

Nature-Inspired Population-Based Evolution of Large Language Models

Yiqun Zhang^{1 2} Peng Ye² Xiaocui Yang¹ Shi Feng^{1 †}
Shufei Zhang² Lei Bai² Wanli Ouyang² Shuyue Hu^{2 †}

¹ School of Computer Science and Engineering, Northeastern University
Shenyang 110819, China

² Shanghai Artificial Intelligence Laboratory

yiqunzhang@stumail.neu.edu.cn {yangxiaocui, fengshi}@cse.neu.edu.cn
{yepeng, zhangshufei, bailei, ouyangwanli, hushuyue}@pjlab.org.cn

Abstract

Evolution, the engine behind the survival and growth of life on Earth, operates through the population-based process of reproduction. Inspired by this principle, this paper formally defines a newly emerging problem: the population-based evolution of large language models (LLMs). We introduce a novel framework that starts with a population of parent LLMs and allows this population to evolve through four key operations: (i) crossover, merging the weights of different parents to create offspring LLMs, (ii) mutation, introducing small, random changes to model weights to foster diversity, (iii) selection, prioritizing high-performing models, and (iv) succession, transferring the learned experience from parent to offspring LLMs. With only 200 samples per new task, the LLM population evolves rapidly to adapt to the task at hand, without any gradients. Experiments on 12 datasets show that our framework consistently outperforms existing multi-LLM merging and adaptation methods, achieving relative performance gains of up to 54.8% over the best LLM in the initial population. Moreover, our framework allows for (i) the evolution of LLMs across multiple new tasks simultaneously, (ii) scaling effectively with populations of up to 40 LLMs, and (iii) even zero-shot generalization to unseen held-out tasks. Code: <https://github.com/ZhangYiqun018/GENOME>

1 Introduction

For billions of years, nature has taught us a profound lesson: *evolution* drives the survival and flourishing of all life on Earth, from the simplest single-celled organisms to complex ecosystems and human civilizations. Evolution works not on individuals in isolation but on entire populations, where “*survival of the fittest*” shapes the

story of life. Now, imagine channeling these very principles not to biological species but to cutting-edge AI technology—what if we could harness population-based evolution to advance large language models (LLMs)?

There are multiple reasons for the population-based evolution of LLMs to be attractive. First, there is a growing proliferation of fine-tuned “expert” models that are specialized to specific tasks; for example, there are more than 170,000 expert models, trained through parameter-efficient fine-tuning techniques, available on Hugging Face (Muqeeth et al., 2024), which lies a solid foundation for conceptualization an LLM population. Second, recent studies have shown that merging the weights of LLMs with different expertise can produce versatile models (Yu et al., 2024; Goddard et al., 2024; Mavromatis et al., 2024), sometimes even resulting in emergence of new capabilities (Yadav et al., 2024). Third, enabling the evolution of an LLM population could allow these models to dynamically adapt to new, unseen tasks that have not been explicitly trained on (Huang et al., 2024; Mavromatis et al., 2024). This effectively recycles the substantial efforts and computing resources invested in training existing models. Last but not least, it also opens the door to democratizing foundation model development (Akiba et al., 2024), encouraging a broader community to jointly advance the field.

In this paper, we formalize the LLM population-based evolution problem, and propose a novel solution inspired by the very basic units of life—genes. Just as genes define biological traits, LLM weights define a model’s capabilities and behaviors. This naturally leads us to the genetic algorithm (GA) (Holland, 1992), a well-established evolutionary algorithm for optimization. We generalize GA to solve the LLM population-based evolution problem, and introduce **GENOME** (**GENetic Optimization for Model Evolution**).

[†]Corresponding author.

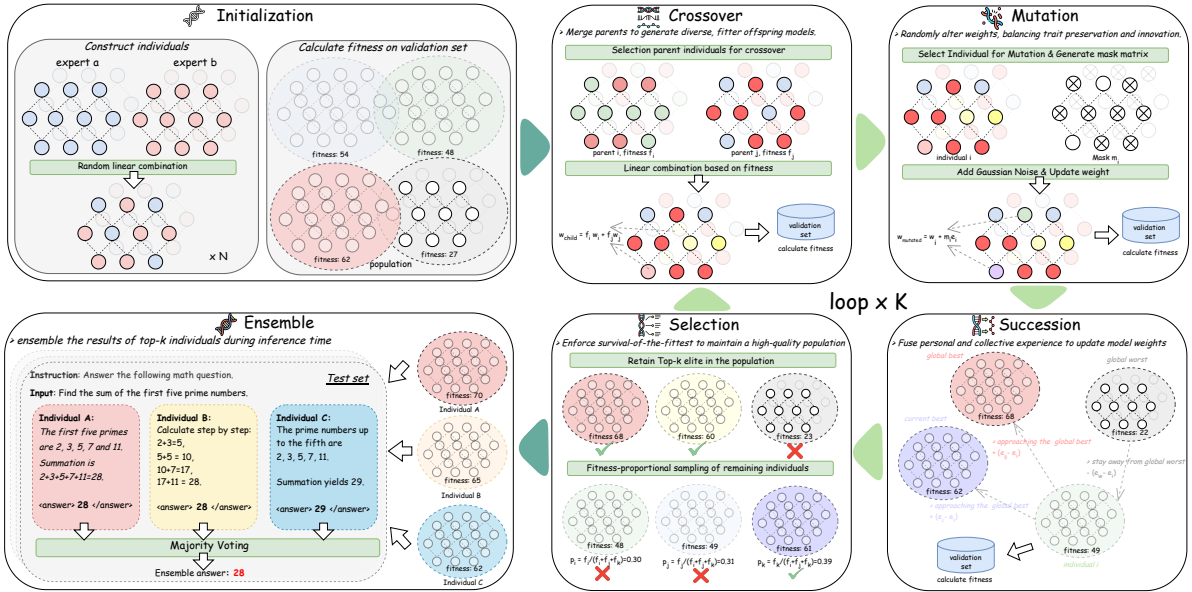


Figure 1: **GENOME(+): an easily accessible, gradient-free solution to the LLM population-based evolution problem.** We generalize the well-established genetic algorithms to solve the LLM population-based evolution problem, and additionally introduce two new operators. Starting with a pool of task-specific expert LLMs, the *fitness* of each model is evaluated based on its performance on a validation set (comprising only 200 samples). Through *crossover*, we combine the parameter weights of different parent LLMs to generate offspring LLMs. *Mutation* introduces small random perturbations to inject controlled variability, while *selection* retains top-performing LLMs, ensuring population diversity. *Succession* allows LLMs to learn from the best-performing models, and *ensemble* aggregates answers within the population during inference. This process repeats for multiple generations, enabling gradient-free adaptation and consistent improvements across tasks.

GENOME adapts the key concepts of GA as follows: (i) **fitness** measures how well a LLM performs (or fits) on a specific task, (ii) **crossover** merges the weights of different LLMs, generating offspring models and enabling them to inherit advantageous features from parent models, (iii) **mutation** adds small random perturbations to the model weights, enabling the model to acquire new abilities, and **selection** selects high-performing models to survive, effectively driving the LLM population to evolve in a better direction.

While connecting GA to LLMs is conceptually appealing, we recognize that GA heavily relies on reproduction, overlooking broader evolutionary principles, such as learning from peers and group decision-making. Thus, we additionally introduce two operations: (i) **succession**, which allows LLMs to learn from the best-performing LLM in the population, and (ii) **ensemble**, which integrates the outputs of top- k best performing models into a final answer. This results in **GENOME+** (as shown in Figure 1). Compared to GENOME, GENOME+ makes better use of the diversity within a LLM population. Despite their differ-

ences, both methods (collectively referred to as **GENOME(+)**) are easy to implement. Moreover, they allow for gradient-free weight updates, eliminating the need for backpropagation process. This provides an easily accessible solution to the LLM population-based evolution problem.

We conduct extensive evaluations on 12 datasets, showing that GENOME(+) consistently outperforms 8 methods, with only 200 samples per dataset. Unless otherwise stated, all percentage gains reported below are relative improvements in performance over the corresponding baseline. Notably, compared to the state-of-the-art Model Swarms (Feng et al., 2024), GENOME+ averages +10.75% and excels in reasoning-intensive tasks, improving by 36.3% on DROP (Dua et al., 2019) and 19.3% on MGSM (Shi et al., 2022). When evolving to simultaneously adapt to two tasks, GENOME+ surpasses Model Swarms by 13.46% on average, 50.9% on DROP, 17.6% on EmoryNLP (Zahiri and Choi, 2018), and 16% on GSM8k (Cobbe et al., 2021). Furthermore, GENOME+ generalizes zero-shot across tasks (App A.1.1, Table 5), achieving an av-

erage +11.72% cross-dataset gain over *Expert Fusion* (other baselines are $\sim 0-1\%$ or negative). On unseen languages (App A.1.1, Table 6), GENOME/GENOME+ improve over the *Best Single* by +12.83%/+15.98% on Flores-101 (Goyal et al., 2022). Even as the population scales up to 40 LLMs, both methods continue to demonstrate performance improvements. To assess computational efficiency, we repeat our key experiments on a single 4090 GPU with 24GB of memory. The results show that our method maintains consistent performance without requiring extensive computational resources. Last but not least, an ablation study confirms the positive contribution of each evolutionary operation.

Our key contributions are as follows:

- We present a novel conceptual framework that connects LLM weights to genes, task performance to individual fitness, and weight updates to evolutionary operations.
- We propose GENOME and GENOME+, which incorporate the key principles of traditional genetic algorithms, namely, crossover, selection, and mutation, as well as two operations: succession and ensemble. This showcases how traditional evolutionary algorithms can shed light on the population-based evolution of LLMs.
- We conduct extensive experiments in single task, multi-task domain, and zero-shot generalization settings, demonstrating GENOME+’s advantages, verifying its scalability to 40-model populations, and confirming each operation’s necessity through ablation studies. Experiments show that our results can be reproduced on a 4090 GPU with 24GB of memory.

