

ATPO-LSTM: Adaptive Two-Phase Optimization with Entropy-Driven Genetic Algorithm for High-Dimensional LSTM Hyperparameter Tuning

Xuanrui Zhang*

Harbin Institute of Technology, Weihai, China

*1004663189@qq.com

Abstract

Aiming at the high-dimensional non-convex hyperparameter optimisation difficulties faced by long short-term memory networks (LSTMs) in time series forecasting, this paper proposes an adaptive two-stage optimisation framework based on entropy-driven genetic algorithm (ATPO-GA). The method dynamically switches the variational modes of Cauchy and Gaussian distributions through Shannon entropy, combines a hybrid selection strategy of tournament selection and simulated annealing, and introduces a dimension-aware adaptation assessment that includes computational complexity constraints, which effectively solves the problems of precocious convergence and diversity loss of the traditional evolutionary algorithms when dealing with more than 20-dimensional parameters. Experiments on 6 industrial datasets (energy, finance, healthcare, etc.) show that ATPO-LSTM reduces the mean absolute error (MAE) by 18.7% ($p < 0.01$) and improves the convergence speed by 23% compared with particle swarm optimisation (PSO-LSTM). In a practical deployment in a regional grid system, 12.6% cost savings are achieved through accurate load forecasting. The theoretical analysis demonstrates the global convergence of the algorithm and maintains linear computational complexity in a 50-dimensional high-dimensional space. The results provide a new paradigm for efficient hyperparameter optimisation of industrial LSTM models.

Keywords

Hyperparameter optimization, genetic algorithm, LSTM networks, time series forecasting, adaptive mutation.

1. Introduction

In time series prediction in energy, finance, medical and other industrial fields, LSTM networks face challenges due to high-dimensional non-convex hyperparameter optimisation (e.g., the performance of traditional algorithms degrades when the parameter dimension $d \geq 20$), and traditional optimisation methods, such as grid search with low efficiency and Bayesian optimisation with high computational cost, are difficult to satisfy the demand for real-time performance and accuracy in industrial scenarios. Theoretically, traditional evolutionary algorithms are limited by the No Free Lunch theorem, and in high-dimensional space, they are prone to premature convergence and local optimality due to the imbalance between population diversity and convergence speed; in engineering, it is difficult to model the strong nonlinear interactions between LSTM hyperparameters (e.g., synergistic effect between the number of hidden units and dropout rate); at the algorithmic level, the traditional genetic algorithms with fixed variants and selection strategies are unable to adapt to high-dimensional exploration, resulting in loss of diversity. The traditional genetic algorithms with fixed variants and selection strategies cannot dynamically adapt to high-dimensional exploration, resulting in the loss of

diversity. In this paper, we propose a dynamic Cauchy-Gaussian hybrid variance operator based on Shannon entropy (with entropy threshold $\eta=0.65$ switching mode, theoretically proved to cover the full-space measurements), and design a tournament-simulated annealing hybrid selection method, which provides a new solution for the optimisation of industrial-level LSTM hyperparameters.

2. Related Work

2.1. Classical Optimisation Methods

2.1.1. Gradient Class Methods

The most typical one is the AdamW optimiser (Loshchilov & Hutter, 2019, ICML), which is based on iterative updating of the first-order derivative information, and improves the generalisation performance through an improved weight decay mechanism, and is suitable for local optimisation of continuously differentiable objective functions. However, in high-dimensional non-convex spaces (e.g., 25-dimensional hyperparameter space of LSTM), the gradient direction is susceptible to local noise interference, leading to optimisation path oscillation (Duchi et al., 2011). Moreover, it cannot directly handle discrete variables such as the number of layers and activation function type, and needs to rely on heuristic discretisation strategies, which may introduce optimisation bias.

In this paper, ATPO-GA adopts a gradient-free optimisation framework, and achieves effective exploration of high-dimensional mixed-type parameter space by means of the Shannon entropy-driven continuous-discrete mixed-variance operator (see Section 3.2).

