



HAL
open science

Stochastic learning control of inhomogeneous quantum ensembles

Gabriel Turinici

► **To cite this version:**

Gabriel Turinici. Stochastic learning control of inhomogeneous quantum ensembles. 2019. hal-02149374v2

HAL Id: hal-02149374

<https://hal.science/hal-02149374v2>

Preprint submitted on 9 Sep 2019 (v2), last revised 25 Nov 2019 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stochastic learning control of inhomogeneous quantum ensembles

Gabriel Turinici

IUF - Institut Universitaire de France

CEREMADE, Université Paris Dauphine - PSL Research University

September 9, 2019

Abstract

In quantum control, the robustness with respect to uncertainties in the system's parameters or driving field characteristics is of paramount importance and has been studied theoretically, numerically and experimentally. We test in this paper stochastic search procedures (Stochastic gradient descent and the Adam algorithm) that sample, at each iteration, from the distribution of the parameter uncertainty, as opposed to previous approaches that use a fixed grid. We show that both algorithms behave well with respect to benchmarks and discuss their relative merits. In addition the methodology allows to address high dimensional parameter uncertainty; we implement numerically, with good results, a 3D and a 6D case.

PACS number(s): 03.67.a, 02.30.Yy, 03.67.Pp

1 Introduction

Quantum control is a promising technology with many applications ranging from NMR [12] to quantum computing [15] and laser control of quantum dynamics [7]. The controlling field encounters many molecules which although identical in nature may interact differently with the incoming field because of e.g., different Larmor frequencies or rf attenuation factors (in NMR spin control or quantum computing, see [19, 29, 35, 22, 13, 17]), different spatial profile (see [24]) or other parameters (see [36, 8, 10]). For obvious practical

reasons, it is of paramount importance to ensure that the control quality is robust with respect to this heterogeneity. Thus the quantum control problem involves a unique set of driving fields $u(t) \in \mathbb{R}^L$, the same for all molecules in the ensemble, however each molecule is described by a set of parameters $\theta \in \Theta \subset \mathbb{R}^d$ and the control outcome depends on both u and θ ; the goal can be expressed as the maximization of the control quality averaged over θ . A different view is when the variability is not due to the presence of many different molecules but when uncertainties in the control implementation require to devise a field robust to fluctuations in those parameters.

A first natural question is whether this is at all possible, i.e., if a single field can drive several distinct molecules to a common target; the answer is given by the theory of ensemble control controllability, see [31, 5, 19, 4, 6] and is in general positive. However the theory does not explain how to find the control (except under specific regimes, see [2]). To do so, different algorithms have been proposed: the pseudo-spectral approach of Li et al. [20, 28, 21] consider spectral and/or polynomial representations of the control problem in 2D ($d = 2$); Wang considers iterative procedures based on sampling [32]; the learning approach of Chen et al. [8] and Kuang et al. [18] (the latter in the context of time-optimal control) consider a fixed uniform grid over the inhomogeneous parameter space and was tested for $d = 2$. Finally, Wu and al. [33] find robust controls using uniform grids in 2D and 3D ($d = 3$).

In all these works there is always a fixed grid (or fixed sampling) involved when the control is searched. The rationale behind this idea is that a fixed grid makes the search more stable and a good choice of the grid is enough to describe efficiently the mean performance of the control over the parameter space in the spirit of a quadrature formula for the average over θ . This is coherent with results from the approximation theory which inform that convergence is of order $e^{-\sqrt[d]{\mathcal{N}}}$, with respect to the number \mathcal{N} of grid points; however the same formula indicates a bad scaling with respect to d . To address this *curse of dimensionality* and also explore the nature of the search landscape, we take here a different view: at each control iteration we use a new sampling in the spirit of Monte Carlo methods (see [23, Section 7.7]) for computing high dimensional integrals. This will induce slight oscillations in the average but has the advantage to cover the space Θ of inhomogeneity even in high dimensions d . A similar approach has been tested independently in a very recent work by Wu and al. [34] for a two-dimensional example and promising results were obtained; see Section 2.2 for comments on the

differences between the two approaches. The procedure we propose is detailed in the next section and the numerical results are the object of Section 3.

2 Algorithms for ensemble quantum control

We consider a control $u(t) = (u_1(t), \dots, u_L(t)) \in \mathbb{R}^L$ acting on a molecule part of a larger ensemble. Each molecule is completely characterized by some inhomogeneity parameter $\theta \in \Theta \subset \mathbb{R}^d$ obeying a distribution law $P(\theta)$ on Θ (which can be the uniform distribution or any other). All molecules are subjected to the same control $u(t)$ during the time interval $[0, T]$ in order to reach some target.

2.1 Evolution equations

The dynamics of each molecule in the sample is governed by the Hamiltonian $H(\theta, u) = H_0(\theta) + \sum_{\ell=1}^L u_\ell(t)H_\ell(\theta)$ through the Schrödinger equation:

$$i \frac{d}{dt} \psi(t; \theta) = H(\theta, u) \psi(t; \theta), \quad (1)$$

where ψ is the wave-function of the molecule (here and below we set $\hbar = 1$). Of course, ψ depends on u but for notational convenience we omit to write explicitly this dependence from now on. Once a finite dimensional basis $\{|j\rangle, j = 1, \dots, N\}$ is chosen, the state of the quantum system can be represented as

$$|\psi(t; \theta)\rangle = \sum_{j=1}^N c_j(t; \theta) |j\rangle. \quad (2)$$

Denoting $C(t; \theta) = (c_0(t; \theta), \dots, c_N(t; \theta))^T$ the vector of coefficients C satisfies the equation:

$$\frac{d}{dt} C(t; \theta) = X(\theta, u) C(t; \theta), \quad (3)$$

where X is the representation of the Hamiltonian H (including the $1/i$ factor) in the basis $|j\rangle, j = 1, \dots, N$.

Note that same setting also applies to non-linear Hamiltonians e.g. Bose-Einstein condensates (nonlinearity in ψ), or high order control terms [11, 9] (nonlinearity in u).

