

VOYAGER: A Training Free Approach for Generating Diverse Datasets using LLMs

Anonymous ACL submission

Abstract

Large language models (LLMs) are increasingly being used to generate synthetic datasets for the evaluation and training of downstream models. However, prior work has noted that such generated data lacks diversity. In this paper, we propose VOYAGER a novel principled approach to generate diverse datasets. Our approach is iterative and directly optimizes a mathematical quantity that optimizes the diversity of the dataset using the machinery of determinantal point processes. Furthermore, our approach is training-free, applicable to closed-source models, and scalable. In addition to providing theoretical justification for the working of our method, we also demonstrate through comprehensive experiments that VOYAGER improves diversity by $1.5 - 3x$ compared to popular baseline approaches.

1 Introduction

Large language models (LLMs) are widely used to generate synthetic data for scenarios where real world data is scarce. Although such data are valuable for training downstream models, post-trained LLMs used to generate such data often exhibit poor diversity in generation (Wright et al., 2025; Padmakumar and He, 2023). Common strategies to address this include sampling methods, such as temperature, top- p sampling and min- p sampling (Kool et al., 2019; Holtzman et al., 2020). However, these operate on the next token distribution and often fail to prevent mode collapse or semantic clustering, even at high temperatures (Jiang et al., 2025); as they lack global view of diversity between generations. Alternatively, “prompt based diversity control” instructs LLMs to cover specific topics. Although applicable to closed-weight models, this approach is not scalable and requires significant domain expertise. Finally, while diversity can be encoded into post-training rewards (Li et al., 2025; Chen et al., 2025), such methods are computation-

ally expensive and restricted to open-weight models. We introduce VOYAGER, a novel training-free framework for diverse synthetic data generation. Grounded in mathematical theory, our approach optimizes a geometric property intrinsic to dataset diversity. Crucially, VOYAGER operates without accessing model parameters, ensuring computational efficiency and compatibility with closed-weight (black-box) models.

VOYAGER is inspired by cartography, where explorers guided by a central command are encouraged to explore uncharted regions to maximize coverage. Analogously, our framework explores the data manifold, accepting generations that differ significantly from prior samples. To ensure computational efficiency, we maintain a fixed-size “anchor set” of representative regions rather than comparing against the full history. We dynamically prune this set to maximize its geometric volume, instructing the model to avoid clustered areas and expand diversity. VOYAGER formalizes this intuition using Determinantal Point Processes. Empirical evaluations show that VOYAGER consistently outperforms strong baselines across a suite of tasks.

To summarize, our main contributions are:

- **Principled approach to maximizing diversity:** Our approach leverages the machinery of determinantal point processes to maximize diversity by formulating it in terms of a geometric metric (volume). This provides theoretical justification to our approach.
- **Applicable to both open and closed weight models:** Our approach is training-free and requires no access to model weights.
- **Strong empirical performance:** We improve diversity by $1.5 - 3x$ over baselines.

2 Related Work

Existing diversity promotion methods fall into two categories: training-free or training-based.

2.1 Training-free Methods

These methods do not update model weights and use decoding strategies or prompt engineering.

Sampling and Decoding Strategies Some approaches manipulate output probabilities and decoding strategies via techniques like temperature, nucleus, min- p , or top- p (Kool et al., 2019; Holtzman et al., 2020; Amballa et al., 2025; Hewitt et al., 2022; Minh et al.; Chang et al., 2024), or utilize diversity-rewarding beam search (Vijayarumar et al., 2016). However, these methods require access to raw logits, making them unsuitable for closed-source models and are prone to incoherent or grammatically flawed outputs. In contrast, our method optimizes a global diversity proxy without requiring access to model weights or logits.

Prompt Diversity Control Prompt-based methods guide generation by explicitly instructing the model to vary dimensions such as topic or style (Wong et al., 2024; Samvelyan et al., 2024); or refine outputs across multiple turns (Mehrotra et al., 2024; Tian et al., 2024; Lu et al., 2025; Wong et al., 2024). While effective for open and closed models, these approaches rely on pre-defined and granular topics. Consequently, they optimize for topical diversity rather than semantic diversity; generations from distinct topics may still share significant vocabulary. Unlike these heuristics, our method directly maximizes the semantic volume of the data, ensuring diversity beyond topical differences.

2.2 Training-based Methods

Training-based approaches fine-tune model weights to maximize diversity-centric reward functions (Li et al., 2025; Yao et al., 2025; Chen et al., 2025; Ismayilzada et al., 2025; Lanchantin et al., 2025). Although effective, these methods incur significant computational overhead and are inapplicable to closed-source models.