2 Related Work

Model Merging Leveraging the complementary strengths of multiple LLMs has recently attracted much interest (Zhao et al., 2024; Wang et al., 2024a; Chen et al., 2024). Within this landscape, our work is most closely related to the line of research on model parameter merging. Earlier approaches, such as, Model Soup (Wortsman et al., 2022), Model Stock (Jang et al., 2024), TIES (Yadav et al., 2023), and DARE (Yu et al., 2024), typically merge the parameters of multiple models into those of a single one based on predefined rules, and maintain the model parameters *unchanged* after merging. More recent studies have extended to explore *dynamic* parameter merging

for adapting to new tasks (Hu et al., 2024; Akiba et al., 2024; Yang et al., 2024; Huang et al., 2024). Akiba et al. (2024) propose evolving the mixing weights at each layer within the parameter space, as well as evolving layer permutations within the data flow space. LoRAHub (Huang et al., 2024) formulates dynamic parameter merging as a linear-combination problem and optimizes linear coefficients. MERGE³ (Mencattini et al., 2025) similarly applies multi-objective evolutionary search over full-parameter mixing coefficients, but is instantiated in a full-weight setting rather than the LoRA-based multi-domain regime. Likewise, PackLLM (Mavromatis et al., 2024) introduces a test-time merging strategy for combining by solving an optimization problem to minimize perplexity without requiring training. Model Kinship (Hu et al., 2024) assesses relatedness between LLMs, guiding merges of top-performing models with those of distinct task capabilities. Model Swarms (Feng et al., 2024) is the closest dynamic baseline: it adapts PSO to iteratively update particles toward strong solutions. By contrast, GENOME optimizes the population-level objective, namely the aggregated fitness of the top-k individuals in the final population, and applies GA-style operators directly in LoRA parameter space.

Unlike these, our work differs in two key aspects: (i) instead of a single final model, we maintain and evolve a diverse population of LoRA adapters over time; (ii) we apply GA-style operations (e.g., mutation, succession) directly to LoRA parameters rather than optimizing low-dimensional mixing coefficients. This direct evolution consistently outperforms existing static and dynamic baselines.

3 Methodology

3.1 LLM Population-based Evolution

We define the **LLM population-based evolution** as an iterative optimization problem over a population of large language models (LLMs). Given a population: $P^{(t)} = \{x_i^{(t)}\}_{i=1}^N$, $x_i^{(t)} \leftrightarrow \mathbf{w}_i^{(t)} \in \mathbb{R}^d$, where each model $x_i^{(t)}$ is represented by learnable parameters $\mathbf{w}_i^{(t)}$, which may encompass either the full LLM weights or a subset thereof (such as LoRA or adapter parameters). The performance of each individual model on a given task T is quantified by a function $f(\mathbf{w}_i^{(t)}; T)$. The goal is to evolve the population $P^{(t)}$ over iterations $t = 0, 1, \dots, K$

to optimize the aggregated fitness of the top- k individuals in the final population $P^{(K)}$:

$$\max_{\{P^{(t)}\}_{t=0}^K} \frac{1}{k} \sum_{i=1}^k f(\mathbf{w}_{(i)}^{(K)}; T) \quad (1)$$

where $\mathbf{w}_{(i)}^{(K)}$ denotes the parameters of the i -th best-performing individual in $P^{(K)}$ ranked by fitness. The integer $k \in [1, N]$ controls whether the optimization focuses solely on the best individual ($k = 1$), or collectively improves a larger subset or the entire population ($k > 1$).

3.2 GENOME

In nature, biological evolution operates through an iterative process of **crossover**, **mutation**, and natural **selection**, where genetic variation drives diversity and selection preserves traits with higher adaptability. This cycle enables species to refine survival strategies in dynamic environments. Viewed through this lens, LLM performance reveals similar dynamics: each model represents an individual with unique traits, its weight parameters forming a digital genome, and task performance informs the selection mechanism that governs its chances of “survival” and “reproduction”.

We propose GENOME (GENetic Optimization for Model Evolution). Here, the target task T serves as the “environment”, and the **fitness function** f is the performance on the validation set of T . Starting with a population of n homogeneous LLM expert models (derived from the same base model), we propose to create an initial population $P^{(0)} = \{x_i\}_{i=1}^N$ via an **initialization** operation, where each individual’s genes are represented by its **LoRA** weights \mathbf{w}_i . The **crossover** operation is employed to create new offspring LLMs through merging parameters from parent LLMs, while the **mutation** operation is utilized to introduce variations to the parameters. The **selection** operation simulates natural selection pressure by retaining advantageous individuals. Algorithm 1 outlines GENOME, which includes initialization, crossover, mutation, and selection, executed over K iterations.

Initialization Starting with n LLM expert models $\{x_i\}_i^n$, we construct a diverse initial population through random linear combinations of expert model weights. For each new individual i , we combine the LoRA parameters \mathbf{w}_a and \mathbf{w}_b from

two expert models:

$$\mathbf{w}_i \leftarrow t \cdot \mathbf{w}_a + (1 - t) \cdot \mathbf{w}_b, \quad t \sim U(0, 1) \quad (2)$$

This process is repeated N times to form a population of size N . We then evaluate each individual on a validation set using the fitness function to obtain fitness scores f_i for the evolution process.

Crossover During sexual reproduction, chromosomal crossover creates novel genetic combinations that increase population diversity, providing raw material for natural selection to enhance adaptability. GENOME draws on this mechanism by combining the weights of different models. Specifically, we first design a selection probability based on individual fitness: $p_i = f_i / \sum f_k$, where $k = 1, 2, \dots, N$. Using this probability, we select parent pairs (x_{p_1}, x_{p_2}) with their corresponding weights $(\mathbf{w}_{p_1}, \mathbf{w}_{p_2})$. The offspring is generated by combining these weights:

$$\mathbf{w}_{child} \leftarrow t \cdot \mathbf{w}_{p_1} + (1 - t) \cdot \mathbf{w}_{p_2} \quad (3)$$

where $t = f_{p_1} / (f_{p_1} + f_{p_2})$ is the normalized weight. The crossover rate (c_r) controls the proportion of population undergoing this process, preserving model diversity while improving performance.

Mutation In biological evolution, gene mutations introduce random variations that maintain population diversity. Following this principle, GENOME applies mutation in two stages. First, each individual x_i is selected for mutation with probability im_r . Then, for each selected model, a binary mask \mathbf{m}_i is generated with probability gm_r to determine which weight parameters will be mutated. The mutation is executed as follows:

$$\mathbf{w}'_i \leftarrow \mathbf{w}_i + \mathbf{m}_i \cdot \mathbf{E}_i, \quad \mathbf{E} \sim \mathcal{N}(0, \sigma^2) \quad (4)$$

where the binary mask \mathbf{m}_i identifies the parameters to change, while the Gaussian noise \mathbf{E}_i and its standard deviation σ control the mutation intensity. The resulting mutated weights \mathbf{w}'_i form new individuals x' in the population, balancing the preservation of successful traits with the introduction of innovative variations. Crossover and mutation play complementary roles in population evolution. Fitness-weighted crossover recombines high-fitness traits but also tends to contract diversity over generations, whereas mutation restores exploration on selected dimensions and helps prevent premature homogenization. We therefore use

Gaussian perturbations as a simple continuous-space exploration operator: they do not guarantee monotonic improvement, but they assign positive probability to new directions that can be filtered by subsequent selection.

Selection Selection in GENOME serves to simulate natural selection pressure while maintaining a fixed population size N . After crossover and mutation operations, the selection process first retains the top αN individuals with the highest fitness scores ($\alpha \in (0, 1)$). Then, to restore the population size to N , individuals are selected from the entire current population with probabilities proportional to their fitness scores, where the selection probability for individual i is calculated as $p_i = f_i / \sum_{k=1}^N f_k$. This process effectively maintains population stability while favoring better-performing individuals.

3.3 GENOME+

While GA traditionally emphasizes evolution at the genetic level, the evolutionary history of intelligent species highlights the importance of knowledge transfer and collective decision-making within social groups. When models are viewed as a population, their evolution should involve not only parameter crossover and mutation, but also mechanisms for experience sharing and group decision-making. Motivated by this, we introduce GENOME+, which adds two new operations to the GENOME to achieve these capabilities. Through **succession**, individuals inherit successful strategies from the population’s evolutionary history, while avoiding prior failures. **Ensemble** methods combine the outputs of diverse models, leveraging their collective intelligence. Algorithm 2 outlines GENOME+.

Succession GENOME+ introduces succession as an experience-sharing operator for population evolution. Although it is inspired by the broader idea of social learning, it is not a PSO update rule: unlike Model Swarms (Feng et al., 2024), which iteratively updates particles toward best solutions, succession operates within GENOME’s GA pipeline and serves the population-level objective. Building on this insight, succession tracks learning patterns via an experience vector \mathbf{e}_i (matching LoRA weight dimensions). For each individual, \mathbf{e}_i is updated by integrating four sources: global best \mathbf{e}_g , current best \mathbf{e}_c , global worst \mathbf{e}_w (as

negative feedback), and self-experience \mathbf{e}_i . The update follows:

$$\mathbf{e}_i \leftarrow \frac{1}{C} [\phi_e \mathbf{e}_i + \phi_g (\mathbf{e}_g - \mathbf{e}_i) + \phi_c (\mathbf{e}_c - \mathbf{e}_i) - \phi_w (\mathbf{e}_w - \mathbf{e}_i)] \quad (5)$$

where ϕ_* are source weights and C is a normalization constant. Weights are then adjusted via $\mathbf{w}_i \leftarrow \mathbf{w}_i + \lambda \mathbf{e}_i$ with learning rate λ . This mechanism balances individual characteristics with collective wisdom during evolution.

