PHILOSOPHICAL TRANSACTIONS A

rsta.royalsocietypublishing.org

Research

Cite this article: Russell B, Rabitz H. 2017 Common foundations of optimal control across the sciences: evidence of a free lunch. *Phil. Trans. R. Soc. A* **375**: 20160210.
http://dx.doi.org/10.1098/rsta.2016.0210

Accepted: 20 November 2016

One contribution of 15 to a theme issue 'Horizons of cybernetical physics'.

Subject Areas:
physical chemistry, quantum engineering, applied mathematics, differential equations, quantum physics, mathematical physics

Keywords:
optimization, control, control landscapes

Author for correspondence:
Benjamin Russell
e-mail: br6@princeton.edu

THE ROYAL SOCIETY PUBLISHING

Common foundations of optimal control across the sciences: evidence of a free lunch

Benjamin Russell and Herschel Rabitz

Department of Chemistry, Princeton University, Princeton, NJ 08540, USA

BR, 0000-0003-2038-0453

A common goal in the sciences is optimization of an objective function by selecting control variables such that a desired outcome is achieved. This scenario can be expressed in terms of a control landscape of an objective considered as a function of the control variables. At the most basic level, it is known that the vast majority of quantum control landscapes possess no traps, whose presence would hinder reaching the objective. This paper reviews and extends the quantum control landscape assessment, presenting evidence that the same highly favourable landscape features exist in many other domains of science. The implications of this broader evidence are discussed. Specifically, control landscape examples from quantum mechanics, chemistry and evolutionary biology are presented. Despite the obvious differences, commonalities between these areas are highlighted within a unified mathematical framework. This mathematical framework is driven by the wide-ranging experimental evidence on the ease of finding optimal controls (in terms of the required algorithmic search effort beyond the laboratory set-up overhead). The full scope and implications of this observed common control behaviour pose an open question for assessment in further work.

This article is part of the themed issue 'Horizons of cybernetical physics'.

1. Introduction to control landscapes

Cybernetics is often referred to as 'the science of communications and automatic control systems in both machines and living things' [1]. This paper, based on

© 2017 The Author(s) Published by the Royal Society. All rights reserved.

a presentation at the 2016 IEEE Conference on Norbert Wiener in the 21st Century: Thinking Machines in the Physical World, fully embraces the spirit of cybernetics by considering control throughout the sciences and beyond in the twenty-first century.

Since Weiner wrote his original work, optimization has become ubiquitous throughout the sciences and engineering. In the laboratory and in many industrial settings, one seeks to optimally control a physical, chemical or biological system to induce a specific transformation using an adequate resource set. In natural evolution, it is generally accepted that Nature is performing a stochastic optimization expressed in the survival of the fittest. Although not humanly driven, we also discuss evolution in a control context.

A control landscape is an objective function, or *goal*, subject to optimization by choosing control parameters in a system. The landscape metaphor is established by considering the graph, or higher dimensional analogue, of the objective as a function of the control parameters. The concept of a fitness landscape in evolutionary biology was introduced [2], where it is prevalent. In recent years, the study of control landscapes has focused extensively on quantum systems which evolve according to Schrödinger's equation within a well-defined mathematical structure. As such, this work will primarily consider the inherent simplicity of control landscapes of quantum systems and further present evidence that these features exist also outside of quantum mechanics and propose an underlying mathematical basis for this commonality. In particular, evidence abounds in chemistry and material science, evolution and engineering that control landscape behaviour exists like that found in quantum mechanics.

Control systems come in many forms; however, there are generic mathematical structures underlying many practical cases. In order to set the framework for the remainder of this paper, we first describe a general class of mathematical problems given by a system of differential equations comprising a physical model of the system under control

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = F(x(t), w(t)), \tag{1.1}$$

where $x(t)$ is the state vector at time t and $w(t)$ is the control chosen by an experimenter or by Nature as in biological evolution. This class of equations is vast and includes diverse examples from many domains. The Schrödinger equation is of this form, as are some models used in molecular dynamics and evolutionary biology [3]. In some cases, the variables defining the structural form of F may also be considered part of the control.

In order to enquire about *optimal* control, one must have an objective function to optimize. This function is often represented by: $\mathcal{F}[w] = J(x(T))$, a function of the state at some given final time T, which we seek to maximize by choosing (possibly time-dependent) control variables, or simply *controls*, represented by w. Such functions are known as terminal cost/pay-off functions, as they only depend on the state at some final time without additional 'run-time' costs depending on $w(t)$.

In such scenarios, it becomes prescient to ask about the structure of the control landscape, which is crucial in establishing the feasibility of finding optimal controls. We will render this feasibility assessment in terms of two issues: (1) the existence of landscape traps and (2) the landscape optimization convergence rate. Both of the issues relate strongly to the nature of the algorithms seeking an optimal control. The presence of landscape traps, that is, sub-optimal extrema, would call for a stochastic algorithm to 'step over' such features. The second issue (2) addresses the search effort required to locate an optimal control. In laboratory cases, where there can be hundreds of control variables present, out to biological evolution for which there are approximately 10^9 gnomic sites for Nature to select nucleic acids, the success of a multitude of experimental and natural optimization processes indicates that the structure of these landscapes is favourable (i.e. free from traps under reasonable assumptions), permitting rapid convergence despite the curse of dimensionality.

