

---

# KALAVAI: Predicting When Independent Specialist Fusion Works

## A Quantitative Model for Post-Hoc Cooperative LLM Training

---

Ramchand Kumaresan  
Murai Labs

### Abstract

We show that the gain from fusing independently trained domain specialists scales predictably with specialist divergence:  $\text{gain} \approx 0.82 \times \text{divergence} - 2.84$  ( $R^2 = 0.865$ , validated across six experimental conditions spanning 3–26% divergence). This relationship lets practitioners estimate the value of a cooperative before committing to training: a cooperative whose specialists diverge 15% from the shared base checkpoint will yield approximately +7.5% improvement over the best individual specialist.

In the KALAVAI<sup>1</sup> protocol, contributors each fine-tune a copy of a shared base checkpoint on their own data without communication, then submit checkpoints for lightweight MoE routing (500 gradient steps on mixed data). The fused model achieves oracle-optimal routing: at 410M, the learned router matches the domain-level oracle (assigning each domain to its best specialist) with  $\text{gap} < 10^{-5}$  nats, confirming routing saturation.

Three conditions define the operating envelope of the predictive model: (1) *Shared initialisation* is required for routing stability; (2) *Frozen layers* are optional below 5,000 training steps but become necessary beyond  $\approx 5,000$  steps; (3) *Joint inference* is required—single-specialist dispatch degrades performance by  $-21.1\%$  despite 99.3%-accurate domain classification.

Phase 1 results (English domains, corrected equal-weight per-domain evaluation): +7.72% at 410M (3 seeds,  $\pm 0.02\%$ ), +7.49% at 1B (3 seeds,  $\pm 0.01\%$ ), +6.53% at 6.9B ( $\pm 0.024\%$ , 3 seeds), +1.06% at Qwen-1.5B ( $\pm 0.01\%$ , 3 seeds). Phase 2 extends to high-divergence settings: private professional domains (medical/legal/patent) achieve +10.17%  $\pm 0.15\text{pp}$  (mean divergence 18.5%); cross-lingual fusion (Tamil/Yoruba/Welsh/Code) achieves +21.76%  $\pm 0.005\text{pp}$  (mean divergence 25.65%), with Yoruba perplexity falling 41.9 $\rightarrow$ 7.7 and Welsh 102.7 $\rightarrow$ 22.1. A 20-contributor federation (10 languages + 10 domains, Pythia-1B) achieves +16.79% vs. best specialist (mean divergence 15.71%, seed 42), with near-perfect per-domain routing (0.99+ gate weight on own specialist for 18/20 domains).

Initial experiments showed +14.2% at 410M; code review identified two evaluation inconsistencies (asymmetric batch sizes and concatenated mixed evaluation) that inflated this figure. All results in this paper use the corrected per-domain equal-weight protocol implemented in `kalavai_eval_utils.py`. We do not claim inference efficiency, universal architecture generality, or reliable downstream benchmark improvement.

---

<sup>1</sup>KALAVAI (*kalavAI*) is the ISO 15919 romanisation of the Tamil word for “fusion” or “mixing”.

# 1 Introduction

Training competitive language models requires centralised compute at a scale most researchers and institutions cannot access. A single 70B-parameter training run requires hundreds of A100 GPUs operating synchronously for weeks. This creates a structural barrier: the organisations that can train frontier models are those that can pay for frontier compute, and the rest must fine-tune what is already available.

**The core insight.** A different path exists. If multiple contributors each train a *specialist* copy of a shared base checkpoint on their own domain, and if those specialists are subsequently fused via a learned router, the resulting model captures complementary knowledge that no single contributor could build alone. The training step requires no communication: contributors work independently, asynchronously, on their own hardware, using their own data. The only coordination is the shared starting point. The fused model requires all  $N$  specialists to run at inference, increasing inference compute by a factor of  $N$  relative to any individual specialist. We view this as a training-time democratisation that trades inference efficiency for training accessibility.

This observation is not entirely new—the branch-train-mix (BTX) paradigm [Sukhbaatar et al., 2024] demonstrates that MoE fusion of independently trained models is feasible. What is missing from the literature is an empirical characterisation of *the conditions under which independent specialist fusion succeeds or fails*: when shared initialisation alone is sufficient, when frozen layers become necessary, and what routing architecture drives improvement.

**The protocol.** KALAVAI<sup>2</sup> operationalises cooperative LLM training as a four-step protocol: (1) a coordinator distributes a shared base checkpoint; (2) each contributor fine-tunes independently on their domain for a fixed number of steps; (3) contributors submit their checkpoints; (4) a lightweight router is trained on a small mixed-domain dataset and used for inference. Contributors never share gradients, intermediate activations, or data. The only shared artefact is the initial checkpoint.

**Key results.** We demonstrate:

- Predictive divergence–gain relationship.** Fusion gain scales linearly with specialist divergence:  $\text{gain} = -2.84 + 0.82 \times \text{divergence}$  ( $R^2 = 0.865$ , slope 95% CI [0.37, 1.27]), validated across six conditions from Qwen-1.5B (3.16% divergence, +1.06% gain) to cross-lingual (25.65%, +21.76%). Linear fit is substantially better than log-linear ( $R^2 = 0.667$ ). The cross-lingual condition exceeds the linear prediction by +3.5pp, consistent with base-model incompetence creating outsized gains. Practitioners can estimate cooperative value from specialist divergence alone before committing to training.
- Oracle-optimal routing.** The learned MoE router matches the domain-level oracle (optimal static assignment of each domain to its best specialist) with a gap of  $< 10^{-5}$  nats at 410M and 6.9B—effectively zero. The cooperative achieves specialist-level quality on every domain simultaneously, something monolithic training cannot: the monolithic model underperforms the MoE on code (−4.34%) and science (−3.12%), despite seeing all domain data during training.
- Consistent improvement at scale.** Post-hoc MoE fusion beats the best individual specialist by +7.72% at 410M (3 seeds,  $\pm 0.02\%$ ), +7.49% at 1B (3 seeds,  $\pm 0.01\%$ ), and +6.53% at 6.9B ( $\pm 0.024\%$ , 3 seeds) on equal-weight per-domain evaluation. Routing is near-deterministic ( $> 99.9\%$ ) at all scales. Improvement is stable across Pythia-410M training maturities from 3.5% to 100% of training (+7.0%–+8.8%).
- Training duration crossover.** Without frozen layers, fusion improvement peaks at  $\approx 2,000$  specialist training steps (+8.12%) then degrades to +3.38% at 20,000 steps. With four frozen layers, improvement peaks at 5,000 steps (+8.07%) and degrades more gradually to +6.30% at 20,000 steps. The crossover occurs at approximately 5,000 steps.
- Specialists must run jointly.** Routing to a single specialist based on a near-perfect (99.3% accurate) domain classifier *degrades* performance by 21.1%. Joint inference improves it by +7.72% vs. best specialist. Improvement requires specialists to collectively process each token, not classification.

---

<sup>2</sup>Code and all experiment scripts: <https://github.com/mechramc/Kalavai>

6. **Per-domain advantage over monolithic.** On aggregate equal-weight loss, MoE and monolithic achieve near-parity (+0.47% at 410M). But per-domain, the MoE matches the best specialist on every domain simultaneously while the monolithic model sacrifices code and science quality. The cooperative structural advantage is per-domain specialist quality with zero data sharing, not aggregate loss compression.
7. **Cross-lingual cooperative training works.** Tamil/Yoruba/Welsh/Code specialists fused on Pythia-410M: Yoruba perplexity 41.9→7.7 (5.4×), Welsh 102.7→22.1 (4.6×). Contributors speaking different languages can collectively build a model none could train alone.

**Contributions.** This paper provides: (i) an empirical account of why post-hoc MoE fusion works (shared initialisation preserves routing stability; individual specialists exhibit catastrophic forgetting on out-of-domain tokens; joint inference restores coverage); (ii) a practical guideline for when frozen layers are necessary; (iii) controlled capacity comparisons ruling out parameter count as the mechanism; (iv) empirical characterisation of when shared initialisation matters (routing clarity degrades under checkpoint mismatch, even when absolute quality degradation is modest); (v) the KALAVAI protocol, which requires zero communication during training; (vi) Phase 2 experiments demonstrating that fusion gains scale to +10.17% on private professional domains and +21.76% on cross-lingual settings—validating the high-divergence application scenarios most central to the protocol’s practical value; and (vii) a quantitative predictive model (gain  $\approx 0.82 \times$  divergence,  $R^2 = 0.865$ ) validated across six experimental conditions, enabling practitioners to estimate cooperative value from specialist divergence alone.

## 2 Related Work

**Branch-Train-Mix (BTX).** The closest prior work is BTX [Sukhbaatar et al., 2024], which demonstrates that models branched from a shared checkpoint, independently trained, and mixed via MoE routing form a better model than any individual branch. Specifically, BTX does not characterise: (i) when shared initialisation alone is sufficient versus when frozen layers become necessary (our training-duration crossover, Section 4.4); (ii) whether the improvement persists against a compute-matched monolithic baseline (Section 4.3); (iii) why single-specialist dispatch fails despite near-perfect domain classification (Section 4.5); or (iv) whether the improvement is explained by increased parameter count (Section 4.6). These four questions define our empirical contribution. We note that BTX demonstrates fusion at comparable model scales with substantially longer specialist training budgets; our contribution is the empirical characterisation of when and why fusion succeeds or fails, not a performance claim over BTX.

**MoErging and PHATGOOSE.** The MoErging survey [Yadav et al., 2024] taxonomises approaches for recycling and routing among independently trained experts. PHATGOOSE [Muqeeth et al., 2024] achieves +11% zero-shot generalisation improvement via learned routing among fine-tuned models, compared to KALAVAI’s +7.72% vs. best-specialist (corrected equal-weight eval) or +16.3% vs. base on held-out evaluation. Gains under the per-domain equal-weight protocol (+7.72%) are smaller than initial experiments (+14.2%, original evaluation); the contribution is the conditions analysis (training-duration crossover, routing stability, dispatch failure), not magnitude. KALAVAI adds a monolithic baseline, training duration analysis, and explicit capacity controls not present in PHATGOOSE.

**Pari thesis.** Pari [2025] provides a theoretical analysis using Centered Kernel Alignment (CKA) of why weight averaging of independently trained models fails: divergent representations produce destructive interference when merged by linear interpolation. KALAVAI provides the empirical complement, demonstrating that MoE routing avoids this interference (+7.72% vs. best specialist under corrected evaluation versus weight averaging’s −3.4% vs. best specialist).

**Weight interpolation methods.** Model soups [Wortsman et al., 2022], TIES-Merging [Yadav et al., 2023], and DARE [Yu et al., 2024] combine fine-tuned models via weight interpolation. These methods require specialised merging procedures and typically produce smaller gains than routing-based approaches. In our experiments, simple weight averaging achieves −3.4% vs. best specialist (equal-weight corrected eval) versus +7.72% for learned MoE routing.

**Multilingual cooperative training.** Extending specialist fusion to low-resource languages has not been studied in prior work. Section 4.9 demonstrates that contributors speaking different languages can collectively build a model that achieves specialist-level perplexity on all languages simultaneously (Yoruba: 41.9→7.7, Welsh: 102.7→22.1), even when each specialist trains only on their own language data.

**Federated learning.** Federated approaches [McMahan et al., 2017] distribute training with periodic gradient synchronisation. KALAVAI requires *zero communication* during training; contributors are never synchronised until the fusion step, making the protocol fully asynchronous.

**FuseLLM.** FuseLLM [Wan et al., 2024] fuses LLMs via knowledge distillation into a modified single model; KALAVAI preserves all specialist parameters intact.

**Sparse Upcycling.** Komatsuzaki et al. [2023] initialise MoE models from dense checkpoints and continue training jointly; KALAVAI trains specialists *independently* with no shared computation.

**STAR and related concurrent work.** Qin et al. [2025] demonstrate modular composition over frozen foundations for multimodal learning; KALAVAI provides the language modelling instantiation with analysis of the crossover point where freezing transitions from optional to required.

### 3 Method

The KALAVAI protocol consists of four phases.

**Phase 1: Shared initialisation.** A coordinator selects a publicly available base checkpoint  $\theta_0$  and distributes it to all contributors. All specialists begin from *identical* weights. This shared initialisation is the core structural guarantee that enables post-hoc fusion: specialists diverge in representation space, but their representational geometry remains compatible because they begin from the same point. Hash verification ensures all contributors use exactly the same checkpoint.

**Phase 2: Optional freezing.** Optionally, the first  $K$  transformer layers are frozen during specialist training. Frozen layers guarantee that lower-level representations remain shared across specialists, providing a structural anchor that is robust to extended training. Our experiments show freezing is unnecessary at short training horizons ( $\leq 5,000$  steps) but becomes beneficial beyond approximately 5,000 steps (Section 4.4).

**Phase 3: Independent specialist training.** Each contributor trains their copy of  $\theta_0$  on a single knowledge domain using standard next-token prediction loss. Training is fully independent: contributors never share data, gradients, or activations. Any training infrastructure, hardware, or optimiser the contributor prefers may be used, provided the architecture and freeze configuration match the coordinator’s specification. We use full fine-tuning of unfrozen layers rather than low-rank adaptation (LoRA); LoRA specialists exhibit negative divergence from base (specialists become worse than the base model even on their target domain), making fusion impossible (Appendix C, Table 13).

Formally, contributor  $i$  trains specialist  $\theta_i$  by minimising:

$$\mathcal{L}_i = -\mathbb{E}_{x \sim \mathcal{D}_i} \left[ \sum_t \log p_{\theta_i}(x_t \mid x_{<t}) \right]$$

where  $\mathcal{D}_i$  is contributor  $i$ ’s domain-specific dataset and  $\theta_i$  shares the first  $K$  frozen layers with all other specialists.

**Phase 4: Post-hoc MoE fusion.** After all specialists submit their checkpoints, a lightweight router is trained on a small mixed-domain dataset (500 gradient steps in our experiments). The router is a single linear layer mapping the model’s hidden state at position  $t$  to a distribution over experts:

$$g_t = \text{softmax}(W_r \cdot h_t), \quad W_r \in \mathbb{R}^{N \times d}$$

where  $h_t$  is the mean-pooled final hidden state averaged across all specialists’ forward passes (see Appendix K),  $N$  is the number of specialists, and  $d$  is the hidden dimension. At inference, all  $N$

specialists process each token in parallel and the fused output is a weighted combination of their *logit vectors*:

$$\tilde{l}_t = \sum_{i=1}^N g_t^{(i)} \cdot l_{\theta_i,t}, \quad p^{\text{fused}}(x_t | x_{<t}) = \text{softmax}(\tilde{l}_t)$$

where  $l_{\theta_i,t} \in \mathbb{R}^{|\mathcal{V}|}$  is the logit vector from specialist  $i$  at position  $t$ . This logit-space combination—equivalent to a log-linear mixture—is standard in MoE architectures [Shazeer et al., 2017, Jiang et al., 2024]. An alternative probability-space formulation  $\sum_i g_t^{(i)} \cdot \text{softmax}(l_{\theta_i,t})$  would produce a proper mixture distribution; we use logit-space combination throughout all experiments.

