

State-Dependent Temperature Control for Langevin Diffusions¹

Xuefeng Gao², Zuo Quan Xu³, Xun Yu Zhou⁴

June 25, 2021

Abstract

We study the temperature control problem for Langevin diffusions in the context of non-convex optimization. The classical optimal control of such a problem is of the bang-bang type, which is overly sensitive to errors. A remedy is to allow the diffusions to explore other temperature values and hence smooth out the bang-bang control. We accomplish this by a stochastic relaxed control formulation incorporating randomization of the temperature control and regularizing its entropy. We derive a state-dependent, truncated exponential distribution, which can be used to sample temperatures in a Langevin algorithm, in terms of the solution to an HJB partial differential equation. We carry out a numerical experiment on a one-dimensional baseline example, in which the HJB equation can be easily solved, to compare the performance of the algorithm with three other available algorithms in search of a global optimum.

Keywords. Langevin diffusion, non-convex optimization, stochastic relaxed control, entropy regularization, Boltzmann exploration, HJB equation.

1 Introduction

Consider the problem of finding the global minimizer of a *non-convex* function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, where f is assumed to be differentiable. Traditional algorithms such as gradient descent often end up at a local optimum. The simulated annealing (SA) technique (Kirkpatrick et al., 1983) has been developed to resolve the problem of algorithms being trapped in local optima. The main thrust of an SA algorithm is *exploration via randomization*: At each

¹We thank the two reviewers and the Associate Editor for constructive comments which have led to an improved version of the paper.

²Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong, China. Email: xfgao@se.cuhk.edu.hk.

³Department of Applied Mathematics, The Hong Kong Polytechnic University, Hong Kong, China. Email: maxu@polyu.edu.hk.

⁴Department of Industrial Engineering and Operations Research and The Data Science Institute, Columbia University, New York, NY 10027, USA. Email: xz2574@columbia.edu.

iteration, the algorithm *randomly* samples a solution close to the current one and moves to it according to a (possibly time-dependent) probability distribution, which in the literature is mostly exogenous as a pre-defined schedule. This scheme allows for a more extensive search or exploration for the global optima with the risk of moving to worse solutions at some iterations. The risk is however controlled by slowly cooling down over time the “temperature” which is used to characterize the level of exploration or randomization.

The *Langevin algorithm* applies the SA technique by adding an independent series of Gaussian noises to the classical gradient descent algorithm, where the variance of the noises is linearly scaled by a sequence of temperature parameters $\{\beta_k\}$ which control the level of exploration/randomization. The continuous-time version of the Langevin algorithm is the so-called *overdamped* (or *first-order*) *Langevin diffusion* governed by a stochastic differential equation (SDE), where the temperature is now a stochastic process $\{\beta(t) : t \geq 0\}$. There is a large volume of literature on Langevin diffusions for studying non-convex optimization; see, to name just a few, [Chiang et al. \(1987\)](#); [Geman and Hwang \(1986\)](#); [Márquez \(1997\)](#). The advantage of studying a diffusion process instead of a discrete-time iteration lies in the simplicity and tractability of the former thanks to the many available analytical tools such as stochastic calculus, stochastic control and partial differential equation (PDE).

If we fix the temperature process in a Langevin diffusion to be a constant $\beta > 0$, then one can prove ([Chiang et al., 1987](#)) that, under some mild assumptions on f , the Langevin diffusion process converges to a unique stationary distribution whose density is a *Gibbs measure*, $\frac{1}{Z(\beta)} e^{-\frac{f(x)}{\beta}}$, where $Z(\beta)$ is the normalizing factor. As the temperature cools down (i.e. as $\beta \rightarrow 0$), this stationary distribution increasingly concentrates around the global minimizers of f which in turn justifies the use of the Langevin diffusion for non-convex optimization. Indeed, Langevin diffusions and their variants have recently found many applications in data science including large-scale Bayesian inference and stochastic non-convex optimization arising from machine learning; see e.g. [Dalalyan \(2017\)](#); [Gürbüzbalaban et al. \(2020\)](#); [Welling and Teh \(2011\)](#); [Raginsky et al. \(2017\)](#) and the references therein.

A critical question for solving global non-convex optimization using the Langevin diffusions is the design of the temperature schedule $\{\beta(t) : t \geq 0\}$. In the literature, the temperature $\beta(\cdot)$ is typically taken either as a constant (a hyperparameter) or various functions of time t , mostly *exogenously* given; see e.g. [Chiang et al. \(1987\)](#); [Gelfand and Mitter \(1991\)](#); [Geman and Hwang \(1986\)](#); [Holley et al. \(1989\)](#). When $\beta(t) \equiv \beta > 0$, it is well known that the expected transition time between different local minima for the

overdamped Langevin diffusion is exponential in $1/\beta$, a phenomenon known as metastability (Bovier et al., 2004, 2005). Raginsky et al. (2017) and Xu et al. (2018) further upper bound the time of Langevin dynamics converging to an approximate global minimum of non-convex functions, and Zhang et al. (2017) and Chen et al. (2020) analyze the hitting time of Langevin dynamics to a local minimum, where the temperature is all assumed to be a constant. On the other hand, Munakata and Nakamura (2001) formulate a deterministic optimal control problem to investigate optimal temperature schedules and derive an ordinary differential equation (ODE) for the time-dependent temperature $\beta(\cdot)$.

Ideally, in implementing an algorithm for finding the global optimizers of a function, the temperature should be fine tuned based on where the *current* iterate is; namely it should be a function of the state. For instance, in order to quickly escape a deeper local minimum, one should impose a higher temperature on the process. On the other hand, only lower temperatures are needed when the current process is not trapped in local minima in order to stay near a good solution. As a result of the need for this state-dependence, the temperature should be formulated as a *stochastic process* because the state itself follows a stochastic process. Fang et al. (1997) consider a temperature process as a specific, *exogenously* given increasing function of the current function value $f(x)$, so that the trajectory can perform larger steps in the search space when the current solution is seen to be far from optimal. They show numerically that this scheme yields more rapid convergence to global optimizers. Other exogenously devised state-dependent temperatures have also been shown to improve the rate of convergence of the Langevin samplers to the target Gibbs measure using Markov chain Monte Carlo (MCMC); see e.g. Girolami and Calderhead (2011); Ma et al. (2015). When target distributions have non-log-concave densities, several temperature heuristics including parallel tempering (Earl and Deem, 2005; Swendsen and Wang, 1986; Dupuis et al., 2012) and simulated tempering (Marinari and Parisi, 1992; Lee et al., 2018) have been proposed to overcome the multi-modality issue and improve the mixing time of the Langevin Monte Carlo method.

The goal of this paper is to develop a *theoretical* framework (instead of a heuristic approach) for designing *endogenous* state-dependent temperature schedule for Langevin diffusions in the context of non-convex optimization. It is natural to formulate an optimal stochastic control problem in which the controlled dynamic is the Langevin diffusion with the temperature process taken as the control, while the objective is to minimize a cost functional related to the original function f to be minimized. However, because the temperature appears linearly in the variance term of the diffusion, optimal control is of a

“bang-bang” nature, i.e. the temperature ought to repeatedly switch between two extreme points which are the lower and upper bounds of the temperature parameter. Such a bang-bang solution is not desirable in practice because it is *extremely* sensitive to errors: a tiny error may cause the control to take the *wrong* extreme point; see e.g. [Bertrand and Epenoy \(2002\)](#); [Silva and Trélat \(2010\)](#). To address this issue, we take the so-called *exploratory control* formulation, developed by [Wang et al. \(2019\)](#), in which classical controls are randomized and replaced by *distributions* of controls. This allows the control to take other values than just the extreme ones. Moreover, to encourage certain level of randomization, we explicitly incorporate the entropy of the distributions - which measures the extent of randomization - into the cost functional as a *regularization* term. This formulation smoothes out the temperature control and motivates the solution to deviate, if discreetly, from the overly rigid bang-bang strategy.