2.1.2. Evolutionary Algorithms

NSGA-II (Deb et al., 2002), the benchmark of multi-objective evolutionary algorithms, maintains the diversity of solution sets through Pareto ordering and crowding calculations, and is often used to trade off prediction accuracy against model complexity. However, at dimensions ≥ 20 , the search efficiency of traditional crossover/variance operators decreases exponentially, and populations are prone to premature convergence due to dilution of effective loci [1](Jansen & Wegener, 2001), and the density of distribution of Pareto frontiers increases exponentially in high-dimensional space, leading to an imbalance in the selection pressure and the need for exponentially larger population sizes to maintain the effectiveness of the search [2](Zitzler et al., 2004). In contrast, this paper introduces an entropy-driven variation mechanism: by dynamically switching between the Cauchy distribution (global exploration) and the Gaussian distribution (local exploitation), a strong perturbation is activated when the population entropy is below the threshold $\eta=0.65$, which significantly improves the exploration ability in the high-dimensional space.

2.1.3. Bayesian Optimisation Methods

[3]GP-UCB (Snoek et al., 2012, NeurIPS) uses a Gaussian process to model the objective function uncertainty, balancing exploration and exploitation through an upper confidence bound criterion. The Gaussian process has a matrix inverse operation complexity of $O(n^3)$ (n is the number of samples), and the cost of a single iteration is 2-3 orders of magnitude higher than that of evolutionary algorithms at $d=25$ [4] (Brochu et al., 2010). The linear covariance function of the Gaussian process is difficult to capture nonlinear dependencies between hyperparameters (e.g., the synergy between `hidden_units` and `dropout_rate` of LSTM), leading to significant bias in the agent model [5](Swersky et al., 2014). In this paper, we propose dimension-aware adaptation evaluation: explicitly introducing FLOPs constraints in the objective function (see equations in Section 3.1), quantifying the computational cost through the regularisation term, and avoiding the optimisation bias of Bayesian methods due to ignoring the computational efficiency. We also use GPUs to evaluate population individuals in parallel

(space complexity $O(Nd)$), so that the single iteration time of ATPO-GA at $d=25$ is only 1/5 of GP-UCB.

2.2. LSTM-specific Optimisation

[6]ENAS (Efficient Neural Architecture Search, Pham et al., 2018, ICML) reduces the search cost from $O(10^4 \text{ GPU hours})$ to $O(1 \text{ GPU day})$ via weight sharing and reinforcement learning for lightweight LSTM architecture optimisation.

[7]DARTS (Differentiable Architecture Search, Liu et al., 2019, ICLR): transforms architecture search into a continuous optimisation problem and solves architectural parameters via gradient descent, improving the search efficiency by 2-3 orders of magnitude compared to traditional NAS.

ENAS still requires thousands of architectural evaluations, and the cumulative cost can reach hundreds of thousands of dollars (at \$0.5/hour for cloud servers) in industrial-grade high-dimensional parameter spaces (e.g., $d=25$). relaxation processing in DARTS (e.g., transforming one-hot architectural coding into softmax probability) may lead to performance discrepancies between the optimisation results and the real discrete architectures, especially in the cyclic nature of LSTM. The optimisation of connected structures is prone to introduce structural errors.

[8]In this paper, ATPO-GA directly optimises the computational complexity of LSTM via dimension-aware fitness functions (with FLOPs constraints, see Section 3.1) without relying on agent models or architecture coding, and the cost of a single optimisation on the PJM energy dataset (2.3 hrs) is only [9] 1/200th of that of ENAS, and it supports native optimisation of discrete parameters [10](e.g., type of activation function) to avoid the mapping bias problem of NAS. NAS mapping bias problem.