In principle, and perhaps intuitively, one might expect the complexity of many control systems to result in landscapes that possess large numbers of local optima. We shall, however, argue that a specific notion of complexity is favourable for control optimization rather than deleterious, to actually reduce the possibility of traps. Figures 1 and 2 illustrate some low-dimensional (two

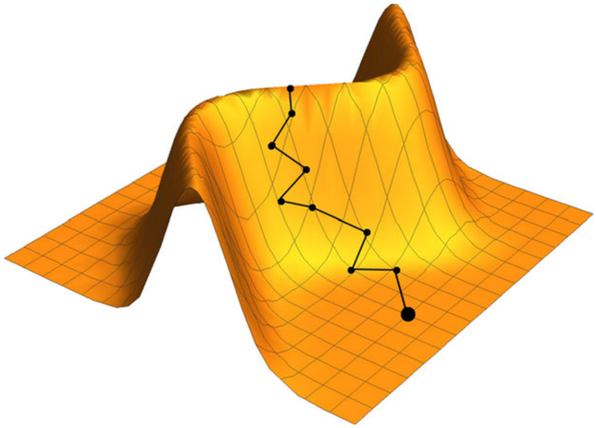

Figure 1. A control landscape with a plateau (showing a one-dimensional critical manifold of global optima) but no with traps. Landscapes of this type, or even with additional saddle features, are highly favourable for finding optimal controls. (Online version in colour.)

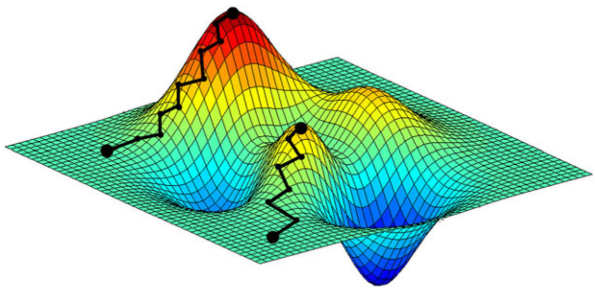

Figure 2. A local gradient search climbing the landscape reaching either dead-end sub-optimal traps or the desired true global optimum, depending on the starting point. Landscapes of this type, especially in high dimensions, are highly problematic for finding optimal controls. (Online version in colour.)

control parameters) examples. In practice, typically, many more control parameters than two are present and control landscapes cannot be directly visualized.

The *critical point topology* of a control landscape determines the behaviour of local optimization algorithms, such as gradient ascent and even non-local stochastic algorithms could be greatly hindered by a rough, high-dimensional landscape. An understanding of the critical point topology, i.e. the set of controls where $\nabla\mathcal{F}=0$, permits assessing the ease of finding optimal controls. Saddle points could exist on some landscapes, but they would only slow down a search rather than stop it from proceeding to full optimization. For an application of a gradient method in laboratory practice, see [4] in which the algorithm is applied to spectrally filtered and integrated second harmonic generation as well as excitation of atomic rubidium. In this work, monotonic convergence to maximum fidelity is seen within practical laboratory time scales. Work on assessing the experimental relevance of saddle points has also been undertaken [5] in which their impact was found to be negligible, as theoretically expected.

The remainder of this paper is organized as follows. Section 2 summarizes the experimental and simulation behaviour seen when seeking optimal controls in the domain of quantum phenomena, chemistry and material science, and natural evolution. Importantly, this evidence stands as a body of empirical facts on the favourable structure of landscapes that demands an explanation, which we claim has a common foundation. Section 3 focuses on that explanation in quantum control, wherein detailed and exacting mathematical models exist. Section 4 starts with a brief summary of the situation in chemistry and material science, as well as evolution, and the reader will be referred to the literature for specific details of these analyses. The most challenging case of equation (1.1) covers such a wide range of behaviour that achieving a full landscape analysis has not yet been achieved. But, we lay out the framework of conditions for finding favourable landscapes in this broad class of systems and offer arguments that the conditions plausibly hold. In this regard, we further offer a well-defined mathematical conjecture to be investigated in the future. The evidence in §2 speaks to the validity of this conjecture across varied domains, and, if further analysis expands the present support, it would offer a highly unified picture of optimization across the sciences and engineering.

2. Evidence of common landscape structure in the sciences

Recent work [6] has identified evidence of common landscape structure from the control of chemical, physical and biological evolutionary processes. Both the success of the optimization throughout these diverse areas and the remarkably efficient search effort required are summarized in table 1, alongside the nature of the sufficient conditions required to theoretically establish this simplicity.

(a) Quantum control

Control of a quantum system using electromagnetic fields generally encompasses one of the following tasks:

(i) Control the wave function (or density matrix ρ), as illustrated in figure 3, of a quantum system to create a specific state transformation $|\psi_{\mathrm{I}}\rangle \mapsto |\psi_{\mathrm{F}}\rangle$ [7,8].
(ii) Maximize the expectation of a particular quantum observable $\hat{O}$ expectation value [9].
(iii) Create a specific quantum gate G, which is a unitary transformation [10].

There is strong numerical [11,12], mathematical [13,14] and some experimental evidence [5,15] that the landscapes of closed, finite-dimensional quantum systems are trap free, corroborating the analytic claims when the assumptions are met. A few cases deliberately constructed to possess traps [16–18] are, however, known. Work in [11] assesses the few known quantum control examples with traps to better understand the volume of the traps. Analysis in [19] goes on further to numerically search for and analyse the typical singularities in quantum control landscapes for four level systems, and concludes that all the identified singularities are saddles. The work in [20] studies the effect upon the landscape of terms beyond linear order in the control within the Hamiltonian. An additional term, quadratic in the control field, is added to the Hamiltonian. It is found that the effect on the control landscape involves the removal of some specific traps known to exist without such a term. This finding is a step towards a general conjecture in §4 that nominal system complexity aids the ability to find optimal controls.

In this work, we will focus on closed quantum systems, i.e. those which do not interact significantly with their environment. However, there is existing work on the landscapes for observable preparation in open quantum systems [21], and a proof that any two-level open system (meeting simple further assumptions) is free from traps [22].

