented in an appendix, but in this section we will explain how we used the PSO algorithm to minimize R . The main goal is to determine whether these stochastically identified initial fields produce more or less linear control trajectories compared to ‘randomly’ chosen initial control fields.

Each field $E_k(t)$ in a family of K trial control fields is normalized to start at a predetermined

initial $P_{i \rightarrow f}^0$ value. If $P_{i \rightarrow f}[E_k(t)] < P_{i \rightarrow f}^0$, then the DMORPH algorithm can be used to drive the initial field to $P_{i \rightarrow f}^0$ where $\gamma(s) = 1$ (c.f., Eq. (8)); otherwise, if $P_{i \rightarrow f}[E_k(t)] > P_{i \rightarrow f}^0$, $\gamma(s)$ can be set to -1 to decrease $P_{i \rightarrow f}$ to $P_{i \rightarrow f}^0$. Once each initial field is identified, it undergoes a DMORPH optimization to achieve $P_{i \rightarrow f} = 1.0$ to specified numerical precision. The corresponding R values for all K optimization trajectories are recorded. The initial field that produced the lowest R value upon DMORPH optimization, denoted by $E_{best}^{swarm}(t)$, is then treated as the ‘best’ field and is used to generate a new set of K trial fields. These new K fields are normalized to $P_{i \rightarrow f}^0$ and then optimized with respect to $P_{i \rightarrow f}$; if a new ‘best’ field is found, E_{best}^{swarm} is updated. The collection of these best fields are used to again generate a new set of trial control fields. The algorithm continues for a predetermined number of generations and is not structured to drive fields to obtain a particular R fitness value (whose practical minimum is not known *a priori*). While the DMORPH optimizations utilize the field value at each time point as the set of control variables, the PSO algorithm (that is, the search for new trial control fields) uses the amplitudes of the fields’ frequency components as the control variables. These comprises a smaller set of controls and place a practical limit on the frequencies accessible to the optimization algorithm. In section 5, the R values obtained through DMORPH optimization using randomly chosen initial fields and those identified using the PSO algorithm will be compared through a statistical analysis.

5 Numerical Illustrations

The following numerical illustrations first explore the behavior of R for state-to-state transition probability optimizations performed using a local gradient-based search algorithm. The necessary structure of the gradient from an identically straight trajectory is examined. Additionally, the ‘straight shot’ algorithm that was introduced in section 2 is compared to a complete gradient-based optimization approach. A stochastic algorithm is used to search for initial control fields that, upon gradient-based optimization, yield near-unity values for R ; the R values of these associated control trajectories are compared to those obtained from gradient-based optimizations that make no effort to linearize their respective control trajectories. The illustrations use elements of an applied electromagnetic field as control variables and consider a quantum system’s Hamiltonian (i.e., internal energies and allowed dipole transitions) to be fixed. For clarity of presentation, the Hamiltonians of the quantum systems utilized are of low dimension, though higher dimensional

systems may also be considered without loss of generality.

5.1 Behavior of R with DMORPH-assisted observable optimization

As a means to understand the complexity of trajectories followed by control variables during an observable optimization, 2,000 individual state-to-state transition probability optimizations were performed and the resulting R ratio values were computed and collected. The optimizations used a fixed five-dimensional Hamiltonian (c.f., Eq. (2)), where

$$H_0 = \begin{pmatrix} -10 & 0 & 0 & 0 & 0 \\ 0 & -7 & 0 & 0 & 0 \\ 0 & 0 & -3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 8 \end{pmatrix}, \quad (28)$$

and

$$\mu = \begin{pmatrix} 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 & \pm 0.5^3 \\ \pm 1 & 0 & \pm 1 & \pm 0.5 & \pm 0.5^2 \\ \pm 0.5 & \pm 1 & 0 & \pm 1 & \pm 0.5 \\ \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 & \pm 1 \\ \pm 0.5^3 & \pm 0.5^2 & \pm 0.5 & \pm 1 & 0 \end{pmatrix} \quad (29)$$

where the signs of the dipole matrix elements were chosen randomly under the constraint that μ remained symmetric. The target observable was $P_{1 \rightarrow N}$, where $N = 5$. Earlier work has shown that larger or smaller values of the system dimension N does not have a significant effect on the values of R observed [13, 12]. The applied control field was parametrized in the form

$$E(t) = \frac{1}{F} \sum_{\omega=1}^{20} \exp[-0.3(t - \frac{T}{2})^2] a_{\omega} \sin(\omega t + \phi_{\omega}), \quad (30)$$

