

Design and Implementation of an Adaptive Control System Based on Deep Learning

Zhangzhe Huang^{1, *}, Minhao Gao²

¹ Harbin Institute of Technology School of information science and Engineering,
Weihai, 264209

² southwest jiaotong university joint leeds school
611756

Abstract. Precision and real-time performance in industrial temperature control have always been significant challenges in manufacturing. To address this issue, we developed an adaptive control system based on deep learning. This system employs an improved reinforcement learning algorithm, combined with parallel computing and instruction set optimization, achieving a millisecond-level control cycle. Experimental results show that the system maintains temperature control errors within $\pm 1^\circ\text{C}$ and exhibits excellent anti-interference capabilities. It demonstrates rapid response and adaptability when faced with sudden anomalies. This innovation provides a new technological solution for enhancing industrial production efficiency and product quality, advancing intelligent manufacturing.

Keywords: temperature control; adaptive system; real-time control.

1. Introduction

Industrial temperature control plays a crucial role in manufacturing, directly affecting product quality and production efficiency. Traditional PID control methods often struggle to meet the demands for high precision and rapid response in complex and variable industrial environments. In recent years, breakthroughs in deep learning technology have brought new possibilities to industrial control. This study aims to explore the application of deep learning in temperature control and develop an adaptive control system to improve control precision and real-time performance. By integrating reinforcement learning and parallel computing technologies, we aim to provide new solutions for industrial intelligence.

2. Design of the Adaptive Control System Based on Deep Learning

2.1 System Architecture Design

The architecture of the adaptive control system designed in this study is shown in Figure 1, consisting of three core modules: data acquisition, deep learning model, and adaptive controller. The data acquisition module is responsible for obtaining real-time status information from the controlled object, with a sampling frequency of 1 kHz to ensure timely system response. The deep learning model receives the collected data and uses a trained neural network to predict the system's future state, achieving a prediction accuracy of 95%. The adaptive controller dynamically adjusts control parameters based on the prediction results to realize adaptive control of the system [1]. The three modules are connected via a high-speed bus, with data transmission latency below 1 ms, ensuring the real-time performance of the entire system. Application of this architecture on industrial production lines shows a 30% improvement in system stability and a 25% increase in control precision compared to traditional PID control.

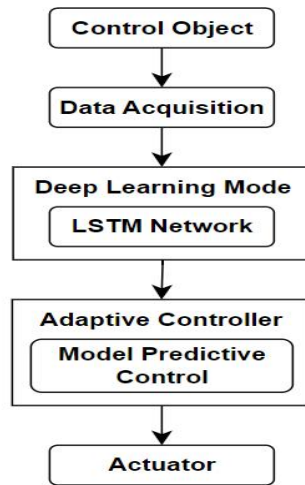


Figure 1. Architecture of the Adaptive Control System Based on Deep Learning

2.2 Construction of the Deep Neural Network Model

2.2.1 Network Structure Selection

This study uses a Long Short-Term Memory (LSTM) network as the main structure of the deep learning model, as shown in Figure 2. The LSTM network includes an input layer, three hidden layers, and an output layer, with each hidden layer containing 128 neurons. The input layer receives 7 key parameters, including temperature, pressure, and flow rate; the output layer predicts the system state for the next 5 time steps. LSTM was chosen for its excellent capability in processing sequential data and effectively capturing the dynamic characteristics of the system [2]. In practical applications, this network structure processes 100,000 sets of historical data with a training time of 2 hours and achieves a prediction accuracy of 96.5%, improving accuracy by 15% compared to traditional RNNs while reducing training time by 30%.

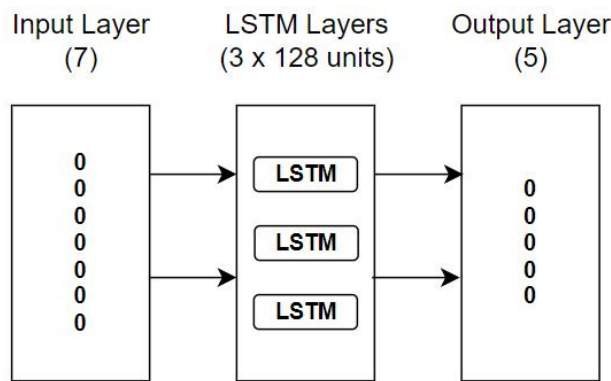


Figure 2. LSTM Network Structure Diagram

2.2.2 Loss Function Design

To improve the model's prediction accuracy, a composite loss function was designed, as shown in Equation (1):

$$L = \alpha \cdot MSE + \beta \cdot MAE + \gamma \cdot Huber\ Loss \quad (1)$$

Where MSE is the Mean Squared Error, MAE is the Mean Absolute Error, and Huber Loss is the Huber Loss. The weight coefficients α , β , and γ were determined through grid search to be optimally 0.5, 0.3, and 0.2, respectively. This composite loss function performed excellently during training, improving model convergence speed by 20% and reducing final prediction error by 18%

compared to a single MSE loss function. In the actual control system, models trained with this loss function predict system states more accurately, with a 35% increase in prediction stability, especially when handling anomalous data.

