

From Instance Selection to Fixed-Pool Data Recipe Search for Supervised Fine-Tuning

Haodong Wu¹ Jiahao Zhang² Lijie Hu² Yongqi Zhang^{1*}

¹The Hong Kong University of Science and Technology (Guangzhou)

²Mohamed bin Zayed University of Artificial Intelligence

hwu315@connect.hkust-gz.edu.cn yongqizhang@hkust-gz.edu.cn

{jiahao.zhang, lijie.hu}@mbzuai.ac.ae

Abstract

Supervised fine-tuning (SFT) data selection is commonly formulated as instance ranking: score each example and retain a top- k subset. However, effective SFT training subsets are often produced through ordered curation recipes, where filtering, mixing, and deduplication operators jointly shape the final data distribution. We formulate this problem as *fixed-pool data recipe search*: given a raw instruction pool and a library of grounded operators, the goal is to discover an executable recipe that constructs a high-quality selected subset under a limited budget of full SFT evaluations, without generating, rewriting, or augmenting training samples. We introduce **AutoSelection**, a two-layer solver that decouples fixed-pool materialization based on cached task-, data-, and model-side signals from expensive full evaluation, using warmup probes, realized subset states, local recipe edits, Gaussian-process-assisted ranking, and stagnation-triggered re-seeding. Experiments on a 90K instruction pool show that AutoSelection achieves the strongest in-distribution reasoning average across three base models, outperforming full-data training, random recipe search, random top- k , and single-operator selectors. Additional held-out reasoning, instruction-following, graph-reasoning, search-stability, structural-ablation, and 1.5B-to-7B transfer results further show that recipe structure matters beyond individual selection operators. Code is available at <https://github.com/w253/AutoSelection>.

1 Introduction

The effectiveness of supervised fine-tuning (SFT) depends critically on the quality, diversity, and composition of the instruction data used for adaptation (Zhou et al., 2023a; Liu et al., 2024; Chen et al., 2024a). Many automatic approaches formulate this problem as *instance-level data selection*

(see Figure 1(a)): they score each sample with quality heuristics, influence estimates, learned representations, or iterative utility signals, and then retain a subset (Liu et al., 2024; Xia et al., 2024; Ma et al., 2025; Lin et al., 2025). This score-and-select view is useful and has produced strong targeted selectors (Xia et al., 2024; Ma et al., 2025; Lin et al., 2025). However, it abstracts away how SFT datasets are often assembled in practice, where filtering, source mixing, deduplication, and light cleaning are applied as a multi-stage curation workflow rather than as a top- k instance-retention decision (Chen et al., 2026).

We study *fixed-pool data recipe search* for SFT. In this setting, the raw data pool is fixed, every operator is grounded in cached signals over that pool, and a candidate recipe is an ordered sequence of filtering, selection, deduplication, or set-composition operations whose execution returns a selected subset (see Figure 1(b)). The goal is not to assign one final score to each instance, but to decide which executable recipe should be evaluated next so that the best selected subset after a small number of full SFT evaluations performs as well as possible. Under this view, a conventional top- k selector is a length-one degenerate recipe: it applies one scoring rule followed by one retention rule. Fixed-pool recipe search strictly enlarges this view by making operator choice, ordering, parameters, and intermediate subset states part of the optimization object.

Fixed-pool data recipe search is complementary to recent workflow-level and LLM-driven systems. DataChef (Chen et al., 2026) formulates end-to-end data recipe generation for LLM adaptation, and LLM-AutoDP (Huang et al., 2026) uses LLM agents to generate and iteratively refine data-processing strategies with feedback from model training and evaluation. These methods show that generate-and-evaluate loops are powerful for automating data workflows. AutoSelection studies a different controlled setting: all candidates operate

*Corresponding author.




Figure 1: Conceptual contrast between instance-level selection and fixed-pool data recipe search for SFT data curation. (a) Instance selection compresses curation into one scoring view and a top- k retention step over individually scored instances. Figure 8 and Table 24 in Appendix D.4 show that subsets with similar one-dimensional metric distributions can still yield different benchmark results. (b) Fixed-pool recipe search evaluates multiple ordered recipes over the same raw pool $\mathcal{D}_0$; different operator choices, parameters, and ordering transform the same sample-ID space into different selected subsets.

on the same raw pool, and language models are not used to rewrite, synthesize, or augment training samples. This boundary is useful because it helps attribute measured differences primarily to grounded operator choices over a fixed raw pool, while reducing confounds from generation, rewriting, or newly introduced samples.

The resulting problem is a structured search problem under expensive validation. The candidate space is compositional because a recipe must choose which operators to activate, where to place them, and how to set their parameters. Operator choices interact through the intermediate subsets they produce, so the value of one step cannot be understood independently of the rest of the recipe. Moreover, each reliable observation requires recipe execution, model fine-tuning, and benchmark evaluation. The central algorithmic challenge is therefore budget allocation over interacting operator choices rather than scoring samples once.

In this work, we present **AutoSelection**, a framework for fixed-pool data recipe search in supervised fine-tuning. AutoSelection separates cheap search-side reasoning from expensive full evaluation: it caches grounded task-, data-, and model-side signals; probes multiple retention regimes during warmup; represents realized subsets with state vectors; edits current seed recipes through a Summarizer, Proposer, and Ranker; and refreshes the seed only after stagnation. When language models are used, they act only as search-side assistants for summarizing history, proposing grounded recipe edits, ranking candidates, and reseeding. Under matched full evaluation budgets, we evaluate whether this design can discover higher-scoring se-

lected subsets than fixed-pool baselines, including Random recipe search and single-operator alternatives.

Our contributions are threefold.

- **Problem formulation.** We formalize SFT data curation as *fixed-pool data recipe search*, a budgeted black-box problem over executable recipes on a fixed raw pool; this framing subsumes top- k instance selection and makes operator composition, ordering, and realized subset states explicit.
- **Method.** We propose **AutoSelection**, a two-layer solver that decouples cached fixed-pool candidate materialization from expensive full SFT evaluation, using retention warmup, state-aware local edits, GP-assisted ranking, and stagnation-triggered reseeding.
- **Empirical evidence.** On a 90K instruction pool across three base models, AutoSelection improves over full-data training, random recipe search, random top- k , and single-operator selectors, with additional OOD, stability, structural-ablation, and 1.5B-to-7B transfer analyses.

2 Related Work

2.1 Instance-level data selection

Many automatic data selection methods for instruction tuning operate on individual samples or fixed subsets. One line of work selects data by task relevance, including gradient-based influence signals in LESS (Xia et al., 2024) and model-centric activation signals in MONA (Ma et al., 2025). Another line estimates instance utility during training, as in

LEAD (Lin et al., 2025). Related preprocessing-style methods use instruction difficulty, exact or semantic redundancy, and instruction-structure diversity as selection signals (Li et al., 2024; Lee et al., 2022; Abbas et al., 2023; Bukharin et al., 2024). Diversity- and coverage-aware methods move beyond naive one-dimensional top- k selection: DEITA combines complexity, quality, and diversity, while D3 and graph-based approaches similarly balance multiple data properties (Liu et al., 2024; Zhang et al., 2025a; Wu et al., 2025). Once their scorer, diversity rule, and budget are fixed, however, they still execute a predetermined subset-construction rule. By contrast, recipe search optimizes the surrounding operator choices, order, parameters, and realized subset states under downstream feedback.

2.2 Recipe and pipeline optimization

Recent work increasingly treats workflows, recipes, and pipelines as first-class optimization objects. Data-Juicer (Chen et al., 2024a) provides configurable operator pipelines, while AutoPipe (Chwa et al., 2026) searches broader LLM post-training pipelines under compute constraints. DataChef (Chen et al., 2026) and LLM-AutoDP (Huang et al., 2026) move toward recipe-level automation, but rely on open-ended recipe generation or LLM-driven processing modules. These methods are adjacent to ours, but many target broader pipeline optimization, external LLM processing, or generated data workflows rather than controlled search over a fixed raw pool. By contrast, our work fixes the raw pool, searches finer-grained sample-level operator compositions, and uses full evaluation as the central search signal. This narrower scope enables controlled budget-matched comparison while preserving the core challenge of fixed-pool data recipe search.

2.3 HPO and AutoML

Our setting is also related to black-box hyperparameter optimization (HPO) and AutoML, where random search (Bergstra and Bengio, 2012), Bayesian optimization (Snoek et al., 2012), Hyperband-style resource allocation (Li et al., 2018), and automated pipeline search allocate limited evaluations over structured configuration spaces (Falkner et al., 2018; Feurer et al., 2015). AutoSelection borrows the budgeted-search perspective, but the optimized configurations are ordered data-curation programs whose execution changes the SFT training subset.

Thus, each observation reflects recipe execution, supervised fine-tuning, and downstream evaluation, rather than validation loss under a fixed training set. This distinction makes subset state, operator ordering, and fixed-pool attribution central to our formulation.

3 Method

We first define fixed-pool data recipe search independently of any particular solver, and then instantiate it with AutoSelection. We use **AutoSelection** to refer to the full budgeted solver described in Sections 3.2-3.3. Conceptually, AutoSelection has two coupled layers: a fixed-pool materialization layer that caches grounded signals and summarizes executed candidate subsets, and a search-controller layer that allocates full SFT evaluations through warmup, local edits, ranking, and reseeding. Figure 2 summarizes this search loop.

Running example. Consider the recipe $\text{N-gram} \rightarrow \text{MONA} \rightarrow \text{SemDedup} \rightarrow \text{random-}k$. N-gram selection first changes the lexical composition, MONA then applies task-relevance signals to the surviving pool, SemDedup removes representation-space redundancy, and random- k controls the final retained scale. After materialization, AutoSelection observes the realized retained-example and token ratios, task-relevance profile, difficulty, uncertainty, and distribution drift before allocating a full evaluation. This separates the recipe syntax from the state of the subset it actually produces.

3.1 Fixed-pool data recipe search

Definition. Let $\mathcal{D}_0 = \{d_1, \dots, d_N\}$ be a fixed raw data pool and let $\mathcal{D}'_0 = \text{Canon}(\mathcal{D}_0)$ be its canonicalized executable representation. Canonicalization normalizes fields, attaches source and execution metadata, and preserves stable sample identifiers, so $\mathcal{D}'_0$ contains the same N underlying samples as $\mathcal{D}_0$. Let f_0 be the base model, $\mathcal{E}$ the evaluation suite, and $\mathcal{O}$ a shared library of grounded operators. Each grounded operator is a subset transformer over the fixed canonicalized pool:

$$o(\cdot; \theta) : 2^{\mathcal{D}'_0} \rightarrow 2^{\mathcal{D}'_0}, \quad o \in \mathcal{O}, \theta \in \Theta_o,$$

where $2^{\mathcal{D}'_0}$ denotes the set of all subsets of $\mathcal{D}'_0$. The parameter θ contains the operator-specific execution choices, such as thresholds and retained sizes. Thus, an operator may filter, select, deduplicate, or recombine samples, but it never creates sample identifiers outside the fixed pool.




Figure 2: AutoSelection as a solver for fixed-pool data recipe search. All candidate recipes operate on the same canonicalized pool derived from the fixed raw pool. Warmup probes initialize the search across retention regimes and set the first seed anchor. During search, the Summarizer converts evaluated history into guidance, the Proposer generates local recipe edits, the Ranker chooses one materialized candidate for full evaluation, and the Reseeder refreshes the seed anchor after stagnation.

A data recipe is a bounded variable-length ordered program

$$r = ((o_\ell, \theta_\ell))_{\ell=1}^{L(r)}, \quad (1)$$

$$o_\ell \in \mathcal{O}, \quad \theta_\ell \in \Theta_{o_\ell}, \quad L(r) \leq L_{\max}.$$

where $L(r)$ is the recipe length and $L_{\max}$ is the maximum admissible length. Let $\mathcal{R}$ denote the recipe space induced by $\mathcal{O}$ and the admissible parameter sets. Executing a recipe produces a subset, fine-tuned model, and observed utility:

$$S(r) = \text{Exec}(\mathcal{D}'_0, r), \quad (2)$$

$$\hat{f}_r = \mathcal{A}_{\text{SFT}}(f_0, S(r)), \quad y(r) = \text{Eval}_{\mathcal{E}}(\hat{f}_r).$$

Here $y(r)$ is the observed downstream utility of the recipe after full evaluation.

The fixed-pool recipe search problem is to adaptively choose at most B recipes to evaluate. For the t -th full evaluation, we write r_t for the queried recipe, $S_t = S(r_t)$ for its selected subset, and $y_t = y(r_t)$ for its observed score. A finite-budget solver returns the best observed recipe and subset:

$$t^* \in \arg \max_{1 \leq t \leq B} y_t, \quad r^* = r_{t^*}, \quad S^* = S_{t^*}. \quad (3)$$

One full evaluation means executing a recipe, fine-tuning f_0 on the selected subset, and evaluating the fine-tuned model on $\mathcal{E}$. The budget B is therefore the primary resource constraint.

