論文 / 著書情報
Article / Book Information

題目(和文)	
Title(English)	Learning Rate Adaptation for Evolution Strategies
著者(和文)	野村将寛
Author(English)	Masahiro Nomura
出典(和文)	学位:博士(工学), 学位授与機関:東京科学大学, 報告番号:甲第366号, 授与年月日:2025年3月26日, 学位の種別:課程博士, 審査員:小野 功,三宅 美博,山村 雅幸,瀧ノ上 正浩,小野 峻佑
Citation(English)	Degree:Doctor (Engineering), Conferring organization: Institute of Science Tokyo, Report number:甲第366号, Conferred date:2025/3/26, Degree Type:Course doctor, Examiner:,,,,
学位種別(和文)	博士論文
Category(English)	Doctoral Thesis
種別(和文)	論文要旨
Type(English)	Summary

（博士課程）
Doctoral Program

論 文 要 旨

THESIS SUMMARY

系・コース： Department of, Graduate major in	情報工学 系 知能情報 コース	申請学位（専攻分野）： Academic Degree Requested	博士 Doctor of （ 工学 ）
学生氏名： Student's Name	野村 将寛	審査員主査： Chief Examiner	小野 功

要旨（英文800語程度）
Thesis Summary（approx.800 English Words ）

Black-box continuous optimization is widely used for real-world problems because it does not require explicit information about the objective function (e.g., gradients or differentiability). This flexibility makes it suitable for complex problems where solution evaluation involves expensive numerical simulations. As a result, techniques to accelerate the optimization process, such as parallel evaluations, are important in practice. Among the factors that make black-box continuous optimization challenging, multimodality and noise are particularly significant.

Evolution strategies (ES) is a promising framework for black-box continuous optimization problems. ES employs a multivariate Gaussian distribution, parameterized by a mean vector and covariance matrix, to guide optimization. In each iteration, ES generates multiple candidate solutions, with the population size indicating the number of solutions generated. These candidate solutions are evaluated on the objective function and ranked. ES then updates the distribution parameters based on this ranking and a learning rate. This process repeats until a predefined stopping criterion is met. The objective function value and distribution parameters are used to define the stopping criterion. For example, the algorithm may stop when the best evaluation value fails to improve over a give period, or when the covariance matrix sufficiently shrinks (indicating convergence) or expands (indicating divergence). Prominent methods within ES include exponential natural evolution strategies (xNES) and covariance matrix adaptation evolution strategy (CMA-ES), both of which are partially explained by the information geometric optimization (IGO) framework.

A key practical issue for ES is determining the appropriate population size. While a smaller population size generally performs well for many unimodal problems, increasing the population size can be beneficial for more difficult tasks, such as those involving multimodal landscapes and additive noise. However, in a black-box scenario, understanding the problem structure of the objective function is challenging, which makes selecting the appropriate population size in advance is also challenging. To address this, online adaptation of the population size has been proposed to address the issue.

It has been observed that, in CMA-ES, increasing the population size has a similar effect to decreasing the learning rate for the mean vector. Inspired by this observation, this study focuses on learning rate adaptation rather than population size adaptation. We argue that the learning rate adaptation is more advantageous than the population size adaptation from a practical perspective, as the former is better suited for parallel implementations. For example, practitioners often want to set the population size to match a specific number of workers to avoid wasting computational resources. However, the population size adaptation may not always utilize resources efficiently, as the population size can fluctuate throughout the optimization process. In contrast, learning rate adaptation enables full exploitation of resources by fixing the population size at the maximum number of workers. Furthermore, with learning rate adaptation, parameter updates occur regularly, whereas ES with population size adaptation does not progress until all evaluations are complete, complicating the determination of an appropriate termination point.

This study consists of two works. In the first work, we propose a learning rate adaptation method aimed at accelerating optimization, using xNES as the optimization method. With a population size sufficient to solve the problem, the proposed method measures tendencies of updates in the distribution parameters and accelerates the optimization by increasing the learning rate appropriately when a sufficient tendency is detected. Our method enables a larger learning rate for relatively easy problems, resulting in faster search. Conversely, for more difficult problems (e.g., multimodal problems), it allows for a conservative learning rate, leading to a robust and stable search. Experimental evaluations on both unimodal and multimodal problems demonstrate that the proposed method works properly depending on a search situation and is effective over existing methods, such as those using a fixed learning rate.

In the second work, we propose a learning rate adaptation method for solving difficult tasks, such as multimodal or noisy problems, using CMA-ES as the optimization method. Unlike the first work, the second work accepts any population size without assuming it is sufficiently large. This study comprehensively explores the learning rate impact on IGO to demonstrate the necessity of a small learning rate by considering ordinary differential equations. Thereafter, it discusses the setting of an ideal learning rate. Based on these discussions, we develop a novel learning rate adaptation mechanism for CMA-ES that maintains a constant signal-to-noise ratio. Additionally, we investigate the behavior of CMA-ES with the proposed learning rate adaptation mechanism through numerical experiments and compare the results with those obtained for CMA-ES with a fixed learning rate and with population size adaptation. The results show that CMA-ES with the proposed learning rate adaptation works well for multimodal and/or noisy problems without extremely expensive learning rate tuning.

While the proposed method in the second work was designed to solve difficult problems safely, we believe it can also be extended to accelerate optimization, aligning with the goals of the first work. This study marks an important step toward developing fully hyperparameter-free ES algorithms for general-purpose optimization.