

Explicit and Implicit Graduated Optimization in Deep Neural Networks

Naoki Sato, Hideaki Iiduka

Meiji University
naoki310303@gmail.com, iiduka@cs.meiji.ac.jp

Abstract

Graduated optimization is a global optimization technique that is used to minimize a multimodal nonconvex function by smoothing the objective function with noise and gradually refining the solution. This paper experimentally evaluates the performance of the explicit graduated optimization algorithm with an optimal noise scheduling derived from a previous study and discusses its limitations. The evaluation uses traditional benchmark functions and empirical loss functions for modern neural network architectures. In addition, this paper extends the implicit graduated optimization algorithm, which is based on the fact that stochastic noise in the optimization process of SGD implicitly smooths the objective function, to SGD with momentum, analyzes its convergence, and demonstrates its effectiveness through experiments on image classification tasks with ResNet architectures.

Code — <https://github.com/iiduka-researches/igo-aaai25>

Introduction

Nonconvex optimization problems are ubiquitous in the machine-learning and artificial-intelligence fields, and those arising in training deep neural networks (DNN) (Bengio 2009) are particularly interesting from multiple aspects. In particular, because the nonconvex functions that appear in DNN training have many local optimal solutions, stochastic gradient descent (SGD) (Robbins and Monro 1951) and its variants, which use gradients to update the sequence, may fall into local optimal solutions and may not minimize losses sufficiently.

Graduated optimization (Blake and Zisserman 1987) is a global optimization method that can be applied to multimodal functions in order to avoid local optimal solutions and converge to the global optimal one. The method first prepares M monotone decreasing noises $(\delta_m)_{m \in [M]}$. Then, by smoothing the original objective function f with that noise, a sequence of M smoothed functions $(\hat{f}_{\delta_m})_{m \in [M]}$ is obtained that gradually approaches the original objective function. The function \hat{f}_{δ_1} smoothed with the largest noise δ_1 is then optimized, and the approximate solution is used as the initial point for optimizing the function \hat{f}_{δ_2} smoothed with the second largest noise δ_2 . After that, the function \hat{f}_{δ_3} smoothed

with the third largest noise δ_3 is optimized with the second approximate solution as the initial point, and the procedure is repeated in order to search for the global optimal solution without falling into a local optimal one. See Figure 1 for a conceptual diagram of this process.

Graduated optimization is often used in machine learning and computer vision for, e.g., semi-supervised learning (Chapelle, Chi, and Zien 2006; Sindhwani, Keerthi, and Chapelle 2006; Chapelle, Sindhwani, and Keerthi 2008) and robust estimation (Yang et al. 2020; Antonante et al. 2022; Peng, Kümmerle, and Vidal 2023). In addition, score-based generative models (Song and Ermon 2019, 2020) and diffusion models (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020; Song et al. 2021; Rombach et al. 2022), which are state-of-the-art models of image generation, use this approach.

Theoretical analysis of graduated optimization began with the pioneering work by (Duchi, Bartlett, and Wainwright 2012) on nonsmooth convex functions, and several papers (Mobahi and Fisher III 2015; Hazan, Yehuda, and Shalev-Shwartz 2016; Iwakiri et al. 2022; Li, Wu, and Zhang 2023; Sato and Iiduka 2023) have produced important results. In particular, (Hazan, Yehuda, and Shalev-Shwartz 2016) defined a family of nonconvex functions that satisfy the conditions for convergence of a graduated optimization algorithm, called σ -nice, and proposed a first-order algorithm based on graduated optimization. They also studied the convergence and convergence rate of the algorithm to a global optimal solution for σ -nice functions. In addition, (Sato and Iiduka 2023) showed that the stochastic noise in the optimization process which SGD naturally has can implicitly smooth the objective function, and they proposed an implicit graduated optimization algorithm exploiting this property. Furthermore, they defined an extension of the σ -nice function, called the new σ -nice function, and provided a convergence analysis of the implicit graduated optimization algorithm for this function. Here, this paper refers to graduated optimization with explicit smoothing operations (see Definition 1) as “explicit graduated optimization” and to graduated optimization with implicit smoothing operations as “implicit graduated optimization”.

This paper analyzes the new σ -nice function and discusses the implicit graduated optimization algorithm, which has so far been studied in relation to stochastic noise in SGD, in

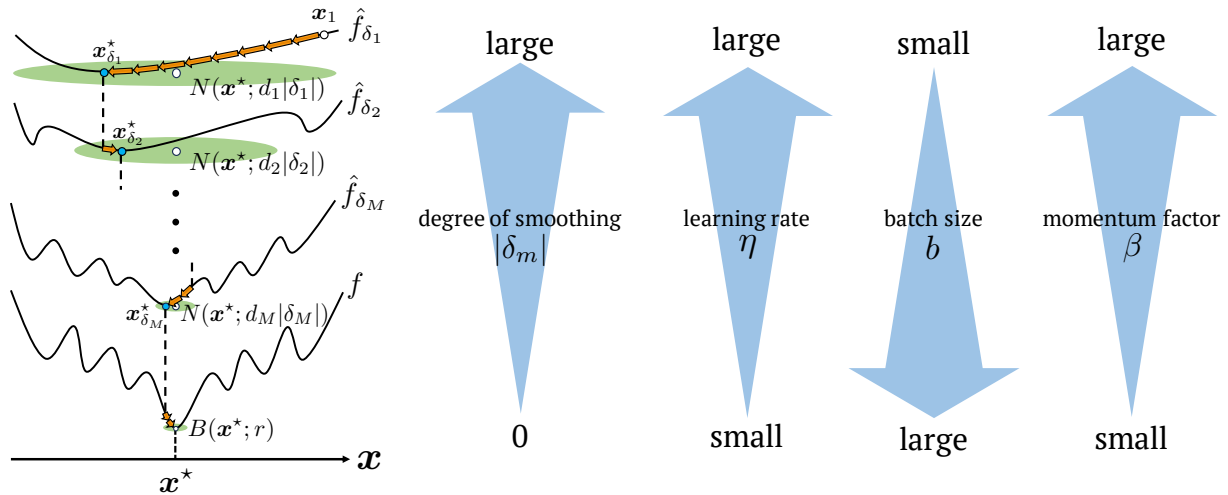


Figure 1: Conceptual diagrams of new σ -nice function and its smoothed versions and implicit graduated optimization. Note that this is a direct quotation from the prior work (Sato and Iiduka 2023) with minor changes made to fit our results.

relation to stochastic noise in SGD with momentum. It complements the prior studies with important theoretical and experimental validations on explicit and implicit graduated optimization.