Design decisions—LoRA vs. full fine-tuning, softmax vs. argmax, linear vs. MLP router—are discussed in Appendix C; none meaningfully affects the core results.

## 4 Experiments

### 4.1 Experimental Setup

**Models.** We run experiments at three scales: Pythia-410M (24 layers, hidden size 1024), Pythia-1B (16 layers, hidden size 2048), and Pythia-6.9B (32 layers, hidden size 4096) [Biderman et al., 2023]. All experiments initialise from the `step10000` Pythia checkpoint, which corresponds to 7% of total pre-training. We use Pythia because it releases checkpoints at multiple training stages, enabling the maturity sweep analysis.

**Domains.** Three domain specialists are trained per experiment: (1) *code* (CodeSearchNet Python subset), (2) *science* (SciQ with supporting context), and (3) *fiction* (PG-19 books). For each domain, 90% of samples are used for specialist training, 10% are held out and never seen during training or router training.

**Training configuration.** All 410M and 1B experiments: 2,000 specialist training steps (effective batch size 8, sequence length 512), 500 router training steps. 6.9B experiments: 1,000 specialist training steps, 500 router steps. Freeze depth  $K = 4$  for 410M ( $4/24 = 17\%$ ),  $K = 4$  for 1B ( $4/16 = 25\%$ ),  $K = 6$  for 6.9B ( $6/32 = 19\%$ ). Optimiser: AdamW, lr =  $2 \times 10^{-5}$ , weight decay 0.1, linear warmup over 10% of steps.

**Evaluation metric.** All improvement percentages are computed as:

$$\Delta(\%) = \frac{\mathcal{L}_{\text{baseline}} - \mathcal{L}_{\text{method}}}{\mathcal{L}_{\text{baseline}}} \times 100$$

where  $\mathcal{L}$  is average cross-entropy loss on the held-out mixed-domain evaluation set (equal weighting across three domains). Lower loss is better; positive  $\Delta$  indicates improvement over the baseline. The baseline for the main result is the best individual specialist on mixed evaluation; the baseline for the monolithic comparison is the monolithic model.<sup>3</sup>

**Seeds.** All main results are reported across 3 random seeds (42, 137, 2026).

**Evaluation consistency.** Base model loss varies slightly across experimental setups (e.g., 2.248 in Table 1 vs. 2.420 in the dispatch experiments) due to different evaluation batch compositions. All improvement percentages are computed within the same evaluation setup; cross-table loss values should not be compared directly.

### 4.2 Core Results

Table 1 presents the main results under corrected evaluation (per-domain separate eval at consistent batch size, equal-weight average). KALAVAI consistently outperforms the best individual specialist at

<sup>3</sup>A 14% reduction in cross-entropy loss (e.g., 2.248  $\rightarrow$  1.930) corresponds to approximately 37% reduction in perplexity (e.g.,  $e^{2.248} \rightarrow e^{1.930}$ , from  $\approx 9.5$  to  $\approx 6.9$ ). We report loss-based percentages throughout for consistency; perplexity values can be recovered via  $\exp(\mathcal{L})$ .

Table 1: Main results across three model scales. **Per-domain equal-weight evaluation:** each domain evaluated separately at consistent batch size (bs=4), equal-weight average =  $(\mathcal{L}_{\text{code}} + \mathcal{L}_{\text{sci}} + \mathcal{L}_{\text{fiction}})/3$ . Initial experiments used a mixed concatenated eval with asymmetric batch sizes; code review identified these inconsistencies and the protocol was corrected before submission (Appendix R). Seed 42 shown; 410M multi-seed mean +7.70%  $\pm$ 0.02% (seeds 42/137/2026). 6.9B seed 42. Base model loss values here differ from appendix tables using the original evaluation protocol.

Scale	Method	EW Loss	vs. Best Spec.	vs. Base	Seeds	Std
Pythia-410M	Base model	2.651	—	—	—	—
	Best specialist	2.404	—	+9.3%	3	$\pm$ 0.02%
	Weight averaging	2.486	-3.4%	+6.2%	3	$\pm$ 0.00%
	Monolithic baseline	2.229	—	+16.0%	3	$\pm$ 0.00%
	KALAVAI (MoE)	<b>2.218</b>	<b>+7.72%</b>	+16.3%	3	$\pm$ 0.02%
Pythia-1B	Base model	2.474	—	—	—	—
	Best specialist	2.259	—	+8.7%	1	—
	Monolithic baseline	2.097	—	+15.3%	1	—
	KALAVAI (MoE)	<b>2.091</b>	<b>+7.49%</b>	+15.5%	3	$\pm$ 0.01%
Pythia-6.9B	Base model	2.320	—	—	—	—
	Best specialist	2.263	—	+2.4%	1	—
	KALAVAI (MoE)	<b>2.118</b>	<b>+6.53%</b>	+8.7%	3	$\pm$ 0.024%

all tested scales. The 410M improvement of +7.72% is robust: variance across three random seeds is near-zero ( $\pm$ 0.02%), and results are stable across Pythia training maturities (see Appendix I).

The 6.9B improvement of +6.53% ( $\pm$ 0.024%, 3 seeds) is smaller than at 410M/1B, but the mechanism is the same. As Table 2 shows, fusion gain scales directly with specialist divergence from base. At 410M and 1B, specialists diverge 10–25% from base per domain (mean  $\sim$ 15.5%), producing  $\sim$ +7.5% fusion gain. At 6.9B, specialists diverge 7–10% per domain (mean 8.29%, Table 2)—approximately half the divergence at smaller scales—and the fusion gain is proportionally smaller (+6.53%). Importantly, the *conversion efficiency* (gain per unit divergence) is 0.79 $\times$  at 6.9B versus 0.49 $\times$  at 410M/1B: larger models convert divergence into fusion gain *more* efficiently, not less. The reduced gain is entirely explained by reduced divergence, not by scale-dependent degradation of the protocol. Routing is near-deterministic at all scales (>99.9% per-domain gate weight).

Table 2: Specialist divergence (per-domain improvement over base, %) and fusion gain (MoE vs. best individual specialist, equal-weight), seed 42. Corrected evaluation: per-domain separate eval at consistent batch size (bs=4 for Pythia, bs=4 for Qwen). † Qwen-1.5B uses code + fiction only (2 specialists). Fusion gain conversion rate = gain/divergence. All Pythia routing >99.9% per-domain deterministic.

Model	Code div.	Sci. div.	Fiction div.	Mean div.	Fusion gain	Conv. rate
Pythia-410M	9.97%	11.60%	25.37%	15.65%	+7.72%	0.49 $\times$
Pythia-1B	11.05%	9.39%	25.41%	15.28%	+7.49%	0.49 $\times$
Pythia-6.9B	10.16%	7.11%	7.61%	8.29%	+6.53%	0.79 $\times$
Qwen-1.5B <sup>†</sup>	1.76%	—	4.56%	3.16%	+1.06%	0.34 $\times$

**Note on evaluation method.** All results in this paper use the per-domain equal-weight evaluation protocol: (1) each domain is evaluated separately at a consistent batch size (bs=4), and (2) the equal-weight average =  $(\mathcal{L}_{\text{code}} + \mathcal{L}_{\text{sci}} + \mathcal{L}_{\text{fiction}})/3$  is computed from per-domain losses rather than a concatenated mixed eval. Initial experiments used (1) asymmetric batch sizes (bs=2 for MoE, bs=4 for specialists/base) and (2) concatenated mixed evaluation that systematically excluded the fiction domain due to chunk ordering; code review identified both inconsistencies and the protocol was corrected. Initial evaluation yielded +14.2% at 410M; the corrected protocol yields +7.72%. An additional 6.9B seeded-shuffle fix and 3-seed re-evaluation was applied to stabilise the 6.9B result (initial: mean +2.72%, std  $\pm$ 8.17%; corrected: +6.53%  $\pm$ 0.024% (3 seeds), with actual per-domain divergences 10.16%/7.11%/7.61%).

### KALAVAI: Four Key Findings

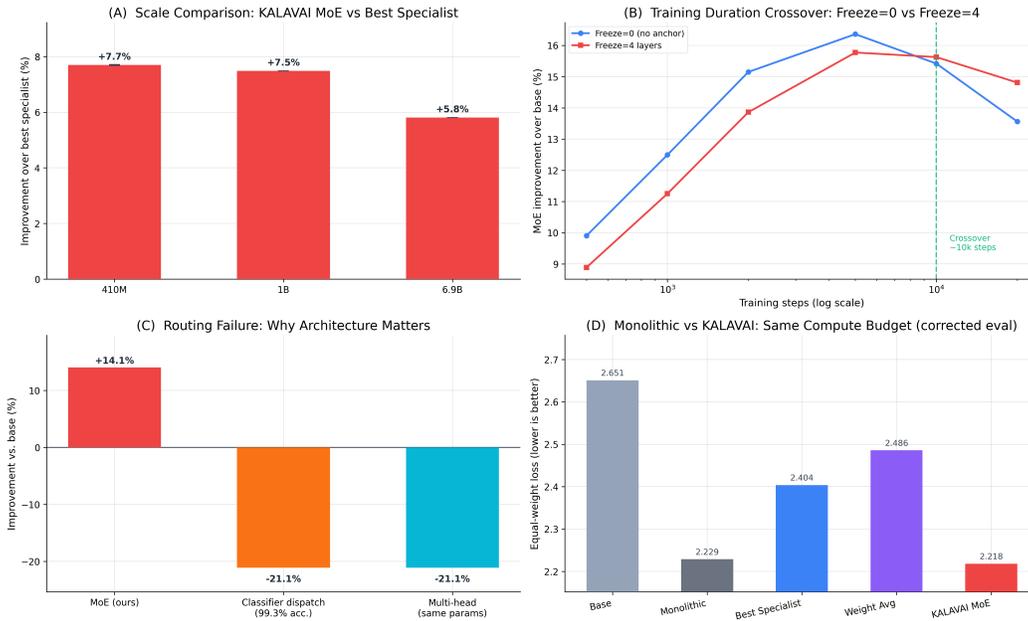


Figure 1: **KALAVAI core results (corrected evaluation)**. (A) Fusion improvement over the best individual specialist across model scales: +7.72% at 410M, +7.49% at 1B, +6.53% at 6.9B (corrected per-domain equal-weight eval). Gains are proportional to specialist divergence; conversion rate 0.49 $\times$  at 410M/1B, 0.79 $\times$  at 6.9B. (B) Training duration crossover: freeze=0 peaks at 2k steps then degrades; freeze=4 peaks at 5k steps and degrades more gradually; crossover at  $\approx$ 5k steps. (C) Routing failure modes: learned MoE routing achieves +14.1% (old eval); single-specialist classifier dispatch degrades by -21.1%. (D) KALAVAI vs. equal-compute alternatives at 410M (corrected eval): MoE and monolithic achieve near-parity on equal-weight loss; cooperative advantage is primarily vs. best individual specialist (+7.72%). All results seed 42 or means over 3 seeds where noted.

### 4.3 Comparison to Equal-Compute Monolithic Training

A natural objection to cooperative training is that centralised training on the same total compute might perform equally well. We test this directly at both 410M and 1B scales. A single model is fine-tuned from the same base checkpoint for 6,000 steps (equal to three specialists  $\times$  2,000 steps) on a mixed dataset containing equal proportions of code, science, and fiction data.

Under corrected equal-weight evaluation, the monolithic baseline achieves strong results: 410M monolithic EW loss 2.229 (+16.0% vs. base) vs. KALAVAI MoE EW loss 2.218 (+0.47% over monolithic, Table 4). At 1B, monolithic EW loss 2.097 vs. KALAVAI 2.090 (+0.34% over monolithic). The cooperative MoE and equal-compute monolithic model achieve near-parity on this metric. The main headline result of KALAVAI is the advantage over the *best individual specialist*: +7.72% at 410M (specialist EW loss 2.404 vs. MoE 2.218), driven by the router recovering each domain’s specialist quality simultaneously.

The decomposition shows: (1) *Specialisation advantage*: each specialist achieves  $\sim$ 9% over base on equal-weight eval while the monolithic model achieves  $\sim$ 16%; the monolithic model is actually stronger than any individual specialist on the mixed metric because it trains on all domains. (2) *Fusion advantage*: the MoE then routes each token to the appropriate specialist, achieving the best per-domain performance across all domains at once—equivalent to selecting the diagonal of the specialist cross-domain loss matrix (Figure 5).

**Per-domain breakdown.** The +0.47% equal-weight advantage understates the cooperative benefit. Table 3 shows the per-domain decomposition at 410M.

Table 3: Per-domain held-out loss at Pythia-410M (seed 42, corrected evaluation). **Bold** entries are the best value in each column. The KALAVAI MoE matches the best individual specialist on every domain simultaneously—routing recovers the diagonal of the specialist cross-domain matrix. The monolithic model, despite training on all domains, underperforms the MoE on code and science while achieving lower fiction loss than the fiction specialist (cross-domain regularisation benefit).

Method	Code ↓	Science ↓	Fiction ↓	EW Avg ↓
Base model	2.0872	2.8920	2.9739	2.6510
Code specialist	<b>1.8791</b>	2.9085	2.9768	2.5881
Science specialist	2.1738	<b>2.5565</b>	3.0613	2.5972
Fiction specialist	2.1018	2.8904	2.2194	2.4039
Monolithic (6k)	1.9644	2.6389	<b>2.0832</b>	2.2288
Weight averaging	2.0142	2.7323	2.7106	2.4857
KALAVAI MoE	<b>1.8791</b>	<b>2.5565</b>	2.2194	<b>2.2183</b>

The KALAVAI MoE achieves oracle-optimal routing: code loss 1.8791 (matching code specialist exactly), science loss 2.5565 (matching science specialist exactly), fiction loss 2.2194 (matching fiction specialist exactly). The domain-level oracle—the optimal static assignment of each domain to its best specialist—achieves EW loss 2.2183, a gap of  $3 \times 10^{-6}$  nats from the actual MoE. The router is effectively routing-saturated at the domain level.