Entropy regularization is a widely used heuristic in reinforcement learning (RL) in discrete time and space; see e.g. [Haarnoja et al. \(2018\)](#). [Wang et al. \(2019\)](#) extend it to continuous time/space using a stochastic relaxed control formulation. The motivation behind [Wang et al. \(2019\)](#)’s formulation is *repeated learning* in unknown environments. Policy evaluation is achieved by repeatedly sampling controls from a same distribution and applying law of large numbers. In the present paper there is not really an issue of *learning* (as in reinforcement *learning*) because we can assume that f is a known function. The commonality that prompts us to use the same formulation of [Wang et al. \(2019\)](#) is *exploration* aiming at escaping from possible “traps”. Exploration is to broaden search in order to get rid of over-fitted solutions in RL, to jump out of local optima in SA, and to deprive of too rigid schedules in temperature control.

Under the infinite time horizon setting, we show that the optimal exploratory control is a truncated exponential distribution whose domain is the pre-specified range of the temperature. The parameter of this distribution is state dependent, which is determined by a nonlinear elliptic PDE in the general multi-dimensional case and an ODE in the one-dimensional case. The distribution is a continuous version of the *Boltzmann* (or *softmax*) *exploration*, a widely used heuristic strategy in the discrete-time RL literature ([Bridle, 1990](#); [Sutton and Barto, 2018](#)) that applies exponential weighting of the state-action value function (a.k.a. *Q-function*) to balance exploration and exploitation. In our setting, the *Q-function* needs to be replaced by the (generalized) *Hamiltonian*. This, however, is not surprising because in classical stochastic control theory optimal feedback control is to *deterministically* maximize the Hamiltonian ([Yong and Zhou, 1999](#)).

As discretization of the optimal state process in our exploratory framework, which satisfies an SDE, naturally leads to a Langevin algorithm with state-dependent noise, our results have algorithmic implications whenever the HJB equations involved can be easily solved. In particular, the HJB equation in our setting is an ODE when the state space is one-dimensional, which can be efficiently solved. For numerical demonstration, we compare the performance of this new algorithm with three other existing algorithms for the global minimization of a one-dimensional baseline non-convex function. The first method is a Langevin algorithm with a constant temperature (Raginsky et al., 2017), the second one is a Langevin algorithm where the temperature decays with time in a prescribed power-law form (Welling and Teh, 2011; Neelakantan et al., 2015), and the last one is a replica exchange algorithm (Dong and Tong, 2020). The experiment indicates that, at least in the one-dimensional case, our Langevin algorithm with state-dependent noise finds the global minimizer faster and outperforms the other three methods.

We must, however, emphasize that the main contribution of this paper is *theoretical*, rather than *algorithmic*. It establishes and develops a theoretical framework for studying state-dependent temperature control in SA, in which a non-convex optimization problem is connected to an HJB PDE. Among other theoretical interests in this connection, one observation is that solving the HJB requires *global* information on the underlying function f , whereas most existing SA algorithms use only *local* information of f at any given iteration. Intuitively, the former is more advantageous other things being equal. This is actually demonstrated by the outperformance of our algorithm in the one-dimensional case when solving the HJB equation poses no numerical challenge. The insight is that one should try to make use of the global information of f as much as possible. That said, for high dimensional non-convex optimization, we are not advocating our theory for actual algorithmic implementation before the curse of dimensionality for PDEs has been resolved. Recently there has been some encouraging progress using deep neural networks to numerically and efficiently solve high-dimensional PDEs (Han et al., 2018; Beck et al., 2019); so hopefully our results could also contribute to devising SA algorithms in the future when that line of research has come to full fruition.

The rest of the paper proceeds as follows. In Section 2, we describe the problem motivation and formulation. Section 3 presents the optimal temperature control. In Section 4, we report numerical results comparing the performance of our algorithm with three other methods. Finally, we conclude in Section 5.

2 Problem Background and Formulation

2.1 Non-convex optimization and Langevin algorithm

Consider a non-convex optimization problem:

$$\min_{x \in \mathbb{R}^d} f(x), \quad (2.1)$$

where $f: \mathbb{R}^d \rightarrow [0, \infty)$ is a continuously differentiable, *non-convex* function. The Langevin algorithm aims to have global convergence guarantees and has the following iterative scheme:

$$X_{k+1} = X_k - \eta f_x(X_k) + \sqrt{2\eta\beta_k}\xi_k, \quad k = 0, 1, 2, \dots, \quad (2.2)$$

where f_x is the gradient of f , $\eta > 0$ is the step size, $\{\xi_k\}$ is i.i.d Gaussian noise and $\{\beta_k\}$ is a sequence of the temperature parameters (also referred to as a cooling or annealing schedule) that typically decays over time to zero. This algorithm is based on the discretization of the overdamped Langevin diffusion:

$$dX(t) = -f_x(X(t))dt + \sqrt{2\beta(t)}dW(t), \quad X(0) = x \quad (2.3)$$

where x is an initialization, $\{W(t) : t \geq 0\}$ is a standard d -dimensional Brownian motion with $W(0) = 0$, and $\{\beta(t) : t \geq 0\}$ is an adapted, nonnegative stochastic process, both defined on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}; \{\mathcal{F}_t\}_{t \geq 0})$ satisfying the usual conditions.

When $\beta(t) \equiv \beta > 0$, under some mild assumptions on f , the solution of (2.3) admits a unique stationary distribution with the density $\pi(x) \propto e^{-\frac{1}{\beta}f(x)}$. Moreover, Raginsky et al. (2017) show that for a finite $\beta > 0$,

$$\mathbb{E}_{X \sim \pi} f(X) - \min_{x \in \mathbb{R}^d} f(x) \leq \mathcal{I}(\beta) := \frac{d\beta}{2} \log \left(\frac{eM}{m} \left(\frac{b}{d\beta} + 1 \right) \right), \quad (2.4)$$

where $\mathbb{E}_{X \sim \pi} f(X) := \int_{\mathbb{R}^d} f(x) \pi(dx)$, M, m, b are constants associated with f . It is clear that $\mathcal{I}(\beta) \rightarrow 0$ when $\beta \rightarrow 0$.

Our problem is to control the temperature process $\{\beta(t) : t \geq 0\}$ so that the performance of the continuous Langevin algorithm (2.3) is optimized. We measure the performance using the expected total discounted values of the iterate $\{X(t) : t \geq 0\}$, which is $\mathbb{E} \int_0^\infty e^{-\rho t} f(X(t)) dt$ where $\rho > 0$ is a discount factor. The discounting puts smaller weights on function values that happen in later iterations; hence in effect it dictates the computational budget (i.e. the number of iterations budgeted) to run the

algorithm (2.3). Clearly, if this performance or cost functional value (which is always nonnegative) is small, then it implies that, on average, the algorithm strives to achieve smaller values of f over iterations, and terminates at a budgeted time T with an iterate $X(T)$ that is close to the global optimum.