Ensemble GENOME+ introduces an ensemble operation during the inference phase, selecting the top- k individuals with the highest fitness from the evolved population and aggregating their prediction results. This ensemble decision-making mechanism integrates the outputs of multiple high fitness individuals, while fully leveraging the advantages accumulated by the population during the evolutionary process. Through result aggregation in the inference phase, GENOME+ can achieve more robust predictive performance.

4 Experiments

In the following discussion, **GENOME(+)** refers to GENOME and GENOME+ together. Our experiments investigate the following key questions: (i) How effectively GENOME(+) adapts to a single task compared to existing methods? (ii) Can the framework maintain stable performance when handling multiple tasks simultaneously? (iii) How well does our framework scale with increasing population size? (iv) How does hardware configuration affect time efficiency and overall performance? (v) What is the impact of each operator on overall performance?

4.1 Experimental Setup

Models We employ *gemma-2-2b-it* (Team et al., 2024) as our foundation model, and construct a set of domain-specific experts by fine-tuning the foundation model on 10 distinct domains extracted from the *Tulu-v2-SFT-mixture* dataset (Iverson et al., 2023a). The fine-tuning process is implemented using the llama-factory framework (Zheng et al., 2024), incorporating the low-rank adaptation (LoRA) technique (Hu et al., 2021). We further demonstrate the capabilities of these 10 experts across different domains (see Figure 3), which confirms that our training process

yields expert models with varying proficiencies. Complete training hyperparameters and configurations are detailed in Table 18.

Datasets We consider 12 datasets covering 7 key capabilities of LLMs, including **General Knowledge** (MMLU (Hendrycks et al., 2021a), MMLUPro (Wang et al., 2024b)), **Mathematics** (MATH (Hendrycks et al., 2021b), GSM8K (Cobbe et al., 2021), and MGSM (Shi et al., 2022)), **Code Generation** (MBPP (Austin et al., 2021)), **Logical Reasoning** (DROP (Dua et al., 2019), BBH (Suzgun et al., 2022)), **Multilingual Processing** (MGSM (Shi et al., 2022) and FLORES-101 (Goyal et al., 2022)), **Affective Computing** (EmoryNLP (Zahiri and Choi, 2018), MELD (Poria et al., 2018)), and **Question Answering** (ARC_C (Clark et al., 2018) and CSQA (Talmor et al., 2019)). MGSM is a multilingual extension of mathematics problems, thus appearing in both Mathematics and Multilingual Processing. Each dataset is divided into a 200-sample validation set and a test set comprising approximately 1,000 samples. Detailed descriptions of the data splits and metrics can be found in Appendix A.2 and Table 17. These datasets do not overlap with the training data used for the expert models.

Implementation of GENOME(+) Unless otherwise stated, we use the following setting: crossover rate (c_r) of 0.3, individual mutation rate (im_r) of 0.3, gene mutation rate (gm_r) of 0.2, and standard deviation (σ) of 0.001. Additionally, the maximum number of iterations is set to 10, the population size is 10, ϕ_* is (0.95, 0.2, 0.2, 0.1), λ is 0.95. For ensemble operation, we select the top-3 individuals with the highest fitness scores on the validation set.

Baselines We compare GENOME(+) to 8 baselines as follows:

- **Best Single**, which refers to the best-performing expert model for a given dataset.
- **Data Merge**, which combines the 10 sub-datasets of Tulu-v2-sft-mixture into one complete training dataset for training a single LoRA.
- **Expert Fusion**, which merges the 10 expert models through $\mathbf{w}_{\text{fusion}} = \sum_{i=1}^n t_i \mathbf{w}_i$, where t_i denotes the normalized fitness of each expert i .
- **DARE_TIES** (Yu et al., 2024), which merges multiple LoRA models into a single one, while effectively reducing parameter interference.
- **MERGE³-LoRA** (Mencattini et al., 2025), originally designed for full-parameter model merging; in our experiments, we use a straightforward *LoRA-based adaptation* that applies the same merging procedure to the 10 LoRA experts on top of a shared frozen base model.
- **LoraHub** (Huang et al., 2024), which dynamically selecting and merging LoRA models, uses GA to optimize the combination weights.
- **Pack of LLMs** (Mavromatis et al., 2024), which dynamically merges LoRA models, assigning weights based on each model’s perplexity for given input prompts.
- **Model Swarms** (Feng et al., 2024), which employs PSO to iteratively optimize the set of expert models and adapt to a given dataset.

Note that Expert Fusion, LoraHub, Pack of LLMs, Model Swarms, and our framework all require a few samples for dynamic adaptation (*dynamic methods*); for a fair comparison, we fix the number of samples to 200. Additionally, all aforementioned methods—except for Data Merge, which does not require multiple expert models—utilize the same set of 10 expert models trained on the Tulu-v2-SFT-mixture. The implementation details of these methods are illustrated in Appendix A.3.

4.2 Results and Analysis

All experiments are conducted using **five random seeds**, and the reported results are averaged.

4.2.1 Single Task

Table 1 demonstrates that the GENOME and GENOME+ have attained superior performance across **12 datasets**. Specifically, GENOME+ exhibits the most robust performance—on average, enhancing by 24.06% relative to Best Single and by 10.75% in comparison to Model Swarms, thoroughly substantiating its efficacy. Our frameworks have demonstrated consistent performance enhancements across multiple categories, particularly in tasks necessitating intricate reasoning skills, such as BBH and DROP, where the framework attains increases of 36% and 54.80%, respectively. In Mathematics, it increases by 23.74% relative to the Best Single and by 13.37% in comparison to Model Swarms.

Nonetheless, we notice that in Question Answering (ARC_C, CSQA) and General Knowledge (MMLU, MMLUPro), the enhancements in

Method	MMLU	MMLUPro	GSM8k	MATH	MGSM	Flores-37	ARC_C	CSQA	BBH	DROP	EmoryNLP	MBPP
Best Single	52.90	26.57	40.80	14.30	30.22	21.51	57.34	64.30	30.22	30.40	32.53	31.78
Data Merge	13.60	25.57	33.00	12.20	25.98	22.51	46.08	37.40	25.98	27.40	32.37	15.37
DARE_TIES	48.30	28.57	35.50	12.50	31.85	22.33	69.80	23.90	31.85	19.22	34.22	37.77
Expert Fusion	55.80	27.67	39.50	13.10	31.82	21.74	69.03	65.00	31.82	22.60	32.66	28.55
LoraHub	53.00	27.17	46.47	14.90	34.70	21.83	69.97	66.10	39.30	35.32	31.53	42.63
MERGE ³ -LoRA	53.16	28.09	48.03	14.90	37.81	22.19	70.01	67.68	39.40	33.28	32.77	42.48
Pack of LLMs	54.30	27.67	38.13	11.50	31.06	19.24	70.22	63.30	38.90	21.30	30.46	42.56
Model Swarms	55.91	27.77	45.82	15.06	33.15	21.76	68.53	<u>68.76</u>	38.80	34.52	33.98	42.56
GENOME	56.66	<u>27.77</u>	<u>49.34</u>	<u>15.72</u>	<u>37.19</u>	<u>22.78</u>	<u>70.69</u>	<u>68.22</u>	<u>40.00</u>	<u>35.28</u>	<u>34.97</u>	43.60
- vs Best Single	+7.11%	+4.52%	+20.93%	+9.93%	+23.06%	+5.90%	+23.29%	+6.10%	+32.36%	+16.05%	+7.50%	+37.19%
- vs Model Swarms	+1.34%	—	+7.68%	+4.38%	+12.19%	+4.69%	+3.15%	-0.79%	+3.09%	+2.20%	+2.91%	+2.44%
GENOME+	<u>56.44</u>	30.98	51.24	16.41	39.55	23.38	74.38	69.89	41.10	47.06	38.81	<u>43.54</u>
- vs Best Single	+6.69%	+16.60%	+25.59%	+14.76%	+30.87%	+8.69%	+29.72%	+8.69%	+36.00%	+54.80%	+19.31%	+37.00%
- vs Model Swarms	+0.95%	+11.56%	+11.83%	+8.96%	+19.31%	+7.44%	+8.54%	+1.64%	+5.93%	+36.33%	+14.21%	+2.30%

Table 1: Performance comparison of different methods across 12 datasets (averaged over 5 runs with different random seeds). Best results are in bold and second-best results are underlined. In comparison rows, red/green percentages indicate relative improvements/decreases respectively, and ‘—’ denotes no change. GENOME+ achieves average relative improvements of **24.06%/10.75%** over Best Single/Model Swarms respectively (max **54.80%/36.33%** on DROP).