The quantum system can also be described in terms of a density matrix $\rho(t; \theta)$; this matrix is expressed in some basis for operators. Same happens when the molecule is coupled to a bath or when relaxation phenomena are at work, see [1]; in both cases the coefficients of this expansion follow an equation similar to (3).

2.2 Optimization by stochastic gradient descent and Adam algorithms

The control goal is encoded as the minimization, with respect to u , of an error, or "loss" functional $\mathcal{L}(u, \theta)$ depending on the control u and the Hamiltonian parameters θ . When all the ensemble is considered, the following loss functional is to be minimized:

$$\mathcal{J}(u) = \int_{\Theta} \mathcal{L}(u, \theta) P(d\theta). \quad (4)$$

The stochastic optimization algorithms described below construct an iterative process in order to find the u that minimizes (4).

Historically the first to be considered, the stochastic gradient descent algorithm [26] (henceforth called SGD) consists in the following procedure:

Algorithm 1 SGD

- 1: Choose a learning rate $\alpha > 0$, a mini-batch size $M > 0$ and the initial control u^0 .
 - 2: Set iteration counter $k = 0$.
 - 3: **repeat**
 - 4: Draw M independent parameters $\theta_1^k, \dots, \theta_M^k$ from the distribution $P(\theta)$ and compute the approximation $g^k := \frac{1}{M} \sum_{m=1}^M \nabla_u \mathcal{L}(u^k; \theta_m^k)$ of the gradient $\nabla_u \mathcal{J}(u^k)$ of $\mathcal{J}(\cdot)$ at u^k .
 - 5: set $u^{k+1} = u^k - \alpha g^k$ and $k = k + 1$.
 - 6: **until** some stopping criterion is satisfied.
-

In order to accelerate the convergence of the SGD algorithm, several improvements have been proposed (see [27]) among which the Adam [16] variant which proved to be one of the most efficient and very scalable. The difference between Adam and SGD is that Adam uses a different learning rate for each parameter which is tuned as follows: when the uncertainty

in the gradient is large the learning rate is taken to be small and contrary otherwise. In order to have a robust estimation for the gradient (in absolute value) a Exponential Moving Average is computed on the fly (see below). It can be described as:

Algorithm 2 Adam

- 1: Choose the learning rate $\alpha > 0$, the EMA parameters β_1 and β_2 , the mini-batch size $M > 0$, the epsilon $\varepsilon > 0$ and the initial control u^0 .
 - 2: Set iteration counter $k = 0$, first moment estimate $\mu = 0$, second moment estimate $v = 0$.
 - 3: Set $k = k + 1$.
 - 4: **repeat**
 - 5: Draw M independent parameters $\theta_1^k, \dots, \theta_M^k$ from the distribution $P(\theta)$ and compute the approximation $g^k := \frac{1}{M} \sum_{m=1}^M \nabla_u \mathcal{L}(u^{k-1}; \theta_m^k)$ of the gradient $\nabla_u \mathcal{J}(u^{k-1})$ of $\mathcal{J}(\cdot)$ at u^{k-1} .
 - 6: Compute the moving averages $\mu^k := \beta_1 \mu^{k-1} + (1 - \beta_1) g^k$, $v^k := \beta_2 v^{k-1} + (1 - \beta_2) |g^k|^2$.
 - 7: Compute bias-corrected moment estimates: $\hat{\mu}^k = \mu^k / (1 - (\beta_1)^k)$, $\hat{v}^k = v^k / (1 - (\beta_2)^k)$.
 - 8: set $u^k = u^{k-1} - \alpha \hat{\mu}^k / (\sqrt{\hat{v}^k} + \varepsilon)$.
 - 9: **until** some stopping criterion is satisfied.
-

The momentum algorithm used in [34] can be seen as being halfway between SGD and Adam; it is formally a special case of the Adam algorithm for $\beta_1 = \lambda$, $\beta_2 = 1$, $v^0 = 1$ and no bias correction step 7 (that is $\hat{\mu}^k = \mu^k$, $\hat{v}^k = v^k$). In practice the numerical results are very similar and point in the same direction; in particular we expect that the momentum algorithm is also relevant to high dimensional robust control problems.

3 Numerical results

We test the performance of the algorithms in Section 2.2 for several benchmarks from the literature (or that generalize cases from the literature).

In sections 3.1 and 3.1 we compare the SGD algorithm with a fixed grid sampling method from the literature. Then in section sections 3.3 and 3.4 we compare the SGD with the Adam algorithm and in Section 3.5 we draw further conclusions concerning stochastic optimization.

In the situations considered below, the goal is to maximize the so-called *fidelity* denoted $\mathcal{F}(u; \theta)$. For sections 3.1 and 3.2 this has the formula $\mathcal{F}(u; \theta) = |\langle C(T; \theta), C_{\text{target}} \rangle|$ where C_{target} is a prescribed target state. But this expression is not differentiable everywhere and numerically it is easier to replace it with its square. Moreover, to express the problem as a minimization, a -1 multiplicative constant is introduced and 1 added to the result in order to have it positive. So the cost functional \mathcal{J} will be the mean, over $\theta \in \Theta$ of the error in the fidelity squared as in formula (8). On the contrary, when the fidelity is more well behaved as in section 3.3 where $\mathcal{F}(u; \theta) = c_4(T, \theta)$ or in section 3.4 where $\mathcal{F}(u; \theta) = c_6(T, \theta)$ the square operation is useless and the cost functional has the form in (11) or (13). However, in all sections, we will plot the error in the fidelity itself; the reason why not plotting the fidelity (instead of the error) is that the error can be very small (as in Section 3.1) and the results are more visible on a logarithmic scale. Note that in some cases the best control cannot attain the target with 100% quality (even for a single molecule). However, for any given value of the parameter θ , the best attainable performance is known (see [12, 30, 14]) and is denoted $F_{\text{max}}(\theta)$. We will therefore consider the fidelity relative to $F_{\text{max}}(\theta)$. In all cases the error is computed as the average over $M_{\text{test}} = 300$ random independent parameters $\theta_1^{\text{test}}, \theta_2^{\text{test}}, \dots, \theta_{M_{\text{test}}}^{\text{test}}$ drawn (once for all) from the distribution $P(\theta)$ and has the following expression:

$$\sum_{k=1}^{M_{\text{test}}} \left(1 - \frac{\mathcal{F}(u; \theta_k^{\text{test}})}{F_{\text{max}}(\theta_k^{\text{test}})} \right). \quad (5)$$