(b) Chemical yield and material property optimization

Evidence relevant to the quality of discovered optimal controls, and the optimization search effort required to find them, abounds in chemical and material sciences [6,23]. These data are from several different objectives corresponding to common goals in chemistry [23,24] of chemical yield and material property optimization. A clear pattern is present, especially within robotized experiments, that the number of samples required to obtain effective controls is

rsta.royalsocietypublishing.org Phil. Trans. R. Soc. A 375: 20160210

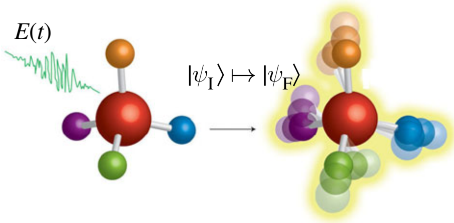

Figure 3. A molecule under vibrational control by a laser field $E(t)$. (Online version in colour.)

Table 1. Summary of landscape analysis evidence and nature of the associated sufficient conditions in different areas. Importantly, the sufficient conditions which support the existence of trap-free landscapes in all three categories are essentially the same, although the analysis methodology is different in each domain.

	quantum control	chemistry and material science	natural evolution
evidence implying simplicity or demonstrating a trap-free landscape	extensive successful numerical simulations and some experimentally observed landscapes	many successful experiments, including direct landscape observations	many diverse perturb and observe experiments free from evident traps, and the existence of complex life forms
sufficient assumptions for proving the structure of a trap-free landscape	three precise physical assumptions	the same assumptions as in quantum mechanics, but also the ability to manipulate the environment	assumption that all probability distributions over a species population can be created by appropriate genomic variations

typically infinitesimal compared with the full search space of chemical and operating conditions. This behaviour strongly suggests that the processes being optimized possess landscapes rife with many global optima, with few or no traps, thereby permitting highly rapid ascent to a global optimum in control space.

(c) Optimization in biological evolution

The concept of a fitness landscape is pervasive in natural evolution. In this context, the fitness landscape in genetics can be thought of as a graph of genotypic allele frequency in a population, against phenotype fitness [25], which is further developed in [26]. In this sense, the genes of a species population (or even over multi-species populations) are the 'controls', and the fitness is the optimality condition. However, many subtle variants exist in the exact nature of the fitness studied including using population growth rate as a measure of fitness, or average number of individual offspring.

One important means of assessing genetic landscapes for traps in the laboratory are 'disturb and observe' experiments. In one form of such experiments, a small change is made to the genome of an organism with a short life cycle, such as *Drosophila* fruit flies. Then the population of such genetically 'disturbed' organisms is allowed to freely evolve and the average population fitness is monitored to assess if it returns to the original, undisturbed value. If the fitness returns to its original value, this is evidence that the fitness landscape, at least in the vicinity of the initial genomic distribution, is free from local optima. In many known cases [26], the fitness was seen to return to its original value but with a different genome indicative of a level set at the top of the landscape, as shown in figure 2. Beyond such laboratory experiments lies the most stark evidence—that complex organisms exist. The period of time taken for the appearance of bacteria, and beyond to higher forms of life, as only a few billion years, is very short given their inherent complexity and the required multitude of elaborate evolutionary steps. These findings, particularly the latter, have left a puzzle especially considering that intuition would suggest that evolutionary landscapes with approximately 10^9 nucleic acid sites would be very rugged.

The evidence in the scientific domains of quantum mechanics, chemistry and material science, and evolution exists in a vast body of the literature, which must be taken as foundational facts to be reconciled, perhaps individually in each domain, or more enticingly in a sweeping consistent manner as indicated in the second portion of table 1 and the remainder of this paper. The details in treating the evidence differ in swinging between domains, but the analyses are similar, which eventually leads to conjecturing that the scope of observed landscape behaviour extends to the general structure of equation (1.1) upon satisfaction of three assumptions, enumerated in the remainder of the paper.

3. Specification of landscape assumptions for quantum control

As the above evidence points towards the existence of common features in the optimal control of systems arising in diverse areas of science, this finding provides an enticing incentive to seek a unifying explanation for the origin of this simplicity. The focus of the analysis in the literature thus far has been in quantum control, which, of all the domains, has the most detailed mathematical foundation to rest on. This foundation has aided in assessing the way in which the key assumptions (i.e. sufficient conditions), and the conclusions they imply about trap-free control landscapes, generalize to systems outside of quantum mechanics.

In [27], the first essential principles of the analysis of quantum control landscapes were set out, as were a nascent version of the three assumptions sufficient for trap-free quantum control landscapes which are summarized below, and later generalized in §4. Shortly after the introduction of the 'photonic reagent' control concept [28], the search effort required for the discovery of effective quantum controls was quickly identified to be dramatically less than expected, both in simulation and in mounting numbers of experiments.

While the case of a quantum system controlling another [29] does fall within the class of nonlinear control problems studied in this work (1.1), here we restrict our attention to quantum systems under control by semi-classical fields, as is typical throughout the vast majority of the quantum control literature. Closed systems evolve in time according to the Schrödinger equation:

$$\frac{\mathrm{d}U_t}{\mathrm{d}t} = \mathrm{i}H[w(t)]U_t, \tag{3.1}$$

which is an example system of (1.1) for which the propagator U_t evolves on the Lie group of $n \times n$ unitary matrices $U(n)$. Here, the Hamiltonian $H[w(t)]$ depends on the time-dependent control function $w(t)$, which often represents an external laser field (i.e. a photonic reagent). A common approximation in physical scenarios where the quantum state of a molecule is controlled by a laser is the dipole approximation:

$$H[w(t)] = H_0 + w(t)H_c. \tag{3.2}$$

For dipole systems of the form in (3.2) (which includes the control of nuclear spins in nuclear magnetic resonance), an in-depth landscape analysis has been performed for a control which implements a desired unitary transformation [30], quantum state transfer [31] or maximizes a desired quantum observable [32]. The satisfaction of a set of three simple assumptions is known to guarantee that the control landscape in the dipole approximation is free from traps in all these cases.

In the discussion of quantum control in this work, the cost function,

$$J(U) := |\mathrm{Tr}(G^{\dagger}U)|^2, \tag{3.3}$$

will be used to measure how well the unitary goal gate G was implemented by a system for which the end-time propagator is $U_T \equiv U$. Given a system where the assumptions hold, there may still be sub-optimal critical points introduced into the landscape by the cost function J itself (3.3); however, these are known analytically to all be 'benign' saddles [33] rather than true local optima.

