Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 3, No. 7, 2024

Semantic Context Modeling for Fine-Grained Access Control Using Large Language Models

Yuting Peng
New York University, New York, USA
yp2212@nyu.edu

Abstract: This paper proposes a context-aware fine-grained access control framework based on a large language model. It addresses the limitations of existing access control methods in understanding contextual information under complex and dynamic environments. The method is built on the LLaMA3 model. It encodes user, resource, and context information in natural language format, enabling deep semantic modeling of access requests and dynamic policy generation. To enhance the model's understanding of domain knowledge, an external knowledge graph embedding mechanism is introduced. This integrates structured security knowledge with contextual semantics, improving classification accuracy in long-tail access scenarios. In addition, the system includes an access explanation module. It generates natural language justifications for access decisions, enhancing both interpretability and auditability of the model. Extensive experiments on a real-world platform dataset demonstrate that the proposed method significantly outperforms traditional approaches and baseline language models across multiple evaluation metrics. It shows superior robustness and accuracy, especially in complex contexts and cross-role access scenarios.

Keywords: Access control, large language models, context awareness, knowledge graphs.

1. Introduction

With the rapid development of information technology, data security and privacy protection have become major societal concerns [1]. Emerging technologies such as cloud computing, the Internet of Things, and artificial intelligence have led to high-frequency data circulation across domains and platforms [2]. As a result, traditional static and coarse-grained access control mechanisms are increasingly inadequate for meeting the diverse and dynamic security demands of modern systems. Ensuring data usability and flexible access while protecting sensitive information has become a key direction in information security research. Fine-grained access control, which allows precise and flexible authorization of data resources, has emerged as a research focus in both academia and industry. However, existing approaches often rely on static attribute matching and rule definition, lacking in-depth modeling of dynamic contextual factors. This limits their ability to meet the complex and multidimensional requirements of real-world access control scenarios [3].

Against this backdrop, Context-Aware Access Control (CAAC) models have been proposed and gradually adopted in practice. These models introduce dynamic context factors such as user behavior, environmental conditions, and device information to analyze and evaluate access requests more accurately. They aim to provide more flexible and secure data protection strategies [4]. However, due to the unstructured and diverse nature of contextual data, as well as semantic inconsistencies caused by dynamic changes, effectively understanding and reasoning about complex context information remains a core challenge in CAAC research. In addition, current methods often depend on manually defined rules or shallow semantic analysis for context modeling and decision logic. These limitations hinder their scalability and intelligence in complex environments.

In recent years, Large Language Models (LLMs) have demonstrated unprecedented capabilities in natural language processing, knowledge reasoning, and semantic understanding. Trained on massive text corpora, LLMs are capable of interpreting contextual semantics, capturing relational patterns, and generating natural language responses. These strengths offer a new perspective for solving the challenges of context awareness and semantic reasoning in access control. By deeply modeling context and abstracting semantics, LLMs can automatically understand access scenarios without relying on predefined rules. This enables the dynamic generation of more adaptive control strategies. In fine-grained access control, LLMs can provide intelligent and interpretable decisions based on user behavior patterns, resource types, and usage intentions across multiple dimensions.

Therefore, integrating LLMs into context-aware fine-grained access control is not only a natural progression of technological advancement but also a key strategy to enhance system security and flexibility. On one hand, this integration can overcome limitations in traditional systems related to context modeling, rule configuration, and policy generation, enabling dynamic response and precise control in complex scenarios. On the other hand, the strong generalization and transfer capabilities of LLMs allow for rapid adaptation across various security domains. This reduces system deployment and maintenance costs while improving universality and practicality. Furthermore, this research can be extended to emerging areas such as multi-tenant cloud environments, edge computing platforms, and zero-trust architectures, offering

theoretical and technical support for building smarter and more reliable data protection mechanisms.

