

SEE: Continual Fine-tuning with Sequential Ensemble of Experts

Zhilin Wang^{♣*}, Yafu Li[♣], Xiaoye Qu^{♣†}, Yu Cheng^{♡†}
♣ Shanghai AI Laboratory ♡ Chinese University of Hong Kong
{linzwc, yafuly}@gmail.com, quxiaoye@pjlab.org.cn
chengyu@cse.cuhk.edu.hk

Abstract

Continual fine-tuning of large language models (LLMs) suffers from catastrophic forgetting. Rehearsal-based methods mitigate this problem by retaining a small set of old data. Nevertheless, they still suffer inevitable performance loss. Although training separate experts for each task can help prevent forgetting, effectively assembling them remains a challenge. Some approaches use routers to assign tasks to experts, but in continual learning, they often require retraining for optimal performance. To address these challenges, we introduce the Sequential Ensemble of Experts (SEE) framework. SEE removes the need for an additional router, allowing each expert to independently decide whether a query should be handled. The framework employs distributed routing, and during continual fine-tuning, SEE only requires the training of new experts for incoming tasks rather than retraining the entire system. Experiments reveal that the SEE outperforms prior approaches, including multi-task learning, in continual fine-tuning. It also demonstrates remarkable generalization ability, as the expert can effectively identify out-of-distribution queries, which can then be directed to a more generalized model for resolution. This work highlights the promising potential of integrating routing and response mechanisms within each expert, paving the way for the future of distributed model ensembling. The code is available at [url](https://github.com/linzwc/SEE).

1 Introduction

Large language models (LLMs) have demonstrated impressive performance across various scenarios (Touvron et al., 2023; OpenAI et al., 2024; Guo et al., 2025; Li et al., 2025; Qu et al., 2025). However, they often require further refinement or the ability to learn new tasks for real-world applications. In these scenarios, LLMs are typically

* Work was done during Zhilin Wang’s internship at Shanghai AI Laboratory.

† Corresponding authors.

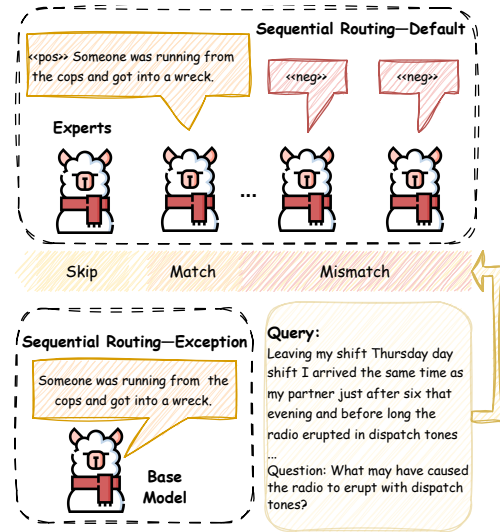


Figure 1: The inference process in the SEE framework involves a query being passed sequentially through a series of experts until it matches one, which then generates a response. If an expert fails to produce a special indicator, the query is routed to the base model, which is considered to possess the best generalization ability.

trained through a sequence of tasks performed successively, a process known as continual fine-tuning (Luo et al., 2024). As a subclass of continual learning, continual fine-tuning also suffers catastrophic forgetting, whereby LLMs tend to lose previously acquired knowledge as they assimilate new information (Kirkpatrick et al., 2017; Wu et al., 2024; Shi et al., 2024; Qu et al., 2025).

While there have been many studies on alleviating catastrophic forgetting, rehearsal-based methods prove to be the most effective way in continual fine-tuning (Zhang et al., 2023). These methods work by combining new data with a subset of previous data (de Masson d’Autume et al., 2019; Rolnick et al., 2019), or synthesized instances from the prior data distribution (Huang et al., 2024). However, despite their effectiveness in preserving prior knowledge, these methods still experience some

inevitable performance loss.

Although training separate experts for each task can prevent forgetting, organizing these experts effectively remains a challenge. The methods of organizing multiple models to create a more powerful model are commonly called Mixture of Experts (MoE) (Zhu et al., 2024; Qu et al., 2024). However, these methods typically operate on a fixed set of tasks and domains (Lu et al., 2023; Jiang et al., 2023) and introduce additional routers to manage the models (Jang et al., 2023; Lu et al., 2024). Except for their inefficiency in incremental scenarios, recent work (Lv et al., 2025) also suggests that the separations between the router’s decision-making and experts’ execution may lead to suboptimal expert selection and ineffective learning.

To this end, we propose a framework referred to as the **Sequential Ensemble of Experts (SEE)**, which combines the rehearsal-based method with MoE. SEE comprises a base model and a sequence of experts, each proficient in recognizing questions within its specialized domain and generating precise responses accordingly. These experts are derived from the base model, such as Llama2, and are further specialized for specific tasks through LoRA training. When a new task arrives, SEE combines the new data with a small subset of instances from previous tasks, similar to rehearsal-based methods, and reconstructs the dataset to generate both positive and negative instances. A new expert is then trained through supervised fine-tuning (SFT) on these instances and appended to the end of the expert sequence. SEE integrates all experts via a distributed routing mechanism called sequential routing, as illustrated in Figure 1.

Empirical experiments on task sequences from the SuperNI dataset show that SEE outperforms previous rehearsal-based methods and matches or exceeds multi-task learning (MTL) performance. Besides, SEE further outperforms rehearsal-based methods on MMLU (Hendrycks et al., 2021) and exhibits astonishing knowledge preservation and out-of-distribution (OOD) generalization ability. Moreover, when regarded as a single model, the perplexity of SEE proves lower than MTL on tasks during continual learning. Additionally, while SEE uses special indicators to identify a query, our analysis demonstrates that introducing additional tokens as indicators is more efficient than using tokens from the vocabulary for constructing indicators. Although SEE can theoretically utilize sam-

ples not included in the continual learning process to construct negative instances, this would significantly degrade routing accuracy, thus highlighting the essence of rehearsal. Finally, we demonstrate that the additional latency introduced by sequential routing is affordable.