The three assumptions of this analysis are described below and a proof that they are sufficient for a trap-free landscape is included in appendix A.

(a) Assumption I: the system is controllable

The definition of controllability is: the final time T and Hamiltonian H are such that for every objective evolution of the system there exists at least one control w such that $U_T = G$. For systems in the dipole approximation (3.2), a criterion ensuring controllability exists [34,35]. This criterion, known as the *Lie algebra rank condition*, is necessary and sufficient for controllability. It states that H_0 and H_c must 'generate' the whole algebra of observables (Hermitian matrices), in the sense that a basis can be created from a succession of Lie bracket expressions of the two generators, $\{H_0, H_c, [H_0, H_c], [[H_0, H_c], H_0] \ldots\}$, in order for the system to be controllable. For a visual representation of checking this criterion, see [36]. It is further known that for typical quantum systems of the form (3.2) the criterion holds [35]. Specifically, only a null set of Hamiltonians (a set of zero measure, or equivalently probability zero, if cases are chosen at random) fail to be controllable.

(b) Assumption II: the system is locally controllable in the direction of the cost function gradient

Local controllability is defined as the ability, given any control, to arbitrarily vary the system's state at the final time by varying the control infinitesimally. The required landscape condition is a special case of this property, as we only require that the state at the final time can be varied in *one* specific direction. More explicitly, the second assumption is: for all controls $w(t)$ there exists a small change in the control δw such that δU_T has some component in the direction ∇J. That is,

$$\langle \nabla J|_{U_T}, \delta U_T \rangle \neq 0, \quad \forall \delta w. \tag{3.4}$$

This expression means that the control can be varied slightly to increase the objective function J by 'steering' U_T in the direction ∇J and thus towards the goal (i.e. in the direction of increasing J). Informally, this criterion is simply that it is possible to 'steer' the system up the landscape by a small variation of the control. This property can also be proved to typically hold, even well beyond systems in the dipole approximation. Specifically, it has been shown that, in analogy with the results of [35], this second assumption only fails for a null set of Hamiltonians. For a detailed discussion of the somewhat technical underlying mathematics of this result, which is based on an application of the parametric transversality theorem [37], see [38].

(c) Assumption III: the control resources are unrestricted

This assumption means that the control w is not limited in its form. However, in practice there are always restrictions on the control in the laboratory. Although this assumption technically needs to be adopted in order to mathematically prove that quantum control landscapes are almost all trap free, in practice adequate freedom in the control often suffices. Pragmatically, this reduces to allowing the control to have *sufficient* freedom to exploit assumptions I and II [39]. For a review of quantum control landscapes in the presence of severe constraints, see [40–44]. A commonly applied constraint is to limit the total field fluence: $Q = \int_0^T |E(t)|^2 \, dt$; this constraint is known to be able to introduce traps into the landscape above a certain critical value. We generally expect that traps will appear in quantum control landscapes when severe constraints are imposed upon the control fields. However, the exact conditions when constraints introduce traps into the landscape

rsta.royalsocietypublishing.org Phil. Trans. R. Soc. A 375: 20160210

is not known analytically and could form the basis of future work as it is a topic with practical significance.

To reiterate the main conclusion: upon satisfaction of all three assumptions about quantum control, it is known that almost all landscapes are free from traps.

4. Landscapes beyond quantum mechanics

Building on the previous results, this section draws a general conclusion and conjectures about about the behaviour of a wide class of controlled phenomena under optimization. Initially, the landscapes of linear time-invariant (LTI) systems are analysed and compared with the quantum system case, then some analysis and a conjecture are offered on the fully nonlinear case (1.1). We simply remark that the landscape analysis for chemical and material science and for evolutionary biology follow a somewhat different track due to their complexity and the general nature of being open to the environment. See the sufficient conditions in table 1, and details in [24,26].

There is large freedom to choose a cost function J for both LTI (see (4.2)) and nonlinear systems. For control systems, both LTI and nonlinear, which evolve on $\mathbb{R}^N$, a common and favourable choice is

$$J(X) = \|\boldsymbol{X}(T) - \boldsymbol{X}_{\mathrm{G}}\|^2, \tag{4.1}$$

where $\boldsymbol{X}(T)$ is the final time state and $\boldsymbol{X}_{\mathrm{G}}$ is the goal (to be minimized in this case). This function is straightforwardly confirmed to possess no saddles or local optima by a gradient calculation. For systems which evolve on more general state spaces than $\mathbb{R}^N$, other cost functions are appropriate and a fully general canonical choice is not widely agreed upon. However, the conclusions of this work apply to any cost function J which possesses no local optima *of its own* as a function of the state (rather than the control), even in the fully nonlinear case.

(a) Autonomous linear systems with unrestricted controls are almost all trap free

The study of linear control systems is pervasive in theory [45,46] and applications [47], which motivates the study of their landscapes, as does the potential for using them as a stepping stone to analysing the landscape of the fully nonlinear case (1.1).

LTI systems are defined by the equation:

$$\frac{\mathrm{d}\boldsymbol{X}}{\mathrm{d}t} = A\boldsymbol{X}(t) + w(t)\boldsymbol{b}, \tag{4.2}$$

where $\boldsymbol{X}(t) \in \mathbb{R}^N$ is the state vector, A is an $N \times N$ matrix, $\boldsymbol{b} \in \mathbb{R}^N$ is a vector coupling to the control. Equation (4.2) warrants explicit comparison with (3.1), which is a *bilinear* control system as a product of the state and the control is present.

For system (4.2), it is possible to check the three assumptions of landscape analysis directly, showing that they almost always hold in the space of all systems of form (4.2).