In conclusion, the study of a context-aware fine-grained access control framework based on LLMs holds significant practical relevance and cutting-edge value. It offers both theoretical depth and application breadth. This research will effectively address existing gaps in complex semantic modeling and dynamic policy reasoning. It provides an innovative technological path for data security governance in the modern era. Moreover, the outcomes will lay a solid foundation for future exploration in intelligent security decision-making, semantic risk assessment, and human-machine collaborative control.

2. Related work

As a core technology in the field of information security, access control has long received widespread attention. Traditional access control models such as Discretionary Access Control (DAC), Mandatory Access Control (MAC), and Role-Based Access Control (RBAC) have been widely used due to their clear rules and ease of implementation. However, these models mostly rely on static attributes and predefined rules. They lack the ability to perceive dynamic environmental changes and semantic information. As a result, they are difficult to apply in modern distributed systems where user roles are ambiguous and environments are constantly changing. To improve the flexibility and adaptability of control strategies, researchers proposed Attribute-Based Access Control (ABAC). ABAC introduces user, resource, and environment attributes to achieve more fine-grained access control. However, in practice, ABAC still faces challenges such as complex attribute configuration, weak context interpretation, and high policy update costs. These issues are particularly significant in multi-source and heterogeneous environments [5].

To address these challenges, Context-Aware Access Control (CAAC) has been introduced. CAAC emphasizes the inclusion of dynamic context factors such as time, location, device, and user behavior in access decisions. This helps improve the accuracy and timeliness of policies. Existing research mainly focuses on context modeling, policy description language design, and reasoning mechanism construction. For example, some studies use ontology to build semantic context models, aiming to enhance semantic consistency and cross-domain sharing. Other studies apply machine learning methods to train models on historical behavior data to strengthen context awareness. However, due to the limitations of traditional modeling approaches in semantic abstraction and logical reasoning, these methods still lack generalization ability and interpretability in complex, dynamic environments. Therefore, there is an urgent need for a more intelligent approach with stronger semantic understanding to enhance context modeling and policy decision-making capabilities.

In recent years, Large Language Models (LLMs) have made groundbreaking progress in natural language processing and are gradually being applied to security decision-making and access control. Some studies have explored the use of pretrained language models for tasks such as automatic policy generation, user intent parsing, and context semantic extraction. These works have confirmed the advantages of LLMs in handling unstructured semantic data. For example, GPT-based models have been used to generate policy templates and assist in policy recommendation. BERT-based models have been applied to understand natural language requests and infer user access intentions. These methods show the potential of LLMs to understand complex contexts and model relationships between contextual elements. However, most studies remain at the level of policy representation and have not formed a systematic CAAC framework. In addition, current research often overlooks the problem of constructing fine-grained policies. There is still a lack of systematic exploration that integrates LLM capabilities with fine-grained and dynamic decision-making mechanisms. Therefore, building a fine-grained access control framework that incorporates LLM capabilities and supports context understanding and policy reasoning can not only address existing gaps in modeling granularity and reasoning power, but also provide a feasible path for the intelligent evolution of security systems.

3. Method

This study proposes a context-aware fine-grained access control method based on a large language model, aiming to realize the intelligent and semantic decision-making process for user access requests in a dynamic environment. The core idea is to uniformly encode the three key elements of access control: user requests, resource attributes, and context information as natural language input, and perform deep semantic modeling and policy reasoning through a large language model. The model architecture is shown in Figure 1.