Contributions

1. An example of a new σ -nice function. (Sato and Iiduka 2023) proposed new σ -nice functions, i.e., a family of non-convex functions that satisfy the necessary conditions for analysis and convergence of graduated optimization algorithms, but did not provide an example. Here, we theoretically prove that Rastrigin’s function (Törn and Zilinskas 1989; Rudolph 1990), one of the traditional benchmark functions for global optimization, is a new σ -nice function (Theorem 1).

2. Experimental validation in explicit graduated optimization. (Sato and Iiduka 2023) theoretically derived the optimal decay rate of the noise in both explicit and implicit graduated optimization, but did not provide numerical experiments to determine whether it is valid for explicit graduated optimization (Algorithm 1). Here, we applied the explicit graduated optimization algorithm and general global optimization methods to traditional benchmark functions for global optimization and found that the explicit graduated optimization algorithm with the optimal noise schedule outperforms the other methods on some problems (Table 1). We also applied the explicit graduated optimization algorithm to DNNs and found that its performance is not superior to other methods (Figure 2).

3. Experimental validation in implicit graduated optimization with SGD. (Sato and Iiduka 2023) proposed an implicit graduated optimization algorithm using stochastic noise in SGD (Algorithm 3), but did not provide numerical experimental results with the noise decay rate set as per theory. In this paper, we provide numerical experimental results that are consistent with theory as far as computational

complexity permits, confirming that their algorithm works effectively (Figure 3).

4. Theoretical development of implicit graduated optimization algorithm with SGD with momentum. On the basis that the stochastic noise of SGD is determined by the ratio of learning rate and batch size and that this noise helps to smooth the objective function, (Sato and Iiduka 2023) proposed an algorithm that achieves implicit graduated optimization by varying those parameters during training so that the noise decreases gradually. In this paper, we extend this argument to SGD with momentum, and on the basis that the stochastic noise of SGD with momentum is determined by the learning rate, batch size, and momentum factor (Sato and Iiduka 2024), we construct an algorithm (Algorithm 6) to perform implicit graduated optimization in the same way (see Figure 1). We also analyze the algorithm’s convergence and show that it reaches an ϵ -neighborhood of the global optimal solution of the new σ -nice function f in $\mathcal{O}\left(1/\epsilon^{\frac{1}{p}}\right)$ ($p \in (0, 1]$) rounds (Theorem 2).

5. Experimental validation of implicit graduated optimization with SGD with momentum. We tested the effectiveness of the proposed algorithm by using it to train ResNet18 and WideResNet-28-10 on the CIFAR100 dataset for image classification tasks. We confirmed that the algorithm has a lower loss function value and higher test accuracy than those of vanilla SGD with momentum (Figure 4). In addition, based on the theory of optimal noise decay rate in graduated optimization, the optimal decaying learning rate for SGD with momentum was tested by training ResNet34 on the ImageNet dataset. The results confirmed that the theoretically optimal learning rate scheduler, a polynomial decay scheduler with a power less than 1, achieves the lowest loss function value (Figure 5).

Preliminaries

Notation. Let \mathbb{N} denote the set of non-negative integers. For $m \in \mathbb{N} \setminus \{0\}$, define $[m] := \{1, 2, \dots, m\}$. The space \mathbb{R}^d is a d -dimensional Euclidean space with an inner product $\langle \cdot, \cdot \rangle$, which induces the norm $\|\cdot\|$. The Euclidean closed ball of radius r centered at $\hat{\mathbf{x}}$ is denoted by $B(\hat{\mathbf{x}}; r) := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \hat{\mathbf{x}}\| \leq r\}$. The DNN is parameterized by a vector $\mathbf{x} \in \mathbb{R}^d$, which is optimized by minimizing the empirical loss function $f(\mathbf{x}) := \frac{1}{n} \sum_{i \in [n]} f_i(\mathbf{x})$, where $f_i(\mathbf{x})$ is the loss function for $\mathbf{x} \in \mathbb{R}^d$ and the sample \mathbf{z}_i ($i \in [n]$). Let ξ be a random variable independent of $\mathbf{x} \in \mathbb{R}^d$, and let $\mathbb{E}_\xi[X]$ denote the expectation with respect to ξ of a random variable X . The random variable $\xi_{t,i}$ is generated from the i -th sampling at time t , and $\boldsymbol{\xi}_t := (\xi_{t,1}, \xi_{t,2}, \dots, \xi_{t,b})$ is independent of the sequence $(\mathbf{x}_k)_{k=0}^t \subset \mathbb{R}^d$, where b ($\leq n$) represents the batch size. The independence of $\boldsymbol{\xi}_0, \boldsymbol{\xi}_1, \dots$ allows us to define the total expectation \mathbb{E} as $\mathbb{E} = \mathbb{E}_{\xi_0} \mathbb{E}_{\xi_1} \dots \mathbb{E}_{\xi_t}$. Let $\mathbf{G}_{\xi_t}(\mathbf{x})$ be the stochastic gradient of $f(\cdot)$ at $\mathbf{x} \in \mathbb{R}^d$. The mini-batch \mathcal{S}_t consists of b samples \mathbf{z}_i at time t , and the mini-batch stochastic gradient of $f(\mathbf{x}_t)$ for \mathcal{S}_t is defined as $\nabla f_{\mathcal{S}_t}(\mathbf{x}_t) := \frac{1}{b} \sum_{i \in [b]} \mathbf{G}_{\xi_{t,i}}(\mathbf{x}_t)$.

Definition 1 (Smoothed function). *Given an L_f -Lipschitz function f , define \hat{f}_δ to be the function obtained by smoothing f ,*

$$\hat{f}_\delta(\mathbf{x}) := \mathbb{E}_{\mathbf{u} \sim B(\mathbf{0}; 1)} [f(\mathbf{x} - \delta \mathbf{u})], \quad (1)$$

where $\delta \in \mathbb{R}$ represents the degree of smoothing and \mathbf{u} is a random variable distributed uniformly over $B(\mathbf{0}; 1)$. Also,

$$\mathbf{x}^* := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) \text{ and } \mathbf{x}_\delta^* := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^d} \hat{f}_\delta(\mathbf{x}).$$

There are a total of M smoothed functions in this paper. The largest noise level is δ_1 and the smallest is $\delta_{M+1} = 0$. Thus, $\hat{f}_{\delta_{M+1}} = f$.