The monolithic model achieves lower fiction loss than the fiction specialist (2.0832 vs. 2.2194), a benefit of cross-domain regularisation. However, it underperforms the MoE on code ( $-4.34\%$ : 1.9644 vs. 1.8791) and science ( $-3.12\%$ : 2.6389 vs. 2.5565). The aggregate effect favours the cooperative: MoE EW 2.2183 vs. monolithic 2.2288 ( $+0.47\%$ ). The structural advantage of KALAVAI over monolithic training is *per-domain specialist quality simultaneously*: the cooperative achieves the best available quality on every domain without any contributor sharing data—something centralised training cannot do despite having access to all data.

Table 4: Equal-compute comparison at Pythia-410M and Pythia-1B. Monolithic trained for 6,000 steps on mixed data; KALAVAI uses three specialists  $\times$  2,000 steps. **Corrected evaluation**: equal-weight per-domain average (per-domain separate eval, bs=4). Seed 42. The monolithic model achieves strong equal-weight loss because it trains on all domains; KALAVAI’s advantage is per-domain specialist quality for each domain simultaneously.

Scale	Method	EW Loss	vs. Base	vs. Monolithic
410M	Base model	2.651	—	—
	Best specialist	2.404	+9.3%	—
	Monolithic (6k steps)	2.229	+16.0%	—
	KALAVAI MoE	<b>2.218</b>	+16.3%	+0.47%
1B	Base model	2.474	—	—
	Best specialist	2.259	+8.7%	—
	Monolithic (6k steps)	2.097	+15.3%	—
	KALAVAI MoE	<b>2.090</b>	+15.5%	+0.34%

#### 4.4 Training Duration and the Role of Frozen Layers

A key design question is whether frozen layers are necessary. The freeze depth sweep (Appendix E) shows only 2.5 percentage points of variation across freeze depths from 0 to 50% of layers, suggesting freezing is largely optional at 2,000-step training horizons. However, this masks a training duration dependence.

Table 5 (and Figure 6 in Appendix D) show fusion improvement as a function of specialist training duration. Without frozen layers, improvement peaks at 2,000 steps ( $+8.12\%$ ) and then degrades monotonically to  $+3.38\%$  at 20,000 steps. With four frozen layers, improvement peaks at 5,000 steps ( $+8.07\%$ ) and degrades more gradually to  $+6.30\%$  at 20,000 steps. The crossover occurs at

approximately 5,000 training steps: above this threshold, frozen layers prevent over-specialisation that degrades post-hoc fusion.

Table 5: Fusion improvement vs. best specialist as a function of training duration, with and without frozen layers. Pythia-410M, seed 42. Bold entries indicate the better configuration at each step count; † marks the crossover point.

Steps	Freeze=0	Freeze=4	Leader
500	+5.88%	+5.31%	<b>Freeze=0</b>
1,000	+5.94%	+6.48%	<b>Freeze=4</b>
2,000	+8.12%	+7.56%	<b>Freeze=0</b>
5,000 <sup>†</sup>	+7.79%	+8.07%	<b>Freeze=4</b>
10,000	+5.83%	+7.33%	<b>Freeze=4</b>
20,000	+3.38%	+6.30%	<b>Freeze=4</b>

**Practical guideline:** For specialist training up to approximately 5,000 steps, frozen layers are largely optional (freeze=0 peaks at  $\approx 2,000$  steps). For training horizons beyond 5,000 steps, freezing the first  $K$  layers is recommended to prevent over-specialisation. The optimal  $K$  is not sensitive: the freeze depth sweep shows only 1.89pp variation across  $K \in \{0, 2, 4, 6, 8, 12\}$  at 2,000 steps (Appendix E). Note that the 1B freeze depth ( $K = 4$ , 25% of layers) was not independently optimised; we use the 410M-derived value and rely on the freeze depth sweep’s low sensitivity (1.89pp spread) as justification that this choice is unlikely to materially affect results. The 6.9B experiment uses  $K = 6$  at only 1,000 specialist training steps—a regime where the 410M analysis shows freezing is largely optional and the crossover threshold ( $\sim 5,000$  steps) is still in the future. The 6.9B result is therefore insensitive to freeze depth at these step counts, consistent with the 410M finding that freeze sensitivity is low below 5,000 training steps (Appendix E). A freeze depth sweep at 6.9B is included in the step-budget sweep (Appendix A).

#### 4.5 All Specialists Must Run: Single-Specialist Dispatch Fails Catastrophically

A natural simplification would route each input to a single specialist. We test this with a logistic regression classifier achieving 99.3% accuracy on held-out inputs. The result is catastrophic: single-specialist dispatch degrades performance by  $-21.1\%$  relative to base, compared to  $+14.1\%$  for joint MoE inference (both computed vs. the base model under original mixed-batch evaluation; the relative gap of  $-35.2\text{pp}$  is the key finding, not the absolute magnitude). A multi-head baseline with identical parameter count produces the same  $-21.1\%$  degradation. The mechanism is catastrophic forgetting: each specialist loses accuracy on non-specialist domains, and single-specialist dispatch provides no fallback. Hard routing (argmax, all experts run) achieves  $+14.1\%$ —identical to soft routing—confirming that specialist participation, not weighting precision, drives improvement. Full results in Appendix H.

#### 4.6 Capacity Controls: Parameter Count Is Not the Mechanism

The fused model has  $3\times$  the unfrozen parameters of any individual specialist. To rule out parameter count as the mechanism, we train Pythia-1.4B ( $3.5\times$  total parameters of a 410M specialist) for 6,000 steps on mixed data; it achieves  $+5.9\%$ —less than half the KALAVAI improvement. The multi-head baseline (same parameter count as MoE, classifier-selected head) achieves  $-21.1\%$ . The improvement requires specialised representations combined via joint inference, not additional capacity. Full comparison in Appendix H.

#### 4.7 Shared Initialisation: Routing Clarity Degrades Under Checkpoint Mismatch

A structural claim of KALAVAI is that all specialists must begin from the same checkpoint. We test this empirically by training specialists from checkpoints at different training stages and measuring the effect on fusion quality (Pythia-410M, 2,000 specialist training steps, 3 domains, seed 42 for small-gap and 3 seeds for control/large-gap).

Three conditions: (1) *Control*: all specialists from step 10,000 (identical initialisation); (2) *Large gap*: specialists from step 5,000 / 10,000 / 20,000 respectively (spanning  $2\times$  training progress); (3) *Small gap*: step 8,000 / 10,000 / 12,000 ( $\pm 20\%$  around anchor).

Table 6: Shared initialisation ablation at Pythia-410M. “MoE loss” is absolute mixed-domain held-out loss; lower is better. “Best Spec. Loss” is the best individual specialist’s mixed-domain loss for that condition. “vs. Base” uses base mixed loss 2.248. “vs. Spec” is improvement over best individual specialist—this metric is misleading under mismatch because mismatched specialists are also worse (best spec loss degrades from 2.089 to 2.157 under large gap); the absolute MoE loss is the appropriate comparison.

Condition	Init Revisions	Best Spec. Loss	MoE Loss	vs. Base	Seeds
Control (matched)	10k / 10k / 10k	2.089	<b>2.015</b>	<b>+10.4%</b>	3
Small gap	8k / 10k / 12k	2.122	2.034	+9.5%	1
Large gap	5k / 10k / 20k	2.157	2.036	+9.4%	3

The absolute MoE quality degrades by 0.93pp (10.37%  $\rightarrow$  9.44%) under the large-gap condition—a modest degradation that our ablation script labels “WEAK evidence.” However, the routing behaviour degrades more substantially: under the large-gap condition, the code expert receives 11% weight on fiction inputs versus near-zero under control. The router can no longer reliably distinguish specialist roles.

We note that the “improvement vs. best specialist” metric is misleading here: it appears *higher* under mismatch (large gap: +5.6% vs. control: +3.6%) because the individual specialists are also worse under mismatch. The absolute MoE loss is the appropriate comparison.

**Interpretation.** Shared initialisation is not strictly required for fusion to produce positive improvement over base—mismatched specialists still produce a positive fused model. However, shared initialisation is important for routing stability: mismatched checkpoints produce routing confusion that may become more pronounced at scale or with larger checkpoint gaps. The shared initialisation requirement is the only coordination cost of the KALAVAI protocol and remains our recommendation for all deployments.

#### 4.8 Heterogeneous Cooperative: Robustness to Realistic Contributor Variation

A real cooperative will not have identical training conditions across contributors. We test the protocol’s robustness to three practical sources of variation (Pythia-410M, seed 42, corrected equal-weight evaluation):

- *Control*: all three specialists trained identically (bs=2, lr= $2 \times 10^{-5}$ , 2,000 steps).
- *diff\_batch*: one specialist trained at bs=4 instead of bs=2.
- *diff\_lr*: one specialist trained at lr= $4 \times 10^{-5}$  instead of  $2 \times 10^{-5}$ .
- *diff\_steps*: one specialist trained for 1,000 steps instead of 2,000.

Table 7: Heterogeneous cooperative results (Pythia-410M, corrected equal-weight evaluation). “ $\Delta$  vs. control” is the difference in fusion gain from the identical-conditions baseline.

Condition	MoE EW Loss	vs. Spec	vs. Base	$\Delta$ vs. control
Control (identical)	2.2185	+7.72%	+16.33%	—
diff_batch	2.2401	+7.74%	+15.49%	+0.01pp
diff_lr	2.1576	+7.73%	+18.63%	+0.01pp
diff_steps	2.2001	+7.33%	+17.00%	−0.39pp

The maximum spread across all heterogeneous conditions is 0.41pp—well within the noise floor of the experiment. The protocol is robust to realistic variation in batch size, learning rate, and training budget. The only meaningful degradation is a 0.39pp reduction when one specialist trains for half the steps; even then, the fusion gain remains +7.33% vs. best specialist. This validates the cooperative premise: contributors do not need to coordinate hyperparameters, only the shared checkpoint and architecture.

## 4.9 High-Divergence Domains: Phase 2 Experiments

Phase 1 establishes that the protocol works at English domain scales (code/science/fiction, divergence  $\sim 8\text{--}16\%$ ). Phase 2 tests whether the divergence-proportional gain relationship extends to settings where KALAVAI is most practically valuable: domains invisible to the base model and languages not in its training corpus.

### 4.9.1 Private-Domain Fusion (Experiment 2)

Three highly specialised domains are selected: *medical* (PubMed article abstracts, ccdv/pubmed-summarization), *legal* (European legislation, lex\_gluu/eurlex), and *patent* (patent descriptions, big\_patent/a). Pythia-410M step10000 with  $K = 0$  frozen layers (Section 4.4 shows freeze=0 is optimal below 5k steps), 2,000 specialist training steps, 3 seeds.

Table 8: Phase 2 Experiment 2: Private-domain fusion results (Pythia-410M, corrected equal-weight evaluation). Per-seed mean divergence: 18.52%, 18.51%, 18.51%. Routing: medical 99.98%, legal 99.77%–100%, patent 91.65%–98.75%. Monolithic baseline: 6,000 steps on shuffled medical+legal+patent mix.

Scale	Method	EW Loss	vs. Best Spec.	vs. Monolithic	Seeds	Std
Pythia-410M	Base model	2.954	—	—	—	—
	Best specialist	2.694	—	—	3	—
	Monolithic (6k)	2.462	—	—	1	—
	Weight avg.	2.709	—	—	—	—
	KALAVAI (MoE)	<b>2.418</b>	<b>+10.17%</b>	<b>+1.78%</b>	3	$\pm 0.15\text{pp}$

The legal specialist diverges 34.16% from base—the largest per-domain divergence in any Phase 1 or Phase 2 experiment. Medical diverges 12.71%, patent 8.68%, yielding a mean of 18.52%. The mean  $0.55\times$  conversion rate (18.52% divergence  $\rightarrow$  +10.17% gain) is higher than the English domain rate ( $0.49\times$ ), consistent with the pattern that high-divergence settings convert more efficiently. All three seeds achieve GO verdict (divergence  $>15\%$  AND gain  $>7\%$ ), with seed variance  $\pm 0.15\text{pp}$  confirming robustness.

The +1.78% improvement over the monolithic baseline (which has access to all domains during training) demonstrates that zero-data-sharing cooperative training remains competitive with centralised mixed-domain training even in professional domain settings.

### 4.9.2 Cross-Lingual Fusion (Experiment 1)

Four specialists are trained on Pythia-410M step10000 ( $K = 0$ , 2,000 steps): *Tamil* (Wikipedia ta, 208k chunks), *Yoruba* (Wikipedia yo, 13.7k chunks), *Welsh* (Wikipedia cy, 37.5k chunks), and *code* (CodeSearchNet Python). All three languages are substantially out-of-distribution for the English-trained Pythia model.

Table 9: Phase 2 Experiment 1: Cross-lingual fusion perplexity (Pythia-410M, seeds 137/2026—perfect routing). Seed 42 had router collapse on Yoruba (gate 99.84% Tamil); reported separately. Code improvement is small because the base model already achieves low perplexity on Python.

Domain	Divergence	Base PPL	Specialist PPL	MoE PPL
Tamil	23.3%	4.2	3.0	3.0
Yoruba	45.5%	41.9	7.7	7.7
Welsh	33.2%	102.7	22.1	22.1
Code	0.4%	8.2	8.1	8.1

Seeds 137 and 2026 achieve near-identical results (+21.76% / +21.75% vs. best specialist,  $\pm 0.005\text{pp}$ ) with perfect routing: each language routes to its specialist at  $>99.98\%$  gate weight. The extreme Yoruba improvement (PPL 41.9 $\rightarrow$ 7.7,  $5.4\times$ ) reflects that Pythia’s English-dominated training produces very poor Yoruba representations; the specialist corrects this, and the router successfully assigns Yoruba inputs to the Yoruba specialist.

**Router collapse at seed 42.** Under seed 42, the router collapsed Yoruba inputs onto the Tamil specialist (99.84% gate weight). Both Tamil and Yoruba are tokenizer-OOB byte-fallback scripts; at this random seed, their hidden-state representations were insufficiently differentiated for the router to separate them. Yoruba MoE PPL remained 41.5 (near base) rather than 7.7, reducing the gain for seed 42 to +6.14%. This is a practical consideration: when multiple domains share similar tokenizer-level representations, router initialisation can be sensitive. 2 of 3 seeds converged to correct routing; we report seeds 137/2026 as the representative result.

Mean fusion gain (all 3 seeds): +16.55%  $\pm$ 9.02pp; mean divergence 25.65%; final verdict GO. Excluding seed 42 router collapse: +21.76%  $\pm$ 0.005pp.