Single-operator selection as a degenerate recipe.

Let $\mathcal{U}$ be the set of single-operator selectors included in $\mathcal{O}$, and let Θ_u be the admissible parameter set for $u \in \mathcal{U}$. If $\mathcal{R}$ contains all length-one recipes, then each single-operator baseline is a recipe $r_{u,\theta} = ((u, \theta)) \in \mathcal{R}$. Therefore,

$$\max_{r \in \mathcal{R}} y(r) \geq \max_{u \in \mathcal{U}, \theta \in \Theta_u} y(r_{u,\theta}).$$

This containment statement does not imply that AutoSelection attains the global maximum under a finite budget B ; it only formalizes that the recipe space is at least as expressive as the included single-operator selectors. Appendix B.3 provides the proof and a concrete top- k example.

3.2 Making the fixed pool searchable

Given the formulation above, a practical solver must address two issues. First, many candidate recipes need to be materialized without repeatedly recomputing expensive sample-level signals. Second, candidate recipes that look similar syntactically can produce very different realized subsets. AutoSelection addresses these issues through a cold-start cache and a state-vector abstraction.

During cold start, the raw samples are canonicalized into $\mathcal{D}'_0$ with normalized instruction-response fields, source metadata, and stable identifiers. The solver then precomputes reusable task-, data-, and model-side signals over this canonicalized pool. Task-side signals include benchmark-conditioned

activation-similarity statistics following MONA (Ma et al., 2025). Data-side signals summarize intrinsic properties such as lexical or instruction-structure diversity. Model-side signals summarize cached responses of the base model, such as instruction-following difficulty (Li et al., 2024), varentropy (Kontoyiannis and Verdú, 2014), and sparse-activation statistics. Because these signals are grounded in $\mathcal{D}'_0$, different recipes can reuse the same cached measurements while still producing different selected subsets.

For any materialized candidate recipe r , evaluated or not, AutoSelection summarizes the resulting subset with a state vector

$$\begin{aligned} z(r) &= \phi(S(r)) \\ &= \left[z^{\text{task}}(r); z^{\text{data}}(r); z^{\text{model}}(r) \right]. \end{aligned} \quad (4)$$

The task block contains benchmark-conditioned relevance statistics, the data block contains realized scale statistics such as retained-example and token ratios, and the model block contains cached model-side summaries such as IFD, varentropy, and sparse-activation distribution drift. For an evaluated recipe r_t , we write $z_t = z(r_t)$. State vectors let the search controller compare realized subset properties before spending a full SFT evaluation on a candidate. Tables 11 and 12 in Appendix C.3 give the field-level definition and diagnostics.

3.3 Navigating the recipe space

The search-controller layer addresses the exploration–exploitation trade-off induced by expensive full evaluations. Given a limited evaluation budget, AutoSelection should not exploit around an arbitrary initial recipe too early, but it also cannot spend the budget on uniform exploration over the entire compositional recipe space. Therefore, it begins with a small warmup exploration stage that probes 3 data-scale regimes before switching to seed-centered local refinement.

This warmup stage resolves an early high-impact uncertainty: prior instruction-tuning studies show that downstream behavior can change substantially with the amount of retained data, and that smaller, better-curated subsets can sometimes match or outperform much larger ones (Jha et al., 2023; Liu et al., 2024; Iverson et al., 2025). We sample candidate recipes from $\mathcal{R}$, monitor the retained-example ratio after execution on $\mathcal{D}'_0$, and keep three probe recipes that fall into low-, medium-, and high-retention bins. Evaluating these probes provides

initial evidence about which retention regime is promising, sets the first seed recipe r_1^{seed} , and initializes the search history $\mathcal{H}_3 = \{(r_i, z_i, y_i)\}_{i=1}^3$. The initial seed anchor is set to the best warmup recipe,

$$r_1^{\text{seed}} \in \arg \max_{(r, z, y) \in \mathcal{H}_3} y.$$

For each later evaluation step $t > 3$, AutoSelection uses the accumulated history $\mathcal{H}_{t-1} = \{(r_i, z_i, y_i)\}_{i=1}^{t-1}$. The *Summarizer* reads $\mathcal{H}_{t-1}$ and produces search guidance $g_{t-1} = \text{Summarizer}(\mathcal{H}_{t-1})$, a short set of data-backed hypotheses about which operators, retention levels, or compositions appear promising or risky. These findings are summaries of the evaluated recipes, not edits to the raw pool, and they only steer the next local proposal.

The Proposer then samples a sibling candidate set around the current seed anchor:

$$\begin{aligned} \mathcal{C}_t &= \{r'_{t,j}\}_{j=1}^{M_t}, \\ r'_{t,j} &\sim Q_{\text{prop}}(\cdot \mid r_m^{\text{seed}}, g_{t-1}, \mathcal{H}_{t-1}). \end{aligned}$$

Here $M_t = |\mathcal{C}_t|$, r' denotes a candidate recipe, and Q_{prop} is a local-edit proposal policy over operations such as inserting, deleting, swapping, or retuning recipe steps. All proposed recipes are constrained to the operator catalog and are validated before execution. Each candidate recipe $r' \in \mathcal{C}_t$ is executed on the cached pool to obtain $S(r')$ and $z(r')$, but it is not yet used for SFT. Following standard Bayesian optimization practice for expensive ML evaluations (Snoek et al., 2012), a GP surrogate fitted on previous recipe encodings provides a cheap score prior; Appendix B.1 gives the compact recipe-vector example used for this surrogate feature.

$$\hat{\mu}_{t-1}(r'), \hat{\sigma}_{t-1}(r') = \text{GP}_{t-1}(\psi(r')). \quad (5)$$

The surrogate uses only the recipe encoding; realized state summaries are passed to the Ranker instead of entering the surrogate. The *Ranker* combines recipe structure, state-vector information, the GP prior, and search history to choose one candidate for full evaluation; for compactness in Eq. 6, $\hat{\mu}$ and $\hat{\sigma}$ denote $\hat{\mu}_{t-1}(r')$ and $\hat{\sigma}_{t-1}(r')$:

$$\begin{aligned} \rho_t(r') &= f_{\text{rank}}(r', z(r'), \hat{\mu}, \hat{\sigma}, \mathcal{H}_{t-1}), \\ r_t &\in \arg \max_{r' \in \mathcal{C}_t} \rho_t(r'). \end{aligned} \quad (6)$$

Only r_t is evaluated, and the resulting triple (r_t, z_t, y_t) is appended to the history.

The seed anchor and incumbent are tracked separately. The incumbent (r^*, S^*, y^*) is the best observed result, whereas r_m^{seed} defines the local proposal neighborhood for the current search phase. Keeping the seed fixed within a phase avoids moving the neighborhood after every noisy single evaluation. When the search has not improved the incumbent for P consecutive evaluations, the Reseeder refreshes the anchor:

$$r_{m+1}^{\text{seed}} \sim Q_{\text{seed}}(\cdot \mid \mathcal{H}_t),$$

where Q_{seed} is a history-conditioned policy that selects a new promising motif or recipe region. This mechanism provides exploration after stagnation while keeping most full evaluations focused on local exploitation. Algorithm 1 in Appendix B.2 gives the full pseudocode, and Appendix B.3 summarizes the search-side complexity.

4 Experiments

4.1 Experimental setup

We evaluate AutoSelection in the fixed-pool setting: all methods receive the same raw SFT pool, use no synthetic data generation, LLM-based sample rewriting, or pool augmentation, and are compared by selected-subset quality. The raw pool is a 90K-sample merged instruction-tuning pool with 30K samples each from OpenHermes-2.5 (Teknium, 2023), the LESS instruction-tuning data pool (Xia et al., 2024), and Alpaca-52K (Wang et al., 2023b). We use GPQA (Rein et al., 2024), GSM8K (Cobbe et al., 2021), BBH (Srivastava et al., 2023), and MMLU (Hendrycks et al., 2021) as the validation suite, and evaluate three base models—Qwen2.5-1.5B (Yang et al., 2025), Llama3.2-1B (Grattafiori et al., 2024), and Qwen2.5-3B (Yang et al., 2025)—to cover model families and scales. All methods and models use the same fixed randomly sampled 10% MMLU draw. The search budget is $B = 15$ full evaluations, with stagnation patience $P = 4$ for reseeding and proposer candidate set size $|\mathcal{M}_t| = 5$. The main table decomposes the first independent AutoSelection run by benchmark; aggregate comparisons use the three-run median. All full evaluations use one fixed SFT configuration (Table 5); implementation details and exact search and evaluation prompts are in Appendices A.3 and E. We also test the selected subsets and their recipes on two OOD graph benchmarks, GraphWiz (Chen et al., 2024b) and NLGraph (Wang et al., 2023a).

We use a shared operator library for all search methods. The library covers task relevance (MONA (Ma et al., 2025)), model-internal difficulty and uncertainty (IFD (Li et al., 2024) and varentropy (Kontoyiannis and Verdú, 2014; Maadani et al., 2020; Li et al., 2026; Zhang et al., 2025b)), intrinsic data diversity (N-gram entropy (Shannon, 1948; Wu et al., 2025) and action-object branching (Zhao et al., 2023; Bukharin et al., 2024)), redundancy reduction (SemDedup (Abbas et al., 2023)), stochastic exploration (Random top- k (Bergstra and Bengio, 2012; Diddee and Ippolito, 2025; Ivison et al., 2025; Nayak et al., 2026)), and set composition (Mix). The same library defines the search space and the Random recipe search baseline; Table 4 in Appendix A.2 summarizes its operators and parameters. To avoid comparing against weak default settings, each in-distribution single-operator baseline uses its best tested per-model operating point over 17 full SFT evaluations: three settings for each of five top- k selectors and two SemDedup thresholds. Tables 6–8 in Appendix A.4 report the complete sweeps. The OOD graph columns retain the original shared-operating-point evaluations and are therefore a separate, untuned evaluation slice. To probe a neighboring LLM-driven pipeline, we also run an early LLM-AutoDP (Huang et al., 2026) pilot; its 25.98 average and protocol limitations are reported in Table 22 of Appendix D.2 rather than the main fixed-pool comparison.

4.2 Main results

As shown in Table 1, AutoSelection achieves the highest in-distribution reasoning average for all three base models. Using the three-run median and the strongest per-model-tuned fixed selector, the gains are 0.70, 1.69, and 1.23 points on Llama3.2-1B, Qwen2.5-1.5B, and Qwen2.5-3B, respectively. The preferred selector itself changes from MONA to AO to Varentropy, indicating model-dependent operating points. To compare recipe search with an existing selector that jointly scores complexity, quality, and diversity, we further evaluate DEITA (Liu et al., 2024) at two same-pool budgets. Its best scores of 14.78, 39.79, and 44.19 remain below the AutoSelection medians of 19.73, 42.23, and 49.01; Table 9 in Appendix A.5 gives the full comparison.