Definition 2. *Let $\delta_1 \in \mathbb{R}$. A function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is said to be “new σ -nice” if the following conditions hold:*

(i) *For all $m \in [M]$ and all $\gamma_m \in (0, 1)$, there exist $\delta_m \in \mathbb{R}$ with $|\delta_{m+1}| := \gamma_m |\delta_m|$ and $\mathbf{x}_{\delta_m}^*$ such that*

$$\left\| \mathbf{x}_{\delta_m}^* - \mathbf{x}_{\delta_{m+1}}^* \right\| \leq |\delta_m| - |\delta_{m+1}|.$$

(ii) *For all $m \in [M]$ and all $\gamma_m \in (0, 1)$, there exist $\delta_m \in \mathbb{R}$ with $|\delta_{m+1}| := \gamma_m |\delta_m|$ and $d_m > 1$ such that the function $\hat{f}_{\delta_m}(\mathbf{x})$ is σ -strongly convex on $N(\mathbf{x}^*; d_m \delta_m)$.*

Assumptions and Lemmas

We make the following assumptions:

Assumption 1. (A1) $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is continuously differentiable and L_g -smooth, i.e., for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L_g \|\mathbf{x} - \mathbf{y}\|$. (A2) $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is an L_f -Lipschitz function, i.e., for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $|f(\mathbf{x}) - f(\mathbf{y})| \leq L_f \|\mathbf{x} - \mathbf{y}\|$. (A3) Let $(\mathbf{x}_t)_{t \in \mathbb{N}} \subset \mathbb{R}^d$ be the sequence generated by an optimizer. (i) For each iteration t , $\mathbb{E}_{\xi_t} [\mathbf{G}_{\xi_t}(\mathbf{x}_t)] = \nabla f(\mathbf{x}_t)$. (ii) There exists a nonnegative constant C_{opt}^2 for an optimizer such that $\mathbb{E}_{\xi_t} [\|\mathbf{G}_{\xi_t}(\mathbf{x}_t) - \nabla f(\mathbf{x}_t)\|^2] \leq C_{\text{opt}}^2$. (A4) For

each iteration t , the optimizer samples a mini-batch $\mathcal{S}_t \subset \mathcal{S}$ and estimates the full gradient ∇f as

$$\nabla f_{\mathcal{S}_t}(\mathbf{x}_t) := \frac{1}{b} \sum_{i \in [b]} \mathbf{G}_{\xi_{t,i}}(\mathbf{x}_t) = \frac{1}{b} \sum_{\{i: \mathbf{z}_i \in \mathcal{S}_t\}} \nabla f_i(\mathbf{x}_t).$$

(A5) *There exists a positive constant K_{opt} for an optimizer, for all $t \in \mathbb{N}$, $\mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] \leq K_{\text{opt}}^2$.*

Remark 1. Three algorithms appear in this paper: SGD (Algorithm 2), SHB (Algorithm 4), and NSHB (Algorithm 5). In assumptions (A3)(ii) and (A5), C_{opt}^2 and K_{opt} represent optimizer-specific constants. For example, C_{SGD}^2 is the variance of the stochastic gradient when $(\mathbf{x}_t)_{t \in \mathbb{N}}$ is the sequence generated by SGD.

Explicit Graduated Optimization

Algorithm 1 (Sato and Iiduka 2023) below is an embodiment of explicit graduated optimization, where the noise decay rate $\gamma_m := \frac{(M-m)^p}{\{M-(m-1)\}^p}$ ($p \in (0, 1]$) is that of a polynomial decay scheduler with power p , which is theoretically optimal in the sense that it satisfies the conditions necessary to ensure that the graduated optimization algorithm does not leave the strongly convex region of the new σ -nice function. Algorithm 2 is used to optimize each smoothed function.

Remark 2. In explicit graduated optimization, each smoothed function can be optimized with any algorithm as long as its convergence to a local strongly convex function is guaranteed. We follow the prior work (Sato and Iiduka 2023) and use SGD, which includes GD, in Algorithm 2.

Algorithm 1: Explicit Graduated Optimization

Require: $\epsilon, r, \bar{d}, B_2, H_3, H_4 > 0, p \in (0, 1]$,
 $\mathbf{x}_1 \in \mathbb{R}^d, b \in [n], \eta > 0$
 $\delta_1 := \frac{2L_f}{\sigma r}$
 $\alpha_0 := \min \left\{ \frac{\sigma r}{8L_f^2(1+\bar{d})}, \frac{\sqrt{\sigma r}}{2\sqrt{2}L_f} \right\}, M^p := \frac{1}{\alpha_0 \epsilon}$
for $m = 1$ to $M + 1$ **do**
 if $m \neq M + 1$ **then**
 $\epsilon_m := \sigma \delta_m^2, T_F := H_4 / (\epsilon_m - H_3 \eta m)$
 $\gamma_m := \frac{(M-m)^p}{\{M-(m-1)\}^p}$
 end if
 $\mathbf{x}_{m+1} := \text{SGD}(T_F, \mathbf{x}_m, \hat{f}_{\delta_m}, b, \eta)$
 $\delta_{m+1} := \gamma_m \delta_m$
end for
return \mathbf{x}_{M+2}

