



(REVIEW ARTICLE)



## Machine learning mathematical optimization techniques for enhancing model efficiency and convergence

Mamatha N<sup>1</sup> and Bhuvaneshwari Shetty<sup>2,\*</sup>

<sup>1</sup> Lecturer in Science Department, Karnataka (Govt) Polytechnic Mangalore, Karnataka, India.

<sup>2</sup> Lecturer in Computer Science Department, Government Polytechnic for Women, Bondel Mangalore, Karnataka, India.

World Journal of Advanced Research and Reviews, 2021, 10(03), 471-481

Publication history: Received on 17 April 2021; revised on 30 May 2021; accepted on 02 June 2021

Article DOI: <https://doi.org/10.30574/wjarr.2021.10.3.0235>

### Abstract

Mathematical optimization plays a crucial role in machine learning, providing the foundation for efficient model training, parameter tuning, and convergence improvement. Effective optimization techniques enhance the learning process by minimizing loss functions, improving generalization, and accelerating convergence rates. This paper explores various optimization methods, including convex and non-convex optimization, gradient-based approaches such as stochastic gradient descent (SGD), Adam, and RMSprop, as well as gradient-free techniques like evolutionary algorithms and Bayesian optimization. We analyze their theoretical foundations, computational complexity, and practical implications in different machine learning tasks, including supervised and unsupervised learning, reinforcement learning, and deep learning. Furthermore, we present a comparative analysis of these optimization strategies, supported by mathematical formulations, empirical results, and visual representations such as equations, figures, and tables. The findings of this study aim to provide insights into selecting appropriate optimization techniques based on problem characteristics, model architecture, and computational constraints. Finally, we discuss emerging trends in optimization, including second-order methods, meta-learning approaches, and hybrid optimization frameworks, highlighting their potential to further enhance model efficiency and convergence in real-world applications.

**Keywords:** Mathematical Optimization; Machine Learning, Gradient Descent; Stochastic Gradient Descent (SGD); Adam Optimizer; Convex and Non-Convex Optimization; Gradient-Free Optimization

### 1. Introduction

Optimization is a fundamental aspect of machine learning, significantly influencing model performance, generalization, and computational efficiency. In machine learning, optimization refers to the process of finding the best parameters for a model to minimize a predefined objective function, such as a loss function in supervised learning. Efficient optimization techniques enable models to learn meaningful patterns from data, improve predictive accuracy, and enhance robustness against overfitting. As machine learning models grow in complexity, selecting and implementing appropriate optimization strategies becomes increasingly critical to ensuring efficient training and deployment. Optimization techniques in machine learning can be broadly classified into convex and non-convex optimization methods. Convex optimization problems, where the objective function has a single global minimum, allow for efficient and deterministic solutions. In contrast, non-convex optimization problems, which are more common in deep learning, contain multiple local minima and saddle points, making optimization more challenging. Gradient-based methods, such as stochastic gradient descent (SGD) and its variants, are widely used due to their efficiency in handling large datasets and high-dimensional parameter spaces. However, they often require careful tuning of hyperparameters such as learning rate and momentum to achieve optimal performance. Beyond gradient-based methods, gradient-free techniques have also gained attention for optimizing machine learning models, particularly in scenarios where

\* Corresponding author: Bhuvaneshwari Shetty

gradients are difficult to compute or unreliable. These methods, including genetic algorithms, Bayesian optimization, and particle swarm optimization, offer alternative approaches for optimizing complex functions. Additionally, second-order optimization methods, such as Newton's method and quasi-Newton methods, leverage second-order derivative information to improve convergence speed, though they often come at a higher computational cost. The choice of optimization technique depends on various factors, including model architecture, dataset size, and computational constraints. Despite the advancements in optimization methods, several challenges persist in training complex machine learning models. Issues such as vanishing and exploding gradients, slow convergence rates, and sensitivity to hyperparameter selection remain active areas of research. For instance, deep neural networks with multiple layers often suffer from vanishing gradients, where gradients become too small to facilitate effective weight updates. Techniques like batch normalization, adaptive learning rates, and advanced initialization methods have been developed to mitigate these challenges. Additionally, optimizing models in resource-constrained environments, such as edge computing and mobile applications, requires lightweight and energy-efficient optimization strategies[1].