The held-out graph columns show that AutoSelection remains competitive: it ranks first on Qwen2.5-1.5B, nearly ties the best baseline

Group	Method	In-distribution reasoning					OOD graph		
		GPQA	GSM8K	BBH	MMLU	Avg	GraphWiz	NLGraph	Avg
Llama3.2-1B									
Full	Full data	17.19	16.60	4.78	33.06	17.91	33.62	57.25	45.44
Recipe	Random	15.62	13.49	9.45	32.33	17.72	32.16	53.55	42.86
Top-k	Random	19.41	14.10	8.47	33.33	18.83	32.43	53.22	42.83
Top-k	MONA	23.90	10.50	8.20	33.50	19.03	32.09	54.67	43.38
Top-k	AO	20.30	8.60	11.50	34.60	18.75	31.21	55.00	43.11
Top-k	IFD	19.90	5.60	5.10	30.20	15.20	29.50	50.80	40.15
Top-k	N-gram	22.30	10.20	2.70	34.60	17.45	32.65	50.48	41.57
Dedup	SemDedup	13.80	6.20	11.10	27.60	14.68	30.56	49.35	39.96
Top-k	Varentropy	23.30	7.70	10.40	33.80	18.80	32.43	48.70	40.57
Recipe	AutoSelection	26.78	11.52	6.08	34.55	19.73	33.59	55.48	44.54
Qwen2.5-1.5B									
Full	Full data	21.20	55.26	22.93	55.94	38.83	36.96	54.03	45.50
Recipe	Random	23.43	52.23	27.39	56.83	39.97	36.88	50.16	43.52
Top-k	Random	21.20	52.16	27.82	56.38	39.39	38.71	53.22	45.97
Top-k	MONA	21.20	54.35	24.89	55.77	39.05	37.18	56.45	46.82
Top-k	AO	23.66	49.81	33.36	55.33	40.54	37.65	57.74	47.70
Top-k	IFD	23.21	23.88	24.45	56.22	31.94	37.93	57.09	47.51
Top-k	N-gram	24.10	51.63	19.13	57.38	38.06	37.28	52.25	44.77
Dedup	SemDedup	14.73	40.40	20.43	56.88	33.11	38.46	53.54	46.00
Top-k	Varentropy	20.31	38.51	25.76	55.50	35.02	35.71	52.90	44.31
Recipe	AutoSelection	29.01	54.58	30.00	55.33	42.23	38.91	58.54	48.73
Qwen2.5-3B									
Full	Full data	23.66	64.82	31.95	61.00	45.36	37.31	64.03	50.67
Recipe	Random	21.43	64.44	30.22	62.00	44.52	35.71	66.29	51.00
Top-k	Random	21.87	64.59	31.73	61.94	45.03	36.75	56.29	46.52
Top-k	MONA	21.21	67.10	33.91	62.94	46.29	38.56	59.19	48.88
Top-k	AO	22.77	60.58	34.89	58.78	44.26	34.62	66.12	50.37
Top-k	IFD	27.46	28.66	31.85	61.17	37.29	36.25	55.80	46.03
Top-k	N-gram	26.34	61.71	31.20	63.06	45.58	36.15	57.41	46.78
Dedup	SemDedup	27.23	61.11	35.65	61.67	46.42	35.87	61.93	48.90
Top-k	Varentropy	25.00	68.30	36.20	61.60	47.78	35.81	56.61	46.21
Recipe	AutoSelection	22.99	72.78	36.84	63.72	49.08	36.28	65.65	50.97

Table 1: Main results across model scales. *Group* separates full-pool training, recipe-level methods, top-*k* single-operator selection, and deduplication baselines. The best score in each metric column is bolded and the second-best score is underlined.

on Qwen2.5-3B, and is the strongest non-full-data method on Llama3.2-1B. We further evaluate Qwen2.5-1.5B on ARC-Challenge (Clark et al., 2018), SVAMP (Patel et al., 2021), and IFEval (Zhou et al., 2023b), none of which enters the search objective or the original MONA targets. AutoSelection obtains a held-out macro of 53.35, matching Full data at 53.33 and exceeding the strongest evaluated selection baseline at 51.89. Removing MONA from the operator catalog still reaches 41.33, above AO at 40.54 and Full data at 38.83. Tables 20 and 21 in Appendix D.1 give the full results. Overall, recipe search finds stronger compositions without relying on one universally best selector or collapsing on the held-out checks.

4.3 Search stability under randomness

To check whether the gain is caused by lucky stochastic decisions, we run three independent 15-

evaluation searches for every model. The resulting means are 19.72 ± 0.04 , 42.07 ± 0.33 , and 49.03 ± 0.04 for Llama3.2-1B, Qwen2.5-1.5B, and Qwen2.5-3B. Even the lowest run remains above the strongest per-model-tuned selector by 0.64, 1.15, and 1.22 points, respectively. Table 16 in Appendix C.7 reports all runs, with detailed Qwen2.5-1.5B trajectories in Table 17 and Figure 7.

4.4 Search-side ablations

The ablation in Table 2 is intended as a search-policy diagnostic rather than the main evidence for AutoSelection. We use the strongest complete 1.5B AutoSelection run as the full-reference setting. Except for w/o Warmup, which removes the initial retention-regime probes to test warmup sensitivity, each ablation starts from the same warmup recipes and then removes one post-warmup search component. For w/o Ranker, candidates are selected

Method	Primary objective: best observed recipe					Trajectory diagnostics	
	GPQA	GSM8K	BBH	MMLU	Best@4–15	Mean@4–15	GapArea@4–15
Full AutoSelection	24.55	58.00	29.02	57.56	42.28	39.85	1.02
w/o Warmup	25.45	53.90	29.24	56.56	41.29	38.07	0.51
w/o Reseeder	23.66	56.18	29.13	57.00	41.49	40.36	0.71
w/o Ranker	23.44	53.30	33.59	56.83	41.79	38.97	2.00
w/o State Vectors	28.35	50.57	30.33	56.72	41.49	39.57	1.79
w/o Summarizer	23.66	54.66	29.67	56.33	41.08	39.87	0.48
Random Select	25.89	53.90	28.37	56.50	41.17	39.66	1.27

Table 2: Search-side ablations on the 1.5B setting. The primary objective is the best observed post-warmup recipe under the fixed budget; benchmark columns decompose this selected recipe. Mean@4–15 and GapArea@4–15 are trajectory diagnostics over post-warmup steps 4–15, not optimization targets.

Variant	Code	Avg	Std	Drop
Reference	A→B→C→D	41.61	0.47	–
Swap MONA/N-gram	B→A→C→D	40.68	0.59	0.93
SemDedup early	A→C→B→D	41.10	0.57	0.51
No SemDedup	A→B→D	40.94	0.26	0.67
Mix replacement	A→B→E→D	40.72	0.53	0.89

Table 3: Matched-seed structural ablation of the selected Qwen2.5-1.5B recipe. Each variant uses the same random- k seeds as the reference recipe while keeping operator parameters and the evaluation protocol fixed. Drop is the mean paired decrease; recipe codes are defined in Appendix C.6.

using only GP scores. Best@4–15 is the primary metric, while Mean@4–15 and GapArea@4–15 summarize trajectory smoothness and exposure to low-scoring edits. Under this objective, the full setting reaches the highest post-warmup score; the remaining columns summarize how different components affect the search trajectory, with the corresponding curves shown in Figure 3 in Appendix C.1. The diagnostics suggest three main component roles rather than monotonic improvements on every trajectory statistic. 1) Warmup and summarization improve budget use in different ways: removing warmup makes the search spend early evaluations in a low-scoring region, while removing the Summarizer keeps a relatively smooth trajectory but finds a weaker best recipe, suggesting that history summaries help turn past outcomes into sharper local guidance. 2) Reseeding trades short-term smoothness for peak discovery: w/o Reseeder is comparatively stable, but its best-so-far curve flattens below the full method, indicating that reseeding helps escape saturated local regions. 3) Ranking and state vectors help screen risky edits: GP-only selection outperforms random candidate choice on Best@4–15, suggesting that the surrogate provides useful coarse direction, but the full

Ranker is more stable; removing state vectors exposes the search to candidates whose recipe form may look promising but whose realized subset state is weaker. The extended Random Select curve remains volatile (Figure 4), the rolling GP fit is mixed (Figure 5), and a five-decision Ranker audit places the true best candidate in its top two every time (Tables 18 and 19); these supporting diagnostics are in Appendices C.2–C.8. Overall, these ablations suggest that AutoSelection’s components mainly improve budget allocation and candidate screening, while full SFT evaluation remains necessary for confirming the best recipe.

4.5 Matched-seed structural ablations

We further examine whether the selected recipe is sensitive to structure beyond its operator set. Because the reference Qwen2.5-1.5B recipe contains a random- k operator, we use a matched-seed design: the reference recipe and each structural variant are executed with the same random seeds (42,256,1024), while keeping operator parameters and the evaluation protocol fixed. This controls the stochastic component of the random- k step and makes the paired drop from the reference recipe the relevant comparison. As shown in Table 3, the reference recipe obtains the highest mean score, while all structural variants show positive mean paired drops. These results suggest that, for this reference recipe, ordering and composition can affect performance beyond the operator set alone, supporting our recipe-level formulation of fixed-pool data optimization. Scale-matched recipe cases and the 12-point score ranges within large retention bins show that retained size alone does not explain the differences (Figure 6 and Tables 13 and 15 in Appendix C.5). Across 181 full-evaluation records from nine searches, MONA→SemDedup and N-gram→MONA are also more frequent in the top

score tertile (Table 14); this is a retrospective diagnostic rather than evidence that either motif is universally optimal.

4.6 Transferability to larger models

Motivated by prior evidence that small models can provide useful signals for larger-model data selection and mixture optimization (Mekala et al., 2024; Yang et al., 2024; Xie et al., 2023), we conduct a 7B transfer check with Qwen2.5-7B (Yang et al., 2025) as supporting analysis. Specifically, selected subsets obtained during the Qwen2.5-1.5B search are used to fine-tune a larger model, covering subsets produced by both strong and weaker runs. The observed trend is coarse but informative: selected subsets from stronger 1.5B runs tend to remain competitive after transfer, while selected subsets from weaker runs do not reliably become strong solely because the target model is larger (see Table 23 in Appendix D.3). The 1.5B and 7B selected-subset rankings have a positive Spearman rank correlation of $\rho = 0.82$, suggesting moderate cross-scale similarity but not exact rank preservation. We therefore treat cross-scale transfer as an approximate data-quality signal. This one 1.5B-to-7B transfer check suggests that some selected-subset motifs may remain useful across scale.

5 Conclusion

We introduce AutoSelection, a fixed-pool data recipe search framework for supervised fine-tuning. Rather than ranking individual examples once, AutoSelection searches ordered combinations of filtering, mixing, deduplication, and selection operators under a limited full-evaluation budget. Cached task-, data-, and model-side signals, warmup probes, history summarization, candidate ranking, and re-seeding enable efficient exploration of compositional curation strategies. Across model scales and independent runs, AutoSelection improves the in-distribution reasoning average over full-data training, tuned single-operator selection, Random top- k , Random recipe search, and a same-pool diversity-aware selector. It remains competitive on held-out graph benchmarks and matches full-data training on the added held-out reasoning and instruction-following macro. Structural ablations and case studies further show that operator ordering and composition matter, supporting recipe-level optimization for fixed-pool data selection.

Limitations

This work focuses on fixed-pool data recipe search under a controlled evaluation protocol. While AutoSelection improves SFT performance across the studied settings, our evaluation is still limited to a moderate-size instruction pool, several model scales, and a finite set of reasoning benchmarks. Future work can extend the evaluation to larger raw pools, more heterogeneous data sources, and domain-specific SFT tasks to better understand the generality of recipe-level data optimization. In addition, AutoSelection relies on full evaluation, where each evaluation requires recipe execution, model fine-tuning, and benchmark evaluation. This makes scaling to larger models or longer search budgets computationally expensive. Developing cheaper proxy evaluations, multi-fidelity search strategies, or transferable recipe priors is therefore an important direction for making fixed-pool data recipe search more scalable.

Ethical Considerations

This work uses publicly available instruction-tuning datasets, open-weight models, and standard academic benchmarks. We do not collect private or sensitive user data, nor do we generate new training samples during the fixed-pool search process. The main ethical considerations concern the responsible use of publicly released datasets and models, including potential biases inherited from the underlying data sources and base models. We will release code and derived artifacts with appropriate documentation to support transparent and reproducible use.

Acknowledgments

This work was partially supported by HKUST(GZ) - Kunpeng&Ascend Center of Cultivation, Guangdong Basic and Applied Basic Research Foundation Project 2025A1515010304, Guangdong Province Project 2024QN11X088, Guangzhou Science and Technology Planning Project 2025A03J4491.

References

Amro Abbas, Kushal Tirumala, Dániel Simig, Surya Ganguli, and Ari S. Morcos. 2023. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *arXiv preprint arXiv:2303.09540*.