(i) Assumption I: the system is controllable

The first assumption of controllability has a clear analogue to that of assumption I in quantum control, i.e. all states $\boldsymbol{X}$ can be obtained by applying some control $w(t)$. A sufficient condition for this to hold is known. One must check that the *controllability matrix* is full rank [46], particularly that the columns of

$$[\boldsymbol{b}, A\boldsymbol{b}, \ldots, A^{(N-1)}\boldsymbol{b}] \tag{4.3}$$

form a basis, or equivalently that $\boldsymbol{b}$ is a cyclic vector of A. One can readily check that this property holds for typical A and $\boldsymbol{b}$ (see appendix B) in a similar sense to the quantum case. An analogous condition is developed for the linear time-varying case in [48].

rsta.royalsocietypublishing.org Phil. Trans. R. Soc. A 375: 20160210

(ii) Assumption II: the system is locally controllable in the direction of the cost function gradient

LTI systems uniquely have the property that controllability is equivalent to local controllability (see appendix C), and thus it follows that assumption II holds whenever assumption I holds.

(iii) Assumption III: the control resources are unrestricted

In the LTI context, this assumption remains unchanged compared with the quantum control case.

(b) Control of nonlinear systems

In this section, landscapes for systems of the form (1.1) will be discussed. See also the paper by A. Fradkov [49], which introduces this theme issue on Norbert Wiener.

We first note that the Schrödinger equation and the LTI equation are special cases of (1.1), as well as a Schrödinger equation nonlinear in the control [20,50]. The primary open question is that of the fullest possible scope of the landscape observations in this paper. To tentatively assess this topic, we investigate the same three assumptions in this broader context of the control of nonlinear systems.

We see that the first assumption of global controllability can be assessed in a wide variety of cases and that it is known to hold for a large and practically relevant class of nonlinear control systems. We further argue that the third assumption is on an identical footing to the LTI and quantum control cases. The conclusion of this section is that gaining a deeper understanding of the second assumption is central to a fuller analysis of the landscapes of nonlinear systems in general.

In the quantum case, it was shown that *almost all* (in a specific sense that the failing set is null) quantum control systems are locally controllable. We argue that the same mathematical tool set as that used in quantum mechanics [38] can be deployed in the assessment of the status of the second assumption in the nonlinear case, but we do not give a full proof that the analogous conclusion holds in the nonlinear case. We highlight that assessing the status of the second assumption in the nonlinear case is the key open issue in understanding their control landscapes.

(i) Assumption I: the system is controllable

Criteria for the controllability of several large classes of nonlinear systems are known [51–53]. Specifically and importantly, the closest generalization of the controllability conditions applied to quantum (bilinear) and LTI control systems are developed in [51,54], wherein it is shown that a sufficiently small nonlinear term added to an LTI system will not violate global controllability. This circumstance allows the conclusions about assumption I—that almost all LTI and quantum systems are globally controllable—to be pushed directly through to nonlinear systems.

In [54], several 'rank' conditions are described, which are analogous to the results in [35] for quantum systems and to [55] for LTI systems. In [51], it is shown that if a system has a controllable linear part and the nonlinear part is 'small enough' in a specific and unrestrictive sense, then the overall system is also globally controllable. It follows from the observation that almost all LTI systems are controllable and the conclusions of [51] that a very rich class of nonlinear control systems are globally controllable. In this sense, understanding the linear case serves as a gateway to understanding which of the nonlinear family of systems are globally controllable, and to constructing a multitude of globally controllable examples.

(ii) Assumption II: the system is locally controllable in the direction of the cost function gradient

One measure of the *complexity* of a control system is its degree of local controllability (i.e. the number of linearly independent directions the state can be steered in by infinitesimally varying the control). These quantities can be expressed as:

$$\left\{ \frac{\partial \mathbf{X}}{\partial w_k} \right\} \tag{4.4}$$

for each control variable w_k or, if the control is a function of time, as the functional derivative:

$$\left\{ \frac{\delta \mathbf{X}}{\delta w(t)} \right\}. \tag{4.5}$$

If the span of the set of *all* such derivatives contains ∇J for every control, then it is possible to infinitesimally vary the control in a way which increases J everywhere on the system's landscape. The gradient of fidelity with respect to the control is then non-zero, and gradient ascent can continue to increase fidelity. The presence of more variables (i.e. nominally, additional system complexity) can aid the simplicity of optimal control by better ensuring the satisfaction of assumption II. For important work offering conditions on the local controllability of nonlinear systems, see [56], and for an assessment of their local controllability specifically at equilibrium points, see [57].

(iii) Assumption III: sufficient resources

The overall nature of the required resources in the control of nonlinear systems is parallel to that in LTI and quantum control. In the case of unconstrained resources, the argument that the control landscape will be trap free (given the satisfaction of the other two assumptions) is identical to the quantum and LTI cases (appendix A). Further work is needed to investigate the effect of nonlinearity on which resource constraints introduce landscape traps, as this remains an open issue.

Importantly, satisfaction of the three assumptions for a nonlinear system is sufficient to ensure a trap-free landscape. However, we at this time cannot assess what is *typical* for nonlinear systems in the same sense as the quantum or LTI cases. The fullest range with which the three assumptions hold remains open for investigation and is a topic of prime importance.

5. Conclusion and future work

We have reviewed numerical and experimental evidence motivating a unifying mathematical framework for understanding control landscapes. We presented the state of the field of quantum control landscapes, and explained the generalization of these ideas to the control landscapes of more general systems. It has further been argued that an explanation for this commonality may be found in the control landscapes of general nonlinear control systems. It is very interesting that system complexity should be conducive to the simplicity of landscape critical point topology, as singular critical points are known to be very unlikely to exist in this scenario. Furthermore, §§3 and 4 clearly indicate that the same set of primary assumptions apply to establishing control landscape topology for a vast expanse of systems and domains.