For sections 3.1 and 3.2 we will also plot the max relative error:

$$\max_{k=1, \dots, M_{\text{test}}} \left(1 - \frac{\mathcal{F}(u; \theta_k^{\text{test}})}{F_{\text{max}}(\theta_k^{\text{test}})} \right). \quad (6)$$

Finally, in order to compare our algorithm with those from the literature, we take as indicator of the numerical effort the number of gradient $\nabla_u \mathcal{L}(u; \theta)$ evaluations; for instance one iteration of SGD or Adam algorithms count as M gradient evaluations. In all situations we used for the Adam algorithm the standard values $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\varepsilon = 10^{-8}$.

3.1 Two level inhomogeneous ensemble

Consider an ensemble of spins as in [8, section III.]. The spins have different Larmor frequencies ω in the range $[0.8, 1.2]$ and the controls ($L = 2$) have multiplicative inhomogeneity $\epsilon \in [0.8, 1.2]$; we set $\theta = (\omega, \epsilon)$ and with the previous notations the dynamics corresponds to the equation:

$$\begin{pmatrix} \dot{c}_1(t; \theta) \\ \dot{c}_2(t; \theta) \end{pmatrix} = \begin{pmatrix} 0.5\omega i & 0.5\epsilon(u_2(t) - iu_1(t)) \\ -0.5\epsilon(u_2(t) - iu_1(t)) & -0.5\omega i \end{pmatrix} \begin{pmatrix} c_1(t; \theta) \\ c_2(t; \theta) \end{pmatrix}, \quad (7)$$

where c_1, c_2 are the coefficients of the wavefunction of the spin system in the canonical basis, as detailed in equation (2).

The initial state of each member of the quantum ensemble is set to $|\psi_0\rangle = |0\rangle$; i.e., $C_0 = (1, 0)^T$, and the goal is to reach the target state $|\psi_{\text{target}}\rangle = |1\rangle$; i.e., $C_{\text{target}} = (0, 1)^T$. The objective is encoded as the requirement to minimize:

$$J(u) = \frac{1}{2} \left(1 - \int_{\Theta} |\langle C(T; \theta), C_{\text{target}} \rangle|^2 P(d\theta) \right). \quad (8)$$

Here $F_{\text{max}}(\theta) = 1$. The total time is $T = 2$ is divided into $Q = 200$ time steps, of length $\Delta t = T/Q = 0.01$ each. The initial choice for the control u is $u^{k=0}(t) = \{u_1^0(t) = \sin t, u_2^0(t) = \sin t\}$.

Several mini-batch sizes $M = 1, 4, 8, 16$ and 32 are tested and compared with implementation in [8, section III.A.] where a 2D uniform grid of 5×5 values for θ is chosen. In all cases very good convergence results are attained. We plot in Figure 1 the results for $M = 1, M = 4$ relative to the convergence with the uniform 5×5 grid. In all cases ($M = 1, 4$, uniform grid) we set $\alpha = 500$; note that the learning rate α was optimized to obtain the best possible results for the fixed grid algorithm and indeed the results are better than those in [8, section III.A.]. But similar conclusions are reached for any value of α . An acceleration by a factor of 5 is obtained for both $M = 1$ and $M = 4$, essentially due to the fact that each SGD iteration uses only M gradient evaluations. Note that the SGD algorithm oscillates but these oscillations can be cured by lowering α (or stopping the search) as soon as a good result is obtained. The question of which is the best choice among $M = 1$ and $M = 4$ is a matter of striking a balance between speed and uncertainty: for $M = 4$ the convergence is slightly slower but oscillations are diminished. This behavior is observed, to a larger or lesser extent, in all test cases.

Note that in order to compare our learning rate α (for the fixed uniform grid) with that in [8, section III.A.] a multiplicative factor of $\Delta t/2$ has to

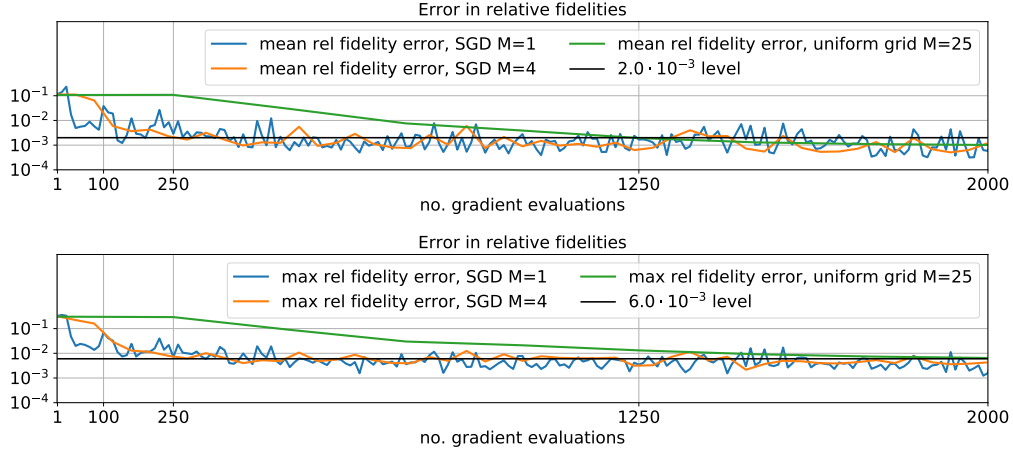


Figure 1: Convergence for the numerical case in Section 3.1. Top image: mean fidelity error (as defined in equation (5)). Bottom image: maximum (over the sample) fidelity error (as defined in equation (6)). We consider three simulations: a fixed uniform 2D grid ($M = 25$) as in [8, section III.A.] and the SGD algorithm with $M = 1$ and $M = 4$. This SGD converges about 5 times faster: the mean fidelity error of $2.0 \cdot 10^{-3}$ is obtained after 1250 gradient evaluations of the fixed grid algorithm and after 250 evaluations of the SGD algorithm with $M = 1, 4$. Same for other levels of errors.

be introduced because our gradient (see Appendix A) contains an extra Δt factor and the coefficient $1/2$. Thus one should transform $\alpha = 500$ to $1/2 * 0.01 * 500 = 2.5$ to compare with 0.2 used in [8].

