

How Many Experts Are Enough?

Towards Optimal Semantic Specialization for Mixture-of-Experts

Sumin Park¹, Noseong Park¹

¹Korea Advanced Institute of Science and Technology (KAIST)
{psmiz, noseong}@kaist.ac.kr

Abstract

Finding the optimal configuration of Sparse Mixture-of-Experts (SMoE) that maximizes semantic differentiation among experts is essential for exploiting the full potential of MoE architectures. However, existing SMoE frameworks either heavily rely on hyperparameter tuning or overlook the importance of diversifying semantic roles across experts when adapting the expert pool size. We propose **Mixture-of-Experts for Adaptive Semantic Specialization (MASS)**, a semantic-aware MoE framework for adaptive expert expansion and dynamic routing. MASS introduces two key advancements: (i) a gradient-based semantic drift detector that prompts targeted expert expansion when the existing expert pool lacks capacity to capture the full semantic diversity of the data, and (ii) an integration of adaptive routing strategy that dynamically adjusts expert usage based on token-level routing confidence mass. We first demonstrate that MASS reliably converges to the point of optimal balance between cost-performance trade-off with notably improved semantic specialization in a highly controlled synthetic setup. Further empirical results on real-world datasets across language and vision domains show that MASS consistently outperforms a range of strong MoE baselines, demonstrating its domain robustness and enhanced expert specialization.

Introduction

The Sparse Mixture-of-Experts (SMoE) approach has recently emerged as an effective mechanism that substantially improves model capacity with minimal computational overhead (Shazeer et al. 2017; Lepikhin et al. 2020; Fedus, Zoph, and Shazeer 2022). Incorporating SMoE structure into deep neural networks, particularly large-scale Transformer-based architectures, has significantly boosted performance across diverse tasks and modalities (Li et al. 2023; Qiu, Huang, and Fu 2024). A fundamental advantage of SMoE is its inherent scalability and flexibility, achieved by selectively routing input tokens to a sparse subset of specialized subnetworks (experts). This sparsity-based selective routing mechanism is especially appealing for large-scale neural architectures, such as Large Language Models (LLMs), which commonly struggle with substantial computational demands, handling massive model sizes and diverse datasets (Zhu et al. 2024; Yang et al. 2024; Jiang et al. 2024; Dai et al. 2024).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

However, a crucial issue when employing SMoE lies in determining the optimal size of the expert pool, commonly denoted as K , and the number of active experts per token, referred to as top- k selection (Shazeer et al. 2017; Fedus, Zoph, and Shazeer 2022; Zhou et al. 2022, 2024). Current methods remain largely heuristic and reliant on resource-intensive hyperparameter searches.

Recently, **DynMoE** (Guo et al. 2025) has been proposed as a promising solution that adaptively adjusts the expert set size based on token-expert routing coverage. However, DynMoE relies solely on token-level expert activation pattern and does not explicitly assess whether the current expert pool becomes functionally saturated to capture semantic variations in input data. This overlooks a critical aspect of MoE modeling: fine-grained semantic specialization across experts, which is essential for minimizing functional redundancy in experts and thus, promoting complementary mixture of subnetworks.

To address this limitation, we propose **MASS (Mixture-of-Experts for Adaptive Semantic Specialization)**, a semantic-aware MoE framework that adaptively expands the expert pool based on gradient-based indicators of functional insufficiency. MASS identifies when an expert is overloaded beyond its semantic capacity by combining two signals: (1) a sustained increase in gradient magnitude detected by probabilistic *Change Point Detection (CPD)* approach, and (2) a semantic drift indicated by misalignment between the current gradient update and the expert’s historical weight representation. If both conditions are met, a new expert is instantiated to specialize on the emerging role. Figure 1 provides a visual comparison of expert expansion mechanism between DynMoE and MASS. While DynMoE heuristically expands experts based on unassigned tokens, MASS uses a principled two-stage mechanism that directly targets semantic insufficiency, enabling finer-grained and functionally grounded expert specialization.

We complement this expansion mechanism with an integration of a probabilistic routing strategy inspired by Top- p routing (Huang et al. 2024). Instead of selecting a fixed number of experts (as in top- k gating), the gating module activates the minimal subset of experts whose cumulative routing probability exceeds a confidence threshold p . Under this strategy, confident tokens naturally use fewer experts, reducing computation, while uncertain tokens are routed to a larger set of experts for richer processing. This yields a

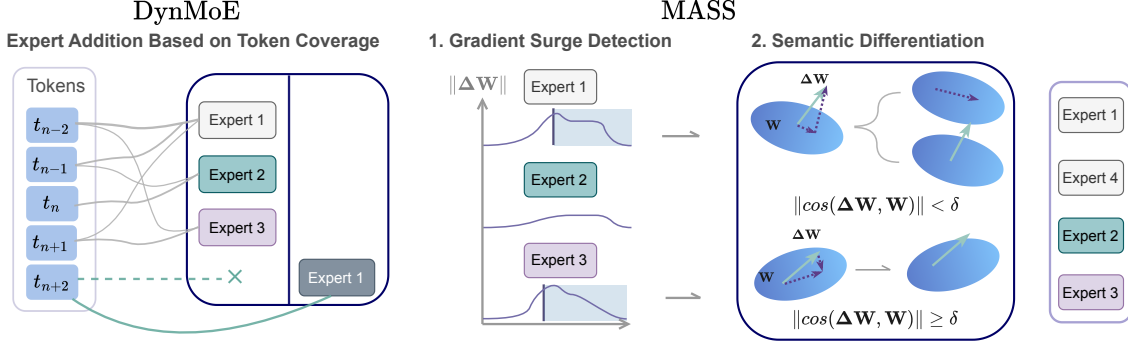


Figure 1: Comparison of Adaptive Expert Expansion Mechanisms. Left (DynMoE): Experts are added when tokens fail to activate any existing expert, indicating insufficient token coverage. This strategy relies on token-level activation statistics, without explicitly assessing semantic adequacy. Right (MASS): Expert expansion is triggered through a two-step process: (1) MASS statistically detects upward shifts in gradient norm over training steps to identify experts under high update pressure, and (2) it evaluates semantic drift by measuring misalignment between the gradient update $\Delta \mathbf{W}$ and the expert’s current representation \mathbf{W} . If the alignment $\|\cos(\Delta \mathbf{W}, \mathbf{W})\|$ falls below a threshold δ , a new expert is instantiated to specialize on the diverging semantics.

balanced and adaptive expert allocation that complements MASS’s principled expansion criterion.