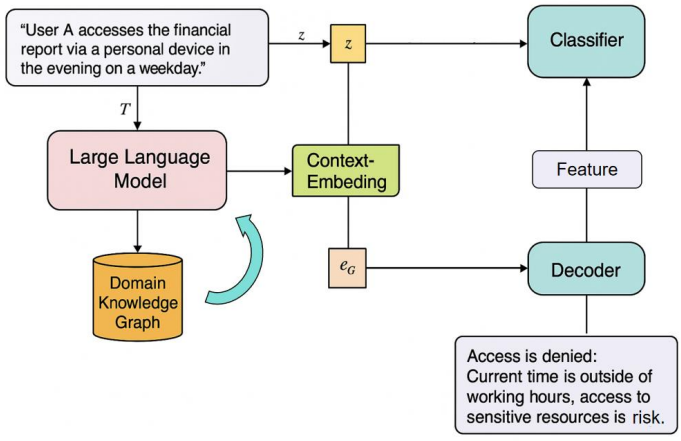

Figure 1. Overall model architecture

Let the user access request be $R = \{u, r, ctx\}$, where u represents the user ID, r represents the target resource, and ctx represents the current context state (such as time, location, device type, historical behavior, etc.). We first convert the above triples into a language representation T and input it into the large language model L_θ to obtain the context

embedding vector $z = L_\theta(T) \in R^d$ as the semantic basis for access policy reasoning.

In the policy generation process, the traditional method relies on a static rule set $\prod = \{p_1, p_2, ..., p_n\}$, while in this method, the policy set is no longer predefined, but is dynamically constructed based on context semantics through a language model. We model the policy as a conditional probability distribution, that is, whether access is allowed is the target variable $y \in \{0,1\}$, then the goal of the model is to learn function $f : z \to y$ and satisfy:

$$P(y = 1 \mid u, r, ctx) = \sigma(W_z + b)$$

$W \in R^{1\times d}, b \in R$ is a trainable parameter and $\sigma(\cdot)$ is a Sigmoid function that outputs the probability value of access permission. The probability threshold is set to $r \in (0,1)$ by the system policy. When $P(y = 1) \geq r$, the access request is allowed; otherwise, it is rejected.

In order to improve the interpretability and policy consistency of access control, we introduce an auxiliary decoder based on policy generation to generate access control explanation text S. This module takes z as input and decodes it into natural language output through the language generation head G_ϕ:

$$S = G_\phi = Decoder(z)$$

This explanation text can be used for system audits and also helps security administrators understand the basis for access decisions, enhancing the credibility and transparency of the model in actual deployment. In addition, to improve the training effect, we introduce a joint loss function:

$$L = L_{acc} + \lambda L_{gen}$$

Where L_{acc} is the binary cross entropy loss for access decision, L_{gen} is the language modeling loss for text generation, and λ is the weight hyperparameter used to balance the importance of the two.

To optimize the model parameters, we use the standard gradient descent strategy to minimize the joint loss function L:

$$\theta \leftarrow \theta - \eta \cdot \nabla_\theta L$$

Among them, η is the learning rate, and θ represents the total gradient of the parameters of the two sub-modules, which ensures that the access prediction and language interpretation modules are optimized synchronously to enhance the overall expressiveness.

In addition, considering that contextual factors in actual environments are highly dynamic and heterogeneous, this method also integrates a small amount of external knowledge enhancement mechanisms to supplement the judgment ability of large language models when domain knowledge is insufficient or input context is missing. Specifically, we introduce the domain security knowledge graph G, extract the entity relationship vector e_G through the embedding module, and concatenate it with the context semantic vector z to form an extended vector $z' = [z; e_G]$, which is further used for decision prediction. This mechanism improves the robustness and generalization ability of the model in fine-grained, long-tail access request scenarios, and effectively expands its scope of application.

4. Experiment

4.1 Datasets

The dataset used in this study is constructed from access control logs and contextual information collected in a real platform environment. It includes records of access requests to various sensitive resources by multiple users under different times, locations, devices, and roles. The data collection period spans three months. The platform is an internal enterprise information management system, covering business modules such as finance, administration, and human resources. It features typical characteristics of multi-user, multi-resource, and multi-context scenarios. During data collection, strict privacy protection protocols were followed. All records were anonymized, and only key fields necessary for structured attributes and contextual semantic representation were retained.