where $T = 10$ and the interval $[0, T] = [0, 10]$ was discretized into 1,001 time points. The field frequency components, ω , are treated as integers from 1 to 20, which permits substantial resonance with the system's internal energy level spacings described by H_0 . The amplitudes a_{ω} were chosen from a uniform distribution between 0 and 1, and the phases ϕ_{ω} were chosen from the interval $[0, 2\pi]$. F equals the maximum of the field's absolute value and is included to limit the strength of the field. In order to 'normalize' the starting point of the $P_{1 \rightarrow 5}$ optimizations, the fields were

first optimized to reach a starting observable value of $P_{1 \rightarrow 5}^0 = 0.01$. This was done with DMORPH where $\gamma(s) = \pm 1$ (c.f., Eq. (8)) for $P_{1 \rightarrow 5} < 0.01$ and $P_{1 \rightarrow 5} > 0.01$, respectively. The resulting field was then used as an initial field for the main optimization, where DMORPH was used with $\gamma(s) = 1$ to push the target observable to $P_{1 \rightarrow 5} = 0.99$. The control variables were treated as the field values at each of the 1,001 time points, though other controls (i.e., phases and/or amplitudes) can be considered as well.

Figure 3 shows the behavior of the computed R values from the 2,000 simulations, where a histogram distribution of $\log_{10}(R - 1)$ is shown for clarity of presentation. The inset plot shows the distribution of $R - 1$. Out of all simulations, the smallest R obtained was 1.02 and the largest was 1.26; the mean was 1.12 and the mean of the skewed $\log_{10}(R - 1)$ distribution was -0.954. As seen by the positively skewed $R - 1$ distribution, most of the R values lie closer to unity. This implies that in general, the control variables do not need to take a highly gnarled trajectory in control space to achieve excellent control.

The dipole matrix μ that was used in these simulations (whose structure provides a meaningful description of many physical systems [8]) permitted the control field to access any $|j\rangle \rightarrow |k\rangle, j \neq k$ state-to-state transition, which in principle makes a target state-to-state transition easier to achieve by providing more accessible intermediate transition pathways. To determine if the behavior of R is strongly dependent upon the allowed state transitions, the two following dipole matrices were also used in subsequent DMORPH optimizations, where a highly restrictive dipole is described by

$$\mu = \begin{pmatrix} 0 & \pm 1 & 0 & 0 & 0 \\ \pm 1 & 0 & \pm 1 & 0 & 0 \\ 0 & \pm 1 & 0 & \pm 1 & 0 \\ 0 & 0 & \pm 1 & 0 & \pm 1 \\ 0 & 0 & 0 & \pm 1 & 0 \end{pmatrix} \quad (31)$$

and a less restrictive one by

$$\mu = \begin{pmatrix} 0 & \pm 1 & \pm 1 & \pm 1 & \pm 1 \\ \pm 1 & 0 & \pm 1 & \pm 1 & \pm 1 \\ \pm 1 & \pm 1 & 0 & \pm 1 & \pm 1 \\ \pm 1 & \pm 1 & \pm 1 & 0 & \pm 1 \\ \pm 1 & \pm 1 & \pm 1 & \pm 1 & 0 \end{pmatrix} \quad (32)$$

where the same restriction on the random signs applies. While the structure of the matrix in Eq. (32) is the same as that in Eq. (29) in that any $|j\rangle \rightarrow |k\rangle$ transition is permitted, the strength of the off-diagonal elements in Eq. (32) is higher, particularly for the target $|1\rangle \rightarrow |5\rangle$ transition. The same procedure as before was used to climb the landscape and calculate R ; 1000 trials were performed for each of the additional dipole matrices. Using the dipole of Eq. (32), the minimum, mean, and maximum of the $R - 1$ values found were .005, .05, and .17. Using the dipole of Eq. (31), the minimum, mean, and maximum of the $R - 1$ values found were found to be 0.07, 0.24, and 0.64, respectively. The effect of the dipole matrix on R is evident; a restricted dipole results in control fields being harder to find and drives R up; a dipole matrix that allows more transitions results in control fields being easier to find and lets R become lower.