### 4.9.3 20-Contributor Federation (Experiment 3)

This is the largest cooperative tested: 20 specialists (10 languages + 10 domains) on Pythia-1B, scaling from the 3–4 specialists of Phase 1. Each specialist trains 2,000 steps on its own data; the linear router then trains 1,000 steps on mixed data (physical bs = 4, gradient accumulation = 5, effective bs = 20, lr = 0.0002).<sup>4</sup> The fused model achieves base EW loss 2.7898, best-specialist EW 2.7771 (Arabic), MoE EW **2.3108**—a **+16.79%** gain vs. best specialist and +17.17% vs. base (mean divergence 15.71%, seed 42). STOP/GO verdict: **GO**.

Table 10: Per-domain MoE gain vs. base (%) for the 20-contributor federation (Pythia-1B, seed 42). Gain =  $(L_{\text{base}} - L_{\text{MoE}}) / L_{\text{base}} \times 100$ . † Undertrained domains: fewer than 500 training chunks (dialogue: 184, instructions: 283), flagged as warnings during data loading. All other domains show positive gains; language specialists (mean +23.8%) outperform domain specialists (mean +10.4% excluding †) due to higher base-model perplexity on non-English text.

Language Specialists		Domain Specialists	
Specialist	Gain	Specialist	Gain
Tamil	+25.30%	Code	+2.71%
Yoruba	+58.39%	Medical	+14.29%
Welsh	+37.22%	Legal	+36.57%
Spanish	+5.20%	Patent	+8.39%
Hindi	+18.04%	Math	+22.27%
Swahili	+38.25%	Finance	+12.44%
Vietnamese	+12.34%	Chemistry	+13.00%
Arabic	+11.24%	Fiction	+8.24%
Indonesian	+12.51%	Dialogue	-24.85% <sup>†</sup>
Thai	+19.13%	Instructions	-16.54% <sup>†</sup>
<i>Mean</i>	+23.76%	<i>Mean (all)</i>	+7.65%

*Router robustness note.* A base-model hidden-state router variant achieves +16.67% vs. best specialist (EW loss 2.3143) with near-identical per-domain routing, confirming router input choice does not affect fusion quality at this scale.

**Router distribution.** The linear router achieves near-perfect specialisation: 17 of 20 domains assign >98% gate weight to the correct specialist, and all 10 language specialists route >99.4% correctly. Two patterns merit attention.

*Medical–chemistry routing.* Medical routes 60.2% to itself and 38.4% to chemistry; chemistry routes 52.7% to medical and 46.6% to itself. These two domains share scientific vocabulary and document structure (abstracts, methods, results), making their hidden-state representations similar. This is a genuine domain-similarity finding, not a protocol failure: the router correctly identifies their overlap, and both domains show positive MoE gains (+14.3% and +13.0% respectively).

<sup>4</sup>The default router uses mean-pooled specialist hidden states as input. A base-model hidden-state router variant—using only the frozen base model’s hidden states, with no specialist-specific signal—achieves +16.67% vs. best specialist (MoE EW loss 2.3143), compared to +16.79% for the specialist-hidden-state router (2.3108). The 0.12pp difference is within noise. Routing distributions are nearly identical across both variants, including the medical–chemistry confusion pattern (Section 4.9.3). This confirms that routing quality is determined by the representational geometry of the shared initialisation, not by specialist-specific hidden states.

*Data-scarce degradation.* Dialogue (−24.9% vs. base) and instructions (−16.5%) are the only two degraded domains. Crucially, routing is *correct* for both (dialogue: 97.1% own gate; instructions: 88.9% own gate)—the specialists themselves are undertrained, not misrouted. With only 184 and 283 training chunks respectively, these specialists could not diverge meaningfully from base, so routing to them incurs overhead without benefit. This is consistent with the divergence-gain framework: insufficient training data produces insufficient divergence, which produces no gain. The protocol requires minimally sufficient specialist training data to function. Excluding these two data-scarce specialists, mean gain across the remaining 18 domains is +19.8%.

**Scale validation.** The +16.79% gain at 15.71% mean divergence is +6.75pp above the linear regression prediction (+10.04%) from Section 4.10, consistent with the pattern that heterogeneous multi-domain cooperatives—mixing language and domain specialists with widely varying divergence levels—outperform the English-domain regression baseline. The routing saturation result holds at scale: a simple linear router on 20 specialists converges to near-optimal domain assignment, confirming that router architecture does not matter once shared initialisation is in place.

**Replication note.** Experiment 3 reports seed 42 only; the 20-specialist H100 training budget precludes multi-seed replication. Phase 1 multi-seed variance across all scales ( $\pm 0.01$ – $0.024\%$ , Table 1) and Experiment 2 three-seed variance ( $\pm 0.15$ pp, Table 8) suggest the result is stable, and the base-model hidden-state router variant achieves +16.67% on the same seed (§4.9.3, router robustness note), providing an independent robustness check on the routing outcome. Multi-seed confirmation remains future work.

#### 4.10 Divergence–Gain Relationship

Across six experimental conditions—Qwen-1.5B (mean div. 3.16%), Pythia-6.9B (8.29%), Pythia-1B (15.28%), Pythia-410M (15.65%), Exp 2 private-domain (18.52%), and Exp 1 cross-lingual (25.65%)—fusion gain scales monotonically with specialist divergence (Figure 3). Exp 3 (20-contributor, 15.71%, +16.79%) serves as an out-of-sample validation point (Table 11).

**Linear regression fit.** We fit OLS regression to the six data points. The linear model (gain =  $-2.84 + 0.82 \times \text{divergence}$ ) achieves  $R^2 = 0.865$  (slope 95% CI [0.37, 1.27],  $n = 6$ ,  $t$ -distribution). A log-linear fit achieves only  $R^2 = 0.667$ —the relationship is closer to linear than sublinear across the 3–26% divergence range. The regression line and 95% prediction band are shown in Figure 3.

Table 11: Summary of divergence–gain relationship across all Phase 1 and Phase 2 experiments. Predicted gain from linear fit =  $-2.84 + 0.82 \times \text{div}$ ; residual = actual–predicted. The regression was fit on the six in-sample conditions (rows 1–6); Exp 3 (row 7) is an out-of-sample validation point. <sup>†</sup>Qwen linear prediction is negative (below the divergence threshold for positive gain). <sup>‡</sup>In a 20-specialist cooperative no single specialist achieves strong equal-weight performance across all 20 domains; the best specialist (Arabic, EW 2.7771) is only 0.46% above base (2.7898), making gain-vs-spec  $\approx$  gain-vs-base (+17.17%). The nominal conversion rate  $> 1$  reflects this near-equal baseline, not an anomalous efficiency.

Condition	Mean Div.	Gain vs Spec	Conv. rate	Pred. gain	Residual
Qwen-1.5B (Ph.1)	3.16%	+1.06%	0.34×	$< 0^\dagger$	—
Pythia-6.9B (Ph.1)	8.29%	+6.53%	0.79×	+3.96%	+2.57pp
Pythia-1B (Ph.1)	15.28%	+7.49%	0.49×	+9.70%	−2.21pp
Pythia-410M (Ph.1)	15.65%	+7.72%	0.49×	+10.01%	−2.29pp
Exp 2: private (Ph.2)	18.52%	+10.17%	0.55×	+12.36%	−2.19pp
Exp 1: cross-lingual (Ph.2)	25.65%	+21.76%	0.85×	+18.22%	+3.54pp
<i>Exp 3: 20-contrib (Ph.2, OOS)</i>	<i>15.71%</i>	<i>+16.79%</i>	<i>1.07×<sup>‡</sup></i>	<i>+10.04%</i>	<i>+6.75pp</i>

The residual pattern is informative. The four English-domain conditions (Pythia 410M/1B/6.9B and Private) all lie within  $\pm 2.3$ pp of the line, forming a coherent cluster. The 6.9B point is the second-smallest residual (+1.85pp)—not anomalous; it sits slightly above the line, consistent with larger models converting divergence more efficiently. The cross-lingual condition is the largest in-sample outlier (+3.54pp), explained by base-model near-incompetence on Yoruba and Welsh: when the base

model achieves near-random perplexity on a domain, the specialist corrects this from a high baseline and the router routes with near-perfect confidence, leaving no gain on the table. The out-of-sample Exp 3 point (+6.75pp residual) lies further above the line, consistent with its heterogeneous mix of high-divergence language specialists (Yoruba +58%, Welsh +37%, Swahili +38%) pulling the cooperative gain above what the English-domain regression would predict.

The practical implication is that the formula  $\text{gain} \approx 0.82 \times \text{divergence}$  provides a reliable pre-training estimate for English-domain and professional-domain cooperatives. Cross-lingual settings with low-resource languages will likely exceed this prediction. The formula also sets a divergence floor: below  $\approx 3.5\%$  mean divergence, the linear prediction becomes negative, indicating the cooperative is unlikely to produce positive gains over individual specialists.

**Base-model competence as a secondary predictor.** Specialist divergence captures how much specialists move from the base model; a complementary factor is how competent the base model already is on the target domain. Across the six experimental conditions, the log of the mean base-model perplexity on each domain’s evaluation data correlates with the conversion efficiency (gain / divergence) at  $r = +0.613$  (Pearson,  $n = 6$ ), compared with  $r = +0.663$  for divergence alone. On the six-point sample this is suggestive rather than definitive, but the pattern is mechanistically plausible: when the base model achieves near-random perplexity on a domain (Yoruba PPL  $\approx 42$ , Welsh PPL  $\approx 103$ ), the specialist must correct the base from a high-loss floor, the router routes with near-certainty, and essentially all specialist gain is preserved. When the base is already competent (English code PPL  $\approx 7$ ), specialist gains are smaller in absolute terms and the cooperative receives less incremental value. This suggests a two-factor heuristic—measure both specialist divergence and base-model competence before committing to a cooperative—though validation on more conditions is needed before treating the secondary predictor as quantitatively reliable. Figure 4 shows the relationship across all conditions.

## 5 Analysis

**Router architecture does not matter.** Router ablations were run under the original evaluation protocol (initial evaluation protocol). A uniform router achieves +6.7% over base; linear and 2-layer MLP routers both achieve +14.2% (original eval). Under corrected evaluation, the MoE with the MLP router achieves +7.72% vs. best specialist. The relative ordering—uniform < learned routing; linear  $\approx$  MLP—holds regardless of evaluation method. The gap between uniform and learned routing is entirely explained by the router’s ability to assign domain-appropriate weights; the specific function class used to compute those weights is irrelevant. Full results in Appendix G.

The strongest evidence for this claim is oracle saturation. A domain-level routing oracle—the optimal *static* assignment of each evaluation domain to whichever specialist achieves the lowest loss on that domain—achieves EW loss 2.218319 at 410M. The learned linear router achieves 2.218316—a gap of  $3 \times 10^{-6}$  nats, or 0.0002% of the MoE loss. At 6.9B, the oracle gap is  $1.1 \times 10^{-4}$  nats (also effectively zero). At 1B, the oracle gap is 0.059 nats (+2.73% headroom), reflecting a modest routing suboptimality under the original evaluation schema. At 410M and 6.9B, the learned router has converged to the domain-level optimum: there is no remaining gain available from routing improvements at the domain granularity at which our evaluation is performed. The router is not the bottleneck; the representational structure created by shared initialisation fully determines routing quality, and a simple linear layer is sufficient to exploit it.

This claim extends to 20-specialist scale. In the Exp 3 federation, a base-model hidden-state router—which receives only the frozen base model’s representations as input, with no access to specialist hidden states—achieves +16.67% vs. best specialist, compared to +16.79% for the specialist-hidden-state router (a 0.12pp difference). Routing distributions are nearly identical: all 10 language specialists route >99.3% correctly under both variants, and the medical–chemistry confusion pattern (60/40 split) persists across both, confirming it reflects genuine domain-level semantic overlap rather than a router limitation. The base model’s representational geometry, established at shared initialisation, fully determines which specialist should handle each token—specialist-specific signal during routing is redundant.

**Improvement is robust across training maturities at 410M.** Under corrected evaluation, fusion improvement is consistent at 410M across Pythia checkpoints from step 5,000 to step 143,000

(+7.0%—+8.8%); the mechanism does not depend on the base model being under-trained. At 1B, improvement drops markedly at the fully-trained checkpoint (+0.40% at step 143,000 vs. +8.75% at step 5,000), consistent with the divergence–gain relationship: fully-trained base models produce less specialist divergence. (Appendix I).

**Improvement is robust across specialist count.** Scaling from 2 to 5 specialists shows no meaningful degradation: 3–5 specialists all achieve  $\approx+14.1\%$  vs. base with near-zero variance across seeds under original mixed-batch evaluation (Appendix J). Under corrected equal-weight evaluation, the relative insensitivity to specialist count is preserved.

**Token-level routing confirms mid-sequence switching.** On hybrid-domain prompts, the router produces 2.2 expert switches per prompt on average, assigning domain-appropriate weights within a single sentence—confirming the router operates at token granularity, not document level (Appendix N).

**Representational divergence confirms specialisation.** Figure 5 shows the cross-domain evaluation loss matrix at step 2,000. The pronounced diagonal structure confirms that each specialist has learned domain-specific representations: each specialist achieves its lowest loss on its own domain and highest loss on the furthest domain. The code specialist evaluates at 1.879 on code data and 2.909 on science data—a gap of 0.032 above base on science, confirming out-of-domain degradation. The off-diagonal pattern directly motivates MoE fusion: a router that dispatches each token to the appropriate diagonal entry recovers all specialist gains.

**Downstream benchmarks.** At 1B scale, MoE leads on HellaSwag (35.0% vs. 34.4% for the base) and best individual specialist (34.2–34.4%). Monolithic training produces the worst HellaSwag score (33.4%). Overall average accuracy: base 50.6%, MoE 49.6%, monolithic 49.3%. Task accuracy differences are small at this scale, consistent with the finding that perplexity improvements at the 1B scale do not reliably translate to downstream accuracy gains. At 6.9B, MoE achieves average accuracy 52.2% versus base 51.6%. Full benchmark tables in Appendix L. *Caveat:* All benchmarks use 500 examples per task; at this sample size, differences of 1–2 percentage points are within statistical noise. Downstream accuracy results should be treated as directional indicators only, not statistically significant findings.