2 Preliminaries

Continual Fine-tuning Let \mathcal{M} denote the LLM. The model undergoes continual fine-tuning across N stages, with each stage i involving updates based on the instruction dataset $D^{(i)}$ corresponding to task T_i (Luo et al., 2024). Formally, the update process at each stage is defined as:

$$\mathcal{M}_i = \text{Update}(\mathcal{M}_{i-1}, D^{(i)}) \quad (1)$$

where \mathcal{M}_0 denotes the initial model.

Rehearsal-based Methods As extensively discussed in the literature (de Masson d’Autume et al., 2019; Rolnick et al., 2019), Rehearsal-based methods involve sampling a subset of instances from earlier stages to expand the training data for the current stage. Formally, the augmented training set is expressed as:

$$D^{(t)} \cup \sum_{i=1}^{t-1} \left(\tau D^{(i)} \right) \quad (2)$$

where τ denotes the rehearsal ratio, representing the proportion of training instances sampled from previous stages. At stage i , the model \mathcal{M}_i is updated as follows:

$$\mathcal{M}_i = \text{Update}(\mathcal{M}_{i-1}, D^{(t)} \cup \sum_{i=1}^{t-1} \left(\tau D^{(i)} \right)) \quad (3)$$

The rehearsal-based methods can effectively mitigate the issue of catastrophic forgetting in LLMs, as demonstrated in prior studies (Scialom et al., 2022; Mok et al., 2023).

3 Methods

In contrast to previous research, our work integrates the concept of MoE with rehearsal-based methods. For each task T_i , we assign a specific expert E_i to handle it. These experts are able to assess whether a query falls within their scope of responsibility. If so, they proceed to generate the response. SEE integrates the experts and the base model through a process of sequential routing. An overview of SEE is shown in Figure 2, and it can be divided into three main steps: (1) Task Reconstruction, (2) Expert Training, and (3) Inference.

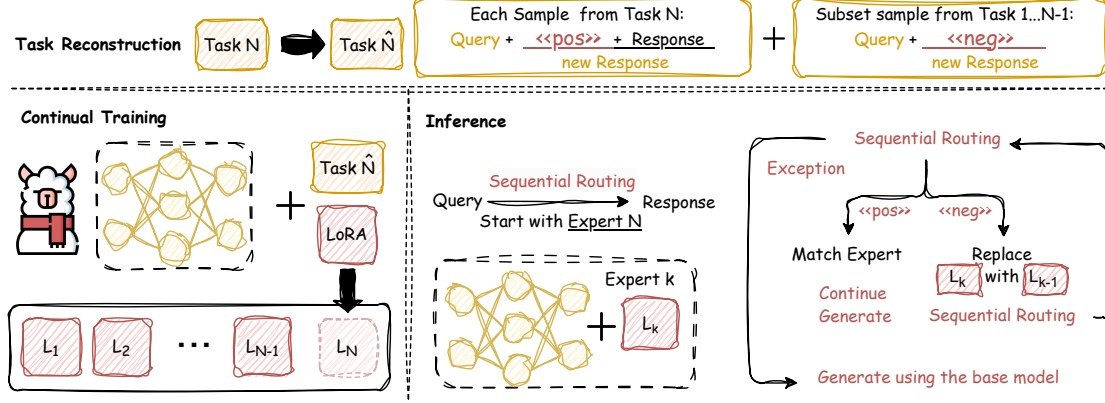


Figure 2: **Overview of the SEE Framework:** The SEE framework operates in three steps when a new task is introduced: (1) Task Reconstruction: Current data are combined with sampled instances from previous tasks to guide expert routing and responses. (2) SFT: A new expert is trained using a new LoRA on the reconstructed task. (3) Inference: All experts are integrated into a MoE system through sequential routing, enabling powerful inferences by leveraging the entire system.

Task Reconstruction At stage i , we reconstruct each task dataset D_i to enable the experts to clarify their responsibilities. Initially, each sample x in D_i can be represented as a tuple (q, r) , where q is the query and r is the corresponding response. We transform each x into \hat{x} , which consists of the tuple (q, o_{pos}, r) , by adding a special positive indicator o_{pos} . Additionally, we sample $\tau\%$ of the data from previous tasks, as in rehearsal-based methods, but ignore the response r and add a negative indicator o_{neg} . The final constructed \hat{D}_i consists of positive samples $\{(q, o_{pos}, r)_j\}_{j=1}^n$ from the current task and a small portion of negative samples $\{(q, o_{neg})_j\}_{j=1}^m$ from previous tasks:

$$\hat{D}_i = \{(q, o_{pos}, r)_j\}_{j=1}^n \cup \{(q, o_{neg})_j\}_{j=1}^m \quad (4)$$

SFT We frame both the decision-making and generation processes of the experts within the paradigm of standard supervised fine-tuning. Multiple LoRA adapters are used as expert weights, enabling the base model \mathcal{M}_0 to specialize into task-specific experts. Starting with the initial model \mathcal{M}_0 , a new LoRA adapter \mathcal{L}_i is introduced whenever a new task T_i arises. The expert for this task is represented as $E_{(\mathcal{M}_0, \mathcal{L}_i)}$. Next, we obtain the dataset \hat{D}_i and maximize the following optimization goal:

$$\max_{\mathcal{L}_i} \frac{1}{n+m} \left(\sum_{i=1}^n E_{(\mathcal{M}_0, \mathcal{L}_i)} \left(r^{(i)}, o_{pos}^{(i)} \mid q^{(i)} \right) + \sum_{i=1}^m E_{(\mathcal{M}_0, \mathcal{L}_i)} \left(o_{neg}^{(i)} \mid q^{(i)} \right) \right) \quad (5)$$

Given a query, the expert first determines whether the question falls under its responsibility and decides whether to provide an answer. By sequentially ensembling all experts, we can handle various tasks through sequential routing.

Inference After continuing fine-tuning on N tasks, we obtain N experts, denoted as $(E_{(\mathcal{M}_0, \mathcal{L}_1)}, \dots, E_{(\mathcal{M}_0, \mathcal{L}_N)})$, where the i -th expert specializes in task T_i . When a question is posed, the expert $E_{(\mathcal{M}_0, \mathcal{L}_N)}$ generates an indicator to decide whether to continue generating a response. If the indicator is o_{neg} , the expert $E_{(\mathcal{M}_0, \mathcal{L}_N)}$ stop generating and routes the query to the expert $E_{(\mathcal{M}_0, \mathcal{L}_{N-1})}$. We iteratively repeat this process until an expert generates o_{pos} . Once this happens, the expert is chosen to generate the final response. If an expert fails to produce an indicator, we route the query to the base model \mathcal{M}_0 and let it to answer the question. This entire routing process is referred to as Sequential Routing. By applying sequential routing, we combine all individual experts into a more powerful model.

