Algorithms for closed loop control of quantum dynamics

Herschel Rabitz
Department of Chemistry
Princeton University
Princeton, New Jersey, U.S.A., 08544
E-Mail: hrabitz@princeton.edu

ABSTRACT.

Most quantum systems considered for control by external fields are plagued by a serious lack of complete information about the underlying Hamiltonian. Traditional feedback control techniques are generally not appropriate due to the latter problem, as well as the ultrafast nature of typical quantum dynamics phenomena and the fact that observations of the quantum system will inevitably lead to a disturbance which may often be contradictory to the desired control. In contrast, learning control techniques have a special role to play in the manipulation of quantum dynamics phenomena. The unique capabilities of quantum systems making them amenable to learning control are (a) the ability to have very large numbers of identical systems for submission to control, (b) the high duty cycle of laboratory laser controls, and (c) the ability to observe the impact of trial controls at ultrafast time scales. Various learning algorithms have been proposed to guide this control process. The present paper will discuss these proposals, as well as some new perspectives.

1. INTRODUCTION.

There is rapidly growing interest in the control of quantum systems for a variety of applications with atoms, molecules, and condensed phases[1]. Most of these applications aim at altering the material, either transiently or permanently changing its state. These tasks may be viewed as microworld engineering in a domain where the laws of physics are quantum mechanical, rather than classical mechanics, as in the traditional engineering disciplines. Most quantum mechanical cases under consideration employ tailored ultrafast laser electric fields $\varepsilon(t)$ as the control, with hopefully, the capability of driving the system from a chosen initial state $\psi(0)$ to a desired final state $\psi(T)$ at time T. As quantum systems naturally exhibit complex motions, it is exceedingly difficult to intuit the necessary structure of the control field $\varepsilon(t)$ capable of producing successful manipulations of the quantum system. Much work has gone into open loop optimal designs[1,2] for the laser fields $\varepsilon(t)$, based on the assumptions (1) that the system Hamiltonian H_0 is known to high precision, (2) that the multidimensional Schrödinger equation $i\hbar\frac{\partial\psi}{\partial t} = \left[H_0 - \mu\varepsilon(t)\right]\psi$ may be reliably solved (here, μ is the molecular dipole function, also assumed known), and (3) that the resultant control field design $\varepsilon(t)$ may be faithfully reproduced in the laboratory. Except for the simplest of quantum systems, all of these assumptions are seriously flawed, and this paper aims to address the nature of suitable closed loop learning algorithms for direct application in the laboratory to circumvent these problems. Despite the difficulties with theoretically designing quantum controls, much insight has been learned from such efforts, with the most important point being that successful control demands full cooperation between the temporal field structure $\varepsilon(t)$ and the dynamical capabilities of the quantum system.

The difficulties outlined above for controlling quantum systems are not unusual in the engineering disciplines, and a natural solution is to consider closed loop control in the laboratory. In this context, two forms may be considered: feedback control[3] and learning control[4]. Feedback control refers to a real-time process where the controller utilizes observed data on the state of the quantum system and suggests a control field update for immediate implementation in the laboratory. This process is generally difficult to implement in quantum systems for a variety of reasons, including that observation of the state will surely disturb the system, and Hamiltonians are rarely known accurately. On the other hand, quantum control has some especially attractive features when approached as a learning problem. First, it is often easy to produce many (even $\sim 10^{23}$) samples, with each starting in the same state, and the duty cycle of going from one control to the next can be many thousands or more per second[5]. The action of each control, in turn, can be observed by a second rapid probe, and the overall process repeated with a new sample at an exceedingly high rate under computer control[6]. Thus, one can envision performing thousands or more distinct learning control experiments on comfortable laboratory time scales. This circumstance is unprecedented and it offers an enormous advantage to the quantum system learning control process, where each excursion around the control loop involves a new sample. An enabling technical feature of quantum learning control is the ability to create rather arbitrarily shaped control fields $\varepsilon(t)$ through advanced laser pulse shaping technology[5]. Although the details of this process are not important for the purposes of this paper, the capability is significant, including the ability to switch from one control field to another at the rate of $\sim 10^3 - 10^6$/sec. The last key element of the control process is the nature of the learning algorithm in the closed loop, which is the focus of this paper. In principle, any of a variety of algorithms could be considered, but one requirement is that they be operationally fast, as they will be directly employed in the laboratory with real-time ongoing experiments.