Several studies have explored optimization strategies in machine learning, contributing to the development of robust algorithms. Early work on convex optimization by Boyd and Vandenberghe (2004) laid the foundation for modern optimization techniques in machine learning. More recent studies have focused on adaptive gradient methods such as Adam (Kingma & Ba, 2015) and RMSprop (Tieleman & Hinton, 2012), which dynamically adjust learning rates to improve convergence. Research on evolutionary algorithms for hyperparameter tuning, such as Google's AutoML framework, has demonstrated the effectiveness of automated optimization in model selection and performance enhancement. Additionally, advances in reinforcement learning-based optimization strategies have opened new avenues for optimizing complex, high-dimensional models. Another critical area of research in optimization for machine learning is meta-learning, where models learn to optimize themselves by adapting optimization strategies based on prior experience. Meta-learning techniques, such as model-agnostic meta-learning (MAML) proposed by Finn et al. (2017), enable faster convergence and improved generalization across tasks. Similarly, hybrid optimization methods that combine gradient-based and gradient-free techniques have shown promise in achieving better performance in complex machine learning scenarios. These approaches leverage the strengths of different optimization paradigms to enhance efficiency and robustness. With the rapid evolution of machine learning models, optimization techniques continue to evolve to meet growing demands for scalability and efficiency. The integration of distributed optimization techniques has become essential for training large-scale models on cloud and high-performance computing platforms. Techniques such as decentralized SGD and federated learning enable collaborative optimization across multiple devices while preserving data privacy. Furthermore, hardware-aware optimization methods, which tailor optimization strategies to specific hardware architectures (e.g., GPUs, TPUs, and neuromorphic processors), are becoming increasingly relevant in accelerating model training and inference. In this paper, we provide a comprehensive analysis of mathematical optimization techniques in machine learning, focusing on their theoretical foundations, practical applications, and comparative performance. We examine the strengths and limitations of different optimization strategies and discuss emerging trends in optimization research. By presenting a structured evaluation of various techniques, supported by mathematical formulations and empirical analysis, we aim to offer valuable insights into selecting appropriate optimization methods for different machine learning applications. The subsequent sections delve into the core principles of optimization, specific techniques, and their applications in real-world machine learning problems[2].

## 2. Optimization in Machine Learning

Optimization plays a critical role in machine learning by minimizing objective functions, such as loss functions, to improve model performance. This section categorizes optimization techniques into convex vs. non-convex optimization, gradient-based methods, and gradient-free approaches[3].

### 2.1. Convex vs. Non-Convex Optimization

#### 2.1.1. Convex Optimization

Convex optimization deals with problems where the objective function is convex, meaning any local minimum is also a global minimum. These problems are easier to solve using well-established optimization techniques such as gradient descent and second-order methods. Mathematically, a function  $f(x)$  is convex if:

$$f(\lambda x_1 + (1 - \lambda)x_2) \leq \lambda f(x_1) + (1 - \lambda)f(x_2), \quad \forall x_1, x_2 \in \mathbb{R}^n, \quad \lambda \in [0, 1]$$

One example is linear regression, where the objective function is convex and can be minimized efficiently using gradient descent.

### 2.1.2. Non-Convex Optimization

Non-convex optimization is common in deep learning, where objective functions have multiple local minima, saddle points, and potentially chaotic behavior. A typical deep learning loss function, such as cross-entropy for classification, is non-convex:

$$L = - \sum_{i=1}^N y_i \log(\hat{y}_i)$$

where  $y_i$  represents the true labels, and  $\hat{y}_i$  represents predicted probabilities. Due to the presence of multiple local minima, optimization requires advanced techniques such as momentum-based methods and adaptive learning rate strategies.