Algorithm 2: SGD with constant learning rate

Require: $T_F, \hat{\mathbf{x}}_1^{(m)} \in \mathbb{R}^d, \hat{f}_{\delta_m}, b \in [n], \eta > 0$
for $t = 1$ to T_F **do**
 $\hat{\mathbf{x}}_{t+1}^{(m)} := \hat{\mathbf{x}}_t^{(m)} - \eta \nabla \hat{f}_{\delta_m, \mathcal{S}_t}(\hat{\mathbf{x}}_t^{(m)})$
end for
return $\hat{\mathbf{x}}_{T_F+1}^{(m)} = \text{SGD}(T_F, \hat{\mathbf{x}}_1^{(m)}, \hat{f}_{\delta_m}, b, \eta)$

function's reference	function	EGO(geo)	EGO(nice)	GA	PSO
(Ackley 1987; Bäck and Schwefel 1993)	Ackley's	1.80E+01	4.04E-03	1.52E+00	5.37E+00
(Rahnamayan, Tizhoosh, and Salama 2007)	Alpine1	2.34E-01	1.25E-01	2.59E-01	2.72E-04
(Marcin Molga 2005)	Drop-Wave	9.93E-01	9.94E-01	9.30E-01	9.10E-01
(Yang 2010)	Ellipsoid	4.22E-01	4.82E-04	5.35E+00	1.12E-04
(Griewank 1981)	Griewank	2.55E+00	2.32E-03	9.88E-03	3.86E-02
(Beyer and Finck 2012)	HappyCat	2.39E+01	1.76E+00	1.12E+00	5.81E-01
(Tan 2016)	HGBat	5.07E-01	5.04E-01	1.25E+00	5.86E-01
(Plevris and Solorzano 2022)	Modified Ridge	6.63E+00	6.63E+00	1.73E+00	4.57E-01
(Törn and Zilinskas 1989; Rudolph 1990)	Rastrigin's	1.85E+00	2.26E-02	2.44E+01	1.51E+02
(Rosenbrock 1960; Dixon and Mills 1994)	Rosenbrock's	6.00E+98	2.44E+127	9.74E+01	9.53E+01
(Marcin Molga 2005)	Rotated Hyper-ellipsoid	5.20E-01	4.95E-04	6.73E+00	2.53E-05
(Salomon 1996)	Salomon	7.24E-01	2.02E-01	7.54E-01	2.11E+00
(Schaffer 1985; Schaffer et al. 1989)	Schaffer's F7	2.68E+01	1.12E+01	6.83E-01	1.80E+01
(Schwefel 1981)	Schwefel	8.33E+03	8.33E+03	7.10E+03	7.20E+03
(Schwefel 1981)	Schwefel 2.21	3.76E+01	2.06E-02	2.17E+00	3.34E+01
(Schumer and Steiglitz 1968)	Sphere	6.98E-07	1.58E-05	2.34E-01	3.46E-06

Table 1: Average values of the optimum results over 50 runs, for Algorithm 1, GA, and PSO, for dimensions $D = 50$.

If the objective function f is a new σ -nice function, it is guaranteed that Algorithm 1 will converge to an ϵ -neighborhood of the global optimal solution of f in $\mathcal{O}\left(1/\epsilon^{\frac{1}{p}}\right)$ rounds. Whether Algorithm 1 works effectively with test functions for global optimization has not been considered in a previous study. This paper deals with 16 traditional benchmark functions to verify the performance of global optimization algorithms (see Table 2 for the function's name and Appendix B for their definitions). For example, for all $\mathbf{x} := (x_1, \dots, x_D)^\top \in \mathbb{R}^D$, Rastrigin's function is defined as follows:

$$f(\mathbf{x}) := \sum_{i=1}^D \{x_i^2 - 10 \cos(2\pi x_i)\} + 10D.$$

Figure 6 in Appendix B plots Rastrigin's function for $D = 2$. Theorem 1 shows that this function, which has many local optimal solutions, is a new σ -nice function. The proof of the theorem is in Appendix A.

Theorem 1. *Rastrigin's function is a new σ -nice function.*

Theorem 1 implies that Algorithm 1 can achieve global optimum of Rastrigin's function. To test the algorithm's performance, we optimized the 16 benchmark functions with it (EGO) with optimal decay rates (nice) and with suboptimal decay rates (geo), where $\gamma_m := c$ ($c \in (0, 1)$), i.e., the noise sequence is a geometric progression, and did the same with two well-known global optimization algorithms, i.e., the genetic algorithm (GA) (Holland 1975) and particle swarm optimization (PSO) (Kennedy and Eberhart 1995). The dimension D of the objective functions was set to 50. The experimental environment was Intel Core i9 13900KF CPU.

Table 1 shows the average of the optimum results of 50 runs starting from a random initial point. EGO with both decay rates used GD to optimize the smoothed function. The results for GA and PSO are taken from (Plevris and Solorzano 2022). Note that the optimal value for all func-

tions is 0. See Appendix B for the search range of each objective function and the parameters of the Algorithms. The results in the table indicate that Algorithm 1 (EGO) significantly outperforms the other methods for several objective functions and that EGO with theoretically optimal decay rates (nice) generally performs better than EGO with suboptimal decay rates (geo). Thus, Algorithm 1 with optimal decay rates is effective for traditional global optimization benchmark functions.

However, explicit graduated optimization is not effective for DNNs, such as ResNet (He et al. 2016). Figure 2 shows the results of using Algorithm 1 to train ResNet18 on the CIFAR100 dataset for 200 epochs. Note that it used SGD with $\nabla f_{S_i}(\mathbf{x}_t + \delta_m \mathbf{u}_t)$ as the search direction for decreasing $(\delta_m)_{m \in [M]}$, where $\mathbf{u}_t \in \mathbb{R}^d$ is Gaussian noise and $M = 200$; i.e., the noise δ_m was decreased each epoch. The experimental environment for training the DNN was as follows: NVIDIA GeForce RTX 4090 \times 1GPU.




Figure 2: Loss function value for training versus the number of epochs in training ResNet18 on the CIFAR100 dataset. The solid lines represent the mean value, and the shaded areas represent the maximum and minimum values over three runs.

Compared with vanilla SGD without noise, Algorithm 1

with initial noise values δ_1 of 0.1, 0.01, and 0.001 does not achieve significantly lower loss function values; instead, the noise simply interferes with training. In fact, the loss function values are comparable to those of vanilla SGD only when the noise is nearly zero, and it does not fulfill its role of avoiding locally optimal solutions and yielding lower loss function values. The cause is theoretically unknown, but the experimental validity of similar methods in the previous study (Hazan, Yehuda, and Shalev-Shwartz 2016) suggests that the number of parameters in the model is a factor, i.e., that explicit graduated optimization may not be effective for very high dimensional functions. The number of model parameters in the previous study is 23,860, while ours is about 11.2M. Because of these limitations of explicit graduated optimization, we do not see any benefit to applying it to DNNs.