The dataset contains approximately 100,000 access request records. Each sample includes contextual information such as user identity, resource type, timestamp, device type, geographic location, and a summary of the user's historical behavior. Each entry is labeled to indicate whether the access was authorized by the system. To match the input structure of language models, all fields were encoded into natural language format. The semantic relationships of the original context were preserved. The dataset shows a balanced distribution of access types, covering scenarios such as office document access, database queries, and sensitive report viewing. It is suitable for training and evaluating fine-grained permission classification and policy reasoning tasks.

To ensure generalizability and scalability, we specifically considered the impact of contextual variation on access control during dataset construction. For example, a user may receive different authorization outcomes when accessing the same resource at different times, or when using personal versus corporate devices. This design enhances the dataset's sensitivity to context. It also provides a strong foundation for evaluating the generalization ability of models in complex and dynamic environments. Overall, this dataset offers realistic and comprehensive support for exploring context-aware fine-grained access control strategies based on large language models.

4.2 Experimental setup

In the experiments of this study, we selected Meta's open-source large language model LLaMA3 as the core architecture

for semantic modeling and policy reasoning. To ensure experimental stability and reproducibility, all model training and inference tasks were conducted in a unified computing environment. The model was adapted to the tasks through supervised fine-tuning. The training objectives included joint optimization of access authorization classification and access explanation generation. The experiments were implemented using the PyTorch framework. The Hugging Face Transformers library was used to load pretrained weights, integrated with customized data processing and task decoding modules to complete the full experimental pipeline.

Regarding the hardware environment, training was conducted on a high-performance workstation equipped with an NVIDIA RTX 4090 GPU. The CPU was an Intel Xeon Platinum 8474C, with 80GB of RAM. The operating system was Ubuntu 22.04 LTS. To ensure training efficiency and model performance, we tested multiple hyperparameter combinations. The optimal configuration was selected for the main experimental analysis. Detailed experimental settings are shown in Table 1.

Table 1: Experimental Settings and Hyperparameters

Parameter name	Settings
GPU	NVIDIA RTX 4090 (24GB)
CPU	Intel Xeon Platinum 8474C (15 vCPU)
Memory	80GB
System	Ubuntu 22.04 LTS
Deep Learning Frameworks	PyTorch 2.1.0
Python Version	Python 3.10
Pre-trained models	LLaMA3-7B
Maximum input length	1024 tokens
Optimizer	AdamW
Learning rate (initial)	0.001
Batch Size	16

4.3Experimental Results

1) Experiments on accuracy of different models

This paper first gives the access authorization accuracy evaluation experiment of different models, and the experimental results are shown in Table 2.

Table 2: Experiments on access authorization accuracy evaluation of different models.

Model	Acc	Precision	Recall	F1-Score
Bert[6]	78.4	75.6	71.2	73.3
Qwen-7B[7]	84.1	82.3	79.7	81.0
ChatGLM3[8]	89.5	90.1	87.8	88.9
LLAMA3 [9]	91.3	92.5	89.4	90.3
LLAMA3+Ours	94.2	95.0	92.8	93.9

The experimental results show significant performance differences across models in the access authorization task. Improvements in model capability have a direct impact on final classification accuracy. The baseline BERT model performs relatively poorly across all metrics, with an accuracy of 78.4% and an F1-score of 73.3%. This indicates its limitations in semantic understanding and context modeling. It struggles to capture the implicit relationships and multi-dimensional contextual semantics in complex access requests.

In contrast, more advanced Chinese language models such as Qwen-7B and ChatGLM3 demonstrate clear improvements in both accuracy and robustness. Notably, ChatGLM3 exceeds 90% in both precision and recall. This shows its strong ability to understand Chinese language context and infer user access intent. Although Qwen-7B performs slightly lower overall, its stable output makes it suitable for medium-complexity access scenarios.

LLAMA3, as a pretrained English language model, also shows strong performance in this experimental setting. It achieves an F1-score of 90.3%, reflecting its high capability in modeling complex contexts and access patterns in access control tasks. This result suggests that large language models with extensive parameters and context window support can maintain strong transferability and generalization, even in non-native language settings.