It is interesting to note that, within the optimal control problems discussed herein, a function of the following form is being optimized:

$$F(w) := J(x_T(w)). \tag{5.1}$$

The 'no free lunch' theorem [58,59] states that, averaged over all functions to optimize, no one algorithm outperforms any other. It is notable that this result is not an issue in control as only a special functional form is being optimized. This is consistent with the results and observations in this work, and we further speculate about the existence of a free lunch upon seeking optimal control.

There are many additional directions to explore in the domain of general system control landscape analysis. These include systems having time-dependent drift dynamics and stochastic systems. For the systems analysed in this paper, all time dependence enters the systems through the control.

Based on the experimental, analytic and numerical evidence discussed, the authors conjecture that systems of the form (1.1) are almost all trap free and that their convergence rate to a globally optimal control can be shown to imply that only a small fraction of the control space needs to be explored to discover such an optimal control. Yet, the fullest extent to which there is a free lunch in control remains to be resolved and stands as a challenge to assess.

Authors' contributions. The authors have each contributed equally to the work.
Competing interests. The authors feel there are no conflict of interest issues pertaining to this work.
Funding. B.R. thanks the NSF (grant no. CHE-1464569) and H.R. thanks the Templeton foundation (grant no. 52265) for funding.
Acknowledgements. The authors would like to thank Shanon Vuglar and Robert Kosut for helpful discussions.

Appendix A. Proof that assumptions are sufficient for a trap-free landscape

Given a control system of the form (1.1), here an informal proof is given that the three assumptions are sufficient for a trap-free landscape. Firstly, we assume that the cost function J is trap free itself. Secondly, we denote the manifold of states by M and the control space by C.

We will denote the endpoint map by V_T, which maps a control w to the corresponding solution to (1.1), and the overall fidelity of a control w to be $\mathcal{F}[w]=J(V_T[w])$. The overall derivative of fidelity with respect to the control can be obtained by the functional derivative chain rule

$$\frac{\delta\mathcal{F}}{\delta w}=\frac{\delta\mathcal{J}}{\delta x_T}\circ\frac{\delta x_T}{\delta w}. \tag{A 1}$$

Here the three assumptions are rephrased in geometric terms. In these terms assumption I, which is controllability, is that V_T is a surjective function $V_T: C\to M$ when all controls are permitted, which is itself exactly assumption III. Assumption II, which is sufficient local controllability, is that the image of the derivative (push forward) $\delta V_T|_{x_T}$ is not orthogonal to $\nabla J|_{x_T}$ for all controls $w\in C$. The three assumptions together yield two geometric properties of V_T and subsequently of $\mathcal{F}$. Firstly, $\forall g\in M$, $\exists w\in C$ s.t. $V_T[w]=g$. Secondly, there does not exist a control $w\in C$ s.t. $\langle \mathrm{d}V_T|_{x_T}[\delta w], \nabla J|_{x_T}\rangle=0$.

Examining equation (A 1), one sees that there are two types of critical points $\delta\mathcal{F}/\delta w=0$. The case $\delta J/\delta x_T=0$ is excluded by assumption that J does not possess traps of its own. The second case $\delta x_T/\delta w=0$ is excluded by assumption II—that $\langle \mathrm{d}V_T|_{x_T}[\delta w], \nabla J|_{x_T}\rangle=0$.

Assumptions I and II together imply that the image of V_T is the full space of states. The assumption II $\langle \mathrm{d}V_T|_{x_T}[\delta w], \nabla J|_{x_T}\rangle=0$ implies that the image of $\delta x_T/\delta w$ is not orthogonal to $\nabla J|_{x_T}$. Together these two statements imply that gradient ascent can continue until the objective is maximized, as away from the global maxima and minima of $\mathcal{F}$ the gradient is not zero, and the objective can be maximized due to controllability. Together these statements imply that gradient ascent can maximize J, and that the landscape is free from traps.

We further note that if J is known to possess saddles (as is the typical case in quantum control [60]), then these only translate into saddles on the control landscape rather than true traps.

Appendix B. Linear time-invariant systems are almost all controllable

One can readily check that, over all $A, \boldsymbol{b}$, the probability of selecting a non-controllable system is zero or, equivalently, that the set of $A, \boldsymbol{b}$ which corresponds to uncontrollable systems forms a null set. First note that the probability of selecting diagonalizable A is one when A is chosen at random. Thus, it will have distinct, non-zero eigenvalues from which it follows that

$$[\boldsymbol{b}, A\boldsymbol{b}, A^2\boldsymbol{b},\dots,A^N\boldsymbol{b}]=Q[\boldsymbol{c}, D\boldsymbol{c}, D^2\boldsymbol{c},\dots,D^N\boldsymbol{c}],$$

where $\boldsymbol{c}=Q^{-1}\boldsymbol{b}$. Now observe that the determinant of the controllability matrix (4.3) is equal to the determinant of the Vandermonde matrix generated by the eigenvalues of A (a_n say) and the matrix diag($\boldsymbol{c}$), i.e. $\det(V)\mathrm{diag}(\boldsymbol{c})=c_1\dots c_N\det(V)$. This determinant is clearly non-zero unless any a_n is zero, which itself happens with probability zero.

Appendix C. Equivalence of assumptions I and II for linear time-invariant systems

The variation of the endpoint map can be found directly from the LTI defining equation to also be an LTI equation,

$$\delta X(t) = A(\delta X)(t) + (\delta w)(t)\boldsymbol{b}.$$

We also note that the endpoint map for an LTI system, unlike almost all nonlinear systems, can be found in closed form,

$$X(T) = Ce^{tA}X(0) + \int_0^T w(t)e^{(T-t)A}\boldsymbol{b}\,\mathrm{d}t.$$

See any text book on LTI control systems for this derivation.