An important question to consider is whether control trajectories that result in a low R value correspond to fields that are ‘distributed’ in control space differently than those that result in a high R value. As a means to examine the ‘distance distribution’ of control fields, pairwise distances were computed for (1) all initial fields, (2) all final fields, and (3) all initial-final field pairs. From the 2,000 simulations described above that used the dipole of Eq.??, the distance distributions of the control trajectories from the 500 simulations with the lowest R values (mean $R = 1.071$) are shown in plot A of figure 4. The distances between initial fields (blue curve) is on average less and have a smaller deviation than those corresponding to the optimal fields (red curve), which implies that optimal fields are spread more widely across control space than initial fields. The yellow curve corresponds to the pairwise distances between all initial and optimal fields, and on average, the distances between initial and optimal fields is less than the distances between optimal fields. For comparison, the same distance distributions were computed for the 500 control trajectories associated with the highest R values (mean $R = 1.175$). The results are shown in plot B of figure 4 and show virtually indistinguishable behavior from the distributions in plot A. This implies that fields corresponding to low R trajectories are distributed in a similar manner to those corresponding to higher R trajectories; consequently, low R trajectories are as prevalent as relatively high R trajectories, which helps explain why near-straight control trajectories do not appear difficult to find [13, 11, 12].

In section 4, an upper bound for the path length d_{PL} was shown to depend on the strength of the transition dipole matrix as well as the value for s_{max} . In principle, achieving perfect $P_{i \rightarrow f} = 1$

control would require $s_{max} \rightarrow \infty$. In this view, it is useful to consider the effect of increasing the required degree of control on R . In particular, three optimization scenarios were considered: (1) $P_{1 \rightarrow 5}^0 = 0.001 \rightarrow P_{1 \rightarrow 5} = 0.999$, (2) $P_{1 \rightarrow 5}^0 = 0.0001 \rightarrow P_{1 \rightarrow 5} = 0.9999$, and (3) $P_{1 \rightarrow 5}^0 = 0.00001 \rightarrow P_{1 \rightarrow 5} = 0.99999$. For each case, 1,000 DMORPH simulations were performed and the resulting R values were computed. Figure 5 shows the distribution of $R - 1$ values for the three scenarios. For case (1), the mean of the $R - 1$ values was 0.159. For case (2), the mean was 0.178, and for case (3), the mean was 0.184. This trend agrees with our expectation that demanding more extreme values for the start and end of a DMORPH optimization increases the difficulty of the search and causes R to go up.

Additionally, in order to examine the nature of a straight line optimization trajectory, we examined a climb with a very low R value of 1.0036, using the Hamiltonian of Eq. (28) and the dipole matrix of Eq. (29), and starting and ending points of $P_{1 \rightarrow 5} = 0.01$ and 0.99, respectively. The resulting optimization trajectory approximates a straight line very well. To verify the behavior of $E(s, t)$ predicted by Theorem 1, we plotted $\frac{\delta P_{i \rightarrow f}}{E(s, t)} = \frac{\partial E(s, t)}{\partial s}$ along this optimization trajectory in Figure 7. From the figure, it can be seen that $\frac{\partial E(s, t)}{\partial s}$ nearly does factor into a function $\beta(t)$ that is independent of s and a bell-shaped function $\alpha(s)$ that governs how quickly the climb proceeds. The empirical fact that $\alpha(s)$ peaks in the middle of the climb and becomes small at either end of the climb is in accord with observations that a gradient-based optimization will proceed much more quickly in the middle than at the beginning and end of the climb.

5.2 ‘Straight shot’ optimizations

In addition to the DMORPH gradient-based optimization method, a ‘straight shot’ method was introduced in section 2 as a means to reduce computational optimization effort by eliminating the time-consuming gradient calculation required at each optimization step and instead following the ‘initial’ gradient computed at $P_{i \rightarrow f}^0$. However, this simpler method will likely not produce the same final field compared to a full DMORPH simulation, and in fact complete $P_{i \rightarrow f}$ optimization is not guaranteed. For the 2,000 DMORPH simulations performed using the dipole matrix of Eq. (29), the initial fields (i.e., where $P_{1 \rightarrow 5} = 0.01$) were treated as initial conditions for straight shot optimizations. Intuitively, using an initial field from a DMORPH trajectory that produces an R value very close to unity should allow the straight shot method to provide a reasonable