After integrating our method (LLAMA3+Ours), all metrics further improve. In particular, recall and F1-score reach 92.8% and 93.9%, respectively. This verifies the effectiveness of our proposed context enhancement mechanism and policy optimization module. Our method builds on the semantic modeling strengths of LLAMA3, while further enhancing the model's ability to perceive fine-grained contextual features. As a result, the system achieves higher decision accuracy and robustness when handling complex and dynamic access requests.

2) Generalization capability test under multi-scenario and multi-role access requests

Furthermore, this paper presents a generalization capability test under multi-scenario and multi-role access requests, and the experimental results are shown in Figure 2.

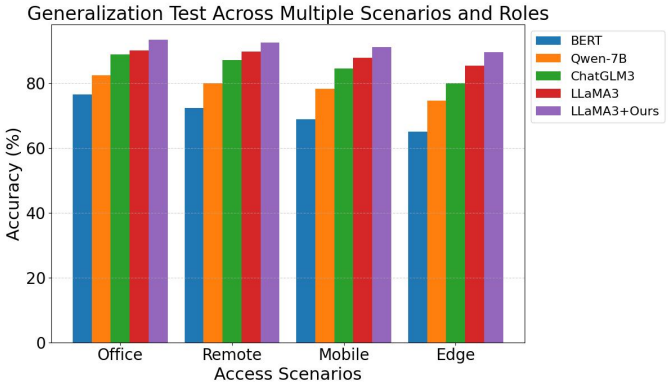

Figure 2. Generalization Test Across Multiple Scenarios and Roles

As shown in the results in Figure 2, different models exhibit noticeable performance differences under multi-scenario and multi-role access requests. BERT shows the lowest accuracy across all scenarios. It performs especially poorly in the “Edge” scenario, with accuracy dropping to around 66%. This indicates a clear weakness in modeling

contextual information under complex environments. It also suggests that traditional pretrained models struggle to adapt effectively to access control tasks in distributed and edge scenarios.

Qwen-7B and ChatGLM3 perform relatively stably across various scenarios, maintaining accuracy above 80% in most cases. They show particular strength in the "Office" and "Remote" scenarios. Between the two, ChatGLM3 demonstrates stronger adaptability in more dynamic contexts such as "Mobile" and "Edge." This confirms its overall advantage in complex semantic modeling and access intent recognition.

LLaMA3 and LLaMA3+Ours outperform all other models across every scenario. Their advantage is especially evident in the "Edge" and "Mobile" settings. LLaMA3+Ours reaches nearly 90% accuracy in the "Edge" scenario. This validates the effectiveness of our proposed context enhancement mechanism and knowledge graph integration method in improving model generalization. These results suggest that the method better captures contextual changes in dynamic environments, enabling more robust fine-grained access control.

3) Comparative experiment on the performance of different model sizes in control tasks

Next, this paper also gives a comparative experiment on the performance of different model sizes in control tasks, and the experimental results are shown in Figure 3.