3.2 A three level Λ atomic ensemble

In this section we test a Λ atomic ensemble from [8, Section IV] which can be written as a 3-level system with the following dynamics:

$$\begin{pmatrix} \dot{c}_1(t; \theta) \\ \dot{c}_2(t; \theta) \\ \dot{c}_3(t; \theta) \end{pmatrix} = \begin{pmatrix} -1.5\omega i & 0 & -i\epsilon u_2(t) \\ 0 & -\omega i & -i\epsilon u_1(t) \\ -i\epsilon u_2(t) & -i\epsilon u_1(t) & 0 \end{pmatrix} \begin{pmatrix} c_1(t; \theta) \\ c_2(t; \theta) \\ c_3(t; \theta) \end{pmatrix}, \quad (9)$$

where ω and ϵ have uniform distributions in $[0.8, 1.2]$ and c_1, c_2, c_3 are the coefficients of the wavefunction of the spin system in the canonical basis, as

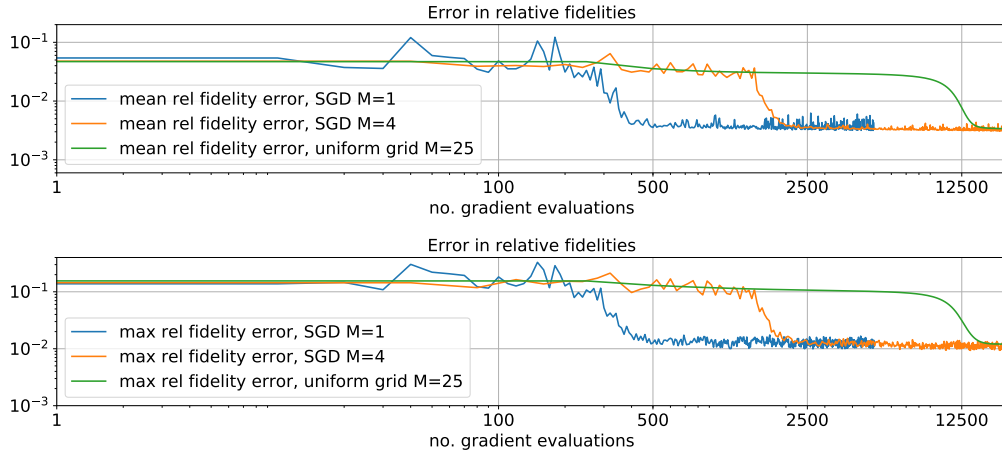


Figure 2: Convergence for the numerical case in Section 3.2. Top image: mean fidelity error – as defined in equation (5). Bottom image: maximum (over the sample) fidelity error (as defined in equation (6)). We consider two algorithms: a fixed uniform 2D grid ($M = 25$) as in [8, section IV.] and the SGD algorithm with $M = 1$ and $M = 4$. This latter approach converges about 7 times faster: the convergence settles in after 17'500 gradient evaluations of the fixed grid algorithm compared with cca. 2'500 evaluations of the SGD algorithm. This acceleration factor is even more important for $M = 1$ but at the price of larger oscillations.

detailed in equation (2).

The objective is to find a control $u(t) = (u_1(t), u_2(t))$ which drives all the inhomogeneous members from $|\psi_0\rangle = \frac{1}{\sqrt{3}}(|1\rangle + |2\rangle + |3\rangle)$ (i.e., $C_0 = (\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})$) to $|\psi_{\text{target}}\rangle = |3\rangle$ (i.e., $C_{\text{target}} = (0, 0, 1)$); the objective is encoded as the minimization of (8). Here $F_{\text{max}}(\theta) = 1$.

We plot in Figure 2 the results for $M = 1$ and $M = 4$ relative to the convergence with an uniform grid as in [8, section IV.]. In all cases ($M = 1, 4$, uniform grid) we set $\alpha = 100$. The acceleration factor is around 7 for $M = 4$ and even larger for $M = 1$ (but at the price of larger oscillations too).

3.3 A 3D example: two spin systems without cross-correlated relaxation

As argued before, methods from the literature may have difficulties to address high dimensional parameters, and often limit to two dimensional ($d = 2$) inhomogeneity (see however [33, Sec. V.B] for a 3D case). In order to test the full power of our method, we consider two situations that extend cases treated in the literature but have never been treated before. The first test is a three dimensional ($d = 3$) example which addresses the coherence transfer between two spins without cross-correlated relaxation, taken from [28, Section III.B.1. eq(15)] (but with an additional inhomogeneity dimension). An example of such a system is an isolated hetero-nuclear spin system composed of two coupled spins $1/2$ corresponding to atoms 1H and ^{15}N . For a general treatment of the relaxation terms and the formulation of this equation see [1]. The spins display control inhomogeneity described by the parameter ϵ as above but there is also variation in the relaxation rate and coupling constant, which, denoting $\theta = (\epsilon, J, \xi)$ results in the dynamical system:

$$\begin{pmatrix} \dot{c}_1(t; \theta) \\ \dot{c}_2(t; \theta) \\ \dot{c}_3(t; \theta) \\ \dot{c}_4(t; \theta) \end{pmatrix} = \begin{pmatrix} 0 & -\epsilon u_1(t) & 0 & 0 \\ \epsilon u_1(t) & -\xi & -J & 0 \\ 0 & J & -\xi & -\epsilon u_2 \\ 0 & 0 & \epsilon u_2 & 0 \end{pmatrix} \begin{pmatrix} c_1(t; \theta) \\ c_2(t; \theta) \\ c_3(t; \theta) \\ c_4(t; \theta) \end{pmatrix}. \quad (10)$$

Here c_1, c_2, c_3, c_4 are not any more related to relation (2) but are the coefficients of the density matrix of the two spins system (without cross-correlated relaxation) in a basis for operators which is composed of Kronecker products of Pauli operators for a each single system; the exact derivation of this equation is beyond the scope of this work, see [1, 12, 14] for details. On the other hand also note that the dynamics is not reversible (relaxation is present) and the equations do not correspond to a unitary evolution.