Method	Affective Computing		Mathematics		General Knowledge		Question Answering		Logical Reasoning	
	MELD	EmoryNLP	GSM8k	MATH	MMLU	MMLUPro	ARC_C	CSQA	DROP	BBH
Expert Fusion	50.35	31.91	41.77	12.60	53.10	25.97	68.94	63.00	22.10	38.30
LoraHub	50.77	32.67	41.77	10.30	52.90	26.37	69.37	62.40	34.10	38.30
MERGE ³ -LoRA	51.23	33.12	44.52	13.24	53.02	27.41	69.45	64.82	32.84	38.52
Pack of LLMs	51.70	29.81	37.90	11.90	53.40	27.07	69.03	63.10	21.20	38.60
Model Swarms	52.31	33.53	43.84	<u>15.10</u>	55.50	27.99	67.76	65.78	31.80	38.40
GENOME	<u>52.68</u>	<u>35.08</u>	<u>46.64</u>	14.48	<u>55.52</u>	<u>28.31</u>	<u>69.73</u>	<u>67.92</u>	<u>37.90</u>	<u>39.15</u>
- vs Model Swarms	+0.71%	+4.62%	+6.39%	-4.11%	+0.04%	+1.14%	+2.91%	+3.25%	+19.18%	+1.95%
GENOME+	56.05	39.42	50.87	17.24	56.40	30.13	73.60	70.35	48.00	39.90
- vs Model Swarms	+7.15%	+17.57%	+16.04%	+14.17%	+1.62%	+7.65%	+8.62%	+6.95%	+50.94%	+3.91%

Table 2: The results of multi-task domain.

performance are less significant. To further investigate this, we specifically analyze and compare performance on reasoning-intensive versus knowledge-intensive task subsets, where we observe a more pronounced improvement on reasoning tasks; detailed analysis results are provided in Appendix A.1.4 and Table 7.

4.2.2 Multi-Task Domain

In our second experiment, we consider five domains each consists of two tasks. To allow for simultaneous adaptation to two tasks we consider the fitness value to be the average of the two tasks’ performance. We only compare dynamic methods because the performance of static methods remains constant across tasks. Compared to single task, this scenario is more challenging.

As shown in Table 2, GENOME and GENOME+ are still able to achieve performance at the level comparable to single-task scenario, whereas other dynamic methods suffer from various degrees of performance degradation. Interestingly, GENOME+ achieves the best results in every category, proving the stability of our framework in navigating challenging situations.

4.2.3 Scalability

We conduct comparison experiments using Model Swarms, GENOME, and GENOME+ to examine their scalability, increasing the population size from 10 to 40 by 10. In the MMLUpro, MATH, and MMLUProReasoning tasks, all approaches exhibit enhanced performance with an increase in population size, with GENOME+ attaining the highest performance across all sizes. As shown in Figure 2, GENOME consistently outperforms Model Swarms, while GENOME+ demonstrates the strongest performance. These results confirm that GENOME+ can scale to larger population sizes up to 40.

4.2.4 Ablation Study

We perform ablation experiments on six typical dataset to confirm the contribution of each operation in the GENOME+. As shown in Table 3, every operator of GENOME+ has a positive impact on performance. For comparison, we use *random selection* in place of selection because eliminating it completely would result in an unlimited growth of the population size. The replacement of our selection with random selection results in an average performance loss of 10.57%. For the en-

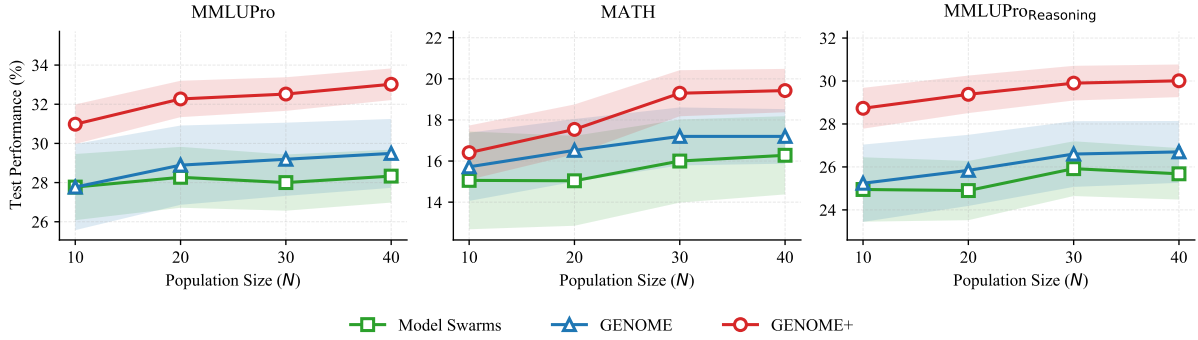


Figure 2: Performance trends with increasing population sizes (N) across different methods.

Setting	MMLUPro	GSM8k	DROP	BBH	EmoryNLP	MGSM
GENOME+	30.98	51.24	47.06	41.10	38.81	39.55
<i>w/o initialization</i>	29.72 (-4.07%)	46.02 (-10.19%)	42.70 (-9.26%)	39.65 (-3.53%)	38.29 (-1.34%)	36.61 (-7.43%)
<i>w/o crossover</i>	30.32 (-2.13%)	48.86 (-4.64%)	45.10 (-4.16%)	40.00 (-2.68%)	38.01 (-2.06%)	35.72 (-9.68%)
<i>w/o mutation</i>	29.72 (-4.07%)	49.51 (-3.38%)	43.90 (-6.71%)	40.10 (-2.43%)	37.17 (-4.23%)	38.41 (-2.88%)
<i>w/o cro. & mut.</i>	27.87 (-10.04%)	37.00 (-9.98%)	38.60 (-17.98%)	35.29 (-9.07%)	33.11 (-16.28%)	45.38 (-11.44%)
<i>random selection</i>	26.77 (-13.59%)	44.96 (-12.26%)	38.50 (-18.19%)	38.91 (-5.33%)	35.21 (-9.28%)	37.67 (-4.75%)
<i>w/o succession</i>	28.76 (-7.17%)	50.15 (-2.13%)	40.22 (-14.53%)	39.72 (-3.36%)	36.52 (-5.90%)	38.70 (-2.15%)
<i>w/o ensemble</i>	28.81 (-7.00%)	48.46 (-5.43%)	37.00 (-21.38%)	38.31 (-6.79%)	34.76 (-10.44%)	32.51 (-17.80%)

Table 3: Ablation study of different operations

semble operation, its removal results in an average performance drop of 11.47%, especially a steep reduction in 21.38% in the DROP. The removal of initialization and succession operations lead to performance drops of 5.87% and 5.97%, respectively. The crossover and mutation operations are crucial for maintaining population diversity, despite their modest consequences (performance decreases of 4.23% and 3.95%, respectively). When both operations are removed simultaneously, the performance drop enlarges to 12.47%, indicating a synergistic effect between them, eliminating both reduces population diversity and thus leads to a greater decline in overall performance.

4.2.5 Time Efficiency on Different GPUs

Method	Task	Score		Time (s)	
		4090	A100	4090	A100
GENOME	MMLUPro	28.97	28.67	~3600	~1600
	DROP	35.80	35.32	~1000	~800
	MATH	15.56	15.78	~4000	~2600
GENOME+	MMLUPro	31.47	30.98	~5000	~3000
	DROP	46.20	46.92	~2000	~1200
	MATH	16.79	16.81	~6000	~3800

Table 4: Time and score on different GPUs

To evaluate the time efficiency of our method, we conduct experiments on two different GPUs: **RTX 4090 (24GB VRAM)** and **A100 (80GB VRAM)**. All experiments run on a single GPU and

are repeated five times to obtain the average result. The base model is *gemma-2-2b-it*. The population size is set in 10, and other hyperparameters remain consistent with those used in the single-task experiments. As shown in Table 4, the inference speed of the 4090 and A100 varies across different tasks. Despite having less VRAM, the 4090 maintains high inference performance on MMLUPro, DROP, and MATH tasks, further validating the applicability of our method.

4.2.6 Extended Analysis

Due to space constraints, we provide additional experimental results in the Appendix A.1: (i) **Zero-shot and Language Generalization:** GENOME(+) demonstrates robust cross-dataset transfer and adaptation to unseen languages (Tables 5, 6); (ii) **Validation Set Sensitivity:** Performance remains stable across varying validation set sizes (Tables 9–11); (iii) **Model Size Scalability:** GENOME(+) maintains superior performance over baselines when scaled to larger models, including *Llama-3.1-8B* and *Qwen-2.5-14B* (Tables 12, 13); (iv) **Comparison with Distilled Fine-tuning:** GENOME+ remains competitive even against baselines fine-tuned on GPT-5 reasoning traces, highlighting sample efficiency without external model supervision (Table 14); (v) **Hyperparameter Analysis:** Detailed studies (Tables 15, 16) confirm that GENOME(+) is insensi-

tive to hyperparameter configurations, eliminating the need for extensive search.

5 Conclusion

This paper formally defines the population-based evolution problem for LLMs and presents two novel solutions: GENOME and GENOME+. These two frameworks leverage principles of biological evolution to adapt LLMs to new tasks with the use of 200 samples. Experiments show that GENOME+ consistently surpasses existing methods across various tasks, showing notable enhancements in reasoning-intensive tasks. The framework demonstrates strong multi-task domain capabilities, while ensuring scalability with larger population sizes (up to 40). Ablations demonstrate the essential role of each evolutionary component, with selection and ensemble operations exhibiting the greatest influence on performance. In addition, our experiments confirm that GENOME(+) can operate on NVIDIA RTX 4090. This study illustrates that population-based evolution presents a promising approach to enhance LLMs' capabilities through evolutionary optimization.

Limitations

Our proposed method, GENOME(+), is gradient-free and computationally efficient, capable of running on consumer-grade hardware such as an NVIDIA RTX 4090. However, its optimization requires a pool of trained LoRA adapters. Furthermore, similar to most model-merging techniques, our approach currently supports only the fusion of homogeneous models. This is primarily due to the structural discrepancies in vocabulary sizes, hidden dimensionalities, and the specific dimensions of Query-Key-Value (Q, K, V) projections, which preclude direct, lossless weight synchronization across heterogeneous architectures. Nevertheless, this limitation does not hinder the practical utility of our framework, as the ecosystem provides a vast array of experts derived from identical base models. As of January 2026, the Hugging Face community hosts tens of thousands of specialized adapters for dominant homogeneous families like Qwen and Llama, providing a rich and sustainable foundation for population-based model evolution research. Moreover, our multi-task experiments use a scalarized average fitness and do not explicitly model pairwise task conflicts; extending population evolution to conflict-aware multi-task set-

tings is an important direction for future work.