Implicit Graduated Optimization

Implicit graduated optimization with SGD. As mentioned above, (Sato and Iiduka 2023) showed that simply using SGD generates stochastic noise of level $\delta^{\text{SGD}} = \eta C_{\text{SGD}}^2 / \sqrt{b}$, which smooths the objective function f .

Here, let \mathbf{y}_t be the parameter updated by gradient descent (GD) at time t , and let \mathbf{x}_{t+1} be the parameter updated by SGD, i.e.,

$$\begin{aligned} \mathbf{y}_t &:= \mathbf{x}_t - \eta \nabla f(\mathbf{x}_t), \\ \mathbf{x}_{t+1} &:= \mathbf{x}_t - \eta \nabla f_{\mathcal{S}_t}(\mathbf{x}_t) \\ &= \mathbf{x}_t - \eta (\nabla f(\mathbf{x}_t) + \boldsymbol{\omega}_t^{\text{SGD}}), \end{aligned}$$

where $\boldsymbol{\omega}_t^{\text{SGD}} := \nabla f_{\mathcal{S}_t}(\mathbf{x}_t) - \nabla f(\mathbf{x}_t)$ is stochastic noise of SGD. Then, using Definition 1 and $\mathbb{E}_{\epsilon_t} [\|\boldsymbol{\omega}_t^{\text{SGD}}\|^2] \leq \frac{C_{\text{SGD}}}{\sqrt{b}}$ which holds under Assumptions (A3)(ii) and (A4), we have

$$\mathbb{E}_{\omega_t} [\mathbf{y}_{t+1}] = \mathbb{E}_{\omega_t} [\mathbf{y}_t] - \eta \nabla \hat{f}_{\frac{\eta C}{\sqrt{b}}}(\mathbf{y}_t). \quad (2)$$

See Section 3.3 and Appendix A of (Sato and Iiduka 2023) for details on the derivation of equation (2).

Therefore, optimizing the objective function f by SGD is equivalent to optimizing its smoothed version $\hat{f}_{\delta^{\text{SGD}}}$ by GD in the sense of expectation. On the basis of this fact, (Sato and Iiduka 2023) proposed the following algorithm that implicitly achieves graduated optimization by varying the learning rate η and batch size b so as to decrease the degree of smoothing δ^{SGD} during training. That is, they decreased the learning rate η and increased the batch size b during training, thereby decreasing the degree of smoothing δ^{SGD} . Algorithm 3 is an implicit graduated optimization algorithm that exploits the natural stochastic noise of SGD. The decay rate of δ^{SGD} is set so that the combined effect of decreasing the learning rate $\kappa_m \in (0, 1]$ and increasing the batch size $\lambda_m \geq 1$ is $\gamma_m := \frac{(M-m)^p}{\{M-(m-1)\}^p}$. Algorithm 2 (SGD) is used to optimize each smoothed function.

Algorithm 3: Implicit Graduated Optimization with SGD

Require: $\epsilon, r, \bar{d}, B_2, H_3, H_4 > 0, p \in (0, 1]$,
 $\mathbf{x}_1 \in \mathbb{R}^d, b_1 \in [n], \eta_1 > 0$
 $\delta_1 := \frac{\eta_1 C}{\sqrt{b_1}}$
 $\alpha_0 := \min \left\{ \frac{\sqrt{b_1}}{4L_f \eta_1 C(1+\bar{d})}, \frac{\sqrt{b_1}}{\sqrt{2\sigma \eta_1 C}} \right\}, MP := \frac{1}{\alpha_0 \epsilon}$
for $m = 1$ to $M + 1$ **do**
 if $m \neq M + 1$ **then**
 $\epsilon_m := \sigma^2 \delta_m^2, T_F := H_4 / (\epsilon_m - H_3 \eta_m)$
 $\gamma_m := \frac{(M-m)^p}{\{M-(m-1)\}^p}$
 $\kappa_m / \sqrt{\lambda_m} = \gamma_m (\kappa_m \in (0, 1], \lambda_m \geq 1)$
 end if
 $\mathbf{x}_{m+1} := \text{SGD}(T_F, \mathbf{x}_m, \hat{f}_{\delta_m}, \eta_m, b_m)$
 $\eta_{m+1} := \kappa_m \eta_m, b_{m+1} := \lambda_m b_m$
 $\delta_{m+1} := \frac{\eta_{m+1} C}{\sqrt{b_{m+1}}}$
end for
return \mathbf{x}_{M+2}

Remark 3. In implicit graduated optimization, the algorithm that optimizes each smoothed function must be GD or SGD. Unlike explicit graduated optimization, which explicitly provides a smoothed function, based on (2), the smoothed function can only be optimized by a GD- or SGD-like algorithm. It is unclear which, GD or SGD, is better. Our analysis is valid in both cases, since we provide an analysis of SGD, an extension of GD.

If the objective function f is a new σ -nice function, it is guaranteed that Algorithm 3 converges to an ϵ -neighborhood of the global optimal solution of f in $\mathcal{O}\left(1/\epsilon^{\frac{1}{p}}\right)$ rounds. Since the prior study lacked numerical experiments to verify the performance of Algorithm 3, here, we provide experimental results that are faithful to the decay rate γ_m employed in Algorithm 3. Figure 3 shows the results of training ResNet18 on the CIFAR100 dataset with Algorithm 3 with a constant learning rate and constant batch size, i.e. $\kappa_m = 1, \lambda_m = 1$ (method 1, i.e., vanilla SGD), a decaying learning rate and constant batch size, i.e., $\kappa_m = \gamma_m$ and $\lambda_m = 1$ (method 2), a constant learning rate and increasing batch size, $\kappa_m = 1, \lambda_m = 1/\gamma_m^{13/10}$ (method 3), and a hybrid setting, i.e., $\kappa_m = \gamma_m^{4/9}$ and $\lambda_m = 1/\gamma_m$, where the decay rate is set to $\gamma_m = \frac{(M-m)^{0.9}}{\{M-(m-1)\}^{0.9}}$ ($M = 200, m \in [M]$) (method 4). In a 200-epoch training, methods 2, 3, and 4 update the hyperparameters every epoch. In method 1, the learning rate and the batch size are fixed at 0.1 and 128, respectively. In method 2, the initial learning rate is 0.1 and the batch size is fixed at 128. In method 3, the learning rate is fixed at 0.1 and the initial batch size is 32. In method 4, the initial learning rate is 0.1 and the initial batch size is 32.

