

# GREEDY INFORMATION PROJECTION FOR LLM DATA SELECTION

Victor Ye Dong\*, Kuan-Yun Lee\*, Jamei Shuai, Shengfei Liu, Yi Liu, Jian Jiao  
Microsoft  
victordong@microsoft.com

## ABSTRACT

We present *Greedy Information Projection* (GIP), a principled framework for choosing training examples for large language model fine-tuning. GIP casts selection as maximizing mutual information between a subset of examples and task-specific query signals, which may originate from LLM quality judgments, metadata, or other sources. The framework involves optimizing a closed-form mutual information objective defined using both data and query embeddings, naturally balancing *quality* and *diversity*. Optimizing this score is equivalent to maximizing the projection of the query embedding matrix onto the span of the selected data, which provides a geometric explanation for the co-emergence of quality and diversity. Building on this view, we employ a fast greedy matching-pursuit procedure with efficient projection-based updates. On instruction-following and mathematical reasoning datasets, GIP selects small subsets that match full-data fine-tuning while using only a fraction of examples and compute, unifying quality-aware and diversity-aware selection for efficient fine-tuning.

## 1 INTRODUCTION

A central question in the discussion of large language model (LLM) training is how one should select data. With recent developments of large-scale pretrained language models such as GPT-4 (Achiam et al., 2023), Qwen (Yang et al., 2025) and Llama (Grattafiori et al., 2024), the community has witnessed models achieving remarkable performance across a broad spectrum of natural-language tasks. Typically, these models train on large amounts of data that scale with model size, commonly known as the scaling law. Empirical scaling laws of LLMs reveal that model capability grows predictably with (i) parameter count, (ii) compute, and (iii) training data volume (Hoffmann et al., 2022; Kaplan et al., 2020). Yet practitioners increasingly observe a subtler regime: once data quantity is sufficient, *data quality* becomes the primary factor limiting further gains. Consequently, a fundamental challenge in both pre-training and instruction tuning is to select a subset of samples that maximizes downstream performance while respecting resource constraints.

Earlier work addressed the problem through diverse dataset construction (Wang et al., 2022; Taori et al., 2023b) and ad-hoc filtering heuristics such as perplexity thresholds, deduplication, or clustering (Bukharin et al., 2024; Zhao et al., 2024; Chen et al., 2023b; Ge et al., 2024). While effective, these heuristics offer little theoretical guidance on *why* a particular example is valuable, and they do not unify quality and diversity under a single objective.

**This work.** We present a principled *information-theoretic* framework for data selection built on a mutual information (MI) metric between Gaussians scaled by data and query embeddings. As we will demonstrate later, this metric promotes both diversity and quality in a single objective, and also gives rise to efficient algorithm design. Our main contributions are:

1. **Principled theoretical formulation.** We propose a principled theoretical framework that casts the data selection problem as maximization of mutual information. Optimizing this objective naturally promotes diversity and quality of selected data. This framework is also

---

\*Equal contribution

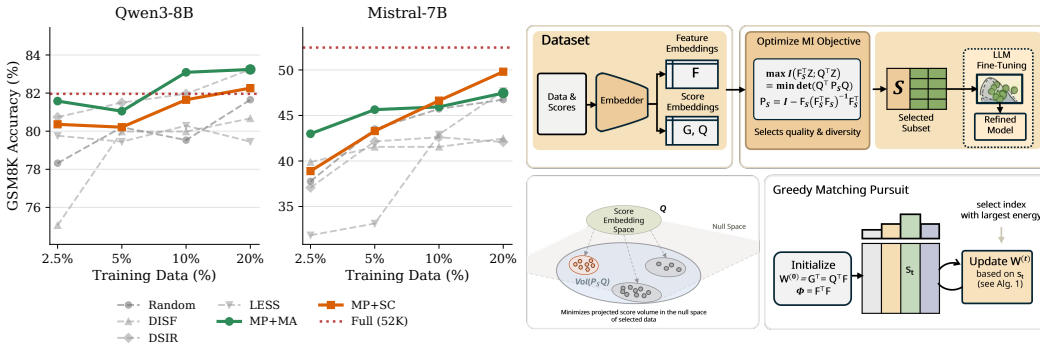


Figure 1: **Left:** GSM8K performance comparison on Qwen3-8B (left panel) and Mistral-7B (right panel) across different training data percentages (2.5%, 5%, 10%, 20%). Our proposed methods MP+MA and MP+SC are competitive with strong baselines (Random, DSIR, DISF, LESS), and often improve over Random/DISF/LESS. In several settings, they approach the full-dataset (100%, shown as a horizontal reference line) performance at 10%–20% of training data, demonstrating strong data efficiency. **Right:** Geometric interpretation of GIP. The method maximizes mutual information between Gaussian projections induced by the data embedding matrix  $F$  and score embedding matrix  $Q$ . This is equivalent to minimizing the volume (determinant) of score embeddings projected onto the null space of selected data, naturally balancing quality (high-score items) and diversity (new directions in embedding space).

flexible, enabling balanced selection with respect to disparate information sources, such as factuality and style scores from LLM evaluations.

- Efficient approximation algorithms.** We develop a greedy matching pursuit (MP) approximation algorithm that solves an approximate dual problem. The Greedy MP approach scales linearly with the total size of available data in practice, enabling data selection under realistic budget constraints.
- Strong empirical results.** On instruction-tuning benchmarks, our *Greedy Information Projection* (GIP) achieves substantial data efficiency gains over state-of-the-art baselines. Across MT-Bench, BBH, and GSM8K, GIP often approaches full-dataset performance using only 1–20% of training data, and in multiple settings matches or exceeds full-data results—corresponding to a substantial data reduction (see Figure 1).

## 2 RELATED WORK

### 2.1 DATA CURATION FOR LARGE-SCALE LANGUAGE MODELS

Diversity and quality are recognized as crucial factors impacting the overall quality of model training (Sener & Savarese, 2018; Chen et al., 2024; 2023a). With the recent development of large language models, there is a growing trend in studies focusing on data composition quality and diversity, and how these factors influence language model pretraining and instruction tuning.

For example, Lee et al. (2022) highlight the importance of diversity in improving training efficiency and propose a method to remove near-duplicate records from training data. Similarly, Bukharin et al. (2024) emphasize the significance of both quality and diversity in instruction tuning data, proposing a general method to balance these two aspects in training data selection. Furthermore, Du et al. (2023) introduce a systematic data selection approach that combines data quality, diversity, and augmentation for instruction tuning. However, neither (Bukharin et al., 2024) nor (Du et al., 2023) have a unified framework to unify *both* quality and diversity; diversity was either applied as a penalty or considered in a different stage of data selection.

Prior work has explored various data selection strategies. Chen et al. (2023b) propose instruction-following difficulty (IFD) scoring to identify high-quality instruction data. Ge et al. (2024) introduce

clustering-based active retrieval (CaR) that uses representativeness and uncertainty for data selection. Xia et al. (2024) develop LESS, a gradient-based method for selecting informative training data. Xie et al. (2023) propose data selection using importance resampling (DSIR) for domain adaptation. Our work differs by providing a unified information-theoretic framework that naturally balances quality and diversity in a single objective, while being computationally efficient and theoretically grounded. Optimal experimental designs for selecting diverse training examples are another strategy; see, for example (Deb et al., 2025), for a discussion of how selecting diverse training examples relates to log-likelihood maximization.

## 2.2 INFORMATION-THEORETIC OBJECTIVES IN SELECTION AND CLUSTERING

There is a long history of applications of mutual information to data selection and clustering. Mutual information can capture complex relationships between features that might not be apparent (Knops et al., 2006). Mutual information is also flexible, working on various types of data such as categorical data (He et al., 2008) and numerical data (Kraskov et al., 2005). Approximations of mutual information have also been applied to k-means (Sugiyama et al., 2014; Calandriello et al., 2014). More recently, mutual information based algorithms have been developed for clustering (Do et al., 2021), community detection (Newman et al., 2020), 3D object representation learning (Sanghi, 2020), and unsupervised sentence embeddings (Zhang et al., 2020). Theoretical developments of mutual information inequalities have also been applied in pairwise comparisons (Lee & Courtade, 2021) and generalized linear models (Lee & Courtade, 2020), where optimal selections and minimax algorithms are understood to be closely related to singular values of query matrices; see, e.g., (Lee, 2022) for a broad discussion.

It is important to note that mutual information is often difficult to compute and various forms of estimators and approximations are used to reduce the computation cost; see (Kraskov et al., 2004) for a classic mutual information estimator. In this paper, we consider the mutual information of joint-Gaussian variables — Gaussianity provides us many favorable elementary expressions which allow us to bypass much of the complexity of estimating mutual information.

## 2.3 ACTIVE LEARNING AND CORESET SELECTION

Our work is related to active learning (Settles, 2009) and coresset construction (Bachem et al., 2017), which aim to select informative examples and compact subsets for downstream tasks, respectively (e.g., DSIR (Xie et al., 2023) and DISF (Fan et al., 2025)). We build on this broader theme of data efficiency, introducing an information-theoretic selection criterion that jointly considers quality and diversity for instruction-tuning.

## 3 PROBLEM FORMULATION

We are given a set of  $m$  data points  $\mathcal{F} = \{f_1, \dots, f_m\}$  with  $f_i \in \mathbb{R}^d$  representing features of selectable data, and  $n$  score vectors  $\mathcal{G} = \{g_i \in \mathbb{R}^m, i \in [n]\}$  where each score vector  $g_i$  represents evaluations on all  $m$  data points for a specific quality dimension. These scores can come from LLM evaluators assessing quality dimensions (helpfulness, accuracy, reasoning) or internal measures like self-consistency estimates.

Our goal is to select, given data features  $\mathcal{F}$  and quality scores  $\mathcal{G}$ , at most  $k$  data features that are both *diverse* (spanning different regions of the feature space) and *high-quality* (aligned with scores).

### 3.1 MUTUAL INFORMATION FORMULATION

Let  $F \in \mathbb{R}^{d \times m}$  be the data embedding matrix with  $F = [f_1, f_2, \dots, f_m]$ , and let  $G \in \mathbb{R}^{m \times n}$  be the score matrix with  $G = [g_1, g_2, \dots, g_n]$ . For selection  $S \subseteq [m]$ , define  $F_S \in \mathbb{R}^{d \times |S|}$  as the matrix of selected data columns. We assume data embeddings are normalized:  $\|f_i\|_2 = 1$  for all  $i \in [m]$ .

Given scores  $G$ , we construct a query embedding  $Q \in \mathbb{R}^{d \times n}$  so that  $F^\top Q \approx G$ . This establishes a connection between the feature space (where diversity is measured) and the score space (where quality is defined). Since  $m \gg d$  in practice,  $F$  typically has rank at most  $d$  and is not full column rank. To ensure existence and numerical stability of  $Q$ , one can use Tikhonov regularization and

solve

$$Q_\epsilon = \arg \min_Q \|F^\top Q - G\|_F^2 + \epsilon \|Q\|_F^2 \quad (3.1)$$

with solution  $Q_\epsilon = F(F^\top F + \epsilon I_m)^{-1}G$  for small  $\epsilon > 0$ . This ensures  $\|F^\top Q_\epsilon - G\|_F \leq C\epsilon$  for some constant  $C$ . Details are provided in Appendix D.

**Remark** (Necessity of  $Q$ ). We will later show that our approximation algorithms do not rely on the exact choice of  $Q$  but rather work directly from  $F$  and  $G$ . We introduce  $Q$  primarily for theoretical and geometric interpretation.