References

1. Wiener N. 1961 *Cybernetics or control and communication in the animal and the machine*. Cambridge, MA: MIT Press.
2. Wright S. 1932 The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proc. Sixth Int. Congress Genet.* **1**, 356–366.
3. Jones DS, Plank M, Sleeman BD. 2009 *Differential equations and mathematical biology*, 2nd edn. Chapman & Hall/CRC Mathematical and Computational Biology. Boca Raton, FL: CRC Press.
4. Roslund J, Rabitz H. 2009 Gradient algorithm applied to laboratory quantum control. *Phys. Rev. A* **79**, 053417. (doi:10.1103/PhysRevA.79.053417)
5. Sun Q, Pelczer I, Riviello G, Wu R-B, Rabitz H. 2015 Experimental observation of saddle points over the quantum control landscape of a two-spin system. *Phys. Rev. A* **91**, 043412. (doi:10.1103/PhysRevA.91.043412)
6. Rabitz H. 2012 Control in the sciences over vast length and time scales. *Quant. Phys. Lett.* **1**, 1–19.
7. Turinici G, Rabitz H. 2003 Wavefunction controllability for finite-dimensional bilinear quantum systems. *J. Phys. A Math. Gen.* **36**, 2565–2576. (doi:10.1088/0305-4470/36/10/316)
8. Weinacht TC, Ahn J, Bucksbaum PH. 1999 Controlling the shape of a quantum wavefunction. *Nature* **397**, 233–235. (doi:10.1038/16654)
9. Hsieh M, Wu R, Rabitz H. 2009 Topology of the quantum control landscape for observables. *J. Chem. Phys.* **130**, 104109. (doi:10.1063/1.2981796)
10. Hsieh M, Rabitz H. 2008 Optimal control landscape for the generation of unitary transformations. *Phys. Rev. A* **77**, 042306. (doi:10.1103/PhysRevA.77.042306)
11. Riviello G, Brif C, Long R, Wu R-B, Tibbetts KM, Ho T, Rabitz H. 2014 Searching for quantum optimal control fields in the presence of singular critical points. *Phys. Rev. A* **90**, 013404. (doi:10.1103/PhysRevA.90.013404)
12. Sun Q, Riviello G, Wu R, Rabitz H. 2015 Measuring the distance from saddle points and driving to locate them over quantum control landscapes. *J. Phys. A Math. Theor.* **48**, 465305. (doi:10.1088/1751-8113/48/46/465305)
13. Raj C, Herschel R. 2007 Quantum control landscapes. *Int. Rev. Phys. Chem.* **26**, 671–735. (doi:10.1080/01442350701633300)
14. Dong D, Petersen IR. 2010 Quantum control theory and applications: a survey. *IET Control Theory Appl.* **4**, 2651–2671. (doi:10.1049/iet-cta.2009.0508)
15. Sun Q, Pelczer I, Riviello G, Wu R-B, Rabitz H. 2014 Experimental exploration over a quantum control landscape through nuclear magnetic resonance. *Phys. Rev. A* **89**, 033413. (doi:10.1103/PhysRevA.89.033413)
16. De Fouquieres P, Schirmer SG. 2013 A closer look at quantum control landscapes and their implication for control optimization. *Infinite Dimens. Anal. Quant. Probab. Related Top.* **16**, 1350021. (doi:10.1142/S0219025713500215)
17. Rabitz H, Ho T, Long R, Wu R, Brif C. 2012 Comment on 'are there traps in quantum control landscapes?' *Phys. Rev. Lett.* **108**, 198901. (doi:10.1103/PhysRevLett.108.198901)
18. Pechen AN, Tannor DJ. 2011 Are there traps in quantum control landscapes? *Phys. Rev. Lett.* **106**, 120402. (doi:10.1103/PhysRevLett.106.120402)