Figure 3 indicates that methods 2, 3, and 4 yield a lower loss function and higher test accuracy than those of vanilla SGD (method 1). In particular, the fact that methods 3 and 4, which increase the batch size, are superior in both loss function value and test accuracy suggests that a high learning rate is necessary for successful training. In fact, the loss curve for




Figure 3: Accuracy score in testing and loss function value in training versus the number of epochs in training ResNet18 on the CIFAR100 dataset with SGD. The solid lines represent the mean value, and the shaded areas represent the maximum and minimum values over three runs.

method 2, which decreases the learning rate, is nearly identical to that of methods 3 and 4 until the middle of training.

Implicit graduated optimization with SGD with momentum. There are a number of variants of SGD with momentum; this paper focuses on the simplest ones, i.e., the stochastic heavy ball (SHB) method and normalized stochastic heavy ball (NSHB) method, which are defined as follows:

Algorithm 4: Stochastic Heavy Ball (SHB)

Require: $x_0, \eta > 0, \beta \in [0, 1), m_{-1} := \mathbf{0}$
for $t = 0$ to $T - 1$ **do**
 $\mathbf{m}_t := \nabla f_{S_t}(\mathbf{x}_t) + \beta \mathbf{m}_{t-1}$
 $\mathbf{x}_{t+1} := \mathbf{x}_t - \eta \mathbf{m}_t$
end for
return \mathbf{x}_T

Algorithm 5: Normalized Stochastic Heavy Ball (NSHB)

Require: $x_0, \eta > 0, \beta \in [0, 1), d_{-1} := \mathbf{0}$
for $t = 0$ to $T - 1$ **do**
 $\mathbf{d}_t := (1 - \beta) \nabla f_{S_t}(\mathbf{x}_t) + \beta \mathbf{d}_{t-1}$
 $\mathbf{x}_{t+1} := \mathbf{x}_t - \eta \mathbf{d}_t$
end for
return \mathbf{x}_T

(Sato and Iiduka 2024) showed that the SHB and NSHB methods each generate stochastic noise of level,

$$\delta^{\text{SHB}} = \eta \sqrt{\left(1 + \hat{\beta}\right) \frac{C_{\text{SHB}}^2}{b} + \hat{\beta} K_{\text{SHB}}^2},$$

$$\delta^{\text{NSHB}} = \eta \sqrt{\frac{1}{1 - \beta} \frac{C_{\text{NSHB}}^2}{b}},$$

which smooths the objective function f , where $\hat{\beta} := \frac{\beta(\beta^2 - \beta + 1)}{(1 - \beta)^2}$. That is, at time t , let \mathbf{y}_t be the parameter updated by GD and \mathbf{z}_{t+1} be the parameter updated by SHB.

Then, the following holds:

$$\mathbb{E}_{\omega_t^{\text{SHB}}} [\mathbf{y}_{t+1}] = \mathbb{E}_{\omega_t^{\text{SHB}}} [\mathbf{y}_t] - \eta \nabla f_{\delta^{\text{SHB}}}(\mathbf{y}_t), \quad (3)$$

where $\omega_t^{\text{SHB}} := \mathbf{m}_t - \nabla f(\mathbf{x}_t)$ is stochastic noise of SHB. Since optimizing the objective function f with SHB or NSHB is equivalent (in the sense of expectation) to optimizing its smoothed versions, $\hat{f}_{\delta^{\text{SHB}}}$ and $\hat{f}_{\delta^{\text{NSHB}}}$, with GD, we can use the natural stochastic noise of SHB and NSHB to construct an implicit graduated optimization algorithm. That is, the degree of smoothing, δ^{SHB} and δ^{NSHB} , is decreased by decreasing the learning rate η and momentum factor β and increasing the batch size b during training.

Algorithm 6: Implicit Graduated Optimization with SHB

Require: $\epsilon > 0, p \in (0, 1], \bar{d} > 0, \mathbf{x}_1 \in \mathbb{R}^d, \eta_1 > 0,$
 $b_1 \in [n], \beta_1 \in [0, 1), C_{\text{SHB}}^2 > 0, K_{\text{SHB}} > 0$
 $\hat{\beta}_t := \frac{\beta_t(\beta_t^2 - \beta_t + 1)}{(1 - \beta_t)^2} (\forall t \in \mathbb{N})$
 $\delta_1^{\text{SHB}} := \eta_1 \sqrt{\left(1 + \hat{\beta}_1\right) \frac{C_{\text{SHB}}^2}{b_1} + \hat{\beta}_1 K_{\text{SHB}}^2}$
 $\alpha_0 := \min \left\{ \frac{1}{4L_f |\delta_1^{\text{SHB}}| (1 + \bar{d})}, \frac{1}{\sqrt{2\sigma} |\delta_1^{\text{SHB}}|} \right\}, M^p := \frac{1}{\alpha_0 \epsilon}$
for $m = 1$ to $M + 1$ **do**
if $m \neq M + 1$ **then**
 $\epsilon_m := \sigma^2 \delta_m^{\text{SHB}^2}, T_F := H_4 / (\epsilon_m - H_3 \eta_m)$
 $\gamma_m := \frac{(M - m)^p}{\{M - (m - 1)\}^p}$
 $\frac{\kappa_m \eta_m \sqrt{\left(1 + \rho_m \hat{\beta}_m\right) \frac{C_{\text{SHB}}^2}{\lambda_m b_m} + \rho_m \hat{\beta}_m K_{\text{SHB}}^2}}{\eta_m \sqrt{\left(1 + \hat{\beta}_m\right) \frac{C_{\text{SHB}}^2}{b_m} + \hat{\beta}_m K_{\text{SHB}}^2}} = \gamma_m$
 $(\kappa_m, \rho_m \in (0, 1], \lambda_m \geq 1)$
end if
 $\mathbf{x}_{m+1} := \text{SGD}(T_F, \mathbf{x}_m, \hat{f}_{\delta_m}, \eta_m, b_m)$
 $\eta_{m+1} := \kappa_m \eta_m, b_{m+1} := \lambda_m b_m, \hat{\beta}_{m+1} := \rho_m \hat{\beta}_m$
 $\delta_{m+1}^{\text{SHB}} := \eta_{m+1} \sqrt{\left(1 + \hat{\beta}_{m+1}\right) \frac{C_{\text{SHB}}^2}{b_{m+1}} + \hat{\beta}_{m+1} K_{\text{SHB}}^2}$
end for
return \mathbf{x}_{M+2}

Algorithm 6 is an implicit graduated optimization that exploits the natural stochastic noise of SHB. The decay rate of δ^{SHB} is set so that the combined effect of decreasing the learning rate $\kappa_m \in (0, 1]$ and momentum factor $\rho \in (0, 1]$ and increasing the batch size $\lambda_m \geq 1$ is $\gamma_m := \frac{(M - m)^p}{\{M - (m - 1)\}^p}$.