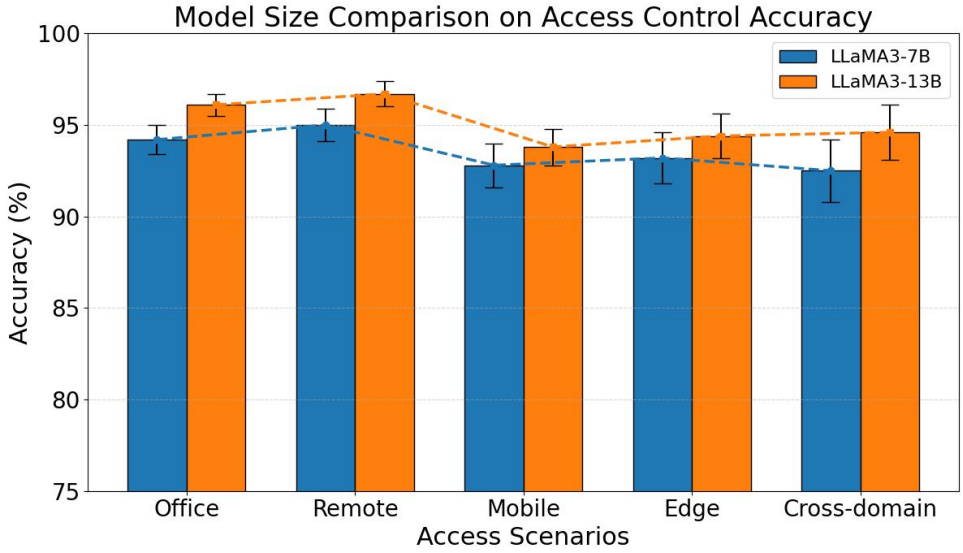

Figure 3. Model Size Comparison on Access Control Accuracy

As shown in Figure 3, LLaMA3-13B outperforms LLaMA3-7B across all access scenarios. It demonstrates stronger overall modeling and context understanding capabilities. The performance gap is particularly evident in semi-structured or semantically variable scenarios such as "Remote" and "Mobile." In these cases, LLaMA3-13B reaches or exceeds 96% accuracy, while LLaMA3-7B performs slightly lower. This indicates that larger models are more effective at capturing complex semantic relations and context-dependent features.

Across all scenarios, LLaMA3-13B shows more stable accuracy. Its error bars (standard deviation) are significantly smaller than those of LLaMA3-7B. This suggests that the larger model not only has better predictive ability but also exhibits stronger robustness and generalization. In the "Cross-domain" scenario, the performance gap between the two models remains stable. This further confirms the suitability of larger models for access control across roles and contextual boundaries.

In addition, the trend lines indicate that LLaMA3-13B performs more consistently across different access policy scenarios. In contrast, LLaMA3-7B shows noticeable fluctuations under complex conditions. These results support the strategy of prioritizing large-parameter models in highly dynamic and security-sensitive systems, especially in tasks requiring high-confidence access control decisions.

4) Hyperparameter sensitivity experiments

This paper also conducts experiments on hyperparameter sensitivity, mainly adjusting and analyzing the learning rate and optimizer. First, the experimental results on the learning rate are given, as shown in Table 3.

Table 3: Hyperparameter sensitivity experiment results (learning rate)

Lr	Acc	Precision	Recall	F1-Score
0.005	89.3	87.9	90.1	89.0
0.004	91.0	91.8	91.2	91.5
0.003	92.6	93.0	91.9	92.4
0.002	93.5	94.1	92.3	93.2
0.001	94.2	95.0	92.8	93.9

As shown in Table 3, model performance varies significantly under different learning rates (Lr). Overall, there is a clear trend: smaller learning rates lead to better performance. When the learning rate is set to 0.005, the model achieves only 89.3% accuracy. Both precision and recall are

relatively low. This suggests that a large step size may cause training instability, making it difficult for the model to fit the fine-grained access control task effectively.

As the learning rate decreases, the model shows continuous improvement across all metrics. In particular, within the 0.002 to 0.001 range, metrics such as accuracy and F1-score approach saturation. This indicates that the model is converging and can better capture contextual differences in access requests. The increase in precision suggests more stable decisions on permitted accesses. The improvement in recall reflects greater sensitivity in identifying valid requests.

Ultimately, with a learning rate of 0.001, the model achieves optimal performance across all four metrics. The F1-score reaches 93.9%, showing a good balance between accuracy and robustness. Overall, the results demonstrate that learning rate, as a key hyperparameter, has a significant impact on training large language models for access control tasks. Proper tuning can effectively enhance system performance.

Next, the experimental results of different optimizers are given, as shown in Table 4.