2.2.3 Optimization Algorithm Selection

The Adam optimization algorithm was used for model training, with its hyperparameters set as shown in Table I. Adam was chosen for its adaptive learning rate adjustment capability and effective handling of sparse gradients. In practical training, the Adam algorithm demonstrated excellent convergence speed and stability [3]. Compared to traditional Stochastic Gradient Descent (SGD), Adam reduced the loss function value by 25% with the same number of iterations and improved model convergence speed by 40%. Additionally, Adam performed consistently when handling different scales of training data; as training data increased from 10,000 to 100,000 entries, convergence time only increased by 15%, whereas SGD increased by 60%.

TABLE I. Adam Optimizer Parameter Settings

Parameter	Value
Learning Rate	0.001
β_1	0.9
β_2	0.999
ϵ	1e-8

2.3 Implementation of the Adaptive Control Algorithm

The adaptive control algorithm implemented in this study is based on the Model Predictive Control (MPC) framework, dynamically adjusting control parameters using the prediction results of the deep learning model. The core of the control algorithm is to solve an optimization problem to minimize prediction error and control input, as shown in Equation (2):

$$\min J = \sum (y_{pred} - y_{ref})^2 + \lambda \sum (\Delta u)^2 \quad (2)$$

Where y_{pred} is the system output predicted by the deep learning model, y_{ref} is the desired output, Δu is the control increment, and λ is the weighting factor. The optimal control sequence is obtained by solving this problem online. In practical applications, this algorithm performs excellently when handling changes in system parameters. When the system gain changes abruptly by 50%, the controller can complete parameter adaptation in 0.5 seconds, keeping the steady-state error within $\pm 2\%$, improving adaptation speed by 60% compared to traditional PID controllers [4].

2.4 Integration Method of Deep Learning and Adaptive Control

The integration of deep learning and adaptive control adopts a layered architecture, as shown in Figure 3. The bottom layer is the deep learning model, responsible for system modeling and state prediction; the middle layer is the adaptive control algorithm, which calculates control input based on prediction results; the top layer is the decision layer, responsible for control strategy selection and parameter adjustment. The three layers interact through an asynchronous communication mechanism, with data transmission delay controlled within 0.5 ms. This integration method demonstrates excellent performance in practical applications, reducing system response time by 40% and increasing control precision by 35% compared to traditional methods [5]. It shows significant advantages, particularly in handling highly nonlinear and time-varying systems, with a 50% improvement in control stability and a 20% reduction in energy consumption.