- James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2).
- Alexander Bukharin, Shiyang Li, Zhengyang Wang, Jingfeng Yang, Bing Yin, Xian Li, Chao Zhang, Tuo Zhao, and Haoming Jiang. 2024. Data diversity matters for robust instruction tuning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3411–3425.
- Daoyuan Chen, Yilun Huang, Zhijian Ma, Hesun Chen, Xuchen Pan, Ce Ge, Dawei Gao, Yuexiang Xie, Zhaoyang Liu, Jinyang Gao, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024a. Data-juicer: A one-stop data processing system for large language models. In *SIGMOD*.
- Nuo Chen, Yuhua Li, Jianheng Tang, and Jia Li. 2024b. Graphviz: An instruction-following language model for graph computational problems. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 353–364.
- Yicheng Chen, Zerun Ma, Xinchun Xie, Yining Li, and Kai Chen. 2026. Datachef: Cooking up optimal data recipes for llm adaptation via reinforcement learning. *arXiv preprint arXiv:2602.11089*.
- Channe Chwa, Xinle Wu, and Yao Lu. 2026. Automatic configuration of llm post-training pipelines. *arXiv preprint arXiv:2603.18773*.
- 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](#). *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Harshita Diddee and Daphne Ippolito. 2025. Chasing random: Instruction selection strategies fail to generalize. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1943–1957.
- Stefan Falkner, Aaron Klein, and Frank Hutter. 2018. [BOHB: Robust and efficient hyperparameter optimization at scale](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1437–1446. PMLR.
- Matthias Feurer, Aaron Klein, Katharina Eggenberger, Jost Springenberg, Manuel Blum, and Frank Hutter. 2015. [Efficient and robust automated machine learning](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Aaron Grattafiori, Abhimanyu Dubey, and 1 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Wei Huang, Anda Cheng, Yingui Wang, Lei Wang, and Tao Wei. 2026. Llm-autodp: Automatic data processing via LLM agents for model fine-tuning. *arXiv preprint arXiv:2601.20375*.
- Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. 2024. [Sparse autoencoders find highly interpretable features in language models](#). In *The Twelfth International Conference on Learning Representations*.
- Hamish Ivison, Muru Zhang, Faeze Brahman, Pang Wei Koh, and Pradeep Dasigi. 2025. Large-scale data selection for instruction tuning. *arXiv preprint arXiv:2503.01807*.
- Aditi Jha, Sam Havens, Jeremy Dohmann, Alexander Trott, and Jacob Portes. 2023. Limit: Less is more for instruction tuning across evaluation paradigms. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following*.
- Ioannis Kontoyiannis and Sergio Verdú. 2014. [Optimal lossless data compression: Non-asymptotics and asymptotics](#). *IEEE Transactions on Information Theory*, 60(2):777–795.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. Deduplicating training data makes language models better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Roshtamizadeh, and Ameet Talwalkar. 2018. [Hyperband: A novel bandit-based approach to hyperparameter optimization](#). *Journal of Machine Learning Research*, 18(185):1–52.
- Ming Li, Yong Zhang, Zhitao Li, Jiahui Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. 2024. From quantity to quality: Boosting LLM performance with self-guided data selection

- for instruction tuning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7602–7635.
- Xianzhi Li, Ethan Callanan, Abdellah Ghassel, and Xiaodan Zhu. 2026. [Entropy-gated branching for efficient test-time reasoning](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5054–5069, Rabat, Morocco. Association for Computational Linguistics.
- Xiaotian Lin, Yanlin Qi, Yizhang Zhu, Themis Palpanas, Chengliang Chai, Nan Tang, and Yuyu Luo. 2025. Lead: iterative data selection for efficient llm instruction tuning. *Proceedings of the VLDB Endowment*, 19(3):426–439.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning. In *The Twelfth International Conference on Learning Representations*.
- Da Ma, Gonghu Shang, Zhi Chen, Libo Qin, Yijie LUO, Hongshen Xu, Lei Pan, Shuai Fan, Kai Yu, and Lu Chen. 2025. Task-specific data selection for instruction tuning via monosemantic neuronal activations. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Saeid Maadani, Gholam Reza Mohtashami Borzadaran, and Abdolhamid Rezaei Roknabadi. 2020. [A new generalized varentropy and its properties](#). Preprint.
- Dheeraj Mekala, Alex Nguyen, and Jingbo Shang. 2024. [Smaller language models are capable of selecting instruction-tuning training data for larger language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10456–10470, Bangkok, Thailand. Association for Computational Linguistics.
- Nihal V Nayak, Paula Rodriguez-Diaz, Neha Hulkund, Sara Beery, and David Alvarez-Melis. 2026. A critical look at targeted instruction selection: Disentangling what matters (and what doesn’t). *arXiv preprint arXiv:2602.14696*.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094, Online. Association for Computational Linguistics.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- Claude E. Shannon. 1948. [A mathematical theory of communication](#). *Bell Syst. Tech. J.*, 27:623–656.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, and 1 others. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*. Featured Certification.
- Teknum. 2023. [Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants](#).
- Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. 2023a. Can language models solve graph problems in natural language? *Advances in Neural Information Processing Systems*, 36:30840–30861.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khoshdel, and Hannaneh Hajishirzi. 2023b. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 13484–13508.
- Minghao Wu, Thuy-Trang Vu, Lizhen Qu, and Gholamreza Haffari. 2025. [The best of both worlds: Bridging quality and diversity in data selection with bipartite graph](#). In *Forty-second International Conference on Machine Learning*.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. 2024. [LESS: Selecting influential data for targeted instruction tuning](#). In *Forty-first International Conference on Machine Learning*.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy S Liang, Quoc V Le, Tengyu Ma, and Adams Wei Yu. 2023. [Doremi: Optimizing data mixtures speeds up language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 69798–69818. Curran Associates, Inc.
- 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, and 1 others. 2025. [Qwen2.5 technical report](#). Preprint, arXiv:2412.15115.
- Yu Yang, Siddhartha Mishra, Jeffrey Chiang, and Baharan Mirzasoleiman. 2024. [Smalltolarge \(s2l\): Scalable data selection for fine-tuning large language models by summarizing training trajectories of small models](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 83465–83496. Curran Associates, Inc.

- Jia Zhang, Chen-Xi Zhang, Yao Liu, Yi-Xuan Jin, Xiao-Wen Yang, Bo Zheng, Yi Liu, and Lan-Zhe Guo. 2025a. [D3: Diversity, difficulty, and dependability-aware data selection for sample-efficient llm instruction tuning](#). In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 8348–8356. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Jinghan Zhang, Xiting Wang, Fengran Mo, Yeyang Zhou, Wanfu Gao, and Kunpeng Liu. 2025b. [Entropy-based exploration conduction for multi-step reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 3895–3906, Vienna, Austria. Association for Computational Linguistics.
- Yingxiu Zhao, Bowen Yu, Binyuan Hui, Haiyang Yu, Fei Huang, Yongbin Li, and Nevin L Zhang. 2023. A preliminary study of the intrinsic relationship between complexity and alignment. *arXiv preprint arXiv:2308.05696*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. 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)*, pages 400–410, 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, and 1 others. 2023a. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36:55006–55021.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023b. [Instruction-following evaluation for large language models](#).

A Experimental Protocol and Operator Details

A.1 Use of LLMs

We used LLMs in two distinct ways. For manuscript preparation, an LLM was used only for language editing. As part of the research pipeline, we used DeepSeek-R1 (Guo et al., 2025) as the backend of our system. All methodological decisions, experiments, analyses, and final claims were made by the authors.

A.2 Operator Library

Table 4 summarizes the operators used in our experiments. The library is designed to cover heterogeneous selection signals rather than a single definition of data quality. These operators include task-relevance signals, model-internal difficulty and uncertainty signals, intrinsic diversity signals, redundancy-reduction operations, and recipe-level combinators. For single-operator baselines, we apply one operator at a time. For AutoSelection, the same operators form the search space from which candidate recipes are constructed.

Operator	Signal / Criterion
MONA top- k	Benchmark-conditioned activation similarity
IFD top- k	Instruction-data difficulty
Varentropy top- k	Token-level uncertainty fluctuation
N-gram top- k	N-gram entropy / lexical diversity
AO top- k (action-object)	Action-object pattern diversity
SemDedup	Embedding-space near-duplicate similarity
Random top- k	Uniform sampling from the current candidate pool
Mix	Set union of two selected subsets

Table 4: Summary of the operator library used in AutoSelection. Each operator provides a distinct signal or composition rule for constructing candidate subsets.

Operator roles. MONA provides a direct task-relevance signal from benchmark-conditioned model representations. For multiple validation tasks, we compute task-specific similarity scores, select top- k samples for each task, and merge the selected sets by union. IFD and varentropy are model-internal signals: IFD measures how difficult an instruction-response pair is for the base model, while varentropy captures fluctuations in token-level predictive uncertainty. N-gram entropy and action-object branching (AO) are lightweight

data-side diversity signals, targeting lexical diversity and instruction-structure diversity respectively; they are intended as interpretable axes of variation rather than standalone claims about data quality. SemDedup reduces near-duplicate samples in the embedding space.

Random top- k acts as a stochastic escape operator. Because deterministic filters can repeatedly return overlapping subsets, random sampling allows the search to jump to alternative data regions under the same budget. This follows the motivation of random search in structured configuration spaces (Bergstra and Bengio, 2012) and is further supported by recent instruction-data selection studies where random subsets remain strong and difficult to consistently outperform (Diddee and Ippolito, 2025; Ivison et al., 2025).

In addition to filtering operators, AutoSelection includes a union-style combinator for recipe construction. When this operation is selected, the current candidate subset is mixed with a previously strong subset from the search history. The operation is implemented as a set union over sample identifiers, so duplicate samples are retained only once. This allows the search process to reuse effective data regions while still exploring new operator compositions.

Operator computation. Given a score $s(x)$ over pool D , top-fraction α keeps $\text{Top}_{\lceil \alpha|D| \rceil}(D; s)$, top- k keeps $\text{Top}_k(D; s)$, and thresholding keeps $\{x : s(x) \geq \tau\}$. MONA computes a benchmark-conditioned relevance score in sparse-autoencoder (SAE) space,

$$s_b(x) = \frac{\sum_j \min(a_j(x), t_{b,j})}{\sum_j \max(a_j(x), t_{b,j})},$$

where $a(x)$ is the sparse activation vector of sample x and t_b is the target vector for benchmark b ; for multiple benchmarks, we keep top samples per benchmark and take their union. IFD uses the cached instruction-following difficulty score

$$s_{\text{IFD}}(x) = \frac{\mathcal{L}(\text{response} \mid \text{instruction})}{\mathcal{L}(\text{response})}.$$

Varentropy is computed from the token distribution p_j at token position j as $H_j = -\sum_v p_j(v) \log p_j(v)$ and $V_j = \sum_v p_j(v)(\log p_j(v) + H_j)^2$, then averaged over valid tokens. N-gram entropy is implemented as unigram Shannon entropy over normalized

word tokens, $H(x) = -\sum_w p_x(w) \log_2 p_x(w)$. AO scores dependency branching by $s_{AO}(x) = 0.5 |\text{verbs}(x)| + 0.5 \text{mean}_{v \in \text{verbs}(x)} |\text{subtree}(v)|$. SemDedup builds L2-normalized SAE sparse vectors, clusters them with MiniBatch K-Means, and greedily drops x_i within its cluster if $\max_{x_j \in K_c} \cos(\hat{a}_i, \hat{a}_j) \geq \tau$ for already kept samples K_c .

Random top- k draws k samples uniformly from the current candidate pool, which acts as a stochastic escape operator and follows the motivation of random search in structured configuration spaces (Bergstra and Bengio, 2012). The Mix operator is a set union over sample identifiers, $\text{Mix}(C, S) = C \cup (S \setminus C)$, where C is the current subset and S is a selected subset from search history.

Use in baselines and AutoSelection. The single-operator baselines evaluate whether any individual signal is sufficient on its own. Random top- k evaluates whether improvements exceed simple random subset selection. Random recipe search samples compositions from the same operator space without the proposed search strategy. AutoSelection uses the full operator library to search over data recipes, allowing heterogeneous signals and recipe-level combinators to be combined under a fixed validation budget.

A.3 Implementation Details

All reported evaluations use the same fixed raw pool described in Section 4.1. Before search, samples are canonicalized into a shared representation with stable sample identifiers, normalized instruction-response fields, source metadata, and cached operator-side signals. Candidate recipes therefore differ only in how they transform this fixed pool; they do not add synthetic samples, rewrite responses, or query external generators during data construction.

For MONA-style task-relevance features, we follow the task-vector construction and related feature-extraction settings from MONA (Ma et al., 2025). The sparse-autoencoder (SAE) model used to obtain sparse activation features is trained with the EleutherAI sparsify library¹ on RedPajama-Data-1T.² All other MONA-specific hyperparameters are kept consistent with the original MONA setting.

¹<https://github.com/EleutherAI/sparsify>

²<https://huggingface.co/datasets/togethercomputer/RedPajama-Data-1T>

One full evaluation consists of executing a recipe on the canonicalized pool, fine-tuning the selected base model on the selected subset, and evaluating the fine-tuned model on the validation suite. We run SFT training with LLaMA-Factory (Zheng et al., 2024) and use vLLM (Kwon et al., 2023) as the inference backend for benchmark evaluation. The primary search budget is $B = 15$ full evaluations. Full evaluations are conducted on a single-node server equipped with 16 Ascend 910C accelerators. Warmup occupies the first three evaluations by probing low-, medium-, and high-retention regimes; the later search-side ablations align comparisons on the post-warmup budget, i.e., steps 4–15. In the ablation table, metrics named Best@4–15, Mean@4–15, and GapArea@4–15 are computed over this post-warmup window rather than over all 15 evaluations.

All SFT jobs use the same training configuration summarized in Table 5. We keep these hyperparameters fixed across candidate recipes so that measured differences are attributable to the selected subset rather than optimizer or systems settings.

Hyperparameter	Value
Training stage	SFT
Precision	bf16
Maximum sequence length	2048
Epochs	3
Learning rate	2.0×10^{-5}
Scheduler	Cosine
Warmup ratio	0.1
Train batch size	256
Flash attention	FA2

Table 5: Main SFT training hyperparameters used for full evaluation.