Table 4: Hyperparameter sensitivity experiment results (Optimizer)

Optimizer	Acc	Precision	Recall	F1-Score
AdaGrad [10]	89.1	87.4	88.7	88.0
SGD [11]	90.3	89.8	89.1	89.4
Adam [12]	92.5	93.1	91.0	92.0
AdamW [13]	94.2	95.0	92.8	93.9

As shown in Table 4, different optimizers have a significant impact on model performance in the access control task. Among them, the AdamW optimizer achieves the best results, leading across all evaluation metrics. It reaches an F1-score of 93.9%, indicating that its advantages in weight regularization and gradient updates help improve model generalization and training stability.

In comparison, the Adam optimizer also performs well, with an F1-score of 92.0%, but slightly lower than AdamW. This may be due to its limitations in handling long-term dependencies and parameter normalization. Although SGD and AdaGrad work well on simpler tasks, their performance is limited in high-dimensional semantic modeling. In particular, AdaGrad achieves only 88.0% F1-score, reflecting slow convergence and weak optimization ability under complex contextual scenarios.

Overall, the results confirm that choosing an appropriate optimizer is crucial for improving the performance of large language models in access control tasks. AdamW, with its enhanced weight decay strategy, captures subtle contextual variations more effectively and provides more stable support for fine-grained access decisions.

5) Loss function changes with epoch results

Finally, the loss function drop graph is given, as shown in Figure 4.

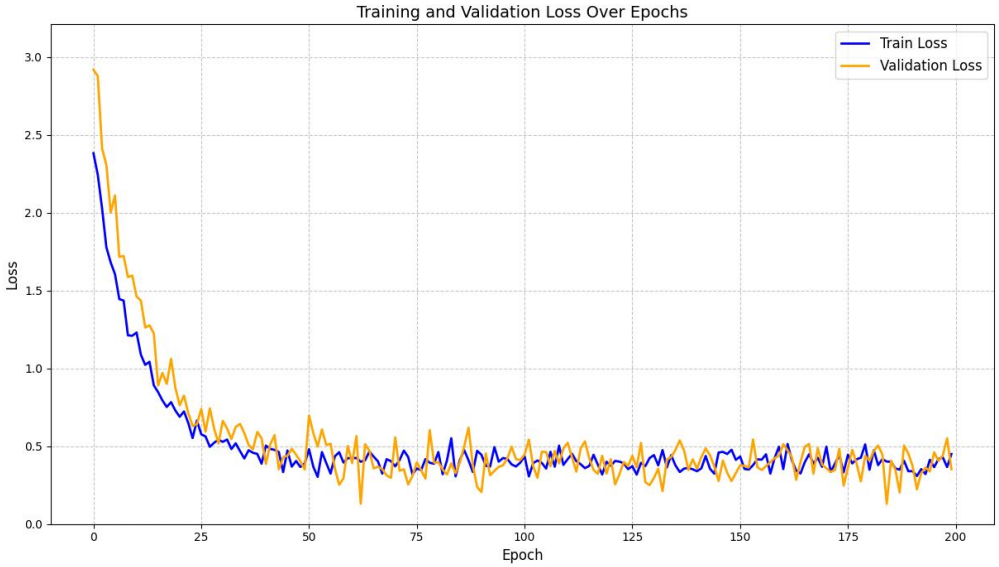

Figure 4. Loss function drop graph

As shown in Figure 4, both the training loss and validation loss drop rapidly within the first 20 epochs. This indicates that the model's fitting ability improves significantly in the early stages. The parameter updates are effective, and the model quickly learns the basic patterns and semantic structures of the access control task.

Around the 30th epoch, the loss curves begin to stabilize. Both training and validation loss fluctuate around 0.4, suggesting that the model has generally converged without obvious overfitting. Although there are slight fluctuations at individual points on the validation curve, the overall trend remains consistent with the training loss. This indicates good generalization capability.