Another critical element in the learning control process is the ability to measure the outcome resulting from a given control, which is usually expressed as the expectation value $\langle\psi(T)|O|\psi(T)\rangle$ of some operator O that is being considered for control at a terminal time T. Thus, one can envision a closed loop process, making repetitive excursions to identify a suitable control field to meet the particular objective expectation value. Under rather general conditions, the non-linear nature of controlled quantum dynamics suggests that there will also be multiple control fields $\varepsilon(t)$ giving comparable results. The overall learning control process aims to discover at least one of these good fields.

Section 2 of this paper will discuss the nature of the currently considered learning algorithms, and open up the prospect of a new family with special features. Section 3 will introduce a new non-linear learning algorithm. Finally, Section 4 presents brief concluding remarks.

2. QUANTUM LEARNING CONTROL ALGORITHMS.

The control of a quantum system is normally expressed[2] as an optimization problem by specification of a cost functional $J = J_Q + J_\varepsilon$. Here, J_Q refers to the costs associated with steering the quantum system as close as possible to a desired outcome for the target operator O, expectation value $\langle\psi(T)|O|\psi(T)\rangle$, and J_ε refers to any relevant costs associated with the magnitude or form of the control field. The particular structure of J_Q and J_ε are problem-dependent, and many cases may be considered. The goal is minimization of J with respect to the control $\varepsilon(t)$, which is a functional minimization problem. The field is often expressed in terms of its Fourier transform $\tilde{\varepsilon}(\omega)$, which is more easily amenable to manipulation in the laboratory[5]. At present, there is no information suggesting that the control of quantum systems lends itself to a special class of algorithms, and as such, a variety have been considered as summarized below.

(2.1) Gradient Algorithms. Nominally, the employment of gradient algorithms in laboratory optimization calls for caution, as the gradient $\delta J/\delta\varepsilon(t)$ (or its parametric analog, if the field is discretized in some fashion) must be measured in the laboratory where inherent errors will be introduced. Although laboratory noise of sufficient amplitude will make this approach inappropriate, a saving grace in quantum applications is the fact that a very large number of experiments may be performed with nominally the same variation $\delta\varepsilon(t)$ of the functional response J. Thus, the ensemble average $\langle\delta J/\delta\varepsilon(t)\rangle$ may be observed to drive down the noise contributions. This procedure has been explored in simulations[7], where it was shown to be quite effective. Observational error contributions to J can be readily averaged away (assuming they are random), but errors in the control field will generally not average away, as the quantum system observables can be highly nonlinear functionals of the control. Nevertheless, good robustness has been shown in simulations, and this technique deserves further analysis and attention.

(2.2) Genetic Algorithms. The use of genetic algorithms (GA) for quantum learning control[6,8] introduces a stochastic-type element in the search process, to more thoroughly explore for viable control field $\varepsilon(t)$ solutions. However, a potential difficulty with these algorithms is the need to work with large numbers of control variables, ranging in the hundreds to thousands, typically coming from a discretization of the field Fourier transform $\tilde{\varepsilon}(\omega)$ in keeping with how the fields are generated in the laboratory[5]. Thus, the optimization algorithm (genetic or otherwise) must deal with a very high dimensional space of control variables. Under normal conditions, GA's have been reported to critically slow down with such large numbers of variables. All of the current laboratory learning experiments in quantum systems have employed GA's involving approximately 100 generations and families of 50 members. Despite the very large number of control variables, good quality results have been obtained[9,10] with GA's. The reason for this success has yet to be explained, but it may occur because a high degree of readily identified cooperativity exists amongst the many control variables in $\tilde{\varepsilon}(\omega)$.

