ML Lab Research Proposal

Param-Mix: Parameter-Space Interpolation for Mixed-Data Optimization in Large Language Models

Team:	Olivia (solo)
Date:	May 24, 2026
Project:	Param-Mix: Parameter-Space Interpolation for Mixed-Data Optimization in LLMs
Target Compute:	2x AMD Instinct MI300X GPU accelerators, approx. 2 weeks

1. PROBLEM STATEMENT

Optimizing data mixtures for large language models remains computationally prohibitive because identifying the best mixture ratios typically requires repeated full-scale training runs guided by heuristic tuning. Existing methods are primarily optimized for language modeling loss such as perplexity, but reductions in perplexity do not reliably translate to improved downstream reasoning, coding, and mathematical performance.

This project investigates whether parameter-space interpolation can serve as a high-fidelity proxy for mixed-data optimization by exploiting the shared first-order optimization dynamics between independently trained domain experts and mixed-domain training trajectories.

2. PROPOSED APPROACH

We propose a framework that approximates mixed-data training using linear interpolation between independently trained domain-specific expert models. Given interpolation coefficients $\alpha \in \Delta^{K-1}$, the merged parameters are constructed as:

$$\Theta_{merge}(\alpha) = \Theta_{base} + \Sigma_k \alpha_k(\Theta_k - \Theta_{base})$$

This formulation allows efficient exploration of the mixture simplex without retraining full models. The method is motivated by a second-order Taylor expansion analysis showing that mixed-data optimization and parameter interpolation share identical first-order gradient terms, while discrepancies are confined to second-order curvature interactions involving Hessian-gradient products $H_k g_j$. We additionally fit LightGBM regressors to model the performance surface across downstream benchmark domains and optimize a utility function over predicted benchmark scores.

3. DATASET

We use a large-scale mid-training corpus partitioned into four primary domains: mathematics, code, supervised fine-tuning (SFT) data, and web/other data. The benchmark suite spans reasoning, mathematical problem solving, coding, language understanding, and professional knowledge tasks, including GSM8K, MATH, HumanEval, MBPP, GPQA, MMLU, ARC, BBH, and TriviaQA.

Preprocessing

- Domain clustering and token filtering to enforce structural purity.
- Hierarchical decomposition into fine-grained sub-clusters (e.g., formal mathematics, synthetic study notes, code repositories, long chain-of-thought supervision, and academic papers).
- Normalization of text distributions across all training segments.

4. EVALUATION METRICS

Category / Experiment	Metric	Target / Interpretation
Proxy Validation	Spearman rank correlation coefficient (ρ)	Target $\rho > 0.9$ between predicted interpolation performance and full training trajectories
Downstream Capability	Average benchmark accuracy	Evaluated across 5 key capability domains: math, code, knowledge, language, and reasoning
Interference Analysis	Hessian interaction matrices & task-vector cosine similarities	Quantify cross-domain interference and optimization orthogonality to identify breakdown points

5. COMPUTE REQUIREMENTS

The project requires approximately 2x AMD Instinct MI300X GPU accelerators with a combined memory budget near 384GB HBM memory to support distributed sparse Mixture-of-Experts mid-training experiments.

Phase	Hardware	Estimated Wall-Clock	Notes / Scope
Domain-Specific Expert Training	2x AMD Instinct MI300X	~7-9 days	Multiple 5B-25B token expert training runs with shared initialization
Interpolation Sweeps & Surface Fitting	2x AMD Instinct MI300X	~2-3 days	Exploring mixture simplex and training LightGBM performance regressors
Benchmark Evaluation & Verification	1x AMD Instinct MI300X	~1-2 days	Validation verification runs on sampled mixture configurations

Software Dependencies

- PyTorch, DeepSpeed distributed framework, FlashAttention
- Distributed checkpoint interpolation utilities and custom Hessian-vector product modules
- LightGBM-based regression models for downstream capability prediction
- ROCm 6.x + official PyTorch ROCm wheel environment

6. EXPECTED DELIVERABLES

- A working framework for efficient data mixture optimization using parameter interpolation without full retraining loops.
- Trained domain-specialized expert model checkpoints and a final optimized mid-trained language model.
- Comprehensive benchmark evaluation reports and detailed visualization maps of capability landscapes projected over the mixture simplex.
- Empirical analysis of first-order versus second-order optimization behavior during interpolation-based search sweeps.
- Reproducible scripts for interpolation search, utility optimization, and downstream benchmark testing.

7. RISK & MITIGATION

Risk	Likelihood	Mitigation
Higher-order curvature interactions cause parameter interpolation to diverge from mixed-data training	Medium	Constrain expert training to short horizons with constant learning rates, preserve shared initialization across all experts, and empirically validate via Hessian interaction analysis.
Domain-specific task vectors do not remain sufficiently orthogonal, reducing rank correlation	Medium	Monitor cosine similarity measurements continuously and perform rapid full-training verification runs on sampled mixture boundaries.