The following theorem guarantees the convergence of Algorithm 6 for the new σ -nice function (See Appendix C for details on the proof of Theorem 2).

Theorem 2 (Convergence analysis of Algorithm 6). *Let $\epsilon \in (0, 1]$ and f be an L_f -Lipschitz new σ -nice function. Suppose that we apply Algorithm 3; after $\mathcal{O}\left(1/\epsilon^{\frac{1}{p}}\right)$ rounds. Then, the algorithm reaches an ϵ -neighborhood of the global optimal solution \mathbf{x}^* .*

To test the ability of Algorithm 6 to reduce stochastic noise, we compared it with a vanilla SHB method in which

the learning rate, batch size, and momentum are all constant. Here, Algorithm 6 used a noise reduction method in which the hyperparameters are updated to reduce the degree of smoothing δ^{SHB} .




Figure 4: Accuracy score in testing and loss function value in training ResNet18 on the CIFAR100 dataset with Algorithm 6 versus the number of epochs. The blue plot represents vanilla SHB, and the other five plots represent Algorithm 6. “lr” means the learning rate. The solid lines represent the mean value, and the shaded areas represent the maximum and minimum values over three runs.

Figure 4 plots the accuracy in testing and the loss function value in training ResNet18 on the CIFAR100 dataset with SHB versus the number of epochs. Here, Algorithm 6 outperformed vanilla SHB in both test accuracy and loss function value, thereby demonstrating that it is superior to SGD with momentum using constant parameters on image classification tasks. We found similar results in training WideResNet-28-10 on the CIFAR100 dataset with SHB (see Figure 8 in Appendix D). See Figure 9 for the results of comparative experiments with warmup and other methods.

Appendix E analyzes the implicit graduated optimization algorithm using stochastic noise in NSHB and presents the results of numerical experiments.

Optimal decaying learning rates for SGD with momentum. In graduated optimization, the optimal decay rate of noise is $\gamma_m := \frac{(M-1)^p}{\{M-(m-1)\}^p}$ ($p \in (0, 1]$) (Sato and Iiduka 2023). In implicit graduated optimization, the stochastic noise δ^{opt} is proportional to the learning rate η , so the optimal learning rate decay rate is likewise γ_m , i.e., a polynomial decay with a power less than $p = 1$. (Sato and Iiduka 2023) conducted experiments on the optimal learning rate for SGD and found that a polynomial decay with a power less than $p = 1$ achieves the smallest loss function value and yields the highest test accuracy, as theory suggests.

In order to see which learning rate scheduler gives the smallest loss function value for SGD with momentum, we trained ResNet34 (He et al. 2016) on the ImageNet dataset (Deng et al. 2009) with SHB for 100 epochs. The results in Figure 5 indicate that a polynomial decay with a power less than or equal to 1 achieves the smallest training loss function. We confirmed similar results in training ResNet18 and WideResNet-28-10 on the CIFAR100 dataset with SHB and NSHB (see Figure 14-17 in Appendix F).




Figure 5: Accuracy score in testing and loss function value in training ResNet34 on the ImageNet dataset with SHB (Algorithm 4) versus the number of epochs. The solid lines represent the mean value, and the shaded areas represent the maximum and minimum values over three runs.

Conclusion

We showed that Rastrigin’s function, one of the classical benchmark functions, is a new σ -nice function and that explicit graduated optimization with optimal noise scheduling is valid for traditional benchmark functions, but not effective for DNNs. Then, we developed an implicit graduated optimization algorithm using stochastic noise in SGD with momentum and analyzed its convergence. Finally, we conducted experiments on ImageNet on the optimal decaying learning rate scheduler for SGD with momentum and found that the theoretically optimal polynomial decay scheduler with a small power achieves the lowest loss function value. The most important advantage of our approach is that we can theoretically guarantee that SGD with momentum will converge to a global optimal solution only by increasing or decreasing hyperparameters. This has not been achieved by any other method. The contributions that were made by extending the implicit graduated optimization approach of SGD to that of SGD with momentum with stochastic noise include the following:

- Previous theories have been able to guarantee convergence to a local optimal solution for both $\beta \rightarrow 0$ and $\beta \rightarrow 1$ (Gitman et al. 2019). By introducing the momentum factor into the implicit graduated optimization approach, we theoretically show that $\beta \rightarrow 0$ contributes to convergence to the global optimal solution, and we demonstrated its effectiveness in experiments.
- From the fact that the degree of smoothing in SGD with momentum is determined by hyperparameters (see Figure 1), when graduated optimization is introduced, the knowledge of the optimal decay rate of the degree of smoothing leads directly to knowledge on the optimal decay of hyperparameters such as the learning rate and momentum factor.

These findings can only be derived from an implicit graduated optimization perspective, and they should be of help to users in setting hyperparameters intuitively.

Acknowledgements

We are sincerely grateful to Program Chairs, Area Chairs, and anonymous reviewers for helping us improve the original manuscript. This research is partly supported by the computational resources of the DGX A100 named TAIHO at Meiji University. Naoki Sato was supported by Tsuchiya Cultural Foundation. This work was supported by the Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Number 24K14846 awarded to Hideaki Iiduka.