The inhomogeneity $\theta = (\epsilon, J, \xi)$ is uniformly distributed in $\Theta = [0.9, 1.1] \times [0.5, 1.5] \times [0, 2]$. The final time $T = 7\pi/6$ is discretized with $Q = 200$ uniform time steps. The control is initialized as before. The initial state is encoded as $c_0 = (1, 0, 0, 0)$ and the target is to minimize the three-dimensional integral:

$$\mathcal{J}(u) = 1 - \int_{\Theta} c_4(T; \theta) P(d\theta). \quad (11)$$

Recall that here the fidelity is $\mathcal{F}(u; \theta) = c_4(T, \theta)$; in this case (see [12, 14]) $F_{max}(\theta) = \sqrt{1 + (\xi/J)^2} - \xi/J$ (the worse performance being $-F_{max}(\theta)$).




Figure 3: Convergence for the numerical case in Section 3.3. The quantity plotted is given in equation (5). We set $M = 4$; for the SGD algorithm we choose $\alpha = 10.0$ and for the Adam algorithm we set $\alpha = 0.01$. The continuous (—) and dotted (·) curves stand for the mean fidelity errors of the plain SGD and Adam algorithm respectively; the convergence is similar and a 95% mean target relative fidelity (or equivalently 5% mean target relative fidelity error) is obtained. For the controls see Figure 4.

The results are in Figures 3 and 4. Note that although for each θ taken individually the figure $F_{max}(\theta)$ can be attained with a pair (recall $L = 2$) of suitable control fields, it is unknown whether a unique control pair exists ensuring 100% (relative to $F_{max}(\theta)$) target yield simultaneously for all $\theta \in \Theta$. In practice we did not find any, irrespective of the algorithm hyper-parameters such as α , the maximum number of iterations etc.; we conclude on one hand that this ensemble is not 100% simultaneously controllable and on the other hand that our procedure improves significantly the robustness of the control with respect to $\theta \in \Theta$ from an initial value of 67% up to 95%. Note that the results from the literature (which for this case only consider 2 dimensional inhomogeneity) do not obtain 100% control either (exact figure is not reported).

3.4 A 6D example: two spin systems with cross-correlated relaxation

We continue here to address new systems that previous methods could not treat. We consider an ensemble of two spin systems with cross-correlated relaxation as in [20, Section III.A.2.], [28, Section III.B.2 eq. (16)] and also [32, Example 3], [1].

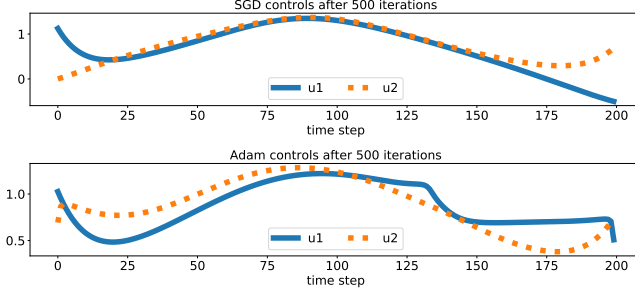


Figure 4: Converged controls for the SGD (up) and Adam (bottom) for the situation in in Section 3.3 (for the convergence see Figure 3). Controls obtained with the SGD algorithm are smoother than those from the Adam algorithm.

The spins display control inhomogeneity described by the parameters ϵ_1 and ϵ_2 and there is also variation in the auto-correlated relaxation rate ξ_a , the quotient ξ_c/ξ_a of the cross-correlation relaxation rate ξ_c with respect to the auto-correlated relaxation rate ξ_a and finally, a dispersion in the Larmor frequencies of each spin. Denoting $\theta = (\epsilon_1, \epsilon_2, \omega_1, \omega_2, \xi_a, \xi_c/\xi_a) \in \Theta = [0.9, 1.1]^2 \times [0, 1]^2 \times [0.75, 1.25] \times [0.7, 0.9]$, the dynamical system can be written:

$$\begin{pmatrix} \dot{c}_1(t; \theta) \\ \dot{c}_2(t; \theta) \\ \dot{c}_3(t; \theta) \\ \dot{c}_4(t; \theta) \\ \dot{c}_5(t; \theta) \\ \dot{c}_6(t; \theta) \end{pmatrix} = \begin{pmatrix} 0 & -\epsilon_1 u_1(t) & \epsilon_2 u_2(t) & 0 & 0 & 0 \\ \epsilon_1 u_1(t) & -\xi_a & \omega_1 & -J & -\xi_c & 0 \\ -\epsilon_2 u_2(t) & -\omega_1 & -\xi_a & -\xi_c & J & 0 \\ 0 & J & -\xi_c & -\xi_a & \omega_2 & -\epsilon_2 u_2(t) \\ 0 & -\xi_c & -J & -\omega_2 & -\xi_a & \epsilon_1 u_1(t) \\ 0 & 0 & 0 & \epsilon_2 u_2(t) & -\epsilon_1 u_1(t) & 0 \end{pmatrix} \begin{pmatrix} c_1(t; \theta) \\ c_2(t; \theta) \\ c_3(t; \theta) \\ c_4(t; \theta) \\ c_5(t; \theta) \\ c_6(t; \theta) \end{pmatrix}. \quad (12)$$

The vector $C = (c_1, \dots, c_6)$ has real entries and represents the density matrix of the two spin system expressed in a suitable basis of Kronecker products, similar to that in section 3.3, see [12, 14, 1] for details of the derivation of the model; the dynamics is not reversible (relaxation is present) nor unitary.