The validation score aggregates GPQA, GSM8K, BBH, and MMLU. GraphWiz and NLGraph are held out from the search objective and are used as OOD graph benchmarks in the main result table. To simplify the evaluation pipeline while preserving broad benchmark coverage, we evaluate MMLU on a fixed randomly sampled 10% subset shared by all methods. For NLGraph, we filter out topology tasks and retain 3,200 samples for the final graph-reasoning evaluation set, again using the same filtered set for all methods. The prompt templates used for benchmark evaluation are reported in Appendix E.

Language models used inside AutoSelection act only as search-side assistants: they summarize previous evaluations, propose grounded recipe edits,

Method	Param	Kept	GPQA	GSM8K	BBH	MMLU	Avg
Random top- k	0.5	45,000	21.20	52.16	27.82	56.38	39.39
MONA top- k	0.05	–	21.20	54.35	24.89	55.77	39.05
MONA top- k	0.1	–	23.88	51.02	26.30	55.00	39.05
MONA top- k	0.3	–	22.32	53.44	23.58	55.77	38.78
AO top- k	0.1	9,000	18.08	48.80	26.30	51.94	36.28
AO top- k	0.3	27,000	20.75	48.29	27.71	52.77	37.38
AO top- k	0.5	45,000	23.66	49.81	33.36	55.33	40.54
IFD top- k	0.1	8,999	12.50	17.05	20.86	53.16	25.89
IFD top- k	0.3	26,995	16.51	17.28	14.89	54.27	25.74
IFD top- k	0.5	44,991	23.21	23.88	24.45	56.22	31.94
N-gram top- k	0.1	9,000	16.74	36.39	19.23	56.05	32.10
N-gram top- k	0.3	27,000	22.76	51.25	16.63	56.72	36.84
N-gram top- k	0.5	45,000	24.10	51.63	19.13	57.38	38.06
SemDedup	0.975	–	17.41	35.10	23.26	52.88	32.16
SemDedup	0.985	–	14.73	40.40	20.43	56.88	33.11
Varentropy top- k	0.1	8,999	16.29	43.36	17.17	50.88	31.93
Varentropy top- k	0.3	26,995	21.20	40.71	22.17	55.05	34.78
Varentropy top- k	0.5	44,991	20.31	38.51	25.76	55.50	35.02

Table 6: Complete Qwen2.5-1.5B operating-point sweep for the ID single-operator baselines. The reasoning average is computed over GPQA, GSM8K, BBH, and MMLU. Bold parameter and Avg entries indicate the rows used in Table 1; “–” marks unavailable retained-example counts.

rank candidate recipes after recipe execution and state-vector extraction but before full evaluation, and reseed after stagnation. In our implementation, these search-side LLM calls use DeepSeek-R1 (Guo et al., 2025) as the backend model. They do not inspect held-out answers, generate new training samples, or modify the raw pool.

Subject to the source-dataset licenses, the release will include the search and recipe-execution code, exact recipes and operator parameters, selected sample IDs, retained-example and token-ratio statistics from the execution logs, realized state vectors, search trajectories and logs, and the training and evaluation configurations. Thus, the release covers the derived artifacts needed to reconstruct each evaluated subset rather than only the controller implementation.

A.4 Baseline Parameter Setting

Tables 6–8 report all 17 tested operating points for each model: three settings for each of MONA, AO, IFD, N-gram, and Varentropy, plus two SemDedup thresholds. The ID columns in Table 1 use the highest-Avg row for each operator and model. This per-model selection avoids favoring AutoSelection by reusing a weak operating point across scales. The OOD graph columns in the main table remain the original shared-operating-point evalua-

tions and are not recomputed from these ID sweeps. Retained-example counts are reported wherever they were recorded; they need not equal the numeric parameter because some operators use thresholds rather than direct size control.

A.5 Additional Comparison with DEITA

To complement the single-operator sweeps, Table 9 evaluates DEITA at 9K and 20K budgets on the same raw pool. This comparison tests a stronger fixed selector that jointly scores complexity, quality, and diversity, while AutoSelection searches executable compositions over the shared operator catalog.

B Search Procedure and Formal Details

B.1 Recipe Vector

Suppose the recipe space contains three operators: task-relevance filtering, deduplication, and size control. A recipe that enables the first and third operators, disables deduplication, sets a relevance threshold of 0.70, and targets a retained-example ratio of 0.30 can be encoded as

$$\psi(r) = [1, 0.70, 0, 0.00, 1, 0.30],$$

where each operator contributes an on/off indicator followed by its normalized θ value. In the implementation, $\psi(r)$ is a compact fixed-dimensional

Method	Param	Kept	GPQA	GSM8K	BBH	MMLU	Avg
MONA top- k	0.05	14,437	19.00	12.10	5.50	32.10	17.18
MONA top- k	0.1	26,634	19.00	11.60	9.50	32.80	18.23
MONA top- k	0.3	67,343	23.90	10.50	8.20	33.50	19.03
AO top- k	0.1	9,000	20.80	5.50	13.30	31.10	17.68
AO top- k	0.3	27,000	20.30	8.60	11.50	34.60	18.75
AO top- k	0.5	45,000	22.10	9.10	10.40	32.80	18.60
IFD top- k	0.1	8,999	19.90	5.60	5.10	30.20	15.20
IFD top- k	0.3	26,995	15.00	8.00	4.50	30.30	14.45
IFD top- k	0.5	44,991	15.20	9.50	4.40	29.10	14.55
N-gram top- k	0.1	9,000	16.70	6.30	5.20	28.80	14.25
N-gram top- k	0.3	27,000	16.70	7.20	3.50	33.70	15.28
N-gram top- k	0.5	45,000	22.30	10.20	2.70	34.60	17.45
SemDedup	0.975	2,591	19.90	5.60	0.90	24.40	12.70
SemDedup	0.985	10,770	13.80	6.20	11.10	27.60	14.68
Varentropy top- k	0.1	8,999	20.50	2.30	10.00	26.40	14.80
Varentropy top- k	0.3	26,995	12.90	4.40	9.20	29.60	14.03
Varentropy top- k	0.5	44,991	23.30	7.70	10.40	33.80	18.80

Table 7: Complete Llama3.2-1B operating-point sweep for the ID single-operator baselines. Bold parameter and Avg entries identify the rows reported in the ID columns of Table 1.

surrogate feature derived from operator presence and normalized operator parameters. The ordered recipe itself is still supplied to the search controller and Ranker, so order-sensitive reasoning is not claimed to be fully represented by $\psi(r)$ alone.

B.2 Search Pseudocode

B.3 Formal Notes and Search-Side Complexity

Proof for the single-operator containment. Every recipe $r_{u,\theta} = ((u, \theta))$ on the right-hand side of the containment statement in Section 3 is an element of $\mathcal{R}$ by assumption. Therefore, maximizing over all recipes in $\mathcal{R}$ cannot be worse than maximizing over this subset of length-one recipes.

For a concrete example, a scoring-based selector u with score s and parameter $\theta = \alpha$ maps a subset $\mathcal{D} \subseteq \mathcal{D}'_0$ to

$$u(\mathcal{D}; \alpha) = \text{Top}_{\lceil \alpha |\mathcal{D}| \rceil}(\mathcal{D}; s).$$

When used alone, this selector corresponds to the length-one recipe $((u, \alpha))$.

Search-side complexity. Let W be the number of warmup probes, $M = \max_t |\mathcal{C}_t|$ the maximum number of sibling candidates generated per search step, and $L_{\max}$ the maximum recipe length. AutoS-election performs exactly B full evaluations. Apart from these expensive evaluations, it executes at most

$$W + M(B - W)$$

candidate recipes on the cached pool for subset materialization and state-vector extraction, treating warmup probes as singleton candidate sets. If $c_{o,\theta}(n)$ denotes the cost of applying operator o with parameter θ to a subset of size n , the search-side

Method	Param	Kept	GPQA	GSM8K	BBH	MMLU	Avg
MONA top- k	0.05	9,315	21.21	67.10	33.91	62.94	46.29
MONA top- k	0.1	15,527	20.31	61.18	37.61	63.33	45.61
MONA top- k	0.3	35,763	23.44	46.40	33.15	62.78	41.44
AO top- k	0.1	9,000	18.75	57.85	34.57	60.83	43.00
AO top- k	0.3	27,000	22.77	60.58	34.89	58.78	44.26
AO top- k	0.5	45,000	23.88	59.82	33.26	58.56	43.88
IFD top- k	0.1	8,999	20.09	25.17	32.17	57.11	33.64
IFD top- k	0.3	26,995	21.65	20.47	35.54	58.00	33.92
IFD top- k	0.5	44,991	27.46	28.66	31.85	61.17	37.29
N-gram top- k	0.1	9,000	18.53	36.47	32.39	62.17	37.39
N-gram top- k	0.3	27,000	20.09	51.33	31.41	63.06	41.47
N-gram top- k	0.5	45,000	26.34	61.71	31.20	63.06	45.58
SemDedup	0.975	80,973	27.23	61.11	35.65	61.67	46.42
SemDedup	0.985	87,319	22.32	52.77	37.72	62.89	43.93
Varentropy top- k	0.1	8,999	18.75	35.94	19.57	55.67	32.48
Varentropy top- k	0.3	26,995	26.56	52.62	33.91	60.17	43.32
Varentropy top- k	0.5	44,991	25.00	68.30	36.20	61.60	47.78

Table 8: Complete Qwen2.5-3B operating-point sweep for the ID single-operator baselines. Bold parameter and Avg entries identify the rows reported in the ID columns of Table 1.

Model	Method	Budget	Threshold	ID Avg	Selection rule
Llama3.2-1B	DEITA	9K	0.90	11.98	Complexity, quality, diversity
	DEITA	20K	0.90	14.78	Complexity, quality, diversity
	Best single operator	—	—	19.03	MONA (0.3)
	AutoSelection	15 evals	—	19.73	Three-run median
Qwen2.5-1.5B	DEITA	9K	0.95	38.90	Complexity, quality, diversity
	DEITA	20K	0.97	39.79	Complexity, quality, diversity
	Best single operator	—	—	40.54	AO (0.5)
	AutoSelection	15 evals	—	42.23	Three-run median
Qwen2.5-3B	DEITA	9K	0.98	39.56	Complexity, quality, diversity
	DEITA	20K	0.99	44.19	Complexity, quality, diversity
	Best single operator	—	—	47.78	Varentropy (0.5)
	AutoSelection	15 evals	—	49.01	Three-run median

Table 9: Same-pool comparison with DEITA (Liu et al., 2024). DEITA is evaluated at two retained-example budgets with the best tested threshold shown. AutoSelection values are three-run medians; the single-operator values are the best rows in the complete sweeps.

execution cost is bounded by

$$O\left(\sum_{t=1}^B \sum_{r \in \mathcal{C}_t} \sum_{\ell=1}^{L(r)} c_{o_\ell, \theta_\ell}(|S_{\ell-1}|)\right),$$

where $S_{\ell-1}$ is the intermediate subset before the ℓ -th operator. Exact GP refitting over t evaluated recipes costs $O(t^3)$ per step, and prediction over $|\mathcal{C}_t|$ candidates costs $O(|\mathcal{C}_t|t^2)$, which is negligible in our setting because $B = 15$. The dominant cost remains

$$\sum_{t=1}^B [C_{\text{SFT}}(|S_t|) + C_{\text{Eval}}(\mathcal{E})],$$

namely fine-tuning and benchmark evaluation. The cold-start cache is therefore important not because

it removes full evaluations, but because it prevents repeated recomputation of task-, data-, and model-side signals across many candidate recipes.

The empirical accounting in Table 10 is consistent with this complexity picture. We summarize representative 15-evaluation search runs and report the consistently available component-level accounting: accumulated recipe execution time, combined SFT-and-evaluation time, and search-side LLM time. The available records do not reliably split SFT training from benchmark inference, so the two are reported jointly. Some logs contain broader wall-clock totals that include additional overhead, but these fields are not available uniformly across runs; Table 10 therefore uses the shared compo-

Algorithm 1 AutoSelection search.