4 Experiments

4.1 Setup

The setup of our experiments mainly follows the paper (Huang et al., 2024).

Dataset Our experiments are conducted using the SuperNI (Wang et al., 2022) dataset, a comprehensive and extensive benchmark designed for instruction tuning. To simulate a typical continual learning process, we select a subset of 10 tasks

from SuperNI, representing a range of domains and tasks. Each task includes 2,000 instances for training and 500 instances for evaluation. Additional details can be found in Appendix A or the referenced paper (Huang et al., 2024). For simplicity, the default continual learning order for the {5, 10} SuperNI tasks is as follows: QA → QG → SA → Sum. → Trans. (→ DSG → Expl. → Para. → PE → POS).

Base LLMs Following (Huang et al., 2024), we utilize three base LLMs: Llama-2-7B (Touvron et al., 2023), Llama-2-7B-Chat (Touvron et al., 2023), and Alpaca-7B (Taori et al., 2023) in our experiments. However, we also provide the results of the latest models, including Llama3.1-8B, Qwen2.5-7B, and Mistral-7B-v0.3, in Appendix C.

Baselines We primarily compare the SEE method with the following baselines:

- **Multi-task Learning (MTL)**: The most commonly used baseline, where all tasks are trained simultaneously.
- **Average Single-task Learning (AvgSTL)**: For each task, AvgSTL trains a dedicated model and assesses its performance. The final results are derived by averaging the performance across all tasks.
- **Non-rehearsal**: A naive baseline where the LLM is fine-tuned with only the instruction data $T^{(t)}$ at each stage t .
- **RandSel(τ)** (Scialom et al., 2022): A baseline where we randomly sample $\tau = \{1, 10\}\%$ of the original instruction data for each previous task. As noted in (Scialom et al., 2022), the capabilities of language models can be effectively preserved with $\tau = 1\%$.

Additionally, we include the **KmeansSel** (Huang et al., 2024) and **SSR** (Huang et al., 2024) methods in our experiments.

Evaluation Metrics Following Huang et al. (2024), we evaluate LLM performance on SuperNI tasks using the ROUGE-L metric (Lin, 2004), which aligns closely with human evaluation (Wang et al., 2022). The following ROUGE-L-based metrics are selected for our experiments, where $a_t^{(i)j}$ denotes the ROUGE-L score for task j at stage i :

- **Average ROUGE-L (AR)**: The average LLM performance across all T tasks at stage T :

$$\text{AR} = \frac{1}{T} \sum_{i=1}^T a_i^{(T)} \quad (6)$$

- **Backward Transfer (BWT)**: Assesses the impact of learning new tasks on previous ones by comparing final performance $a_i^{(T)}$ to performance during stage i :

$$\text{BWT} = \frac{1}{T-1} \sum_{i=1}^{T-1} (a_i^{(T)} - a_i^{(i)}) \quad (7)$$

Negative BWT indicates catastrophic forgetting of previously learned knowledge.

On the other hand, we do not use Forward transfer metric (Lopez-Paz and Ranzato, 2017) to evaluate generalization abilities. In our design, this ability is tied to the base model. Thus, We employ a different instruction tuning dataset and evaluate how effectively it can be routed to the base model.

Training Details We set the LoRA rank to 8 and the dropout rate to 0.1. The Adam optimizer is used with an initial learning rate of 2×10^{-4} . The global batch size is 32 in our all experiments. The input and output lengths are configured to 1,024 and 512, respectively. Each LLM is trained for 3 epochs, with evaluation conducted using the final checkpoint.

4.2 Results on 5 SuperNI Tasks

We present the experimental results on five SuperNI tasks in Table 1. The table clearly demonstrates that SEE methods consistently outperform previous rehearsal-based approaches in both AR and BWT metrics across nearly all model configurations and task sequences, even when only 1% of queries from previous tasks are retained. Moreover, the additional parameters enable SEE to not only compete with, but potentially surpass MTL in the AR metric—traditionally considered the upper bound for continual learning in a single LLM. In contrast, AvgSTL exhibits lower performance compared to MTL. This trend aligns with previous research (Caruana, 1997; Ruder, 2017), which suggests that models trained on multiple tasks simultaneously can improve generalization and outperform those trained on tasks in isolation.

While higher BWT scores generally indicate reduced catastrophic forgetting, for the SEE model,