rsta.royalsocietypublishing.org Phil. Trans. R. Soc. A **375**: 20160210

19. Wu R, Long R, Dominy J, Ho T, Rabitz H. 2012 Singularities of quantum control landscapes. *Phys. Rev. A* **86**, 013405. (doi:10.1103/PhysRevA.86.013405)
20. Russell B, Rabitz H, Wu R. 2016 Quantum control landscapes beyond the dipole approximation. (http://arxiv.org/abs/1602.06250)
21. Wu R, Pechen A, Rabitz H, Hsieh M, Tsou B. 2008 Control landscapes for observable preparation with open quantum systems. *J. Math. Phys.* **49**, 022108. (doi:10.1063/1.2883738)
22. Yang F, Cong S, Long R, Ho T-S, Wu R, Rabitz H. 2013 Exploring the transition-probability-control landscape of open quantum systems: application to a two-level case. *Phys. Rev. A* **88**, 033420. (doi:10.1103/PhysRevA.88.033420)
23. Moore KW, Pechen A, Feng X, Dominy J, Beltrani V, Rabitz H. 2011 Universal characteristics of chemical synthesis and property optimization. *Chem. Sci.* **2**, 417–424. (doi:10.1039/c0sc00425a)
24. Moore KW, Pechen A, Feng X, Dominy J, Beltrani VJ, Rabitz H. 2011 Why is chemical synthesis and property optimization easier than expected? *Phys. Chem. Chem. Phys.* **13**, 10048–10070. (doi:10.1039/c1cp20353c)
25. Wright S. 1984 *Evolution and the genetics of populations, volume 1: genetic and biometric foundations*. Chicago, IL: University of Chicago Press.
26. Feng X, Pechen A, Jha A, Wu R, Rabitz H. 2012 Global optimality of fitness landscapes in evolution. *Chem. Sci.* **3**, 900–906. (doi:10.1039/C1SC00648G)
27. Rabitz HA, Hsieh MM, Rosenthal CM. 2004 Quantum optimally controlled transition landscapes. *Science* **303**, 1998–2001. (doi:10.1126/science.1093649)
28. Rabitz H. 2003 Shaped laser pulses as reagents. *Science* **299**, 525–527. (doi:10.1126/science.1080683s)
29. Sun Q, Wu R-B, Ho T-S, Rabitz H. 2016 Inherently trap-free convex landscapes for full quantum optimal control. (https://arxiv.org/abs/1612.03988)
30. Rabitz H, Hsieh M, Rosenthal C. 2005 Landscape for optimal control of quantum-mechanical unitary transformations. *Phys. Rev. A* **72**, 052337. (doi:10.1103/PhysRevA.72.052337)
31. Rabitz HA, Hsieh MM, Rosenthal C. 2004 Quantum optimally controlled transition landscapes. *Science* **303**, 1998–2001. (doi:10.1126/science.1093649)
32. Hsieh M, Wu R, Rabitz H. 2009 Topology of the quantum control landscape for observables. *J. Chem. Phys.* **130**, 104109. (doi:10.1063/1.2981796)
33. Dominy J, Rabitz H. 2012 Dynamic homotopy and landscape dynamical set topology in quantum control. *J. Math. Phys.* **53**, 082201. (doi:10.1063/1.4742375)
34. Jurdjevic V. 1997 *Geometric control theory*. Cambridge Studies in Advanced Mathematics. Cambridge, UK: Cambridge University Press.
35. Altafini C. 2002 Controllability of quantum mechanical systems by root space decomposition of su(n). *J. Math. Phys.* **43**, 2051. (doi:10.1063/1.1467611)
36. Burgarth D, Arenz C, Gualdi G. 2014 Control of open quantum systems: case study of the central spin model. *New J. Phys.* **16**, 065023. (doi:10.1088/1367-2630/16/6/065023)
37. Thom R. 1954 Quelques propriétés globales des variétés différentiables. *Commentarii Mathematici Helvetici* **28**, 17–86. (doi:10.1007/BF02566923)
38. Russell B, Rabitz H, Wu R. 2016 Quantum control landscapes are almost always trap free. (http://arxiv.org/abs/1608.06198)
39. Moore KW, Chakrabarti R, Riviello G, Rabitz H. 2011 Search complexity and resource scaling for the quantum optimal control of unitary transformations. *Phys. Rev. A* **83**, 012326. (doi:10.1103/PhysRevA.83.012326)
40. Moore KW, Rabitz H. 2012 Exploring constrained quantum control landscapes. *J. Chem. Phys.* **137**, 134113. (doi:10.1063/1.4757133)
41. Hsieh M, Wu R, Rosenthal C, Rabitz H. 2008 Topological and statistical properties of quantum control transition landscapes. *J. Phys. B Atom. Mol. Opt. Phys.* **41**, 074020. (doi:10.1088/0953-4075/41/7/074020)
42. Donovan A, Rabitz H. 2015 Systematically altering the apparent topology of constrained quantum control landscapes. *J. Math. Chem.* **53**, 718–736. (doi:10.1007/s10910-014-0453-7)
43. Zhdanov DV, Seideman T. 2015 Role of control constraints in quantum optimal control. *Phys. Rev. A* **92**, 052109. (doi:10.1103/PhysRevA.92.052109)
44. Riviello G, Tibbetts KM, Brif C, Long R, Wu R, Ho T, Rabitz H. 2015 Searching for quantum optimal controls under severe constraints. *Phys. Rev. A* **91**, 043401. (doi:10.1103/PhysRevA.91.043401)

45. Antosiewicz HA. 1963 Linear control systems. *Arch. Ration. Mech. Anal.* **12**, 313–324. (doi:10.1007/BF00281230)
46. Kwakernaak H, Sivan R. 1972 *Linear optimal control systems*, vol. 1. New York, NY: Wiley Interscience.
47. Driels M. 1995 *Linear control systems engineering*. New York, NY: McGraw-Hill Higher Education.
48. Leiva H, Zambrano H. 1999 Rank condition for the controllability of a linear time-varying system. *Int. J. Control* **72**, 929–931. (doi:10.1080/002071799220669)
49. Fradkov AL. 2017 Horizons of cybernetical physics. *Phil. Trans. R. Soc. A* **375**, 20160223. (doi:10.1098/rsta.2016.0223)
50. Jean-Michel C, Andreea G, Catalin L, Gabriel T. 2009 Quantum control design by Lyapunov trajectory tracking for dipole and polarizability coupling. *New J. Phys.* **11**, 105034. (doi:10.1088/1367-2630/11/10/105034)
51. Lukes DL. 1972 Global controllability of nonlinear systems. *SIAM J. Control* **10**, 112–126. (doi:10.1137/0310011)
52. Sussmann HJ, Jurdjevic V. 1972 Controllability of nonlinear systems. *J. Differ. Equ.* **12**, 95–116. (doi:10.1016/0022-0396(72)90007-1)
53. Hirschorn RM. *Controllability in nonlinear systems*, pp. 205–214. Dordrecht, The Netherlands: Springer.
54. Sontag ED. 1991 *Kalman's controllability rank condition: from linear to nonlinear*, pp. 453–462. Berlin, Germany: Springer.
55. Zabczyk J. 2009 *Mathematical control theory: an introduction*. Modern Birkhäuser Classics. Boston, MA: Birkhäuser.
56. Davison EJ, Kunze EG. 1970 Some sufficient conditions for the global and local controllability of nonlinear time-varying systems. *SIAM J. Control* **8**, 489–497. (doi:10.1137/0308035)
57. Sussmann HJ. 1987 A general theorem on local controllability. *SIAM J. Control Optimiz.* **25**, 158–194. (doi:10.1137/0325011)
58. Wolpert DH, Macready WG. 1997 No free lunch theorems for optimization. *Trans. Evol. Comp.* **1**, 67–82. (doi:10.1109/4235.585893)
59. Wolpert DH, Macready WG. 1995 No free lunch theorems for search. Technical Report SFI-TR-95-02-010. Santa Fe Institute, Santa Fe, NM, USA.
60. Ho T, Dominy J, Rabitz H. 2009 Landscape of unitary transformations in controlled quantum dynamics. *Phys. Rev. A* **79**, 013422. (doi:10.1103/PhysRevA.79.013422)