## 2.2. Gradient-Based Optimization

Gradient-based optimization methods use gradients of the loss function to update parameters iteratively, allowing the model to converge towards an optimal solution.

### 2.2.1. Gradient Descent (GD)

Gradient Descent updates the model parameters  $\theta$  iteratively by moving in the direction of the negative gradient of the loss function  $J(\theta)$ :

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla J(\theta^{(t)})$$

where:

- $\eta$  is the learning rate,
- $\nabla J(\theta)$  is the gradient of the objective function.

Though effective, standard GD requires computing gradients over the entire dataset, making it computationally expensive for large-scale problems.

### 2.2.2. Stochastic Gradient Descent (SGD)

SGD improves efficiency by updating weights after each mini-batch rather than the entire dataset:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla J(\theta^{(t)}; x_i, y_i)$$

Where  $(x_i, y_i)$  represents a single sample or mini-batch, reducing computational cost but introducing stochasticity.

## 2.3. Momentum-Based Optimization

Momentum methods accelerate convergence by incorporating past gradients into weight updates:

$$v^{(t+1)} = \beta v^{(t)} + (1 - \beta) \nabla J(\theta^{(t)})$$

$$\theta^{(t+1)} = \theta^{(t)} - \eta v^{(t+1)}$$

where:

- $v^{(t)}$  is the velocity term,
- $\beta$  (e.g., 0.9) determines how much past gradients influence updates.

This helps models escape saddle points and smooth out noisy updates in SGD.

### 2.3.1. Adam Optimizer

Adam (Adaptive Moment Estimation) combines momentum with adaptive learning rates:

$$m^{(t)} = \beta_1 m^{(t-1)} + (1 - \beta_1) \nabla J(\theta^{(t)})$$

$$v^{(t)} = \beta_2 v^{(t-1)} + (1 - \beta_2) (\nabla J(\theta^{(t)}))^2$$

$$\hat{m}^{(t)} = \frac{m^{(t)}}{1 - \beta_1^t}, \quad \hat{v}^{(t)} = \frac{v^{(t)}}{1 - \beta_2^t}$$

$$\theta^{(t+1)} = \theta^{(t)} - \frac{\eta}{\sqrt{\hat{v}^{(t)} + \epsilon}} \hat{m}^{(t)}$$

where

- $m$  and  $v$  are the first and second moment estimates,
- $\beta_1, \beta_2$  control exponential moving averages,
- $\epsilon$  prevents division by zero.

Adam adapts learning rates per parameter, making it well-suited for deep learning models.

## 2.4. Gradient-Free Optimization

Gradient-free optimization methods are useful when gradients are expensive or difficult to compute, such as in reinforcement learning or black-box optimization[4].

### 2.4.1. Evolutionary Algorithms

Evolutionary algorithms mimic natural selection processes to optimize machine learning models. Genetic algorithms (GA) use crossover, mutation, and selection to evolve better solutions. A simplified update step is:

$$P_{t+1} = \text{Selection}(\text{Mutation}(\text{Crossover}(P_t)))$$

Where  $P_{t+1}$  represents the population at iteration  $t+1$ .

### Particle Swarm Optimization (PSO)

PSO optimizes functions by simulating a swarm of particles moving through the search space:

$$v_i^{(t+1)} = \omega v_i^{(t)} + c_1 r_1 (p_i - x_i^{(t)}) + c_2 r_2 (g - x_i^{(t)})$$

$$x_i^{(t+1)} = x_i^{(t)} + v_i^{(t+1)}$$

where:

- $x_i$  and  $v_i$  are the position and velocity of particle  $i$ ,
- $p_i$  is the particle's best position found so far,
- $g$  is the global best position,
- $\omega, c_1, c_2$  are hyperparameters controlling exploration and exploitation.