Regardless of the algorithm employed for quantum learning control, an ancillary desire is to learn as much as possible about the dynamical mechanism of control from the structure of the optimal field. While such an analysis will always be limited due to its inherent kinematic nature (i.e., the actual observable may draw on the field in a more complex fashion than evident from the kinematics alone due to the non-linear nature of quantum control), seeking such information is important. In order to ensure a reliable physical interpretation of the control in terms of mechanism, it is essential that the field $\varepsilon(t)$ or $\tilde{\varepsilon}(\omega)$ contain no extraneous features that have essentially no impact on the quantum observable[11]. In the context of GA's, this implies that "genetic pressure" needs to be applied against the turning of those quantum control knobs that have little impact on the control output. Simulations have shown dramatic field structure differences, arising from placing genetic pressure against extraneous field components[6].

(2.3) Linear Control Maps. Although quantum observables typically draw on the control field $\varepsilon(t)$ in a non-linear fashion, one may consider the generation of approximate input-output (IO) relationships based on a *sequence* of laboratory-generated linear mappings. Such a mapping procedure has been considered for engineering control[4], and it has a natural extension to the quantum mechanical domain[12]. As an illustration, the simplest case occurs with temporal experimental data $O(t) = \langle \psi(t) | O | \psi(t) \rangle$, $0 \le t \le T$. Thus, let $\varepsilon^s(t)$ be the current best estimate of the control field at the s-th iterative map level producing the data track $O^s(t)$ as an approximation to the optimal or desired track $O^d(t)$. As a further step towards the desired track, an improved field $\varepsilon^{s+1}(t) = \varepsilon^s(t) + \delta\varepsilon^s(t)$ can be obtained from the mapping

$$\delta O^s(t) = \int_0^t M^s(t,t')\delta\varepsilon^s(t')dt' \qquad (1)$$

where $\delta O^s(t) = O^d(t) - O^s(t)$. The correction of the control field $\delta\varepsilon^s(t)$ can be determined by inverting the relation in Eq. (1) to form

$$\delta\varepsilon^s(t') = \int_{t'}^T N^s(t',t)[O^d(t) - O^s(t)]dt \qquad (2)$$

In Eq. (2), the inverse map, $N^s(t',t)$, would generally be identified from the forward map $M^s(t,t')$ by singular value decomposition. This process would then be repeated to form $\varepsilon^{s+2}(t)$, etc., until acceptable convergence is reached approaching the desired track $O^d(t)$. The forward map $M^s(t,t')$ may be determined by inputting a set of trial fields $\varepsilon_i^s(t)$, $i = 1,2,\ldots$, and observing their impact $O_i^s(t)$, i = 1,2,…. The use of singular value decomposition in executing the linear mapping control process is an especially attractive feature, as it corresponds to projecting both the input control and the observable signal into lower dimensional spaces. This projection operation makes the determination of the inverse map $N^s(t',t)$ practical. At this juncture, quantum learning control with linear maps has been successfully tested under simulation for quantum systems[12].

3. NON-LINEAR MAPS FOR LEARNING CONTROL AND THE IDENTIFICATION OF COOPERATIVE CONTROL FEATURES.

Section 2 summarized the current algorithmic technology being considered for the control of quantum systems in the laboratory. Thus far, only genetic-type algorithms have been implemented[9,10] in the laboratory, and some very encouraging results have been obtained. However, there is no indication that a GA alone provides the best algorithm, especially for cases with large numbers of significant control variables. Secondly, GA's, when successful, only provide the control field, and not much insight into the physical basis behind the control field structure. A step towards addressing the latter important issue is evident in the learning algorithm involving the generation of linear maps discussed in subsection 2.3 However, the linearity in these maps prevents a proper identification of the cooperative role amongst the variables of the control field. Furthermore, the presence of very large numbers of control variables implies that better optimization procedures may be necessary to effectively find the controls. These concerns suggest that going to the next step of forming non-linear mappings may prove to be valuable for the field of quantum control[13]. This non-linear mapping approach may also be effectively combined with genetic or other types of algorithms, to perform global searches for the best possible solutions.