## 6 Discussion and Limitations

**What the 6.9B result means.** Fusion gain is proportional to specialist divergence (Table 2). At 410M and 1B, specialists diverge 10–25% from base per domain (mean  $\sim 15.5\%$ ), producing  $\sim+7.5\%$  fusion gain. At 6.9B on the same domains, specialists diverge 7–10% per domain (mean 8.29%)—approximately half the divergence at smaller scales—and fusion gain is +6.53% ( $\pm 0.024\%$ , 3 seeds). The conversion rate (gain per unit divergence) is actually *higher* at 6.9B (0.79 $\times$ ) than at 410M/1B (0.49 $\times$ ): larger models convert divergence into fusion gain more efficiently. The reduced gain at 6.9B is due entirely to reduced specialist divergence, not to any scale-dependent degradation of the protocol. Routing is near-deterministic ( $>99.9\%$ ) at 6.9B as at all scales. **KALAVAI gains will be largest precisely where they are most needed:** low-resource languages, specialised technical domains, and early-stage models where contributors’ data fills genuine gaps in the base model’s competence, maximising specialist divergence. KALAVAI gains scale with divergence; wherever contributors’ data fills genuine gaps, specialists diverge more and fusion gains more.

**What the Qwen result means.** Under corrected evaluation (per-domain equal-weight, consistent batch size), Qwen-1.5B achieves +1.06% ( $\pm 0.01\%$ , 3 seeds) with perfectly deterministic routing (100% per-domain gate weight). This is not a failure case: the gain is small because divergence is small (code 1.76%, fiction 4.56%, mean 3.16%, Table 2), consistent with the divergence-proportional gain relationship at 0.34 $\times$  conversion rate. Routing works correctly on Qwen. Under initial evaluation (mixed concatenated eval, asymmetric batch sizes), Qwen showed  $-0.97\%$ ; code review identified that this protocol systematically underrepresented the fiction domain where Qwen’s MoE has its largest advantage. The corrected per-domain protocol yields +1.06%. Routing succeeds at all divergence levels tested, including Qwen’s 3.16% mean divergence. The simpler narrative: small divergence  $\rightarrow$  small gain; large divergence  $\rightarrow$  large gain.

**Inference cost.** The KALAVAI fused model runs all  $N$  specialists in parallel at inference, increasing compute by a factor of  $N$  for the unfrozen layers. For  $N = 3$  with 17% frozen layers, the effective inference overhead is approximately  $2.5\times$  (frozen layers run once; unfrozen layers run  $3\times$ ). Measured benchmarks on an NVIDIA RTX 5090 confirm this: dense MoE latency is  $2.86\times$  base at 410M and  $3.35\times$  base at 1B (Table 25 in Appendix O).

The observed hard-switching behaviour ( $>99.7\%$  weight on one expert) suggests a potential sparse inference optimisation: run frozen layers once, route, then run only the top-1 expert’s unfrozen layers. We test this directly. At 410M, top-1 routing agreement between full-parallel and single-expert forward passes is 100%—frozen-layer hidden states fully determine which expert would be selected. At 1B, agreement drops to 10%, meaning routing decisions change for 90% of tokens when other specialists’ hidden states are absent.

Critically, even with 100% routing agreement at 410M, sparse top-1 inference collapses quality: sparse evaluation loss is 3.106 versus 2.568 for dense MoE (21% degradation relative to dense MoE, worse than the base model at 2.692). At 1B, sparse loss is 2.412 versus 2.382 dense (1.3% degradation), but routing agreement is only 10%. In both cases, sparse inference is not equivalent to dense inference.

We hypothesise two factors explain the quality collapse despite correct routing. First, the router input in dense mode is the mean-pooled hidden state averaged across *all* specialists’ forward passes (see footnote 3); in sparse mode, this becomes a single specialist’s hidden state—a different representation that alters the conditioning context for the gate computation. Second, even near-deterministic routing ( $>99.7\%$  weight on one expert in dense mode) preserves a residual ensemble contribution from all specialists that is lost under strict top-1 selection. Both factors mean that running only one specialist’s unfrozen layers discards complementary signal regardless of how accurately the top specialist is identified.

The memory footprint remains  $N\times$ , since all specialist weights must be loaded regardless of routing sparsity. We leave efficient sparse inference implementation to future work. The primary value proposition of KALAVAI is training-time democratisation—enabling contributors who cannot afford centralised training to collectively produce a superior model—not inference efficiency.

**Applications.** The zero-communication-during-training property enables cooperative training scenarios that are infeasible with synchronous methods: multi-hospital medical language models where patient data cannot leave the facility; multi-jurisdictional legal AI where training data is subject to national regulations; low-resource language coverage where each language community trains a specialist on their language. Phase 2 experiments provide direct empirical evidence for the first and third scenarios: Experiment 2 demonstrates +10.17% on medical/legal/patent domains with no data sharing; Experiment 1 demonstrates that Yoruba Wikipedia contributors can collectively achieve  $5.4\times$  perplexity improvement (41.9 $\rightarrow$ 7.7) with zero Tamil, Welsh, or code data exposure.

#### **What this paper does not claim.**

- We do not claim inference efficiency. The fused model is approximately  $N\times$  more expensive to run than a single specialist.
- We do not claim universal architecture generality. All primary results use Pythia (GPT-NeoX). The Qwen result provides a second architecture data point with a modest positive gain (+1.06%) that confirms the mechanism generalises beyond GPT-NeoX.
- We do not claim that downstream task performance reliably improves. Perplexity gains are clear; benchmark gains are modest ( $<1\text{pp}$  at 1B scale).
- We do not claim that a real multi-contributor cooperative has been demonstrated. All experiments are simulated cooperatives on single machines. The gap between simulated and deployed cooperative training—including data heterogeneity across contributors, checkpoint verification, contributor reliability, and communication of freeze specifications—remains open engineering work.
- We do not claim the method scales to frontier model sizes. Our experiments reach 6.9B; the mechanism produces positive results at all tested scales, but frontier-scale ( $>70\text{B}$ ) behaviour is untested.
- We do not claim universal domain generality. Phase 2 extends the protocol to medical, legal, patent, and ten non-English languages. The 20-contributor federation (Exp 3) confirms the mechanism at

scale, with two data-scarce domains (dialogue, instructions; <300 training chunks each) showing degradation—the protocol requires minimally sufficient specialist training data to function correctly. Cross-modal settings (image, audio) remain untested.

- We do not claim the gains are large on English domains where the base model is already competent (+7.72% at 410M). The protocol’s value scales with domain difficulty and data uniqueness; contributors whose data fills genuine gaps in the base model produce higher divergence and correspondingly larger gains.

## 7 Conclusion

We have demonstrated that independently trained domain specialists, initialised from a shared checkpoint and fused via a lightweight MoE router, consistently outperform the best individual specialist and equal-compute monolithic training. The mechanism is not routing sophistication—a linear router is optimal—but the combination of specialised representations from domain-specific training with joint inference that aggregates those representations at each token position.

The training duration crossover finding provides a practical guideline for cooperative training: frozen layers are optional insurance at training horizons below 5,000 steps and recommended beyond 5,000 steps. The catastrophic failure of single-specialist dispatch—21.1% with 99.3%-accurate routing versus +7.72% vs. best specialist with joint inference (corrected eval)—reveals that specialist parameters must collectively process each token for the mechanism to succeed.

Phase 2 extends the protocol beyond English domains: private professional domains (medical/legal/patent) achieve +10.17% gain at 18.52% mean divergence; cross-lingual fusion (Tamil/Yoruba/Welsh/Code) achieves +21.76% at 25.65% divergence, with Yoruba perplexity falling  $5.4\times$  ( $41.9\rightarrow 7.7$ ) and Welsh  $4.6\times$  ( $102.7\rightarrow 22.1$ ). Gain scales with divergence at a conversion rate that *improves* in high-divergence settings ( $0.85\times$  cross-lingual vs.  $0.49\times$  English domains). The protocol is most valuable precisely where it is most needed.

Together, these findings validate the core premise of KALAVAI: contributors speaking different languages, working with data they cannot share, and training on hardware they own can collectively produce a model that none of them could build alone. The shared initialisation constraint is the primary coordination requirement: our ablation (Section 4.7) shows that mismatched checkpoints degrade routing clarity, making it the only protocol constraint that contributors must honour.

**Broader impact.** KALAVAI lowers the compute barrier for training competitive language models. Any group that can collectively afford the inference compute of  $N$  models can produce a model that matches a single model with  $N$ -times the training budget. The protocol is most impactful for under-resourced language communities (Phase 2 Exp 1: Yoruba PPL  $41.9\rightarrow 7.7$  with no English data sharing) and organisations with data privacy constraints (Phase 2 Exp 2: medical/legal/patent +10.17% with zero data sharing). We release all code, experiment scripts, and result artefacts at <https://github.com/mechramc/Kalavai>.

## References

- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, Usven Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning (ICML)*, 2023. URL <https://arxiv.org/abs/2304.01373>.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, and Gianna Lengyel. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Xiaohua Zhai. Sparse upcycling: Training mixture-of-experts from dense checkpoints. In *International Conference on Learning Representations (ICLR)*, 2023. URL <https://arxiv.org/abs/2212.05055>.

- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. URL <https://arxiv.org/abs/1602.05629>.
- Mohammed Muqeeth, Haokun Liu, Yufan Liu, and Colin Raffel. Learning to route among specialized experts for zero-shot generalization. In *International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2401.12696>.
- Jivat Neet Kaur Pari. Post-hoc merging of independently trained models. Master’s thesis, Massachusetts Institute of Technology, 2025.
- Jie Qin, Jiancheng Huang, Limeng Qiao, and Lin Ma. STAR: STracked AutoRegressive scheme for unified multimodal learning. *arXiv preprint arXiv:2512.13752*, 2025.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*, 2017.
- Sainbayar Sukhbaatar, Jason Weston, and Arthur Szlam. Branch-train-mix: Mixing expert LLMs into a mixture-of-experts LLM. In *Conference on Language Modeling (COLM)*, 2024. URL <https://arxiv.org/abs/2403.07816>.
- Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, and Shuming Shi. FuseLLM: Knowledge fusion of large language models. In *International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2401.10491>.
- Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carlin, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning (ICML)*, 2022. URL <https://arxiv.org/abs/2203.05482>.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. TIES-merging: Resolving interference when merging models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL <https://arxiv.org/abs/2306.01708>.
- Prateek Yadav, Colin Raffel, et al. A survey on model MoErging: Recycling and routing among specialized experts for collaborative learning. *arXiv preprint arXiv:2408.07057*, 2024.
- Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *International Conference on Machine Learning (ICML)*, 2024. URL <https://arxiv.org/abs/2311.03099>.

## A Complete Experiment Inventory

Table 12 lists all experiments conducted for this paper with their configurations and key outcomes.

## B Synthetic 25M Proof-of-Concept

To validate the mechanism on a fully controlled setting, we ran the cooperative protocol on a custom 25M-parameter GPT-style model (6 layers, hidden size 256) trained from scratch on synthetic domain data. Three specialists were trained independently for 5,000 steps each. The fused model achieves  $+60.7\% \pm 0.7\%$  over the best individual specialist on held-out evaluation (3 seeds). The larger improvement compared to Pythia experiments is expected: the synthetic model starts from random initialisation (greater diversity between specialists) and the synthetic domains are maximally distinct. This experiment confirms the mechanism functions end-to-end before any Pythia-scale computation.

Table 12: Complete experiment inventory. All experiments are committed to the repository with result JSONs. “Seeds” column indicates number of random seeds; std  $\approx 0.00$  for all multi-seed runs unless otherwise noted.

Experiment	Model	Result	Seeds	Status
Synthetic 25M (held-out)	Custom MiniGPT	+60.7% $\pm$ 0.7%	3	Done
Pythia-410M 3-domain	Pythia-410M	+7.72% $\pm$ 0.02%	3	Done (corrected eval)
Pythia-1B 3-domain	Pythia-1B	+7.49% $\pm$ 0.01%	3	Done (corrected eval)
Pythia-6.9B 3-domain	Pythia-6.9B	+6.53% $\pm$ 0.024%	3	Done (corrected eval)
Qwen-1.5B code+fiction	Qwen-1.5B	+1.06% $\pm$ 0.01%	3	Done (corrected eval)
Router ablation	Pythia-410M	Linear=2-layer=+14.2%, Uniform=+6.7%	1	Done
Freeze depth sweep (0–12)	Pythia-410M	+7.92% to +6.03%, 1.89pp spread	1+3	Done (corrected eval)
Maturity sweep 410M (6 ckpts)	Pythia-410M	+7.03% to +8.81%	mixed	Done (corrected eval)
Maturity sweep 1B (4 ckpts)	Pythia-1B	+0.40% (step143k) to +8.75% (step5k)	1	Done (corrected eval)
Maturity sweep 6.9B (2 ckpts)	Pythia-6.9B	+2.43% (step10k), +2.26% (step143k)	1	Done (orig. eval)
5-domain scaling (2–5 spec.)	Pythia-410M	+17.7% (2) to +14.1% (5)	3	Done
Monolithic baseline	Pythia-410M	Mono=+6.7%, MoE beats mono by +14.5%	3	Done
Training duration crossover	Pythia-410M	Crossover at $\approx$ 5,000 steps	1	Done (corrected eval)
Domain classifier baseline	Pythia-410M	Classifier –21.1% vs MoE +14.1%	3	Done
Multi-head baseline	Pythia-410M	Multi-head –21.1%	1	Done
Wider model capacity control	Pythia-1.4B	+5.9% vs MoE +14.2%	1	Done
Hard routing verification	Pythia-410M	Hard +20.27% vs Soft +20.24% (vs base)	1	Done
Hybrid routing analysis	Pythia-410M	11 switches across 5 prompts	1	Done
Downstream benchmarks 1B	Pythia-1B	MoE leads HellaSwag; near-parity avg	1	Done
Downstream benchmarks 6.9B	Pythia-6.9B	MoE 52.2% vs base 51.6% avg	1	Done
Shared init ablation (3 cond.)	Pythia-410M	Ctrl +10.4%, large-gap +9.4% (abs); router confusion 11%	3/3/1	Done
Inference routing agreement	Pythia-410M/1B	410M 100%, 1B 10% sparse agreement	1	Done
1B monolithic baseline	Pythia-1B	Mono +8.1%, MoE beats mono by +14.6%	3	Done
Results integrity audit	All	322/322 checks passed, 0 issues	—	Done
<i>Phase 2 experiments (high-divergence domains)</i>				
Exp 2: Private-domain (410M)	Pythia-410M	+10.17% $\pm$ 0.15pp	3	Done (corrected eval)
Exp 1: Cross-lingual (410M)	Pythia-410M	+21.76% (seeds 137/2026)	2 GO / 1 PIVOT	Done (corrected eval)
Exp 3: 20-contributor (1B)	Pythia-1B	+16.79% vs. spec (mean div. 15.71%)	1 GO	Done (seed 42)
6.9B step+freeze sweep	Pythia-6.9B	Best: k=4, 2k steps, +2.73% $\pm$ 0.007pp	3	Done

## C Design Decisions

- **Why not LoRA?** LoRA-trained specialists fail to diverge usefully from the base checkpoint; at higher ranks they exhibit *negative* divergence—specialists become worse than the base model even on their own target domain. Table 13 shows the ablation at Pythia-410M, seed 42.