approximation of the maximum observable value and corresponding control field since it should replicate a near-straight trajectory; by contrast, using an initial field that produces a high R value from a DMORPH optimization should yield a dramatically different maximum $P_{i \rightarrow f}$ with the straight shot approach. Plot A of figure 6 shows the $P_{1 \rightarrow 5}$ trajectories from three different straight shot simulations that used initial fields with low, middle and high R values obtained from the corresponding DMORPH optimization. The first maximum is treated as the optimal $P_{1 \rightarrow 5}$ value as a means to approximate a monotonic optimization. For the case where $R = 1.003$, the optimal $P_{1 \rightarrow 5}$ was 0.9888, which is an excellent result close to the $P_{1 \rightarrow 5} = 0.99$ obtained using DMORPH. With $R = 1.045$, the optimal $P_{1 \rightarrow 5}$ was 0.7940. Finally, the initial field corresponding to $R = 1.707$ yielded an optimal $P_{1 \rightarrow 5}$ value of 0.3697, far below the optimal yield obtained using DMORPH and reflective of the fact that the DMORPH trajectory did not produce a near-straight control trajectory. Plot B of figure 6 plots the optimal $P_{1 \rightarrow 5}$ values obtained from the straight shot method using the 2,000 initial fields from the DMORPH simulations versus $\log_{10}(R-1)$. An inverse relationship appears to be present, indicating that as R increases, the optimal $P_{i \rightarrow f}$ produced from the straight shot method value decreases.

If we are able to identify the conditions under which almost all optimization pathways on the quantum control landscape have sufficiently low values of R , then when working under these special conditions optimizing control fields will be much easier when using the algorithm presented here. One criterion to distinguish what ‘sufficiently low’ means is pathways with sufficiently low values of R will on average yield target observable values of $P_{i \rightarrow f} \approx 0.9$, for example.

The results of the straight shot optimization method can also act as a criterion by which ‘low’ R values may be distinguished from ‘high’ R values, even if both values are close to unity, as are most of the data presented in this paper. In such cases, a distinction between low and high R values may seem artificial and arbitrary, but the results of this section shows that there exists a tangible difference between high R and low R optimization pathways.

5.3 Optimizing R

Section 4 described an implementation of a particle swarm algorithm that can be used to search for control fields that lead to near-unity values of R . A set of 50 initial fields interacting with the Hamiltonian from Eq. (28), (29) were optimized using DMORPH from $P_{1 \rightarrow 5} = 0.01$ to $P_{1 \rightarrow 5} =$

0.99. From these 50 optimizations, the trajectory that yielded the lowest R value (computed after optimizations were completed) was treated as the ‘best’ field and was then implemented into a PSO algorithm to identify a new set of 50 initial fields through a variation of 20 amplitudes and 20 phases. The amplitudes were uniformly chosen between $[0,1]$ and the phases between $[0, 2\pi]$. The new set of 50 initial fields are then optimized using DMORPH (reverting back to using the time-point control variables), and the best control trajectory is identified *post facto*. The best control trajectory each individual particle has found is also recorded at this point, and is updated after every subsequent iteration. This procedure occurs over a predetermined 50 generations and results in a set of amplitudes and phases that yield optimized (minimized) R values. More details are available in the appendix.

Figure 8 shows the R (or more specifically, $R - 1$) values obtained from three control fields at each generation of the PSO algorithm. In particular, the $R - 1$ values from the fields that yield the best (lowest), worst (highest) and median value of R are plotted for each of the 50 generations. The figure shows a general downward trend in $R - 1$, indicating that as the optimization proceeds, it continues to find initial control fields that upon DMORPH optimization of $P_{1 \rightarrow 5}$ will produce straighter control trajectories. The various spikes (increases in $R - 1$) are a result of the stochastic nature of the PSO algorithm. After the 50 generations were completed, the lowest $R - 1$ values were found to be on the order of 10^{-4} , which is two orders of magnitude lower than the best field found from the 2,000 DMORPH simulations described above where $R_{smallest} - 1 = 0.02$.

While so far the ‘best’ fields have been those that produce the lowest R values, it is possible to amend the PSO algorithm to search for ‘best’ fields that have the highest R values. This can be viewed as a means to determine if there is a numerical upper bound on R . Using the same algorithm parameters, three separate PSO optimizations were performed. The resulting values for $R - 1$ were 0.702, 0.608, and 0.718, which are significantly larger than even the highest R found from the DMORPH simulations ($R_{largest} - 1 = 0.3$). The fact that the three individual stochastic optimizations resulted in similar R values hints at a practical upper bound on R .