Require: Raw pool $\mathcal{D}_0$, base model f_0 , recipe space $\mathcal{R}$, suite $\mathcal{E}$, budget B , patience P
Ensure: Best subset S^* , recipe r^* , and score y^*

- 1: Canonicalize $\mathcal{D}_0$ into $\mathcal{D}'_0$ and cache reusable signals
- 2: Sample warmup recipes from $\mathcal{R}$ across retention regimes
- 3: Evaluate warmup recipes to initialize $\mathcal{H}_3$
- 4: $r_1^{\text{seed}} \leftarrow \arg \max_{(r,z,y) \in \mathcal{H}_3} y$
- 5: Initialize r^*, S^*, y^* from the best warmup recipe
- 6: $m \leftarrow 1, t \leftarrow |\mathcal{H}_3| + 1, c \leftarrow 0$
- 7: **while** $t \leq B$ **do**
- 8: $g_{t-1} \leftarrow \text{Summarizer}(\mathcal{H}_{t-1})$
- 9: $\mathcal{C}_t \sim Q_{\text{prop}}(\cdot \mid r_m^{\text{seed}}, g_{t-1}, \mathcal{H}_{t-1})$
- 10: **for each** $r' \in \mathcal{C}_t$ **do**
- 11: Execute r' to compute $S(r')$ and $z(r') = \phi(S(r'))$
- 12: Estimate $\hat{\mu}_{t-1}(r'), \hat{\sigma}_{t-1}(r')$ from $\psi(r')$
- 13: $\rho_t(r') \leftarrow f_{\text{rank}}(r', z(r'), \hat{\mu}, \hat{\sigma}, \mathcal{H}_{t-1})$
- 14: **end for**
- 15: $r_t \leftarrow \arg \max_{r' \in \mathcal{C}_t} \rho_t(r')$
- 16: $S_t \leftarrow S(r_t), z_t \leftarrow z(r_t)$
- 17: Fine-tune f_0 on S_t and evaluate on $\mathcal{E}$ to get y_t
- 18: Update $\mathcal{H}_t \leftarrow \mathcal{H}_{t-1} \cup \{(r_t, z_t, y_t)\}$
- 19: **if** $y_t > y^*$ **then**
- 20: $r^* \leftarrow r_t, S^* \leftarrow S_t, y^* \leftarrow y_t, c \leftarrow 0$
- 21: **else**
- 22: $c \leftarrow c + 1$
- 23: **end if**
- 24: **if** $c \geq P$ **then**
- 25: $r_{m+1}^{\text{seed}} \sim Q_{\text{seed}}(\cdot \mid \mathcal{H}_t)$
- 26: $m \leftarrow m + 1, c \leftarrow 0$
- 27: **end if**
- 28: $t \leftarrow t + 1$
- 29: **end while**
- 30: **return** S^*, r^* , and y^*

nent fields only. Under this accounting, search-side LLM calls remain small relative to full evaluation: LLM time is 0.37–0.61 hours, whereas recipe execution plus SFT-and-evaluation accounts for 8.25–12.15 hours over the same 15 evaluated recipes.

Run	Eval	Exec h	SFT/eval h	LLM h	Total h
A	15	1.35	6.90	0.52	8.77
B	15	2.85	7.95	0.52	11.32
C	15	3.75	8.40	0.61	12.76
D	15	4.35	7.50	0.37	12.22

Table 10: Representative component-level compute accounting for 15-evaluation search runs. All times are accumulated hours over the first 15 full evaluations.

C Search Diagnostics and Ablation Details

C.1 Search-Side Ablation Curves

Figure 3 complements Table 2 by showing both raw-score and best-so-far curves after the warmup budget. The Full AutoSelection reference is the strongest complete 1.5B search run, and the component ablations share its warmup-derived context except for the explicit w/o Warmup condition. The figure is used to distinguish variants that reach similar peak values from variants that spend evaluations less efficiently or suffer larger temporary collapses.

C.2 Random Select Curve

Random Select is included in Table 2 as a policy-control ablation because it changes candidate selection rather than removing a single internal module. After the same warmup and candidate-generation stages, this control chooses one candidate uniformly at random for full evaluation instead of using the GP prior and Ranker to select the next candidate. Figure 4 shows all 25 full evaluation records available in the Random Select run, including the three warmup probes and the later restart seeds. Extending the budget gives Random Select more chances to improve its best result: the best-so-far score rises to 41.17 by step 11. However, the raw curve remains highly volatile, with large drops after apparently strong evaluations and another collapse after the later restart. Thus, additional random budget can raise the best-so-far curve, but it does not provide a stable search policy.

C.3 State Vector Definition and Diagnostics

The state vector summarizes the realized subset S after a candidate recipe has been executed on the fixed reference pool $\mathcal{D}'_0$. We group the fields into task, data, and model components,

$$z(S) = [z^{\text{task}}(S); z^{\text{data}}(S); z^{\text{model}}(S)].$$

Table 11 lists the fields used in our search. SNAR denotes sparse-neuron activation rate.

For the distribution-drift row, $\text{SNAR}(\cdot)$ denotes a dataset-level sparse-neuron activation-rate vector over cached SAE/MONA features. Let $A(x)$ be the active SAE feature-index set of sample x and let S_{valid} be samples with cached sparse features. We compute

$$\text{SNAR}_j(S) = \frac{1}{|S_{\text{valid}}|} \sum_{x \in S_{\text{valid}}} \mathbf{1}[j \in A(x)].$$

Thus, SNAR records how often each SAE feature appears in the subset, ignoring activation magnitudes. This follows the use of sparse autoencoder features as sparse activation coordinates in prior interpretability work (Huben et al., 2024) and in MONA task relevance modeling (Ma et al., 2025). Here s_{IFD} is the IFD score, $V(x)$ is sample-level varentropy, and $s_b(x)$ is the MONA similarity to benchmark b . Candidate recipes are executed cheaply before SFT, so these fields let the Ranker inspect realized subset scale, distribution shift, difficulty, uncertainty, and task relevance before selecting one candidate for full evaluation. As an additional diagnostic, Table 12 reports correlations between saved state-vector fields and downstream full evaluation scores over 107 state-score points from six complete search runs. The strongest all-point Spearman correlations are token ratio, retain ratio, GSM8K task relevance, and score mean, while model-side uncertainty and difficulty fields are negatively associated with score in this set. A leave-one-out ridge probe over the state vector gives Pearson 0.798 and Spearman 0.406. These results support the use of state vectors as search context for screening candidates, but they should not be read as evidence that state vectors are standalone predictors that replace full evaluation.




Figure 3: Trajectory diagnostics for the search-side ablations on the 1.5B setting. Curves show post-warmup raw scores and best-so-far scores over full-evaluation steps 4–15; shaded regions indicate the gap-area diagnostic used to summarize exposure to low-scoring edits.




Figure 4: Extended-budget Random Select curve on the 1.5B setting. Shaded steps denote warmup probes, and dashed vertical markers denote restart seeds.

Grp.	Field	Comp.	Meaning
Task	score_mean	$\text{mean}_{x,b} s_b(x)$	Average MONA task relevance.
	score_std	$\text{std}_{x,b} s_b(x)$	Heterogeneity of task-relevance scores.
	score_per_task	$\{\text{mean}_x s_b(x)\}_{b \in \mathcal{E}}$	Benchmark-wise relevance profile.
Data	retain_ratio	$ S / \mathcal{D}'_0 $	Retained-example scale change after filtering.
	token_ratio	$T(S)/T(\mathcal{D}'_0)$	Training-token scale retained by the recipe.
Model	distribution_drift	$\ \text{SNAR}(S) - \text{SNAR}(\mathcal{D}'_0)\ _2 / \sqrt{D_{\text{SAE}}}$	Sparse-activation shift from the reference pool.
	mean_ifd	$\bar{s}_{\text{IFD}}(S) / \bar{s}_{\text{IFD}}(\mathcal{D}'_0)$	Relative instruction-following difficulty.
	mean_varentropy	$\bar{V}(S) / \bar{V}(\mathcal{D}'_0)$	Relative predictive-uncertainty complexity.

Table 11: State-vector fields used to summarize an executed candidate subset.

Feature	All ρ	Within ρ	Centered r
token_ratio	0.420	0.436	0.716
retain_ratio	0.374	0.362	0.676
score_per_task.gsm8k	0.343	0.274	0.675
mean_varentropy	-0.339	-0.263	-0.456
score_mean	0.320	0.275	0.649
mean_ifd	-0.211	-0.249	-0.584

Table 12: State-feature correlations with downstream full evaluation score over 107 saved state-score points.

C.4 GP Fit

The GP surrogate is used to estimate candidate quality from recipe encodings before full evaluation is run. Its role is to provide a cheap historical prior over the search space, not to replace full evaluation. Realized state vectors are used by the Ranker, not by the GP surrogate. Figure 5 shows the rolling fit between GP-predicted scores and observed full evaluation scores over the analyzed search runs. The predictions provide a coarse historical signal over evaluated recipes, but the fit remains mixed and should not be treated as a replacement for full evaluation. This is expected because the search has few observations and each full evaluation is noisy. We therefore use the GP score only as a directional signal for candidate prioritization while still relying on full evaluation for final selection.

C.5 Operator-Composition Case Studies

Before the paired case studies, we first summarize two run-level diagnostics over 181 full evaluation recipe records from nine evaluated search runs. Figure 6 visualizes both the retained-example-scale scatter and the adjacent-operator motif comparison. These diagnostics are descriptive, not causal ablations, but they help contextualize why retained scale and operator composition are both treated as search variables.

Retention scale is not sufficient. Table 13 shows that retained subset scale matters, but does not explain downstream score by itself. For example, the 40K–60K output range contains 51 records with a 12.98-point score range, while the 60K–80K range contains 67 records with a 12.27-point range. Thus, retention scale is a useful state variable, but recipes with comparable realized sizes can still differ substantially.

Examples	N	Min	Mean	Max	Range
0–20K	26	23.30	33.42	41.20	17.90
20–40K	8	34.62	38.74	40.79	6.17
40–60K	51	29.26	39.34	42.23	12.98
60–80K	67	30.01	40.03	42.29	12.27
80–100K	29	30.87	39.48	41.49	10.62

Table 13: Binned relationship between retained-example scale and validation score over 181 full evaluation recipe records.

Operator motifs recur in stronger records. We split the same records into top- and bottom-tertile groups by validation score and compare adjacent

operator pairs. Table 14 lists the largest positive differences. The enrichment of composed motifs such as MONA→SemDedup and N-gram→MONA supports the fixed-pool data recipe search framing: high-scoring records are not explained only by the presence of one isolated operator. These motif counts should be read as retrospective search diagnostics rather than proof that any pair is universally optimal.

Pair	Top	Bottom	Delta
MONA→SemDedup	11	5	+6
N-gram→MONA	10	5	+5
MONA→N-gram	7	3	+4
SemDedup→Mix	4	2	+2
SemDedup→N-gram	4	2	+2

Table 14: Adjacent-operator motif differences between top- and bottom-tertile full evaluation records.

The search runs provide qualitative evidence that operator effects are context-dependent even when retained scale is roughly controlled. Table 15 therefore reports paired recipes with comparable realized subset sizes, so the comparisons focus on recipe structure after coarse retained-scale matching. These comparisons should be read as search-run case studies, not isolated one-factor ablations: they support the fixed-pool data recipe search formulation by showing that order, thresholds, deduplication choices, and seed composition can change subset quality beyond retained scale alone. Several patterns are visible from the scale-matched pairs: deduplication thresholds are not monotonic, task-relevance and diversity filters interact with their order, and mixing is most useful when followed by an additional diversity-oriented operator rather than used alone. For compactness, the recipe strings use operator names consistently: N-gram, random- k , and Mix denote N-gram top- k , Random top- k , and Mix, respectively.




Figure 5: Rolling GP fit over the three analyzed runs. Each panel compares saved surrogate predictions with downstream full evaluation scores where both are available, then reindexes the comparable points onto the same x-axis.




Figure 6: Run-level search diagnostics. Left: retained-example scale versus validation score. Right: adjacent-operator motif differences between top- and bottom-tertile recipe records.

Case	Result	Recipe	n	Avg
1	High	MONA(0.70)→SemDedup(0.73)	71.1K	41.29
	Low	MONA(0.65)→SemDedup(0.70)	67.6K	39.01
2	High	SemDedup(0.75)→random- k (54K)	54.0K	40.78
	Low	SemDedup(0.85)→random- k (54K)	54.0K	38.60
3	High	MONA(0.87)→Varentropy(0.87)→SemDedup(0.79)→random- k (50K)	50.0K	41.17
	Low	MONA(0.85)→Varentropy(0.85)→SemDedup(0.81)→random- k (50K)	50.0K	39.13
4	High	N-gram(0.90)→MONA(0.85)→SemDedup(0.88)→random- k (45K)	45.0K	42.23
	Low	N-gram(0.85)→MONA(0.80)→SemDedup(0.85)→random- k (42K)	42.0K	38.78
5	High	random- k (80K)→MONA(0.92)→N-gram(0.92)	69.6K	42.28
	Low	random- k (80K)→N-gram(0.90)→MONA(0.90)	67.4K	40.80
6	High	random- k (68K)→SemDedup(0.80)→N-gram(0.90)	60.9K	41.69
	Low	random- k (59.6K)→SemDedup(0.82)	59.3K	40.51
7	High	random- k (48K)→SemDedup(1600, 0.96)→MONA(0.58)→Mix	63.1K	42.17
	Low	random- k (48K)→SemDedup(1600, 0.96)→MONA(0.70)→Mix	63.3K	38.85

Table 15: Scale-matched recipe case studies from the search runs. Scores are validation averages on the 1.5B setting, and n is the realized number of retained examples.