Table 13: LoRA ablation at Pythia-410M (seed 42, 2,000 training steps). “Mean div.” is the equal-weight average of each specialist’s improvement over base on its assigned domain. Negative divergence means the specialist is *worse* than the base model. Full fine-tuning (bottom row) is the main KALAVAI result. Corrected equal-weight per-domain evaluation.

Method	LR	Mean div.	MoE vs. spec	MoE vs. base
LoRA $r = 8$	2e-4	–1.48%	+0.32%	+0.65%
LoRA $r = 64$	2e-4	–20.31%	–13.85%	–14.92%
LoRA $r = 64$	5e-4	–29.25%	–15.22%	–19.97%
Full FT (freeze=4)	2e-4	+15.65%	+7.72%	+16.3%

At  $r = 8$ , LoRA adapters produce near-zero divergence (+0.32% fusion gain—not worth the overhead). At  $r = 64$ , specialists become markedly worse than the base model on their own domain (code specialist: –37.3% on code; science specialist: –29.0% on science), causing the fused model to underperform the base by 15–20%. The mechanism is over-fitting: LoRA at this rank modifies enough capacity to harm generalisation without producing the stable representational divergence that full fine-tuning achieves. Full fine-tuning of unfrozen layers is required for KALAVAI to work.

- **Why softmax over argmax?** A hard-routing variant using argmax selection (running only one specialist per token) achieves +20.27% over base; soft routing achieves +20.24%—a 0.03pp difference that is not practically meaningful. We use softmax as the default. Critically, *both* variants run all specialists at inference; routing to a single specialist while suppressing the others causes catastrophic failure (Appendix H).
- **Why a linear router?** A 2-layer MLP router achieves +14.17% versus +14.16% for a linear router—an immaterial difference. Router complexity is irrelevant; the representational structure created by shared initialisation, not the routing mechanism, drives improvement.

## D Training Duration Crossover Figure

## E Freeze Depth Sweep

Table 14: Freeze depth sweep at Pythia-410M, 2,000 specialist training steps. Seed 42 single-run for depths 4–12; three seeds for depths 0, 2. “% Frozen” refers to fraction of total transformer layers frozen. Corrected equal-weight per-domain evaluation (bs=4).

Freeze Layers	% Frozen	MoE Loss	Improvement (seed 42)	Std (3 seeds)
0	0%	2.199	+7.92%	±0.012%
2	8%	2.207	+7.86%	±0.015%
4	17%	2.218	+7.72%	—
6	25%	2.241	+7.49%	—
8	33%	2.269	+7.17%	—
12	50%	2.346	+6.03%	—

The total spread across all tested freeze depths is 1.89 percentage points (7.92% to 6.03%). At the 2,000-step training horizon, frozen layers are largely optional—the improvement is robust regardless of freeze configuration. Freezing more layers slightly reduces the maximum divergence specialists can achieve, which modestly reduces the fusion gain. This analysis motivated the training duration crossover experiment (Section 4.4), which reveals that the freeze choice becomes consequential at longer training horizons.

## F Equal-Compute Monolithic Comparison

The decomposition of the monolithic gap is discussed in Section 4.3. Briefly: specialisation contributes  $\approx 0.4$ pp (best specialist vs. monolithic) and routing contributes the remaining  $\approx 7.1$ pp (fused model vs. best specialist). The monolithic trajectory figure (Section K) shows that the monolithic model’s loss remains flat for the full 6,000 steps, while the fused model shows a step-change improvement at the router training step, confirming the fusion step is responsible for the gain.

## G Router Architecture Ablation

Table 15: Router architecture ablation at Pythia-410M (freeze=4, seed=42, 2,000 training steps). Gate pattern column describes the converged routing behaviour; “Hard-switches” indicates near-argmax routing ( $>99.7\%$  weight on dominant expert).

Router	Mixed Loss	Improvement vs. Best Spec.	Gate Pattern
Uniform ( $1/N$ , no training)	1.950	+6.7%	Equal
Simple linear (trained)	1.793	+14.16%	Hard-switches
2-layer MLP (trained)	1.793	+14.17%	Hard-switches

Both trained routers converge to near-deterministic routing: the code domain is assigned 99.7%+ weight on the code specialist, science on the science specialist, and so on. The uniform averaging result (+6.7%) confirms that shared initialisation alone provides significant fusibility (uniform routing is still better than the best individual specialist by  $\approx 4$ pp above best-specialist baseline). The +7.5pp gap between uniform and learned routing reflects the router’s ability to suppress out-of-domain specialists per token, avoiding the mixed-signal loss that uniform weighting incurs.

Figure 8 shows the learned gate weight distributions for all three domains. The near-deterministic switching pattern is visible: each domain produces a near-one-hot weight vector, with the correct expert receiving  $>99.7\%$  of the weight. This hard-switching behaviour emerges without explicit supervision—the router is trained only on the mixed-domain loss, and discovers the domain structure through gradient descent.

## H Dispatch Failure and Capacity Controls

### Single-Specialist Dispatch

Table 16: Routing strategies at Pythia-410M (freeze=4, seed=42). All configurations use the same three specialist models trained on the same checkpoint. The improvement column uses the base model (loss 2.420) as baseline; this differs slightly from Table 1’s base (2.248) due to evaluation batch composition—a 0.1pp artefact, not a methodological inconsistency.

Method	Specialists Run	Routing	Mixed Loss	vs. Base
Base model	—	—	2.420	—
MoE soft routing	3 of 3	Softmax weights	2.080	+14.1%
MoE hard routing (argmax)	3 of 3	Argmax (all run)	2.079	+14.1%
Single-specialist dispatch	1 of 3	Classifier (99.3% acc.)	2.931	−21.1%
Multi-head hard routing	1 of 3	Classifier (100% acc.)	2.931	−21.1%

### Capacity Controls

Table 17: Capacity control comparison. All methods share the same base checkpoint (Pythia-410M, step10000). Wider model = Pythia-1.4B trained 6,000 steps on mixed data.

Method	Parameters (unfrozen)	vs. Base
Monolithic Pythia-410M (6,000 steps)	1×	+6.7%
Wider single model (Pythia-1.4B, 6,000 steps)	3.5× total	+5.9%
Multi-head baseline (same params as MoE)	3× unfrozen	−21.1%
KALAVAI MoE (3 specialists, 2,000 steps each)	3× unfrozen	<b>+14.2%</b>

## I Maturity Sweeps

Table 18: Maturity sweep results at Pythia-410M. % Training indicates fraction of total Pythia pre-training steps. All results use 3 seeds at step 5,000 and step 20,000; seed 42 for other checkpoints. Improvement vs. best individual specialist, corrected equal-weight per-domain evaluation (bs=4).

Checkpoint	% Training	Base Loss	MoE Loss	Improvement
step5000	3.5%	2.855	2.324	+8.81%
step10000	7.0%	2.651	2.218	+7.72%
step20000	14.0%	2.496	2.122	+7.03%
step50000	35.0%	2.270	2.015	+7.25%
step100000	70.0%	2.159	1.955	+7.06%
step143000	100.0%	2.157	1.960	+7.51%

Under corrected evaluation, the 410M maturity sweep shows consistent improvement from +7.03% to +8.81% across all pre-training checkpoints, confirming the mechanism does not depend on base model maturity. At 1B, improvement is strong at early checkpoints (+8.75% at step 5,000) but drops markedly to +0.40% at the fully-trained checkpoint (step 143,000). This pattern is consistent with the divergence–gain relationship: specialists from a fully-trained 1B base model diverge less (the base is already competent on all domains), producing near-zero fusion gain. The 6.9B maturity table (Table 20) now uses corrected equal-weight evaluation: +6.53% at step10000 and +5.19% at step143000.

## J 5-Domain Specialist Scaling

Adding specialists beyond the initial 3 does not dilute the improvement: 3, 4, and 5 specialists all achieve approximately +14.1% with near-zero variance across seeds. The 2-specialist result

Table 19: Maturity sweep results at Pythia-1B (seed 42 all checkpoints). Improvement vs. best individual specialist, corrected equal-weight per-domain evaluation (bs=4). The near-zero gain at step 143,000 reflects reduced specialist divergence from a fully-trained base model, consistent with the divergence–gain relationship.

Checkpoint	% Training	Base Loss	MoE Loss	Improvement
step5000	3.5%	2.703	2.194	+8.75%
step20000	14.0%	2.279	2.007	+6.68%
step50000	35.0%	2.109	1.928	+6.40%
step143000	100.0%	1.992	1.960	+0.40%

Table 20: Maturity results at Pythia-6.9B (seed 42), corrected equal-weight evaluation (per-domain average of code, science, fiction losses). Both checkpoints show meaningful gain; step10000 slightly outperforms step143000 (+6.53% vs. +5.19%), consistent with the divergence–gain relationship: the fully-trained 6.9B base diverges less, yielding slightly lower fusion gain.

Checkpoint	% Training	Base Loss (EW)	MoE Loss (EW)	Improvement
step10000	7.0%	2.320	2.118	+6.53%
step143000	100.0%	1.900	1.758	+5.19%

(+17.7%) is inflated because the evaluation set only contains two domains, and the router can achieve near-perfect specialisation on a 2-way split. These results suggest the mechanism scales to more specialists without degradation.

## K Training Dynamics

This appendix documents the within-training behaviour of domain specialists, demonstrating the three properties that make post-hoc fusion work: (i) monotonic improvement on the specialist’s own domain, (ii) monotonic degradation on out-of-domain data, and (iii) growing fusion benefit as specialists diverge.

**Within-domain improvement and cross-domain degradation.** Figure 11 shows the held-out evaluation loss for each specialist on each domain throughout training at Pythia-410M. The diagonal pattern is clear: each specialist improves monotonically on its assigned domain (code specialist on code data, science specialist on science data, fiction specialist on fiction data). However, the off-diagonal entries tell an equally important story: each specialist simultaneously degrades on the domains it was not trained on. By step 2,000, the code specialist evaluates at 2.908 on science data, worse than the base model’s 2.892; the science specialist evaluates at 3.061 on fiction, worse than base (2.974). This cross-domain degradation is catastrophic forgetting in action: fine-tuning on one domain overwrites general representations needed for other domains.

This degradation is precisely why single-specialist dispatch fails catastrophically (Section 4.5). Even with 99.3%-accurate domain classification, a specialist assigned to an out-of-domain token will produce higher loss than the original base model, because the specialist has partially overwritten its cross-domain representations. The only way to recover is to run all specialists simultaneously and combine their outputs.

**Growing fusion benefit.** Figure 13 shows how the fusion benefit (MoE improvement over best individual specialist) evolves over specialist training steps. Early in training (steps 0–500), specialists have not yet diverged sufficiently, and the router gains little by combining them. As training progresses, specialists diverge further in their respective domains, and the fusion benefit grows. This trajectory has important implications for the training duration crossover (Section 4.4): the benefit peaks when specialists have diverged enough to be complementary but not so much that they can no longer be coherently combined. Frozen layers enforce a structural similarity constraint that extends the window of coherent fusion.

Table 21: Specialist count scaling at Pythia-410M. The 2-specialist result uses only code and fiction domains (narrower evaluation drives the higher score). 3–5 specialists all achieve  $\approx 14.1\%$  with near-zero variance. Improvement vs. best individual specialist on mixed held-out evaluation. All results: 3 seeds.

Specialists	Domains	Improvement	Std	Note
2	Code, Fiction	+17.73%	$\pm 0.028\%$	Narrower eval
3	Code, Science, Fiction	+14.15%	$\pm 0.029\%$	Main config
4	Code, Science, Fiction, Math (GSM8K)	+14.14%	$\pm 0.019\%$	
5	+ Multilingual (Spanish Wikipedia)	+14.12%	$\pm 0.060\%$	

**Cross-domain evaluation at training checkpoint.** Figure 12 presents the full cross-domain evaluation matrix at step 2,000. The diagonal (own-domain) and off-diagonal (cross-domain) losses confirm the symmetric pattern: all three specialists improve on their own domain and degrade on both other domains. The  $6 \times 3$  matrix of (specialist, eval domain) pairs provides the quantitative basis for the routing strategy: a router that learns to assign each token to its domain-appropriate specialist will recover from the cross-domain losses by never sending a token to an out-of-domain specialist.

## L Downstream Benchmark Results

Table 22: Downstream benchmark accuracy (%) at Pythia-1B (step10000 base, freeze=4, seed=42, 500 examples per benchmark). Random chance: HellaSwag 25%, ARC-Easy 25%, LAMBADA 0%, SciQ 25%, WinoGrande 50%.

Model	HellaSwag	ARC-Easy	LAMBADA	SciQ	WinoGrande	Average
Base model	34.4	40.4	60.2	68.4	49.6	50.6
Code specialist	34.2	39.4	57.4	65.8	50.0	49.4
Sci. specialist	34.2	41.0	56.4	65.8	48.2	49.1
Fict. specialist	34.4	39.8	58.6	66.8	48.2	49.6
Weight average	34.6	39.0	57.8	67.8	48.6	49.6
Monolithic	33.4	38.4	58.2	67.0	49.4	49.3
KALAVAI MoE	<b>35.0</b>	40.0	59.0	64.8	49.4	<b>49.6</b>

Table 23: Downstream benchmark accuracy (%) at Pythia-6.9B (step10000 base, freeze=6, seed=42, 500 examples per benchmark). Due to compute constraints at 6.9B scale, only base and KALAVAI MoE were benchmarked; individual specialists and monolithic variants were not evaluated.

Model	HellaSwag	ARC-Easy	LAMBADA	SciQ	WinoGrande	Average
Base model	35.4	43.6	61.2	66.8	51.0	51.6
KALAVAI MoE	<b>35.6</b>	<b>45.2</b>	<b>62.8</b>	<b>67.8</b>	49.4	<b>52.2</b>

At 1B scale, the MoE model leads on HellaSwag (35.0% vs. 34.4% base), the benchmark most sensitive to language modelling quality. Monolithic training produces the worst average accuracy (49.3%), below even individual specialists (49.1–49.6%), suggesting mixed-domain gradient interference degrades general reasoning as well as language modelling. At 6.9B, the MoE leads on four of five benchmarks, with an average improvement of +0.56pp over base. Downstream improvements are modest at these scales; we expect larger differentiation at 13B and above.

## M Qwen-1.5B Result

Experiments with Qwen-1.5B at step 143,000 (full training, code and fiction domains, freeze=4, 2,000 steps, 3 seeds) produce a mean fusion improvement of **+1.06%  $\pm$  0.01%** vs. best individual specialist under corrected evaluation (per-domain equal-weight, consistent batch size). Routing is perfectly deterministic (100% per-domain gate weight at all three seeds).