To validate the effectiveness of our approach, we first rigorously tested MASS in a controlled synthetic setting with language-like data modeled by multinomial Hidden Markov Models (HMMs), following the same settings in GINC dataset generation (Xie et al. 2022). This setup allows us to examine whether MASS can adaptively converge to an optimal balance on the trade-off between the performance and computational cost, i.e., the number of experts. Beyond synthetic validation, we further evaluate MASS on real-world datasets spanning both language and vision tasks in comparison to strong MoE baselines, including DynMoE.

Contributions Our main contributions are as follows:

- We propose MASS, a semantic-aware adaptive MoE expansion framework that expands the expert pool using gradient-based semantic drift detection.
- We adopt a confidence-driven Top- p routing, enabling sufficient and balanced per-token expert attention.
- We empirically validate MASS on both synthetic and real-world tasks, demonstrating its ability to determine the optimal expert pool size with consistent performance gains over diverse baselines across multiple domains.

Background

Sparse Mixture of Experts Recent works have actively employed Sparse Mixture-of-Experts (SMoE) structure in modern deep learning architectures, demonstrating its wide applicability across multiple domains (Li et al. 2023; Qiu, Huang, and Fu 2024; Puigcerver et al. 2024). SMoE layers are commonly integrated into Transformer-based architectures by replacing the standard feed-forward network with a sparse, gated mixture of expert sub-networks. A gating module routes each token to a small subset of experts, typically

via top- k token-wise selection, enabling increased model capacity with reduced computational cost (Fedus, Zoph, and Shazeer 2022; Dai et al. 2024).

Beyond foundational design, SMoE has also been explored in various fine-tuning scenarios (Zadouri et al. 2023; Li et al. 2024; Wu et al. 2024). Among them, MoEfication (Zhang et al. 2022) restructures pre-trained feed-forward layers into sparse expert modules for efficient inference and GMoE (Li et al. 2023), initialized from a pretrained vision transformers, enhance domain generalization in vision tasks.

Multinomial HMM The Hidden Markov Model (HMM) (Baum and Petrie 1966) is a classical generative model designed for sequential data, where a sequence of observable outputs is generated by an underlying sequence of latent variables that evolve according to a first-order Markov process. In this framework, each latent state generates an observation independently based on a state-dependent emission distribution. Let $\{h_t\}_{t=1}^T$ represent the sequence of latent states and $\{o_t\}_{t=1}^T$ the corresponding observed outputs. The model is governed by three components: an initial state distribution $\pi = p(h_1)$, a state transition function $\mathcal{T} = p(h_t | h_{t-1})$, and an emission distribution $\mathcal{E} = p(o_t | h_t)$. The joint probability of the full sequence is expressed as:

$$p(o_{1:T}, h_{1:T} | \theta) = \pi(h_1) \prod_{t=2}^T \mathcal{T}(h_t | h_{t-1}) \prod_{t=1}^T \mathcal{E}(o_t | h_t),$$

where $\theta = \{\pi, \mathcal{T}, \mathcal{E}\}$ denotes the full set of model parameters. A categorical HMM is a variant where each hidden state emits observations drawn from a categorical distribution over a discrete set of outcomes. Each latent state $h_t \in \{1, \dots, N\}$ is tied to a categorical distribution over a finite observation set $\mathcal{O} = \{1, \dots, M\}$. The emission probabilities are parameterized by a matrix $\mathcal{E} \in \mathbb{R}^{N \times M}$, where each row represents the probability of producing each outcome at a given hidden

state. Along with a transition matrix $\mathcal{H} \in \mathbb{R}^{N \times N}$ and an initial state distribution $\pi \in \mathbb{R}^N$, this model effectively captures sequential patterns in symbolic data streams.

Probabilistic Change Point Detection Change Point Detection (CPD) refers to the task of identifying abrupt distributional shifts in time series. A classical approach relies on the Cumulative Sum (CUSUM) algorithm, which tracks cumulative sum of deviations from a reference distribution and triggers a change point when this drift exceeds a fixed threshold. Probabilistic CPD builds on this by recasting the drift detection as a statistical test. Instead of fixed thresholds, it quantifies how likely the observed cumulative deviation is under the null hypothesis of no distributional change.

Formally, given a time series $\{X_t\}_{t=1}^T$, we standardize observations as $Z_t = \frac{X_t - \hat{\mu}_X}{\hat{\sigma}_X}$, where $\hat{\mu}_X$ and $\hat{\sigma}_X$ are estimates of the mean and standard deviation of the series. The cumulative sum $S_T = \sum_{t=1}^T Z_t$ approximates $\mathcal{N}(0, T)$ under the assumption that the Z_t is i.i.d. and satisfies the Central Limit Theorem (CLT) conditions, yielding a normalized statistic

$$\tilde{S}_T = S_T / \sqrt{T} \sim \mathcal{N}(0, 1).$$

Given a test statistic \tilde{s}_T of the random variable \tilde{S}_T , the probability of observing a value less than or equal to \tilde{s}_T is as:

$$\Phi(\tilde{s}_T) \approx P(\tilde{S}_T \leq \tilde{s}_T),$$

where $\Phi(\cdot)$ denotes the cumulative distribution function (c.d.f.) of the standard normal distribution. The right-tail probability $1 - \Phi(\tilde{s}_T)$ serves as a p-value in hypothesis testing. A change point is flagged when it falls below a predefined threshold, enabling robust and underlying distribution-free detection of drift.