Acknowledgements

The work was supported by the National Natural Science Foundation of China (Nos. 62272092, 62172086, 62506186), and the Fundamental Research Funds for the Central Universities under Grants (N25XQD004). This work was supported by Shanghai Artificial Intelligence Laboratory. This work was done during Yiqun Zhang's internship at Shanghai Artificial Intelligence Laboratory.

References

- Takuya Akiba, Makoto Shing, Yujin Tang, Qi Sun, and David Ha. 2024. Evolutionary optimization of model merging recipes. *arXiv preprint arXiv:2403.13187*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>.
- Shuhao Chen, Weisen Jiang, Baijiong Lin, James T. Kwok, and Yu Zhang. 2024. Routerdc: Query-based router by dual contrastive learning for assembling large language models. *Preprint*, arXiv:2409.19886.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, and 1 others. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *Preprint*, arXiv:1803.05457.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. *Preprint*, arXiv:1903.00161.

- Shangbin Feng, Zifeng Wang, Yike Wang, Sayna Ebrahimi, Hamid Palangi, Lesly Miculicich, Achin Kulshrestha, Nathalie Rauschmayr, Yejin Choi, Yulia Tsvetkov, and 1 others. 2024. Model swarms: Collaborative search to adapt llm experts via swarm intelligence. *arXiv preprint arXiv:2410.11163*.
- Charles Goddard, Shamane Siriwardhana, Malikeh Ehghaghi, Luke Meyers, Vlad Karpukhin, Brian Benedict, Mark McQuade, and Jacob Solawetz. 2024. Arcee’s mergekit: A toolkit for merging large language models. *arXiv preprint arXiv:2403.13257*.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. *NeurIPS*.
- John H Holland. 1992. Genetic algorithms. *Scientific american*, 267(1):66–73.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Yedi Hu, Yunzhi Yao, Ningyu Zhang, Shumin Deng, and Huajun Chen. 2024. Exploring model kinship for merging large language models. *arXiv preprint arXiv:2410.12613*.
- Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. 2024. [Lorahub: Efficient cross-task generalization via dynamic lora composition](#). *Preprint*, arXiv:2307.13269.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, and 1 others. 2023a. Camels in a changing climate: Enhancing llm adaptation with tulu 2. *arXiv preprint arXiv:2311.10702*.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, and 1 others. 2023b. Camels in a changing climate: Enhancing llm adaptation with tulu 2. *arXiv preprint arXiv:2311.10702*.
- Dong-Hwan Jang, Sangdoon Yun, and Dongyoon Han. 2024. Model stock: All we need is just a few fine-tuned models. In *European Conference on Computer Vision*, pages 207–223. Springer.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, and 1 others. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.
- Costas Mavromatis, Petros Karypis, and George Karypis. 2024. [Pack of llms: Model fusion at test-time via perplexity optimization](#). *Preprint*, arXiv:2404.11531.
- Tommaso Mencattini, Adrian Robert Minut, Donato Crisostomi, Andrea Santilli, and Emanuele Rodolà. 2025. Merge³: Efficient evolutionary merging on consumer-grade gpus. *arXiv preprint arXiv:2502.10436*.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of gpt-4](#). *Preprint*, arXiv:2306.02707.
- Mohammed Muqeeth, Haokun Liu, Yufan Liu, and Colin Raffel. 2024. Learning to route among specialized experts for zero-shot generalization. *arXiv preprint arXiv:2402.05859*.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2022. [Language models are multilingual chain-of-thought reasoners](#). *Preprint*, arXiv:2210.03057.

- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. **CommonsenseQA: A question answering challenge targeting commonsense knowledge**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhatipatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- TechCrunch. 2023. **Sharegpt website**.
- Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2024a. Mixture-of-agents enhances large language model capabilities. *arXiv preprint arXiv:2406.04692*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhramil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024b. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. **Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time**. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. Resolving interference when merging models. *arXiv preprint arXiv:2306.01708*, 1.
- Prateek Yadav, Tu Vu, Jonathan Lai, Alexandra Chronopoulou, Manaal Faruqui, Mohit Bansal, and Tsendsuren Munkhdalai. 2024. What matters for model merging at scale? *arXiv preprint arXiv:2410.03617*.
- Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. 2024. **Adamerging: Adaptive model merging for multi-task learning**. *Preprint*, arXiv:2310.02575.
- Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. 2024. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*.
- Sayed M Zahiri and Jinho D Choi. 2018. Emotion detection on tv show transcripts with sequence-based convolutional neural networks. In *Workshops at the thirty-second aaai conference on artificial intelligence*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Ziyu Zhao, Leilei Gan, Guoyin Wang, Wangchunshu Zhou, Hongxia Yang, Kun Kuang, and Fei Wu. 2024. **Loraretriever: Input-aware lora retrieval and composition for mixed tasks in the wild**. *Preprint*, arXiv:2402.09997.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. **Llamafactory: Unified efficient fine-tuning of 100+ language models**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. **Lima: Less is more for alignment**. In *Advances in Neural Information Processing Systems*, volume 36, pages 55006–55021. Curran Associates, Inc.

A Appendix

A.1 Further Experiments

Method	MMLUPro → MMLU	MATH → GSM8k	EmoryNLP → MELD	Macro Avg. %Δ vs Expert Fusion
Expert Fusion	53.40	42.53	50.75	–
LoraHub	53.50	43.37	49.92	+0.18%
Pack of LLMs	52.70	37.83	51.93	-3.35%
Model Swarms	52.10	44.05	51.29	+0.74%
GENOME	55.40	46.32	52.07	+5.09%
GENOME+	56.65	51.10	55.28	+11.72%

Table 5: Generalization performance across different domains.

A.1.1 Zero-shot Generalization

We assess the zero-shot generalization capability through two experiments: cross-dataset trans-

fer and unseen language adaptation. We identify three groups of related task pairs for cross-dataset transfer: MMLUPro to MMLU, MATH to GSM8k, and EmoryNLP to MELD. For each pair, we initially enable a dynamic method to adapt to one task with 200 samples, and then test its performance on the other task without any samples. Table 5 shows that all dynamic methods transfer, with GENOME/GENOME+ strongest; GENOME+ achieves a macro gain of **+11.72%** over Expert Fusion (vs. GENOME **+5.09%**), while other baselines are \leq **+0.74%** or negative.

Method	Flores-37	Flores-101
Best Single	21.51	12.70
Expert Fusion	21.74 (+1.07%)	13.74 (+8.19%)
LoraHub	21.83 (+1.49%)	13.82 (+8.82%)
Pack of LLMs	19.24 (-10.55%)	12.40 (-2.36%)
Model Swarms	21.76 (+1.16%)	13.64 (+7.40%)
GENOME	22.78 (+5.90%)	14.33 (+12.83%)
GENOME+	23.38 (+8.69%)	14.73 (+15.98%)

Table 6: Performance on seen (**Flores-37**) and unseen languages (**Flores-101**).

Additionally, we assess our framework using the Flores101 dataset, comprising 64 languages not included in the training phase of gemma-2¹. Table 6 indicates that the majority of methods demonstrate performance improvements, with the exception of the Pack of LLMs. GENOME and GENOME+ demonstrate significant enhancements of 12.83% and 15.98%, respectively, exceeding their performance improvements on Flores37, which includes languages encountered during pre-training. The results indicate a robust generalization capability of our framework, especially in the context of low-resource languages.

Method	MMLUProKnowledge	MMLUProReasoning
Best Single	27.90	23.20
Model Swarms	27.36 (-1.94%)	24.95 (+7.54%)
GENOME	27.48 (+0.44%)	25.24 (+8.79%)
GENOME+	29.05 (+6.18%)	28.73 (+23.84%)

Table 7: Performance comparison of different methods on MMLUPro knowledge and reasoning subsets. The percentages in parentheses indicate the relative improvement (red) or degradation (green) compared to the Best Single.

Variance Across Random Seeds For completeness, we report the variability of all main meth-

¹<https://cloud.google.com/vertex-ai/generative-ai/docs/learn/models>

ods across random seeds. All numbers in the main text are means over five runs with different random seeds {41, 42, 47, 53, 3407}. To keep the main tables readable (12 datasets and multiple methods), we only show mean scores there, which is also consistent with prior work we compare to (e.g., LoRAHub, Pack of LLMs, MERGE³, and Model Swarms). Table 8 provides the full “mean \pm standard deviation” statistics for all methods and datasets.

A.1.2 Robustness to Validation Set Size

We examine how sensitive our evolution procedure is to the choice and size of the validation set. Following standard practice in adapter-based model-merging work, we use a relatively small held-out set (200 examples per task) to guide search/merging, and a larger, disjoint test set for final reporting. In our setup, for each task: (i) the validation set is randomly sampled from the benchmark and used only for evolution/selection; and (ii) the test set is strictly disjoint and all main-table results are computed on this test set. To directly address concerns about potential overfitting to a particular 200-example split, we vary the validation size $N \in \{20, 50, 100, 150, 200\}$ and resample the validation set while keeping the test set fixed. For three representative tasks (MMLUPro, MATH, DROP), we run GENOME and GENOME+ and report mean \pm standard deviation over 5 seeds (Table 9–11). The $N = 200$ columns are directly copied from the main experiments. We observe that (i) very small validation sets (e.g., $N = 20$) lead to somewhat less stable performance, which is expected; and (ii) as N increases, both GENOME and GENOME+ show only mild variation, with GENOME+ consistently outperforming GENOME across all N . Overall, we see no evidence that our results depend on a particular 200-example validation split: the final test-set performance is robust to both validation set size and random resampling.

