Noname manuscript No.
(will be inserted by the editor)

On Choosing the Form of the Objective Functional for Optimal Control of Molecules

Carlee Joe-Wong · Tak-San Ho · Herschel Rabitz

Received: 20 May 2015 / Accepted: 15 September 2015

Abstract Optimal control techniques can be used to direct molecular dynamics to meet specified physical goals. However, the effectiveness of finding an optimal control field depends on the nature of the control landscape, defined as the objective as a functional of the control. Extensive analysis has considered the prospect of such landscapes being free of suboptimal traps for particular cases of different objective functions in both classical and quantum systems. While many typical objective functions have been shown to yield trap-free landscapes upon satisfaction of certain assumptions, this work more broadly considers the freedom in the choice of objective functionals. The latter freedom can be exploited to possibly accelerate the search for an optimal control, but we also show that the choice of functional needs to be made carefully to avoid inducing artificial landscape traps.

Keywords molecular control · control landscape · objective functional

1 Introduction

Laser-driven optimal control of molecular dynamics phenomena can be expressed in either a quantum or classical context, depending on the particular circumstance. Such optimal control efforts aim to maximize a specified objective function to meet a physical goal. A common quantum mechanical goal is to maximize the expectation value of a particular observable [3], e.g., the transition probability from one state to another. Similarly, a classically-specified optimal control objective may be for certain atoms in a molecule to reach specified positions and momenta [6]. The exact objective function(al) used, however, is not uniquely specified by the physical control goal. In this work, we consider the appropriate specification of the objective function to not only enforce reaching the control goal, but also ensure that an optimal control field can be found without encountering hindrances, possibly accelerating the search for the control. We aim to highlight this freedom in the choice of objective and to provide some initial guidelines for choosing the objective; we expect that future simulation and laboratory studies will shed further light on this choice.

Most computational algorithms for finding optimal controls are myopically gradient-based, and use local information to iteratively evolve the control field. In doing so, they attempt to climb the control landscape, which we define as the objective as a functional of the control field. The ability of these algorithms to find an optimal control field is thus determined in large part by the topology of the control

C. Joe-Wong
Program in Applied and Computational Mathematics
Princeton University

T.-S. Ho
Department of Chemistry
Princeton University

H. Rabitz
Department of Chemistry
Princeton University E-mail: hrabitz@princeton.edu

landscape. In particular, the presence of traps, or local extrema with suboptimal landscape values, can hinder or prevent the search for optimal control fields.

In the quantum context, a natural choice of objective function is the expectation value of an appropriate Hermitian operator (i.e., based on one of the postulates of quantum mechanics), which yields a trap-free control landscape upon satisfaction of three reasonable assumptions. These assumptions of controllability, surjectivity, and unconstrained control resources [3,5,9,10,12] are sufficient conditions for establishing landscape topology. Landscapes arising from other objective functions, including optimizing the Frobenius norm between a target unitary transformation matrix and that achieved under control, have also been shown to be trap-free [2,11]. However, the topology of complex ratios or other forms using multiple observables remains to be assessed [4]. In the classical context, there is no obviously natural choice for the objective function, which nominally can be an arbitrary function of the positions and momenta of the constituent atoms in the molecule of interest (see Section 2). Interestingly, satisfaction of the same three assumptions underlies the analogous topological analysis of classical control landscapes [6,7]. However, we show in this work that a poor choice of objective function can lead to landscape traps in either the classical or quantum case. Concomitantly, carefully choosing the objective function and possibly adjusting the objective throughout the search could accelerate reaching the landscape maximum, while avoiding the presence of traps.

The assessment here will be particularly focused towards gradient-based simulations, but the conclusions are relevant as well with other algorithms (e.g., stochastic methods are commonly performed in the laboratory), as a potentially high density of traps could be a confounding factor in virtually any circumstances. Section 2 considers the control landscape in classical mechanical scenarios for different types of objective functions. More extensive treatment will be given to classical system landscape analysis in this section, as it is a newer development and will permit introducing freedom in the choice of objective function, which also has a close parallel in the quantum case of Section 3. We then consider the quantum scenario in Section 3. Section 4 concludes the paper.