Method

Overview of MASS

In this section, we provide full algorithmic details of MASS.

- Adaptive Expert Expansion.** During the early phase of training, MASS dynamically adjusts the expert pool size by detecting semantic drift through gradient-based monitoring.
- Standard Training.** The second phase fixes the expert set derived from the first phase and performs conventional supervised training.

Throughout the entire training process, expert routing is governed by a routing-mass-based adaptive Top- p strategy.

MoE Architecture

We define the structure of the MoE layer shared across all subsequent experiments. The layer consists of a token-wise gating module and a pool of K experts, where K denotes the current expert pool size at a given point in training. For each input token representation $\mathbf{x} \in \mathbb{R}^d$, a subset of experts $\mathcal{N}(\mathbf{x}) \subseteq \{1, \dots, K\}$ is selected by the router, and the output is computed as:

$$\mathbf{y} = \sum_{k \in \mathcal{N}(\mathbf{x})} r_k(\mathbf{x}) \cdot e_k(\mathbf{x}).$$

Here, $e_k(\cdot)$ denotes the k -th expert function, and $r_k(\mathbf{x}) = [\mathbf{r}(\mathbf{x})]_k$ is the routing score assigned to expert k . The full routing distribution is computed as $\mathbf{r}(\mathbf{x}) = \text{Softmax}(\mathbf{x}^\top \mathbf{W}_r)$, where $\mathbf{W}_r \in \mathbb{R}^{d \times K}$ is the gating weight matrix.

Routing-Mass Based Expert Selection To determine $\mathcal{N}(\mathbf{x})$ per token, the model adopts the Top- p strategy. Let $r^{(1)} \geq r^{(2)} \geq \dots \geq r^{(K)}$ denote the sorted routing scores in descending order, and let \mathcal{I}_j be the original expert index corresponding to the j -th largest score $r^{(j)} = r_{\mathcal{I}_j}(\mathbf{x})$. The number of active experts k^* is chosen as the smallest index such that:

$$\sum_{j=1}^{k^*} r^{(j)} \geq p, \quad \mathcal{N}(\mathbf{x}) = \{\mathcal{I}_1, \dots, \mathcal{I}_{k^*}\},$$

where $p \in (0, 1)$ is a hyperparameter controlling the routing confidence threshold. Instead of assigning a hard number of experts to every token regardless of certainty, this confidence-aware strategy dynamically allocates expert capacity according to the model's routing uncertainty.

Adaptive Expert Expansion

Training begins with an initial number of experts K_{init} and expands the expert pool up to a maximum expert limit K_{max} during the first 10% of total training steps.

Gradient Monitoring via Probabilistic CPD For each expert e_k parameterized by θ_k , the L2 norm of the task gradient is tracked at each training step t as

$$g_t^{(k)} := \|\nabla_{\theta_k} \mathcal{L}_t\|_2,$$

where \mathcal{L}_t denotes the loss at step t . To ensure stability, each expert undergoes a warmup phase of T_{warmup} steps before any CPD statistics are computed. After warmup, a sliding window of size ω maintains the ω -most recent gradient norms:

$$\mathcal{G}_t^{(k)} := \{g_{t-\omega+1}^{(k)}, \dots, g_t^{(k)}\}.$$

Within this window, the sample mean and unbiased standard deviation are computed as

$$\mu_t^{(k)} := \frac{1}{\omega} \sum_{g \in \mathcal{G}_t^{(k)}} g, \quad \sigma_t^{(k)} := \sqrt{\frac{1}{\omega-1} \sum_{g \in \mathcal{G}_t^{(k)}} (g - \mu_t^{(k)})^2}.$$

The cumulative z-score within window is then computed as:

$$s_t^{(k)} := \sum_{i=t-\omega+1}^t z_i^{(k)}, \quad z_t^{(k)} := \frac{g_t^{(k)} - \mu_t^{(k)}}{\sigma_t^{(k)}}.$$

yielding the normalized test statistic: $\tilde{s}_t^{(k)} := s_t^{(k)} / \sqrt{\omega}$.

Assuming normality under the null hypothesis, i.e., $\tilde{s}_t^{(k)} \sim \mathcal{N}(0, 1)$, the corresponding right-tailed p-value is given by:

$$p_t^{(k)} := 1 - \Phi(\tilde{s}_t^{(k)}),$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. To decide whether to reject the null hypothesis of no distributional shift, we compare the p-value against a predefined significance level $\alpha \in (0, 1)$, and reject when $p_t^{(k)} \leq \alpha$. This right-tailed test identifies statistically significant upward shifts in the gradient magnitude distribution, indicating that expert e_k may require substantial adaptation in its representation.

Semantic Alignment Test For a flagged expert e_k by prior gradient monitoring, semantic misalignment is diagnosed by evaluating the cosine similarity between the current gradient matrix $\nabla^{(k)}$ and the expert’s weight matrix $\mathbf{W}_e^{(k)}$:

$$\cos(\nabla^{(k)}, \mathbf{W}_e^{(k)}) := \frac{\langle \nabla^{(k)}, \mathbf{W}_e^{(k)} \rangle}{\|\nabla^{(k)}\|_2 \cdot \|\mathbf{W}_e^{(k)}\|_2},$$

where both matrices are flattened into vectors before computing similarity. Here, $\mathbf{W}_e^{(k)}$ refers to the projection weight of the first linear layer in the expert function. If the similarity satisfies $\|\cos(\nabla^{(k)}, \mathbf{W}_e^{(k)})\| < \delta$, with $\delta = 0.001$ as a fixed threshold indicating effective orthogonality, the expert is considered to have undergone semantic drift, prompting the addition of a new expert for diverging semantic role.