We set $J = 1$; the total time $T = 5$ is discretized with $Q = 200$ uniform time steps. The control is initialized as before. The initial state is encoded as $c_0 = (1, 0, 0, 0, 0, 0)$ and the target is to minimize the six-dimensional

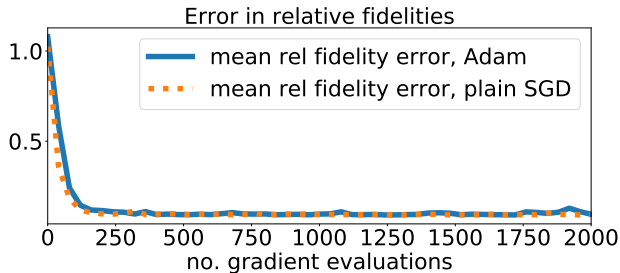


Figure 5: Convergence for the numerical case in Section 3.4. The quantity plotted is defined as in the Figure 3. We set $M = 4$; for the SGD algorithm we choose $\alpha = 10.0$ and for the Adam algorithm we set $\alpha = 0.01$. The continuous (—) and dotted (·) curves stand for the mean fidelity errors of the plain SGD and Adam algorithm respectively; the convergence is similar and 91% mean relative fidelity is obtained. For the controls see Figure 6.

integral:

$$\mathcal{J}(u) = 1 - \int_{\Theta} c_6(T; \theta) P(d\theta). \quad (13)$$

Recall that here the fidelity is $\mathcal{F}(u; \theta) = c_6(T, \theta)$. In this case too, the best attainable performance for a single molecule is known (see [12, 14]) and defined by $F_{max}(\theta) = \sqrt{1 + \eta^2} - \eta$ where $\eta = \sqrt{\frac{\xi_a^2 - \xi_c^2}{J^2 + \xi_c^2}}$.

The simulation results are in Figures 5 and 6. Same conventions are kept as in the previous section (fidelity is relative to maximum attainable figure) and same considerations still apply: 100% simultaneous controllability does not seem attainable but significant improvement in the robustness is obtained (91% up from -8%).

3.5 Stochastic convergence behaviors

The convergence of the stochastic algorithms can have two important regimes:

1. first, when all members of the ensemble can be simultaneously optimized to 100%; in our situation this is equivalent to simultaneous controllability. In this case convergence is "easier" because it is "enough" to follow the gradient for each parameter value in order to converge;

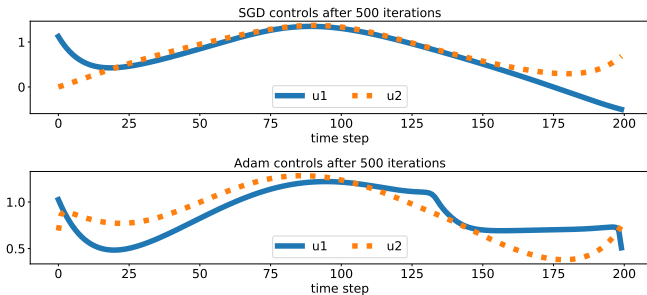


Figure 6: Converged controls for the SGD (up) and Adam (bottom) for the situation in in Section 3.4 (for the convergence see Figure 5). Controls obtained with the SGD algorithm are smoother than those from the Adam algorithm.

at convergence all gradients (as distribution with respect to ω), will collapse to (in practice will be close to) a Dirac mass.

2. secondly, when members of the ensemble cannot be simultaneously optimized; in this case, reaching full control for some θ value will harm the quality of some other parameter values $\theta' \neq \theta$. At convergence gradients will not be distributed as a Dirac mass any more, but the average with respect to theta will be zero (in practice small).

We illustrate this behavior in figures 7 and 8 where we plot the histograms of the gradient (with respect to the first field) $\nabla_{u_1(t)} \mathcal{J}(u(t_n), \theta)$ as random variables of θ at some time snapshots t . It is noticed that while in the first example it is possible to reduce significantly the gradient absolute value for all members of the sample (because simultaneous controllability holds true), in the second test case this reduction reaches a limit and the algorithm tries instead to center the gradients on zero so that the average be as low as possible.

4 Discussion and conclusion

We proposed and tested in this work a stochastic approach to compute the optimal controls of inhomogeneous quantum ensembles. The algorithms have been employed before in other areas of stochastic optimization but not tested

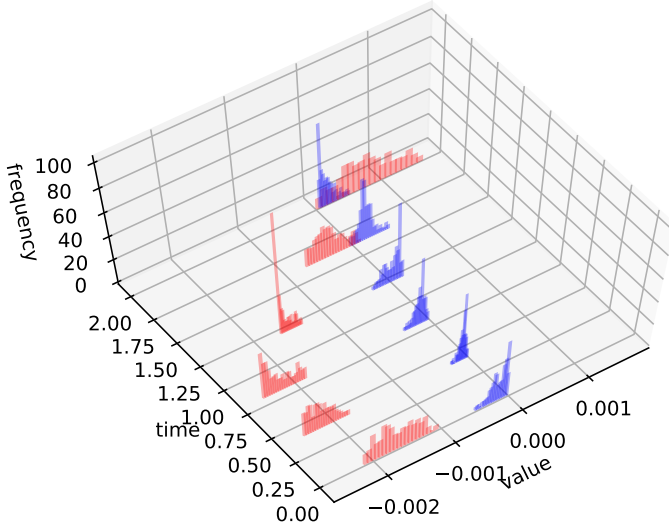


Figure 7: Histogram of the gradients $\nabla_{u_1(t)}\mathcal{J}(u(t_n),\theta)$ computed over the test sample $\theta_1^{test}, \theta_2^{test}, \dots, \theta_{M_{test}}^{test}$ (recall $M_{test} = 300$). Six time instants t are chosen uniformly in $[0, T]$: $t = 0, T/5, 2T/5, \dots, T$. In red are the gradients at $u = u^1$ (iteration $k = 1$) and in blue the gradients at $u = u^{500}$ (iteration $k = 500$). Here we consider the case in Section 3.1, see Figure 8 for the test case in Section 3.4.