Mathematically, given $\rho > 0$, an arbitrary initialization $X(0) = x \in \mathbb{R}^d$ and the range of the temperature $U = [a, 1]$ where $0 < a < 1$, we aim to solve the following stochastic control problem where the temperature process is taken as the control:

$$\begin{aligned} & \text{Minimize } \mathbb{E} \int_0^\infty e^{-\rho t} f(X(t)) dt, \\ & \text{subject to } \begin{cases} \text{equation (2.3),} \\ \{\beta(t) : t \geq 0\} \text{ is adapted, and } \beta(t) \in U \text{ a.e. } t \geq 0, \text{ a.s.} \end{cases} \end{aligned} \quad (2.5)$$

Remark 1 (Choice of the range of the temperature). Naturally the temperature has a nonnegative lower bound. We suppose that it also has an upper bound, assumed to be 1 without loss of generality. In the Langevin algorithm and SA literature, one usually uses a determinist temperature schedule $\beta(t)$ that is either a constant or decays with time t . Hence there is a natural upper bound which is the initial temperature. This quantity is generally problem dependent and can be a hyperparameter chosen by the user. The upper bound of the temperature should be tuned to be reasonably large for the Langevin algorithm to overcome all the potential barriers and to avoid early trapping into a bad local minimum.

2.2 Solving problem (2.5) classically

Using a standard dynamic programming argument, see, e.g. [Yong and Zhou \(1999\)](#), we can write down the Hamilton–Jacobi–Bellman (HJB) equation associated with the problem (2.5):

$$-\rho v(x) + f(x) + \min_{\beta \in [a, 1]} [\beta \text{tr}(v_{xx}(x)) - f_x(x) \cdot v_x(x)] = 0, \quad x \in \mathbb{R}^d, \quad (2.6)$$

where $\text{tr}(A)$ denotes the trace of a square matrix A , and “ $x \cdot y$ ” the inner product between two vectors x and y .

The standard verification theorem in stochastic control theory ([Yong and Zhou, 1999](#), Chapter 5) yields that an optimal *feedback* control policy should achieve the minimum in the above equation. However, the term inside the min operator is linear in the control variable β ; hence the optimal policy has the following bang-bang form: $\beta^*(x) = 1$ if $\text{tr}(v_{xx}(x)) < 0$, and $\beta^*(x) = a$ otherwise. In economics terms, intuitively, the value $v(x)$

can be regarded as the disutility of the resource x (here, it is *disutility*, instead of utility, because the objective is to *minimize*). When $\text{tr}(v_{xx}(x)) < 0$, v is locally concave around x suggesting a risk-seeking preference. Hence one should randomize at the maximum possible level. A symmetric intuition applies to the opposite case when $\text{tr}(v_{xx}(x)) > 0$. This temperature control scheme shows that one should in some states heat at the highest possible temperature, while in other states cool down completely, depending on the sign of $\text{tr}(v_{xx}(x))$. This is *theoretically* the optimal scheme; but practically it is just too *rigid* to achieve good performance as it concentrates on two actions only, thereby leaves no room for errors which are inevitable in real life. This motivates the introduction of relaxed controls and entropy regularization in order to *smooth out* the temperature process in the next subsection.

Remark 2 (Bang-bang control and parallel/simulated tempering). The bang-bang control policy bears some resemblance to the so-called parallel tempering (or replica exchange) and simulated tempering in the MCMC literature ([Tawn et al., 2020](#)), although there are important differences. Both tempering algorithms swap processes with different temperatures based on certain mechanisms in order to sample the desired target distribution; so there is switching at work between different pre-specified temperatures like the bang-bang control. However, these algorithms are mostly designed for probabilistic *sampling* with Metropolis–Hastings rule used to accept or reject a swap between processes, while our focus is (non-convex) *optimization*. Moreover, the tempering algorithms use state space augmentation. Specifically, parallel tempering runs $n \geq 2$ copies of Langevin dynamics with different constant temperature parameters; hence the state space is \mathbb{R}^{dn} . Simulated tempering treats the temperature $\{\beta(t) : t \geq 0\}$ as a stochastic process taking only finite values and augments it to the original dynamics so that the state space is $(d + 1)$ –dimensional. In contrast, if we apply the bang-bang control, we obtain only one copy of the Langevin dynamics where the temperature parameter is state-dependent which is not a constant. Recently, [Dong and Tong \(2020\)](#) propose to use the replica exchange algorithm for non-convex optimization. In the numerical experiments reported in [Section 4](#), we will compare the performance of our algorithm with that of [Dong and Tong \(2020\)](#).

2.3 Relaxed control and entropy regularization

We now present our entropy-regularized relaxed control formulation of the problem. Instead of a classical control $\{\beta(t) : t \geq 0\}$ where $\beta(t) \in U = [a, 1]$ for $t \geq 0$, we consider a *relaxed* control $\pi = \{\pi(t, \cdot) : t \geq 0\}$, which is a *distribution* (or randomization) of classical controls over the control space U where a temperature $\beta(t) \in U$ can be sampled from this distribution whose probability density function is $\pi(t, \cdot)$ at time t .

Specifically, we study the following *entropy-regularized stochastic relaxed control problem*, also termed as *exploratory control problem*, following Wang et al. (2019):

$$V(x) = \inf_{\pi \in \mathcal{A}(x)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} f(X^\pi(t)) dt - \lambda \int_0^\infty e^{-\rho t} \int_U -\pi(t, u) \ln \pi(t, u) du dt \right], \quad (2.7)$$

where $x \in \mathbb{R}^d$, the term $\int_U -\pi(t, u) \ln \pi(t, u) du$ is the differential entropy of the probability density function $\pi(t, \cdot)$ of the randomized temperature at time t , $\lambda > 0$ is a weighting parameter for entropy regularization, $\mathcal{A}(x)$ is the set of admissible distributional controls to be precisely defined below, and the state process X^π is governed by

$$dX^\pi(t) = -f_x(X^\pi(t))dt + \tilde{\sigma}(\pi(t))dW(t), \quad X^\pi(0) = x, \quad (2.8)$$

where

$$\tilde{\sigma}(\pi) := \sqrt{\int_U 2u\pi(u)du}. \quad (2.9)$$

The form of $\tilde{\sigma}(\pi)$ follows from the general formulation presented in Wang et al. (2019). The function V is called the (optimal) *value function*.

Denote by $\mathcal{P}(U)$ the set of probability density functions over U , and by $\mathfrak{B}(U)$ the Borel algebra on U .

Definition 3. We say a density-function-valued stochastic process $\pi = \{\pi(t, \cdot) : t \geq 0\}$, defined on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}; \{\mathcal{F}_t\}_{t \geq 0})$ along with a standard d -dimensional $\{\mathcal{F}_t\}_{t \geq 0}$ -adapted Brownian motion $\{W(t) : t \geq 0\}$, is an *admissible (distributional) control*, denoted simply by $\pi \in \mathcal{A}(x)$, if

- (i) For each $t \geq 0$, $\pi(t, \cdot) \in \mathcal{P}(U)$ a.s.;
- (ii) For any $A \in \mathfrak{B}(U)$, the process $(t, \omega) \mapsto \int_A \pi(t, u, \omega) du$ is \mathcal{F}_t -progressively measurable;

- (iii) The SDE (2.8) admits solutions on the same filtered probability space, whose distributions are all identical.

So an admissible distributional control is defined in the *weak* sense, namely, the filtered probability space and the Brownian motion are also *part* of the control. Condition (iii) ensures that the performance measure in our problem (2.7) is well defined even under the weak formulation. For notational simplicity, however, we will henceforth only write $\pi \in \mathcal{A}(x)$.

