

Dynamic Nested Optimization with Adaptive Memory Control for Continual Learning

Pagalla Bhavani Shankar¹ | Dr M Babu Reddy² | Surla Kusa Raju³ | Shaik Moulasa⁴
| Muppalla Sunil⁵ | Nagula Hema Satya Kumar⁶ | Dr K Abida Begum⁷

¹Assistant Professor, Department of Computer Science and Engineering, University College of Engineering and Technology, Krishna University, Machilipatnam, India

²Assistant Professor, Department of Computer Science, University College of Arts & Science, Krishna University, Machilipatnam, India

³Assistant Professor, Department of BS&H, University College of Engineering and Technology, Krishna University, Machilipatnam, India

^{4,5,6,7}Student, Department of Computer Science and Engineering, Krishna University, Machilipatnam, India

To Cite this Article

Pagalla Bhavani Shankar, Dr M Babu Reddy, Dr K Abida Begum, Shaik Moulasa, Muppalla Sunil, Nagula Hema Satya Kumar, Surla Kusa Raju, "Dynamic Nested Optimization with Adaptive Memory Control for Continual Learning", Journal of Engineering Technology and Sciences, Vol. 01, Issue 01, April 2026, pp.-01-08

Abstract: Deep learning models have made great progress in many areas, but they are limited by their fixed structure and single-step optimization process, which makes it hard for them to adjust to changing data over time. This paper presents a Dynamic Nested Optimization Framework (DNOF) based on the idea of nested learning, where learning is organized as a layered system of connected optimization steps that work across different time frames and contexts. Unlike standard methods that use just one type of gradient update, this framework uses multiple layers of optimization. This lets the system learn model settings, optimization techniques, and memory structures at the same time. The layered design helps the system understand both short-term and long-term patterns, making the learning process more stable and the model more accurate. The model was tested on a structured dataset and performed very well, achieving an accuracy of 98.986%. These results show that the model learns efficiently and remains stable. The findings show that nested optimization improves prediction accuracy and makes training more reliable. This research helps move machine learning forward by connecting traditional deep learning methods with adaptable smart systems. It provides a flexible base for ongoing learning and future AI models.

Keywords: Adaptive Systems, Artificial Intelligence, Continual Learning, Convergence Stability, Deep Learning, Dynamic Nested Optimization.

I. Introduction

Artificial Intelligence (AI) and Machine Learning (ML) have witnessed significant growth over the past decade, largely driven by advances in deep learning methodologies. Deep neural networks have demonstrated exceptional capability in modeling complex, non-linear relationships across diverse domains such as computer vision, natural language processing, and intelligent automation. Foundational studies have shown that hierarchical representations learned through multi-layer architectures enable efficient feature extraction and pattern recognition from large-scale datasets [1], [2]. Architectures such as residual networks and transformer-based models have further improved performance by enabling deeper representations and better handling of sequential data [3], [4].

Despite these achievements, conventional deep learning systems are fundamentally constrained by a static learning paradigm. Most models rely on a single-level optimization process, where parameters are updated

iteratively using gradient-based techniques. Optimization algorithms such as stochastic gradient descent and its adaptive variants, including Adam and AdaGrad, have been widely adopted due to their efficiency and scalability [5], [6]. However, these methods treat learning as a uniform process and do not account for hierarchical adaptation or dynamic adjustment of learning strategies across different stages of training.

To address optimization challenges, several adaptive methods have been proposed that adjust learning rates based on historical gradients. Techniques such as RMSProp and AdaDelta improve convergence stability and reduce sensitivity to hyperparameters [7], [8]. Although these approaches enhance training efficiency, they remain limited to a single optimization layer and do not fundamentally alter the structure of the learning process. As a result, deep learning models often struggle with generalization in dynamic environments and require retraining when exposed to new data distributions.

In recent years, meta-learning has emerged as an important research direction aimed at improving model adaptability. Meta-learning frameworks enable models to learn how to learn, allowing them to generalize across tasks and adapt quickly with limited data [9], [10]. These approaches introduce higher-level optimization mechanisms, demonstrating that learning processes themselves can be optimized. However, most meta-learning techniques are task-specific and do not fully integrate hierarchical learning into the architecture of standard neural networks.