2 Control of Classical Systems

In this section, we consider the classical control landscape for different types of objective functions. The landscape analysis is separated into two parts: first, we summarize recent work [6] assessing landscape topologies by characterizing suboptimal critical points of the control landscape. For ease of exposition, we consider the control of a single molecule's phase space trajectory, but the results can be readily extended to an ensemble of trajectories in phase space [7].

2.1 Molecular Dynamics

Consider a molecule of n atoms, driven by a linearly polarized electric control field $\epsilon(t)$ over a finite time interval $[0,T]$. We use $\mathbf{z}(t)$ to denote the phase space state vector of the n atoms at any time $t \in [0,T]$. Here $\mathbf{z}(t)$ is a concatenated vector of the molecule's constituent atomic positions $\mathbf{q}(t)$ and momenta $\mathbf{p}(t)$: $\mathbf{z}(t) = \left[\mathbf{q}^T(t)\ \mathbf{p}^T(t)\right]^T$ with $\mathbf{q}(t)$ and $\mathbf{p}(t) \in \mathbb{R}^N$ and $N = 3n$, since each of the n atoms moves in three-dimensional space. These variables evolve according to the Hamiltonian

$$H^{\mathrm{cl}} = \frac{1}{2}\mathbf{p}^T\mathbf{M}^{-1}\mathbf{p} + V(\mathbf{q}) - D(\mathbf{q})\epsilon(t), \tag{1}$$

where the functions V and D represent the potential and dipole moment respectively; we assume that they are twice differentiable and finite at any $\mathbf{q} \in \mathbb{R}^N$. The mass matrix $\mathbf{M}$ is defined as the $N \times N$ matrix $\mathrm{diag}\{\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_n\}$, where $\mathbf{m}_i = \mathrm{diag}\{m_i, m_i, m_i\}$ and m_i is the mass of the i'th atom.

2.2 Critical Points of the Classical Control Landscape

We suppose that the control field ϵ is chosen so as to maximize an objective function of the final state, which we denote as $O(\mathbf{z}(T))$. We assume that O is twice-differentiable and define the control landscape to be the functional

$$J[\epsilon(\cdot)] = O(\mathbf{z}(T)), \tag{2}$$

where $O(\mathbf{z}(t))$ satisfies the dynamics $dO(\mathbf{z}(t))/dt = \left\{H^{\mathrm{cl}}, O(\mathbf{z}(t))\right\}$ with $\{\cdot,\cdot\}$ denoting the Poisson bracket. In principle, O can be freely chosen among all twice-differentiable functions of both position and momentum in classical mechanics. It is possible to add other terms to the objective, e.g., defining $J \equiv O(\mathbf{z}(T)) + \alpha \int_0^T \epsilon(t)^2\, dt,\ \alpha < 0$ to bias the chosen control field. However, in this work we focus on the landscapes defined in Eq. (2) and in particular on the choice of the objective function O.

Critical points of the control landscape can be found by differentiating Eq. (2) with respect to $\epsilon(t)$, $t \in [0,T]$:

$$\frac{\delta J}{\delta \epsilon(t)} = \frac{dO}{d\mathbf{z}(T)} \frac{\delta \mathbf{z}(T)}{\delta \epsilon(t)}. \tag{3}$$

A critical point corresponds to $\delta J/\delta\epsilon(t) = 0$ for all $t \in [0,T]$. We now make the assumption (1) that $\{\delta\mathbf{z}(T)/\delta\epsilon(t),\, t \in [0,T]\}$ is surjective, i.e., it spans $\mathbb{R}^{2N}$. Thus, a critical point is achieved if and only if $dO/d\mathbf{z}(T) = 0$. Given that $\{\delta\mathbf{z}(T)/\delta\epsilon(t)\}$ consists of $2N$-dimensional time-dependent vectors, each of infinite length in $t \in [0,T]$, it is likely that surjectivity holds. This assumption has been verified as satisfied in a number of numerical simulations [6,7].