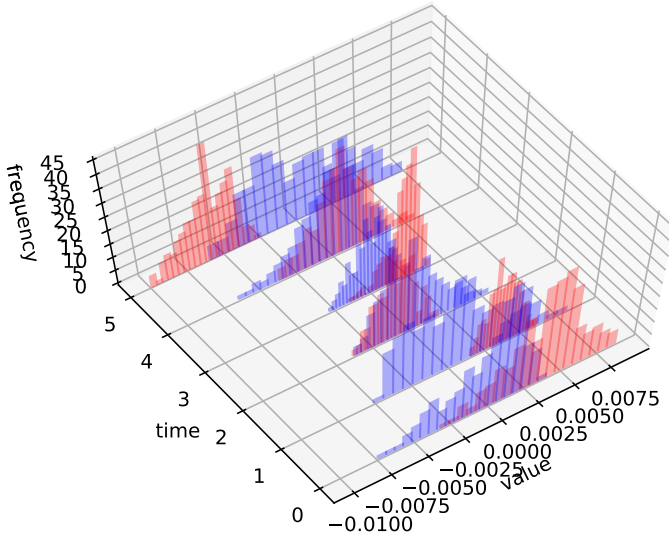


Figure 8: Histogram of the gradients as in Figure 7 except that here the results correspond to the test case in Section 3.4.

in this context (see [34] for similar algorithms). Their specificity is to draw at each iteration a new set of parameters from the inhomogeneous distribution. Although at first the intuition may not recommend such an approach, the numerical results indicate not only convergence but also faster convergence than methods based on fixed samples. In addition the method can address situations when the space of parameters is large and was tested successfully on a 6-dimensional example.

For lower dimensional examples (as in Sections 3.1 and 3.2) the acceleration of the stochastic algorithms (SGD, Adam) is due essentially to the lower effort per iteration compared to a fixed grid sampling (both being proportional to the number of samples used). In higher dimensions the fixed grid approach is inherently less efficient due to the *curse of dimensionality* and may even be prohibitively large.

On the other hand, compared with SGD, the Adam algorithm has the advantage to be more robust with respect to the choice of the learning rate α , but the controls are less regular.

Finally, one of the limitations of this work is to use constant learning rates. Variable learning rates are potentially interesting as it could speed up convergence in the initial phases by using large values of α and avoid oscillations in the end by lowering α . Several schedules are proposed in the stochastic optimization literature (inverse linear, piecewise constant, ...) but their analysis remains for future work.

*

A Gradient computation

We detail below the computation of the gradient for a single parameter θ , the general case being just a mean over θ . Consider the so-called adjoint state $\lambda(t; \theta)$; it is defined at the final time as the derivative of the outcome with respect to $C(T; \theta)$. For instance, for sections 3.1 - 3.2: $\lambda(T; \theta) = -\langle C_{\text{target}}, C(T, \theta) \rangle C_{\text{target}}$ while for sections 3.3 - 3.4 we set $\lambda(T; \theta) = -1$. Then for $t < T$, $\lambda(t; \theta)$ is the solution of the (backward) equation $\frac{d}{dt} \lambda(t; \theta) = X(t, \theta)^\dagger \lambda(t; \theta)$, where $X(t, \theta)^\dagger$ is the transpose conjugate of X when X has complex entries (examples 3.1 and 3.2) and reduces to the transpose when X is a real matrix (examples 3.3 and 3.4). Then $\nabla_{u(t)} \mathcal{J} = \langle \lambda(t; \theta), \frac{\partial X(t; \theta)}{\partial u(t)} C(t; \theta) \rangle$. In practice, given that u is discretized, the state C and the adjoint state λ are also discretized at time instants $t_n = n\Delta t$: $C_n(\theta) \simeq C(t_n; \theta)$, $\lambda_n(\theta) \simeq \lambda(t_n; \theta)$

which satisfy $C_{n+1}(\theta) = e^{\Delta t X(u(t_n); \theta)} C_n(\theta)$ and $\lambda_n(\theta) = e^{\Delta t X(u(t_n); \theta)^\dagger} \lambda_{n+1}(\theta)$ and the exact discrete gradient is $\nabla_{u(t_n)} \mathcal{J} = \langle \lambda_{n+1}(\theta), \frac{\partial e^{\Delta t X(u(t_n); \theta)}}{\partial u(t_n)} C_n(\theta) \rangle$.

Finally, in order to compute $\frac{\partial e^{\Delta t X(u(t_n); \theta)}}{\partial u(t_n)}$ we use a "divide and conquer" approach coupled with a 8-th order expansion as in [3, formula (11)] to obtain at the same time the exponential and the gradient ([25, Chapter VI]) from the knowledge of the inputs $X(u(t_n); \theta)$ and $\frac{\partial X(u(t_n); \theta)}{\partial u_k(t_n)}$.

References

- [1] Peter Allard, Magnus Helgstrand, and Torleif Hard. The complete homogeneous master equation for a heteronuclear two-spin system in the basis of cartesian product operators. *Journal of Magnetic Resonance*, 134(1):7 – 16, 1998.
- [2] Nicolas Augier, Ugo Boscain, and Mario Sigalotti. Adiabatic ensemble control of a continuum of quantum systems. *SIAM J. Control Optim.*, 56(6):4045–4068, 2018.
- [3] Philipp Bader, Sergio Blanes, and Fernando Casas. An improved algorithm to compute the exponential of a matrix. *arXiv e-prints*, page arXiv:1710.10989, Oct 2017.
- [4] Karine Beauchard, Jean-Michel Coron, and Pierre Rouchon. Controllability issues for continuous-spectrum systems and ensemble controllability of Bloch equations. *Comm. Math. Phys.*, 296(2):525–557, 2010.
- [5] Mohamed Belhadj, Julien Salomon, and Gabriel Turinici. Ensemble controllability and discrimination of perturbed bilinear control systems on connected, simple, compact Lie groups. *European Journal of Control*, 22(0):23 – 29, 2015.
- [6] Alfio Borzi, Gabriele Ciaramella, and Martin Sprengel. *Formulation and numerical solution of quantum control problems.*, volume 16. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2017.
- [7] Constantin Brif, Raj Chakrabarti, and Herschel Rabitz. Control of quantum phenomena: past, present and future. *New Journal of Physics*, 12(7):075008, 2010.