Expert Duplication with Gradient Decomposition Upon confirming a semantic drift by two-step detection, the expert e_k is duplicated to form a new expert, denoted as e'_k . e'_k inherits the full gradient update, while the original expert e_k receives only the component aligned with its current weight representation.

An analogous update is applied to the gating network. Let $\mathbf{w}_k = [\mathbf{W}_r]_{:,k}$ denote the gating column vector corresponding to expert e_k . To support the new expert e'_k , a corresponding gating vector \mathbf{w}'_k is initialized by cloning: $\mathbf{w}'_k \leftarrow \mathbf{w}_k$. Then, the original gating vector receives only the aligned gradient, while the new one takes the full gradient.

This selective update helps preserve the original semantics of e_k while allowing e'_k to explore a divergent role. To mitigate redundancy afterwards, a regularization term is applied to encourage divergence in their routing behavior.

Redundancy Regularization Following expert duplication, experts and their gating vectors are initialized to be nearly identical, which risks the functional collapse of the duplicated experts, presenting redundant behavior. To mitigate this, we introduce a regularization term that penalizes high alignment between the gating vectors of duplicated expert pairs. Formally, the regularization loss is defined as:

$$\mathcal{L}_{\text{red}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} (\cos(\mathbf{w}_i, \mathbf{w}_j))^2,$$

where \mathcal{P} is the set of all duplicated expert pairs, and \mathbf{w}_i denotes the gating vector of expert i . By penalizing similarity in gating vectors, we encourage input tokens to be routed differently between the pair, promoting eventual divergence in their learned semantic roles. Note that this regularization is only applied during the expansion phase to avoid severe distortion of natural evolution of the routing behavior.

Expansion Stopping Criteria Adaptive expansion terminates under either of the following conditions:

- The number of experts reaches the maximum allowable value: $K = K_{\text{max}}$.
- The previously added expert fails to improve the loss, as evaluated by negative log-likelihood (NLL) comparison at the following expansion step. Expansion is terminated once no improvement is observed over γ times. See Algorithm 1 for full details.

Algorithm 1: NLL Comparison for Expansion Stopping

Require: Patience counter $\gamma > 0$, expert set \mathcal{K} , list of duplicated experts \mathcal{P}

- 1: **if** the $(m+1)$ -th expert has just been added **and** $|\mathcal{P}| \geq 2$ **then**
- 2: Let e_m be the expert added in the m -th expansion step
- 3: Temporarily disable e_m and compute the masked loss $\mathcal{L}_{\text{mask}}$
- 4: Restore e_m and compute original loss $\mathcal{L}_{\text{orig}}$
- 5: Compute loss gap: $\Delta\mathcal{L} := \mathcal{L}_{\text{mask}} - \mathcal{L}_{\text{orig}}$
- 6: **if** $\Delta\mathcal{L} \leq 0$ **then**
- 7: Decrement patience counter: $\gamma \leftarrow \gamma - 1$
- 8: Remove e_m from the expert pool
- 9: **if** $\gamma = 0$ **then**
- 10: Stop expansion
- 11: **end if**
- 12: **else**
- 13: Retain e_m
- 14: **end if**
- 15: **end if**

Once expansion stops satisfying due to either criterion, training during expansion phase continues as standard training with active redundancy regularization. Restricting the expansion phase to the initial 10% of training ensures training stability and minimizes the computational overhead associated with expansion algorithms, allowing the remainder of training to proceed with a static expert configuration. The full algorithm of MASS is summarized in Appendix A.

Standard Training

Following the adaptive expansion phase, the expert pool has converged to an optimal size with enough semantic coverage for training data. For the remaining 90% of steps, the model is trained upon a classical gradient-guided training scheme.

Results on Synthetic-Data

We firstly evaluate MASS in a controlled synthetic setting based on a structured language modeling scenario. The use of a synthetic setup allows us to isolate the effects of adaptive expert expansion and dynamic routing, without confounding factors present in large-scale network and datasets.

Data Generation

Here we describe full data generation process using multinomial Hidden Markov Models (HMMs), following the GINC dataset generation procedure (Xie et al. 2022). Each hidden state is a tuple $h_t = (v_t, s_t)$, where $v_t \in \mathcal{V}$ is an entity index and $s_t \in \mathcal{S}$ is a property type. A global memory matrix $M \in \mathcal{O}^{|\mathcal{V}| \times |\mathcal{S}|}$ deterministically maps each entity-property pair to a single token in the output vocabulary \mathcal{O} . The transition dynamics are defined as:

$$s_{t+1} \sim P(s_{t+1} | s_t; \theta), \quad \theta \sim p(\theta),$$

$$v_{t+1} \sim 0.9 \cdot \delta(v_{t+1} = v_t) + 0.1 \cdot I(v_{t+1} | v_t),$$

where the model parameter $\theta \in \Theta$ serves as a concept parameter governing distinct property transitions, each sampled as