6 Conclusions

The present work is the first to investigate in detail the near linearity of trajectories taken by control variables during quantum optimal control simulations. This work was motivated by results

from [11] that indicated a quasi-linear trajectory is often taken by control variables during quantum control experiments, as well as other works that found surprisingly linear control trajectory behavior in simulations [12, 13]. We considered the optimal control of a state-to-state transition probability and utilized control variables from an applied external electromagnetic field. To optimize these fields, we used a gradient-based algorithm; this ‘myopic’ search technique permits the consideration of a continuous trajectory that is assumed by the controls during the optimization. A formal derivation of the lower bound of a path length to Euclidean distance ratio R was presented and it was proved that control field optimization trajectories which yield identically straight ($R = 1$) paths must factor into separate functions of s and t .

The numerical illustrations presented in this work showed that in general, the R values calculated for control trajectories are surprisingly close to unity, and a stochastic search for initial control fields is capable of identifying trajectories that yield $R - 1$ values at least two orders of magnitude lower than randomly chosen initial fields. The distribution of low R (close to unity) and high R control trajectory fields are virtually indistinguishable, indicating that there is no preference for the controls to follow a highly gnarled trajectory to achieve optimal control and indeed supporting the findings of linear trajectories [11, 12, 13]. For quasi-linear trajectories, a ‘straight shot’ algorithm, which only requires one gradient calculation, shows good agreement with a more expensive gradient ascent algorithm and could possibly be used to greatly simplify laboratory experiments when operating in a regime where R is almost always sufficiently low.

In general, a quantum control experiment addresses hundreds of control variables in an attempt to extremize a control objective. The high dimensionality of the corresponding control landscape and control search space presumably induces an untenable complexity into the structure of these search spaces, so the identification of control trajectories that are almost linear was surprising and hints that in fact the landscape structure may be much simpler than previously thought. This paper has laid the foundation to further investigate the structure of the control landscape and control space. The fact that a precisely linear trajectory requires that the gradient $\delta P_{i \rightarrow f} / \delta E(s, t)$ be separable into an s -dependent and a t -dependent term is a surprising and intriguing result, since nominally the gradient is a highly complex quantity that, for example, depends upon multiplicative interactions between the unitary propagator $U(t, 0)$ and the transition dipole matrix μ (c.f., Eq. (9)). A future work [17] will further explore this result, in addition to developing a method to

perform a dual optimization that extremizes an observable while simultaneously extremizing R , and should provide novel insight as to the structure of the quantum control landscape.

A Mechanics of the particle swarm optimization algorithm

A particle swarm optimization (PSO) algorithm searches for an optimal solution using a stochastic approach that mimics an evolutionary strategy. In the present work, the PSO algorithm searches for an initial control field that, after a DMORPH-assisted $P_{i \rightarrow f}$ optimization, yields a near-unity value for the associated control trajectory's path length to Euclidean ratio R . The PSO algorithm begins by creating a set of K trial control fields (referred to as the 'swarm' of candidate solutions), where each field is referred to as a 'particle.' The PSO moves each particle in the swarm according to its velocity at the inception of each new generation. In the present work, $K = 50$, and the control variables that were accessible for the PSO were field amplitudes and phases from the domains $[0,1]$ and $[0,2\pi]$, respectively. Using a five-dimensional Hamiltonian (c.f., Eq. (28) and (29)), each of these trial fields was then optimized to an initial $P_{1 \rightarrow 5}^0 = 0.01$ using the time-point DMORPH optimization technique presented in section 2. The resulting 50 fields that yield $P_{1 \rightarrow 5} = 0.01$ are known as initial fields and are to be differentiated from the 50 trial fields that are chosen by the PSO. The initial fields are then optimized to reach $P_{1 \rightarrow 5} = 0.99$ again using the DMORPH optimization technique, and the path length to Euclidean distance ratios R can be computed for each of the 50 optimizations. The initial field that corresponds to the trajectory with the lowest computed R value is denoted by $E_{swarm}^{best,g}(t)$, where g represents the generation index. In the present work, 50 total generations were used. Each particle stores its own history (that is, each trial field associated with a particular particle is stored), and the corresponding initial field (the field obtained after the particle is normalized to $P_{1 \rightarrow 5}^0 = 0.01$) that yields the lowest R value over the total number of generations that have passed is also recorded and denoted as E_k^{best} . Using this information, the velocity of each particle is computed by using

$$v_k^g = C_0 v_k^{g-1} + C_1 R_1 (E_k^{best} - E_{k,g-1}) + C_2 R_2 (E_{swarm,g-1}^{best} - E_{k,g-1}), \quad (33)$$