A.1.3 Model Size Scalability

We used Llama-3.1-8B-Instruct (Grattafiori et al., 2024) and Qwen2.5-14B-Instruct (Qwen et al., 2025) as the base model (fine-tuned with the same SFT dataset for LoRA weights) on three representative tasks (MMLUPro, MATH, DROP) to compare GENOME(+). The results are shown in Table 12 and 13. GENOME(+) still achieved the best performance. This indicates that

Method	MMLU	MMLUPro	GSM8K	MATH	MGSM	Flores-37	ARC_C	CSQA	BBH	DROP	EmoryNLP	MBPP
LoRAHub	53.00±0.78	27.17±1.57	46.47±3.35	14.90±0.89	34.70±2.80	21.83±1.12	69.97±3.58	66.10±1.01	39.30±3.13	35.32±5.59	31.53±0.89	42.63±1.79
MERGE ³ -LoRA	53.16±0.70	28.09±2.08	48.03±0.84	14.90±1.57	37.81±5.20	22.19±2.00	70.01±1.37	67.68±4.53	39.40±2.40	33.28±4.68	32.77±2.09	42.48±1.03
Pack of LLMs	54.30±2.46	27.67±0.78	38.13±1.01	11.50±2.91	31.06±1.12	19.24±2.35	70.22±1.12	63.30±3.13	38.90±0.89	21.30±4.25	30.46±2.80	42.56±1.01
Model Swarms	55.91±1.07	27.77±1.69	45.82±2.49	15.06±2.38	33.15±2.70	21.76±2.57	68.53±2.94	68.76±2.09	38.80±3.23	34.52±5.90	33.98±2.91	42.56±2.65
GENOME	56.66±0.93	27.77±2.20	49.34±2.04	15.72±1.66	37.19±1.79	22.78±1.91	70.69±2.50	68.22±2.71	40.00±2.10	35.28±6.31	34.97±2.51	43.60±1.86
GENOME+	56.44±0.63	30.98±1.01	51.24±1.34	16.41±1.32	39.55±1.57	23.38±1.24	74.38±1.45	69.89±1.12	41.10±1.57	47.06±5.37	38.81±1.45	43.54±1.34

Table 8: Per-dataset mean \pm standard deviation over five random seeds for all main methods.

Method	N=20	N=50	N=100	N=150	N=200
GENOME	26.27±1.72	26.57±2.52	27.50±1.74	27.97±1.75	27.77±2.20
GENOME+	29.97±1.26	30.09±1.10	30.42±0.69	30.74±1.02	30.98±1.01

Table 9: Validation-set size study on MMLUPro: mean \pm std over 5 seeds.

Method	N=20	N=50	N=100	N=150	N=200
GENOME	15.83±1.32	16.13±0.90	15.63±1.45	15.07±1.07	15.72±1.66
GENOME+	15.38±0.88	15.83±1.70	16.48±1.35	17.67±0.73	16.41±1.32

Table 10: Validation-set size study on MATH: mean \pm std over 5 seeds.

Method	N=20	N=50	N=100	N=150	N=200
GENOME	33.75±6.75	34.27±2.41	35.03±7.23	33.70±3.76	35.28±6.31
GENOME+	45.73±5.75	47.97±5.28	45.03±5.73	45.40±5.15	47.06±5.37

Table 11: Validation-set size study on DROP: mean \pm std over 5 seeds.

Method	MMLUPro	MATH	DROP
Best Single	45.00	24.00	50.00
Data Merge	28.47	24.80	33.99
DARE_TIES	46.45	30.00	32.30
Expert Fusion	45.85	31.60	34.10
LoraHub	43.46	31.60	49.60
Pack of LLMs	45.45	31.40	21.90
Model Swarms	45.50	32.20	53.70
GENOME	<u>48.15</u>	<u>32.65</u>	<u>56.20</u>
- vs Best Single	+7.00%	+36.04%	+12.40%
- vs Model Swarms	+5.82%	+1.40%	+4.64%
GENOME+	49.15	34.55	60.10
- vs Best Single	+9.22%	+43.96%	+20.20%
- vs Model Swarms	+8.02%	+7.30%	+11.92%

Table 12: The performance of single task (base model: Llama-3.1-8B-Instruct).

Method	MMLUPro	MATH	DROP	MGSM
Best Single	58.74	36.60	61.20	65.49
Data Merge	53.50	35.50	19.80	59.54
Expert Fusion	59.14	34.30	54.30	53.20
LoraHub	58.27	39.20	66.27	64.95
Pack of LLMs	58.44	28.60	42.33	51.02
Model Swarms	60.24	42.00	64.87	66.64
GENOME	<u>61.02</u>	<u>42.23</u>	<u>66.60</u>	<u>66.77</u>
- vs Best Single	+3.88%	+15.38%	+8.82%	+1.95%
- vs Model Swarms	+1.29%	+0.55%	+2.67%	+0.20%
GENOME+	64.09	43.20	79.20	67.48
- vs Best Single	+9.11%	+18.03%	+29.41%	+3.04%
- vs Model Swarms	+6.39%	+2.86%	+22.09%	+1.26%

Table 13: The performance of single task (base model: Qwen2.5-14B-Instruct).

GENOME(+) is not only effective for 2B models, but also works well on 8B model and 14B model.

A.1.4 Subtask

To verify whether this difference is influenced by task type, we construct subsets of knowledge-based and reasoning-based tasks within the MMLUPro. The former includes tasks that rely on domain knowledge such as law, history, and psychology, while the latter comprises fields that require logical reasoning, such as mathematics, physics, and computer science. As shown in Table 7, the improvement rates of all methods on reasoning-based tasks are much higher than those on knowledge-based tasks. Notably, GENOME+ achieves a performance gain of 23.84% on the reasoning subset compared to the Best Single, whereas the gain is 6.18% on the knowledge subset. Other methods exhibit a similar trend. These results reveal that existing dynamic methods, including our population-based evolution framework, excels in reasoning-intensive tasks. On the other hand, knowledge-intensive tasks may require the integration of external knowledge bases or knowledge enhancement strategies, which is beyond the capability of dynamic methods.

A.1.5 Additional Baselines: Fine-tuning and Direct Ensembles

To complement the main baselines, we further evaluate several additional settings for instruction tuning. The benchmarks we use are standard LLM evaluation suites that only provide inputs and final answers, and generally lack intermediate reasoning traces or consistent output formats. Direct LoRA fine-tuning on such raw (input, answer) pairs is known to be weak supervision for reasoning-heavy tasks and also leads to format-parsing issues at evaluation time. We therefore construct two fine-tuning datasets and train LoRA baselines on Gemma-2-2B-it under the same hyperparameters as our original experts:

- **Raw benchmark fine-tuning:** LoRA fine-tuning directly on the original (question, an-

Setting	MMLU	MMLUPro	GSM8K	MATH	MGSM	Flores-37	ARC_C	CSQA	BBH	DROP	EmoryNLP	MBPP
Raw fine-tuning	26.92	10.00	0.00	0.00	0.00	0.00	26.09	20.00	0.00	47.50	1.33	39.50
Distilled fine-tuning	56.73	<u>28.50</u>	46.50	12.00	32.50	<u>22.75</u>	74.58	<u>68.50</u>	43.00	54.50	36.72	35.00
Direct Top-3 Ensemble	48.80	25.82	42.45	12.55	27.68	19.36	68.98	67.32	32.43	19.75	<u>36.89</u>	31.78
Expert Fusion	55.80	27.67	39.50	13.10	31.82	21.74	69.03	65.00	31.82	22.60	32.66	28.55
GENOME	56.66	27.77	49.34	15.72	37.19	22.78	70.69	68.22	40.00	35.28	34.97	43.60
GENOME+	56.44	30.98	51.24	16.41	39.55	23.38	<u>74.38</u>	69.89	<u>41.10</u>	<u>47.06</u>	38.81	<u>43.54</u>

Table 14: Additional baselines on Gemma-2-2B-it. All numbers are single-seed results (seed = 42) except for GENOME / GENOME+, which are averaged over 5 seeds.

swer) benchmark pairs (with 200 validation examples per task).

- **Distilled benchmark fine-tuning:** a strong teacher model (GPT-5) is used to (i) generate intermediate reasoning and (ii) normalize answer formats; Gemma-2-2B-it is then fine-tuned on this enhanced dataset. Note that this baseline benefits from additional teacher supervision, while our method and all model-merging baselines do not use any external model supervision.

For the requested ensemble-style baselines:

- **Direct Top-3 Ensemble:** a top-3 ensemble of the original experts, matching the main text setup (same k and prompts), without any evolutionary search.
- **Expert Fusion (naive averaging):** this corresponds exactly to the expert-fusion baseline reported in Tables 1–4 (naive parameter averaging of top experts, without evolution).

Since these additional baselines do not involve stochastic evolutionary operations, we report them under a single seed (seed = 42). For comparison, we also include the 5-seed averages of GENOME and GENOME+. Results are summarized in Table 14. We observe three main trends. (i) *Raw fine-tuning* performs the worst overall and even collapses to 0.00 on several datasets, consistent with the intuition that using only (input, final answer) pairs without reasoning traces and with noisy formats provides a very weak supervision signal. (ii) Even with *GPT-5 distillation*, gains over GENOME+ are limited: the distilled fine-tuning baseline benefits from an extra powerful teacher, yet it surpasses GENOME+ on only a few datasets, while GENOME+ never relies on any external model. (iii) *Direct aggregation* (Top-3 Ensemble and naive Expert Fusion) is highly task-dependent: it helps on some benchmarks but degrades sharply on others, illustrating that naive

expert aggregation is unstable compared with the evolutionary GENOME/GENOME+ approach.