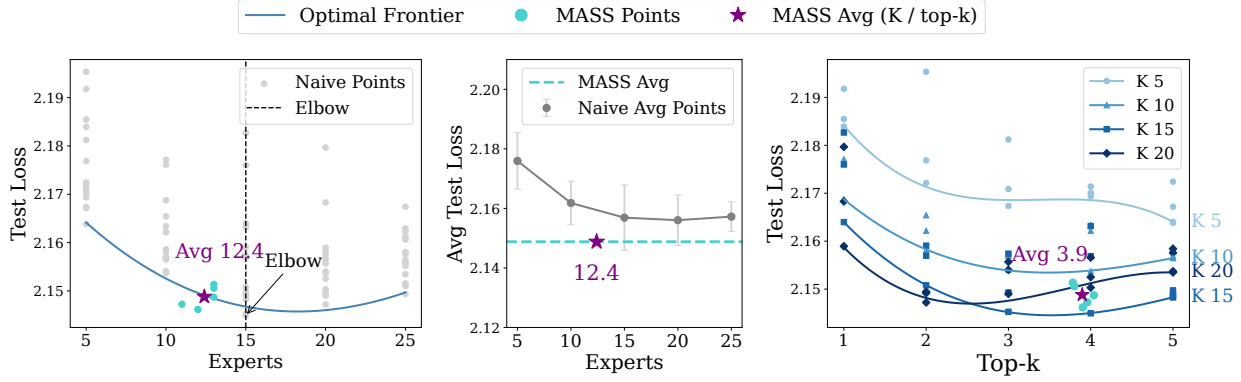


Figure 2: Synthetic Results for Identifying Optimal MoE Configurations with MASS. Left: Test loss vs. expert count for Naive and MASS. Gray dots represent individual runs of Naive with mixed top- k across different K , and the solid blue curve connects the best performing run per K , forming the optimal frontier. Elbow represented by a gray vertical dashed line denotes the point of the most cost-effective K . Green dots show five MASS runs, and the purple star shows their average (Avg 12.4 experts), which aligns closely with the elbow. Center: Average test loss comparison between Naive with varying expert counts (gray bars) and MASS average (dashed green line), showing consistently lower loss for MASS (Avg 2.15). Right: Test loss vs. top- k for fixed $K \in \{5, 10, 15, 20\}$. Colored curves represent per- K optimal frontier for Naive. The 5 MASS samples from the left plot are shared here as well (green circles), with the average active experts marked by the purple star (Avg 3.9).

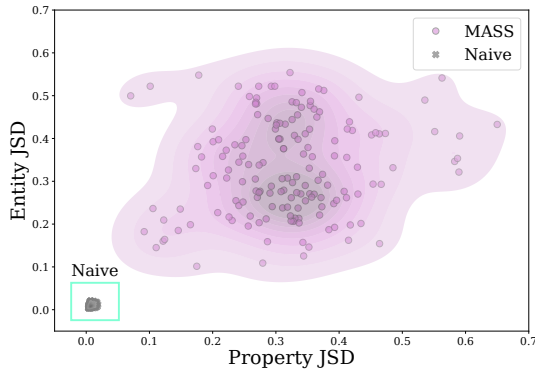


Figure 3: Expert Specialization via Routing Divergence by Different Semantics. Contour plot comparing pairwise JSD of average routing distributions grouped by semantic labels (x =entity and y =property) defined in the synthetic data.

a distinct HMM from a uniform mixture of HMMs over a family Θ of 5 concepts. The $\delta(v_{t+1} = v_t)$ denotes the Dirac delta function, assigning probability mass 1 to $v_{t+1} = v_t$ (i.e., entity persistence). This defines a sticky Markov process that maintains the same entity with high probability (0.9), and transitions to a new entity with a small probability (0.1) according to uniform $I(v_{t+1} | v_t)$. At each time step, the emitted token is deterministic as $o_t = M[v_t, s_t]$ with $p(o_t | (v_t, s_t)) = 1$. Let T denote the total number of tokens in the generated sequence, n the number of training examples, and t' the length of each input sub-sequence x_i . Each example (x_i, y_i) consists of an input $x_i = [o_{t_i}, o_{t_i+1}, \dots, o_{t_i+t'-1}]$ and a target token $y_i = o_{t_i+t'}$. Therefore, each training example consumes $t' + 1$ tokens and total sequence length $T = n \cdot (t' + 1) + t'$. The final training sequence is represented as $[x_1, y_1, o_{\text{delim}}, \dots, x_n, y_n, o_{\text{delim}}, x_{\text{test}}]$ where x_{test}

is used to evaluate the model’s prediction of $y_{\text{test}} = o_T$.

Model

For the synthetic experiment, we consider a single-layer Transformer architecture augmented with a Mixture-of-Experts (MoE) module. The model includes a token embedding layer, sinusoidal positional encoding, and a Transformer encoder block. The MoE module replaces the standard feed-forward network (FFN), where each expert is implemented as a single linear transformation followed by a SiLU activation, mapping from \mathbb{R}^d to $\mathbb{R}^{d'}$. A gating network, implemented as a simple linear layer followed by a softmax activation, computes the routing score distribution over experts. The MoE output is computed as a weighted combination of routed expert outputs and passed through a shared decoder to produce token predictions. Models are trained end-to-end using cross-entropy loss.

We compare two models in this controlled setting:

- **Naive MoE.** This baseline adopts a fixed number of experts, along with standard top- k routing with fixed k .
- **MASS.** Sharing the same network architecture as Naive, it grows experts adaptively by expansion algorithm of MASS and uses Top- p gating.

This setup allows us to assess: (1) whether the MASS expansion algorithm successfully identifies an optimal expert pool size that balances semantic capacity and computational cost, and (2) whether MASS—via improved semantic differentiation and dynamic expert allocation—yields better performance compared to baseline.

Results We evaluate MASS against static MoE baselines by varying the number of experts $K \in \{5, 10, 15, 20, 25\}$ and top- $k \in \{1, 2, 3, 4, 5\}$ in Naive configurations, each repeated by three times. For each K , we select the best-performing

Naive run with lowest test loss and connect these points to form the optimal frontier, a curve that represents the minimal test loss achievable per model capacity. While this is not a strict Pareto front in the multi-objective optimization sense, it serves as an empirical optimality boundary for assessing the performance–cost trade-off. To ensure a fair comparison, we evaluate MASS on five independent runs.