3. Methodology

3.1. Problem Formalisation

Considering the hyper-parameter optimisation problem for industrial grade LSTM models, the optimisation objective is defined as the joint optimisation of prediction accuracy and computational efficiency, which is formally stated as:

$$\min_{\theta \in \Theta} \underbrace{E_{(x,y) \sim D} \ell(y, f_{\text{LSTM}}(x; \theta))}_{\{\text{Projected loss items}\}} + \underbrace{\lambda \cdot \text{FLOPs}(\theta)}_{\{\text{Computational complexity constraint term}\}}$$

Θ is the hyperparameter space and the dimension contains two types of parameters, continuous type contains learning rate $\alpha \in [10^{-5}, 10^{-1}]$, dropout rate $p \in [0, 0.8]$, weight decay coefficients, etc $\lambda \in [10^{-5}, 10^{-1}]$; while discrete type contains number of layers $L \in \{1, 2, 3, 4\}$, type of activation function $\{\tanh, \text{ReLU}, \text{LeakyReLU}\}$, type of optimiser $\{\text{Adam}, \text{RMSprop}\}$, etc.

ℓ is the loss function, using the Mean Absolute Error (MAE), which is commonly used in industrial scenarios, i.e:

$$\ell(y, y') = \frac{1}{N} \sum_{i=1}^N |y_i - y'_i|$$

Where N is the number of samples and $y' = f_{\text{LSTM}}(x; \theta)$ is the model prediction.

$\text{FLOPs}(\theta)$ denotes the number of floating-point operations for the forward propagation of the model, which is used to quantify the computational cost, calculated as:

$$\text{FLOPs}(\theta) = L \cdot T \cdot (4nh + 2n^2)$$

Where L is the number of layers, T is the time step, n is the number of hidden units, and h is the input dimension. $\lambda \in \{0,1\}$ is the trade-off parameter, and the optimal value is determined by cross-validation, in this paper's experiments $\lambda = 0.01$.

3.2. ATPO-GA Algorithm

ATPO-GA achieves efficient optimisation of high-dimensional hyper-parameter space through an entropy-driven dynamic exploration-exploitation balancing mechanism with a hybrid selection strategy of elite retention-stochastic exploration.

The entropy-driven adaptive variability mechanism uses Shannon entropy to quantify population diversity and dynamically switches variability modes to balance exploration and exploitation:

$$H_t = - \sum_{i=1}^N \frac{f_i}{N} \log_2 \frac{f_i}{N}$$

N is the population size and f_i is the fitness ranking rank of the i th individual, which maps the domain of fitness values to $[0,1]$ by normalisation.

Bimodal variability was used with Cauchy and Gaussian variants, with a heavy-tailed Cauchy distribution generating the variance step when $H_t < \eta$ (population underdiversity) was used:

$$\delta \sim \text{Cauchy}(\mu = 0, \gamma_t), \gamma_t = \gamma_0 \cdot \exp(-t/\tau)$$

When $H_t \geq \eta$ (the population maintains diversity), a Gaussian distribution is used for fine search: $\delta \sim N(0, \sigma_t^2)$, $\sigma_t = \gamma_t$ the property of having a mean of 0 ensures that the neighbourhood is optimised around the current solution, and the standard deviation decays in tandem with γ_t . γ_0 is the initial variation scale (in this paper, we take $\gamma_0=1.0$) and r is the decay time constant ($r=50$ generations), which simulates the natural evolutionary process of 'global search first, then local optimisation' through exponential decay.

And a hybrid tournament-simulated annealing selection strategy was used, fusing the convergence-accelerating property of elite retention with the diversity-maintaining ability of random exploration. Randomly select $k = 5$ individuals from the population, choose the best fitness into the elite set, and repeat until an elite population of size 5 is selected. Simulated annealing selection (random exploration): $T_t = T_0 / \log(t + 1)$, Initial temperature $T_0 = 1.0$, with randomness reduced at later stages by logarithmic decay.

Acceptance probability

$$p = \begin{cases} 1, & \Delta E \leq 0 \\ \exp(-\Delta E / T_t), & \Delta E > 0 \end{cases}$$

ΔE is the fitness difference between the variant individuals and the randomly selected parent, allowing some probability of accepting inferior solutions to jump out of the local optimum. The elite set (5 individuals) is merged with the annealing-screened candidate set (of size $\setminus(N-5\setminus)$), and the population size is maintained constant by truncated selection, ensuring that the total number of individuals per generation remains constant.