Model	Order 1		Order 2		Order 3		Avg.	
	AR	BWT	AR	BWT	AR	BWT	AR	BWT
<i>Llama-2-7B</i>								
AvgSTL	52.95	-	52.95	-	52.95	-	52.95	-
MTL	53.05	-	53.05	-	53.05	-	53.05	-
Non-rehearsal	17.67	-44.09	15.25	-47.09	24.16	-35.99	19.03	-42.39
RandSel(1%)	51.16	-2.34	49.21	-4.36	48.63	-5.37	49.67	-4.02
KMeansSel(1%)	50.20	-3.12	49.75	-4.11	50.12	-3.61	50.02	-3.61
RandSel(10%)	50.81	-2.32	50.04	-3.31	50.11	-3.42	50.32	-3.02
KMeansSel(10%)	50.44	-3.03	50.61	-2.32	49.89	-3.53	50.31	-2.96
SSR	52.61	-0.23	51.70	-1.22	52.16	-0.93	52.16	-0.79
SEE(1%)	52.86	-0.47	53.19	-0.38	53.37	-0.13	53.14	-0.33
SEE(10%)	53.40	-0.00	52.94	-0.01	53.44	-0.00	53.26	-0.00
<i>Llama-2-7B-chat</i>								
AvgSTL	52.13	-	52.13	-	52.13	-	52.13	-
MTL	52.81	-	52.81	-	52.81	-	52.81	-
Non-rehearsal	23.87	-36.31	30.96	-27.41	42.06	-13.50	32.30	-25.74
RandSel(1%)	51.28	-1.96	49.77	-3.70	49.41	-4.29	50.15	-3.32
KMeansSel(1%)	51.82	-1.25	50.71	-2.44	50.22	-3.42	50.92	-2.37
RandSel(10%)	50.59	-2.57	50.72	-2.45	50.24	-2.87	50.52	-2.63
KMeansSel(10%)	50.81	-2.55	51.39	-1.42	50.22	-2.84	50.81	-2.27
SSR	52.52	-0.23	52.49	-0.35	52.73	0.05	52.58	-0.18
SEE(1%)	52.85	-0.09	52.77	-0.04	53.03	-0.05	52.88	-0.06
SEE(10%)	53.13	-0.00	53.15	-0.00	52.93	-0.01	53.07	-0.00
<i>Alpaca-7B</i>								
AvgSTL	51.78	-	51.78	-	51.78	-	51.78	-
MTL	52.79	-	52.79	-	52.79	-	52.79	-
Non-rehearsal	17.24	-44.21	45.40	-9.03	35.60	-21.45	32.75	-24.90
RandSel(1%)	51.61	-0.93	49.08	-4.68	49.01	-4.85	49.90	-3.49
KMeansSel(1%)	51.37	-1.53	50.53	-2.68	50.15	-3.17	50.68	-2.46
RandSel(10%)	50.91	-1.82	50.88	-2.11	49.98	-3.59	50.59	-2.51
KMeansSel(10%)	50.78	-2.05	51.20	-1.76	49.76	-3.48	50.58	-2.43
SSR	52.52	-0.14	51.74	-1.21	52.33	-0.51	52.20	-0.62
SEE(1%)	52.50	-0.00	52.71	-0.01	52.44	-0.00	52.55	-0.00
SEE(10%)	52.50	-0.00	52.66	-0.00	52.55	-0.00	52.57	-0.00

Table 1: Final results on 5 SuperNI tasks under 3 types of continual learning orders. More details about task orders can be found in Appendix B. The results of the latest models, including Llama3.1-8B, Qwen2.5-7B, and Mistral-7B-v0.3, are presented in Appendix C.

which does not inherently ‘forget’ knowledge, they reflect improved accuracy in routing instances. As shown in Table 1, even with a limited number of instances from previous tasks, the SEE model demonstrates its ability to effectively learn to route queries. This success can be attributed to the distinct patterns often present in the instructions of different tasks. Consequently, the SEE model proves to be particularly well-suited for applications in Continued Fine-tuning.

The value of τ plays a crucial role in both the performance of each expert within the SEE model and the extent of forgetting across the entire system. As illustrated in Table 1, the BWT values for the SEE(1%) model are lower than those for SEE(10%), suggesting that providing a greater number of instances from previous tasks helps SEE

develop a more effective routing strategy. However, while SEE(10%) achieves near ‘zero forgetting,’ its AR values do not show as significant an improvement as those of SEE(1%), and in certain cases, they even experience a slight decline. This indicates that incorporating more instances from previous tasks may negatively impact the performance of experts on their respective tasks.

4.3 Results on 10 SuperNI Tasks

We further investigate the performance of SEE in continuous learning across 10 tasks sequentially, and the results based on Llama2 are presented in Table 2. Despite the increase in tasks, the table shows that SEE outperforms previous rehearsal-based methods in both AR and BWT metrics. Besides, in the AR metric, SEE(1%) proves to be

Model	AR	BWT
<i>Llama-2-7b</i>		
AvgSTL	65.63	-
MTL	64.69	-
Non-rehearsal	17.33	-53.64
RandSel(1%)	60.64	-5.69
KMeansSel(1%)	60.51	-5.39
RandSel(10%)	61.49	-4.03
KMeansSel(10%)	60.93	-3.90
SSR	63.23	-1.56
SEE (1%)	64.64	-0.56
SEE (10%)	66.19	-0.00

Table 2: Final results for Llama-2-7B on 10 SuperNI tasks. The results of the latest models are presented in Appendix C.




Figure 3: The ROUGE-L scores for SEE (10%), MTL, and AvgSTL across the 10 SuperNI tasks are presented in two figures, separated according to the magnitude of the values for clearer comparison.

competitive with MTL, while SEE(10%) delivers a 1.5% improvement over the MTL baseline.

Comparing the BWT values in Table 1 and Table 2, it is evident that increasing the number of tasks leads to a decrease in the BWT value, indicating a greater challenge in retaining knowledge. Despite this, SEE(10%) still achieves “zero forgetting”. It demonstrates the potential of SEE to effectively extend to continuous learning with a larger number of tasks.

Figure 3 presents an analysis of the performance across individual tasks. This figure compares the ROUGE-L scores of SEE(10%), MTL, and AvgSTL. As shown, MTL demonstrates lower performance on the sa and pos tasks, both of which are classification tasks. This may be due to task conflicts (Javaloy and Valera, 2022; Mueller et al., 2022), where learning one task negatively impacts performance on another. Additionally, SEE(10%) outperforms AvgSTL in most of the generation tasks, highlighting that the introduction of routing capabilities enhances the generative performance for these tasks.

5 Analysis

Generalization While the SEE model has demonstrated outstanding performance in tasks during continuous fine-tuning, it is also essential to assess its generalization. In our design, the generalization ability of the SEE depends on two key factors: the average generalization ability of the experts and the accuracy with which the SEE routes out-of-distribution (OOD) queries to the base model. The former refers to the scenario where an OOD query is incorrectly routed to an expert, in which case SEE’s generalization ability is determined by the average generalization ability of the experts. The latter reflects SEE’s ability to correctly identify OOD tasks. To assess this, we evaluate SEE’s average expert performance on MMLU (Hendrycks et al., 2021) and the standard performance of SEE. We compare their performance against that of the base model, MTL, and RandSel. The experiment is based on Llama-2-7B and the results are presented in Figure 4.




Figure 4: The performance of different methods after continuous learning of 10 SuperNI tasks on the MMLU benchmark. SEE(10%)-AE and SEE(1%)-AE represent the average performance of experts in SEE.