Remark 4 (Choice of the weighting parameter λ). The weighting parameter λ dictates a trade-off between optimization (exploitation) and randomization (exploration), and is also called a temperature parameter in Wang et al. (2019). However, this temperature λ is different from the temperature β in the original Langevin diffusion. These are temperatures at *different* levels: β is at a lower, local level related to injecting Gaussian noises to the gradient descent algorithm, whereas λ is at a higher, global level associated with adding (not necessarily Gaussian) noises to the bang-bang control. It is certainly an open interesting question to endogenize λ , which is yet another temperature control problem at a different level. In implementing the Langevin algorithm, however, λ can be fine tuned as a hyperparameter.

3 Optimal Temperature Control

3.1 A formal derivation

To solve the control problem (2.7), we apply Bellman's principle of optimality to derive the following HJB equation:

$$-\rho v(x) - f_x(x) \cdot v_x(x) + f(x) + \inf_{\pi \in \mathcal{P}(U)} \int_U [\text{tr}(v_{xx}(x))u + \lambda \ln \pi(u)] \pi(u) du = 0. \quad (3.1)$$

The verification approach then yields that the optimal distributional control can be obtained by solving the minimization problem in π in the HJB equation. This problem has a constraint $\pi \in \mathcal{P}(U)$, which amounts to

$$\int_U \pi(u) du = 1, \quad \pi(u) \geq 0 \quad \text{a.e. on } U. \quad (3.2)$$

For any $k \in \mathbb{R}$,

$$\int_U [\text{tr}(v_{xx}(x))u + \lambda \ln \pi(u)] \pi(u) du = \int_U [\text{tr}(v_{xx}(x))u + \lambda \ln \pi(u) + k] \pi(u) du - k.$$

The integrand on the right hand side is a convex function of $\pi(u)$; so the first-order condition gives its unique global minimum at

$$\begin{aligned} \bar{\pi}(u; x) &= \exp \left(-\frac{1}{\lambda} [\text{tr}(v_{xx}(x))u] - \frac{k}{\lambda} - 1 \right) \\ &:= \frac{1}{Z(\lambda, v_{xx}(x))} \exp \left(-\frac{1}{\lambda} [\text{tr}(v_{xx}(x))u] \right), \quad u \in U, \quad x \in \mathbb{R}^d. \end{aligned} \quad (3.3)$$

Since $\bar{\pi}(\cdot; x)$ must satisfy (3.2), we deduce

$$Z(\lambda, v_{xx}(x)) = \int_U \exp \left(-\frac{1}{\lambda} [\text{tr}(v_{xx}(x))u] \right) du > 0.$$

This gives the optimal *feedback law* $\bar{\pi}(\cdot; x)$, which is a *truncated* (in U) *exponential distribution* with the (state-dependent) parameter $c(x) := \frac{\text{tr}(v_{xx}(x))}{\lambda}$. Note we do not require $\text{tr}(v_{xx}(x)) > 0$ (i.e. v is in general non-convex) or $c(x) > 0$ here.

Substituting (3.3) back to the (3.1), and noting

$$\begin{aligned} &\int_U [\text{tr}(v_{xx}(x))u + \lambda \ln \bar{\pi}(u; x)] \bar{\pi}(u; x) du \\ &= \int_U [\text{tr}(v_{xx}(x))u - \text{tr}(v_{xx}(x))u - \lambda \ln(Z(\lambda, v_{xx}(x)))] \bar{\pi}(u; x) du \\ &= -\lambda \ln(Z(\lambda, v_{xx}(x))), \end{aligned}$$

we obtain the following equivalent form of the HJB equation, a nonlinear elliptic PDE:

$$-\rho v(x) - f_x(x) \cdot v_x(x) + f(x) - \lambda \ln(Z(\lambda, v_{xx}(x))) = 0, \quad x \in \mathbb{R}^d. \quad (3.4)$$

Remark 5 (Gibbs Measure and Boltzmann Exploration). In reinforcement learning there is a widely used *heuristic* exploration strategy called the *Boltzmann exploration*. In the setting of *maximizing* cumulative rewards, Boltzmann exploration uses the Boltzmann distribution (or the Gibbs measure) to assign a probability $p(s_t, a)$ to action a when in state s_t at time t :

$$p(s_t, a) = \frac{e^{Q_t(s_t, a)/\lambda}}{\sum_{a=1}^m e^{Q_t(s_t, a)/\lambda}}, \quad a = 1, 2, \dots, m, \quad (3.5)$$

where $Q_t(s, a)$ is the Q-function value of a state-action pair (s, a) , and $\lambda > 0$ is a parameter that controls the level of exploration; see e.g. [Bridle \(1990\)](#); [Cesa-Bianchi](#)

et al. (2017); Sutton and Barto (2018). When λ is high, the actions are chosen in almost equal probabilities. When λ is low, the highest-valued actions are more likely to be chosen, giving rise to something resembling the so-called *epsilon-greedy policies* in multi-armed bandit problems. When $\lambda = 0$, (3.5) degenerates into the Dirac measure concentrating on the action that maximizes the Q-function; namely the best action is always chosen. In the setting of continuous time and continuous state/control, the Q-function $Q_t(s, a)$ is not well defined and can not be used to rank and select actions (Tallec et al., 2019). However, we can use the (generalized) *Hamiltonian* (Yong and Zhou, 1999), which in the special case of the classical problem (2.5) is

$$H(x, u, v_x, v_{xx}) := -f_x(x) \cdot v_x(x) + f(x) + \text{tr}(v_{xx}(x))u.$$

Then it is easy to see that (3.3) is equivalent to the following:

$$\bar{\pi}(u; x) = \frac{\exp(-\frac{1}{\lambda}[H(x, u, v_x, v_{xx})])}{\int_{u \in U} \exp(-\frac{1}{\lambda}[H(x, u, v_x, v_{xx})])du}, \quad u \in U, \quad x \in \mathbb{R}^d, \quad (3.6)$$

which is analogous to (3.5) with x being the state and u being the action. The negative sign in the above expression is due to the minimization problem instead of a maximization one. The importance of this observation is that here we lay a *theoretical underpinning* of the Boltzmann exploration via entropy regularization, thereby provide interpretability/explainability of a largely heuristic approach.⁵

3.2 Admissibility and optimality

The derivation above is formal, and is legitimate only after we *rigorously* establish the verification theorem and check the admissibility of the distributional controls induced by the feedback law (3.3).

To this end we need to impose some assumptions.

Assumption 6. We assume

- (i) $U = [a, 1]$ for some $0 < a < 1$.
- (ii) $f \in C^1$ and $|f_x(x)| < C$ for some constant $C > 0$.

⁵The formula (3.6) was derived in (Wang et al., 2019, eq. (17)) for a more general setting, but the connection with Boltzmann exploration and Gibbs measure was not noted there.

Assumption (i) imposes some minimal heating, $a > 0$, to the Langevin diffusion, capturing the fact that in practice we would generally not know whether we have reached the global minimum at any given time; therefore we always carry out exploration, however small it might be, until some prescribed stopping criterion is met. On the other hand, the difficulty of the original non-convex optimization lies in the possibility of many *local* optima, while it is relatively easier to know (sometime from the specific context of an applied problem under consideration) the approximate range of the possible location of a global optimum. Therefore, the assumption of the gradient of f in Assumption (ii) is reasonable as the function value outside of the above range is irrelevant.