We also make two additional assumptions: (2) controllability, i.e., that any initial state $\mathbf{z}(0)$ can be steered to any target state $\mathbf{z}^{\mathrm{tar}}$ at some time T ($\mathbf{z}(T) = \mathbf{z}^{\mathrm{tar}}$) with an appropriate control field $\epsilon(t)$, $t \in [0,T]$; and (3) that free access is available to any desired control field resource. The latter assumption is necessary so as to avoid limiting the practical ability to satisfy (1) surjectivity and (2) controllability. In realistic circumstances, control fields available in system-specific physically relevant spectral domain(s) are likely adequate to avoid hindering the dynamics. Numerical studies have shown at least a practical level of satisfaction of the three assumptions [6,7].

2.3 Avoiding Traps and the Choice of Objective Function

We now consider the topological nature of the critical points when the three assumptions hold. We first differentiate Eq. (3) to find the Hessian, which we denote as $\mathcal{H}(t,t') = \delta^2 J/\delta\epsilon(t)\delta\epsilon(t')$, at a critical point of the control landscape:

$$\mathcal{H}(t,t') = \left(\frac{\delta \mathbf{z}(T)}{\delta \epsilon(t')}\right)^T \frac{d^2 O}{d\mathbf{z}(T)^2} \frac{\delta \mathbf{z}(T)}{\delta \epsilon(t)}. \tag{4}$$

Since surjectivity holds, the nature of the critical point is determined by $d^2O/d\mathbf{z}(T)^2$. This matrix is symmetric, so we may write it as $d^2O/d\mathbf{z}(T)^2 = P^T S P$, where S is a diagonal matrix and P is an orthogonal matrix. Thus, traps, or local, sub-optimal maxima, do not occur if $d^2O/d\mathbf{z}(T)^2$ is not a negative-definite matrix at any critical point of the control landscape except the global maximum. For ease of terminology, we term this the *classical trap-free criterion* in the discussion below. The choice of the objective function $O(\mathbf{z}(T))$ can therefore be used to ensure that traps do not exist.

One type of function $O(\mathbf{z}(T))$ that meets the desired criterion is a concave function: all critical points of a concave function are both local and global maxima. Concavity, however, is not a necessary condition: consider, for instance, a monotonically increasing transformation of a concave function $O(\mathbf{z}(T)) = h(g(\mathbf{z}(T)))$, where h is a strictly increasing function and g is concave. It is easy to show that the critical points of $O(\mathbf{z}(T))$ are the same as the critical points of $g(\mathbf{z}(T))$. For instance, the Gaussian function

$$O(\mathbf{z}(T)) = \exp\left(-\left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)\right) \tag{5}$$

is a monotonically increasing transformation of the concave quadratic function $g(\mathbf{z}(T)) = -\left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)$, with a unique maximum located at $\mathbf{z}(T) = \mathbf{z}^{\mathrm{tar}}$.

We now show that even monotonic transformations of concave functions do not span the full class of objective functions for which the control landscape remains trap-free. To do so, we consider a bimodal function:

$$O(\mathbf{z}(T)) = -\left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right) \exp\left(-\frac{\left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)}{2}\right). \tag{6}$$

Taking the first derivative, we find two critical points: one local maximum of value zero at $\mathbf{z}(T) = \mathbf{z}^{\mathrm{tar}}$ and one local minimum at $\left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\mathrm{tar}}\right) = 2$. This function therefore has no local, non-global maxima, though it is neither concave nor a monotonic transformation of a concave function. We