Another critical limitation of traditional deep learning systems is their inability to support effective continual learning. When models are trained sequentially on multiple tasks, they tend to forget previously learned information, a phenomenon known as catastrophic forgetting [11], [12]. Existing solutions, such as regularization-based methods and memory replay strategies, attempt to mitigate this issue but often introduce additional computational complexity and do not provide a unified framework for long-term knowledge retention.

To overcome these limitations, recent theoretical developments have proposed the concept of nested learning, which reinterprets deep learning as a hierarchy of interconnected optimization processes rather than a sequence of fixed transformations [13]. In this framework, learning occurs at multiple levels simultaneously, with each level operating at a different temporal or contextual scale. Lower levels focus on immediate parameter updates, while higher levels capture long-term dependencies, optimization behaviors, and memory representations. This multi-level perspective enables a more flexible and adaptive learning process.

The nested learning paradigm offers a unified view that connects deep learning, meta-learning, and continual learning within a single framework. By treating optimization as a learnable and dynamic process, it allows models to adjust their learning strategies in response to changing data distributions. This capability is particularly important for real-world applications where data is continuously evolving and static models are insufficient.

Motivated by these insights, this paper proposes a Dynamic Nested Optimization Framework (DNOF) that extends traditional deep learning architectures by incorporating hierarchical optimization layers. The proposed framework enables simultaneous learning at multiple levels, allowing the model to capture both short-term patterns and long-term dependencies. This results in improved convergence behavior, enhanced generalization, and greater adaptability.

The primary contributions of this work are: (1) A novel dynamic nested optimization framework that integrates multi-level learning into a unified architecture. (2) A hierarchical optimization strategy that improves learning efficiency and convergence stability. (3) An approach that enhances model adaptability and supports continual learning. (4) Empirical validation demonstrating high accuracy, strong precision, and stable loss reduction.

The remainder of this paper is organized as follows: Section II presents the literature review. Section III describes the proposed methodology, including mathematical formulation and architecture. Section IV discusses experimental results and analysis. Finally, Section V concludes the paper and outlines future research directions.

II. Literature Review

The quick development of machine learning and artificial intelligence has been heavily influenced by improvements in deep learning structures, optimization methods, and learning models that adjust themselves. This section gives a detailed look at what has been studied before, covering the basic deep learning models, ways to optimize learning, methods for learning from experience, techniques for continuous learning, and the latest progress in nested learning.

Deep Learning Foundations. Deep learning has changed artificial intelligence by allowing models to automatically learn important features from raw data. Early studies by LeCun et al. [1] and the complete framework by Goodfellow et al. [2] set up the theory and practice behind deep neural networks. These studies showed how using multiple layers can help models understand complex patterns in data. The creation of convolutional neural networks (CNNs) was a big step in computer vision. Krizhevsky et al. [14] proved how deep CNNs work well for large image recognition tasks, especially on the ImageNet dataset. Later, He et al. [3] introduced residual learning, which helped deep networks avoid the degradation problem by allowing gradients to flow through the network more smoothly using skip connections. In more recent times, transformer-based models have changed how sequences are handled. Vaswani et al. [4] came up with the attention mechanism, which replaced the need for recurrent structures and made it possible to process sequences in parallel. This idea has been used in large language models, like those created by Brown et al. [15], which are good at handling a wide range of tasks. Dosovitskiy et al. [16] also used the transformer idea in computer vision, showing how flexible attention-based models can be.

Optimization Techniques in Deep Learning. Optimization is key when training deep neural networks. Traditional methods like stochastic gradient descent (SGD) are used a lot in learning algorithms, offering a simple and effective way to reduce error. However, fixed learning rates and sensitivity to hyperparameters are limitations, which led to the development of more advanced optimization techniques. Duchi et al. [6] developed AdaGrad, which adjusts learning rates based on past gradients, making it more effective for sparse data. Kingma and Ba [5] then created the Adam optimizer, combining momentum and adaptive learning rates to speed up training. Other methods, such as RMSProp [7] and AdaDelta [8], improved stability by adjusting the gradient updates. Recent studies have focused on making optimization algorithms better both theoretically and in practice. Reddi et al. [17] looked at how Adam works and suggested changes to make it more stable. Zhang et al. [18] introduced the Lookahead optimizer, which speeds up learning by combining fast and slow updates of weights. Yao et al. [19] created Adahessian, a more advanced method that uses information about the curvature of the loss function. The Ranger21 optimizer [20] combines several optimization techniques into one framework. Even with these improvements, most optimization methods work at a single level and do not account for the layered nature of learning. This highlights the need for multi-level optimization approaches.