Apply the feedback law (3.3) to the controlled system (2.8) to obtain the following SDE

$$dX^*(t) = -f_x(X^*(t))dt + h(X^*(t))dW_t, \quad X^*(0) = x \quad (3.7)$$

where

$$h(x) := \sqrt{\frac{2}{Z(\lambda, v_{xx}(x))} \int_a^1 u \exp\left(-\frac{1}{\lambda}[\text{tr}(v_{xx}(x))u]\right) du}. \quad (3.8)$$

It is clear that $0 < \sqrt{2a} \leq h(x) \leq \sqrt{2}$. Thus the drift and diffusion coefficients of (3.7) are both uniformly bounded, while the latter satisfies the uniform elliptic condition. It follows from (Krylov, 1980, p. 87, Theorem 1) that (3.7) has a weak solution. Moreover, (Stroock and Varadhan, 1969, Theorem 6.2) asserts that the weak solution is unique.

Now define

$$\pi^*(t, u) := \bar{\pi}(u; X^*(t)) \equiv \frac{1}{Z(\lambda, v_{xx}(X^*(t)))} \exp\left(-\frac{1}{\lambda}[\text{tr}(v_{xx}(X^*(t)))u]\right), \quad t \geq 0, \quad u \in [a, 1], \quad (3.9)$$

where $\{X^*(t) : t \geq 0\}$ is a weak solution to (3.7). Let $\pi^* := \{\pi^*(t, \cdot) : t \geq 0\}$. Then clearly $\pi^* \in \mathcal{A}(x)$.

Theorem 7. *Suppose $v \in C^2$ is a solution to the HJB equation (3.1), and $|v(x)| < C(1 + |x|^k) \forall x \in \mathbb{R}^d$ for some constants $C > 0$ and $k > 0$. Then $v = V$. Moreover, in this case the distributional control π^* defined by (3.9) is optimal.*

We need a series of technical preliminaries to prove this theorem.

Lemma 8. *For any $\pi \in \mathcal{A}(x)$, we have*

$$\int_a^1 \pi(t, u) \ln \pi(t, u) du \geq a, \quad a.s. \quad (3.10)$$

Proof. Applying the general inequality $x \ln x \geq x - 1$ for $x > 0$, we have

$$\int_a^1 \pi(t, u) \ln \pi(t, u) du \geq \int_a^1 (\pi(t, u) - 1) du = a. \quad (3.11)$$

Lemma 9. Suppose $\{X(t) : t \geq 0\}$ follows

$$dX(t) = b(t)dt + \sigma(t)dW(t), \quad t \geq 0$$

with

$$C_b := \operatorname{ess\,sup}_{(t, \omega)} |b(t)| < \infty, \quad C_\sigma := \operatorname{ess\,sup}_{(t, \omega)} |\sigma(t)| < \infty.$$

Then, for any $k \geq 1$, there exists a constant $C > 0$, which is independent of T and $X(0)$, such that

$$\mathbb{E} \left[\sup_{0 \leq t \leq T} |X(t)|^k \right] < C(1 + T^k + |X(0)|^k), \quad \forall T \geq 0. \quad (3.12)$$

Proof. By the elementary inequality

$$(a + b + c)^k \leq (3 \max\{a, b, c\})^k \leq 3^k a^k + 3^k b^k + 3^k c^k, \quad a, b, c \geq 0,$$

we have

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq t \leq T} |X(t)|^k \right] &\leq \mathbb{E} \left[\left(|X(0)| + \sup_{0 \leq t \leq T} \int_0^t |b(s)| ds + \sup_{0 \leq t \leq T} \left| \int_0^t \sigma(s) dW_s \right| \right)^k \right] \\ &\leq \mathbb{E} \left[\left(|X(0)| + C_b T + \sup_{0 \leq t \leq T} \left| \int_0^t \sigma(s) dW_s \right| \right)^k \right] \\ &\leq 3^k |X(0)|^k + 3^k C_b^k T^k + 3^k \mathbb{E} \left[\sup_{0 \leq t \leq T} \left| \int_0^t \sigma(s) dW_s \right|^k \right] \\ &\leq 3^k |X(0)|^k + 3^k C_b^k T^k + 3^k C_k \left(\mathbb{E} \left[\int_0^T |\sigma(s)|^2 ds \right] \right)^{k/2} \\ &\leq 3^k |X(0)|^k + 3^k C_b^k T^k + 3^k C_k C_\sigma^k T^{k/2}, \end{aligned}$$

where the second to last inequality is due to the Burkholder-Davis-Gundy inequality. This proves (3.12).

Proposition 10. There exists a constant $C_1 > 0$ such that the value function V satisfies

$$|V(x)| \leq C_1(1 + |x|), \quad \forall x \in \mathbb{R}^d. \quad (3.13)$$

Proof. It follows from Assumption 6-(ii) that $|f(x)| \leq C|x| + |f(0)|$. Let π be the uniform distribution on $[a, 1]$. It then follows from Lemma 9 that

$$V(x) \leq \mathbb{E} \left[\int_0^\infty e^{-\rho t} (C|X^\pi(t)| + |f(0)|) dt + \lambda \int_0^\infty e^{-\rho t} \int_a^1 \pi(u) \ln \pi(u) du dt \right] < C'(|x| + 1)$$

for some constant C' independent of x . On the other hand, for any $\pi \in \mathcal{A}(x)$, by virtue of Lemma 8 we have

$$\begin{aligned} & \mathbb{E} \left[\int_0^\infty e^{-\rho t} f(X^\pi(t)) dt - \lambda \int_0^\infty e^{-\rho t} \int_a^1 -\pi_t(u) \ln \pi_t(u) du dt \right] \\ & \geq \mathbb{E} \left[\int_0^\infty -e^{-\rho t} (C|X^\pi(t)| + |f(0)|) dt + \lambda \int_0^\infty e^{-\rho t} a dt \right] > -C''(|x| + 1) \end{aligned}$$

for some constant C'' independent of x . The proof is complete.

We have indeed established in the above that

$$\mathbb{E} \left[\int_0^\infty e^{-\rho t} |f(X^\pi(t))| dt \right] < \infty.$$

To solve problem (2.7), we only need to consider those admissible controls π such that

$$\mathbb{E} \left[\int_0^\infty e^{-\rho t} \int_U \pi(t, u) \ln \pi(t, u) du dt \right] < \infty, \quad (3.14)$$

because any control violating (3.14) renders the infinite value of the minimization problem (2.7) and hence can be excluded from consideration. In view of this, (3.14) is henceforth assumed for any *admissible* distributional control.

Now we prove Theorem 7.

For any $\pi \in \mathcal{A}(x)$, let $\{X^\pi(t) : t \geq 0\}$ be the solution to (2.8). Define $\tau_n := \inf\{s \geq 0 :$

$|X^\pi(s)| > n\}$. For any $T > 0$, by Itô's lemma,

$$\begin{aligned}
& e^{-\rho(\tau_n \wedge T)} v(X^\pi(\tau_n \wedge T)) \\
&= v(X^\pi(0)) + \int_0^{\tau_n \wedge T} e^{-\rho t} v_x(X^\pi(t)) \cdot \tilde{\sigma}(\pi(t)) dW(t) \\
&\quad + \int_0^{\tau_n \wedge T} e^{-\rho t} \left(-\rho v(X^\pi(t)) - f_x(X^\pi(t)) \cdot v_x(X^\pi(t)) + \frac{1}{2} \tilde{\sigma}^2(\pi(t)) \text{tr}(v_{xx}(X^\pi(t))) \right) dt \\
&= v(x) + \int_0^{\tau_n \wedge T} e^{-\rho t} v_x(X^\pi(t)) \cdot \tilde{\sigma}(\pi(t)) dW(t) + \int_0^{\tau_n \wedge T} e^{-\rho t} \left(\int_a^1 \text{tr}(v_{xx}(X^\pi(t))) u \pi(t, u) du \right. \\
&\quad \left. - f(X^\pi(t)) - \inf_{\pi' \in \mathcal{P}([a, 1])} \int_a^1 [\text{tr}(v_{xx}(X^\pi(t))) u + \lambda \ln \pi'(u)] \pi'(u) du \right) dt \\
&\geq v(x) + \int_0^{\tau_n \wedge T} e^{-\rho t} v_x(X^\pi(t)) \cdot \tilde{\sigma}(\pi(t)) dW(t) \\
&\quad + \int_0^{\tau_n \wedge T} e^{-\rho t} \left(-f(X^\pi(t)) - \lambda \int_a^1 \pi(t, u) \ln \pi(t, u) du \right) dt.
\end{aligned}$$