C.6 Structural Recipe Codes

Table 3 uses compact recipe codes to keep the matched-seed structural ablation in one column: A denotes N-gram top- k , B denotes MONA top- k , C denotes SemDedup, D denotes random- k , and E denotes Mix. Thus, $A \rightarrow B \rightarrow C \rightarrow D$ is the reference recipe, while the remaining rows change the order, remove SemDedup, or replace SemDedup with Mix before the same random- k size-control step.

C.7 Search Stability

Table 17 gives the full three-run stability summary used in Section 4.3, including Best@15 scores, benchmark breakdowns, early means, and the trajectory evidence needed for reproducibility. Table 16 gives the corresponding Best@15 summary for all three model scales.

C.8 Recipe Ranking Case Study

Table 18 reports Hit@1, Hit@2, MRR, and the average actual rank of the Ranker’s top-1 choice. The Ranker places the true best candidate at rank 1 in 2/5 cases and within the top 2 in all 5 cases. This is a counterfactual diagnostic rather than a large-scale ranking benchmark, but it supports the claim that the Ranker helps allocate the full evaluation budget toward promising candidates. Table 19 gives the corresponding sibling-candidate ranking study used as diagnostic support for the search-side ablation analysis in Section 4.4. Each row in Table 19 is one sibling candidate from a randomly sampled 1.5B decision point.

Metric	Value
Audited decision points	5
Hit@1	2/5
Hit@2	5/5
MRR	0.70
Mean actual rank of Ranker top-1	1.8

Table 18: Compact Ranker audit metrics over five audited decision points.

Model	Run 1	Run 2	Run 3	Mean $\pm$ SD	Median	Range
Llama3.2-1B	19.73	19.75	19.67	19.72 $\pm$ 0.04	19.73	0.08
Qwen2.5-1.5B	42.23	42.28	41.69	42.07 $\pm$ 0.33	42.23	0.59
Qwen2.5-3B	49.08	49.01	49.00	49.03 $\pm$ 0.04	49.01	0.08

Table 16: Search stability for all three models under three independent 15-full-evaluation runs. SD is the sample standard deviation.




Figure 7: Raw-score and best-so-far curves for three 1.5B AutoSelection runs under the same 15 full evaluation budget.

Run	Pts	Best@15	GPQA	GSM8K	BBH	MMLU	Early Mean @3
Run 1	15	42.23	29.02	54.59	30.00	55.33	36.21
Run 2	15	42.28	24.55	58.00	29.02	57.56	37.16
Run 3	15	41.69	26.12	53.68	28.91	58.06	37.42

Table 17: Three-run search stability under the 15 full evaluation budget. Benchmark columns report the GPQA, GSM8K, BBH, and MMLU scores of the recipe that attains Best@15 in each run.

Dec.	Var.	LLM	Actual	Avg
D1	V3	1	2	39.95
D1	V1	2	1	41.96
D1	V2	3	4	36.84
D1	V4	4	3	37.80
D1	V5	5	5	32.19
D2	V1	1	1	40.62
D2	V4	2	3	39.47
D2	V3	3	2	39.90
D2	V5	4	5	38.65
D2	V2	5	4	39.17
D3	V3	1	2	41.07
D3	V2	2	1	41.43
D3	V4	3	5	37.04
D3	V1	4	3	40.33
D3	V5	5	4	38.57
D4	V2	1	3	38.72
D4	V1	2	1	40.79
D4	V5	3	4	35.42
D4	V3	4	5	30.67
D4	V4	5	2	38.78
D5	V1	1	1	42.20
D5	V5	2	2	41.69
D5	V3	3	3	39.33
D5	V2	4	5	33.16
D5	V4	5	4	37.22

Table 19: LLM Ranker audit on five randomly selected 1.5B decision points. Each decision point contains sibling candidate recipes that were retrospectively evaluated; candidates are anonymized as V1–V5 and rows are ordered by the LLM rank. Lower actual rank is better.

D Additional Empirical Analyses and Baseline Boundaries

D.1 Held-Out Generalization and Operator-Catalog Ablation

Table 20 evaluates Qwen2.5-1.5B on ARC-Challenge (Clark et al., 2018), SVAMP (Patel et al., 2021), and IFEval (Zhou et al., 2023b). These tasks are excluded from the search objective and from the benchmark set used to construct MONA task-relevance features. We use the same task prompts for every method and report strict-prompt IFEval accuracy. The held-out macro is an unweighted mean over the three tasks, so it is not used to select a recipe or tune an operating point.

Table 21 removes MONA from the searchable operator catalog. The resulting 41.33 ID average remains above full-data training and the strongest single-operator baseline, while the gap to full AutoSelection indicates that task relevance is useful but not solely responsible for the recipe-search gain.

D.2 Comparison Boundary for LLM-Driven Data-Processing Agents

We do not include LLM-AutoDP in the main quantitative table because it does not match the controlled fixed-pool data recipe search protocol used in this paper. LLM-AutoDP relies on external LLM instructions embedded in data-processing modules

and was designed around strong domain-specific assumptions. We nevertheless ran an early Qwen2.5-1.5B pilot to understand whether it could serve as a neighboring LLM-driven baseline in our fixed-pool setting. The evaluated pilot output is reported in Table 22.

Method	GPQA	GSM8K	BBH	MMLU	Avg
Pilot	16.74	16.07	22.17	48.94	25.98

Table 22: Early LLM-AutoDP pilot result on the Qwen2.5-1.5B setting. Scores are reported on the same in-distribution validation suite used for the main search objective.

We stopped this pilot before running a complete repeated baseline for two practical reasons. First, applying LLM-AutoDP API calls over a large instruction pool introduced prohibitive wall-clock latency and API overhead compared with the fixed-pool operators used in AutoSelection. Second, the available implementation was not domain-neutral: several LLM-facing modules were written for medical-domain data selection. For example, the LLM optimizer and filtering prompts contain instructions that favor medically relevant samples, and other prompt-based filters ask the LLM to judge data quality through that domain-specific lens. When moved to our mixed instruction pool, these built-in priors can select for the wrong notion of relevance rather than for GPQA/GSM8K/BBH/MMLU performance.

This pilot therefore serves mainly as a boundary case for the comparison and clarifies why the fixed-pool boundary in the main text is useful. Once an LLM is placed directly inside the data-selection stage, the observed result can depend on generator quality, prompt design, domain-specific judging criteria, and API behavior in addition to the underlying data recipe. Highly customized prompts can help in their intended domain, but they also introduce brittle priors when transferred to a different pool, as seen in the medical-relevance prompts above. AutoSelection instead keeps LLMs on the search side, where they summarize evaluated histories and rank grounded recipe edits without rewriting, augmenting, or individually judging every training sample. This design isolates the measured gains as much as possible to choices over grounded operators on a fixed raw pool, rather than to prompt engineering or newly introduced samples.

Method	ARC-C	SVAMP	IFEval strict prompt	IFEval strict instruction	Held-out macro
AO (0.3)	71.76	52.33	23.48	34.77	49.19
Full data	72.27	58.33	29.39	41.37	53.33
IFD (0.1)	74.06	56.67	24.95	36.21	51.89
MONA (0.05)	72.95	54.00	23.84	35.85	50.27
N-gram (0.3)	72.70	56.67	22.18	35.01	50.51
SemDedup (0.985)	73.72	55.33	26.43	37.29	51.83
Varentropy (0.5)	74.66	48.67	28.10	40.05	50.47
AutoSelection	73.72	57.67	28.65	39.81	53.35

Table 20: Qwen2.5-1.5B generalization to held-out ARC-Challenge (ARC-C), SVAMP, and IFEval. None is included in the search objective or original MONA target set. To avoid double counting IFEval, the macro is the unweighted mean of ARC-C, SVAMP, and strict-prompt IFEval; strict-instruction accuracy is reported separately. Bold marks the best selection-method score; Full data is retained as a reference.

Method	ID Avg
Full data	38.83
Best single operator (AO)	40.54
AutoSelection without MONA	41.33
Full AutoSelection	42.23

Table 21: Qwen2.5-1.5B catalog ablation. Removing MONA reduces the best result but still outperforms full-data training and the best single operator.

Recipe	Iter.	Retained	Score
Recipe 1	3	9,136	32.75
Recipe 2	9	40,471	37.70
Recipe 3	19	22,326	38.05
Recipe 4	4	63,252	38.36
Recipe 5	7	24,085	38.72
Recipe 6	13	66,748	39.88
Recipe 7	17	75,044	40.48
Recipe 8	16	73,716	41.06
Recipe 9	15	69,634	42.28

Table 24: Aggregate scores for the nine anonymous recipes visualized in the metric-distribution figure.

D.3 Cross-Scale Transfer Checks

The 7B transfer check is used as a supporting analysis, not as the main evidence for AutoSelection. We take several recipes discovered in the 1.5B search, including recipes that are slightly weaker under the 1.5B objective, and evaluate the corresponding selected subsets on a larger 7B model. As shown in Table 23, the observed behavior is trend-level rather than exact: recipes that are strong at 1.5B tend to remain competitive, but the ordering is not perfectly preserved. This supports a cautious interpretation of transferability: small-model recipe search can reveal useful data-construction motifs, but larger-model validation remains necessary before making final claims.

D.4 Metric-Distribution Comparison

Table 24 reports the aggregate scores for the nine anonymous recipes used in the metric-distribution comparison. The purpose of this analysis is not to identify a single best metric, but to examine whether individual metric distributions are sufficient to explain the downstream performance of a selected subset.

Figure 8 compares the same nine evaluated recipes from the 1.5B search process. Each row corresponds to one evaluated recipe, and the six panels report the marginal distributions of MONA-style relevance scores for GPQA, GSM8K, BBH, and MMLU, followed by IFD and entropy.

The main observation is that several evaluated recipes exhibit highly similar marginal distributions but obtain different aggregate scores after full evaluation. In the MONA-based columns, the GPQA, GSM8K, BBH, and MMLU distributions are all concentrated in narrow ranges and mostly preserve similar unimodal shapes across recipes. The IFD distributions are also strongly concentrated near the lower end of the axis for nearly all recipes. Even entropy, which shows more variation in tail shape than the other metrics, does not provide a clear separation among all recipes. These patterns indicate that data can look similar under individual selection signals while still leading to different downstream outcomes.

The distributional evidence supports the central design choice of AutoSelection: data selection should be treated as multi-view selection rather than single-score filtering. Since similar one-dimensional distributions can correspond to dif-

Recipe	Cross-scale ranking					7B validation scores			
	1.5B score	1.5B rank	7B score	7B rank	Rank change	GPQA	GSM8K	BBH	MMLU
Recipe 1	41.96	<u>2</u>	51.92	1	+1	19.64	76.19	<u>41.63</u>	70.22
Recipe 2	42.23	1	<u>51.90</u>	<u>2</u>	-1	20.09	<u>74.91</u>	43.37	<u>69.22</u>
Recipe 3	36.29	4	51.76	3	+1	25.00	71.95	<u>41.63</u>	68.44
Recipe 4	34.79	5	50.33	4	+1	20.98	73.31	<u>38.48</u>	68.56
Recipe 5	39.06	3	49.22	5	-2	18.53	73.09	37.50	67.78
Recipe 6	25.74	7	44.54	6	+1	14.96	65.96	33.48	63.78
Recipe 7	25.90	6	39.47	7	-1	13.84	56.48	29.46	58.11

Table 23: Transfer from 1.5B recipe search to 7B evaluation after excluding the full-data baseline. Ranks are recomputed over the seven transferred recipes.

ferent downstream performance, the search process needs to consider how operators, thresholds, and ordering jointly shape the selected subset. Full evaluation remains necessary because it observes the actual effect of a recipe after fine-tuning, while the metric distributions serve as contextual signals that guide exploration and exploitation under a limited evaluation budget.



Figure 8: Qualitative metric-distribution comparison for nine anonymous evaluated recipes. Each row corresponds to one recipe, and the six panels in the row correspond to subset-level metric views. Aggregate scores for the same recipes are reported in Table 24.

E Prompt Templates Used in Search and Evaluation

This appendix reports the runtime prompt templates used by AutoSelection.

E.1 Search-agent prompts

The search loop uses four LLM-facing prompt templates. The Summarizer converts verified history into concise search guidance, the Proposer generates a candidate pool from the current recipe and guidance, the Ranker chooses among surrogate-ranked candidates, and the Reseeder retunes restart parameters after stagnation.