- [8] Chunlin Chen, Daoyi Dong, Ruixing Long, Ian R. Petersen, and Herschel A. Rabitz. Sampling-based learning control of inhomogeneous quantum ensembles. *Phys. Rev. A*, 89:023402, Feb 2014.
- [9] Jean-Michel Coron, Andreea Grigoriu, Cătălin Lifter, and Gabriel Turinici. Quantum control design by Lyapunov trajectory tracking for dipole and polarizability coupling. *New Journal of Physics*, 11(10):105034, oct 2009.
- [10] Hai-Jin Ding and Re-Bing Wu. Robust quantum control against clock noises in multiqubit systems. *Phys. Rev. A*, 100:022302, Aug 2019.
- [11] Claude M. Dion, Arne Keller, Osman Atabek, and André D. Bandrauk. Laser-induced alignment dynamics of HCN: Roles of the permanent dipole moment and the polarizability. *Phys. Rev. A*, 59:1382–1391, Feb 1999.
- [12] S. J. Glaser, T. Schulte-Herbrüggen, M. Sieveking, O. Schedletzky, N. C. Nielsen, O. W. Sørensen, and C. Griesinger. Unitary control in quantum ensembles: Maximizing signal intensity in coherent spectroscopy. *Science*, 280(5362):421–424, 1998.
- [13] I. N. Hincks, C. E. Granade, T. W. Borneman, and D. G. Cory. Controlling quantum devices with nonlinear hardware. *Phys. Rev. Applied*, 4:024012, Aug 2015.
- [14] Navin Khaneja, Burkhard Luy, and Steffen J. Glaser. Boundary of quantum evolution under decoherence. *Proceedings of the National Academy of Sciences*, 100(23):13162–13166, 2003.
- [15] Kaveh Khodjasteh and Lorenza Viola. Dynamically error-corrected gates for universal quantum computation. *Phys. Rev. Lett.*, 102:080501, Feb 2009.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. ICLR Proceedings 2015.
- [17] Robert L. Kosut, Matthew D. Grace, and Constantin Brif. Robust control of quantum gates via sequential convex programming. *Phys. Rev. A*, 88:052326, Nov 2013.

- [18] Sen Kuang, Peng Qi, and Shuang Cong. Approximate time-optimal control of quantum ensembles based on sampling and learning. *Physics Letters A*, 382(28):1858 – 1863, 2018.
- [19] J.-S. Li and N. Khaneja. Control of inhomogeneous quantum ensembles. *Phys. Rev. A*, 73:030302, 2006.
- [20] Jr-Shin Li, Justin Ruths, and Dionisis Stefanatos. A pseudospectral method for optimal control of open quantum systems. *The Journal of Chemical Physics*, 131(16):164110, 2009.
- [21] Jr-Shin Li, Justin Ruths, Tsyr-Yan Yu, Haribabu Arthanari, and Gerhard Wagner. Optimal pulse design in quantum control: A unified computational method. *Proceedings of the National Academy of Sciences*, 108(5):1879–1884, 2011.
- [22] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
- [23] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 2007.
- [24] Herschel Rabitz and Gabriel Turinici. Controlling quantum dynamics regardless of laser beam spatial profile and molecular orientation. *Physical review A: Atomic, Molecular and Optical Physics*, 75(4):043409, 2007.
- [25] Louis B. Rall. *Automatic Differentiation: Techniques and Applications*, volume 120 of *Lecture Notes in Computer Science*. Springer, Berlin, 1981.
- [26] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Ann. Math. Statistics*, 22:400–407, 1951.
- [27] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [28] Justin Ruths and Jr-Shin Li. A multidimensional pseudospectral method for optimal control of quantum ensembles. *The Journal of Chemical Physics*, 134(4):044128, 2011.

- [29] Thomas E Skinner, Timo O Reiss, Burkhard Luy, Navin Khaneja, and Steffen J Glaser. Application of optimal control theory to the design of broadband excitation pulses for high-resolution nmr. *Journal of Magnetic Resonance*, 163(1):8 – 15, 2003.
- [30] Gabriel Turinici and Herschel Rabitz. Quantum wavefunction controllability. *Chemical Physics*, 267(1–3):1 – 9, 2001.
- [31] Gabriel Turinici, Viswanath Ramakrishna, Baiqing Li, and Herschel Rabitz. Optimal discrimination of multiple quantum systems: controllability analysis. *Journal of Physics A: Mathematical and General*, 37(1):273, 2004.
- [32] Shuo Wang and Jr-Shin Li. Free-endpoint optimal control of inhomogeneous bilinear ensemble systems. *Automatica*, 95:306 – 315, 2018.
- [33] C. Wu, B. Qi, C. Chen, and D. Dong. Robust learning control design for quantum unitary transformations. *IEEE Transactions on Cybernetics*, 47(12):4405–4417, Dec 2017.
- [34] Re-Bing Wu, Haijin Ding, Daoyi Dong, and Xiaoting Wang. Learning robust and high-precision quantum controls. *Phys. Rev. A*, 99:042327, Apr 2019.
- [35] Bernard Yurke and John S. Denker. Quantum network theory. *Phys. Rev. A*, 29:1419–1437, Mar 1984.
- [36] Y. Zhang, M. Lapert, D. Sugny, M. Braun, and S. J. Glaser. Time-optimal control of spin 1/2 particles in the presence of radiation damping and relaxation. *The Journal of Chemical Physics*, 134(5):054103, 2011.