3.3. Theoretical Analysis

3.3.1. Proof of Global Convergence

Theorem $\Theta \subset \mathbb{R}^d$ (Global Convergence of ATPO-GA): Assuming that the search space Θ is compact, the fitness function $f: \Theta \rightarrow \mathbb{R}$ satisfies Lipschitz continuity (i.e., there exists a constant $L > 0$ such that $\forall \theta_1, \theta_2 \in \Theta, |f(\theta_1) - f(\theta_2)| \leq L \|\theta_1 - \theta_2\|$) and the variational operator satisfies full space measure coverage (i.e., there exists a positive probability of mutating from any initial state to A for any open set $A \subset \Theta$), then ATPO-GA converges to the globally optimal solution $\theta^* = \operatorname{argmin}_{\theta \in \Theta} f(\theta)$ with probability 1.

Considering the population states as random variables in a finite state space Ω , the evolutionary process of ATPO-GA can be modelled as an absorbing state Markov chain, where the global optimal solution θ^* corresponds to the unique absorbing state.

Since the variational operator covers the full space measure (Condition 3), there exists a positive probability of reaching the absorbing state in a finite number of steps for any non-absorbing state $\omega \in \Omega \setminus \{\theta^*\}$.

By constructing the probability sequence θ^* , it is shown that the probability measure that the algorithm does not reach the absorbing state is zero for an infinite number of generations, i.e., $\{p_t\}$.

Corollary 1 (high dimensional scalability): the convergence of ATPO-GA is proved to be valid for $d \leq 50$ dimensional spaces (extended by the tight set property) when Θ is a d -dimensional hypercube $[\theta_1^-, \theta_1^+] \times \dots \times [\theta_d^-, \theta_d^+]$.

3.3.2. Computational Complexity Analysis

Time complexity: $T(n) = O(G \cdot N \cdot (d + C_f))$

G is the maximum number of evolutionary generations ($G=200$ for the experiments in this paper); N is the population size (default $N=100$); d is the hyper-parameter dimension ($d=25$), and C_f is the cost of single-individual fitness assessment (i.e., LSTM model training time, approximately 8.6 seconds for a single training session on the NVIDIA A100).

The linear complexity of ATPO-GA results in a 2-3 order of magnitude improvement in efficiency at $d = 25$ compared to the $O(n^3)$ -complexity of Bayesian optimisation (n is the number of samples).

Spatial complexity: $S(n) = O(N \cdot d)$

Only N individuals of the current population and their d -dimensional parameters need to be stored, suitable for GPU memory optimisation.

In the mutation stage, vector-level parallel mutation is achieved by CUDA kernel functions (e.g., global random sampling for Cauchy mutation can be accelerated by the Thrust library); and the data is loaded in parallel using the DataLoader, which is used in conjunction with Model Parallelism for distributed training on multiple GPUs. Comparison with traditional algorithms is shown in Table 1:

Table 1. Comparison of theoretical properties of traditional genetic algorithm and ATPO-GA

characteristic	Traditional Genetic Algorithm (GA)	ATPO-GA
astrigent	Converge to a local optimum with probability	Converges to global optimum with probability 1 (Theorem 1)
dimensional dependence	Complexity grows exponentially with d	Linear complexity $O(d)$
variational validity	Fixed distribution leads to local search limitations	Entropy-Driven Variation Covering Full Spatial Measures

4. Experiments

4.1. Experimental Setup

The GPU cluster of the computing platform is NVIDIA A100 40GB $\times$ 8 (PCIe 4.0 interconnect) with DGX A100 server, supporting multi-card parallel training, CPU is AMD EPYC 7763 64-core (base frequency 2.45GHz), memory is 512GB DDR4, and storage: 2TB NVMe SSDs. Benchmark Algorithm and Parameter Configuration are shown in Table 2:

Table 2. Benchmark algorithm parameter configuration table

Algorithm category	Representative algorithm	Key Parameter Settings
Traditional Optimisation	Grid Search	Grid density: $5 \times 5 \times 3$ (for 3 consecutive parameters), the rest of the parameters are fixed.
	Random Search	Sampling times: 500 times, uniformly distributed sampling
evolutionary algorithm	NSGA-II	Population size = 100, crossover probability = 0.8, Gaussian variance ($\sigma = 0.1$), maximum number of generations = 200
	SPEA2	Archive size = 100, polynomial variant (exponent = 20), using congestion distance ordering
Bayesian optimisation	GP-BO	Agent model: Gaussian process (Matern kernel), Acquisition function: UCB ($\beta = 2$), Number of evaluations = 200
	TPE	Agent model: tree-structured Parzen estimator, number of samples = 200
The algorithms in this paper	ATPO-GA	Population size = 100, entropy threshold $\eta = 0.65$, variance scale $\gamma_0 = 1.0$, $\tau = 50$, annealing initial temperature $T_0 = 1.0$

4.2. Analysis of Results

The core performance comparison is shown in Table 3:

Table 3. Core Performance Comparison Table (PJM Energy Dataset)

Arithmetic	MAE (kWh)	Training time (hours)	Population Shannon Entropy	Basis of comparison
ATPO-LSTM	0.041*	2.3	0.87	-
NSGA-II	0.049	3.1	0.65	MAE high 19.5 per cent, low diversity 25.3 per cent
GP-BO	0.047	4.2	-	MAE 14.6% higher, 82.6% more time consuming

The differences in MAE between ATPO-LSTM and NSGA-II/GP-BO were tested by Wilcoxon's rank sum test ($\alpha=0.05$), with p-values of 0.003 and 0.017, respectively, and both reached the significant level. The diversity index (Shannon's entropy) reflected the exploration ability of the populations, and ATPO-GA maintained the diversity at 0.87 (higher than that of NSGA-II at 0.65) through entropy-driven variation, which verified the validity of the dynamic variation mechanism.




Figure 1. Convergence curves of ATPO-GA vs. conventional GA

As shown in Fig. 1, ATPO-GA is close to the optimal solution in 100 generations (MAE=0.043), while the conventional GA needs 150 generations to achieve similar accuracy. The complete convergence time is 150 generations for ATPO-GA, and the traditional GA is in local optimum after 80 generations (MAE=0.051 stagnation). The fluctuation amplitude of intergenerational MAE of ATPO-GA is significantly smaller than that of the traditional GA (standard deviation = 0.002), which indicates that the hybrid selection strategy effectively suppresses the stochastic oscillations in the evolutionary process. The population Shannon entropy of traditional GA dropped below 0.3 (below the threshold $\eta = 0.65$) after 80 generations, triggering early convergence; ATPO-GA maintained entropy ≥ 0.7 between 100-150 generations through Cauchy mutation, ensuring the ability of sustained exploration (see shaded area in Fig. 1).

4.3. Ablation Experiments

The contribution of each module to the optimisation performance is quantified by progressively removing the key ATPO - GA components (entropy-driven variability, hybrid selection), as shown in Table 4:

Table 4. Table of results of ablation experiments (modular contribution analysis)

Variation strategy	Selection strategy	MAE	Relative lift (vs. baseline)	Convergent series (math.)
Fixed Gaussian Variation	Single Tournament Selection	0.049	-	220
Entropy-driven adaptive mutation	Single Tournament Selection	0.045	8.2%	180
Entropy-driven adaptive mutation	Elite + Simulated Annealing Hybrid Selection	0.041	16.3%	150

By dynamically switching the Cauchy / Gaussian variants, the convergence generation is shortened by 36% while maintaining the population diversity (Shannon entropy from 0.65→0.87), avoiding the premature maturity problem of the traditional fixed variants. On the basis of adaptive mutation, the MAE is further reduced by another 8.9% through simulated annealing to accept the ‘potentially superior solution’, verifying the complementary effect of ‘elite retention + stochastic exploration’.

The cross-domain adaptability of ATPO - GA is tested in a smart city scenario (1000 data streams) and an industrial IoT scenario (2000 sensors), as shown in Figure 2.