Listing 1: Summarizer prompt template.

```
You are a data science evaluation analyst. Analyze the
following evaluated-recipe history from an automated
data selection search system.

The system is searching for the best data recipe to train an
LLM. Each row is one evaluated recipe: the recipe
produced a selected subset, then a model was trained
and evaluated.

{experiment_history_table}

KEY CONTEXT:
- Higher scores are better (aggregated accuracy across
  benchmarks)
- "Operators" are data filtering/mixing steps applied
  sequentially
- "Samples" is the number of training samples after filtering
- The pool has {estimated_total_pool_size} total samples

TASK: Produce exactly 3-5 concise, actionable findings. Each
finding should be:
1. A specific observation (not vague)
2. Backed by data from the table
3. Actionable (suggests what to try or avoid)

Format each finding as a numbered line. Be direct and
quantitative.
These are HYPOTHESES based on limited data, not proven facts.

Example format:
1. More data consistently helps: recipe_A (12K samples,
  22.2%) > recipe_C (3K samples, 18.5%). Avoid aggressive
  filtering.
2. operator_X at rate 0.3 hurts benchmark_Y: recipe_B dropped
  from 15% to 0.9%. Try higher rates or skip it.
```

Listing 2: Proposer prompt template.

```
You are an expert Data-Centric AI Search Controller
optimizing a data recipe.

YOUR GOAL: Propose {n_candidates} DISTINCT mutated recipe
configurations that resolve current risks and explore
different valid subspaces. Some should be conservative,
some more aggressive.

{operator_catalog}
{registered_operator_note}

=== CURRENT STATE ===
Current Recipe:
{current_recipe_steps}

Current Metric Score: {score}

{state_vector_section}
{benchmark_diagnostic_section}
{pool_context_section}
{search_history_section}
{experiment_insights_section}
```

```
{union_operator_section}

=== INSTRUCTIONS ===
1. Analyze the current recipe, state vector, and search
  history.
2. Select operators and hyperparameters ONLY from the
  OPERATOR CATALOG.
3. Your output MUST be a valid JSON array of objects
  representing the {n_candidates} recipes. Do NOT include
  markdown blocks (````json ````), just raw JSON.
4. Format:
[
  {
    "steps": [
      {
        "operator": "operator_name",
        "params": {"param1": "value", "param2": 123}
      }
    ]
  },
  ... (up to {n_candidates} distinct configurations)
]
```

Listing 3: Ranker prompt template.

```
You are a strategic advisor for an automated data selection
search system. Your task is to select the SINGLE most
promising candidate recipe for real evaluation.

## SEARCH STATE
- Total data pool: {pool_size} samples
- Iterations completed: {n_iterations}
- Budget remaining: {budget_remaining}h / {budget_total}h ({
  budget_pct}%)
- Current best score: {best_score}% (recipe: {best_name})
- Search phase: {phase}

{parent_section}
{detailed_experiment_history_section}
{experiment_insights_section}

## CANDIDATES (ranked by GP surrogate)
NOTE: The GP surrogate predicts expected utility from an 11-D
recipe encoding (operator presence + parameters).
Each candidate's pipeline has been pre-executed to obtain
data state metrics (shown below for reference).

{candidate_table_with_gp_and_state_vectors}

## SELECTION CRITERIA
Consider these factors carefully:

1. Per-Task MONA Scores (PRIMARY SIGNAL):
- score_per_task shows how relevant the selected subset is
  to each benchmark.
- A candidate whose per-task MONA scores improve across
  multiple benchmarks is a strong positive signal, even
  if retain_ratio drops.
- Compare each candidate's score_per_task against the
  parent's and look for improvements on weak benchmarks.
- If a candidate improves some benchmarks but hurts others,
  weigh the magnitude and importance of each.
- score_mean is the aggregate; score_per_task is the
  breakdown. Always prioritize the per-task view.

2. Exploration vs Exploitation Trade-off:
- Early search: prefer high sigma candidates to gather
  information.
- Late search: prefer high mu candidates to refine the
  best.
- Current phase: {phase}.

3. Data Quantity Risk:
- Recipes that aggressively filter data risk producing too
  few samples.
- Historical evidence shows extreme filtering often fails
  catastrophically.
- Union operators can recover data volume and are safer
  exploration choices.
- Refer to the per-benchmark history to see how sample
  count correlates with each benchmark.

4. Operator Synergies and Redundancy:
- Multiple filtering operators in sequence compound data
  loss multiplicatively.
- Operators from the same family are often redundant.
```

```

- Complementary operators tend to work well together.
5. Feedback Alignment:
- Does this candidate address the patterns identified in
  evaluation insights?
- Does it avoid strategies that have been shown to fail?
6. State Vector Patterns:
- High retain_ratio with good score_mean tends to perform
  well.
- High distribution_drift indicates risky distributional
  shift.
- The parent's state vector shows the data profile that
  candidates will modify.
7. GP Model Limitations:
- The GP has only {n_iterations} training points, so
  predictions carry uncertainty.
- Do not blindly trust GP rankings, especially when scores
  are close.
- Qualitative reasoning about operator interactions can
  add value beyond the GP.

## OUTPUT FORMAT
After thorough reasoning, output a full ranking of all
presented candidates as a JSON object.
The ranking list must contain ALL candidate indices (0-based)
sorted from most promising to least:

{
  "ranking": [<best_idx>, <2nd_idx>, ..., <worst_idx>],
  "confidence": "<high|medium|low>",
  "rationale": "<one-sentence explanation of why your top
    choice was chosen>",
  "eval_rationale": "<one-sentence reason for eval decision>"
}

Think carefully before answering. Consider each candidate's
strengths and risks.

```

Listing 4: Reseeder prompt template.

```

You are choosing a restart operator motif for recipe search.

Use the search evidence below to select a small restart motif
that is promising.
You must choose between 1 and 3 NON-TRUNCATE operators.

[OPERATOR_CATALOG]

=== SEARCH HISTORY ===
[SEARCH_HISTORY]

=== POSITIVE OPERATOR SIGNALS ===
[POSITIVE_OPERATOR_SIGNALS_JSON]

=== POSITIVE PAIR SIGNALS ===
[POSITIVE_PAIR_SIGNALS_JSON]

=== HISTORICAL SUCCESSFUL EXAMPLES ===
[HISTORICAL_SUCCESSFUL_EXAMPLES_JSON]

=== RULES ===
1. Select only operators from the allowed catalog above.
2. Prefer operators and combinations supported by the
  evidence.
3. Keep parameters within catalog bounds.
4. Return raw JSON only in this exact format:
[
  {"operator": "mona_filter", "params": {"fraction": 0.5}},
  {"operator": "ngram_entropy", "params": {"fraction": 0.4}}
]

Pool size reference: [POOL_SIZE]

```

E.2 Benchmark evaluation prompts

The validation suite uses GPQA, GSM8K, BBH, and MMLU.

Listing 5: GPQA evaluation prompt scaffold.

```

System:

```

You are an expert assistant. Answer the following multiple choice question by selecting the correct option (A, B, C, or D).

```

Instructions:
1. Read the question carefully.
2. Consider each option.
3. Respond with ONLY the letter of the correct answer (A, B,
  C, or D) on the last line.
4. Format: put your final answer after "Answer:" on the last
  line.

```

```

Example format:
[Your reasoning]
Answer: B

```

```

Few-shot turns:
User: Question: Which of the following is NOT a function of
  the cell membrane?
A. Selective permeability
B. Protein synthesis
C. Cell signaling
D. Cell adhesion
Assistant: The cell membrane has multiple functions including
  selective permeability, cell signaling, and cell
  adhesion. Protein synthesis occurs at ribosomes, not at
  the cell membrane.
Answer: B

```

```

User: Question: What is the SI unit of electrical resistance?
A. Volt
B. Ampere
C. Ohm
D. Watt
Assistant: Electrical resistance is measured in Ohms (Omega),
  named after Georg Ohm. Volts measure potential
  difference, Amperes measure current, and Watts measure
  power.
Answer: C

```

```

User: Question: Which element has the highest
  electronegativity?
A. Oxygen
B. Nitrogen
C. Fluorine
D. Chlorine
Assistant: Fluorine has the highest electronegativity of all
  elements (3.98 on the Pauling scale). It is the most
  electronegative element because it has a small atomic
  radius and high effective nuclear charge.
Answer: C

Final user turn:
Question: {question}
A. {option_A}
B. {option_B}
C. {option_C}
D. {option_D}

```

Listing 6: GSM8K evaluation prompt scaffold.

```

System:
You are a helpful math assistant. Solve the following math
problem step by step. Show your work, then provide the
final answer after '#### ' on the last line.

Few-shot turns:
User: Janet's ducks lay 16 eggs per day. She eats three for
  breakfast every morning and bakes muffins for her
  friends every day with four. She sells every duck egg
  at the farmers' market daily for $2. How much in
  dollars does she make every day at the farmers' market?
Assistant: Janet's ducks lay 16 eggs per day. She eats 3 for
  breakfast and uses 4 for muffins, so she uses 3 + 4 = 7
  eggs. That leaves 16 - 7 = 9 eggs to sell. At $2 each,
  she makes 9 x $2 = $18.
#### 18

User: A robe takes 2 bolts of blue fiber and half that much
  white fiber. How many bolts in total does it take?
Assistant: The robe takes 2 bolts of blue fiber. It takes
  half that much white fiber, so 2 / 2 = 1 bolt of white
  fiber. In total, it takes 2 + 1 = 3 bolts.
#### 3

User: Josh decides to try flipping a house. He buys a house

```

for \$80,000 and puts \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?

Assistant: Josh spent $\$80,000 + \$50,000 = \$130,000$ total. The repairs increased the value by 150%, so the increase is $\$80,000 \times 1.5 = \$120,000$. The new value is $\$80,000 + \$120,000 = \$200,000$. His profit is $\$200,000 - \$130,000 = \$70,000$.

70000

Final user turn:
{question}

Listing 7: BBH evaluation prompt scaffold.

System:
You are an expert reasoning assistant. Answer the following question step by step, then provide your final answer.

Instructions:
1. Think through the problem carefully.
2. Show your reasoning.
3. On the last line, write your final answer after "Answer:" exactly matching the expected format.

Example format:
[Your reasoning]
Answer: (B)

Few-shot turns:
User: not (True) and (True) is
Assistant: not (True) evaluates to False. False and (True) evaluates to False.
Answer: False

User: In the following sentences, explain the antecedent of the pronoun.
Sentence: The nurse notified the patient that his shift would be ending in an hour.
Options:
(A) The nurse's shift
(B) The patient's shift
(C) Ambiguous
Assistant: The pronoun 'his' could refer to either the nurse or the patient. However, given the context of notifying about a shift ending, 'his' most likely refers to the nurse's shift.
Answer: (A)

User: If we list all the natural numbers below 10 that are multiples of 3 or 5, we get 3, 5, 6 and 9. The sum of these multiples is 23.
Find the sum of all the multiples of 3 or 5 below 20.
Assistant: Multiples of 3 below 20: 3, 6, 9, 12, 15, 18
Multiples of 5 below 20: 5, 10, 15
Combined (no duplicates): 3, 5, 6, 9, 10, 12, 15, 18
Sum = $3+5+6+9+10+12+15+18 = 78$
Answer: 78

Final user turn:
{input}

Listing 8: MMLU evaluation prompt scaffold.

System:
You are a knowledgeable assistant. Answer the following multiple choice question by selecting the correct option (A, B, C, or D).

Instructions:
1. Read the question carefully.
2. Consider each option.
3. Respond with your reasoning, then provide the letter of the correct answer after "Answer:" on the last line.

Example format:
[Your reasoning]
Answer: B

Few-shot turns:
User: Question: What is the capital of France?
A. London
B. Berlin
C. Paris
D. Madrid

Assistant: Paris is the capital and largest city of France.
Answer: C

User: Question: Which planet is known as the Red Planet?
A. Venus
B. Mars
C. Jupiter
D. Saturn
Assistant: Mars is commonly known as the Red Planet due to its reddish appearance caused by iron oxide on its surface.
Answer: B

User: Question: What is the powerhouse of the cell?
A. Nucleus
B. Ribosome
C. Mitochondria
D. Golgi apparatus
Assistant: Mitochondria are often called the 'powerhouse of the cell' because they generate most of the cell's supply of ATP, the main energy currency.
Answer: C

Final user turn:
Question: {question}
A. {choice_A}
B. {choice_B}
C. {choice_C}
D. {choice_D}

Listing 9: GraphWiz and NLGraph evaluation prompt scaffold.

You are a graph reasoning assistant.

Solve the graph problem carefully. On the last line, output only:
Answer: Yes
or
Answer: No

Final user turn:
{question}