A.1.6 Hyperparameter Statistical Analysis

To optimize hyperparameter selection and verify the stability of the GENOME and GENOME+ methods, we conduct 500 experiments. The hyperparameters involved in the study include:

- Cross rate (c_r): Range from 0.1 to 1.0.
- Individual mutation rate (im_r): Range from 0.1 to 1.0.
- Gene mutation rate (gm_r): Range from 0.1 to 1.0.

The dependent variables are *test performance*, *Top-3 ensemble performance*, and *optimization time*. The experiments use a fixed random seed of 47, set the population size (N) to 10, σ to 0.001, and selected GENOME as the test method (we believe that the hyperparameter settings of GENOME and GENOME+ are similar). The dataset use in this experiment is MMLUPro, which is a highly comprehensive test dataset. The experimental environment consisted of $4 \times A100$ 80G GPUs. In each experiment, we discretize each hyperparameter range into fixed intervals of 0.1 (e.g., 0.1, 0.2, 0.3, ..., 1.0). From these discrete candidate values, one value was randomly selected for c_r , im_r , and gm_r in each experimental run, subject to the same random seed initialization.

In the correlation analysis (see in Table 15), Spearman and Pearson methods reveal that the c_r had a significant positive correlation with *test performance* and *optimization time*, and a weak correlation with *ensemble performance*. The im_r is found to be weakly negatively correlated with *test performance* but significantly positively correlated with *optimization time*. The gm_r show small correlation coefficients with all dependent variables.

Given the complex coupling relationships between variables (e.g., gene mutation rate is only meaningful when individual mutation rate is

HyperParameter	Metric	Pearson		Spearman	
		r	p-value	ρ	p-value
Cross Rate	Test	0.192***	2.132e-03	0.194***	1.927e-03
	Ensemble	0.119*	5.831e-02	0.102	1.052e-01
	Time	0.382***	3.030e-10	0.408***	1.395e-11
Individual Mutation Rate	Test	-0.134**	3.284e-02	-0.118*	5.946e-02
	Ensemble	-0.091	1.479e-01	-0.094	1.368e-01
	Time	0.775***	4.647e-52	0.796***	9.183e-57
Gene Mutation Rate	Test	-0.093	1.395e-01	-0.092	1.433e-01
	Ensemble	0.077	2.240e-01	0.092	1.423e-01
	Time	0.123*	5.002e-02	0.125**	4.684e-02

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 15: Correlation Analysis between Algorithm Parameters and Performance Metrics

greater than 0), simple correlation analysis may not be sufficient. Therefore, we conduct further multiple linear regression analysis. For example, for *test performance*, we construct the following regression model:

$$\text{test}_i = \beta_0 + \beta_1 \cdot c_{ri} + \beta_2 \cdot im_{ri} + \beta_3 \cdot gm_{ri} + \varepsilon_i \quad (6)$$

We construct corresponding regression models for *ensemble performance* and *optimization time* and estimate the parameters using the least squares method. The significance test results are shown in Table 16.

Analysis results show that c_r has a significant positive impact on *test performance*, and im_r shows a negative impact. The impact of gm_r do not reach significance. For *ensemble performance*, both correlation and multiple regression analyses show that the effects of the three hyperparameters are not significant, indicating that after adopting the ensemble strategy, the GENOME’s sensitivity to hyperparameters decreased. In terms of *optimization time*, increased cross rate and individual mutation rate significantly prolonged total runtime.

Variable	Coefficient	Std. Error	t-statistic	p-value
Panel A: Test Performance				
Constant	0.2881***	0.001	198.057	0.000
Cross Rate	0.0054***	0.002	3.279	0.001
Individual Mutation Rate	-0.0030**	0.002	-1.974	0.049
Gene Mutation Rate	-0.0028*	0.002	-1.826	0.069
Panel B: Ensemble Performance				
Constant	0.2985***	0.002	150.732	0.000
Cross Rate	0.0038*	0.002	1.700	0.090
Individual Mutation Rate	-0.0030	0.002	-1.455	0.147
Gene Mutation Rate	0.002	0.002	1.071	0.285
Panel C: Computational Time (s)				
Constant	507.28***	116.53	4.353	0.000
Cross Rate	1724.37***	130.94	13.169	0.000
Individual Mutation Rate	3159.24***	122.64	25.760	0.000
Gene Mutation Rate	111.91	122.96	0.910	0.364

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 16: Regression Results for Genetic Algorithm Parameters

Combining the results of correlation and regression analysis, it can be concluded that the c_r plays

a significant role in improving the quality of solutions (test performance) generated by GENOME, though it also increases computational runtime. This finding suggests that a higher c_r enhances the combination quality of chromosomes (model weights), leading to better solutions, albeit at the cost of greater computational effort. On the other hand, an excessively high im_r tends to disrupt superior individuals and substantially enlarge the search space, leading to increased time consumption. However, a moderate level of im_r proves beneficial as it helps prevent the population from prematurely converging. In contrast, the gm_r demonstrate only a limited impact within the scope of this study, implying its influence on the GENOME’s performance is relatively minor.

Within a broad range of hyperparameter settings, the performance of GENOME is relatively stable, suggesting it is not sensitive to hyperparameter values. The introduction of the ensemble strategy further reduces sensitivity, effectively balancing different parameter choices while maintaining solution quality.

The recommended hyperparameter settings are a c_r of 0.3 to 0.6, im_r of 0.1, and gm_r of 0.1. This setup can build superior individuals while maintaining population diversity and avoiding the negative effects of excessive mutation. Even in larger-scale or more complex problem scenarios, this configuration, through the ensemble strategy, can smooth out parameter differences and fully leverage the core advantages of GENOME, achieving relatively stable overall performance.

A.1.7 Why Gaussian Mutation?

Crossover and mutation are complementary. Under fitness-weighted crossover,

$$w_c = t w_{p1} + (1-t) w_{p2},$$

if parent perturbations have comparable covariance Σ , then

$$\text{Var}(w_c) \approx (t^2 + (1-t)^2) \Sigma \preceq \Sigma,$$

indicating variance contraction. By contrast, for mutation

$$w' = w + m \odot \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I),$$

exploration is restored on the activated dimensions, with

$$\text{Var}(w') \approx \text{Var}(w) + \text{Pr}(\text{mut}) \cdot gm_r \cdot \sigma^2 I.$$

For a smooth local objective F ,

$$F(w + \delta) = F(w) + \nabla F(w)^\top \delta + \frac{1}{2} \delta^\top H(w) \delta + O(\|\delta\|^3)$$

Thus Gaussian mutation does not guarantee monotonic improvement, but it gives positive probability to exploratory directions, while selection reduces the survival probability of harmful perturbations. This complements the empirical evidence in Table 3, where removing mutation hurts and removing crossover and mutation together hurts more.

A.2 Datasets and Expert Models

Datasets To ensure robust evaluation, we randomly sample 200 instances from each dataset as the validation set for model optimization. For the remaining data, we employ a size-based splitting strategy: for datasets with fewer than 1,000 remaining instances, we use all data as the test set; for those exceeding 1,000 instances, we retain 1,000 instances as the test set, except for multilingual tasks. For multilingual tasks (MGSM, FLORES-101), considering the scarcity of low-resource languages, we retain all remaining data as the test set to avoid potential language bias. Additionally, for FLORES-101, we create a new **FLORES-37** dataset by selecting data corresponding to the 37 languages supported by the *gemma-2* model (Team et al., 2024). Furthermore, compared to the best single experiment, in the multi-task experiment, we utilize the MELD dataset, which, together with EmoryNLP, forms the tasks in the Affective Computing.

Expert Models All models are trained on $8 \times A100$ -80GB GPUs. We visualize the performances (ranks) of our ten initial expert models across all datasets, grouped by category, in Figure 3. It is evident that each model exhibits distinct strengths, demonstrating that we have successfully trained a diverse set of experts.

A.3 Implementation Details

Ensemble For tasks with a single correct answer, we use majority voting to aggregate predictions. For tasks without a unique answer (e.g., MBPP, Flores37/101, DROP), we apply a similarity-based approach. Specifically, we take the top-k outputs, embed each using `text-embedding-3-small`, and compute pairwise similarities via BERTScore (Zhang et al., 2019). The output with the highest overall similarity score is selected as the final ensemble result.

Baseline To ensure fair comparisons, we use the same set of 10 initial expert models for all baselines except for the data merge baseline, and we otherwise align our settings as closely as possible with the original papers. LoRAHub is run for 50 iterations, and Model Swarms uses a population size of 10 for 10 iterations, with other hyperparameters set according to its original recommendations. Unless otherwise specified, all experiments are performed on $4 \times A100$ 80G GPUs and repeated 5 times.

Prompts All methods in our experiments rely on the same prompts for a fair comparison. Following OpenAI’s approach², we design specific prompts for different tasks while maintaining a minimalist design philosophy. All prompts are shown in Figure 4-9. As MMLUPRO questions can have more than four options, we implement a more flexible prompt structure (Figure 5) to accommodate the varying number of choices.

A.4 Algorithm

The pseudocode for GENOME and GENOME(+) can be found in Algorithm 1 and 2.