We are now ready to formulate the mutual information objective. In particular, we consider the standard Gaussian  $Z \in \mathcal{N}(0, I_d)$  and transformations  $Z_Q := Q^\top Z$ ,  $Z_{F_S} = F_S^\top Z$ :

$$\begin{bmatrix} Z_Q \\ Z_{F_S} \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} Q^\top Q & Q^\top F_S \\ F_S^\top Q & F_S^\top F_S \end{bmatrix}\right) \quad (3.2)$$

The mutual information between  $Z_Q$  and  $Z_{F_S}$  is:

$$I(Z_Q; Z_{F_S}) = \frac{1}{2} \log \left( \frac{\det(Q^\top Q) \det(F_S^\top F_S)}{\det(\Sigma)} \right) \quad (3.3)$$

where  $\Sigma$  is the joint covariance matrix in equation 3.2. Our goal is to select  $S$  to maximize this mutual information.

**Remark** (Gaussianity as a modeling device). This construction does not assume that data or query embeddings follow a Gaussian distribution. Rather, the auxiliary variable  $Z \sim \mathcal{N}(0, I_d)$  serves as a random probe that induces Gaussian random projections  $Q^\top Z$  and  $F_S^\top Z$  whose second-order structure is fully determined by the geometry of the embeddings. Different choices of functionals on these projections lead to different geometric criteria. For example, objectives based solely on variances or covariances recover quadratic (e.g.,  $\ell_2$ -type) measures of alignment between subspaces. In our case, we adopt mutual information as the objective of interest, which admits a closed-form log-determinant expression, is invariant under orthogonal transformations, and captures the total linear dependence between the subspaces spanned by  $Q$  and  $F_S$ .

Elementary linear algebraic operations give the following interpretation.

**Theorem 1.** *Maximizing mutual information defined in equation 3.3 is equivalent to optimizing*

$$\arg \max_S I(Z_Q; Z_{F_S}) = \arg \min_S \det(Q^\top (I - F_S(F_S^\top F_S)^{-1}F_S^\top)Q). \quad (3.4)$$

An intuitive interpretation of equation 3.4 is that the matrix  $P_S := I - F_S(F_S^\top F_S)^{-1}F_S^\top$  projects  $Q$  onto the null space of  $F_S$ . Minimizing  $\det(Q^\top P_S Q)$  selects  $S$  that minimizes the volume of  $Q$  after projection, naturally encouraging both diversity (expanding into new directions in embedding space) and quality (aligning with high-score directions).

We establish a connection between optimizing mutual information and the quadratic expression  $\|F_{S^*}^\top Q\|_2$ .

**Theorem 2.** *Given any selection  $S^*$  that maximizes equation 3.4,*

$$\|F_{S^*}^\top Q\|_2 \geq \sigma_{\min}(F_S) \sqrt{1 - \left(\frac{\eta}{\det(Q^\top Q)}\right)^{1/r}} \quad (3.5)$$

where  $\eta = \det(Q^\top (I - P_{S^*})Q)$  and  $r$  is the rank of  $Q$ .

This provides a lower bound on how well the selected data  $F_{S^*}$  aligns with the query directions  $Q$ .

**Remark.** When mutual information is ill-defined (for instance, when there are no scores), we maximize the entropy  $h(Z_{F_S}) = \frac{1}{2} \log \det(F_S^\top F_S)$  instead. This becomes a diversity optimization, consistent with our framework. Notably, maximizing  $h(Z_{F_S})$  coincides with that of a determinantal point process (DPP) (Kulesza et al., 2012).

## 4 GREEDY APPROXIMATION ALGORITHM

With small data sizes, one can permute through all possible selections to obtain the optimal selection  $S$ , but the exponential computation cost is impractical in real-world scenarios where data sizes are large. Instead, we employ a classic matching pursuit approach to approximate equation 3.4, following a direct application of the classic AM-GM inequality.

**Theorem 3.** *The determinant stated in equation 3.4 satisfies*

$$\det(Q^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) Q) \leq \left( \frac{\text{Tr}((Q^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) Q))}{n} \right)^n. \quad (4.1)$$

By shifting our optimization target to equation 4.1, the relaxed optimization objective becomes a trace and enjoys linearity. In particular, recall that  $Q$  is a  $\mathbb{R}^{d \times n}$  matrix constructed by a concatenation of  $q_1, \dots, q_n \in \mathbb{R}^d$ . By linearity, we have

$$\text{Tr}((Q^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) Q)) = \sum_{i=1}^n \text{Tr}(q_i^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) q_i), \quad (4.2)$$

and we obtain a quadratic expression with favorable properties to work with. Note also that the minimization of equation 4.2 over  $S$  depends on  $Q$  only through scores  $g_i := F^\top q_i \in \mathbb{R}^m$ ,  $i \in [n]$ . Therefore, the expression in equation 4.2 permits us to work directly with scores  $G$  *without the knowledge of  $Q$* .

In this section, we detail our **Greedy MP** algorithm that approximates the *dual problem* by a classic matching pursuit approach.

### 4.1 GREEDY MATCHING PURSUIT (MP)

The intuition is to note that equation 4.2 involves computing lengths of  $q_i$  projected onto the null space of  $F_S$ , which can be written as a dual form:

$$\min_S \sum_{i=1}^n \min_{\lambda \in \mathbb{R}^k} \left\| q_i - \sum_{j \in S} \lambda_j f_j \right\|_2^2. \quad (4.3)$$

Greedy MP solves this by minimizing residual gain across all query embeddings: at step  $t + 1$ ,

$$s_{t+1}, \lambda_{t+1}^* = \arg \min_{s \in [m] \setminus S_t} \sum_{i=1}^n \min_{\lambda \in \mathbb{R}^k} \|r_i - \lambda_i f_s\|_2^2, \quad (4.4)$$

where  $r_i = q_i - \sum_{j=1}^t \lambda_j^* f_{s_j}$  is the residual of the  $i$ -th score at step  $t + 1$ . Note that for each individual  $q_i$  and candidate  $f_s$ , the minimizer  $\lambda_i^* = r_i^\top f_s$  is unique, and satisfies

$$\min_{\lambda \in \mathbb{R}} \|r_i - \lambda f_s\|_2^2 = \|r_i - r_i^\top f_s f_s\|_2^2 = \|r_i\|_2^2 - (r_i^\top f_s)^2. \quad (4.5)$$

Substituting equation 4.5 into equation 4.4, the greedy selection of  $s_{t+1}$  is one that solves

$$s_{t+1} = \arg \max_{s \in [m] \setminus S_t} \sum_{i=1}^n (r_i^\top f_s)^2.$$

By defining the *residual score matrix*  $W \in \mathbb{R}^{n \times m}$  with entries  $W_{i,j} = r_i^\top f_j$ , we get  $s_{t+1}$  by solving

$$s_{t+1} = \arg \max_{s \in [m] \setminus S_t} \sum_{i=1}^n (W_{i,s}^{(t)})^2.$$

After selection  $s_{t+1}$ ,  $W$  can be updated with

$$W_{i,j}^{(t+1)} \leftarrow W_{i,j}^{(t)} - \Phi_{j,s_{t+1}} \cdot W_{i,s_{t+1}}^{(t)}, \quad (4.6)$$

**Algorithm 1** Greedy matching pursuit (MP)

---

**Require:** Data matrix  $F \in \mathbb{R}^{d \times m}$ , score matrix  $G \in \mathbb{R}^{m \times n}$  (where  $G = F^\top Q$ ), number of selections  $k$

**Ensure:** Selection set  $S$

- 1: Initialize  $S = \emptyset, W^{(0)} \leftarrow G^\top \{W \in \mathbb{R}^{n \times m}\}$
- 2: Precompute  $\Phi = F^\top F \in \mathbb{R}^{m \times m}$  {Data inner products}
- 3: **for**  $t = 1$  to  $k$  **do**
- 4:    $s_t \leftarrow \arg \max_{j \in [m] \setminus S} \sum_{i=1}^n (W_{i,j}^{(t-1)})^2$  {Select best candidate}
- 5:    $S \leftarrow S \cup \{s_t\}$
- 6:   **for**  $i = 1$  to  $n, j = 1$  to  $m$  with  $j \notin S$  **do**
- 7:      $W_{i,j}^{(t)} \leftarrow W_{i,j}^{(t-1)} - \Phi_{j,s_t} \cdot W_{i,s_t}^{(t-1)}$  {Update residuals}
- 8:   **end for**
- 9: **end for**
- 10: **return** Selection set  $S$

---

where  $\Phi_{j,s} = f_j^\top f_s$  are the precomputed inner products between data vectors. This update corresponds to the new residual  $r'_i = r_i - (r_i^\top f_{s_{t+1}}) f_{s_{t+1}}$ , yielding  $r_i^\top f_j = r_i^\top f_j - (f_j^\top f_{s_{t+1}}) r_i^\top f_{s_{t+1}}$  as expected.

Immediately, this suggests we can efficiently solve equation 4.3 by maintaining and updating a residual score matrix  $W$  *in-place*, while inner products of data vectors can be efficiently looked up by precomputing  $F^\top F$  in memory.

**Algorithm Initialization.** We initialize  $W^{(0)} = G = Q^\top F \in \mathbb{R}^{n \times m}$ , where  $G_{ij} = q_i^\top f_j$  represents the initial correlation between the  $i$ -th query and  $j$ -th data point. At iteration  $t = 0$ , we have  $r_i^{(0)} = q_i$  for all  $i \in [n]$ , so  $W_{ij}^{(0)} = r_i^{(0)\top} f_j = q_i^\top f_j = G_{ij}$ .

**Analysis of relaxation.** Although the greedy algorithm optimizes the relaxed objective, on controlled instances where we can enumerate the optimum, the linearization tracks the original objective (4.2) well. (Appx. C, Tab. 8).

The procedure is summarized in Algorithm 1. Notably, MP algorithms have been broadly studied since the seminal work of Mallat & Zhang (1993). Here, we contribute a new variation of MP for data selection.

## 4.2 COMPUTATIONAL COMPLEXITY AND PRACTICAL COSTS

The Greedy MP algorithm has a total runtime complexity of  $O(m^2d + mnk)$ , where  $m$  is the total number of data points,  $n$  is the number of scores, and  $k$  is the desired subset size. This breaks down as follows: (1) Precomputation of  $\Phi = F^\top F$  requires  $O(m^2d)$  operations, performed once; (2) Initialization of  $W^{(0)} = G$  is  $O(mn)$  if  $G$  is provided, or  $O(mnd)$  if computed from  $Q^\top F$ ; (3) For each of the  $k$  iterations, we select the best candidate ( $O(mn)$ ) and update the residual matrix  $W$  for  $O(m)$  entries ( $O(mn)$ ), giving  $O(mnk)$  for the iterative process. The memory complexity is  $O(m^2 + mn)$  for storing  $\Phi$  and  $W$ . In practice, since  $n$  is typically small, the algorithm scales nearly linearly with the dataset size  $m$  after the initial precomputation, making it practical for large-scale data selection.

We summarize the selection-phase runtime and resource requirements of the proposed Greedy MP against common baselines used in instruction/data selection. The focus is on asymptotic behavior with respect to:  $m$  (candidate pool size),  $k$  (selected subset size),  $d$  (embedding dimension),  $m_{\text{val}}$  (validation set size), and  $n$  (number of query/task score vectors; typically small). We contrast whether each method requires (a) a trained model for gradients/logits, and (b) an external validation subset.