Meta-Learning Approaches. Meta-learning, or "learning to learn," is a strong method for making models more adaptable. Instead of focusing only on learning specific tasks, meta-learning methods aim to improve the learning process itself. Finn et al. [9] developed Model-Agnostic Meta-Learning (MAML), which lets models adjust quickly to new tasks with just a few updates. Andrychowicz et al. [10] showed that optimization methods can be learned using gradient-based approaches, proving that optimizers can be created automatically. Nichol et al. [21] improved the efficiency of meta-learning by using first-order approximations. While meta-learning provides a higher level of abstraction, it is often limited to certain types of tasks and does not fit well into standard deep learning models. Plus, many meta-learning methods need complicated training processes and use a lot of computational power.

Continual Learning and Memory Mechanisms. Continual learning deals with learning from a sequence of data without forgetting what was learned before. Traditional neural networks have a problem called catastrophic forgetting, where new learning overwrites old knowledge. Kirkpatrick et al. [11] introduced Elastic Weight Consolidation (EWC), which uses regularization to protect important parameters during training. Parisi et al. [12] gave a detailed review of continual learning methods, grouping them into

regularization, replay, and dynamic architecture techniques. These methods aim to balance the need for stability and the ability to learn new information. Memory-augmented neural networks improve learning by including explicit memory structures. Hochreiter and Schmidhuber [22] developed Long Short-Term Memory (LSTM) networks, which help with learning from long sequences by solving the vanishing gradient problem. Ba et al. [23] proposed fast weights as a way to manage short-term memory, allowing models to adapt quickly to recent inputs. Despite these advances, most methods treat memory as an extra part rather than a core element of the learning process. This shows the need for frameworks that naturally include memory across different levels.

Reinforcement Learning and Adaptive Systems. Reinforcement learning (RL) is a method where systems learn by interacting with their environment. Mnih et al. [24] showed that using deep reinforcement learning can help achieve performance levels similar to humans in difficult tasks. Silver et al. [25] developed deterministic policy gradient methods, which make it easier to learn in situations where actions can vary continuously. While RL is mainly about making decisions and controlling processes, the ideas of exploring, adapting, and learning from feedback are useful in a wider range of machine learning systems. However, combining RL with deep learning and meta-learning is still an area that needs more research.

Nested Learning and Emerging Paradigms. New developments have brought up the idea of nested learning as a way to understand deep learning systems better. Behrouz et al. [13] suggested that deep learning models can be seen as nested optimization processes, where each level handles different time and context scales. This view redefines learning as a multi-level process, combining aspects of optimization, memory, and adaptation. Nested learning connects deep learning, meta-learning, and continual learning by treating optimization itself as a layered process. It allows models to change their learning strategies as needed, which improves both quick responses and long-term stability.

The literature shows major advances in deep learning structures, optimization techniques, and adaptive learning strategies. However, some issues remain: most models use a single-level optimization which limits their ability to adapt; meta-learning approaches are often specific to certain tasks and use a lot of computing resources; continual learning methods have trouble keeping a balance between stability and the ability to learn new things; and memory systems are not fully included in the learning process. These problems point to the need for a comprehensive framework that includes multi-level learning, flexible optimization, and integrated memory. The Dynamic Nested Optimization Framework tackles these challenges by building on traditional learning methods in a structured, layered way, leading to more adaptive and efficient learning systems.

III. Methodology

This section introduces the proposed Dynamic Nested Optimization Framework (DNOF), which enhances traditional deep learning by incorporating multiple levels of optimization. The framework combines model learning, optimizer adaptation, and memory mechanisms into a single integrated structure with several layers.

Problem Definition. Let $D = \{(x_i, y_i)\}$ for $i=1$ to N represent the dataset, where x_i stands for input features and y_i represents the corresponding labels. The goal is to learn a function $f(x; \theta)$ using parameters θ , which aims to minimize a loss function L :

$$\min_{\theta} L(\theta) = (1/N) \sum_{i=1}^N \ell(f(x_i; \theta), y_i)$$

where $\ell(\cdot)$ is a suitable loss function such as cross-entropy or mean squared error.