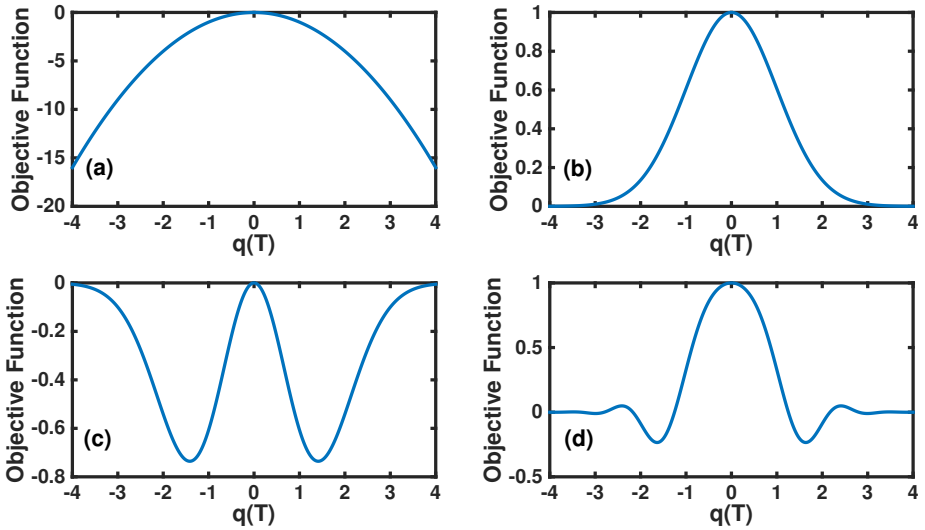

Fig. 1 Possible objective functions for classically-specified dynamics. In order to graphically depict the function, we take $\mathbf{q}$ and $\mathbf{p}$ to be scalars in the projection where $\mathbf{p} = 0$ with $\mathbf{z}^{\text{tar}} = 0$. The (a) concave quadratic and (b) Gaussian functions only have a global maximum at $\mathbf{q} = 0$, while the (c) bimodal function has global maxima at $\mathbf{q} = 0$ and $\mathbf{q} \to \infty$, where the latter condition corresponds to dissociation. The (d) Gaussian-weighted cosine function has a global maximum at $\mathbf{q} = 0$ and local, non-global maxima at $\mathbf{q}^2 = 2\pi k - \pi/4$, $k \in \mathbb{N}$.

note, however, that for the ith coordinate of $\mathbf{z}(T)$ satisfying $\mathbf{z}_i(T) \to \infty$, the function asymptotically approaches its global maximum of 0. The latter maxima correspond to dissociation of the molecule associated with the ith coordinate, which may not be a desirable goal in some cases (i.e., compared to achieving the same function value at $\mathbf{z}(T) = \mathbf{z}^{\text{tar}}$).

Conversely, it is easy to find objective functions that do *not* fit the criterion for a trap-free landscape. For instance, the Gaussian-weighted cosine function

$$O\left(\mathbf{z}(T)\right) = \exp\left(-\frac{\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)}{2}\right) \cos\left(\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)\right) \tag{7}$$

has a unique global maximum at $\mathbf{z}(T) = \mathbf{z}^{\text{tar}}$, but also has local maxima at $\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right) = 2\pi k - \pi/4$, for $k \in \mathbb{Z}$. Thus, the landscape corresponding to this objective function would have traps.

Figure 1 compares the four types of objective functions considered above: (a) a concave quadratic function $-\left(\mathbf{z}^T(T) - \mathbf{z}^{\text{tar}}\right)\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)$, (b) the Gaussian function in Eq. (5), (c) the bimodal function in Eq. (6), and (d) the Gaussian-weighted cosine function in Eq. (7). While all four have a unique global maximum at $\mathbf{z}(T) = \mathbf{z}^{\text{tar}}$, the Gaussian-weighted cosine function has additional local, non-global maxima. Thus, the designer's choice of objective function should take into account not just the global maximum of the objective, but also its overall structure.

An additional practical criterion for choosing the objective function is the efficiency with which gradient-based or other local algorithms can climb the control landscape. In a typical optimization with gradient algorithms, much of the algorithmic effort is often spent in the initial and last portions of the landscape climb as a result of a small gradient near the bottom and top of the landscape, respectively. One could accelerate convergence by increasing the sensitivity of J to changes in the control field (e.g., at the final stages of the search by increasing $\delta J/\delta\epsilon$ near the landscape maximum). For instance, the quadratic function $O\left(\mathbf{z}(T)\right) = -\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)$ has a linearly decreasing first derivative $\partial O/\partial \mathbf{z}(T) = -2\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T$ near its maximum, while the Gaussian $\exp\left(-\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)\right)$ weights the linear term of its first derivative with a quadratic exponential. Thus, choosing a quadratic objective would encourage faster convergence near the landscape maximum due to a steeper landscape gradient. A more dramatic example is provided by the sub-quadratic objective function $J = -\left(\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)\right)^\gamma$ for $0 < \gamma < 1$, whose first derivative approaches

$\pm\infty$ as $\mathbf{z}(T) \to \mathbf{z}^{\text{tar}}$. The gradient $\delta J/\delta\epsilon$ thus increases as the landscape maximum is approached to accelerate the convergence rate.