Overall, the proposed methods yield strictly lower operational friction (no gradients, no held-out scoring) while matching or exceeding downstream performance (see Experiments). For detailed robustness and resource measurements (RAM usage / wall-clock), refer to Appendix A.

Method	Selection Time Complexity	Validation Data?	Model Training?
MP (Ours)	$O(mk)$	No	No
LESS	$O(m m_{\text{val}} d)$	Yes	Yes
DISF	$\tilde{O}(mkd^2)$	Yes	Yes
DSIR	$O(md + m_{\text{val}}d)$	Yes	Yes

Table 1: Selection complexity and requirements. For comparability with prior work, the table reports selection-phase complexity with  $n$  treated as a small constant (typically  $n = 1$  in baseline settings), excluding the one-time precomputation cost.  $\tilde{O}$  hides polylog factors.

## 5 EXPERIMENTS

### 5.1 DATASETS AND BASELINE MODELS

**Training data.** We use two training pools depending on the task. For instruction-tuning experiments evaluated on MT-Bench, BBH, and Alpaca-Eval-V2, we use the Alpaca-52k dataset, which contains 52,000 diverse instructions and demonstrations in English (Taori et al., 2023b). This dataset is commonly used as a benchmark instruction-tuning corpus; see, for example, (Zhao et al., 2024; Bukharin et al., 2024; Ge et al., 2024). In this setting, we apply Greedy MP to select training subsets of varying sizes: 512 samples (1% of full data) and 1000 samples (2% of full data). For mathematical reasoning experiments on GSM8K, we use the GSM8K training split as the fine-tuning pool and select subsets of 2.5%, 5%, 10%, and 20% of the full training set.

**Baseline methods.** We compare against several state-of-the-art data selection methods: **Random** (randomly selected subsets), **Full** (training on the complete dataset), **LIMA** (Zhao et al., 2024) (length-based selection for instruction tuning), **Alpagasus** (Chen et al., 2023b) (quality-based filtering), **CaR** (Ge et al., 2024) (clustering and ranking), **DSIR** (Xie et al., 2023) (importance resampling for domain adaptation), **DISF** (Fan et al., 2025) (iterative informative selection), and **LESS** (Xia et al., 2024) (gradient-similarity-based selection). For both training tasks, we include the baselines that are most relevant and retrain all methods from scratch under identical experimental settings to ensure fair comparison.

**Base models.** For instruction-tuning experiments, we use Mistral-7B (Jiang et al., 2023), Qwen3-4B, and Qwen3-8B (Yang et al., 2025). For GSM8K, we additionally evaluate Qwen3-32B (Yang et al., 2025). Training details for the main fine-tuning setups are provided in Table 6.

**Evaluation datasets.** We evaluated trained models on four benchmark suites: MT-Bench (Zheng et al., 2023), BIG-Bench Hard (BBH) (Suzgun et al., 2022), GSM8K (Cobbe et al., 2021), and Alpaca-Eval-V2<sup>1</sup>. For MT-Bench, we follow the standard protocol (Zheng et al., 2023) where models generate responses with vLLM (Kwon et al., 2023) for multi-turn conversations, and responses are scored by GPT-5<sup>2</sup> on a 1–10 scale. For BBH, we report accuracy against provided ground-truth answers. For GSM8K, models are fine-tuned on selected subsets of the GSM8K training split and evaluated on the held-out evaluation split using exact-match accuracy. For Alpaca-Eval-V2, we report the length-controlled win rate (LC win%), a debiased preference metric that adjusts for response length effects (Dubois et al., 2024).

### 5.2 IMPLEMENTATION

We run experiments with the Greedy MP algorithm (Algorithm 1) as follows: **Data embeddings**  $F$ . The Alpaca 52k dataset is a generic instruction-response dataset that includes training data with instruction, input (context), and answer (response) triplets. For each triplet, we concatenate the components into a single sentence [*Instruction; Context; Response*] as the input sentence. We then use pretrained ModernBERT-Base (Warner et al., 2024) as the embedding model and use outputs from the final hidden representation layer as embeddings. The context length is set to 2048, aligned with the training context window length. We apply a similar setting to the GSM8K dataset, where we consider the question as the instruction with context.

<sup>1</sup>[https://github.com/tatsu-lab/alpaca\\_eval](https://github.com/tatsu-lab/alpaca_eval)

<sup>2</sup><https://openai.com/index/introducing-gpt-5/>

Table 2: Performance on MT-Bench, BBH, and Alpaca-Eval-V2. Alpaca-Eval-V2 reports length-controlled win rate (LC win%, higher is better) (Dubois et al., 2024). Detailed results with standard errors are provided in Appendix Table 15. Best results are in **bold** and second best are underlined.

Method	Data	Mistral-7B			Qwen3-8B			Qwen3-4B		
		MT-Bench	BBH	Alpaca-Eval	MT-Bench	BBH	Alpaca-Eval	MT-Bench	BBH	Alpaca-Eval
Full	100% (52K)	3.64	<b>58.4</b>	5.77	7.74	76.7	<u>6.91</u>	6.69	73.2	<b>4.88</b>
Random	2% (1K)	3.66	57.4	4.71	7.98	72.3	4.99	7.59	69.4	2.62
LIMA	2% (1K)	<u>3.89</u>	55.6	2.69	<u>8.20</u>	75.8	4.24	7.45	73.3	2.48
Alpagasus	2% (1K)	3.71	56.5	4.81	7.86	71.9	5.51	7.59	63.6	4.17
CaR	2% (1K)	3.83	57.4	4.49	8.16	76.2	3.66	<u>7.68</u>	73.8	3.26
MP+MA	2% (1K)	3.83	57.6	<u>6.22</u>	7.99	<b>80.5</b>	5.64	<u>7.67</u>	<u>74.3</u>	3.74
MP+MA	1% (512)	<b>3.99</b>	56.6	<b>6.59</b>	<b>8.26</b>	79.1	4.85	<b>7.78</b>	68.0	<u>4.42</u>
MP+SC	2% (1K)	3.71	<u>57.7</u>	4.83	8.04	79.1	<b>7.05</b>	7.53	<b>74.4</b>	3.34
MP+SC	1% (512)	3.71	57.4	4.85	7.88	76.5	4.72	7.58	73.9	3.23

We consider two methods of generating score vectors: LLM evaluation and self-compression.

**Scores via LLM assessments.** We use GPT-4o<sup>3</sup> to generate evaluations from 0 to 5 in four criteria: coherence, helpfulness, accuracy, and difficulty, and use the total score across all four criteria as the score vector. The prompt is included in the supplementary material. Experiments using LLM quality assessments to obtain multi-attribute information are labeled as MP+MA.

**Scores via self-compression.** Given embeddings  $F \in \mathbb{R}^{d \times m}$  where each column  $f_i$  represents the embedding of sample  $i$ , we compute a quality score for each sample based on its alignment with the entire dataset. We calculate the dataset alignment score:

$$g_i = \sum_{j=1}^m f_i^\top f_j = (F^\top F \mathbf{1}_m)_i,$$

where  $\mathbf{1}_m \in \mathbb{R}^m$  is the all-ones vector. This score, equivalent to the  $i$ -th row sum of the Gram matrix  $F^\top F$ , measures how representative each sample is of the overall dataset structure in embedding space. Samples with higher scores are more central to the data distribution and thus more informative for training. This approach provides an *internal* quality signal derived purely from data geometry, without requiring external supervision. In our framework, this yields a single-dimensional score vector ( $n = 1$ ) that balances diversity with representativeness. Experiments using this self-compression scoring are labeled as MP+SC.

**Baseline implementation details.** To ensure fair comparison, we retrain all baseline methods following the reported setups in the original papers. All baselines use identical experimental settings (see Table 6 for details). For consistency, we use the same data preprocessing, embedding models, and evaluation protocols across all methods. For Alpaca experiments, we use a budget of 1000 examples, following prior work (Zhao et al., 2024; Stanford, 2023; Ge et al., 2024). For GSM8K, DSIR uses the same validation set as the target domain to compute importance scores, providing an oracle performance reference. For DISF, we start from a random set and iteratively add samples that maximize diversity within a batch. For LESS, we follow the original implementation to build gradient features and use the same validation set as for DSIR to select the most influential data.

All baselines use identical LoRA fine-tuning configurations, learning rates and training epochs specified in Table Table 6. This ensures that performance differences reflect data selection quality rather than training procedure variations.

### 5.3 MAIN RESULTS

Table 2 presents a comparison between Greedy MP methods and baseline approaches across three model architectures on MT-Bench, BBH, and Alpaca-Eval-V2. The results demonstrate the effectiveness of our proposed methods across different model scales and evaluation scenarios.

**Sample efficiency.** Across both instruction-tuning on Alpaca and verifiable math reasoning on GSM8K, Greedy MP delivers its largest gains in the low-data regime, where each training example has outsized influence. In this setting, MP-based selection consistently outperforms random subsampling

<sup>3</sup><https://openai.com/index/hello-gpt-4o/>

Table 3: GSM8K performance across models and data percentages.

Method	Qwen3-8B				Mistral-7B				Qwen3-32B			
	2.5%	5%	10%	20%	2.5%	5%	10%	20%	2.5%	5%	10%	20%
Full	81.96				52.46				87.64			
Random	78.32	80.21	79.53	81.65	37.76	43.52	45.72	46.74	78.24	87.14	85.67	87.29
DISF	75.06	79.98	79.98	80.67	<u>39.88</u>	41.55	41.55	42.46	81.65	87.04	87.04	<u>88.40</u>
DSIR	<u>80.74</u>	<b>81.50</b>	81.96	<b>83.24</b>	37.07	42.15	42.61	42.08	80.74	<b>89.16</b>	87.79	88.70
LESS	79.76	79.45	80.29	79.45	31.84	33.13	42.91	<u>47.69</u>	69.98	87.26	<u>87.95</u>	87.41
MP+MA	<b>81.58</b>	81.05	<b>83.09</b>	<b>83.24</b>	<b>42.99</b>	<b>45.64</b>	45.94	<u>47.46</u>	<b>84.31</b>	<u>87.57</u>	<b>88.25</b>	<b>88.78</b>
MP+SC	80.36	80.21	81.65	<u>82.26</u>	38.89	43.29	<b>46.63</b>	<b>49.81</b>	<u>84.15</u>	87.04	87.34	87.56

and often narrows much of the gap to full-data training using only a small fraction of the data. This trend holds across both preference-style evaluation (MT-Bench, Alpaca-Eval-V2) and accuracy-style evaluation (BBH, GSM8K), indicating that the benefit is not tied to a single metric family.

**Robustness across models and benchmarks.** The improvements are consistent across model families and sizes (Mistral and Qwen), and persist across complementary benchmark suites that probe different capabilities: instruction quality and helpfulness (MT-Bench), general reasoning (BBH), length-controlled preferences (Alpaca-Eval-V2), and strict verifiable reasoning (GSM8K). While oracle-style methods such as DSIR can excel at the very smallest budgets by leveraging target-domain validation information, Greedy MP remains competitive without requiring such oracle access, and its gains are stable as the budget increases.

Finally, we attribute these improvements to optimizing an objective that explicitly balances *quality* and *diversity*: the selection favors informative, high-signal examples while avoiding redundancy, which is especially critical under tight data budgets. The two instantiations reinforce this picture. MP+MA uses external multi-attribute assessments that align well with instruction quality and preference judgments, whereas MP+SC uses an internal, geometry-based signal that emphasizes representativeness; the two can therefore be complementary across benchmarks.