Figure 2 summarizes the performance of MASS against these static MoE baselines. As shown in the empirical results in left plot that shows test loss as a function of expert counts, MASS effectively discovers cost-efficient expert configurations. Operating under dynamic expert expansion and routing, it consistently matches or surpasses the best-performing static MoE baselines. Importantly, MASS converges to expert counts near or below the empirical elbow ($K = 15$) with an average of 12.4 experts, demonstrating its ability to identify an optimal balance on cost-performance trade-off. The right plot shows test loss versus top- k for fixed $K \in \{5, 10, 15, 20\}$. MASS with Top- p routing achieves comparable or better results, with 3.9 active experts per token in average. These results highlight that MASS not only discovers an effective expert pool size with enhanced semantic specialization, but also dynamically adapts expert allocation to equally distribute the routing confidence across all tokens.

We further analyze expert specialization behavior using the same set of sample runs for both Naive and MASS. For this, we analyze the diversity of expert routing conditioned on two semantic groups, entity $\in \mathcal{V}$ and property $\in \mathcal{S}$ of tokens in the synthetic data. Specifically, we compute the pairwise Jensen–Shannon Divergence (JSD) between the average expert routing distributions per entity (property) on each of the fully trained models, Naive and MASS. Formally, for a entity pair $(v_i, v_j)_{i,j \in \mathcal{V}}$, let $P_{v_i}, P_{v_j} \in \mathbb{R}^K$ denote the average routing distributions over K experts for v_i and v_j , respectively. The JSD is then computed as:

$$\text{JSD}(P_{v_i} \| P_{v_j}) = \frac{1}{2} \text{KL}(P_{v_i} \| M) + \frac{1}{2} \text{KL}(P_{v_j} \| M),$$

where $M = \frac{1}{2}(P_{v_i} + P_{v_j})$ and KL denotes the Kullback–Leibler divergence. This provides a symmetric and bounded measure in $[0, 1]$ of expert-entity (property) specialization. A larger JSD value implies that the routing distributions for varying input tokens are more diverging—the model assigns different subsets of experts to different semantics. In contrast, low JSD indicates that tokens having different semantics are routed similarly, a lack of specialization.

Figure 3 presents a contour plot of entity-level JSD (x-axis) vs. property-level JSD (y-axis), comparing MASS (purple) and Naive MoE (gray). As shown, MASS exhibits significantly higher JSD in both dimensions, forming a wide spread contour, whereas Naive MoE shows minimal variation clustered near zero in both axes. This suggests that MASS enables more distinct expert usage across different semantic groups, encouraging functional differentiation among experts. In contrast, Naive MoE uses experts uniformly, failing to separate semantic roles across experts.

To verify the sanity of gradient shift detecting CPD algorithm used in MASS, we further provide a qualitative view of MASS’s internal dynamics by visualizing the gradient norm

traces $\|\nabla_{\theta_k} \mathcal{L}_t\|_2$ for experts near their expansion events in Appendix C.

Results on Real-Data

To assess the effectiveness of MASS in realistic settings, we evaluate its performance across two domains, natural language understanding and vision domain generalization. In both domains, we compare MASS against established MoE baselines, including fixed-size sparse MoEs, GMoE, and recent dynamic variant DynMoE.

Language Task For language task, we follow the MoEification setup (Zhang et al. 2022) and fine-tune BERT-large (Devlin et al. 2019) models on the GLUE benchmark (Wang et al. 2019). Specifically, we evaluate on CoLA (Warstadt, Singh, and Bowman 2019), QNLI (Wang et al. 2019), RTE (Bentivogli et al. 2009), MNLI (Xu et al. 2020), and MRPC (Dolan and Brockett 2005), five subtasks of GLUE. We compare the performance of MASS with other baselines, MoE with cosine routers (Li et al. 2023) and DynMoE (Guo et al. 2025), with all results averaged over three random seeds.

As shown in Figure 4, MASS consistently matches or outperforms other MoE variants across all five GLUE tasks, achieving the highest average top-1 accuracy among the compared methods. Importantly, MASS adaptively expands its expert pool to a diverse range of sizes, $K \in [9.5, 12.7]$, depending on task complexity, while DynMoE consistently converges to a fixed size of $K = 9.0$, showing limited responsiveness to task-specific demands.

Moreover, MASS exhibits a substantially lower average activated k usage (2.6–3.2), activating only around 25%–30% of its expert pool per token. In contrast, DynMoE utilizes 70%–90% of its expert capacity with 6.5–8.0 active experts out of 9, reflecting denser and less specialized routing. This clear contrast in expert utilization pattern highlights MASS’s capacity to promote sparse, semantically differentiated expert activation, leading to more efficient and targeted expert usage across input tokens.

Vision Task For vision experiments, we follow the same MoE integration strategy in GMoE (Li et al. 2023), applying it to a pre-trained ViT-S/16 model (Dosovitskiy et al. 2021). We evaluate performance on the DomainBed benchmark (Gulrajani and Lopez-Paz 2020), a widely used framework for domain generalization. DomainBed is constructed to assess a model’s ability to generalize across multiple visual domains by training on a subset of domains and testing on an unseen target domain. We conduct experiments on four diverse datasets—PACS (Li et al. 2017), VLCS (Albuquerque et al. 2021), OfficeHome (Venkateswara et al. 2017), and TerraIncognita (Beery, van Horn, and Perona 2018)—each containing multiple domains with different environments. In our setting, we adopt a leave-one-domain-out protocol: for each dataset, models are trained on all but one domain and evaluated on the held-out target. Model selection follows the standard train-validation scheme, where the checkpoint with the highest validation accuracy is used for evaluation.