- PSO is commonly used in hyperparameter tuning and reinforcement learning applications.
- Applications in Hyperparameter Tuning and Reinforcement Learning
- Gradient-free methods like Bayesian optimization and evolutionary strategies are particularly useful in hyperparameter tuning, where the search space is complex and non-differentiable. In reinforcement learning, policy optimization techniques such as Evolution Strategies (ES) leverage evolutionary algorithms to optimize policy networks without computing policy gradients.

### 3. Mathematical Formulation of Optimization Techniques

Mathematical optimization techniques provide the foundation for training machine learning models by minimizing an objective function, such as a loss function. This section presents the mathematical formulation of key optimization methods, including Gradient Descent and Adam Optimizer[5].

#### 3.1. Gradient Descent Equation

Gradient Descent is an iterative optimization algorithm used to minimize a loss function  $J(\theta)$  by updating the model parameters  $\theta$  in the direction of the negative gradient. The general update rule for batch gradient descent is:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla J(\theta^{(t)})$$

where:

- $\theta$  represents the model parameters (e.g., weights and biases),
- $t$  is the iteration step,
- $\eta$  (learning rate) controls the step size,
- $\nabla J(\theta)$  is the gradient of the loss function with respect to  $\theta$ .

##### 3.1.1. Variants of Gradient Descent

###### Stochastic Gradient Descent (SGD)

Instead of computing the gradient over the entire dataset, SGD updates parameters using a single data point or a mini-batch:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla J(\theta^{(t)}; x_i, y_i)$$

- where  $(x_i, y_i)$  is a randomly selected training example or mini-batch.
- Momentum-Based Gradient Descent
- Momentum helps accelerate convergence by incorporating past gradients:

$$v^{(t+1)} = \beta v^{(t)} + (1 - \beta) \nabla J(\theta^{(t)})$$

$$\theta^{(t+1)} = \theta^{(t)} - \eta v^{(t+1)}$$

where:

- $v$  represents the accumulated momentum,
- $\beta$  (e.g., 0.9) is the momentum coefficient.

Momentum-based optimization helps smooth updates and escape saddle points more effectively.

### 3.2. Adam Update Equations

Adam (Adaptive Moment Estimation) is an advanced optimization algorithm that combines the benefits of momentum and adaptive learning rates. It maintains two moving averages:

- First moment estimate (mean of past gradients):  $m_t$
- Second moment estimate (uncentered variance of past gradients):  $v_t$
- The update rules for Adam are:
- Step 1: Compute Biased First and Second Moment Estimates

$$m^{(t)} = \beta_1 m^{(t-1)} + (1 - \beta_1) \nabla J(\theta^{(t)})$$

$$v^{(t)} = \beta_2 v^{(t-1)} + (1 - \beta_2) (\nabla J(\theta^{(t)}))^2$$

where:

- $m_t$  is the first moment estimate (momentum term),
- $v_t$  is the second moment estimate (variance term),
- $\beta_1$  and  $\beta_2$  are exponential decay rates (commonly set as  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ).

### 3.3. Step 2: Bias Correction for Unbiased Estimates

Since  $m_t$  and  $v_t$  are initialized as zero, they are biased toward zero in early iterations. The bias-corrected versions are:

$$\hat{m}^{(t)} = \frac{m^{(t)}}{1 - \beta_1^t}$$

$$\hat{v}^{(t)} = \frac{v^{(t)}}{1 - \beta_2^t}$$

### 3.4. Step 3: Parameter Update Rule

$$\theta^{(t+1)} = \theta^{(t)} - \frac{\eta}{\sqrt{\hat{v}^{(t)} + \epsilon}} \hat{m}^{(t)}$$

where:

- $\epsilon$  (typically  $10^{-8}$ ) is a small constant to prevent division by zero,
- $\eta$  is the learning rate.