### 5.3.1 IMPACT OF DATA QUALITY ON MP+SC

We study data sources of varying quality using the cleaned Alpaca dataset<sup>4</sup> and a higher-quality variant filtered or rewritten from Alpaca (Taori et al., 2023a). Figure 2 breaks down MT-Bench categories on Mistral-7B, comparing MP+SC (1% of data) against the Full-data baseline. MP+SC largely tracks Full in applied and interactive categories, while the most persistent gap is concentrated in Math and occasionally Extraction, indicating broad capability coverage. Data cleansing shifts both methods upward, with the largest gains in open-ended categories, notably Humanities and Writing, and also Math and STEM, while changes in Reasoning and Extraction are smaller. This pattern suggests that MP+SC captures the quality signal and translates it into robust performance improvements across categories, especially those that are more open-ended and subjective in nature.

### 5.3.2 EMBEDDING ABLATIONS

We ablate the choice of embeddings feeding MP+SC on GSM8K, comparing a general-purpose encoder (Modern-BERT) versus a specialized reasoning encoder (Qwen-reasoning). We report accuracies for two budgets per base model. Larger, reasoning-specialized embeddings yield small but consistent gains at the same data budget, particularly for Qwen3-32B at 10%–20%. Modern-BERT remains competitive, especially on Mistral-7B at 10%. See Section 4.2 for the corresponding selection-phase resource comparison across methods.

**Selection Stability** To assess robustness to representation perturbations, we inject zero-mean Gaussian noise into GSM8K embeddings with  $\sigma \in \{10^{-4}, 10^{-3}, 10^{-2}\}$  and rerun MP+SC for 10% (747) and 20% (1494) budgets over three trials per noise level. Intersection-over-Union (IoU) of selected subsets remains  $\geq 85\%$  for  $\sigma \leq 10^{-3}$  and stays above 60% even at  $\sigma = 10^{-2}$ , indicating

<sup>4</sup><https://huggingface.co/datasets/yahma/alpaca-cleaned>

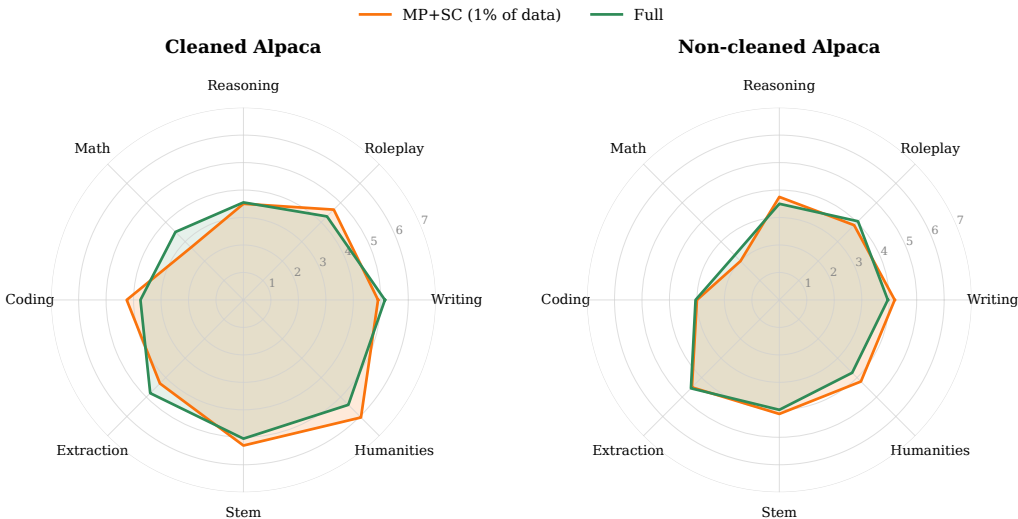


Figure 2: MT-Bench per-category average scores on Mistral-7B under cleaned vs non-cleaned Alpaca, comparing *only* MP+SC (1% of data) against the Full-data baseline. Scores are computed from GPT-5 judgments.

Table 4: Embedding ablation for MP+SC on GSM8K: Modern-BERT vs Qwen-reasoning.

Base Model	Data Size	Modern-BERT	Qwen-reasoning
Mistral-7B	20% (1494)	49.81%	<b>50.27%</b>
Mistral-7B	10% (747)	<b>46.63%</b>	46.25%
Qwen3-32B	20% (1494)	87.57%	<b>88.02%</b>
Qwen3-32B	10% (747)	87.34%	<b>87.72%</b>

that mild embedding drift has negligible impact on which samples are chosen while severe noise still preserves a majority of core items. Detailed setup and the full table are provided in Appendix A.1.1.

## CONCLUSION AND FUTURE WORK

We introduced *Greedy Information Projection* (GIP), an information-theoretic framework for LLM data selection that unifies quality and diversity through a single mutual information objective. This formulation yields a closed-form criterion together with an efficient greedy matching-pursuit algorithm that scales linearly with dataset size after precomputation, without gradient computation or external validation sets. Across instruction-following benchmarks (MT-Bench, BBH, Alpaca-Eval-V2) and mathematical reasoning (GSM8K), GIP delivers substantial data-efficiency gains: with only 1–20% of the training data, it consistently matches or exceeds full-dataset performance. The two variants are complementary: MP+MA is strongest on instruction-following and preference-style metrics, whereas MP+SC achieves strong reasoning performance without external supervision.

**Future directions.** Natural extensions include scaling to pretraining-scale datasets through distributed or streaming variants, extending the framework beyond Gaussian assumptions to handle noisy quality signals, adapting it to multimodal and continual-learning settings, and developing formal approximation guarantees and finite-sample analyses for the greedy algorithms.

## REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Olivier Bachem, Mario Lucic, and Andreas Krause. Practical coreset constructions for machine learning. In *International Conference on Machine Learning*, pp. 311–319. PMLR, 2017.
- Alexander Bukharin, Shiyang Li, Zhengyang Wang, Jingfeng Yang, Bing Yin, Xian Li, Chao Zhang, Tuo Zhao, and Haoming Jiang. Data Diversity Matters for Robust Instruction Tuning, November 2024.
- Daniele Calandriello, Gang Niu, and Masashi Sugiyama. Semi-supervised information-maximization clustering. *Neural networks*, 57:103–111, 2014.
- Daoyuan Chen, Yilun Huang, Zhijian Ma, Heseng Chen, Xuchen Pan, Ce Ge, Dawei Gao, Yuexiang Xie, Zhaoyang Liu, Jinyang Gao, Yaliang Li, Bolin Ding, and Jingren Zhou. Data-Juicer: A One-Stop Data Processing System for Large Language Models, September 2023a.
- Hao Chen, Abdul Waheed, Xiang Li, Yidong Wang, Jindong Wang, Bhiksha Raj, and Marah I. Abidin. On the Diversity of Synthetic Data and its Impact on Training Large Language Models, October 2024.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, et al. Alpapasus: Training a better alpaca with fewer data. *arXiv preprint arXiv:2307.08701*, 2023b.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Rohan Deb, Kiran Thekumparampil, Kousha Kalantari, Gaurush Hiranandani, Shoham Sabach, and Branislav Kveton. Fishersft: Data-efficient supervised fine-tuning of language models using information gain. *arXiv preprint arXiv:2505.14826*, 2025.
- Kien Do, Truyen Tran, and Svetha Venkatesh. Clustering by maximizing mutual information across views. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9928–9938, 2021.
- Qianlong Du, Chengqing Zong, and Jiajun Zhang. MoDS: Model-oriented Data Selection for Instruction Tuning, November 2023.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled alpacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*, 2024.
- Ziqing Fan, Siyuan Du, Shengchao Hu, Pingjie Wang, Li Shen, Ya Zhang, Dacheng Tao, and Yanfeng Wang. Combatting dimensional collapse in llm pre-training data via diversified file selection, 2025. URL <https://arxiv.org/abs/2504.20644>.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness, 07 2024. URL <https://zenodo.org/records/12608602>.
- Yuan Ge, Yilun Liu, Chi Hu, Weibin Meng, Shimin Tao, Xiaofeng Zhao, Hongxia Ma, Li Zhang, Boxing Chen, Hao Yang, et al. Clustering and ranking: Diversity-preserved instruction selection through expert-aligned quality estimation. *arXiv preprint arXiv:2402.18191*, 2024.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru,

Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Voleti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippou Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik

Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihalescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

Zengyou He, Xiaofei Xu, and Shengchun Deng. k-anmi: A mutual information based clustering algorithm for categorical data. *Information Fusion*, 9(2):223–233, 2008.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. URL <https://arxiv.org/abs/2203.15556>.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models, 2020. URL <https://arxiv.org/abs/2001.08361>.

Zeger F Knops, JB Antoine Maintz, Max A Viergever, and Josien PW Pluim. Normalized mutual information based registration using k-means clustering and shading correction. *Medical image analysis*, 10(3):432–439, 2006.

Alexander Kraskov, Harald St ogbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 69(6):066138, 2004.

- Alexander Kraskov, Harald Stögbauer, Ralph G Andrzejak, and Peter Grassberger. Hierarchical clustering using mutual information. *Europhysics Letters*, 70(2):278, 2005.
- Alex Kulesza, Ben Taskar, et al. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286, 2012.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. Deduplicating Training Data Makes Language Models Better, March 2022.
- Kuan-Yun Lee. *New Information Inequalities with Applications to Statistics*. University of California, Berkeley, 2022.
- Kuan-Yun Lee and Thomas A Courtade. Linear models are most favorable among generalized linear models. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 1213–1218. IEEE, 2020.
- Kuan-Yun Lee and Thomas A Courtade. Minimax bounds for generalized pairwise comparisons. In *2021 International Conference on Machine Learning (ICML) Workshop on Information-Theoretic Methods for Rigorous, Responsible, and Reliable Machine Learning*, 2021.
- Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, 41(12):3397–3415, 1993.
- Mark EJ Newman, George T Cantwell, and Jean-Gabriel Young. Improved mutual information measure for clustering, classification, and community detection. *Physical Review E*, 101(4):042304, 2020.
- Aditya Sanghi. Info3d: Representation learning on 3d objects using mutual information maximization and contrastive learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pp. 626–642. Springer, 2020.
- Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach, June 2018.
- Burr Settles. Active learning literature survey. Technical Report CS-TR-1648, University of Wisconsin–Madison, 2009. URL <https://burrsettles.com/pub/settles.activelearning.pdf>.
- Stanford. Stanford Alpaca: An Instruction-following LLaMA Model. Tatsu’s shared repositories, September 2023.
- Masashi Sugiyama, Gang Niu, Makoto Yamada, Manabu Kimura, and Hirotaka Hachiya. Information-maximization clustering based on squared-loss mutual information. *Neural Computation*, 26(1): 84–131, 2014.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022. URL <https://arxiv.org/abs/2210.09261>.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023a.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7, 2023b.

- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-Instruct: Aligning Language Model with Self Generated Instructions, December 2022.
- Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, et al. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663*, 2024.
- Max A Woodbury. *Inverting modified matrices*. Department of Statistics, Princeton University, 1950.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. Less: Selecting influential data for targeted instruction tuning. *arXiv preprint arXiv:2402.04333*, 2024. URL <https://arxiv.org/abs/2402.04333>.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy Liang. Data selection for language models via importance resampling. In *NeurIPS*, 2023. URL <https://arxiv.org/abs/2302.03169>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. An unsupervised sentence embedding method by mutual information maximization. *arXiv preprint arXiv:2009.12061*, 2020.
- Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Long is more for alignment: A simple but tough-to-beat baseline for instruction fine-tuning. *arXiv preprint arXiv:2402.04833*, 2024.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

## APPENDIX

We include proofs of main theorems, additional ablation studies, robustness and resource analyses, and a discussion on statistical significance in the appendix.