It follows

$$\begin{aligned}
v(x) &\leq \mathbb{E} \left[e^{-\rho(\tau_n \wedge T)} v(X^\pi(\tau_n \wedge T)) \right] + \mathbb{E} \left[\int_0^{\tau_n \wedge T} e^{-\rho t} \left(f(X^\pi(t)) + \lambda \int_a^1 \pi(t, u) \ln \pi(t, u) du \right) dt \right] \\
&= \mathbb{E} \left[e^{-\rho(\tau_n \wedge T)} v(X^\pi(\tau_n \wedge T)) \right] + \mathbb{E} \left[\int_0^{\tau_n \wedge T} e^{-\rho t} f(X^\pi(t)) dt \right] \\
&\quad + \lambda \mathbb{E} \left[\int_0^{\tau_n \wedge T} e^{-\rho t} \int_a^1 \pi(t, u) \ln \pi(t, u) du dt \right].
\end{aligned}$$

Sending $n \rightarrow \infty$, we have by Lemma 9 and the dominated convergence theorem that the sum of the first two terms converges to

$$\mathbb{E} \left[e^{-\rho T} v(X^\pi(T)) \right] + \mathbb{E} \left[\int_0^T e^{-\rho t} f(X^\pi(t)) dt \right],$$

while it follows from (3.14) along with the monotone convergence theorem (noticing $\pi(t, u) \ln \pi(t, u) \geq \pi(t, u) - 1$) that the last term converges to

$$\lambda \mathbb{E} \left[\int_0^T e^{-\rho t} \int_a^1 \pi(t, u) \ln \pi(t, u) du dt \right].$$

So

$$\begin{aligned}
v(x) &\leq \mathbb{E} \left[e^{-\rho T} v(X^\pi(T)) \right] + \mathbb{E} \left[\int_0^T e^{-\rho t} f(X^\pi(t)) dt \right] \\
&\quad + \lambda \mathbb{E} \left[\int_0^T e^{-\rho t} \int_a^1 \pi(t, u) \ln \pi(t, u) du dt \right].
\end{aligned}$$

Noting Lemma 9 along with the assumption that v is of polynomial growth, we have

$$\lim_{T \rightarrow \infty} \mathbb{E} [e^{-\rho T} v(X^\pi(T))] = 0. \quad (3.15)$$

Applying the same limiting argument as above we let $T \rightarrow \infty$ to get

$$v(x) \leq \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(f(X^\pi(t)) + \lambda \int_U \pi(t, u) \ln \pi(t, u) du \right) dt \right].$$

Since π is arbitrarily chosen, it follows $v \leq V$; in other words v is a lower bound of V .

All the inequalities in the above analysis become equalities when we take $\pi = \pi^* \in \mathcal{A}(x)$ defined by (3.9), because π^* attains the infimum in the HJB equation (3.1). This yields the aforementioned lower bound is achieved by π^* ; hence π^* is optimal and $v = V$.

Remark 11 (Existence and Uniqueness of Solution to HJB). Due to the linear growth of the value function V (Proposition 10), we can prove by a standard argument (part of it is similar to the argument used in the above proof) that V is a solution to the HJB equation (3.1) as long as $V \in C^2$. Moreover, Theorem 7 yields that the solution is unique among C^2 functions of polynomial growth. The uniqueness follows from (3.15) which is essentially a boundary condition of (3.1). When V is not C^2 , we can apply viscosity solution theory to study the well-posedness of the HJB equation. However, we will not pursue this direction as it is not directly related to the objective of this paper.

4 Algorithm and Numerical Example

As discussed in Introduction, the main contribution of this paper is to develop a theory that connects the state-dependent temperature control problem for SA to an HJB equation, the latter concerning the global information about the underlying function f . The theory also naturally gives rise to an algorithm that can be employed to solve the original non-convex optimization problem to the extent that the HJB equation can be efficiently solved. Specifically, the optimal state process $\{X^*(t) : t \geq 0\}$ is determined by the SDE (3.7). Discretization of this equation then leads to a Langevin algorithm with state-dependent noise. One can use Euler-Maruyama discretization; see, e.g. Mattingly et al. (2002), due to its simplicity and popularity in practice. We now report a numerical experiment for a baseline non-convex optimization example in a one-dimensional state space for which the HJB equation is an ODE.

Consider the problem of finding the global minima of the following asymmetric double-well function:

$$f(x) = \begin{cases} 4x - 20 & \text{for } x > 6, \\ (x - 4)^2 & \text{for } 2 < x \leq 6, \\ 8 - x^2 & \text{for } -2 < x \leq 2, \\ 2(x + 3)^2 + 2 & \text{for } -6 < x \leq -2, \\ -12x - 52 & \text{for } x \leq -6. \end{cases}$$

Note that $f : \mathbb{R} \rightarrow \mathbb{R}$ is a continuously differentiable function with two local minima at $x_1 = -3$ and $x_2 = 4$, where the latter is the (unique) global minimum. See Figure 1 for an illustration.




Figure 1: A double well function f

We now present four algorithms, to be described below, for minimizing f and compare their performances. All the algorithms are initialized at $x = -3$, the suboptimal local minimum of f . We examine the convergence of the performance quantity $\mathbb{E}(f(X_k))$ where X_k is the location of the iterate at iteration k , $k = 1, 2, \dots$. In each of the algorithms, $\{\xi_k\}$ is a sequence of i.i.d standard normal random variable and η is the constant step size which we tune respectively for best performance. In implementation, we change the upper bound of the allowed temperature to be 500, instead of 1 in our theoretical analysis, without loss of validity. This allows the iterates of the algorithms to climb over the barrier between the

two local minimizers quickly.

The four algorithms are

- (1) Langevin algorithm with constant temperature β (Raginsky et al., 2017):

$$X_{k+1} = X_k - \eta f'(X_k) + \sqrt{2\eta\beta}\xi_k. \quad (4.1)$$

We choose the stepsize $\eta \in \{(1/2)^n : n = 0, 1, 2, \dots, 10\}$ and the temperature $\beta \in \{500 \cdot (1/2)^n : n = 0, 1, 2, \dots, 15\}$. The best-tuned step size $\eta = (1/2)^1 = 0.5$ and temperature $\beta = 500 \cdot (1/2)^{10} = 0.4883$.

- (2) Langevin algorithm with the power-law temperature schedule (Neelakantan et al., 2015):

$$X_{k+1} = X_k - \eta f'(X_k) + \sqrt{2\eta\beta_k}\xi_k, \quad (4.2)$$

where we choose $\beta_k = (\frac{d}{1+k})^b$ with $b \in \{0.5, 0.6, \dots, 1\}$, and $d \in \{500 \cdot (1/2)^n : n = 0, 1, 2, \dots, 7\}$; We also choose the stepsize $\eta \in \{(1/2)^n : n = 0, 1, 2, \dots, 10\}$. The best-tuned parameters are given by: $\eta = (1/2)^1 = 0.5, b = 0.9, d = 500 \cdot (1/2)^4 = 31.25$.