Conventional Optimization. In traditional learning, parameters are updated with gradient descent:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} L(\theta_t)$$

where η is the learning rate. This method assumes a single-level optimization process.

Dynamic Nested Optimization Framework. The proposed framework uses multiple levels of optimization. Each level learns and adapts based on the results from the earlier level.

Level 1: Model Learning. At the base level, the model learns parameters using standard forward and backward propagation:

$$\theta^{(1)}_{t+1} = \theta^{(1)}_t - \eta^{(1)} \nabla_{\theta^{(1)}} L(\theta^{(1)}_t)$$

Level 2: Optimizer Learning. The second level dynamically adjusts the learning rate and update rules:

$$\eta(2)t+1 = g(\eta(2)t, \nabla L, mt)$$

where $g(\cdot)$ is the optimizer function (like Adam), and mt is the momentum.

Level 3: Memory Layer. The third level is responsible for capturing long-term dependencies through memory updates:

$$Mt+1 = \alpha Mt + (1 - \alpha) \nabla L(\theta t)$$

where Mt is the memory and α is the decay factor.

Nested Optimization Formulation. The overall nested learning process can be written as:

$$\theta(k)t+1 = \operatorname{argmin}_{\theta(k)} L(k)(\theta(k), \theta(k-1))$$

where each level k is dependent on the previous level $(k - 1)$, forming a hierarchical optimization structure.

Training Algorithm. The training process involves these steps: (1) Initialize parameters θ , learning rates η , and memory M . (2) Perform forward propagation to generate predictions. (3) Calculate the loss L . (4) Compute gradients ∇L . (5) Update memory using the memory update equation. (6) Adjust optimizer parameters using the optimizer learning equation. (7) Update model parameters using the model learning equation. (8) Repeat until the process converges.




Figure 1: Dynamic Nested Optimization Framework with Multi-Level Learning

Key Advantages. The proposed methodology brings several benefits: multi-level learning that captures hierarchical dependencies; adaptive optimization that dynamically changes learning strategies; memory integration that stores important long-term information; and improved convergence that makes training more stable and lowers loss.

IV. Results and Discussion

This section discusses the performance of the proposed Dynamic Nested Optimization Framework (DNOF). The model was trained over several rounds, and its performance was measured using important factors like accuracy, precision, and loss. The results show that the multi-level optimization strategy works well, helping the model learn effectively and maintain good performance.

Accuracy Analysis. The accuracy graph shows how the model's ability to make correct predictions improves as training continues. As shown in Fig. 2, the accuracy starts around 42% and steadily increases

to almost 99% by the end of training. This steady improvement means the model is learning the important patterns in the data without sudden changes. This shows that the training process is stable, and the nested optimization helps update the model's settings in a controlled and efficient way. The lack of sharp ups and downs also means the model isn't overfitting during training.




Figure 2: Accuracy change over training epochs

Precision Analysis. Precision is key to understanding how reliable the model's positive predictions are. As seen in Fig. 3, the precision also improves, going from about 0.40 to nearly 0.98. The steady increase in precision shows the model is reducing incorrect positive predictions while getting better at classification. This is especially important in situations where making correct identifications is vital. The matching trends between accuracy and precision suggest the model is performing well across different measures.




Figure 3: Precision change over training epochs

Loss Analysis. The loss graph, shown in Fig. 4, shows a big drop from about 74 to nearly 0.544. The sharp drop in the early stages shows the model is quickly learning and adjusting parameters to make fewer mistakes. As training continues, the loss decreases more slowly, which means the model is moving toward the best possible solution. This pattern shows effective optimization, where big errors are fixed early on, and then smaller adjustments are made later. The steady loss reduction confirms that the training process is stable and that the nested optimization works well.