- Key Advantages of Adam
- Momentum: Uses first-moment estimates to accelerate convergence.
- Adaptive Learning Rates: Adjusts step size for each parameter individually.
- Robustness: Works well with sparse data and noisy gradients.

#### 4. Comparative Analysis of Optimization Methods

Table 1 presents a comparative analysis of different optimization techniques based on key factors such as convergence speed, stability, computational complexity, adaptability, and suitability for large datasets[6].

##### 4.1. Analysis and Key Takeaways

###### 4.1.1. Gradient-Based Methods (GD, SGD, Momentum, Adam, RMSprop)

- Well-suited for differentiable, convex, or quasi-convex optimization problems.
- SGD and Adam are preferred for deep learning due to their efficiency in handling large datasets.
- Momentum and RMSprop improve stability and convergence speed.
- Newton's Method is the fastest for convex problems but impractical for deep learning due to high computational cost.

###### 4.1.2. Gradient-Free Methods (GA, PSO)

- Used for optimization problems where gradients are unavailable or difficult to compute.
- Genetic Algorithms (GA) are suitable for discrete search spaces and hyperparameter tuning.
- Particle Swarm Optimization (PSO) works well in continuous search spaces but may struggle with high-dimensional problems.
- Both methods have high computational cost and require extensive tuning.

###### 4.1.3. Practical Recommendations

- Deep Learning: Adam and RMSprop are widely used due to adaptive learning rates.
- Convex Problems: Newton's Method (if computational cost allows) or standard Gradient Descent.
- Hyperparameter Optimization: Genetic Algorithms, PSO, or Bayesian Optimization.
- Non-Differentiable Functions: Evolutionary algorithms (GA, PSO) outperform gradient-based approaches.

**Table 1** Presents a comparative analysis of different optimization techniques

Optimization Method	Convergence Speed	Stability	Computational Cost	Adaptability	Suitability for Large Datasets	Advantages	Disadvantages
Gradient Descent (GD)	Slow	Moderate	Low	Low	Poor	Simple, guaranteed convergence for convex problems	Slow for large datasets, sensitive to learning rate
Stochastic Gradient Descent (SGD)	Faster	Less stable	Medium	High	Excellent	Faster convergence, good for large datasets	High variance in updates, noisy convergence
Momentum-Based GD	Faster	More stable	Medium	Medium	Good	Reduces oscillations, accelerates convergence	Needs tuning of momentum hyperparameter
Adam Optimizer	Fast	High	Medium	High	Excellent	Adaptive learning rate, good generalization	Can lead to suboptimal convergence in some cases

RMSprop	Fast	High	Medium	High	Excellent	Suitable for non-stationary problems	Requires careful hyperparameter tuning
Newton's Method	Very Fast	High	High	Low	Poor	Quadratic convergence, effective for convex problems	Computationally expensive, requires second derivatives
Genetic Algorithms (GA)	Variable	High	High	High	Moderate	Good for non-differentiable, complex functions	Computationally expensive, slow convergence
Particle Swarm Optimization (PSO)	Variable	Moderate	High	High	Moderate	No gradient required, good for complex landscapes	Slow for high-dimensional problems, risk of stagnation

## 5. Applications of Optimization in Machine Learning

Optimization techniques play a critical role in various machine learning domains, ensuring efficient training, better generalization, and improved computational efficiency. Below are key applications in deep learning, reinforcement learning, and hyperparameter tuning, along with specific optimization methods used in each domain[7].

### 5.1. Deep Learning: Optimizing Neural Network Weights

Deep learning models, particularly deep neural networks (DNNs), rely heavily on optimization algorithms to minimize the loss function during training. Optimization determines how efficiently a neural network learns patterns from data and converges to an optimal solution[8].