As shown in Figure 4, the base model \mathcal{M}_θ achieves the highest accuracy. The average performance of the experts in SEE is similar to that of RandSel. However, It is important to consider the OOD routing accuracy when evaluating the actual performance of SEE. Specifically, 99.72% of instances are routed to the base model in SEE(10%), while this percentage is 99.52% in SEE(1%). As a result, the actual performance of the SEE framework significantly exceeds that of other methods, demonstrating its superior generalization ability.

Perplexity across 10 SuperNI tasks Treating SEE as a unified model, we compute the perplexity for each instance (q, r) . First, the instance is transformed into (q, o_{pos}, r) , where o_{pos} represents a positive indicator. This transformed instance is then routed to the relevant expert within SEE. Then,

the perplexity is computed by the selected expert. We evaluate the perplexity of SEE(10%) across 10 SuperNI tasks and compare the results with those of MTL, which, in other hand, computes the perplexity using the original (q, r) pairs. Figure 5 illustrates that SEE achieves lower perplexity across the 10 SuperNI tasks, highlighting the effectiveness of the SEE framework.

Impact of τ on Routing and Performance To further investigate the impact of τ on SEE, we introduce two additional SEE models with τ values of 5% and 20% for continuous learning on the 10 SuperNI tasks. We then plot the variations in the AR metric and F1 score of routing with τ , and the results are shown in Figure 6.

The figure shows that as τ increases, both the AR metric and routing accuracy improve steadily. In particular, SEE(20%) significantly outperforms both the MTL and AvgSTL methods. Moreover, SEE(20%) achieves near-perfect routing. A more detailed analysis of the routing performance for each task can be found in Appendix D. These results highlight the efficiency of SEE and its potential to be extended to a larger number of tasks.

Essence of Rehearsal The SEE framework utilizes sequential routing to ensemble all experts. Each expert must determine whether a question falls under its responsibility and decide whether to produce a response. While SEE samples negative instances from previous tasks to acquire this capability, it also brings up the question of whether using data from unrelated tasks (called pseudo-negative sampling) could have the same effect. To test this, following the setup of SEE(10%), we replace the negative samples from previous tasks with an equal number of instances from Alpaca-52k, perform the experiments on the 10 SuperNI tasks setup, and compare their performance before and after the changes. Since the performance of




Figure 5: Comparison of the perplexity distribution between SEE and MTL across 10 SuperNI tasks.




Figure 6: Impact of τ on SEE: The left plot shows how the AR metric increases as τ grows, while the right illustrates the improvement in routing accuracy.

SEE on OOD data mainly depends on the accuracy of instance routing to the base model (RB-Acc.), we evaluate the two versions of SEE on MMLU using this metric.

Model	MMLU	SuperNI	
	RB-Acc.	AR	R-F1(%)
SEE(10%)	99.72	66.19	99.80
- w/pseudo	99.57	46.40	47.22

Table 3: Performance comparison of SEE and SEE with pseudo-Negative sampling.

Table 3 compares SEE(10%) and its pseudo-negative sampling variant on the MMLU and 10 SuperNI tasks, using the Llama2-7B model. It shows that, without instances from previous tasks, the F1-score of routing (R-F1) drops significantly, leading to a substantial degradation in the AR metric. However, the performance of MMLU remains unchanged. These findings highlight the importance of rehearsal for SEE when solving tasks in continuous fine-tuning sequences.

Special token indicators While the SEE framework introduces additional tokens as special indicators, we further investigate whether these can be replaced with in-vocabulary tokens. To explore this, we propose two variants of special indicators: 1) Semantic Indicators, and 2) Non-Semantic Indicators. Semantic indicators are tokens that carry specific semantic meaning, whereas non-semantic indicators lack such meaning. In this experiment, we use "Yes" and "No" as semantic indicators, while "«pos»" and "«neg»" serve as non-semantic indicators. We evaluate the performance of these variants on 10 SuperNI tasks. To further explore the impact of indicator type, we categorize the 10 tasks into two groups: Generation and Classification. As before, we set τ to 10% and use Llama2 as the base model. The results are presented in Table 4.

Indicator Type	Generation	Classification
	AR	Acc.(%)
Additional	58.62	98.25
No-Semantic	58.61	97.25
Semantic	58.41	95.05

Table 4: The performance of SEE(10%) based on Llama2-7B across three indicator types in the 10 SuperNI tasks. The tasks are categorized into Generation and Classification.

The table shows that indicator type primarily affects the performance of classification tasks, with minimal impact on generation tasks. In classification, the SEE framework with a semantic indicator yields the lowest accuracy, while SEE with a no-semantic indicator also performs worse than when additional indicators are used. Therefore, introducing new special tokens as indicators to the vocabulary proves more effective for the SEE framework.

The Advantages of Combining Generation and Routing Compared to using encoder-only models as routers, the expert in SEE framework excels in handling OOD scenarios. When a query does not correspond to a specific learning task, the experts typically avoid generating a specialized indicator. In such cases, SEE directs the query to the base model, regarded as the most generalized model in the system. Furthermore, this base model can be substituted with more powerful models, further improving the overall generalization capability of the SEE framework. This approach, however, is not feasible with encoder-only models such as BERT. Moreover, integrating both generation and routing within a single expert not only reduces the parameter count but also streamlines the training pipeline, resulting in a more efficient and compact solution.

Extra Overhead of SEE While the sequential routing of SEE introduces additional overhead, we demonstrate that this cost remains entirely affordable. To quantify this, we calculate the additional latency introduced by the SEE framework compared to standard LLMs. It can be expressed as:

$$ExtraOverhead = \frac{1}{2} \times \frac{1}{1 + \frac{N}{M-1} \times \frac{TPOT}{TTFT}} \quad (8)$$

where **TTFT** denotes the *Time to First Token*, representing the initial prefix processing latency; **TPOT** is the *Time Per Output Token*, the average time required to generate each token; **N** is the total number of output tokens; and **M** is the number of experts.

As shown in previous work (Zhong et al., 2024), when serving an LLM with 13B parameters under a synthetic workload with an input length of 512 and an output length of $N = 64$ on an NVIDIA A100 (80GB), the ratio TPOT/TTFT is approximately 1/10. Assuming the presence of $M = 10$ experts, the additional overhead introduced by SEE is approximately 0.29.