where C_0 is the particle's inertia (defined below), C_1 is the cognitive attraction and is set to 0.5, C_2 is the social attraction and is set to 1.5, and R_1, R_2 are diagonal matrices with randomly chosen 0 and 1 entries that yield the stochastic nature of the PSO algorithm. The inertia C_0 is defined by

$$C_0 = 0.9 - \frac{g-1}{30} 0.2, \quad (34)$$

where an increase in generation g results in a decreased inertia. The new particles are determined through

$$E_k^g = E_k + v_k^g, \quad (35)$$

where the velocity incorporates the best fields that produced the lowest R values and updates the trial field locations accordingly.

It is important to note that the PSO algorithm is not guaranteed to find the global minimum of R because of the potential for the swarm to collapse around local minima; this occurs when each particle in the swarm converges on the same region in control space and hence the same R value, at the expense of exploring further minima in R that may be more optimal. We must remain cognizant of this phenomenon also because we do not know the character of the functional relationship between R and the trial controls; in particular, we do not know if there exist many local minima of R that may trap the PSO algorithm. We also do not know the numerical value of a lower bound on R , so we cannot be sure of the degree to which the minimal values of R we locate are optimal.

Acknowledgments

A.N. would like to acknowledge Vincent Beltrani for suggesting the use of the Particle Swarm Optimization algorithm. This work was supported by the Program in Plasma Science and Technology at Princeton University.

References

- [1] A. M. Weiner. Femtosecond pulse shaping using spatial light modulators. *Review of Scientific Instruments*, 71(5), 2000.
- [2] Richard S. Judson and Herschel Rabitz. Teaching lasers to control molecules. *Phys. Rev. Lett.*, 68:1500–1503, Mar 1992.
- [3] Herschel Rabitz, Regina de Vivie-Riedle, Marcus Motzkus, and Karl Kompa. Whither the future of controlling quantum phenomena? *Science*, 288(5467):824–828, 2000.

- [4] E. Räsänen, A. Castro, J. Werschnik, A. Rubio, and E. K. U. Gross. Optimal control of quantum rings by terahertz laser pulses. *Phys. Rev. Lett.*, 98:157404, Apr 2007.
- [5] M. Kanno, K. Hoki, H. Kono, and Y. Fujimura. Quantum optimal control of electron ring currents in chiral aromatic molecules. *Journal of Chemical Physics*, 127(20):204314, November 2007.
- [6] Jennifer L. Herek, Wendel Wohlleben, Richard J. Cogdell, Dirk Zeidler, and Marcus Motzkus. Quantum control of energy flow in light harvesting. *Nature*, 417, 2002.
- [7] Raj Chakrabarti and Herschel Rabitz. Quantum control landscapes. *International Reviews in Physical Chemistry*, 26(4), 2007.
- [8] Katharine W. Moore and Herschel Rabitz. Exploring quantum control landscapes: topology, features, and optimization scaling. *Phys. Rev. A*, 84:012109, Jul 2011.
- [9] Vincent Beltrani, Jason Dominy, Tak-San Ho, and Herschel Rabitz. Photonic reagent control of dynamically homologous quantum systems. *The Journal of Chemical Physics*, 126(9), 2007.
- [10] 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, Jul 2006.
- [11] Jonathan Roslund and Herschel Rabitz. Experimental quantum control landscapes: Inherent monotonicity and artificial structure. *Phys. Rev. A*, 80:013408, Jul 2009.
- [12] Katharine Moore, Michael Hsieh, and Herschel Rabitz. On the relationship between quantum control landscape structure and optimization complexity. *The Journal of Chemical Physics*, 128(15), 2008.
- [13] Ashley Donovan, Vincent Beltrani, and Herschel Rabitz. Quantum control by means of hamiltonian structure manipulation. *Phys. Chem. Chem. Phys.*, 13:7348–7362, 2011.
- [14] Adam Rothman, Tak-San Ho, and Herschel Rabitz. Observable-preserving control of quantum dynamics over a family of related systems. *Physical Review A*, 72(2), 2005.
- [15] Adam Rothman, Tak-San Ho, and Herschel Rabitz. Quantum observable homotopy tracking control. *The Journal of Chemical Physics*, 123(13), 2005.

- [16] Herschel A. Rabitz, Michael M. Hsieh, and Carey M. Rosenthal. Quantum optimally controlled transition landscapes. *Science*, 303(5666):1998–2001, 2004.
- [17] Arun Nanduri, Ashley Donovan, Tak-San Ho, and Hershel Rabitz. Second paper. 2012.