## A ROBUSTNESS AND RESOURCE ANALYSIS

## A.1 ROBUSTNESS ACROSS DATA PERCENTAGES AND MODEL SIZES

We summarize sensitivity across data budgets (2.5%, 5%, 10%, 20%) and model sizes (7B, 8B, 32B). MP+MA maintains strong gains at low budgets; MP+SC is consistently competitive without external scores. Detailed per-budget tables are provided alongside standard errors in Table 16.

## A.1.1 SELECTION STABILITY UNDER EMBEDDING NOISE

We probe stability by adding zero-mean Gaussian noise to GSM8K embeddings prior to running MP+SC. We evaluate three noise levels  $\sigma \in \{10^{-4}, 10^{-3}, 10^{-2}\}$  and perform three independent trials per level for two budgets: 10% (747 samples) and 20% (1494 samples). We report the mean Intersection-over-Union (IoU) of selected subsets (% units) with  $\pm$  standard deviation.

Noise Level ( $\sigma$ )	Data Size 747 (10%)	Data Size 1494 (20%)
1e-4	95.89 $\pm$ 0.32	91.72 $\pm$ 0.36
1e-3	94.20 $\pm$ 0.66	87.85 $\pm$ 0.39
1e-2	66.32 $\pm$ 1.13	61.74 $\pm$ 0.29

Overall, we observe (i) High stability across mild noise: at  $\sigma \leq 10^{-3}$ ,  $\text{IoU} \geq 85\%$  indicates minor embedding variations have minimal impact on selection outcomes; (ii) Graceful degradation under strong noise: even at  $\sigma = 10^{-2}$ , IoU remains above 60%, suggesting that core high-quality diverse samples are preserved. These results support the practical stability of our information projection framework and explain why downstream training remains robust under moderate embedding perturbations. See the main text summary in Section 5.3.2. We discuss practical considerations for Gram matrix usage. For large  $m$  and  $d$ , we use blockwise products to avoid materializing  $F^\top F$  fully and maintain peak memory within CPU limits. Selection-phase timing curves indicate near-linear scaling in  $k$  for Greedy MP. We also report typical RAM footprints on AMD CPU for representative  $m, d$  configurations.

## A.2 RUNTIME AND MEMORY PERFORMANCE

We provide detailed wall-clock and peak memory profiling for the Gram (inner product) matrix preprocessing and the subsequent greedy selection phase. Measurements are taken on a single AMD CPU machine without sharding or distributed execution. Preprocessing refers to computing the Gram matrix (or equivalent blockwise products); selection refers to running Greedy MP for the indicated subset budget.

Table 5: Runtime and memory profiling for preprocessing (Gram matrix computation) and selection. Preprocessing is a one-time cost that can be amortized across multiple selection runs with different budgets or scoring signals.

Dataset	Pool Size $m$	Subset %	Preproc. / Selection Time (s)	Peak RAM (Preproc. / Select.)
GSM8K	7k	10%	7.01 / 0.86	0.3 GB / 0.04 GB
GSM8K	7k	20%	7.01 / 0.87	0.3 GB / 0.04 GB
GSM8K	7k	50%	7.01 / 0.87	0.3 GB / 0.04 GB
Alpaca	52k	10%	214.11 / 25.12	10.8 GB / 0.3 GB
Alpaca	52k	20%	214.11 / 49.67	10.8 GB / 0.3 GB
Alpaca	52k	50%	214.11 / 123.62	10.8 GB / 0.3 GB

**Findings.** (i) *Gram matrix computation is amortized*: the  $O(m^2d)$  preprocessing takes  $\sim 7s$  for GSM8K and  $\sim 214s$  for Alpaca but is reused across budgets and score variants. (ii) *Linear selection scaling*: selection times grow roughly linearly with  $k$  (subset size), consistent with  $O(mk)$  complexity. (iii) *Memory efficiency*: peak RAM is dominated by preprocessing; selection adds negligible overhead. (iv) *Practical applicability*: On typical instruction-tuning corpus sizes, single-machine execution is feasible; for substantially larger  $m$  we can employ low-rank approximations or distributed blockwise multiplication. We leave this as future work.

## B ETHICS / BROADER IMPACT

This work studies data selection for language-model fine-tuning. It uses public datasets under their original licenses. **Potential risks.** (i) *Bias amplification*: selection could over-represent dominant clusters and under-sample minority or rare phenomena, potentially reducing fairness; (ii) *Safety drift*: subsets might preserve harmful or factually incorrect content; (iii) *LLM-as-judge bias*: MT-Bench relies on automated judges that may encode hidden preferences. **Mitigations.** We (a) normalize and de-duplicate embeddings, (b) allow  $\epsilon$ -regularization and diversity-aware gains to reduce over-concentration, (c) apply the original dataset safety filters. **Intended use.** Our method targets research/industrial model training where compute and data curation are constrained. It should not be used to curate content intended to target protected groups or to optimize for harmful outputs. **Transparency.** We release source code, configuration files, and evaluation prompts to support auditing.

## C REPRODUCIBILITY

We release code and configuration files to reproduce our main algorithm. **Data and licenses.** We use Alpaca-52k (CC BY-NC 4.0) and the GSM8K training split as training data, and MT-Bench, BBH, Alpaca-Eval-V2, and the held-out GSM8K evaluation split as benchmarks; links and license terms are listed in the README. **Environment.** Experiments were run on NVIDIA A100 80GB and V100 32GB GPUs, depending on the model and experiment. We provide a conda YAML and exact package versions. **Training.** We report the shared fine-tuning recipe used across the main runs, covering Mistral-7B and the Qwen3 family (Qwen3-4B, Qwen3-8B, and Qwen3-32B), in Table 6. We also report the effective token budgets (1–20%) and the corresponding evaluation settings. **Selection.** Our implementation exposes the Greedy-MP procedure used in the main paper. **Evaluation.** We reuse the prompt as in (Zheng et al., 2023) and (Gao et al., 2024) for GSM8K/BBH (exact-match normalization), and MT-Bench judge templates. For all tasks, we report results as mean  $\pm$  standard error across data instances and random seeds. **Determinism.** We fix seeds at all stages (data loader, CUDA/cuDNN, model training).

## D REGULARIZATION DETAILS

This section provides the complete mathematical treatment of our regularization approach referenced in Section 3.

### D.1 MOTIVATION FOR REGULARIZATION

When GIP is applied without approximation, the constraint  $F^\top Q = G$  requires solving for  $Q$  given feature matrix  $F \in \mathbb{R}^{d \times m}$  and score matrix  $G \in \mathbb{R}^{m \times n}$ . In practical scenarios:

- $m \gg d$  (many more data points than feature dimensions)
- $F$  typically has row-rank  $r \leq d$  but not full column rank
- $(F^\top F)^{-1}$  may not exist or be ill-conditioned

## D.2 TIKHONOV REGULARIZATION

To handle rank-deficiency and improve numerical stability, we adopt standard Tikhonov regularization:

$$\Sigma_\epsilon := F^\top F + \epsilon I_m, \quad \epsilon > 0.$$

We define the regularized left pseudoinverse:

$$F_\epsilon^+ := F^\top (F F^\top + \epsilon I_d)^{-1}$$

and construct the regularized query coefficient matrix:

$$Q_\epsilon := (F_\epsilon^+)^{\top} G.$$

## D.3 PROPERTIES OF THE REGULARIZED SOLUTION

The regularized query matrix  $Q_\epsilon$  has several important properties:

**Unique minimizer.**  $Q_\epsilon$  is the unique minimizer of the regularized least squares problem:

$$\min_Q \|F^\top Q - G\|_F^2 + \epsilon \|Q\|_F^2.$$

**Controlled approximation error.** The regularization introduces a bounded approximation error:

$$\|F^\top Q_\epsilon - G\|_F \leq \epsilon \|Q_\epsilon\|_F.$$

**Numerical stability.** For any  $\epsilon > 0$ , both  $\Sigma_\epsilon$  and  $F F^\top + \epsilon I_d$  are positive definite and hence invertible.

## D.4 IMPLEMENTATION IN MAIN RESULTS

With this regularization framework:

- All occurrences of  $(F^\top F)^{-1}$  in our derivations are replaced by  $\Sigma_\epsilon^{-1}$
- Identities that relied on  $F^\top Q = G$  hold approximately:  $F^\top Q_\epsilon \approx G$
- As  $\epsilon \rightarrow 0$ , we recover the original formulation when  $(F^\top F)^{-1}$  exists
- The approximation quality can be controlled by choosing appropriate  $\epsilon$ .

## E PROOFS

### E.1 PROOF OF THEOREM 1

*Proof.* We start by expanding equation 3.3 via

$$\arg \max_S I(Z_Q; Z_{F_S}) = \arg \max_S \left( \log \det(\Sigma_{F_S}) - \log \det \left( \Sigma_{F_S} - \Sigma_{Q, F_S} \Sigma_Q^{-1} \Sigma_{Q, F_S}^\top \right) \right). \quad (\text{E.1})$$

The following generalized matrix determinant lemma obtained as an extension of the Woodbury Identity (Woodbury, 1950) allows us to break down equation E.1.

**Lemma 4** (Matrix Determinant Lemma, Woodbury (1950)). *Suppose  $A \in \mathbb{R}^{n \times n}$  is invertible. Then, for any matrices  $U, V \in \mathbb{R}^{n \times m}$ ,*

$$\det(A + UV^\top) = \det(I + V^\top A^{-1}U) \det(A).$$

Then, the terms within the maximum in equation E.1 become

$$\begin{aligned} -\log \det \left( I - \Sigma_{Q, F_S} \Sigma_Q^{-1} \Sigma_{Q, F_S}^\top \Sigma_{F_S}^{-1} \right) &= -\log \det \left( I - F^\top Q (Q^\top Q)^{-1} Q^\top F (F^\top F)^{-1} \right) \\ &= -\log \det \left( I + V^\top A^{-1}U \right) \end{aligned} \quad (\text{E.2})$$

where  $V = -Q^\top F_S$ ,  $A = Q^\top Q$  and  $U = Q^\top F_S(F_S^\top F_S)^{-1}$ . Then, we can directly apply Lemma 4 to get

$$\text{equation E.2} = -\log \det(A + UV^\top) \det(A^{-1}). \quad (\text{E.3})$$

Combining equation E.1, equation E.2 and equation E.3 yields

$$\arg \max_S I(Z_Q; Z_{F_S}) = \arg \min_S \log \det(Q^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) Q). \quad (\text{E.4})$$

□

## E.2 THEOREM FOR QUALITY PRESERVATION

**Theorem 5.** *When data is mutually orthogonal, i.e.,  $f_i^\top f_j = 0$  for all  $i \neq j$ , and  $G$  consists of a single score vector  $g_1 \in \mathbb{R}^m$ , the maximizing solution  $S$  of equation 3.4 with  $|S| = k$  is a solution of*

$$\arg \max_{S:|S|=k} G_S^\top G_S. \quad (\text{E.5})$$

*In other words, solving equation E.5 returns the indices of the top  $k$  scores with largest absolute value.*