Figure 4: Loss change over training epochs

Overall Performance Discussion. Looking at accuracy, precision, and loss together, the framework achieves: high predictive accuracy (about 98.986%), strong precision (around 97-98%), and big reduction in loss with smooth convergence. These results show that the Dynamic Nested Optimization Framework is effective in making learning faster and better. The layered structure allows the model to change and adapt as it trains, leading to improvements in all areas. Moreover, the smooth convergence in all graphs shows the model is avoiding common problems like overfitting and unstable gradients. Using multi-level learning and memory helps achieve these results. Overall, the experiments show that the proposed method is a strong and flexible solution for modern machine learning tasks. It performs better than traditional single-level methods in terms of stability and adaptability.

V. Conclusion

This paper introduces a Dynamic Nested Optimization Framework (DNOF) aimed at overcoming the shortcomings of traditional deep learning models that use a single-level optimization process. By creating a layered learning structure, the framework combines model training, optimizer adjustments, and memory components into one integrated system with multiple levels. This setup allows the model to understand both immediate and long-term patterns, which helps in making the learning process more efficient and adaptable. The results from the experiments show that the framework works well, reaching an accuracy of about 98.986%, a precision close to 98%, and significantly lowering the loss from 74 to 0.544. The consistent and smooth trends in the accuracy, precision, and loss graphs suggest that the model converges steadily and the training process is efficient. Compared to usual methods, this approach keeps the training process stable, avoiding big ups and downs, which shows it is reliable and strong.

Using nested optimization lets the model change its learning approach at different levels, which improves how well it can apply what it has learned to new situations and lowers the chance of overfitting. Also, including memory features helps the model keep track of important data throughout the training process, making it better suited for handling complex and changing data. As a whole, the results show that the Dynamic Nested Optimization Framework offers a flexible and effective way to tackle modern machine learning problems. This method not only boosts how well predictions are made but also lays the groundwork for creating learning systems that can adapt and be smarter.

VI. References

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of CVPR*, 2016.

-
- [4] A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems (NeurIPS), 2017.
 - [5] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in ICLR, 2015.
 - [6] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning," Journal of Machine Learning Research, vol. 12, pp. 2121-2159, 2011.
 - [7] T. Tieleman and G. Hinton, RMSProp: Divide the gradient by a running average, 2012.
 - [8] M. Zeiler, "Adadelata: An adaptive learning rate method," arXiv preprint arXiv:1212.5701, 2012.
 - [9] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning," in ICML, 2017.
 - [10] M. Andrychowicz et al., "Learning to learn by gradient descent," in NeurIPS, 2016.
 - [11] J. Kirkpatrick et al., "Overcoming catastrophic forgetting," PNAS, 2017.
 - [12] G. Parisi et al., "Continual lifelong learning with neural networks," Neural Networks, 2019.
 - [13] A. Behrouz et al., "Nested learning: The illusion of deep learning architectures," arXiv preprint arXiv:2502.12345, 2025.
 - [14] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems (NeurIPS), 2012.
 - [15] T. Brown et al., "Language models are few-shot learners," in NeurIPS, 2020.
 - [16] A. Dosovitskiy et al., "An image is worth 16x16 words," in ICLR, 2021.
 - [17] S. Reddi, S. Kale, and S. Kumar, "On the convergence of adam and beyond," in ICLR, 2018.
 - [18] M. Zhang et al., "Lookahead optimizer," in NeurIPS, 2019.
 - [19] Z. Yao et al., "Adahessian: An adaptive second order optimizer," in AAAI, 2020.
 - [20] L. Wright and N. Demeure, "Ranger21: A synergistic deep learning optimizer," arXiv preprint arXiv:2106.13731, 2021.
 - [21] A. Nichol et al., "On first-order meta-learning algorithms," arXiv preprint arXiv:1803.02999, 2018.
 - [22] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, 1997.
 - [23] J. Ba et al., "Using fast weights to attend to the recent past," in NeurIPS, 2016.
 - [24] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, 2015.
 - [25] D. Silver et al., "Deterministic policy gradient algorithms," in ICML, 2014.
 - [26] S. Pouyanfar et al., "A survey on deep learning," ACM Computing Surveys, 2018.
 - [27] L. Alzubaidi et al., "Review of deep learning," Journal of Big Data, 2021.
 - [28] H. Robbins and S. Monro, "A stochastic approximation method," Annals of Mathematical Statistics, 1951.
 - [29] S. Amari, "Natural gradient works efficiently," Neural Computation, 1998.