- (3) Replica exchange (GDxLD in Algorithm 1 of Dong and Tong (2020)): one runs a copy of the gradient descent $\{X_k\}$ and a copy of the Langevin algorithm $\{Y_k\}$ with a constant temperature $\gamma > 0$. If $f(X_k) > f(Y_k)$, then the positions of X_k and Y_k are swaped. Output X_N as an optimizer of f when the algorithm terminates at iteration N . We choose $X_0 = Y_0 = -3$, and tune parameters where stepsize $\eta \in \{(1/2)^n : n = 0, 1, 2, \dots, 10\}$, and constant temperature $\gamma \in \{500 \cdot (1/2)^n : n = 0, 1, 2, \dots, 15\}$. The best-tuned parameters are given by $\eta = (1/2)^1 = 0.5, \gamma = 500 \cdot (1/2)^1 = 250$.

- (4) Our Langevin algorithm based on the Euler-Maruyama discretization of the SDE (3.7):

$$X_{k+1}^* = X_k^* - \eta f'(X_k^*) + \sqrt{\eta} h(X_k^*) \xi_k. \quad (4.3)$$

In dimension one we can deduce from (3.8) that $h(x) = g\left(\frac{v_{xx}(x)}{\lambda}\right)$, where v satisfies the ODE in (3.4) with

$$Z(\lambda, v_{xx}(x)) = \frac{e^{-a\frac{v_{xx}(x)}{\lambda}} - e^{-c\frac{v_{xx}(x)}{\lambda}}}{\frac{v_{xx}(x)}{\lambda}}, \quad (4.4)$$

and

$$g(y) := \sqrt{2 \frac{(c+1/y) \cdot e^{-cy} - (a+1/y) \cdot e^{-ya}}{e^{-cy} - e^{-ya}}}, \quad y \in (-\infty, \infty). \quad (4.5)$$

Here, $[a, c]$ is the range of the allowed temperature. The hyperparameter a in (4.5) is fixed at 0.0001 for simplicity, and the hyperparameter c , which is the upper bound on the allowed temperature, is fixed at 500. We choose $\rho, \lambda \in \{5 \cdot (1/2)^n : n = 0, 1, 2, \dots, 8\}$, and stepsize $\eta \in \{(1/2)^n : n = 0, 1, 2, \dots, 10\}$. To implement the algorithm (4.3), we need to solve the second-order nonlinear ODE in (3.4). We use random initial conditions $[v(0), v'(0)]$ where both $v(0)$ and $v'(0)$ are independently sampled from a standard normal distribution. We select 20 random initializations for the ODE. The best-tuned parameters are given by $\eta = (1/2)^3 = 0.125$, $\lambda = 5 \cdot (1/2)^4 = 0.3125$, $\rho = 5 \cdot (1/2)^2 = 1.25$, with the ODE initialization $[v(0), v'(0)] = [-0.2853, 1.1575]$.

Figure 2 shows the performance of the four algorithms, where the expectation $\mathbb{E}(f(X_k))$ is approximated by its sample average with the sample size 500. Each algorithm is terminated if the number of iterations exceeds the allowed number of iterations, which we set to be 1000 in our experiment. In the first 100 iterations, the expected function values $\mathbb{E}(f(X_k))$ from our algorithm are very large, primarily because large noises are injected into the iterates during these initial iterations so as to escape the local minimum at -3 ; see Figure 3 for a zoomed-in version. Hence, for better visualization, we plot in Figure 2 the expected function values $\mathbb{E}(f(X_k))$ from $k = 100$ up to only $k = 500$ iterations. As we can see, Langevin algorithms with constant temperature and power-law temperature schedule have difficulties in locating the global minima within 500 iterations. By its very definition, replica exchange is expected to perform better than the Langevin algorithm with a constant temperature. This is confirmed by Figure 2 in which the replica exchange algorithm finds the global minimum quickly, although it needs to run two algorithms (the gradient descent and a Langevin algorithm) instead of only one algorithm. Our Langevin algorithm with state-dependent noises can find the global minimum faster than the other three algorithms. It is, however, computationally more expensive compared with other three methods, needing to solve a nonlinear ODE.

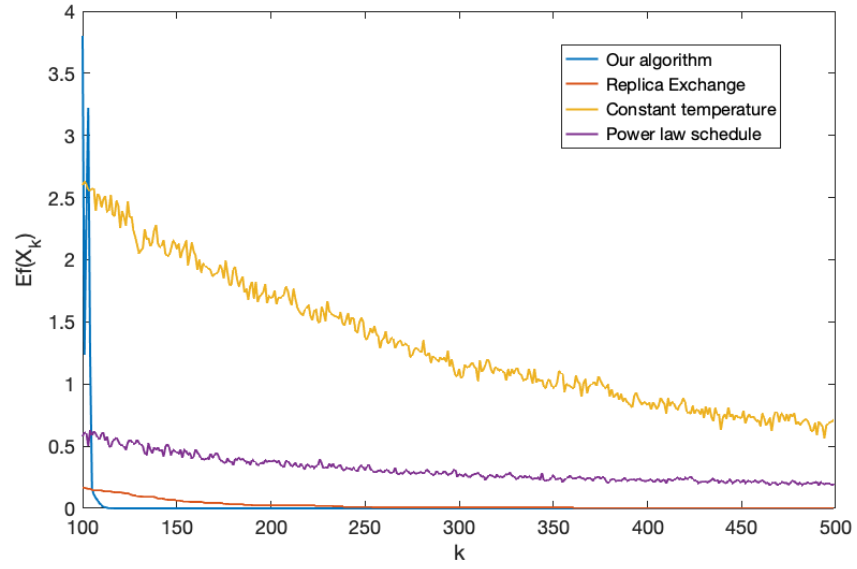


Figure 2: Performance comparison of the four algorithms, all initialized at $X_1 = -3$.




Figure 3: Performance of our algorithm initialized at $X_1 = -3$.

The temperature process for our algorithm depends on the solution v to the ODE (3.4). We plot in Figure 4(a) the function values $v(x)$ where $x \in (-6, 4)$. We observe that when x is close to 4, the global minimum of f , v grows very quickly. This is because much less exploration is needed near the global minimum, in which case the subtraction of the entropy value from the overall objective value in (2.7) is much less, boosting the value of v . We also plot in Figure 4(b) the second order derivative v on the interval $(-6, 2)$.⁶ However, v or v'' affects the temperature process of our algorithm only indirectly, through the function h given by (4.3). So the plots of v and v'' are less informative than that of the state-dependent temperature function $\frac{h^2}{2}$ which is proportional to the variance of the noise in our algorithm; see Figure 5. We see that the temperature is close to zero for $x > 3$, and is mostly large elsewhere. This indicates that our state-dependent algorithm is “intelligent”: it uses the lowest temperature when close to the global minimum, and uses a large temperature (recall the largest temperature allowed is 500) to escape from traps such as suboptimal local minima or saddle points. We also observe that there is a prominent kink at $x = -2$ in Figure 5. This is primarily due to the spike of v'' at -2 in Figure 4(b) and the fact that $h(x) = g\left(\frac{v_{xx}(x)}{\lambda}\right)$ with g given in (4.5).




Figure 4: The HJB solution v and its second order derivative v'' .

⁶We do not plot $v''(x)$ for $x > 2$ in Figure 4(b) for better visualization because $v''(x)$ becomes very large when $x > 2$.




Figure 5: State-dependent temperature function $\frac{h^2}{2}$.