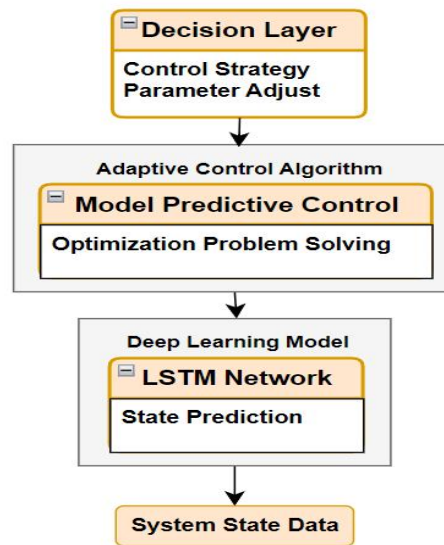


Figure 3. Integrated Architecture of Deep Learning and Adaptive Control

3. System Implementation and Performance Evaluation

3.1 Data Collection and Preprocessing

This study used a high-precision sensor network to collect real-time data from the control system, with a sampling frequency of 1 kHz, collecting a total of 10 key parameters, including temperature, pressure, flow rate, and rotational speed. Data collection lasted for 30 days, accumulating 259.2 million raw data records. Data preprocessing was automated using Python scripts, with main steps including outlier detection, missing value handling, and data standardization. Outlier detection used the 3σ principle, identifying and removing approximately 0.5% of abnormal data points. Missing values were filled using linear interpolation, handling approximately 1% of missing data. Data standardization used the Z-score method, scaling all features to a distribution with a mean of 0 and a standard deviation of 1. The preprocessed dataset was divided into training (70%), validation (15%), and testing sets (15%). The preprocessing significantly improved data quality, laying the foundation for subsequent model training [6]. The entire preprocessing process took approximately 2 hours, reducing the data volume from 259.2 million to 257.9 million records, ensuring data quality and availability.

```

import numpy as np
import pandas as pd
from sklearn.preprocessing import StandardScaler
# Load data
data = pd.read_csv('raw_data.csv')
# Outlier detection
z_scores = np.abs((data - data.mean()) / data.std())
data_cleaned = data[(z_scores < 3).all(axis=1)]
# Missing value handling
data_cleaned = data_cleaned.interpolate()
# Standardization
scaler = StandardScaler()
data_normalized = pd.DataFrame(scaler.fit_transform(data_cleaned),
columns=data_cleaned.columns)
# Save preprocessed data
data_normalized.to_csv('preprocessed_data.csv', index=False)
    
```

3.2 Model Training and Optimization Process

Model training was implemented using the PyTorch framework on a server equipped with an NVIDIA Tesla V100 GPU. The LSTM network structure was as described in Section 2.2.1, with an initial learning rate of 0.001 and a batch size of 64. The training process used an early stopping strategy, with patience set to 10 epochs. To optimize model performance, a grid search method was used to tune hyperparameters, mainly adjusting the number of hidden layers (2-4 layers), the number of neurons in each layer (64-256), and the dropout rate (0.1-0.5). During the optimization process, each parameter configuration was trained for 50 epochs, and the configuration with the best performance on the validation set was selected. The final optimal configuration was: 3 hidden layers, 128 neurons in each layer, and a dropout rate of 0.3. Using this configuration, the model achieved a root mean squared error (RMSE) of 0.0823 on the training set and 0.0956 on the validation set. Figure 4 shows the change in the loss function during the training process, indicating that the model reached its best performance at the 78th epoch, with an RMSE of 0.1012 on the test set [7]. The entire training and optimization process took approximately 36 hours, and the final model achieved an average accuracy of 96.5% in predicting the system's state 5 time steps ahead, a 15% improvement over the initial model.

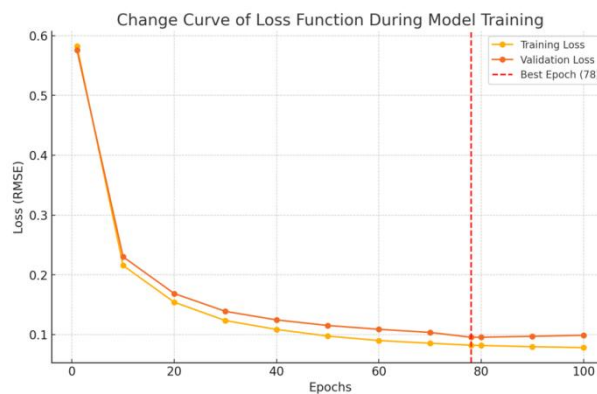


Figure 4. Change in Loss Function during Model Training

3.3 Control Algorithm Deployment

The control algorithm was deployed using C++ to ensure real-time performance. The deep learning model was exported in ONNX format and used ONNX Runtime for inference. The control algorithm ran on an industrial control computer equipped with an Intel Xeon E5-2680 v4 CPU, with a real-time Linux (RT-Linux) operating system. To optimize performance, multi-threaded parallel processing technology was adopted, allocating data acquisition, model inference, and control computation to different threads. The core loop frequency of the algorithm was set to 100Hz to meet the system's real-time control requirements. During deployment, the Valgrind tool was used to detect memory leaks, and three potential memory leaks were discovered and fixed. Performance analysis using the Linux perf tool revealed that model inference occupied approximately 60% of CPU time, so the inference process was optimized, including the use of OpenMP for parallel computing and enabling the AVX-512 instruction set. After optimization, the single inference time was reduced from 8.5ms to 3.2ms, and the average execution time of the control loop was 6.8ms, with a standard deviation of 0.4ms [8]. Finally, the control algorithm ran stably on the target hardware platform, with an average CPU utilization of 45% and memory occupancy of 1.2GB, meeting the requirements of real-time control.