B Figures




Figure 3: Large figure: distribution of $\log_{10}(R-1)$ values for 2,000 $P_{1 \rightarrow 5}$ optimizations performed using the DMORPH technique from section 2. The five-dimensional Hamiltonian from Eq. (28) and (29) was treated as fixed. R values were computed for each optimization trajectory where $P_{1 \rightarrow 5}^0 = 0.01$ and $P_{1 \rightarrow 5} = 0.99$. The mean R from the 2,000 optimizations was 1.12. The inset shows the $R-1$ values; the positive skew indicates that most R values tend to cluster at smaller values.

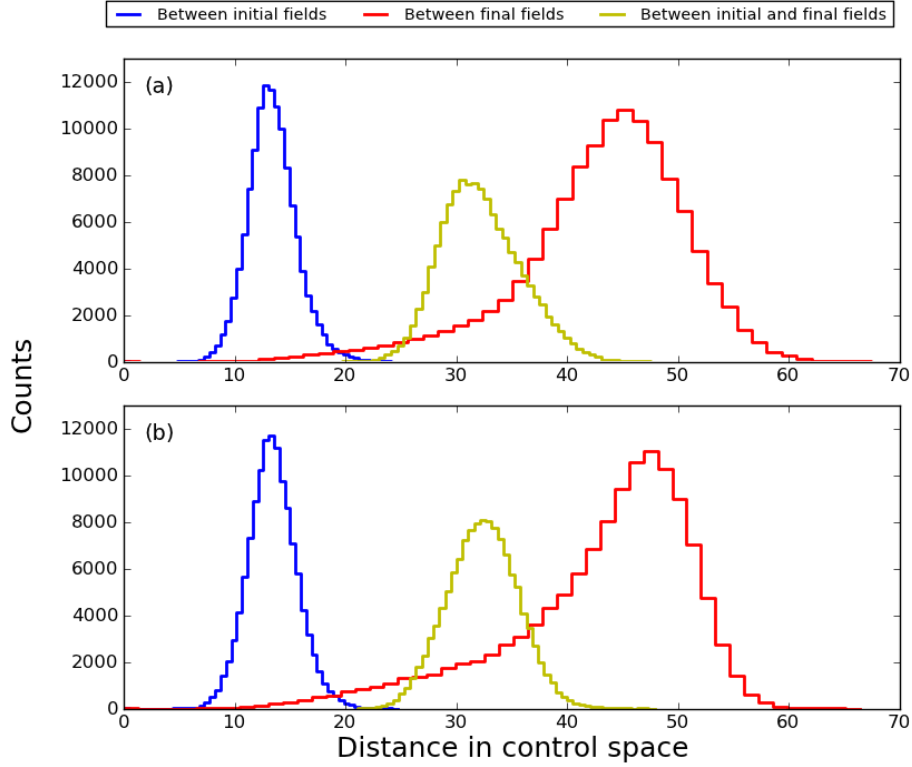


Figure 4: Distributions of pair-wise distances between initial control fields, final control fields, and between all initial-final field pairs. (a) The distance distributions resulting from the 500 optimization trajectories that yielded the lowest values of R from the 2,000 optimizations performed using the Hamiltonian from Eq. (28) and (29), collectively with a mean R value of 1.071. (b) The distance distributions resulting from the 500 optimization trajectories that yielded the highest values of R , which have a mean R value of 1.175. The similarity of the two above plots suggests that low R fields, i.e. initial control fields leading to low values of R when optimized, are distributed in control space in the same way as high R fields.




Figure 5: The distribution of R values resulting from demanding increasing fidelity of the target observable at the beginning and end of the landscape climbs. 1,000 runs were performed for each level of accuracy demanded. A slight increase in R values as the climbs stretch to further extremes of the landscape is visible, but the effect is small.