To appreciate the perspective involved with non-linear control maps, it is useful to consider the control field (either in the frequency or time domains) as broken into a set of N variables $\varepsilon \equiv (\varepsilon_1, \ldots, \varepsilon_N)$. The observable $O(\varepsilon) = O(\varepsilon_1, \ldots, \varepsilon_N) = \langle \psi(T) | O | \psi(T) \rangle$ can be viewed as a function over the N-dimensional space of input controls. The full structure of such spaces in quantum systems is not known; however, it expected that $O(\varepsilon)$ will be a reasonably behaved function over at least some local, or even rather global, domain in the space of controls $\varepsilon \equiv (\varepsilon_1, \ldots, \varepsilon_N)$. The number of variables N can be as large as $\sim 10^2 - 10^3$, thereby making this IO mapping of very high dimensions. The objective is to determine a series of maps $O^\ell(\varepsilon)$, ℓ = 1,2,…, where each is quantitatively valid over some domain $\mathcal{D}^\ell$ in the space of controls. Now, the search for an effective control may be expressed schematically as a march through a series of overlapping domains is depicted in Figure 1. The optimization in each domain $\mathcal{D}^\ell$ could be

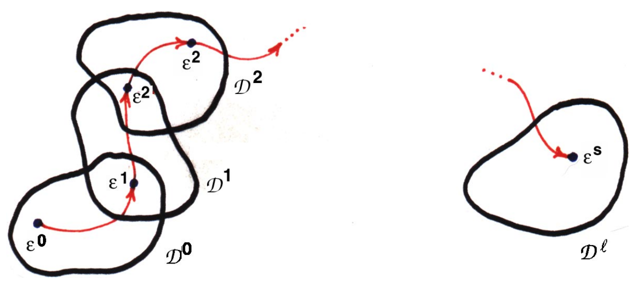

Figure 1. A sequence of non-linear mapping domains, where each of the control points $\varepsilon^\ell(t)$, ℓ = 1,2,… is locally optimal within the indicated domain $\mathcal{D}^\ell$.

performed by a GA or other means. Thus, the mapping problem reduces to determining how to sparsely sample the space $\varepsilon \equiv (\varepsilon_1, \ldots, \varepsilon_N)$ and utilize the resultant laboratory input-output data to produce a reliable map $O^\ell(\varepsilon)$ in the ℓ-th domain.

The analysis above produces a high dimensional function representation problem, which is common in many areas of science and engineering[14]. Such problems are conventionally viewed as being plagued with the curse of dimensionality, calling for the number of necessary experiments to grow exponentially with space dimension N. In practice, this very pessimistic scaling is not likely to occur, and certainly will not be case for relatively small domains. Arguments from the statistical analysis of many diverse problems suggests that even rather large domains may exist, where only a modest number of experiments (i.e., not growing exponentially with N) are sufficient to determine the function representation as a map. In developing such maps, it is useful to consider the degree of cooperativity amongst the control variables $\varepsilon \equiv (\varepsilon_1, ..., \varepsilon_N)$, and analogous statistical analysis problems, suggest that low order cooperativity may be dominant[14]. This perspective naturally leads to considering a hierarchical expansion of the output observable $O^\ell(\varepsilon)$ as

$$O^\ell(\varepsilon) = O_0^\ell + \sum_i O_i^\ell(\varepsilon_i) + \sum_{i<j} O_{ij}^\ell(\varepsilon_i, \varepsilon_j) + ... + O_{1,2,...,N}^\ell(\varepsilon_1, ..., \varepsilon_N) \quad (3)$$

Here, O_0^ℓ is the zero-th order constant mean map output, while $O_i^\ell(\varepsilon_i)$ is a first order independent, and likely non-linear, response of the map output to the action of the i-th control variable ε_i acting alone. Similarly, the term $O_{ij}^\ell(\varepsilon_i, \varepsilon_j)$ gives the second order response due to the non-separable cooperative role of the variables ε_i and ε_j. Finally, the last term $O_{1,2,...,N}^\ell(\varepsilon_1, ..., \varepsilon_N)$ closes the expression with any residual full N-dimensional cooperativity. A special form of this expansion is utilized in statistics for multivariate analysis[14], where it is typically found that terms only up to $\ell \simeq$ 2 or 3 are often significant. Building on this observation, the IO map expansion in Eq. (3) for each domain $\mathcal{D}^\ell$, ℓ = 1,2,… will be defined as adequate when taken to, at most, second order with pair cooperative terms. In turn, this restriction ultimately defines the extent of each domain, $\mathcal{D}^\ell$, where the corresponding map $O^\ell(\varepsilon)$ is reliable. Using the laboratory data on the left-hand of Eq. (3), it is necessary to determine the relevant individual terms on the right-hand side. This may be conveniently done by defining the component functions $O_0, O_i(\varepsilon_i), O_{ij}(\varepsilon_i, \varepsilon_j)$, to be fully orthogonal, both within and between each order through a suitable metric. Giving such a definition, then the minimum L^2 norm between the function on the left-hand side of Eq. (3) and its representation on the right-hand side gives the following unique solution:

$$O_{i_1,...,i_\ell}^\ell(\varepsilon_{i_1}, ..., \varepsilon_{i_\ell}) = \int_{\mathcal{D}^\ell} \prod_{r=1}^{\ell} \left[\delta(\varepsilon'_{i_r} - \varepsilon_{i_r})/\omega_{i_r}(\varepsilon'_{i_r}) - 1\right] \times \prod_{s=1}^{N} \omega_s(\varepsilon'_s) O(\varepsilon') d\varepsilon' \quad (4)$$

where the weights $\omega_s(\varepsilon_s)$, s = 1,2,…N are positive functions and $d\varepsilon'$ is the volume element in the domain $\mathcal{D}^\ell$. Equation (2) gives the component functions as projections on the observed laboratory data $O(\varepsilon)$. The choice of the weights is closely related to how the input space $\varepsilon \equiv (\varepsilon_1, ..., \varepsilon_N)$ is sampled, ranging from ordered to random sampling.

Given the ability to determine the component functions on the right-hand side of Eq. (3), in terms of the laboratory date, an overall optimal learning control algorithm may be set up, having the structure shown in Figure 2. Although there is extra overhead in determining

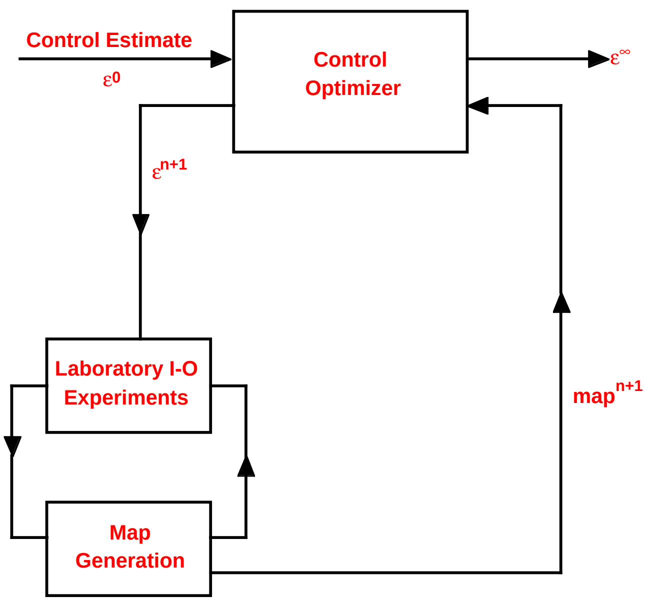

Figure 2. A block diagram indicating how input-output maps may be utilized to produce an overall control algorithm. The control optimizer could employ any of a number of routines, including genetic algorithms. The process may be self-starting or initiated by a control estimate ε^0.