Even with a long input prefix and a short response, this results in only 0.29 additional latency. As the output length increases, this latency overhead can be reduced to a very low level. For instance, when the output length equals the input length (e.g., 512 tokens), the ratio TPOT/TTFT exceeds 1/10, and the additional overhead drops to less than 0.075. In many real-world scenarios, LLMs typically generate outputs significantly longer than the input, the relative overhead introduced by SEE can be expected to further decrease. The detailed calculations can be found in Appendix E.

6 Related Work

Continue learning Existing approaches to continue learning can be divided into regularization-based, architecture-based, and rehearsal-based methods. Regularization-based methods constrain the inner distribution of LLMs to remain close to the original state by adding auxiliary loss or controlling parameter updates (Kirkpatrick et al., 2017; Cha et al., 2021; Huang et al., 2021; Zhang et al., 2022). However, these methods require manually tuned hyperparameters, limiting their applicability. Rehearsal-based methods maintain a history buffer with a subset of previous datasets (de Masson d’Autume et al., 2019; Rolnick et al., 2019) or synthetic instances (Huang et al., 2024), replayed during future training. While effective at mitigating catastrophic forgetting, they cannot fully preserve prior knowledge. Architecture-based methods introduce new parameters for each dataset and learn them independently, potentially avoiding knowledge loss (Xu and Zhu, 2018; Huang et al., 2019; Razdaibiedina et al., 2023). However, managing these additional parameters remains challenging. In our framework, we use rehearsal with queries from previous tasks and organize additional parameters as expert weights, selecting them via sequential routing.

Mixture of experts The basic concept of MoE involves assembling an ensemble of experts to im-

prove performance (Jacobs et al., 1991; Jordan and Jacobs, 1994). Prior work on expert ensembles primarily focuses on fixed domains, tasks, and experts, achieving strong results (Lu et al., 2023; Jiang et al., 2023), but these methods are ill-suited for incremental scenarios or require additional routers (Jang et al., 2023). Furthermore, (Lv et al., 2025) highlights the overlooked issue that separating the router’s decision-making from the experts’ execution may lead to suboptimal expert selection and ineffective learning. Recently, MoE has been applied to transformer architectures (Dai et al., 2024; Jiang et al., 2024), with models proposed to mitigate catastrophic forgetting. Building on GLaM (Du et al., 2022), Lifelong MoE (Chen et al., 2023) expands experts incrementally while freezing previous ones. Some studies replace the standard MoE layer with LoRA-MoE for model editing (Yang et al., 2024; Wang and Li, 2024) or to retain world knowledge in multi-task learning (Dou et al., 2024). However, these methods still struggle with catastrophic forgetting during continual fine-tuning. Our method, based on the MoE concept, sequentially assembles experts in continual fine-tuning, eliminating the need for an additional router and introducing a distributed routing mechanism called sequential routing.

7 Conclusion

In this work, we propose the Sequential Ensemble of Experts (SEE), which is specifically designed to adapt to continuous fine-tuning. SEE resolves the separation between the router’s decision-making and the experts’ execution by integrating routing and response mechanisms within each expert. It employs a distributed routing method called sequential routing, improving the system’s scalability. Our experiments demonstrate SEE’s effectiveness in handling both continual learning tasks and out-of-distribution instances, paving the way for future advancements in distributed model ensembling.

Limitations

While SEE demonstrates remarkable performance in continual fine-tuning, the increasing number of parameters as the task count grows calls for further consideration. Additionally, similar to rehearsal-based approaches, the volume of rehearsal data also expands proportionally with the number of tasks. These challenges are not unique to SEE and have been acknowledged by other methods, underscor-

ing the necessity for continued advancements in this area. Finally, this work does not consider large reasoning models (OpenAI; Guo et al., 2025; Yan et al., 2025), which exhibit distinct behaviors that contribute to improved performance.

Acknowledgement

We extend our gratitude to all the reviewers for their valuable feedback and suggestions, which greatly contributed to enhancing the quality of the paper. This work was supported by the Shanghai Artificial Intelligence Laboratory. We sincerely appreciate their support and resources, which contributed to the successful completion of this research.

References

- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28:41–75.
- Sungmin Cha, Hsiang Hsu, Taebaek Hwang, Flavio P. Calmon, and Taesup Moon. 2021. [CPR: Classifier-Projection Regularization for Continual Learning](#). ArXiv:2006.07326 [cs].
- Wuyang Chen, Yanqi Zhou, Nan Du, Yanping Huang, James Laudon, Zhifeng Chen, and Claire Cu. 2023. [Lifelong Language Pretraining with Distribution-Specialized Experts](#). ArXiv:2305.12281 [cs].
- Damai Dai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y. K. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. 2024. [DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models](#). ArXiv:2401.06066.
- Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. 2019. [Episodic Memory in Lifelong Language Learning](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Wei Shen, Limao Xiong, Yuhao Zhou, Xiao Wang, Zhiheng Xi, Xiaoran Fan, Shiliang Pu, Jiang Zhu, Rui Zheng, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. [LoRAMoE: Alleviating World Knowledge Forgetting in Large Language Models via MoE-Style Plugin](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1932–1945, Bangkok, Thailand. Association for Computational Linguistics.
- Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten Bosma, Zongwei Zhou, Tao Wang, Yu Emma Wang, Kellie Webster, Marie