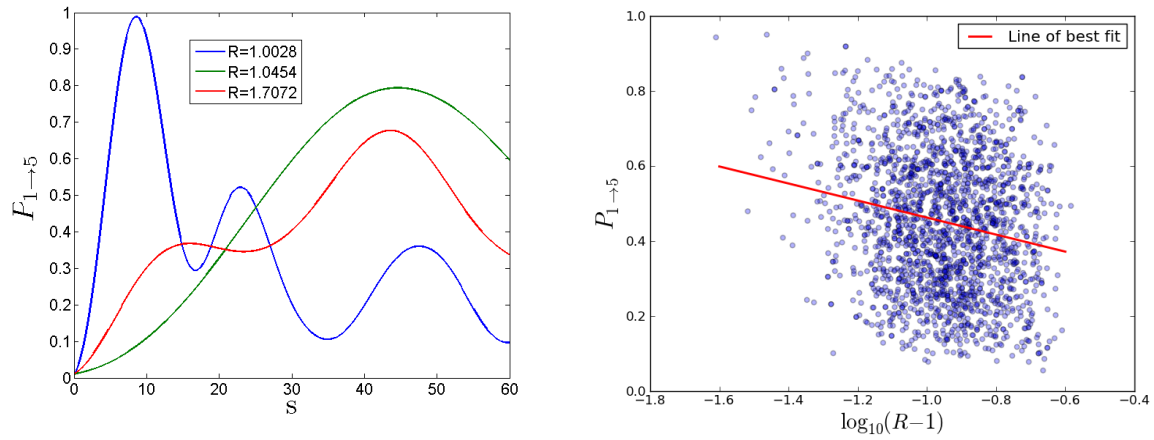


Figure 6: (a) The value $P_{i \rightarrow f}$ as a straight line is traversed in control space in the direction of the initial gradient starting at the initial control field is shown. The first maximum in the curve is taken as the optimal value obtained. Three trials are shown, each starting with an initial control field that yields different values of R . The yield obtained in this way is greater for the fields with lower R . (b) A scatterplot of the $P_{i \rightarrow f}$ values obtained when this method is used to optimize the 2,000 initial control fields used earlier versus the R values these fields yield. As R decreases, the yield obtained increases.




Figure 7: A surface displaying $\frac{\partial E(t)}{\partial s}$, the instantaneous change in the electric field, as a function of s as an optimization trajectory with a low R value is followed. The cross-section at constant s appears to be the same electric field up to a scale factor, which is modulated by the function $\alpha(s)$, whose bell-shaped character can be discerned from the surface.

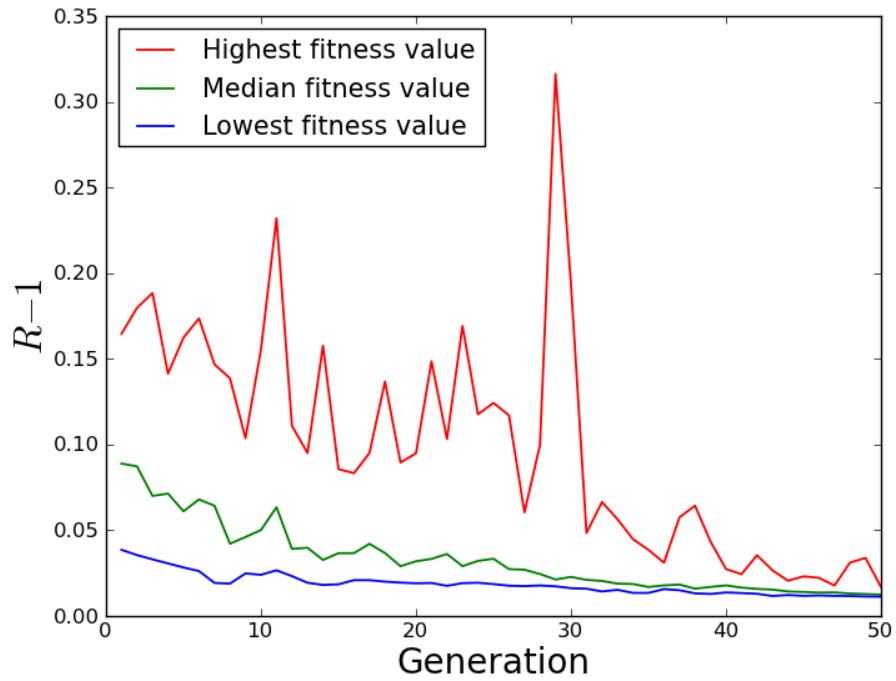


Figure 8: The highest, median, and lowest $R - 1$ values of the collection of initial control fields that make up the particle ‘swarm’ is displayed for each generation of one run of the Particle Swarm Optimization algorithm. The slow decrease to a (possibly suboptimal) low R value is apparent, as well as the convergence of the highest, median, and lowest $R - 1$ values indicating the collapse of the swarm.