In addition, the gap between the training and validation loss remains small throughout the 200 training epochs. This shows that the model maintains consistent optimization

performance across different data distributions. Overall, the loss reduction process validates the effectiveness of the model architecture and training strategy. It provides a solid foundation for subsequent accuracy evaluation and access policy generation.

5. Conclusion

This paper proposes a context-aware fine-grained access control framework based on a large language model. It aims to address the limitations of traditional access control methods, such as weak semantic understanding and rigid policies in dynamic environments. By introducing the LLaMA3 model to jointly model user requests, resource attributes, and contextual information, the framework enables semantic-level access intent reasoning and decision generation. This significantly improves adaptability and accuracy in complex scenarios. In addition, the integration of knowledge graph embeddings and explanation generation enhances both robustness and interpretability, meeting the dual requirements of security and transparency in real-world applications.

In the experimental section, we evaluate the model from multiple perspectives, including generalization ability, hyperparameter sensitivity, and optimizer comparison. Results show that our method outperforms existing models across various access scenarios. It demonstrates higher classification accuracy and stability, especially in complex contexts and cross-role requests. Moreover, the training process shows good convergence, with smooth loss reduction curves, validating the effectiveness of our optimization strategies and model design.

This study not only provides a technical path for intelligent access control systems but also demonstrates the great potential of large language models in security decision-making. Compared to static rule-based systems, our method adapts dynamically to changes in access environments in a data-driven manner. It highlights the unique advantages of large models in real-time policy evolution and multi-dimensional information integration. This capability offers valuable insights for building future-oriented security control architectures. Future work will proceed in two directions. First, we plan to incorporate multimodal contextual data, such as voice commands and image-based identity verification, to further enrich semantic understanding. Second, we aim to deploy the model in edge device environments. By integrating lightweight techniques, we will improve inference efficiency and deployment flexibility, promoting the broader application of large language models in practical security systems.

References

[1] Ahn, Gail-Joon, et al. "Policy administration for language-based access control." IEEE Transactions on Dependable and Secure Computing 4.3 (2007): 201–215.

[2] Hu, Vincent C., et al. "Guide to attribute based access control (ABAC) definition and considerations." NIST Special Publication 800.162 (2014): 1–49.

[3] Brown, Tom B., et al. "Language models are few-shot learners." Advances in Neural Information Processing Systems 33 (2020): 1877–1901.

[4] Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., & Gao, J. (2024). Large Language Models: A Survey. arXiv e-prints, arXiv-2402.

[5] Zhang, Rui, et al. "Blockchain-based data sharing system for AI-powered network security." IEEE Network 34.6 (2020): 24–29.

[6] Taylor, Andrew, et al. "Anomaly detection in automotive controller area networks using neural networks." IEEE Transactions on Intelligent Transportation Systems 20.5 (2019): 1780–1790.

[7] Lee, Jinhyuk, et al. "BioBERT: a pre-trained biomedical language representation model for biomedical text mining." Bioinformatics 36.4 (2020): 1234–1240.

[8] Chen, Hsinchun, et al. "COPLINK: managing law enforcement data and knowledge." Communications of the ACM 46.1 (2003): 28–34.

[9] Touvron, Hugo, et al. "LLaMA: Open and efficient foundation language models." arXiv preprint arXiv:2302.13971 (2023).

[10] Ward, R., Wu, X., & Bottou, L. "Adagrad stepsizes: Sharp convergence over nonconvex landscapes." Journal of Machine Learning Research 21.219 (2020): 1–30.

[11] Gower, R. M., Loizou, N., Qian, X., Sailanbayev, A., Shulgin, E., & Richtárik, P. "SGD: General analysis and improved rates." In International Conference on Machine Learning (2019): 5200–5209.

[12] Zhang, Z. "Improved Adam optimizer for deep neural networks." In 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS) (2018): 1–2. IEEE.

[13] Loshchilov, Ilya, and Frank Hutter. "Decoupled weight decay regularization." arXiv preprint arXiv:1711.05101 (2017).