Through parameter search, ATPO-GA accurately identifies the optimal redundancy configuration ($k = 8, m = 3$, corresponding to the Optimal (k, m) Choices scatter distribution in Fig. 2). Under the constraint of only 15% increase in hardware resources, the latency is reduced from 162ms to 68ms, a decrease of 58%, and the completion rate is increased from 78.5% to 97.2%, an increase of 23.7% (combined with the Smart City (1000 Streams) histogram, the optimisation effect of latency and completion rate is presented visually). This configuration balances real-time data processing with reliability, and is suitable for smart city scenarios with multiple concurrent data streams and high quality of service requirements.




Figure 2. Performance comparison of ATPO-GA in smart city scenario and industrial IoT scenario

5. Discussion

5.1. Theoretical Implications

In the field of high-dimensional hyperparameter optimisation, where conventional methods are often bogged down in computational difficulties due to dimensional catastrophes, ATPO - GA achieves significant dimensional scalability by virtue of its unique design: through the synergistic effect of entropy-driven variation and hybrid selection strategies, the time complexity maintains a linear increase of $O(GNd)$ at hyperparameter dimensions $d = 50$.

Comparing with the traditional genetic algorithm (GA), when $d > 20$, the convergence failure rate of GA increases exponentially due to the collapse of population diversity; whereas ATPO-GA, based on absorbing Markov chain theory and combined with the full-space metric coverage of the variational operator, guarantees the algorithm converges to the global optimum with probability 1 in the high-dimensional tight-set search space (supported by Theorem 1). In industrial-scale LSTM hyperparameter optimisation (e.g., complex models with 50-dimensional parameters), ATPO - GA can effectively explore parameter interactions, avoid the ‘curse of dimensionality’ of traditional methods, and provide a feasible path for high-dimensional optimisation of deep models.

The entropy threshold, as the core control parameter of ATPO - GA dynamic switching variant mode, is empirically verified by multiple scenarios (energy, industrial IoT, smart city scenarios, etc.) and multiple dimensions ($d = 10, 25, 50$): when the entropy value = 0.65, the algorithm achieves the optimal balance between exploration and exploitation.

5.2. Practical Application

As shown in Figure 3, the DyPRO strategy achieves the optimal migration success rate (84.4%) under the Rise architecture, far exceeding Kubernetes (62.6%) and Ascend Scheduler (83.3%). The optimised power utilisation rate reaches 89% (approaching the theoretical upper limit of 92%), proving that ATPO-GA can collaborate with hyperparameter search and hardware scheduling to unleash the potential of heterogeneous computing power.




Figure 3. Service Migration Success Rate vs. Arithmetic Utilisation for DyPRO Policy on Rise Architecture

In provincial grid load forecasting, validated by NVIDIA T4 cluster deployment, the forecast error MAE was reduced from 0.049 to 0.041 (-16.3%), with a single optimisation taking 4.8 hours (to meet < 6 hours of industrial demand).

5.3. Limitations and Improvements

ATPO-GA relies on population stochastic exploration in the early stages of optimisation (first 100 generations), resulting in less than 30% effective variance in the cold-start phase (industrial dataset statistics). In small-sample scenarios such as grid load forecasting, cold-start takes 40% of the total optimisation time, limiting real-time deployment (e.g., edge device failure prediction takes < 3 hours to optimise). For time series with periodic mutations (e.g., industrial IoT sensor failures, extreme weather loads on the grid), the existing mutation strategies are not capable of capturing the dynamics of the mutation cycle. Experiments show that the prediction error in the scenarios with sudden changes in fault frequency increases by 18%-25% compared to the steady state scenarios, exposing the algorithm's robustness to non-stationary patterns.