#### 5.1.1. Optimization Methods for Deep Learning

Adam Optimizer:

- Uses adaptive learning rates and momentum to improve convergence speed.
- Ideal for deep networks, as it adjusts the learning rate for each parameter dynamically.
- Commonly used in Convolutional Neural Networks (CNNs) and Transformers.

#### 5.1.2. Adam weight update equation:

$$\theta^{(t+1)} = \theta^{(t)} - \frac{\eta}{\sqrt{\hat{v}^{(t)} + \epsilon}} \hat{m}^{(t)}$$

#### 5.1.3. RMSprop (Root Mean Square Propagation):

- Controls learning rate by normalizing gradients, preventing drastic updates.
- Works well for recurrent neural networks (RNNs) and speech recognition tasks.

#### 5.1.4. RMSprop weight update:

$$v^{(t)} = \beta v^{(t-1)} + (1 - \beta)(\nabla J(\theta^{(t)}))^2$$

$$\theta^{(t+1)} = \theta^{(t)} - \frac{\eta}{\sqrt{v^{(t)} + \epsilon}} \nabla J(\theta^{(t)})$$



5.1.5. *Stochastic Gradient Descent (SGD):*

- Updates weights per mini-batch, making it computationally efficient.
- Often combined with momentum for accelerated convergence.
- Reinforcement Learning: Policy Optimization
- Reinforcement learning (RL) involves optimizing policy functions to maximize cumulative rewards in an environment. Unlike supervised learning, RL does not rely on labeled data but learns through exploration and experience.
- Optimization Methods for Reinforcement Learning

5.2. **Evolutionary Algorithms (EAs):**

- Used in policy search methods, where policies evolve over multiple generations.
- Examples: Genetic Algorithms (GA), Particle Swarm Optimization (PSO), and Covariance Matrix Adaptation Evolution Strategy (CMA-ES).
- Suitable for cases where gradient-based methods struggle, such as high-dimensional or sparse reward environments.

5.2.1. *Proximal Policy Optimization (PPO):*

- Balances exploration and exploitation by optimizing policy updates.
- Commonly used in deep RL applications such as robotics and game AI.

5.2.2. *PPO objective function:*

$$L(\theta) = \mathbb{E} [\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)]$$

where:

- $r_t(\theta)$  is the probability ratio of new vs. old policy,
- $A_t$  is the **advantage estimate**,
- $\epsilon$  controls how much the policy can change.

5.2.3. *Deep Q-Learning (DQN):*

- Uses gradient-based optimization to update Q-values (expected future rewards).
- Optimizes the Bellman Equation using mean squared error (MSE) loss:

$$L(\theta) = \mathbb{E} [(y_t - Q(s, a; \theta))^2]$$

where:

- $y_t = r + \gamma \max Q(s', a'; \theta^-)$  is the target value,
- $\gamma$  is the discount factor,
- $Q(s, a)$  estimates future rewards.

- Hyperparameter Tuning: Bayesian Optimization for Model Selection
- Hyperparameter tuning is essential for optimizing machine learning models, as poor choices in learning rate, batch size, or network architecture can lead to suboptimal results.
- Optimization Methods for Hyperparameter Tuning

5.3. **Bayesian Optimization:**

- Uses probabilistic models (Gaussian Processes) to model the objective function and select hyperparameters efficiently.

- Instead of exhaustive grid search, it evaluates promising configurations.

#### 5.4. Bayesian Optimization Update Rule:

$$x_{t+1} = \arg \max_x \alpha(x; D_t)$$

where:

- $\alpha(x; D_t)$  is an acquisition function (e.g., Expected Improvement (EI), Upper Confidence Bound (UCB)),
- $D_t$  is the set of previously evaluated hyperparameter configurations.

##### 5.4.1. Grid Search & Random Search:

- Grid Search: Evaluates all possible hyperparameter combinations (computationally expensive).
- Random Search: Randomly selects configurations and performs better than grid search in high-dimensional spaces.