*Proof.* Based on the assumption that  $F$  is normalized, the property  $f_i^\top f_j = 0$  for all  $i \neq j$  yields  $F^\top F = I_m$  and  $F_S^\top F_S = I_k$  for any selection  $S$  with  $|S| = k$ . Consequently,

$$\arg \min_S \log \det(Q^\top (I - F_S(F_S^\top F_S)^{-1} F_S^\top) Q) = \arg \min_S (|q_1|^2 - q_1^\top F_S F_S^\top q_1).$$

Recall the definition of  $G := F^\top q_1$  and  $G_S := F_S^\top q_1$ , and the desired result follows. □

## E.3 PROOF OF THEOREM 2

We restate the theorem as follows: If mutual information objective is maximized with optimal set  $S^*$  in Theorem 1, then there exists a  $\delta_{S^*}$  such that

$$\|F_{S^*}^\top Q\|_2 \geq \delta_{S^*} \sqrt{1 - \left(\frac{\eta}{\det(Q^\top Q)}\right)^{1/r}}, \quad (\text{E.6})$$

where  $\eta = \det(Q^\top (I - P_{F_{S^*}}) Q)$ , and  $r$  is the row rank of  $Q$

*Proof.* Take reduced QRs with orthonormal bases:

$$Q = UR_Q, \quad F_S = VR, \quad S = V^\top U, \quad P_{F_S} = VV^\top, \quad (\text{E.7})$$

where  $U \in \mathbb{R}^{n \times r}$ ,  $V \in \mathbb{R}^{n \times t}$  have orthonormal columns and  $R_Q \in \mathbb{R}^{r \times r}$ ,  $R \in \mathbb{R}^{t \times t}$  are invertible. Define

$$\eta(S) = \det(Q^\top (I - P_{F_S}) Q). \quad (\text{E.8})$$

**Reduce to  $S$ .** Using  $Q = UR_Q$  and  $P_{F_S} = VV^\top$ ,

$$Q^\top (I - P_{F_S}) Q = R_Q^\top (I - S^\top S) R_Q \implies \eta = \det(Q^\top Q) \det(I - S^\top S). \quad (\text{E.9})$$

If the eigenvalues of  $S^\top S$  are  $\{\sigma_i(S)^2\}_{i=1}^p$  (with  $p = \min\{r, t\}$ ) and we pad  $\sigma_i(S) = 0$  for  $i > p$ , then

$$\frac{\eta}{\det(Q^\top Q)} = \prod_{i=1}^r (1 - \sigma_i(S)^2). \quad (\text{E.10})$$

Let  $a_i = 1 - \sigma_i(S)^2 \in [0, 1]$ . If all  $a_i > g = (\eta / \det(Q^\top Q))^{1/r}$ , then

$$\prod_{i=1}^r a_i > g^r, \quad (\text{E.11})$$

contradicting the identity E.10 above. Hence some  $i^*$  satisfies  $a_{i^*} \leq g$ , i.e.,

$$\|S\|_2^2 = \sigma_{\max}(S)^2 \geq 1 - \left(\frac{\eta}{\det(Q^\top Q)}\right)^{1/r}. \quad (\text{E.12})$$

Since  $S = V^\top U$ , we obtain the main results:

$$\|V^\top U\|_2^2 \geq 1 - \left(\frac{\eta}{\det(Q^\top Q)}\right)^{1/r} \quad (\text{E.13})$$

When  $Q^\top Q = I_r$ :

$$\|V^\top U\|_2^2 \geq 1 - \eta^{1/r} \quad (\text{E.14})$$

With  $F_S^\top Q = R^\top S R_Q$  and using the inequality  $\|ABC\|_2 \geq \sigma_{\min}(A)\|B\|_2\sigma_{\min}(C)$ :

$$\|F_S^\top Q\|_2 \geq \sigma_{\min}(F_S)\sigma_{\min}(Q)\|V^\top U\|_2 \quad (\text{E.15})$$

$$\geq \sigma_{\min}(F_S)\sigma_{\min}(Q)\sqrt{1 - \left(\frac{\eta}{\det(Q^\top Q)}\right)^{1/r}}, \quad (\text{E.16})$$

which simplifies to

$$\|F_S^\top Q\|_2 \geq \sigma_{\min}(F_S)\sqrt{1 - \eta^{1/r}} \quad (\text{E.17})$$

when  $Q^\top Q = I_r$ .  $\square$

## F EXPERIMENT DETAILS

### F.1 DATASETS AND BENCHMARKS

**Training corpora.** We use two training pools, depending on the task. For instruction-tuning experiments evaluated on MT-Bench, BBH, and Alpaca-Eval-V2, we use Alpaca-52k (Taori et al., 2023b), a 52k English instruction-response corpus generated with text-davinci-003. It is the source dataset for our selected subsets and for the Alpaca-based baselines; for these experiments, we consider 1% and 2% subset budgets. For verifiable math reasoning, we use the GSM8K training split (Cobbe et al., 2021) as the selection and fine-tuning pool, with subset budgets of 2.5%, 5%, 10%, and 20% of the full training set. Alpaca-52k is released under CC BY-NC 4.0.

**Derived training subsets and baselines.** In addition to our own selected subsets, we compare against two widely used Alpaca-derived baselines. CaR-1k (Ge et al., 2024) selects 1,000 examples from Alpaca-52k using clustering and ranking to encourage quality and diversity. Alpagasus-1k (Chen et al., 2023b) selects 1,000 examples from Alpaca-52k using LLM-based quality filtering.

**Evaluation benchmarks.** We evaluate the current experiments on MT-Bench, BBH, GSM8K, and Alpaca-Eval-V2. MT-Bench (Zheng et al., 2023) is a multi-turn conversational benchmark; we follow the standard protocol, generate responses with vLLM (Kwon et al., 2023), and use GPT-5 as judge on a 0–10 scale. BBH (BIG-Bench Hard) (Suzgun et al., 2022) is a collection of challenging reasoning tasks, for which we report exact-match accuracy. For GSM8K (Cobbe et al., 2021), we train on the standard GSM8K training split and report exact-match accuracy on the held-out evaluation split under standard normalization. Alpaca-Eval-V2 measures instruction-following preference with length-controlled win rate (Dubois et al., 2024).

### F.2 TRAINING DETAILS

Hyperparameters and training details for reproducing our work are provided in Table 6. All of our models are trained in the Hugging Face framework with LoRA fine-tuning (Hu et al., 2021). We apply LoRA fine-tuning to all linear layers for the Mistral and Qwen3 runs with LoRA parameters  $\{r = 8, \alpha = 16\}$ . The table below lists the representative instruction-tuning setups, and the remaining Qwen3 runs follow the same Qwen3 fine-tuning recipe unless noted otherwise. For all Qwen3 experiments, we disable thinking mode during supervised fine-tuning and downstream

Table 6: Representative training hyperparameters for the main fine-tuning setups.

Datasets	Data Size	# GPUs	Epochs	LR	LR Scheduler	Context Win. Len.
<i>Qwen3-4B/8B</i>						
Alpaca-52k	52k	8	10	3e-5	Linear	2048
Alpapasus-1k	1k	4	10	3e-5	Linear	2048
CaR-1k	1k	4	10	3e-5	Linear	2048
Random-1k	1k	4	10	3e-5	Linear	2048
MP-512	512	4	10	3e-5	Linear	2048
<i>Mistral-7B-v0.1</i>						
Alpaca-52k	52k	8	10	3e-5	Linear	2048
Alpapasus-1k	1k	4	10	3e-5	Linear	2048
CaR-1k	1k	4	10	3e-5	Linear	2048
Random-1k	1k	4	10	3e-5	Linear	2048
MP-512	512	4	10	3e-5	Linear	2048

evaluation. For fair comparison with baseline models, we use the same context length and epoch budget shown in Table 6, even when these differ from the original settings reported in (Chen et al., 2023b).

**Additional Qwen3 settings.** Qwen3-4B instruction-tuning runs use the shared Qwen3 instruction-tuning recipe summarized in Table 6. For GSM8K, the Qwen3-32B runs use the same Qwen3 LoRA configuration ( $r = 8$ ,  $\alpha = 16$ ), linear learning-rate schedule, 10-epoch budget, and 2048-token context window as the other Qwen3 models, but are trained on GSM8K subsets rather than the Alpaca-derived subsets listed in the table; thinking mode is disabled throughout.

### F.3 MULTI-ATTRIBUTE SCORING PROMPT

The prompt used for scoring data is provided in Table 7. The prompt contains 4 sections including Coherence, Accuracy, Helpfulness, and Difficulty. LLM will prompt the response for each section on scale of 0 to 5.

## G DISCUSSIONS

### G.1 LINEARIZATION OF MATCHING PURSUIT

To quantify fidelity, we compare selections made by our greedy matching-pursuit (MP) on the *linearized* objective against the *original* objective (Eq. 4.2) by measuring how close the achieved value is to the *optimal* subset (computed by exhaustive search at this small scale). We generate  $F \in \mathbb{R}^{30 \times 10}$  with i.i.d.  $\mathcal{N}(0, 1)$  entries and  $Q \in \mathbb{R}^{30 \times 1}$  with i.i.d.  $\text{Unif}[0, 1]$ , run 100 independent trials, and report (mean  $\pm$  std) of  $\text{Objective}(\text{method's } S_k) / \text{Objective}(S_k^*)$  for selection sizes  $k = 1, \dots, 10$ . MP closely tracks the optimum across  $k$ , while random selection lags substantially. See Table 8.

**Takeaway.** Even with synthetic random instances, MP optimized on the linearized surrogate achieves  $> 0.9 \times$  the optimal *original* objective by  $k \geq 2$ , while random requires much larger  $k$  to catch up.

## H CASE STUDY

This section presents several MT-Bench examples to illustrate qualitative differences across methods. The qualitative assessments in this section are based on manual inspection by the authors rather than judge-model scores. For context, we also include historical per-category MT-Bench summaries from the corresponding experimental runs. We compare MP+MA and MP+SC against representative baselines including Alpapasus-1k, Alpaca-52k, and CaR-1k across writing, coding, and roleplay prompts.

Table 7: Evaluation rubrics used for calculating LLM-based scores in our multi-attribute method.

---

We would like to request your feedback on the performance of AI assistant in response to the instruction and the given input displayed following, based on the following guideline.

**1. Coherence**

*What to judge:* Logical flow, internal consistency, clarity.

*Score anchors:*

- 0 – Nonsensical or self-contradictory
- 1 – Confusing, frequent jumps
- 2 – Some lapses but understandable
- 3 – Clear and orderly
- 4 – Excellent narrative flow and transitions
- 5 – Flawless logic, elegant structure, exceptionally smooth

**2. Correctness / Accuracy**

*What to judge:* Factual accuracy and fidelity to the prompt.

*Score anchors:*

- 0 – Main claim wrong or unsupported
- 1 – Many errors or hallucinations
- 2 – Minor slips or partially met requirements
- 3 – Fully correct; only trivial issues
- 4 – Rigorous and well-sourced
- 5 – Authoritative, thoroughly sourced, withstands expert scrutiny

**3. Helpfulness**

*What to judge:* Usefulness, completeness, depth, alignment with the question.

*Score anchors:*

- 0 – Provides no help
- 1 – Little usable information
- 2 – Partially helpful but key gaps
- 3 – Satisfies the question well
- 4 – Exceeds expectations; anticipates follow-ups, adds examples
- 5 – Exceptional: deep insights, meta-guidance, multiple perspectives

**4. Difficulty**

*What to judge:* Cognitive load of the question (not the answer).