We compare MASS to GMoE with fixed expert sizes ($K = 4, 6, 8$) and DynMoE. Both MASS and DynMoE

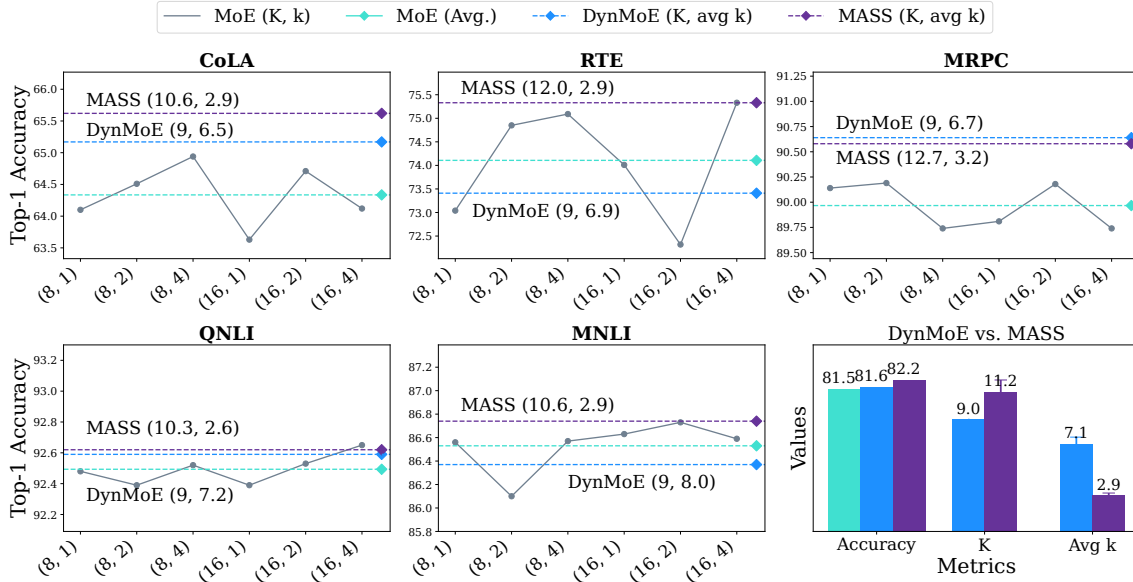


Figure 4: Performance Comparison on GLUE benchmarks. The top five plots show top-1 accuracy for each GLUE task comparing MASS, DynMoE, and static MoE configurations. Gray dots connected by gray line represent the performance of fixed MoE baselines, with horizontal lines marking average performance of DynMoE (blue) and MASS (purple). Numbers indicate the average expert count K and activated k across three random seeds. The bottom-right figure provides summary comparison for top-1 accuracy, K , and k in average across models.

Algorithms	PACS	VLCS	OfficeHome	TerraIncognita	Average
GMoE ($K = 4$) ³	88.2	79.8	73.5	47.8	72.3
GMoE ($K = 6$) ²	88.1	80.2	74.2	48.5	72.8
GMoE ($K = 8$) ³	88.2	80.0	74.2	47.2	72.4
<i>Adaptive Experts</i>					
DynMoE (with Gshard Loss) ¹	88.4	79.4	73.6	46.6 ³	72.0
MASS	88.7	81.1	73.8	47.5	72.8

Table 1: Domain Generalization Performance Comparison Across Multiple Visual Domains. Out-of-domain accuracies with train-validation selection criterion on four DomainBed datasets comparing fixed-size GMoE, DynMoE, and MASS. ¹ from Guo et al. (2025), ² from Li et al. (2023), ³ reproduced by us (DynMoE result for TerraIncognita is not officially reported).

starts with the initial expert size 6. Table 1 reports the domain generalization accuracy across four benchmarks in DomainBed. MASS consistently achieves competitive or superior performance compared to GMoE baselines. Compared to DynMoE, MASS delivers higher generalization across all domains. These results indicate that MASS’s adaptive expansion mechanism is effective beyond the language domain, offering improved generalization in vision tasks.

Conclusion

We introduced MASS, a framework for adaptive expert expansion in MoE models that promotes semantic specialization through gradient-guided expansion and dynamic routing. MASS delivers three key benefits: (1) it discovers an optimal expert pool size without requiring costly hyperparameter search over K , (2) it enhances MoE performance by promoting semantic differentiation across experts, reducing functional redundancy and enabling more complementary expert

behaviors, achieving superior performance even with fewer active experts per token, and (3) it ensures balanced expert attention per token via adaptive expert selection based on routing confidence mass. These advantages are demonstrated across synthetic, language, and vision benchmarks, confirming both the robustness and generality of the approach.

Acknowledgements

This work was partly supported by the Institute for Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korean government (MSIT) (No. RS-2022-II220113, Developing a Sustainable Collaborative Multi-modal Lifelong Learning Framework; No. RS-2025-25442149, LG AI STAR Talent Development Program for Leading Large-Scale Generative AI Models in the Physical AI Domain), and Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT2402-08.