References

- Ackley, D. H. 1987. *A connectionist machine for genetic hillclimbing*. Kluwer Academic Publishers.
- Antonante, P.; Tzoumas, V.; Yang, H.; and Carlone, L. 2022. Outlier-Robust Estimation: Hardness, Minimally Tuned Algorithms, and Applications. *IEEE Transactions on Robotics*, 38(1): 281–301.
- Bäck, T.; and Schwefel, H. 1993. An Overview of Evolutionary Algorithms for Parameter Optimization. *Evolutionary Computation*, 1(1): 1–23.
- Bengio, Y. 2009. Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning*, 2(1): 1–127.
- Beyer, H.-G.; and Finck, S. 2012. HappyCat – A Simple Function Class Where Well-Known Direct Search Algorithms Do Fail. In *Parallel Problem Solving from Nature - PPSN XII*, 367–376. Springer Berlin Heidelberg.
- Blake, A.; and Zisserman, A. 1987. *Visual Reconstruction*. MIT Press.
- Chapelle, O.; Chi, M.; and Zien, A. 2006. A continuation method for semi-supervised SVMs. In *Proceedings of the 23rd International Conference on Machine Learning*, volume 148, 185–192.
- Chapelle, O.; Sindhvani, V.; and Keerthi, S. S. 2008. Optimization Techniques for Semi-Supervised Support Vector Machines. *Journal of Machine Learning Research*, 9: 203–233.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 248–255.
- Dixon, L. C. W.; and Mills, D. J. 1994. Effect of rounding errors on the variable metric method. *Journal of Optimization Theory and Applications*, 80(1): 175–179.
- Duchi, J. C.; Bartlett, P. L.; and Wainwright, M. J. 2012. Randomized Smoothing for Stochastic Optimization. *SIAM Journal on Optimization*, 22(2): 674–701.
- Gitman, I.; Lang, H.; Zhang, P.; and Xiao, L. 2019. Understanding the Role of Momentum in Stochastic Gradient Methods. In *Advances in Neural Information Processing Systems*, volume 32, 9630–9640.
- Griewank, A. O. 1981. Generalized descent for global optimization. *Journal of Optimization Theory and Applications*, 34(1): 11–39.
- Hazan, E.; Yehuda, K.; and Shalev-Shwartz, S. 2016. On Graduated Optimization for Stochastic Non-Convex Problems. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, 1833–1841.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Proceedings of the 34th Conference on Neural Information Processing Systems*.
- Holland, J. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- Iwakiri, H.; Wang, Y.; Ito, S.; and Takeda, A. 2022. Single Loop Gaussian Homotopy Method for Non-convex Optimization. In *Proceedings of the 36th Conference on Neural Information Processing Systems*.
- Kennedy, J.; and Eberhart, R. 1995. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, 1942–1948.
- Li, D.; Wu, J.; and Zhang, Q. 2023. Stochastic Gradient Descent in the Viewpoint of Graduated Optimization. <https://arxiv.org/abs/2308.06775>.
- Marcin Molga, C. S. 2005. Test functions for optimization needs. <https://robertmarks.org/Courses/ENGR5358/Papers/functions.pdf>.
- Mobahi, H.; and Fisher III, J. W. 2015. A Theoretical Analysis of Optimization by Gaussian Continuation. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence*, 1205–1211.
- Peng, L.; Kümmerle, C.; and Vidal, R. 2023. On the Convergence of IRLS and Its Variants in Outlier-Robust Estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17808–17818.
- Plevris, V.; and Solorzano, G. 2022. A Collection of 30 Multidimensional Functions for Global Optimization Benchmarking. *Data*, 7(4).
- Rahnamayan, S.; Tizhoosh, H. R.; and Salama, M. M. A. 2007. A novel population initialization method for accelerating evolutionary algorithms. *Computers and Mathematics with Applications*, 53(10): 1605–1614.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22: 400–407.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Rosenbrock, H. H. 1960. An Automatic Method for Finding the Greatest or Least Value of a Function. *Computer Journal*, 3(3): 175–184.
- Rudolph, G. 1990. *Globale Optimierung mit parallelen evolutionsstrategien*. Ph.D. thesis, Universität Dortmund Fachbereich Informatik.
- Salomon, R. 1996. Re-evaluating genetic algorithm performance under coordinate rotation of benchmark functions. A survey of some theoretical and practical aspects of genetic algorithms. *Biosystems*, 39(3): 263–278.
- Sato, N.; and Iiduka, H. 2023. Using Stochastic Gradient Descent to Smooth Nonconvex Functions: Analysis of Implicit Graduated Optimization with Optimal Noise Scheduling. <https://arxiv.org/abs/2311.08745>.

Sato, N.; and Iiduka, H. 2024. Role of Momentum in Smoothing Objective Function and Generalizability of Deep Neural Networks. <https://arxiv.org/abs/2402.02325>.

Schaffer, J. D. 1985. Multiple Objective Optimization with Vector Evaluated Genetic Algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, 93–100.

Schaffer, J. D.; Caruana, R. A.; Eshelman, L. J.; and Das, R. 1989. A study of control parameters affecting online performance of genetic algorithms for function optimization. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, 51–60.

Schumer, M.; and Steiglitz, K. 1968. Adaptive step size random search. *IEEE Transactions on Automatic Control*, 13(3): 270–276.

Schwefel, H.-P. 1981. *Numerical Optimization of Computer Models*. John Wiley & Sons, Inc.

Sindhwani, V.; Keerthi, S. S.; and Chapelle, O. 2006. Deterministic annealing for semi-supervised kernel machines. In *Proceedings of the 23rd International Conference on Machine Learning*, volume 148, 841–848.

Sohl-Dickstein, J.; Weiss, E. A.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 2256–2265.

Song, Y.; and Ermon, S. 2019. Generative Modeling by Estimating Gradients of the Data Distribution. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 11895–11907.

Song, Y.; and Ermon, S. 2020. Improved Techniques for Training Score-Based Generative Models. In *Proceedings of the 34th Conference on Neural Information Processing Systems*.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *Proceedings of the 9th International Conference on Learning Representations*.

Tan, Y. 2016. *GPU-based Parallel Implementation of Swarm Intelligence Algorithms*. Morgan Kaufmann Publishers Inc.

Törn, A. A.; and Zilinskas, A. 1989. *Global Optimization*, volume 350 of *Lecture Notes in Computer Science*. Springer.

Yang, H.; Antonante, P.; Tzoumas, V.; and Carlone, L. 2020. Graduated Non-Convexity for Robust Spatial Perception: From Non-Minimal Solvers to Global Outlier Rejection. *IEEE Robotics and Automation Letters*, 5(2): 1127–1134.

Yang, X.-S. 2010. *Engineering Optimization: An Introduction with Metaheuristic Applications*. John Wiley & Sons, Inc.