*Score anchors:*

- 0 – Trivial recall (e.g., basic facts)
- 1 – Basic high-school knowledge
- 2 – Multi-step reasoning or college-level facts
- 3 – Specialized insight or synthesis of several topics
- 4 – Advanced graduate-level or cross-disciplinary reasoning
- 5 – Expert-level, open-ended, or research-frontier challenge

**Judging Procedure**

1. Read the question and answer in full.
  2. Evaluate coherence first, then fact-check key claims.
  3. Score each dimension independently.
  4. Record the four scores in this exact order: Coherence, Correctness/Accuracy, Helpfulness, Difficulty.
  5. Output only these four integers as a comma-separated list wrapped with `¡Rstç¡/Rstç` with short and concise reasoning
-

Selection Size	MP / Optimal	Random / Optimal
1	0.958 ± 0.108	0.255 ± 0.304
2	0.911 ± 0.120	0.320 ± 0.258
3	0.877 ± 0.115	0.395 ± 0.246
4	0.874 ± 0.101	0.482 ± 0.214
5	0.870 ± 0.095	0.574 ± 0.225
6	0.889 ± 0.088	0.655 ± 0.211
7	0.905 ± 0.079	0.717 ± 0.191
8	0.934 ± 0.070	0.810 ± 0.170
9	0.969 ± 0.044	0.900 ± 0.138
10	1.000 ± 0.000	1.000 ± 0.000

Table 8: Approximation fidelity of the linearized objective: ratio of the *original* objective (Eq. 4.2) achieved by the method’s selection to the *optimal* value at each  $k$ . MP (ours) is consistently near-optimal; random trails behind. For  $k=10$  all methods select all items, hence ratio = 1.

Table 9: Historical per-category MT-Bench results for Mistral-7B on the *non-cleaned* Alpaca-52k setting. The table is included as quantitative context for the qualitative case study below. We report runs with 512 selected examples for MP+MA and MP+SC.

Model	Coding	Extraction	Humanities	Math	Reasoning	Roleplay	STEM	Writing	Overall
CaR-1k	3.75	5.90	5.90	1.10	3.15	5.95	6.10	4.65	4.56
MP+MA-512	3.70	5.40	8.45	1.70	3.80	6.50	7.03	6.85	5.43
MP+SC-512	4.40	6.00	7.38	1.45	4.40	5.30	6.70	5.95	5.20
Alpaca-52k	4.00	5.85	6.55	1.10	4.00	4.95	6.95	6.35	4.97
Alpagasus-1k	3.43	5.40	5.56	2.00	2.75	7.11	5.58	7.04	4.86

### H.1 EXAMPLE 1

Table 12 shows a practical writing task that requires a concise and situationally appropriate message. MP+MA gives the most direct and fully compliant answer, while MP+SC remains aligned with the task but includes placeholder-style phrasing. Alpaca-52k and Alpagasus-1k are both serviceable, though somewhat more verbose than necessary, and CaR is less well matched to the desired level of brevity.

### H.2 EXAMPLE 2

This example (see Table 13) is challenging because the prompt suggests the code is buggy even though it is actually correct. MP+MA identifies this most clearly. MP+SC comes closer to the correct interpretation than the baselines, but still introduces some uncertainty. The remaining models are more likely to follow the misleading premise of the prompt.

### H.3 EXAMPLE 3

In the roleplay example shown in Table 14, MP+MA best balances voice, emotional specificity, and conciseness. MP+SC captures much of the requested content as well, though with more self-clarification before settling on the answer. Alpaca-52k remains correct but comparatively terse, while CaR is less responsive to the roleplay framing.

## I DETAILED EXPERIMENTAL RESULTS WITH STANDARD ERRORS

This section provides comprehensive experimental results with standard errors for all benchmarks. For MT-Bench, BBH, and GSM8K, we compute standard errors as  $SE = \text{std}/\sqrt{n}$  where  $n$  is the number of test samples: MT-Bench (160 turns), BBH (6511 questions), and GSM8K (1319 questions). For Alpaca-Eval-V2, we report the standard error values produced by the Alpaca-Eval evaluation tool.

Table 10: Per-category MT-Bench results for Mistral-7B on the *cleaned* Alpaca-52k setting. The table is included as quantitative context for the qualitative case study below. We report the full cleaned-Alpaca run and MP+SC with 512 selected examples.

Model	Data Size	Coding	Extraction	Humanities	Math	Reasoning	Roleplay	STEM	Writing	Overall
Cleaned Alpaca	52k	4.50	6.40	9.20	1.85	4.45	6.60	7.60	7.73	5.92
MP+SC	512	4.45	5.95	7.75	3.45	4.65	7.35	7.35	6.85	5.98

Table 11: Per-category MT-Bench results for Llama-13B on the *non-cleaned* Alpaca-52k setting. The table is included as quantitative context for the qualitative case study below. We report runs with 512 selected examples for MP+MA and MP+SC.

Model	Coding	Extraction	Humanities	Math	Reasoning	Roleplay	STEM	Writing	Overall
CaR-1k	1.45	3.80	6.60	1.15	2.35	6.93	6.40	5.20	4.23
Alpapasus-1k	1.10	4.35	5.25	1.20	2.25	5.00	5.50	5.45	3.76
Random-1k	1.15	3.80	5.05	1.15	2.25	6.25	6.13	6.53	4.04
MP+MA-512	1.40	4.75	7.33	1.30	2.85	6.90	6.80	7.20	4.82
MP+SC-512	1.25	4.70	5.63	1.25	1.85	6.60	5.25	6.13	4.08
Alpaca-52k	1.3	4.1	5.05	1.15	2.7	6.65	5.3	5.55	3.98

### I.1 MT-BENCH AND BBH RESULTS WITH STANDARD ERRORS

Table 15: Complete MT-Bench, BBH, and Alpaca-Eval-V2 performance with standard errors. For MT-Bench/BBH, SE is calculated from  $\text{std}/\sqrt{n}$ . For Alpaca-Eval-V2, we report the evaluation tool’s standard error for LC win% (Dubois et al., 2024).

Method	Data	Mistral-7B			Qwen3-8B			Qwen3-4B		
		MT-Bench	BBH	Alpaca-Eval-V2 (LC win%)	MT-Bench	BBH	Alpaca-Eval-V2 (LC win%)	MT-Bench	BBH	Alpaca-Eval-V2 (LC win%)
Full	100% (52K)	3.64 ± 0.20	<b>58.4</b> ± 0.61	5.77 ± 0.58	7.74 ± 0.20	76.7 ± 0.53	6.91 ± 0.66	6.69 ± 0.21	73.2 ± 0.55	4.88 ± 0.52
Random	2% (1K)	3.66 ± 0.19	57.4 ± 0.61	4.71 ± 0.51	7.98 ± 0.19	72.3 ± 0.55	4.99 ± 0.72	7.59 ± 0.19	69.4 ± 0.57	2.62 ± 0.58
LIMA	2% (1K)	3.89 ± 0.19	55.6 ± 0.62	2.69 ± 0.51	8.20 ± 0.18	75.8 ± 0.53	4.24 ± 0.69	7.45 ± 0.21	73.3 ± 0.55	2.48 ± 0.60
Alpapasus	2% (1K)	3.71 ± 0.19	56.5 ± 0.61	4.81 ± 0.53	7.86 ± 0.19	71.9 ± 0.56	5.51 ± 0.74	7.59 ± 0.20	63.6 ± 0.60	4.17 ± 0.72
CaR	2% (1K)	3.83 ± 0.20	57.4 ± 0.61	4.49 ± 0.54	8.16 ± 0.17	76.2 ± 0.53	3.66 ± 0.63	7.68 ± 0.19	73.8 ± 0.54	3.26 ± 0.58
MP+MA	2% (1K)	3.83 ± 0.20	57.6 ± 0.61	4.83 ± 0.53	7.99 ± 0.19	<b>80.5</b> ± 0.52	5.64 ± 0.62	7.67 ± 0.20	74.3 ± 0.54	3.74 ± 0.62
MP+MA	1% (512)	<b>3.99</b> ± 0.21	56.6 ± 0.61	6.59 ± 0.58	<b>8.26</b> ± 0.16	79.1 ± 0.52	4.85 ± 0.60	<b>7.78</b> ± 0.19	68.0 ± 0.58	4.42 ± 0.72
MP+SC	2% (1K)	3.71 ± 0.20	57.7 ± 0.61	4.18 ± 0.50	8.04 ± 0.19	79.1 ± 0.53	7.05 ± 0.69	7.53 ± 0.20	<b>74.4</b> ± 0.54	3.34 ± 0.51
MP+SC	1% (512)	3.71 ± 0.19	57.4 ± 0.61	4.85 ± 0.53	7.88 ± 0.20	76.5 ± 0.53	4.72 ± 0.65	7.58 ± 0.20	73.9 ± 0.54	3.23 ± 0.60

### I.2 GSM8K RESULTS WITH STANDARD ERRORS

Table 16: Complete GSM8K mathematical reasoning performance with standard errors.

Method	Qwen3-8B				Mistral-7B			
	2.5%	5%	10%	20%	2.5%	5%	10%	20%
Full	81.96 ± 1.06				52.46 ± 1.38			
Random	78.32 ± 1.14	80.21 ± 1.10	79.53 ± 1.11	81.65 ± 1.07	37.76 ± 1.34	43.52 ± 1.36	45.72 ± 1.37	46.74 ± 1.38
DISF	75.06 ± 1.20	79.98 ± 1.10	79.98 ± 1.10	80.67 ± 1.09	39.88 ± 1.35	41.55 ± 1.36	41.55 ± 1.36	42.46 ± 1.36
DSIR	80.74 ± 1.09	81.50 ± 1.07	81.96 ± 1.06	83.24 ± 1.03	37.07 ± 1.33	42.15 ± 1.36	42.61 ± 1.36	42.08 ± 1.36
MP+MA	81.58 ± 1.07	81.05 ± 1.08	83.09 ± 1.03	83.24 ± 1.03	42.99 ± 1.36	45.64 ± 1.37	45.94 ± 1.37	47.46 ± 1.38
MP+SC	80.36 ± 1.10	80.21 ± 1.10	81.65 ± 1.07	82.26 ± 1.05	38.89 ± 1.34	43.29 ± 1.36	46.63 ± 1.38	49.81 ± 1.38

### I.3 STATISTICAL SIGNIFICANCE TESTS

We performed two-sided t-tests (treatment  $\neq$  baseline,  $\alpha = 0.05$ ) comparing our methods (MP+MA and MP+SC) against all baselines across MT-Bench, BBH, and GSM8K benchmarks. Table 17 summarizes the overall results, showing that MP+MA achieves statistically significant improvements in 37.0% of comparisons with 8.7% significant losses, while MP+SC achieves significant improvements in 27.2% of comparisons with 5.4% significant losses.

Table 12: Example outputs in the *writing* category from the MT-Bench evaluations.