To ensure quick convergence both near and far from the maximum, it is desirable to define an objective function with relatively large gradients in both regimes. For instance, we might take

$$O\left(\mathbf{z}(T)\right) = -\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right) - \left(\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)\right)^\gamma. \tag{8}$$

When $\left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right)^T \left(\mathbf{z}(T) - \mathbf{z}^{\text{tar}}\right) > 1$, the quadratic term dominates, encouraging a rapid initial landscape climb. Near the landscape maximum, the sub-quadratic term dominates, steepening the landscape gradient relative to that with the quadratic function alone and accelerating convergence to the optimal field. Since this function O has no local, non-global maxima, the resulting control landscape is also trap-free.

3 Control of Quantum Systems

In the quantum control scenario, we suppose that the system evolves over an adequate finite time interval $[0, T]$ and consider a mixed-state quantum ensemble in which the state is described by a $\tilde{N} \times \tilde{N}$ density matrix $\rho(t)$ that satisfies

$$\imath\hbar \frac{\partial \rho}{\partial t} = \left[H^{\text{qu}}(t), \rho(t)\right] = \left[H_0 - \mu\epsilon(t), \rho(t)\right], \tag{9}$$

with Hamiltonian $H^{\text{qu}}(t) = H_0 - \mu\epsilon(t)$, where $[\cdot, \cdot]$ denotes the matrix commutator. Here H_0 and μ are, respectively, the field-free diagonal Hamiltonian and the dipole moment matrices. Equation (9) implies that $\rho(t) = U^\dagger(t, 0)\rho(0)U(t, 0)$, where $U(t, 0)$ is a unitary matrix propagator. Analogues of the three assumptions introduced classically are also adopted in the quantum case. We assume full controllability of U, i.e., that for an appropriate control field $\epsilon(\cdot)$, $U(0, 0)$ can be steered to an arbitrary unitary matrix $U(T, 0)$. Surjectivity desires that $\delta U(T, 0)/\delta\epsilon(t)$, $t \in [0, T]$, has full rank, and adequate control resources are assumed to be available. Extensive works have examined the landscape topology in the quantum case upon satisfaction of these assumptions [3,5,9,10,12], and here we strictly consider issues of the choice of objective function.

In principle, we can specify any objective function $O^{\text{qu}}(\rho)$ and then search for a control field ϵ maximizing O^{qu}. However, in practice the common situation is to express the objective function in terms of the expectation value of a quantum observable A^{qu} for maximization [3]. Extensive analysis has shown that $J = \text{Tr}\left(\rho(T)A^{\text{qu}}\right)$ is trap-free upon satisfaction of the same three assumptions that entered the classical case [3,9,10]. Although this choice of objective function naturally arises from one of the postulates of quantum mechanics, one could also choose various functions of J with similar cautionary guidance as in the classical case. In fact, a similar argument as in Section 2 shows that taking the objective to be $J = F\left(\text{Tr}\left(\rho(T)A^{\text{qu}}\right)\right)$ yields a trap-free landscape if the function F has no local, non-global maxima [4].

Other objectives, besides expectation values, can arise in quantum mechanics; for example, a case of interest is to generate a specific unitary transformation matrix W. One possible objective functional is $J[\epsilon(t)] = \|U(T, 0) - W\|^2$, where $\|\cdot\|$ denotes the Frobenius norm. The resulting control landscape has been shown to be free of traps under the same assumptions as above [2,11]. Another objective of interest is to simply create a desired target state ρ^{tar}; the resulting objective function $J = \|\rho^{\text{tar}} - \rho(T)\|^2$ similarly yields a trap-free landscape. Many additional objective functions have been proposed for the latter purpose [1,8], but the presence of traps in the resulting control landscapes remains to be assessed.