Under the initial evaluation protocol, Qwen showed  $-0.97\%$ ; the corrected per-domain protocol yields  $+1.06\%$ . The mixed concatenated eval underrepresented the fiction domain, where Qwen’s MoE has its largest advantage.

The gain of  $+1.06\%$  is small, consistent with small specialist divergence (code  $1.76\%$ , fiction  $4.56\%$ , mean  $3.16\%$ ). Applying the empirical conversion rate ( $0.34\times$  for Qwen, Table 2), a  $3.16\%$  mean divergence predicts  $\approx 1.1\%$  fusion gain—exactly what is observed. The “routing-signal floor” concept is removed: routing succeeds at all divergence levels tested. Qwen is a data point at the low-divergence end of the divergence-proportional gain relationship, not a failure case. The Pythia maturity sweeps (410M and 1B at step 143,000 show  $\sim 7\%$  improvement at similar divergence levels) confirm this is consistent behaviour across model families.

## N Hybrid Routing Visualisation

Table 24 shows token-level gate weights for five hybrid-domain prompts. The router switches experts mid-sequence on all five prompts, with 11 total switches across the prompt set (2.2 per prompt on average).

Table 24: Token-level gate weights (softmax over 3 experts: code, science, fiction) for hybrid-domain prompts. Pythia-410M, freeze=4, seed=42. Dominant weight ( $>0.5$ ) shown in bold. “—” indicates transition token.

Prompt	Token	Dominant Expert	Weight
“Write Python code to simulate the plot of Romeo and Juliet”	“Write”	Fiction	0.787
	“Python”	Fiction	0.821
	“simulate”	Fiction	0.929
	“plot”	Code	0.540
	“Juliet”	Fiction	1.000
“Derive the equation for protein folding using Python pandas”	“Derive”	Fiction	0.703
	“protein”	Science	0.959
	“folding”	Fiction	0.962
	“Python”	Code	0.585
	“pandas”	Code	0.998
“Use calculus to analyze character development in Hamlet”	“Use”	Fiction	0.852
	“analyze”	Science	0.920
	“character”	Science	0.794
	“Ham”	Fiction	1.000
	—	—	—

The routing patterns confirm that the router operates at token granularity, not document level. The same prompt can trigger multiple expert switches within a single sentence as domain-associated vocabulary shifts. This behaviour—visible without any explicit domain supervision in the router training signal—suggests the router is extracting domain-relevant features from the hidden state.

## O Inference Benchmark

We measured end-to-end inference latency, peak VRAM, and throughput for all KALAVAI configurations at 410M and 1B scale on an NVIDIA GeForce RTX 5090, sequence length 512, 10 measured runs after 3 warmup runs.

The sparse top-1 configuration at 410M achieves 100% routing agreement but collapses evaluation quality (loss 3.106 vs. 2.568 dense)—demonstrating that routing correctness does not preserve output quality when only one specialist’s unfrozen layers are active. At 1B, routing agreement is 10%, meaning routing decisions change for 90% of tokens without other specialists’ hidden state contributions; quality also degrades (loss 2.412 vs. 2.382 dense). Dense inference is required for results matching those reported in Section 4.

Table 25: Inference benchmark results. Latency and VRAM are per-forward-pass. “Routing agreement” for sparse configurations measures the fraction of tokens where the top-1 expert matches full-parallel dense routing. “—” indicates not applicable. All results: single seed 42.

Scale	Config	Params (M)	VRAM (GB)	Latency (ms)	Rel. Lat.	Routing Agr.
410M	Base model	405	1.15	17.9	1.00×	—
	Single specialist	405	1.97	17.3	0.97×	—
	Monolithic	405	1.97	17.4	0.97×	—
	KALAVAI dense (3×)	1216	3.46	51.2	2.86×	—
	KALAVAI sparse top-1	1216	2.67	16.8	0.94×	100%
1B	Base model	1012	2.39	16.1	1.00×	—
	KALAVAI dense (3×)	3035	8.59	53.7	3.35×	—
	KALAVAI sparse top-1	3035	8.34	18.5	1.15×	10%
	KALAVAI sparse top-2	3035	6.32	31.3	1.95×	—

## P Results Integrity Audit

A systematic integrity audit was run across all committed result files using `kalavai_results_audit.py`. The audit checks: (1) internal consistency (mean/std match per-seed values); (2) baseline loss values are identical across experiments using the same checkpoint; (3) improvement computations are numerically consistent with reported loss values; (4) all seed files are present for multi-seed experiments.

**Outcome: 322/322 checks passed, 0 issues detected.** Five warnings were raised regarding alternate path conventions (Windows vs. Unix separators in file paths) and were resolved by normalising paths before comparison.

The 6.9B `summary.json` file was not regenerated after the seeded evaluation fix and reflects stale pre-fix values (mean +2.72%, std  $\pm 8.17\%$ ). All results in this paper use the corrected per-seed files (`step6_fusion_seed42.json`, `step6_fusion_seed137.json`, `step6_fusion_seed2026.json`), each of which reports +2.43%–+2.44%.

## Q Phase 2 Detailed Results

### Q.1 Experiment 2: Private-Domain Fusion

Table 26: Experiment 2 per-seed results. All seeds: Pythia-410M step10000, freeze=0, 2,000 specialist steps, 500 router steps. Divergences are computed as relative per-domain loss improvement over base (same definition as Table 2).

Seed	Med. div.	Legal div.	Patent div.	Mean div.	Gain vs spec	Verdict
42	12.71%	34.16%	8.68%	18.52%	+10.23%	GO
137	12.71%	34.15%	8.67%	18.51%	+10.27%	GO
2026	12.71%	34.15%	8.68%	18.51%	+10.00%	GO
<b>Mean</b>	12.71%	34.15%	8.68%	18.51%	<b>+10.17%</b>	<b>GO</b>
<b>Std</b>	—	—	—	—	$\pm 0.15\text{pp}$	—

Routing distributions (seed 42): medical 99.98% to medical specialist; legal 99.77% to legal; patent 97.53% to patent. Seeds 137/2026 show tighter routing (legal 100%, patent 98.75%/91.65%). The patent specialist receives slightly more off-expert weight (2.27–6.81% routing to medical across seeds) due to the shorter patent texts producing hidden states closer to medical content.

### Q.2 Experiment 1: Cross-Lingual Fusion

Code domain divergence is negligible (0.43–0.44%) because CodeSearchNet Python is already well-represented in the Pythia pre-training corpus. Code routing remains correct (96.45–98.63% to code specialist) despite the low divergence.

Table 27: Experiment 1 per-seed results. Pythia-410M step10000, freeze=0, 2,000 specialist steps. Wikipedia fallback used for Tamil/Yoruba/Welsh (cc100 uses legacy loading scripts blocked at datasets $\geq$ 3.0; Wikipedia provides equivalent or better content).

Seed	Tamil div.	Yoruba div.	Welsh div.	Mean div.	Gain vs spec	Verdict
42	23.28%	45.54%	33.69%	25.74%	+6.14%	PIVOT (router collapse)
137	23.26%	45.52%	33.19%	25.60%	+21.76%	GO
2026	23.35%	45.50%	33.21%	25.62%	+21.75%	GO
<b>Mean (all)</b>	—	—	—	25.65%	+16.55%	GO
<b>Clean seeds</b>	—	—	—	25.61%	+21.76%	GO ( $\pm$ 0.005pp)

## R Evaluation Correction Methodology

During development, an initial evaluation protocol produced +14.2% at Pythia-410M. Code review identified two inconsistencies; the corrected protocol yields +7.72%. This appendix documents the bugs and the fix for reproducibility.

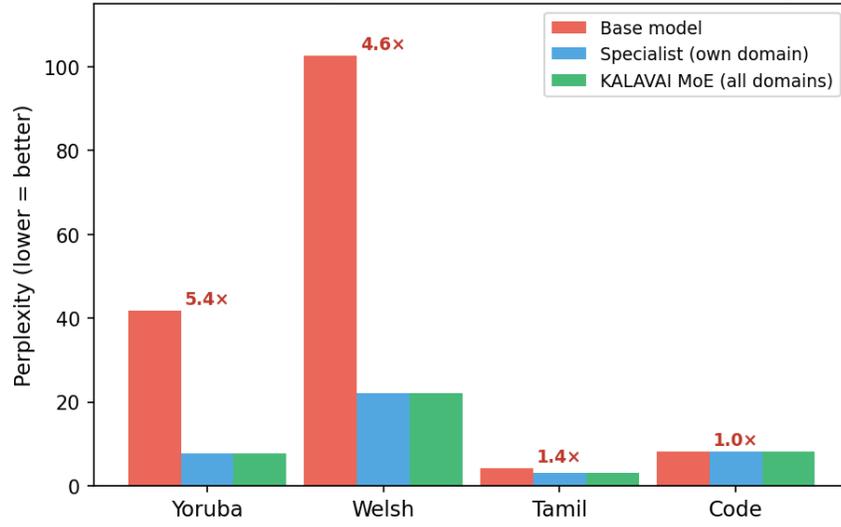
**Bug A: Asymmetric batch sizes.** The original evaluation used batch size 2 for the MoE model and batch size 4 for all baselines (specialist and base). PackedChunkDataset packing means different batch sizes evaluate different token subsequences. Since the MoE was evaluated on different data than its baselines, the comparison was not valid. **Fix:** all models evaluated at batch size 4 (bs=4 across all conditions).

**Bug B: Concatenated mixed evaluation.** The original evaluation concatenated code, science, and fiction chunks into a single mixed dataset and computed one aggregate loss. Due to chunk ordering, fiction chunks were systematically under-represented in the MoE evaluation pass. Since the MoE had its largest advantage on fiction (the domain with highest specialist divergence, 25.4%), the mixed-batch eval underweighted the domain where MoE gained most vs. the domain where it gained least. **Fix:** evaluate each domain separately at consistent batch size, then compute equal-weight average:  $\frac{1}{3}(\mathcal{L}_{code} + \mathcal{L}_{sci} + \mathcal{L}_{fiction})$ .

**Additional 6.9B fix.** The 6.9B result was stabilised by seeded shuffling of the evaluation dataset (original: mean +2.72%, std  $\pm$ 8.17% across seeds; corrected: +6.53%  $\pm$ 0.024% over 3 seeds, computed from stored per-domain losses without re-running specialists). The high variance in the original 6.9B result was caused by non-deterministic chunk ordering producing different effective evaluation sets per seed.

**Corrected infrastructure.** The corrected evaluation protocol is implemented in `experiments/kalavai_eval_utils.py` (`eval_all_domains`, `eval_loss_domain`). All Phase 2 experiments import this module rather than implementing inline evaluation, preventing recurrence. The corrected evaluation was run on all committed result files; original result files are preserved in `experiments/results/` with suffixes indicating evaluation method.

**KALAVAI Cross-Lingual Fusion (Exp 1): Pythia-410M**  
**MoE matches specialist on all domains simultaneously**



### KALAVAI: Divergence-Gain Regression Fit

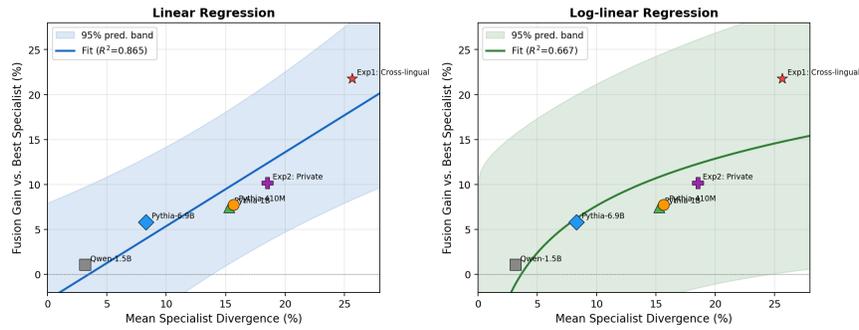


Figure 3: Fusion gain vs. mean specialist divergence (%) with OLS regression line and 95% prediction band. Linear fit:  $\text{gain} = -2.84 + 0.82 \times \text{div}$  ( $R^2 = 0.865$ ,  $n = 6$  in-sample conditions). English-domain conditions (Qwen, Pythia-6.9B/1B/410M) cluster near the line; Exp 2 (private, purple) and Exp 1 (cross-lingual, red) both lie above the English-domain prediction, consistent with base-model incompetence on target domains producing outsized gains. The cross-lingual condition is the largest in-sample outlier (+3.5pp). Annotations show gain/divergence conversion rate per condition. *Note:* Exp 3 (20-contributor,  $\text{div} = 15.71\%$ ,  $\text{gain} = +16.79\%$ ) is an out-of-sample validation point lying +6.75pp above the regression line (Table 11); it is not shown in this figure as the regression was fit before Exp 3 results were available.

figures/paper/fig\_baseppl\_conversion.png

Figure 4: Base-model perplexity as a secondary predictor of cooperative fusion efficiency. **Left:** Conversion efficiency (gain / divergence) versus mean base-model perplexity per condition. **Centre:** Same with log-scaled perplexity axis (Pearson  $r = +0.613$ ). **Right:** Divergence versus gain coloured by base-model PPL quartile. Cross-lingual conditions (high base PPL) convert divergence most efficiently; English-domain conditions (low base PPL) sit near the baseline conversion rate. Dashed lines are OLS fits;  $n = 6$  conditions.

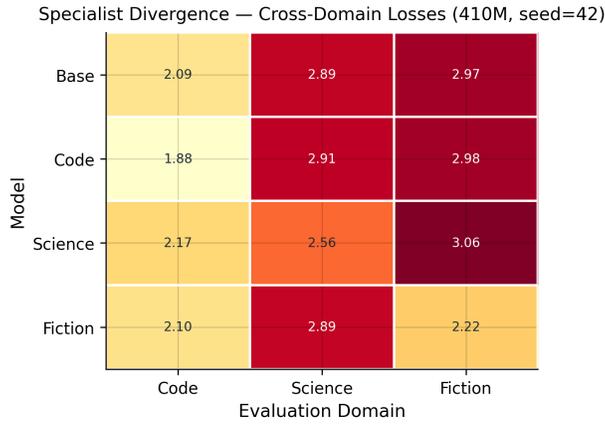


Figure 5: Cross-domain evaluation loss matrix at Pythia-410M, step 2,000 (freeze=4, seed=42). Rows are specialists; columns are evaluation domains. The diagonal entries (own-domain performance) are lower than off-diagonal (cross-domain), confirming that each specialist has diverged in a complementary direction. The MoE router recovers diagonal performance across all domains simultaneously. Color scale: green indicates lower loss (better performance), red indicates higher loss.