References

- Albuquerque, I.; Monteiro, J.; Darvishi, M.; Falk, T. H.; and Mitliagkas, I. 2021. Generalizing to unseen domains via distribution matching. arXiv:1911.00804.
- Baum, L. E.; and Petrie, T. 1966. Statistical inference for probabilistic functions of finite state Markov chains. *The annals of mathematical statistics*, 37(6): 1554–1563.
- Beery, S.; van Horn, G.; and Perona, P. 2018. Recognition in Terra Incognita. arXiv:1807.04975.
- Bentivogli, L.; Clark, P.; Dagan, I.; and Giampiccolo, D. 2009. The Fifth PASCAL Recognizing Textual Entailment Challenge. *TAC*, 7(8): 1.
- Dai, D.; Deng, C.; Zhao, C.; Xu, R. X.; Gao, H.; Chen, D.; Li, J.; Zeng, W.; Yu, X.; Wu, Y.; Xie, Z.; Li, Y. K.; Huang, P.; Luo, F.; Ruan, C.; Sui, Z.; and Liang, W. 2024. DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models. arXiv:2401.06066.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Dolan, W. B.; and Brockett, C. 2005. Automatically Constructing a Corpus of Sentential Paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. arXiv:2101.03961.
- Gulrajani, I.; and Lopez-Paz, D. 2020. In Search of Lost Domain Generalization. arXiv:2007.01434.
- Guo, Y.; Cheng, Z.; Tang, X.; Tu, Z.; and Lin, T. 2025. Dynamic Mixture of Experts: An Auto-Tuning Approach for Efficient Transformer Models. arXiv:2405.14297.
- Huang, Q.; An, Z.; Zhuang, N.; Tao, M.; Zhang, C.; Jin, Y.; Xu, K.; Xu, K.; Chen, L.; Huang, S.; and Feng, Y. 2024. Harder Task Needs More Experts: Dynamic Routing in MoE Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12883–12895. Bangkok, Thailand: Association for Computational Linguistics.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; de las Casas, D.; Hanna, E. B.; Bressand, F.; Lengyel, G.; Bour, G.; Lample, G.; Lavaud, L. R.; Saulnier, L.; Lachaux, M.-A.; Stock, P.; Subramanian, S.; Yang, S.; Antoniak, S.; Scao, T. L.; Gervet, T.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2024. Mixtral of Experts. arXiv:2401.04088.
- Lepikhin, D.; Lee, H.; Xu, Y.; Chen, D.; Firat, O.; Huang, Y.; Krikun, M.; Shazeer, N.; and Chen, Z. 2020. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. arXiv:2006.16668.
- Li, B.; Shen, Y.; Yang, J.; Wang, Y.; Ren, J.; Che, T.; Zhang, J.; and Liu, Z. 2023. Sparse Mixture-of-Experts are Domain Generalizable Learners. arXiv:2206.04046.
- Li, D.; Ma, Y.; Wang, N.; Ye, Z.; Cheng, Z.; Tang, Y.; Zhang, Y.; Duan, L.; Zuo, J.; Yang, C.; and Tang, M. 2024. MixLoRA: Enhancing Large Language Models Fine-Tuning with LoRA-based Mixture of Experts. arXiv:2404.15159.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, Broader and Artier Domain Generalization. arXiv:1710.03077.
- Puigcerver, J.; Riquelme, C.; Mustafa, B.; and Houlsby, N. 2024. From Sparse to Soft Mixtures of Experts. arXiv:2308.00951.
- Qiu, Z.; Huang, Z.; and Fu, J. 2024. Emergent Mixture-of-Experts: Can Dense Pre-trained Transformers Benefit from Emergent Modular Structures?
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. arXiv:1701.06538.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep Hashing Network for Unsupervised Domain Adaptation. arXiv:1706.07522.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. arXiv:1804.07461.
- Warstadt, A.; Singh, A.; and Bowman, S. R. 2019. Neural Network Acceptability Judgments. arXiv:1805.12471.
- Wu, J.; Hu, X.; Wang, Y.; Pang, B.; and Soricut, R. 2024. Omni-SMoLA: Boosting Generalist Multimodal Models with Soft Mixture of Low-rank Experts. arXiv:2312.00968.
- Xie, S. M.; Raghunathan, A.; Liang, P.; and Ma, T. 2022. An Explanation of In-context Learning as Implicit Bayesian Inference. arXiv:2111.02080.
- Xu, L.; Hu, H.; Zhang, X.; Li, L.; Cao, C.; Li, Y.; Xu, Y.; Sun, K.; Yu, D.; Yu, C.; Tian, Y.; Dong, Q.; Liu, W.; Shi, B.; Cui, Y.; Li, J.; Zeng, J.; Wang, R.; Xie, W.; Li, Y.; Patterson, Y.; Tian, Z.; Zhang, Y.; Zhou, H.; Liu, S.; Zhao, Z.; Zhao, Q.; Yue, C.; Zhang, X.; Yang, Z.; Richardson, K.; and Lan, Z. 2020. CLUE: A Chinese Language Understanding Evaluation Benchmark. In Scott, D.; Bel, N.; and Zong, C., eds., *Proceedings of the 28th International Conference on Computational Linguistics*, 4762–4772. Barcelona, Spain (Online): International Committee on Computational Linguistics.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.;

Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024. Qwen2 Technical Report. arXiv:2407.10671.

Zadouri, T.; Üstün, A.; Ahmadian, A.; Ermiş, B.; Locatelli, A.; and Hooker, S. 2023. Pushing Mixture of Experts to the Limit: Extremely Parameter Efficient MoE for Instruction Tuning. arXiv:2309.05444.

Zhang, Z.; Lin, Y.; Liu, Z.; Li, P.; Sun, M.; and Zhou, J. 2022. MoEfication: Transformer Feed-forward Layers are Mixtures of Experts. arXiv:2110.01786.

Zhou, Y.; Du, N.; Huang, Y.; Peng, D.; Lan, C.; Huang, D.; Shakeri, S.; So, D.; Dai, A.; Lu, Y.; Chen, Z.; Le, Q.; Cui, C.; Laudon, J.; and Dean, J. 2024. Brainformers: Trading Simplicity for Efficiency. arXiv:2306.00008.

Zhou, Y.; Lei, T.; Liu, H.; Du, N.; Huang, Y.; Zhao, V.; Dai, A.; Chen, Z.; Le, Q.; and Laudon, J. 2022. Mixture-of-Experts with Expert Choice Routing. arXiv:2202.09368.

Zhu, T.; Qu, X.; Dong, D.; Ruan, J.; Tong, J.; He, C.; and Cheng, Y. 2024. LLaMA-MoE: Building Mixture-of-Experts from LLaMA with Continual Pre-training. arXiv:2406.16554.