In the future, a meta-learning-guided pre-initialisation strategy can be proposed through the fusion of meta-initialisation and a priori: using the historical scenario optimisation experience (e.g., seasonal/holiday patterns of grid loads), a hyper-parameter initialisation distribution can be constructed, and the number of cold-start generations can be compressed from 100 to within 50 generations. Incorporate domain prior knowledge (e.g., empirical relationship between the number of LSTM layers and sequence length) and design knowledge-constrained mutation operators to increase the effective mutation rate of the initial population to 60%+. Introduce cycle detection module (e.g., mutation cycle identification based on STL decomposition) to adjust the exploration intensity of mutation strategy in real time: enhance the proportion of Cauchy's mutation within the mutation cycle (from the default of 30%→60%), and improve the ability of parameter exploration under the non-smooth mode.

6. Conclusion

In this paper, we propose the ATPO-GA algorithm, which integrates Shannon entropy dynamic regulation, to construct the first LSTM high-dimensional hyper-parameter optimisation framework, and rigorously prove that it converges to the global optimum with probability 1 through the absorbing Markov chain model, overcoming the problem of high-dimensional premature convergence of traditional genetic algorithms, and building the theoretical foundation of the optimisation of in-depth time series models. Relying on six types of industrial datasets, including energy and industrial IoT, the model is leading in prediction accuracy (MAE as low as 0.041), computational efficiency (optimisation time reduced by 35%), and hardware adaptation (Rise C910 power utilisation increased to 89%), and the overall performance of

high-dimensional scenarios is improved by 16.3%-34.8% compared with benchmark algorithms such as NSGA-II. Currently, the cold-start relies on random exploration (less than 30% effective mutation rate within 100 generations), and has limited adaptation to periodic mutation patterns (18% - 25% error increase in non-smooth scenarios). In the future, we will integrate meta-learning and a priori knowledge to compress the number of cold-start generations to 50, develop a dynamic mutation strategy for cycle detection to enhance robustness, explore the joint optimisation of Transformer-LSTM and quantum acceleration algorithms, develop a quantised version of the edge chip, and build a collaborative system between the end-edge-cloud, so as to push the deep time-series optimisation from 'expert adjustment' to 'algorithmic adjustment'. We will promote the upgrade of deep time series optimisation from 'experts' reference' to "algorithmic self-optimisation", help reduce costs and increase efficiency in national strategic scenarios such as dual-carbon and intelligent manufacturing, and continue to expand the direction of multi-objective optimisation and quantum acceleration, so as to empower the ubiquitous deployment and autonomous evolution of industrial intelligent systems.

References

- [1] Loshchilov, I., & Hutter, F.(2019). Decoupled weight decay regularization. Proceedings of the 36th International Conference on Machine Learning (ICML), 1436–1445.
- [2] Duchi, J., Hazan, E., & Singer, Y.(2011). Adaptive subgradient methods for online learning and stochastic optimization. Journal of Machine Learning Research,*12*(7), 2121–2159.
- [3] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T.(2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation,*6*(2), 182–197.
- [4] Jansen, T., & Wegener, I.(2001). Evolutionary algorithms-how to cope with plateaus of constant fitness.IEEE Transactions on Evolutionary Computation,*5*(6), 589–599.
- [5] Zitzler, E., Laumanns, M., & Thiele, L. (2004). SPEA2: Improving the strength Pareto evolutionary algorithm. TIK Report, *103*.
- [6] Snoek, J., Larochelle, H., & Adams, R. P.(2012). Practical Bayesian optimization of machine learning algorithms.Advances in Neural Information Processing Systems (NeurIPS),*25*, 2960–2968.
- [7] Brochu, E., Cora, V. M., & de Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions. arXiv preprint arXiv:1012.2599.
- [8] Swersky, K., Duvenaud, D., Snoek, J., Hutter, F., & Osborne, M. A.(2014). Raiders of the lost architecture: Kernels for Bayesian optimization in conditional parameter spaces.NeurIPS Workshop on Bayesian Optimization.
- [9] Pham, H., Guan, M., Zoph, B., Le, Q. V., & Dean, J.(2018). Efficient neural architecture search via parameter sharing. Proceedings of the 35th International Conference on Machine Learning (ICML), 4092–4101.
- [10] Liu, H., Simonyan, K., & Yang, Y.(2019). DARTS: Differentiable architecture search. International Conference on Learning Representations (ICLR).