---

## 6. Conclusion

Mathematical optimization is a cornerstone of machine learning, directly influencing model training efficiency, convergence speed, and computational complexity. Traditional gradient-based methods such as Gradient Descent (GD), Stochastic Gradient Descent (SGD), and Adam have been widely adopted due to their ability to efficiently minimize loss functions and improve generalization. While GD ensures a steady path toward convergence, SGD introduces stochasticity, improving speed but at the cost of higher variance. Adam, combining adaptive learning rates with momentum, has become a standard choice in deep learning due to its faster convergence and better stability across different tasks. However, these methods often struggle with non-convex optimization landscapes, where multiple local minima and saddle points hinder model performance. To address these challenges, gradient-free optimization techniques, such as Genetic Algorithms (GA) and Particle Swarm Optimization (PSO), have gained attention, particularly in scenarios like hyperparameter tuning and reinforcement learning. These techniques operate without requiring gradient information, making them suitable for complex, non-differentiable loss functions. Additionally, hybrid approaches that combine gradient-based optimization with evolutionary strategies have shown promise in accelerating convergence while maintaining robustness. The need for adaptive and self-tuning optimization algorithms is growing, particularly in large-scale machine learning applications, where computational efficiency and scalability are critical concerns. Looking ahead, future research should explore advanced optimization techniques to further enhance machine learning performance. Quantum-inspired optimization methods, self-adaptive learning rate strategies, and decentralized optimization for federated learning are promising areas of development. As AI continues to expand into edge computing and real-time applications, optimization methods must become more energy-efficient and scalable to work within resource-constrained environments. By integrating hybrid techniques, automated model selection, and privacy-preserving optimization, the next generation of optimization strategies will play a crucial role in shaping the future of high-performance, scalable, and adaptive machine learning systems.

---

## Compliance with ethical standards

*Disclosure of conflict of interest*

NO conflict of interest to be disclosed.

---

## References

- [1] Sun, Shiliang, Zehui Cao, Han Zhu, and Jing Zhao. "A survey of optimization methods from a machine learning perspective." *IEEE transactions on cybernetics* 50, no. 8 (2019): 3668-3681.
- [2] Bull, Adam D. "Convergence rates of efficient global optimization algorithms." *Journal of Machine Learning Research* 12, no. 10 (2011).

- [3] Liu, Bo, Hadi Aliakbarian, Zhongkun Ma, Guy AE Vandenbosch, Georges Gielen, and Peter Excell. "An efficient method for antenna design optimization based on evolutionary computation and machine learning techniques." *IEEE transactions on antennas and propagation* 62, no. 1 (2013): 7-18.
- [4] Bottou, Léon, Frank E. Curtis, and Jorge Nocedal. "Optimization methods for large-scale machine learning." *SIAM review* 60, no. 2 (2018): 223-311.
- [5] Khan, Waqar Ahmed, Sai Ho Chung, Muhammad Usman Awan, and Xin Wen. "Machine learning facilitated business intelligence (Part II) Neural networks optimization techniques and applications." *Industrial Management & Data Systems* 120, no. 1 (2020): 128-163.
- [6] Wang, Xiaofei, Yiwen Han, Victor CM Leung, Dusit Niyato, Xueqiang Yan, and Xu Chen. "Convergence of edge computing and deep learning: A comprehensive survey." *IEEE communications surveys & tutorials* 22, no. 2 (2020): 869-904.
- [7] Fazel, Maryam, Rong Ge, Sham Kakade, and Mehran Mesbahi. "Global convergence of policy gradient methods for the linear quadratic regulator." In *International conference on machine learning*, pp. 1467-1476. PMLR, 2018.
- [8] Huang, Chongwen, Ronghong Mo, and Chau Yuen. "Reconfigurable intelligent surface assisted multiuser MISO systems exploiting deep reinforcement learning." *IEEE Journal on Selected Areas in Communications* 38, no. 8 (2020): 1839-1850.