- Pellat, Kevin Robinson, Kathleen Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V. Le, Yonghui Wu, Zhifeng Chen, and Claire Cui. 2022. [GLaM: Efficient Scaling of Language Models with Mixture-of-Experts](#). ArXiv:2112.06905.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Jianheng Huang, Leyang Cui, Ante Wang, Chengyi Yang, Xinting Liao, Linfeng Song, Junfeng Yao, and Jinsong Su. 2024. [Mitigating Catastrophic Forgetting in Large Language Models with Self-Synthesized Rehearsal](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1416–1428, Bangkok, Thailand. Association for Computational Linguistics.
- Shenyang Huang, Vincent François-Lavet, and Guillaume Rabusseau. 2019. [Neural Architecture Search for Class-incremental Learning](#). ArXiv:1909.06686 [cs].
- Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang, and Diyi Yang. 2021. [Continual Learning for Text Classification with Information Disentanglement Based Regularization](#). ArXiv:2104.05489 [cs].
- Robert Jacobs, Michael Jordan, Steven Nowlan, and Geoffrey Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Computation*, 3:79–87.
- Joel Jang, Seungone Kim, Seonghyeon Ye, Doyoung Kim, Lajanugen Logeswaran, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2023. [Exploring the Benefits of Training Expert Language Models over Instruction Tuning](#). ArXiv:2302.03202 [cs].
- Adrián Javaloy and Isabel Valera. 2022. [Rotograd: Gradient homogenization in multitask learning](#). In *International Conference on Learning Representations*.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024. [Mixture of Experts](#). ArXiv:2401.04088.
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. [LLM-Blender: Ensembling Large Language Models with Pairwise Ranking and Generative Fusion](#). ArXiv:2306.02561 [cs].
- Michael I Jordan and Robert A Jacobs. 1994. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the National Academy of Sciences*, 114(13):3521–3526. Publisher: Proceedings of the National Academy of Sciences.
- Yafu Li, Xuyang Hu, Xiaoye Qu, Linjie Li, and Yu Cheng. 2025. [Test-time preference optimization: On-the-fly alignment via iterative textual feedback](#).
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2023. [Routing to the Expert: Efficient Reward-guided Ensemble of Large Language Models](#). ArXiv:2311.08692 [cs].
- Zhenyi Lu, Chenghao Fan, Wei Wei, Xiaoye Qu, Danyang Chen, and Yu Cheng. 2024. Twin-merging: Dynamic integration of modular expertise in model merging. *Advances in Neural Information Processing Systems*, 37:78905–78935.
- Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2024. [An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-tuning](#). ArXiv:2308.08747.
- Ang Lv, Ruobing Xie, Yining Qian, Songhao Wu, Xingwu Sun, Zhanhui Kang, Di Wang, and Rui Yan. 2025. [Autonomy-of-experts models](#).
- Jisoo Mok, Jaeyoung Do, Sungjin Lee, Tara Taghavi, Seunghak Yu, and Sungroh Yoon. 2023. [Large-scale lifelong learning of in-context instructions and how to tackle it](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12573–12589, Toronto, Canada. Association for Computational Linguistics.
- David Mueller, Nicholas Andrews, and Mark Dredze. 2022. [Do text-to-text multi-task learners suffer from task conflict?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2843–2858, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

OpenAI. Introducing openai o3 and o4-mini. <https://openai.com/index/introducing-o3-and-o4-mini/>.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Pow-

ell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#).

Xiaoye Qu, Daize Dong, Xuyang Hu, Tong Zhu, Weigao Sun, and Yu Cheng. 2024. Llama-moe v2: Exploring sparsity of llama from perspective of mixture-of-experts with post-training. *arXiv preprint arXiv:2411.15708*.

Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao Yan, Dongrui Liu, Ganqu Cui, Daizong Liu, Shuxian Liang, Junxian He, et al. 2025. A survey of efficient reasoning for large reasoning models: Language, multimodality, and beyond. *arXiv preprint arXiv:2503.21614*.

Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Madihan Khabisa, Mike Lewis, and Amjad Almahairi. 2023. [Progressive Prompts: Continual Learning for Language Models](#). ArXiv:2301.12314 [cs].

David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Greg Wayne. 2019. [Experience Replay for Continual Learning](#). ArXiv:1811.11682 [cs].

Sebastian Ruder. 2017. [An overview of multi-task learning in deep neural networks](#).

Thomas Scialom, Tuhin Chakrabarty, and Smaranda Muresan. 2022. [Fine-tuned language models are continual learners](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6107–6122, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyan Wang, Yibin Wang, Zifeng Wang, Sayna

- Ebrahimi, and Hao Wang. 2024. **Continual Learning of Large Language Models: A Comprehensive Survey**. ArXiv:2404.16789.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. **Llama: Open and efficient foundation language models**.
- Renzhi Wang and Piji Li. 2024. **LEMoE: Advanced Mixture of Experts Adaptor for Lifelong Model Editing of Large Language Models**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2551–2575, Miami, Florida, USA. Association for Computational Linguistics.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. **Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan, Thuy-Trang Vu, and Gholamreza Haffari. 2024. **Continual Learning for Large Language Models: A Survey**. ArXiv:2402.01364 [cs].
- Ju Xu and Zhanxing Zhu. 2018. **Reinforced Continual Learning**. ArXiv:1805.12369 [cs].
- Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang. 2025. **Learning to reason under off-policy guidance**.
- Shu Yang, Muhammad Asif Ali, Cheng-Long Wang, Lijie Hu, and Di Wang. 2024. **MoRAL: MoE Augmented LoRA for LLMs’ Lifelong Learning**. ArXiv:2402.11260 [cs].
- Han Zhang, Sheng Zhang, Yang Xiang, Bin Liang, Jinsong Su, Zhongjian Miao, Hui Wang, and Ruifeng Xu. 2022. **CLLE: A Benchmark for Continual Language Learning Evaluation in Multilingual Machine Translation**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 428–443, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zihan Zhang, Meng Fang, Ling Chen, and Mohammad-Reza Namazi-Rad. 2023. **CITB: A Benchmark for Continual Instruction Tuning**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9443–9455, Singapore. Association for Computational Linguistics.
- Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. 2024. **Distserve: Disaggregating prefill and decoding for goodput-optimized large language model serving**.
- Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He, and Yu Cheng. 2024. **Llama-moe: Building mixture-of-experts from llama with continual pre-training**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15913–15923.

A Details of the Selected 10 SuperNI Tasks

Table 8 provides an overview of the 10 SuperNI tasks selected for our primary experiments. To simplify the discussion, we use abbreviations to refer to these tasks throughout the paper. The SuperNI dataset can be accessed at <https://github.com/allenai/natural-instructions> for further reference.

B More Details of Experiments on 5 SuperNI Tasks

We assess the performance of various methods across 5 SuperNI tasks, each with different continuous fine-tuning sequences. Table 5 outlines the 3 specific continual learning orders for the five SuperNI tasks used in our experiments.

Order	Task Sequence
1	QA → QG → SA → Sum. → Trans.
2	Trans. → SA → QA → Sum. → QG
3	Sum. → QG → Trans. → QA → SA

Table 5: Continual learning orders on 5 SuperNI tasks.