Algorithm 1 GENOME

Input: Task T , fitness function f , LLM experts $\{x_i\}_{i=1}^N$
Hyperparameters: population size N , crossover rate c_r , mutation rates (im_r, gm_r) , sigma σ , elite ratio α , max iteration K
// Initialization
 $P \leftarrow \{\mathbf{w}_i = t\mathbf{w}_a + (1-t)\mathbf{w}_b \mid t \sim U(0, 1)\}_{i=1}^N$
 evaluate fitness $f(x_i)$ for $x_i \in P$
 $g \leftarrow \arg \max_{x \in P} f(x)$
for iter = 1 to K do
// Crossover
for each pair selected with probability c_r **do**
 $p_i \leftarrow f_i / \sum f_k, \quad k = 1, 2, \dots, N$
 select parents (x_{p_1}, x_{p_2}) according to $\{p_i\}_i^N$
 $\mathbf{w}_{child} \leftarrow t\mathbf{w}_{p_1} + (1-t)\mathbf{w}_{p_2}$, create x_{child} by \mathbf{w}_{child}
 evaluate x_{child} on f
end for
// Mutation
for each individual with probability im_r **do**
 for each weight ω_{ij} with probability gm_r **do**
 $\mathbf{w}'_i \leftarrow \mathbf{w}_i + \mathbf{m}_i \cdot \mathbf{E}_i$, create x'_i by \mathbf{w}_i
 evaluate x'_i on f
 end for
end for
// Selection
 $P \leftarrow \text{elite}(\alpha \cdot N) \cup \text{fitness-based selection}((1 - \alpha) \cdot N)$
 $g \leftarrow \arg \max_{x \in P} f(x)$
end for
 evaluate test fitness $f_{\text{test}}(x_i)$ for $x_i \in P$, update g
return best individual g

²<https://github.com/openai/simple-evals>

Dataset	Category	Metrics	Size	
			validation	test
ARC_C	Question Answering	accuracy, 0-shot	200	1000
BBH	Logical Reasoning	accuracy, 3-shot	200	1000
CSQA	Question Answering	accuracy, 0-shot	200	1000
DROP	Logical Reasoning	exact match, 0-shot	200	1000
EmoryNLP	Affective Computing	accuracy, 0-shot	200	697
Flores-37/101	Multilingual Processing	BLEU, 3-shot	200	1012
GSM8k	Mathematics	accuracy, 0-shot	200	1000
MATH	Mathematics	accuracy, 0-shot	200	1000
MBPP	Code Generation	Pass@1, 0-shot	200	774
MELD	Affective Computing	accuracy, 0-shot	200	1000
MGSM	Multilingual Processing, Mathematics	accuracy, 0-shot	200	2637
MMLU	General Knowledge	accuracy, 0-shot	200	1000
MMLUPro	General Knowledge	accuracy, 0-shot	200	1000

Table 17: Detailed information of the datasets.

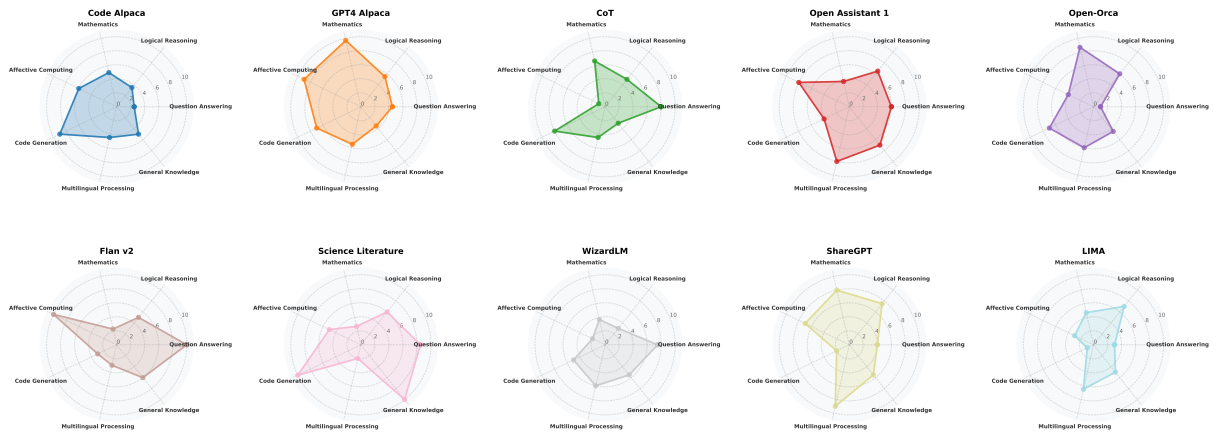


Figure 3: Capability distribution of expert models across seven dimensions, highlighting their specialized strengths.

Subset	Samples	LoRA rank/alpha	Learning rate	Warmup ratio	Batchsize	Epochs
CoT (Chung et al., 2024)	49747	8/16	2.00E-04	0.1	32	5
Code Alpaca (Chaudhary, 2023)	20016	8/16	2.00E-04	0.1	32	5
Flan v2 (Chung et al., 2024)	49123	8/16	2.00E-04	0.1	32	5
GPT4 Alpaca (Peng et al., 2023)	19906	8/16	2.00E-04	0.1	32	5
Open Assistant 1 (Köpf et al., 2024)	7331	8/16	2.00E-04	0.1	32	5
Open-Orca (Mukherjee et al., 2023)	29683	8/16	2.00E-04	0.1	32	5
Science Literature (Iverson et al., 2023b)	7468	8/16	2.00E-04	0.1	32	5
ShareGPT (TechCrunch, 2023)	111912	8/16	2.00E-04	0.1	32	1
WizardLM (Xu et al., 2023)	29810	8/16	2.00E-04	0.1	32	5
LIMA (Zhou et al., 2023)	1018	8/16	2.00E-04	0.1	32	5
Tulu-v2-sft-mixture (Iverson et al., 2023a)	326014	8/16	2.00E-04	0.1	32	1

Table 18: Training configurations and sample sizes for different subsets. The learning rate decay follows a cosine schedule.

Algorithm 2 GENOME+

Input: Task T , fitness function f , LLM experts $\{x_i\}_{i=1}^n$
Hyperparameters: population size N , crossover rate c_r , mutation rates (im_r, gm_r) , sigma σ , elite ratio α , learning weight ϕ_* , learning rate λ , max iteration K
Population Initialization
for $iter = 1$ to K **do**
 Crossover, Mutation
 // *succession*
 for $i = 1$ to N **do**
 $e_i \leftarrow \frac{1}{C}[\phi_e e_i + \phi_g (e_g - e_i) + \phi_c (e_c - e_i) - \phi_w (e_w - e_i)]$
 $w_i \leftarrow w_i + \lambda e_i$, create x_i by w_i
 end for
 Selection
end for
evaluate test fitness f_{test} for $x_i \in P$, update g
// *ensemble*
Select top- k individuals based on validation fitness f_{valid}
Combine selected individuals into an ensemble E
return best individual g and ensemble E

ARC_C

Answer the following question.
The last line of your response should be of the following format:
'Answer: \$LETTER' (without quotes)
where LETTER is one of **{candidates}**.

{question}

{options}

Let's think step by step.

Figure 6: The prompt of ARC_C.

MMLU

Answer the following **{subject}** question.
The last line of your response should be of the following format:
'Answer: \$LETTER' (without quotes)
where LETTER is one of ABCD.

{question}

A) **{A}**
B) **{B}**
C) **{C}**
D) **{D}**

Let's think step by step.

Figure 4: The prompt of MMLU.

DROP

You will be asked to read a passage and answer a question.

Write a line of the form "Answer: \$ANSWER"
at the end of your response.

{context}

Figure 7: The prompt of DROP.

MMLUPro

Answer the following **{subject}** question.
The last line of your response should be of the following format:
'Answer: \$LETTER' (without quotes)
where LETTER is one of **{candidates}**.

{question}

{options}

Let's think step by step.

Figure 5: The prompt of MMLUPro.

CSQA

Answer the following **{question_type}** question step by step.

The last line of your response should be of the following format:
'Answer: \$LETTER' (without quotes)
where LETTER is one of ABCDE.

{question}

{options}

Figure 8: The prompt of CSQA.

EmoryNLP / MELD

Given a conversation history and a current utterance, follow these steps to identify the emotion of the current utterance from the given options.

The emotion should be determined based on both the conversation context and the current utterance.

The last line of your response should be of the following format: 'Answer: \$LETTER' (without quotes) where LETTER is one of ABCDEFG.

Let's think step by step.

History:
{history}

Utterance:
{utterance}

Options:
{options}

Figure 9: The prompt of EmoryNLP and MELD.

GSM8K / MATH

Solve the following math problem step by step.

The last line of your response should be the form Answer: \$ANSWER (without quotes) where \$ANSWER is the answer to the problem.

{Question}

Remember to put your answer on its own line after "Answer:", and you do not need to use a `\boxed` command.

Figure 12: The prompt of GSM8K and MATH.

BBH

Here are some examples to help you understand how to answer the question.

{examples}

Now answer the following question, write a line of the form "Answer: \$Answer" at the end of your response.

{question}

Figure 10: The prompt of BBH.

MGSM

Solve the following math problem step by step.

Provide your explanation and final answer in English, the last line of your response should be the form Answer: \$ANSWER (without quotes) where \$ANSWER is the answer to the problem.

{Question}

Remember to put your answer on its own line after "Answer:", and you do not need to use a `\boxed` command, always respond in English.

Figure 13: The prompt of MGSM.

Flores-37/101

Translate the following sentence from English to *{language}*.

Your output should be formatted as follows:

Translation: \$SENTENCE

(where \$SENTENCE is the translated version of the sentence into *{language}*).

Below are examples to guide the translation task:

{examples}

Now translate the following sentence:

{sentence}

Figure 11: The prompt of Flores-37/101.

MBPP

You are an expert Python programmer, and here is your task:

{question}

Your code should pass these tests:

{test}

Figure 14: The prompt of MBPP.