the maps, they may be utilized a large number of times by the optimizer in the learning algorithm, as the maps can be exceedingly fast to evaluate[13]. Thus, it is anticipated that the extra overhead of determining the maps will be compensated for by obtaining a more thorough optimization. Furthermore, the structure of the maps reveals how the control variables act alone or cooperatively (i.e., by examining $O_i(\varepsilon_i)$ and $O_{ij}(\varepsilon_i, \varepsilon_j)$, respectively), to reach the objective. The latter ancillary benefit, in some cases, might even be the main reason for determining the maps, as control variable cooperativity information can give valuable clues about the control mechanism. Several simulations have been carried out with this algorithm on finite-dimensional quantum systems of up to 10 discrete states and 18 control variables in the frequency domain[13]. Excellent results were obtained with maps using up to either first or second order terms on the right-hand side of Eq. (3), with the allowed domain size $\mathcal{D}^\ell$ increasing for the second order maps. The maps were found to be physically valuable, as they provided information on the underlying control dynamics. In this regard, the story told by the maps, at times varied from that found by simple kinematic analysis of the field structure alone. That is, the non-linear relationship between the input control field and the output quantum observable can obscure the roles of input control variables. Only a proper non-linear analysis, as with the mapping procedure, will reveal the true role of the variables.

In the present illustrations with 18 input variables, a single evaluation of the map took approximately a microsecond on the computer. This high speed facilitated broad excursions over the variable space, including a full statistical analysis on the impact of possible errors in the laboratory controls. The algorithm awaits implementation in the laboratory, but it appears quite promising.

4. CONCLUSION.

This paper presented the current state of learning algorithm development for controlling quantum systems, with special emphasis given to a new algorithm based on non-linear maps in Section 3. The present simulations with the non-linear maps were based on input variables specified in the frequency domain $\tilde{\varepsilon}(\omega)$. Although this is natural from the perspective of the readily available controls in the laboratory[5], it may be better to re-express the map control variables in the time domain. The comparative advantages of operating in the temporal or frequency domains need careful examination. Besides utility in quantum mechanics, the algorithm presented in section 3 also should have application to broader control problems in engineering.

REFERENCES.

[1] P. Gaspard and I. Burghardt, eds., "Chemical Reactions and their Control on the Femtosecond Time Scale", *Adv. Chem. Phys.*, **101** (1997); W.S. Warren, H. Rabitz, M. Dahleh, *Science*, **259**, 1581 (1993).

[2] A.P. Peirce, M. Dahleh, and H. Rabitz, *Phys. Rev. A*, **37**, 4950 (1988); S. Shi and H. Rabitz, *Chem. Phys.*, **139**, 185 (1989); R. Kosloff, S.A. Rice, P. Gaspard, S. Tersigni, D.J. Tannor, *Chem. Phys.*, **139**, 201 (1989).

[3] D. Luenberger, *Introduction to Dynamical Systems: Theory, Models, and Applications* (Wiley, New York, 1979).

[4] K. More, *Iterative Learning Control for Deterministic Systems* (Springer, London, 1993).

[5] H. Kawashima, M.M. Wefers, and K.A. Nelson, *Ann. Rev. Phys. Chem.*, **46**, 627-656 (1995).

[6] R.S. Judson and H. Rabitz, *Phys. Rev. Lett.*, **68**, 1500 (1992).

[7] P. Gross, D. Neuhauser, and H. Rabitz, *J. Chem. Phys.*, **98**, 4557 (1993).

[8] G.J. Tóth, A. Lörincz and H. Rabitz, *J. Chem. Phys.*, **101**, 3715 (1994).

[9] A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, V. Seyfried, M. Strehle, G. Gerber, *Science*, **282**, 919 (1998).

[10] C.J. Bardeen, V.V. Yakovlev, K.R. Wilson, S.D. Carpenter, P.M. Weber, W.S. Warren, *Chem. Phys. Lett.,* **280**, 151 (1997).

[11] H. Rabitz and W. Zhu, Optimal Control of Molecular Motion: Design, Implementation, and Inversion, *Accts. of Chem. Res.*, submitted.

[12] M.Q. Phan and H. Rabitz, *J. Chem. Phys.*, **110**, 34 (1999).

[13] E. Weiss, J. Geremia, and H. Rabitz, Nonlinear Input-Output maps for Quantum Control, *J. Chem. Phys.*, submitted.

[14] H. Rabitz and Ö. Alis, *J. Math. Chem.*, **25**, 197 (1999).

ACKNOWLEDGMENTS.

The author acknowledges support from the Department of Defense (DoD) and the National Science Foundation (NSF). The author would like to thank Emily Weiss and John Geremia for many fruitful discussions and interactions on the development of the non-linear mapping algorithm

.