5 Conclusion

This paper aims to develop an endogenous temperature control scheme for applying Langevin diffusions to find non-convex global minima. We take the exploratory stochastic control framework, originally proposed by Wang et al. (2019) for reinforcement learning, to account for the need of smoothing out the temperature process. We derive a state-dependent, Boltzmann-exploration type distributional control, which can be used to sample temperatures in a Langevin algorithm. Numerical analysis shows that our algorithm outperforms three alternative ones based on a constant temperature, a power decay schedule and a replica exchange method respectively. However, the function used in the numerical example is one-dimensional, for which the HJB equation is an ODE and hence easy to solve. For high-dimensional problems, the HJB equation is a PDE whose numerical solutions may suffer from the curse of dimensionality. Therefore, at least for now, the main contribution of this paper is *not* algorithmic. Rather, it is, generally, to lay a *theoretical* underpinning for smoothing out often overly rigid classical optimal controls (such as bang-bang controls) and, specifically, to provide an interpretable state-dependent temperature control scheme for Langevin diffusions via HJB equations.

Acknowledgments

Xuefeng Gao acknowledges financial support from Hong Kong GRF (No.14201117). Zuo Quan Xu acknowledges financial support from NSFC (No.11971409), Hong Kong GRF (No.15202817), the PolyU-SDU Joint Research Center on Financial Mathematics and the CAS AMSS-POLYU Joint Laboratory of Applied Mathematics, The Hong Kong Polytechnic University. Xun Yu Zhou acknowledges financial supports through a start-up grant at Columbia University and the Nie Center for Intelligent Asset Management. We also thank Mert Gürbüzbalaban and Lingjiong Zhu for their comments, and Yi Xiong for the help with the experiments.

References

- Christian Beck, E Weinan, and Arnulf Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29(4):1563–1619, 2019.
- Régis Bertrand and Richard Epenoy. New smoothing techniques for solving bang–bang optimal control problems – numerical results and statistical interpretation. *Optimal Control Applications and Methods*, 23(4):171–197, 2002.
- Anton Bovier, Véronique Gayrard, and Markus Klein. Metastability in reversible diffusion processes I: Sharp asymptotics for capacities and exit times. *Journal of the European Mathematical Society*, 6(4):399–424, 2004.
- Anton Bovier, Véronique Gayrard, and Markus Klein. Metastability in reversible diffusion processes II: Precise asymptotics for small eigenvalues. *Journal of the European Mathematical Society*, 7(1):69–99, 2005.
- John S Bridle. Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. In *Advances in neural information processing systems*, pages 211–217, 1990.
- Nicolò Cesa-Bianchi, Claudio Gentile, Gábor Lugosi, and Gergely Neu. Boltzmann exploration done right. In *Advances in neural information processing systems*, pages 6284–6293, 2017.
- Xi Chen, Simon S Du, and Xin T Tong. On stationary-point hitting time and ergodicity of stochastic gradient langevin dynamics. *Journal of Machine Learning Research*, 21(68):1–41, 2020.
- Tzoo-Shuh Chiang, Chii-Ruey Hwang, and Shuenn Jyi Sheu. Diffusion for global optimization in \mathbb{R}^n . *SIAM Journal on Control and Optimization*, 25(3):737–753, 1987.
- Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent. In *Conference on Learning Theory*, pages 678–689, 2017.

- Jing Dong and Xin T Tong. Replica exchange for non-convex optimization. *arXiv preprint arXiv:2001.08356*, 2020.
- Paul Dupuis, Yufei Liu, Nuria Plattner, and Jimmie D Doll. On the infinite swapping limit for parallel tempering. *Multiscale Modeling & Simulation*, 10(3):986–1022, 2012.
- David J Earl and Michael W Deem. Parallel tempering: Theory, applications, and new perspectives. *Physical Chemistry Chemical Physics*, 7(23):3910–3916, 2005.
- Haitao Fang, Minping Qian, and Guanglu Gong. An improved annealing method and its large-time behavior. *Stochastic processes and their applications*, 71(1):55–74, 1997.
- Saul B Gelfand and Sanjoy K Mitter. Recursive stochastic algorithms for global optimization in \mathbb{R}^d . *SIAM Journal on Control and Optimization*, 29(5):999–1018, 1991.
- Stuart Geman and Chii-Ruey Hwang. Diffusions for global optimization. *SIAM Journal on Control and Optimization*, 24(5):1031–1043, 1986.
- Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- Mert Gürbüzbalaban, Xuefeng Gao, Yuanhan Hu, and Lingjiong Zhu. Decentralized stochastic gradient langevin dynamics and hamiltonian monte carlo. *arXiv preprint arXiv:2007.00590*, 2020.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Richard A Holley, Shigeo Kusuoka, and Daniel W Stroock. Asymptotics of the spectral gap with applications to the theory of simulated annealing. *Journal of functional analysis*, 83(2):333–347, 1989.
- Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

- N.V. Krylov. *Controlled Diffusion Processes*. Springer, New York, 1980.
- Holden Lee, Andrej Risteski, and Rong Ge. Beyond log-concavity: Provable guarantees for sampling multi-modal distributions using simulated tempering langevin monte carlo. In *Advances in Neural Information Processing Systems*, pages 7847–7856, 2018.
- Yi-An Ma, Tianqi Chen, and Emily Fox. A complete recipe for stochastic gradient mcmc. In *Advances in Neural Information Processing Systems*, pages 2917–2925, 2015.
- Enzo Marinari and Giorgio Parisi. Simulated tempering: a new monte carlo scheme. *EPL (Europhysics Letters)*, 19(6):451, 1992.
- David Márquez. Convergence rates for annealing diffusion processes. *The Annals of Applied Probability*, pages 1118–1139, 1997.
- Jonathan C Mattingly, Andrew M Stuart, and Desmond J Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic Processes and their Applications*, 101(2):185–232, 2002.
- Toyonori Munakata and Yasuyuki Nakamura. Temperature control for simulated annealing. *Physical Review E*, 64(4):046127, 2001.
- Arvind Neelakantan, Luke Vilnis, Quoc V Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. Adding gradient noise improves learning for very deep networks. *arXiv preprint arXiv:1511.06807*, 2015.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703, 2017.
- Cristiana Silva and Emmanuel Trélat. Smooth regularization of bang-bang optimal control problems. *IEEE Transactions on Automatic Control*, 55(11):2488–2499, 2010.
- D.W. Stroock and S.R.S. Varadhan. Diffusion processes with continuous coefficients, i. *Communications On Pure And Applied Mathematics*, 22(22):345–400, 1969.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- Robert H Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Physical review letters*, 57(21):2607, 1986.
- Corentin Tallec, Léonard Blier, and Yann Ollivier. Making deep q-learning methods robust to time discretization. *arXiv preprint arXiv:1901.09732*, 2019.
- Nicholas G Tawn, Gareth O Roberts, and Jeffrey S Rosenthal. Weight-preserving simulated tempering. *Statistics and Computing*, 30(1):27–41, 2020.
- Haoran Wang, Thaleia Zariphopoulou, and Xun Yu Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 2019.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.
- Pan Xu, Jinghui Chen, Difan Zou, and Quanquan Gu. Global convergence of Langevin dynamics based algorithms for nonconvex optimization. In *Advances in Neural Information Processing Systems*, pages 3122–3133, 2018.
- Jiongmin Yong and Xun Yu Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*, volume 43. Springer Science & Business Media, 1999.
- Yuchen Zhang, Percy Liang, and Moses Charikar. A hitting time analysis of stochastic gradient langevin dynamics. In *Conference on Learning Theory*, pages 1980–2022, 2017.