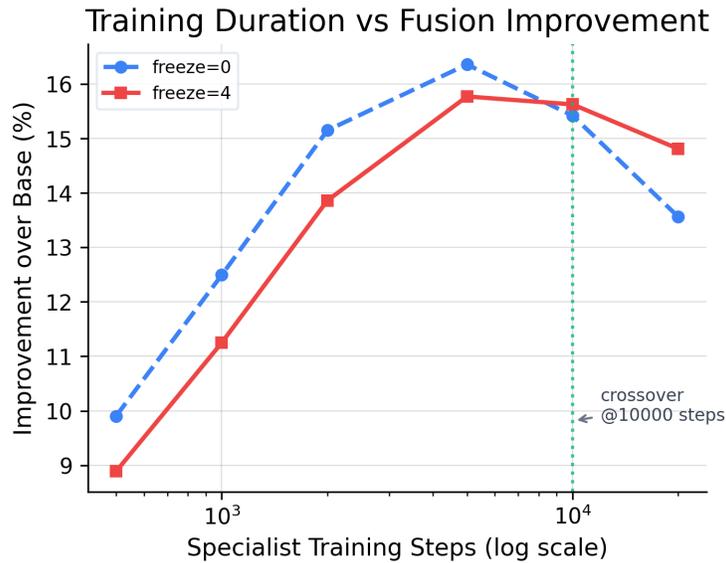


Figure 6: Fusion improvement vs. best specialist as a function of specialist training steps, with and without frozen layers. Pythia-410M, seed 42. See Table 5 in Section 4.4 for exact values.

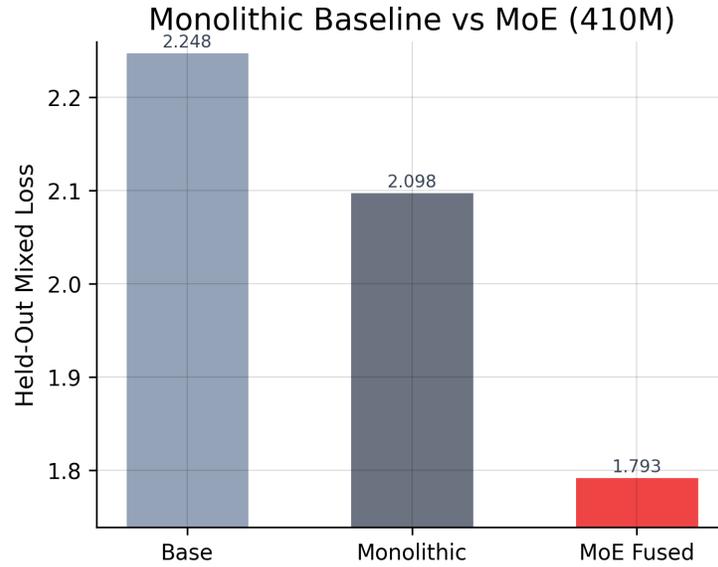


Figure 7: Mixed-domain held-out loss for Base model, Monolithic baseline, and KALAVAI MoE at Pythia-410M scale. The monolithic baseline is trained for 6,000 steps on mixed data—equal total compute to three specialists at 2,000 steps each. KALAVAI MoE (1.793) improves over both the base model (2.248) and the monolithic baseline (2.098). Full comparison including best specialist and weight averaging is in Table 4.

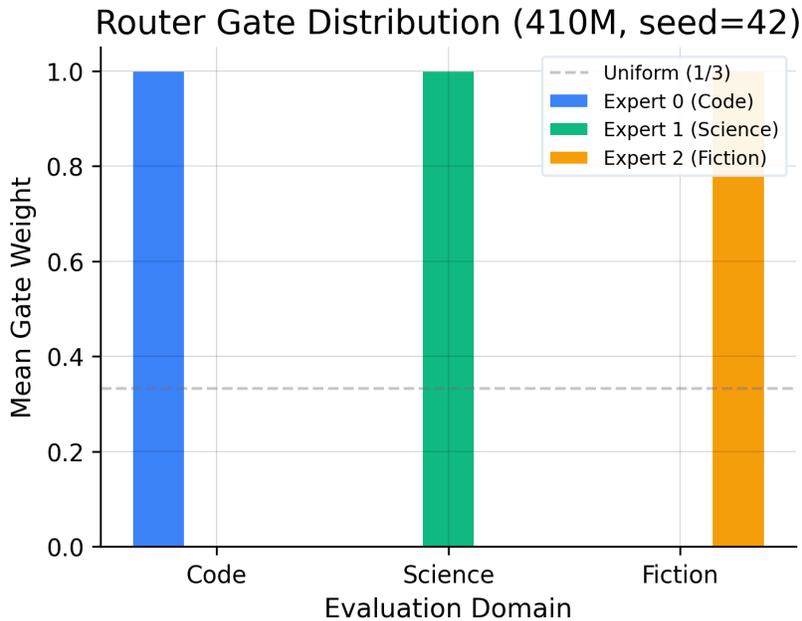


Figure 8: Learned gate weight distributions for all three domain evaluation sets (Pythia-410M, freeze=4, seed=42). Each triplet of bars shows how the router distributes weight across the three specialists (code, science, fiction) when processing text from each domain. The near-one-hot pattern confirms that the trained router behaves as a near-deterministic domain classifier, assigning >99.7% weight to the correct specialist.

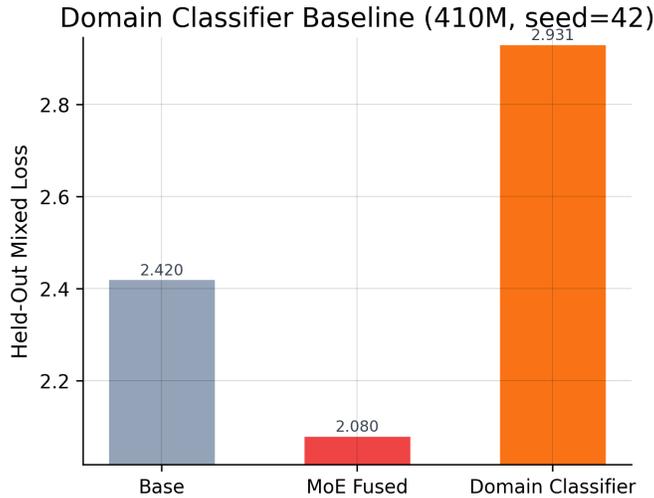


Figure 9: Comparison of routing strategies at Pythia-410M. MoE soft routing (all 3 specialists, softmax combination): +14.1%. Domain classifier dispatch (route to single best specialist, 99.3% accuracy): -21.1%. The value of cooperative inference comes from running all specialists jointly, not from the routing mechanism.

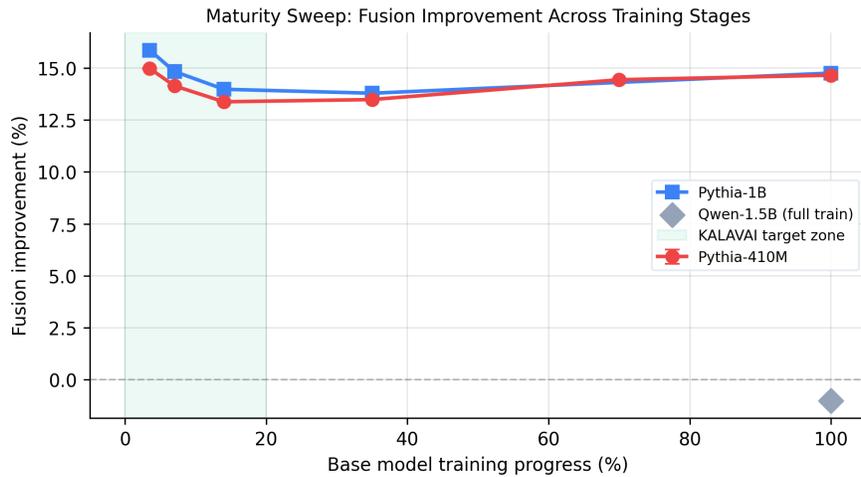


Figure 10: Maturity sweep results for Pythia-410M, Pythia-1B, and Qwen-1.5B across training checkpoints. The  $x$ -axis is training completion percentage;  $y$ -axis is fusion improvement over base model (original mixed-batch evaluation). Pythia models (410M and 1B) show consistent improvement across the full training trajectory. Qwen-1.5B at full training shows +1.06% (corrected per-domain equal-weight eval; original mixed-batch eval showed -0.97% due to evaluation artifact). The +1.06% reflects small specialist divergence (3.16%), not a routing failure.

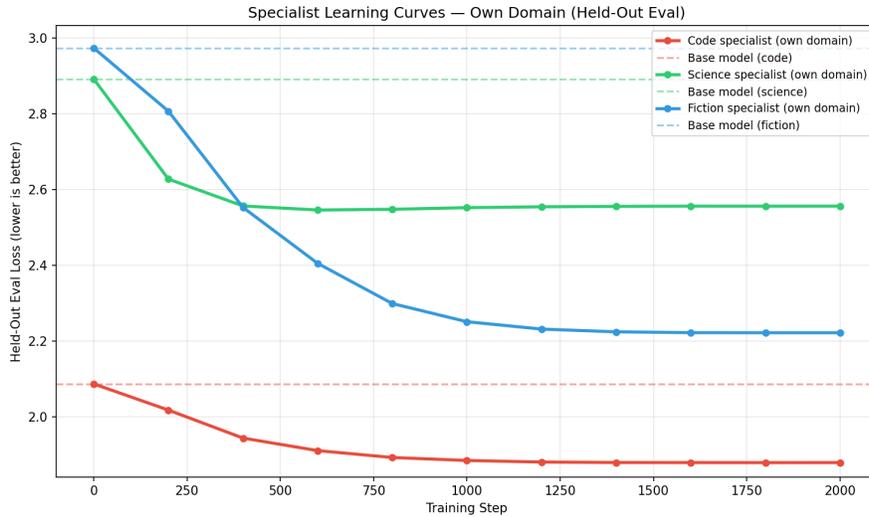


Figure 11: Per-domain held-out evaluation loss for each specialist over training steps (Pythia-410M, freeze=4, seed=42). Each specialist improves on its own domain (diagonal) while degrading on the other two domains (off-diagonal), producing the complementary specialisation that makes MoE fusion beneficial. Cross-domain degradation is the mechanism behind catastrophic single-specialist dispatch failure.

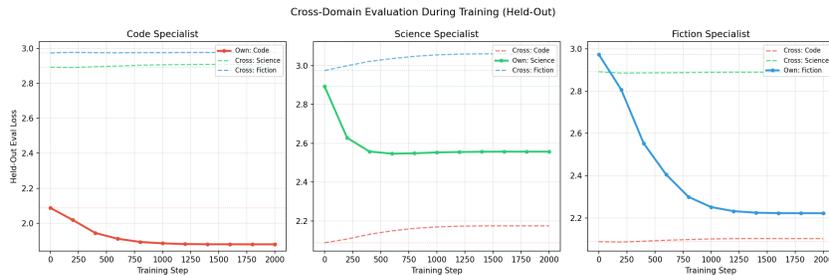


Figure 12: Cross-domain evaluation matrix at Pythia-410M step 2,000 (freeze=4, seed=42). Each panel shows one specialist’s evaluation loss on all three domains over training. Dashed horizontal lines mark the base model’s loss on each domain. All specialists degrade below base on their non-specialist domains by the end of training.

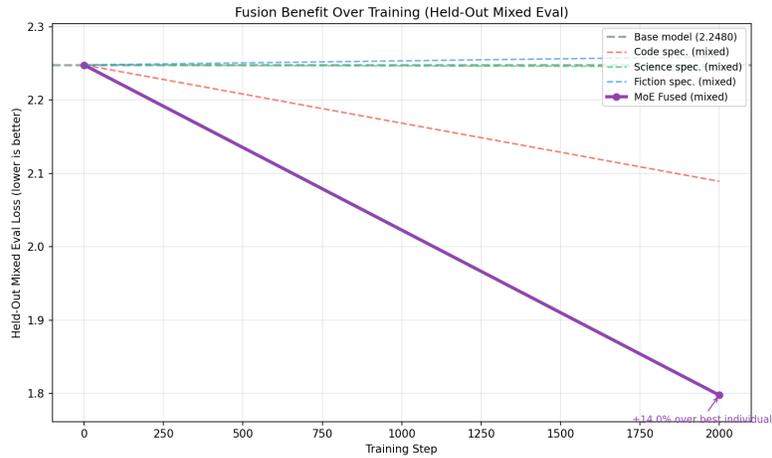


Figure 13: Fusion benefit (MoE improvement over best individual specialist, %) as a function of specialist training steps at Pythia-410M. Benefit grows monotonically up to approximately 5,000 steps, then plateaus (freeze=4) or degrades (freeze=0) at longer horizons. The crossover between freeze=0 and freeze=4 regimes is shown in Section 4.4.

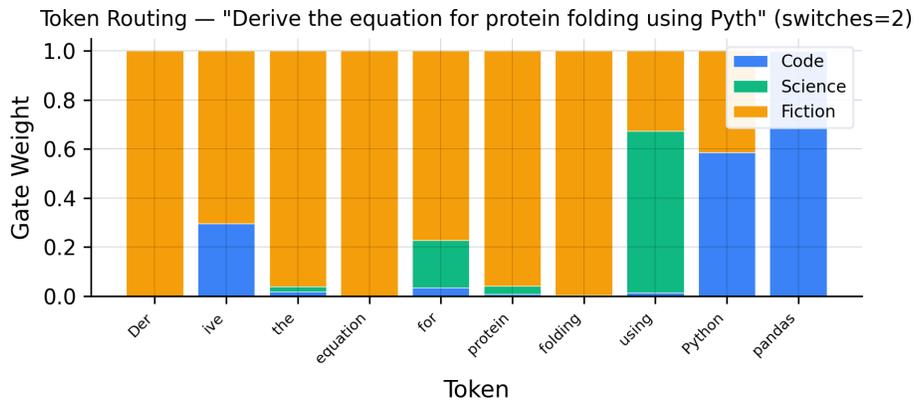


Figure 14: Gate weight heatmap for the prompt “Derive the equation for protein folding using Python pandas” (Pythia-410M, freeze=4, seed=42). Each column is a token; each row is an expert (code, science, fiction). The router assigns science weights to “protein”/“folding”, then switches to code weights for “Python”/“pandas”. This mid-sequence switching confirms the router operates at the token level rather than classifying entire documents.