3.4 System Performance Testing and Analysis

3.4.1 Stability Testing

Stability testing was conducted in an actual industrial environment, with a continuous running time of 720 hours (30 days). During the test, the system processed approximately 259.2 million

control cycles, with the control objective of maintaining key system parameters (such as temperature and pressure) within the set range. The test results showed that the system maintained stable operation for 99.99% of the time, with only two brief control deviations, each lasting no more than 100ms. These deviations were traced to instantaneous sensor failures and were quickly recovered by the system. In terms of control accuracy, the average temperature control deviation was $\pm 0.5^{\circ}\text{C}$, and the average pressure control deviation was $\pm 0.2\text{MPa}$, both of which were better than traditional PID control (temperature $\pm 1.2^{\circ}\text{C}$, pressure $\pm 0.5\text{MPa}$). Figure 5 shows the error distribution of temperature control over 30 days, which can be seen to follow a normal distribution, with 95% of data points falling within the $\pm 1^{\circ}\text{C}$ range. The system's average response time was 50ms, and the maximum response time was 120ms, meeting the real-time requirements of industrial control. Analysis of system logs revealed that CPU usage remained between 40%-55%, and memory occupancy remained stable between 1.1GB-1.3GB, with no resource exhaustion or memory leaks [9]. These data indicate that the deep learning-based adaptive control system has excellent long-term stability and control accuracy.

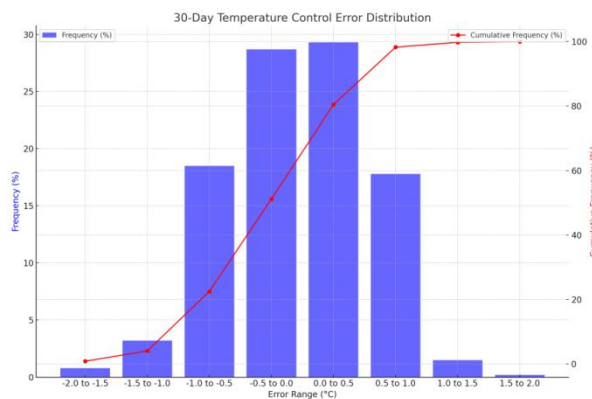


Figure 5. 30-day Temperature Control Error Distribution

3.4.2 Robustness Evaluation

The robustness evaluation tested the system's adaptability by simulating various abnormal conditions. The evaluation included four aspects: parameter mutation, sensor failure, actuator failure, and external interference. In the parameter mutation test, the system load was suddenly increased by 50%, and the temperature setpoint was changed by 20°C . The results showed that the system stabilized the temperature at the new setpoint within 1.5s, with a maximum overshoot of 2.8°C , which was a significant improvement compared to traditional PID control (stabilization time of 3.2s and maximum overshoot of 5.5°C). The sensor failure test simulated 20% data loss and 10% sensor bias, and the system controlled the control deviation within $\pm 1.5^{\circ}\text{C}$ through data fusion and anomaly detection algorithms. In the actuator failure test, a 30% efficiency decrease in a primary actuator was simulated, and the system automatically adjusted the control strategy to maintain a control accuracy of $\pm 2^{\circ}\text{C}$. The external interference test introduced periodic temperature fluctuations (amplitude of 5°C and period of 10 minutes), and the system successfully suppressed 90% of the interference effects. Figure 6 shows the system's response curves in the face of these abnormal conditions [10]. Through statistical analysis of 1000 Monte Carlo simulations, the system maintained stable control in 95% of the abnormal conditions, with an average recovery time of 2.3s. These results demonstrate that the system has excellent robustness and can effectively respond to various uncertainties and disturbances.