C Results of the Latest Models

To demonstrate the efficiency of our method on the latest models, we present the AR scores of the Qwen2.5-7B, Mistral-7B-v0.3, and Llama3.1-8B models. We evaluate them in both 5-task and 10-task settings, with the results shown in Table 6 and Table 7.

Methods	Llama3.1-8B		Mistral-7B-v0.3		Qwen2.5-7B	
	Base	Instruct	Base	Instruct	Base	Instruct
AvgSTL	51.64	52.19	52.38	51.96	51.37	51.43
MTL	51.67	51.92	51.77	52.33	51.38	51.38
SEE(1%)	51.97	52.05	51.48	51.76	51.19	51.65
SEE(10%)	52.14	52.30	51.74	52.31	51.36	51.80

Table 6: AR results of the latest models on five SuperNI tasks in the default order.

Methods	Llama3.1-8B		Mistral-7B-v0.3		Qwen2.5-7B	
	Base	Instruct	Base	Instruct	Base	Instruct
AvgSTL	65.43	65.49	65.77	65.74	64.34	64.67
MTL	65.84	65.54	65.43	65.06	64.62	64.77
SEE(1%)	65.36	65.59	65.80	65.85	64.68	64.94
SEE(10%)	65.68	65.55	65.76	65.86	64.85	64.80

Table 7: AR results of the latest models on ten SuperNI tasks.

D Impact of τ on SEE

We conduct a comprehensive evaluation of the influence of various values of τ on the performance of our SEE framework across 10 distinct SuperNI tasks. The detailed routing F1-scores for each model on every task are presented in Figure 7, offering a thorough analysis of the model’s performance under different configurations.

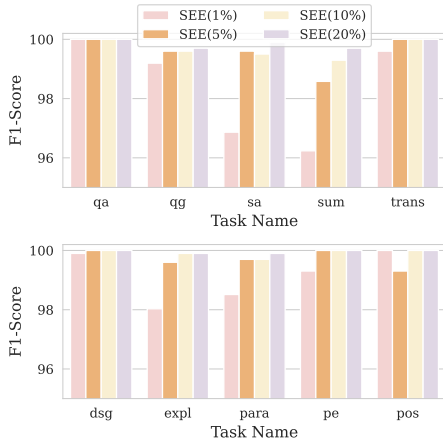


Figure 7: Details of the F1-score for each task on 10 SuperNI tasks.

E Extra Overhead Analysis

Although the increase in the number of tasks introduces additional computational overhead, this extra cost is still affordable. To demonstrate this, we calculate the additional latency of the SEE framework compared to LLMs. The latency of LLMs is given by:

$$Latency_0 = TTFT + TPOT \times N$$

where:

- **TTFT** is the **Time to First Token** (the initial prefix latency),
- **TPOT** is the **Time Per Output Token** (the average time per output token),
- **N** is the number of output tokens.

In the SEE framework, extra latency is introduced due to multiple prefixing operations. Suppose we have **M** experts, labeled as 1, 2, ..., **M**. Assuming the probability of a query hitting any expert or being out-of-distribution (OOD) is uniform, we define **X** as the random variable representing the query routing results.

X	1	2	...	M	Exception
P	$\frac{1}{M+1}$	$\frac{1}{M+1}$	\dots	$\frac{1}{M+1}$	$\frac{1}{M+1}$

To route to *Expert_i*, *i* prefixing operations are required. Additionally, exceptions may occur in each expert, which causes the expected number of routing steps to be $(1 + M)/2$. Therefore, the expected number of routing steps for the SEE framework can be calculated as:

$$E[X] = \sum_{i=1}^M \frac{1}{M+1} \times i + \frac{1}{M+1} \times \frac{1+M}{2}$$

Simplifying this, we get:

$$E[X] = \frac{M}{2} + \frac{1}{2} = \frac{M+1}{2}$$

Thus, the latency for the SEE framework is:

$$Latency_1 = TTFT \times \frac{M+1}{2} + TPOT \times N$$

Now, we can calculate the extra overhead as:

$$ExtraOverhead = \frac{Latency_1 - Latency_0}{Latency_0}$$

Substituting the expressions for $Latency_1$ and $Latency_0$, we get:

$$ExtraOverhead = \frac{TTFT \times \frac{M-1}{2}}{TTFT + TPOT \times N}$$

This simplifies to:

$$ExtraOverhead = \frac{1}{2} \times \frac{1}{1 + \frac{N}{M-1} \times \frac{TPOT}{TTFT}}$$

As shown in (Zhong et al., 2024), when serving an LLM with 13B parameters under a synthetic workload with input length = 512 and output length $N = 64$ on an NVIDIA 80GB A100, the ratio $TPOT/TTFT$ approximates **1/10**. Assuming we have $M = 10$ experts, the extra overhead is:

$$ExtraOverhead = \frac{1}{2} \times \frac{1}{1 + \frac{64}{10-1} \times \frac{1}{10}} = 0.29$$

Even with a long input prefix and a short response, this results in only 0.29 additional latency compared to traditional LLMs.

When the output length increases, this latency can be reduced to a very low level. For example, when the output length is equal to the input length (e.g., 512 tokens), the ratio of $TPOT/TTFT$ exceeds 1/10, and the additional overhead is:

$$ExtraOverhead < \frac{1}{2} \times \frac{1}{1 + \frac{512}{10-1} \times \frac{1}{10}} = 0.075$$

Abbr.	Category	Name	NLU task
QA	Question Answering	task024_cosmosqa_answer_generation	-
QG	Question Generation	task074_squad1.1_question_generation	-
SA	Sentiment Analysis	task1312_amazonreview_polarity_classification	+
Sum.	Summarization	task511_reddit_tifu_long_text_summarization	-
Trans.	Translation	task1219_ted_translation_en_es	-
DSG	Dialogue Sentence Generation	task574_air_dialogue_sentence_generation	-
Expl.	Explanation	task192_hotpotqa_sentence_generation	-
Para.	Paraphrasing	task177_para-nmt_paraphrasing	-
POS	POS Tagging	task346_hybridqa_classification	+
PE	Program Execution	task064_all_elements_except_first_i	-

Table 8: Details of the selected 10 SuperNI tasks.