4 Conclusion

In this work, we consider the role of choosing the objective function in determining the topology of optimal control landscapes while meeting the physical objective and satisfying other criteria. A succinct summary of prior classical and quantum control landscape topology assessments is given as a basis for guiding this choice. The classical context was used to give conditions under which the choice of objective function ensures a trap-free control landscape and to present some examples of functions that satisfy these conditions. Like freedom exists in quantum mechanics, even with the fortunate assurance that the

postulated observable $J = \mathrm{Tr}\,(\rho(T)A^{\mathrm{qu}})$ of quantum mechanics is trap-free. Some alternative function cases in quantum mechanics for state preparation have yet to be assessed for their topology.

The customary freedom available in choosing objective functions could also aid in accelerating the convergence of the gradient or other local algorithms in reaching the landscape maximum. In particular, a typical optimization rises slowly at low and high values of J, while rapidly traversing the intermediate portions of the landscape. This situation might be ameliorated by increasing the "sensitivity" of J to the minimum and maximum regions by a suitable transformation of the objective function. The goal of this paper is to draw attention to the freedom in the choice of objective function, and utilization of this freedom (with duly noted caution about the introduction of traps) will likely call for practical experience gained in future simulations and laboratory studies.

Acknowledgements The authors acknowledge support from ARO (W911NF-13-1-0237 and MURI W911NF-11-1-2068) for the quantum aspects of the research, and DOE (DE-FG02-02ER15344) for the classical analysis.

References

1. Gilchrist, A., Langford, N.K., Nielsen, M.A.: Distance measures to compare real and ideal quantum processes. Physical Review A **71**(6), 062,310 (2005)
2. Ho, T.S., Dominy, J., Rabitz, H.: Landscape of unitary transformations in controlled quantum dynamics. Physical Review A **79**(1), 013,422 (2009)
3. Ho, T.S., Rabitz, H.: Why do effective quantum controls appear easy to find? Journal of Photochemistry and Photobiology A: Chemistry **180**(3), 226–240 (2006)
4. Ho, T.S., Rabitz, H.: Multiple-target landscapes in quantum control dynamics (2009). Unpublished manuscript, revised 2014
5. Hsieh, M., Wu, R., Rabitz, H.: Topology of the quantum control landscape for observables. The Journal of Chemical Physics **130**, 104,109 (2009)
6. Joe-Wong, C., Ho, T.S., Long, R., Rabitz, H., Wu, R.: Topology of Classical Molecular Optimal Control Landscapes in Phase Space. The Journal of Chemical Physics **138**, 124,114 (2013)
7. Joe-Wong, C., Ho, T.S., Rabitz, H., Wu, R.: Topology of classical molecular optimal control landscapes for multi-target objectives. The Journal of Chemical Physics **142**(15), 154,115 (2015)
8. Mendonça, P.E., Napolitano, R.d.J., Marchiolli, M.A., Foster, C.J., Liang, Y.C.: Alternative fidelity measure between quantum states. Physical Review A **78**(5), 052,330 (2008)
9. Rabitz, H., Ho, T.S., Hsieh, M., Kosut, R., Demiralp, M.: Topology of optimally controlled quantum mechanical transition probability landscapes. Physical Review A **74**(1), 12,721 (2006)
10. Rabitz, H., Hsieh, M., Rosenthal, C.: Quantum optimally controlled transition landscapes. Science **303**(5666), 1998–2001 (2004)
11. Rabitz, H., Hsieh, M., Rosenthal, C.: Landscape for optimal control of quantum-mechanical unitary transformations. Physical Review A **72**(5), 052,337 (2005)
12. Wu, R., Rabitz, H., Hsieh, M.: Characterization of the critical submanifolds in quantum ensemble control landscapes. Journal of Physics A: Mathematical and Theoretical **41**, 015,006 (2008)