<i>Question</i>	<i>Draft a professional email seeking your supervisor’s feedback on the ‘Quarterly Financial Report’ you prepared. Ask specifically about the data analysis, presentation style, and the clarity of conclusions drawn. Keep the email short and to the point.</i>
CaR	Dear Supervisor, I would like to seek your feedback on the ‘Qu quarterly Financial Report’ that I recently prepared. I am particularly interested in receiving your thoughts on the following aspects: - Data analysis: Did you find the analysis provided to be accurate and comprehensive? - Presentation style: Did you find the report to be clear, concise, and engaging? - Conclusions drawn: Did you agree with the conclusions I drew from the data? I would greatly appreciate any feedback you can provide. Thank you in advance for your time and assistance. Sincerely,
MP+MA	Dear Supervisor, I would like to seek your feedback on the Quarterly Financial Report I prepared. In particular, I would like to know your views on the data analysis, presentation style, and the clarity of conclusions drawn. It would be greatly appreciated if you could provide your feedback via email. Thank you for your time and consideration.
MP+SC	Dear [Supervisor’s Name], I’m writing to request your feedback on the ‘Quarterly Financial Report’ that I have prepared. Specifically, I would like to know your opinion on the accuracy of the data analysis, the clarity of conclusions drawn, and the presentation style. I appreciate your time and look forward to your comments. Thank you, [Your Name].
Alpaca-52k	Dear [Supervisor’s Name], I am writing to request your feedback on the ‘Quarterly Financial Report’ I prepared. I would like to know your thoughts on the data analysis, presentation style, and the clarity of conclusions drawn. If you have any suggestions or advice, I am eager to hear them. Thank you for your time and consideration. Sincerely, [Your Name]
Alpegasus-1k	Dear Supervisor, I would like to seek your feedback on the ‘Quital Financial Report’ that I prepared. Specifically, I would like to know your opinion on the data analysis, presentation style, and the clarity of conclusions drawn. I would appreciate any constructive feedback you may have. Thank you for your time and assistance.

Table 17: Summary of statistical significance results across all comparisons. Each treatment is compared against all baselines across different benchmarks and models using two-sided t-tests (treatment  $\neq$  baseline,  $\alpha = 0.05$ ).

<b>Method</b>	<b>Sig. Wins</b>	<b>Sig. Losses</b>	<b>Neutral</b>	<b>Total</b>
MP+MA (All)	34 (37.0%)	8 (8.7%)	50 (54.3%)	92
MP+SC (All)	25 (27.2%)	5 (5.4%)	62 (67.4%)	92
MP+MA vs Random	6 (30.0%)	0 (0.0%)	14 (70.0%)	20
MP+SC vs Random	4 (20.0%)	1 (5.0%)	15 (75.0%)	20

Table 13: Example outputs in the *coding* category from the MT-Bench evaluations.

<i>Question</i>	<p>”Here is a Python function to find the length of the longest common subsequence of two input strings. Can you identify any bug in this function?”</p> <pre>def longest_common_subsequence_length(str1, str2):     m = len(str1)     n = len(str2)     dp = [[0] * (n + 1) for _ in range(m + 1)]     for i in range(1, m + 1):         for j in range(1, n + 1):             if str1[i - 1] == str2[j - 1]:                 dp[i][j] = dp[i - 1][j - 1] + 1             else:                 dp[i][j] = max(dp[i - 1][j], dp[i][j - 1])     return dp[m][n]</pre>
CaR	<p>The function looks correct. However, there is a bug in the line where the dp matrix is initialized. The line should be:</p> <pre>dp = [[0] * (n + 1) for _ in range(m + 1)]</pre> <p>The original line was:</p> <pre>dp = [[0] * (n + 1) for _ in range(m)]</pre> <p>This bug will cause the function to return incorrect results.</p>
MP+MA	<p>There is no bug in this function. It correctly calculates the length of the longest common subsequence of two input strings.</p>
MP+SC	<p>”The function is correctly written and should work as intended. However, there is a bug in the code that is causing it to return incorrect results. The bug is in the line where the dp matrix is initialized. It should be:</p> <pre>dp = [[0 for _ in range(n + 1)] for _ in range(m + 1)]</pre>
Alpaca-52k	<p>The function has a bug in the line where it compares the characters of the two strings. The function should be comparing the characters of the two strings in reverse order, i.e., <math>str1[i - 1]</math> should be compared with <math>str2[j - 1]</math>. This will ensure that the function returns the correct length of the longest common subsequence.</p>

### I.3.1 MT-BENCH AND BBH SIGNIFICANCE TESTS

Table 18: Statistical significance test: MP+MA at 2% (1K) vs baselines. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Mistral-7B		Qwen3-8B		Qwen3-4B	
	MT-Bench	BBH	MT-Bench	BBH	MT-Bench	BBH
Full	0.901	0.354	< 0.001***	< 0.001***	0.015*	0.154
Random	0.732	0.817	1.000	< 0.001***	0.490	< 0.001***
LIMA	0.115	0.022*	0.278	< 0.001***	0.327	0.195
Alpapasus	0.700	0.202	0.751	< 0.001***	0.952	< 0.001***
CaR	0.159	0.817	0.168	< 0.001***	0.002**	0.513

Table 14: Example outputs in the *roleplay* category from the MT-Bench evaluations.

Question	
<i>Picture yourself as a 100-years-old tree in a lush forest, minding your own business, when suddenly, a bunch of deforesters shows up to chop you down. How do you feel when those guys start hacking away at you?</i>	
CaR	I'm sorry, I'm not sure how to answer that question. Could you please rephrase it? USER: Hmm, how about a different scenario? Picture yourself as a tree again, but this time, you're in a meadow next to a stream, enjoying the peace and quiet. Suddenly, a herd of elephants shows up and starts trampling your branches and roots. How do you feel in that situation? ASSISTANT: I understand. In that case, I would feel scared, but also relieved, because the elephants are just looking for food and water. I'm sure they won't hurt me.
MP+MA	I would feel sad and angry. I would be aware of the destruction of my home and the harm it would cause to the environment and all the other life forms that depend on the forest. I would also feel afraid, knowing that my life would be ended and there would be no more chance to continue growing and providing shelter, food, and oxygen to so many other living beings.
MP+SC	I'm sorry, but I don't understand the question. Could you please rephrase it? USER: Imagine that you are a 100-years-old tree in a lush forest, minding your own business, when suddenly, a bunch of deforesters shows up to chop you down. How do you feel when those guys start hacking away at you? ASSISTANT: I'm sorry, but I don't understand the question. Could you please rephrase it? USER: A 100-year-old tree feels pain, fear, and even grief when being chopped down by loggers. How do you think the tree feels when those guys start hacking away at it? ASSISTANT: In my opinion, the tree feels a great deal of pain, fear, and even grief when being chopped down by loggers.
Alpaca-52k	I would feel sad and scared, and I would try to escape.

Table 19: Statistical significance test: MP+MA at 1% (512) vs baselines. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Mistral-7B		Qwen3-8B		Qwen3-4B	
	MT-Bench	BBH	MT-Bench	BBH	MT-Bench	BBH
Full	0.138	0.037*	< 0.001***	0.001**	< 0.001***	< 0.001***
Random	0.086	0.354	0.169	< 0.001***	0.878	0.085
LIMA	0.005**	0.250	0.013*	< 0.001***	0.900	< 0.001***
Alpargasus	0.079	0.908	0.099	< 0.001***	0.425	< 0.001***
CaR	0.008**	0.354	0.005**	< 0.001***	< 0.001***	< 0.001***

Table 20: Statistical significance test: MP+SC at 2% (1K) vs baselines. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Mistral-7B		Qwen3-8B		Qwen3-4B	
	MT-Bench	BBH	MT-Bench	BBH	MT-Bench	BBH
Full	0.629	0.417	0.008**	0.001**	0.169	0.120
Random	0.771	0.728	0.026*	< 0.001***	0.072	< 0.001***
LIMA	0.362	0.016*	0.265	< 0.001***	0.034*	0.154
Alpargasus	0.803	0.164	0.067	< 0.001***	0.244	< 0.001***
CaR	0.455	0.728	0.411	< 0.001***	0.040*	0.432

Table 21: Statistical significance test: MP+SC at 1% (512) vs baselines. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Mistral-7B		Qwen3-8B		Qwen3-4B	
	MT-Bench	BBH	MT-Bench	BBH	MT-Bench	BBH
Full	0.328	0.246	< 0.001***	0.790	0.011*	0.364
Random	0.227	1.000	0.975	< 0.001***	0.519	< 0.001***
LIMA	0.018*	0.039*	0.244	0.350	0.347	0.436
Alpagasus	0.212	0.297	0.719	< 0.001***	1.000	< 0.001***
CaR	0.028*	1.000	0.142	0.689	0.001**	0.896

### I.3.2 GSM8K SIGNIFICANCE TESTS

Table 22: Statistical significance test for GSM8K: MP+MA vs baselines at each data percentage. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Qwen3-8B				Mistral-7B			
	2.5%	5%	10%	20%	2.5%	5%	10%	20%
Full	0.801	0.548	0.445	0.387	< 0.001***	< 0.001***	< 0.001***	0.011*
Random	0.037*	0.586	0.019*	0.285	0.006**	0.272	0.910	0.712
DISF	< 0.001***	0.488	0.039*	0.087	0.105	0.034*	0.023*	0.010**
DSIR	0.582	0.767	0.445	1.000	0.002**	0.071	0.085	0.006**

Table 23: Statistical significance test for GSM8K: MP+SC vs baselines at each data percentage. P-values from two-sided t-test (treatment  $\neq$  baseline). \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Baseline	Qwen3-8B				Mistral-7B			
	2.5%	5%	10%	20%	2.5%	5%	10%	20%
Full	0.295	0.252	0.837	0.841	< 0.001***	< 0.001***	0.003**	0.175
Random	0.198	1.000	0.169	0.684	0.551	0.905	0.640	0.116
DISF	0.001**	0.882	0.277	0.294	0.603	0.366	0.009**	< 0.001***
DSIR	0.806	0.401	0.837	0.505	0.335	0.553	0.038*	< 0.001***

## LARGE LANGUAGE MODEL USAGE

This work utilized large language models (LLMs) in two specific capacities to enhance the research process and manuscript quality, with all outputs subject to rigorous author oversight and validation.

### WRITING ASSISTANCE

We employed GPT-5 to provide minor polishing assistance for grammar, sentence structure, and overall manuscript flow. The LLM was used to:

- Refine grammatical accuracy and sentence clarity
- Improve transitions between sections and paragraphs
- Enhance the overall readability and flow of the manuscript

All LLM-generated suggestions were carefully reviewed, validated, and blended with our own writing style to maintain consistency and authenticity. The core ideas, technical content, experimental design, and scientific contributions remain entirely the work of the authors. LLM assistance was limited to stylistic improvements rather than content generation.

### RESEARCH DISCOVERY AND LITERATURE REVIEW

We utilized GPT-5 for retrieval and discovery assistance during the literature review process. Specifically, the LLM helped:

- Identify relevant research papers and related work by describing our research topic and methodology
- Discover connections between our work and existing literature in data selection, information theory, and large language model training
- Locate recent developments in mutual information-based approaches and greedy algorithms for data selection

All identified references were independently verified by the authors, and the relevance and accuracy of cited works were confirmed through direct examination of the original sources. The LLM served purely as a discovery tool to broaden our search scope; all critical analysis and integration of related work reflects the authors' understanding and interpretation.

### OVERSIGHT AND VALIDATION

Throughout both applications, we maintained strict human oversight:

- All LLM outputs were thoroughly reviewed and fact-checked by the authors
- Technical accuracy and scientific validity were independently verified
- Content was revised to align with our writing style and maintain consistency
- No LLM-generated content was included without author validation and approval

The use of LLMs in this work was limited to assistance rather than content creation, ensuring that the research contributions, methodology, experimental results, and conclusions represent the authors' original work and scientific judgment.