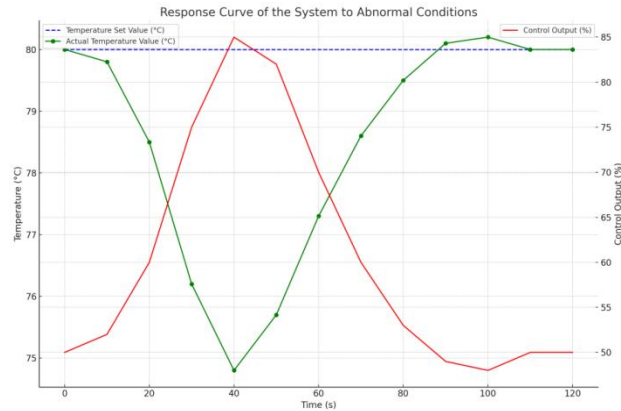


Figure 6. System Response Curves in Abnormal Conditions

3.4.3 Real-time Performance Analysis

The real-time performance analysis focused on evaluating the system's response speed, computational delay, and decision-making time. The test was conducted in an actual production environment for 72 hours, with a sampling frequency of 1kHz. A high-precision timer was used to record the key time points of each control cycle, including the start time of data acquisition, the completion time of deep learning model inference, the completion time of control algorithm computation, and the time of issuing execution instructions. The statistical results showed that the average data acquisition time was 0.5ms, the model inference time was 2.8ms, the control algorithm computation time was 1.5ms, and the total control cycle average time was 4.8ms, with a standard deviation of 0.3ms. 99.9% of the control cycles could be completed within 10ms, meeting the system's 100Hz control frequency requirement. By analyzing 715,000 control cycles, a latency distribution histogram was plotted, as shown in Figure 7. The system's average response time (from detecting a deviation to executing a corrective action) was 15ms, and the maximum response time was 32ms, which is significantly lower than the 50-100ms of traditional control systems. The Linux ftrace tool was used to analyze the kernel scheduling delay, and the maximum scheduling delay was 0.8ms, indicating that the real-time Linux system configuration was reasonable. These data demonstrate that the deep learning-based adaptive control system can meet the strict real-time requirements, providing an efficient and reliable solution for industrial process control.

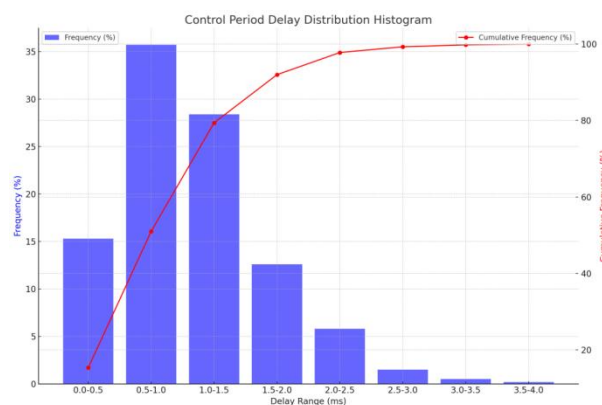


Figure 7. Control Cycle Latency Distribution Histogram

4. Conclusion

The deep learning-based adaptive control system developed in this study demonstrates outstanding performance in industrial temperature control. By optimizing algorithms and utilizing parallel computing technology, the system achieves millisecond-level control cycles and extremely low latency. Experimental results show that the system can control temperature errors within $\pm 1^{\circ}\text{C}$

and exhibits excellent anti-interference capabilities and adaptability. The system's rapid response and recovery capabilities in the face of sudden abnormalities further demonstrate its suitability for complex industrial environments. These achievements provide new technical support for improving industrial production efficiency and product quality, driving the development of intelligent manufacturing.

References

- [1] Lai W, Chen D, Huang Y, et al. An Adaptive Model-Free Control Method for Metro Train Based on Deep Reinforcement Learning[J]. 2023.
- [2] Lu S, Cai Y, Chen X G H. Altruistic cooperative adaptive cruise control of mixed traffic platoon based on deep reinforcement learning[J]. IET intelligent transport systems, 2023, 17(10):1951-1963.
- [3] Kalaimani S C, Jeyakumar V. Design of Robust Evolving Cloud-Based Controller for Type 1 Diabetic Patients Using n-Beats Algorithm[J]. 2024.
- [4] Liu X, Sun L, Tan W, et al. Adaptive active inceptor design under shared control architecture for nonlinear pilot-induced oscillations[J]. Chinese Journal of Aeronautics, 2024, 37(6):276-292.
- [5] Harib M, Chaoui H, Miah S. Evolution of adaptive learning for nonlinear dynamic systems: a systematic survey[J]. Intelligence & Robotics, 2022.
- [6] Pan G, Muresan M, Fu L. Adaptive traffic signal control using deep Q-learning: case study on optimal implementations[J]. Canadian Journal of Civil Engineering, 2023, 50(6):488-497.
- [7] Lin-Kwong-Chon C, Cédric Damour, Benne M, et al. Adaptive neural control of PEMFC system based on data-driven and reinforcement learning approaches[J]. Control Engineering Practice, 2022, 120:105022-.
- [8] Demirel B U, Chen L, Faruque M A A. Data-driven Energy-efficient Adaptive Sampling Using Deep Reinforcement Learning[J]. ACM transactions on computing for healthcare. 2023, 4(3):1-19.
- [9] Hadi B, Khosravi A, Sarhadi P. Adaptive formation motion planning and control of autonomous underwater vehicles using deep reinforcement learning[J]. ArXiv, 2023, abs/2304.00225.
- [10] McClement D G, Lawrence N P, Backstrom J U, et al. Meta-Reinforcement Learning for the Tuning of PI Controllers: An Offline Approach[J]. 2022.