Closely related to our work, Wang et al. (2024) propose using volume to measure dataset diversity for selection purposes. Similarly, Chen et al. (2025) integrate the determinant of the similarity matrix into post-training.¹ However, whereas these works focus on *measuring* diversity or *aligning* models,

¹While Chen et al. (2025) also utilize a determinant-based metric, their approach is a weight-tuning method restricted to open models and focuses on local response diversity. We focus on generating a globally diverse dataset via a training-free iterative algorithm.

we introduce a novel, iterative algorithm that leverages “text-based gradients” to actively *generate* a globally diverse dataset.

3 Background

VOYAGER relies on two key pieces of machinery: (a) Determinantal Point Processes (b) Prompt refinement using “textual” gradients.

Determinantal Point Processes: A Determinantal Point Process (DPP) (Kulesza et al., 2012) is a probabilistic model for subset selection where the probability of selecting a subset S is proportional to the determinant of the corresponding kernel matrix. For a ground set \mathcal{Y} and a positive semi-definite kernel matrix $K \in \mathbb{R}^{|\mathcal{Y}||\mathcal{Y}|}$, the probability of selecting subset $S \subseteq \mathcal{Y}$ is $\mathbb{P}(S \sim \text{DPP}(K)) \propto \det(K_S)$, where K_S is the submatrix of K indexed by elements in S . The determinant $\det(K_S)$ measures the *volume* spanned by the vectors corresponding to items in S , directly encoding diversity: higher determinants (volumes) correspond to more diverse subsets with less redundancy.

Prompt Refinement using Textual Gradients

This is a prompt refinement procedure (analogous to gradient descent, but operates purely in the text space) introduced by Pryzant et al. (2023). The objective here is to improve the prompt with respect to some reward metric. In the first of the two steps, an LLM-judge is asked to judge the prompt and its corresponding generation and suggest changes to the prompt so that the reward metric might be increased. In the second step, these suggestions are provided to another LLM is tasked to incorporate them into the prompt.

4 VOYAGER – Algorithm

VOYAGER mathematically operationalizes the notion of diversity of a set by the determinant of the corresponding similarity matrix of the set (encoded using a suitable kernel) and seeks to approximately maximize this measure iteratively.

4.1 Overview

At a high level, our method (see Algorithm 1, Figure 1) requires a task prompt p that describes the specific data generation task (eg, Generate a poem). The algorithm maintains a fixed-size set of explorers (\mathcal{E}) and a fixed-size set of anchor data-points denoted by Φ that capture representative/diverse regions of the full underlying dataset. In each iteration of the outer loop, we pick the next explorer

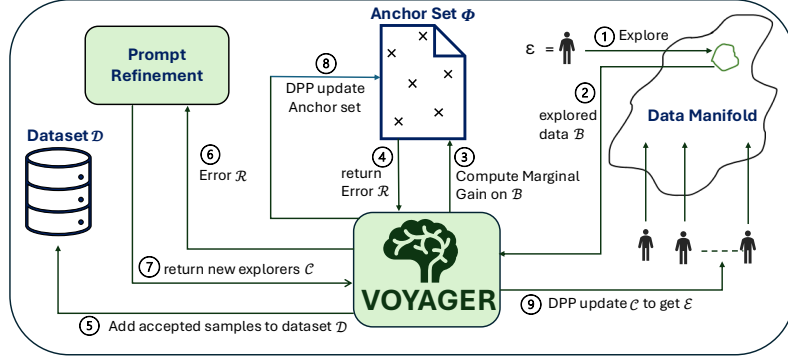


Figure 1: Overview of VOYAGER. We iteratively seek to explore new diverse regions of the data manifold via set of successive voyages carried out by explorers. Each explorer explores a certain region of the manifold (steps 1-2). Regions that are very similar to prior explorations are rejected by the central command which keeps track of a key set of salient regions explored, i.e., the anchor set (steps 3-5). New explorers are encouraged to explore areas different from prior explorations via prompt refinement (steps 6-9).

in the list and let the explorer perform a “voyage” (Line 7). This returns a new batch of data \mathbf{d} , an augmented set of potential anchors \mathbf{A} , and a successor set of candidate explorers \mathbf{s} . We add this new batch of data to the partially constructed dataset \mathcal{D} and also add the set of newly generated candidate explorers to \mathcal{C} (Lines 8-9). Because VOYAGER maintains a fixed-size set of anchor data-points to capture representative and diverse regions, and the newly added data-points could be potential anchor points, we update the anchor set Φ by sampling k diverse points from the augmented set \mathbf{A} (Line 10). Finally, after the current “beam of explorers” is done, we have a new set of candidate explorers \mathcal{C} . Once again, for computational efficiency, we select at most b explorers for the next iteration. Because we favor diverse explorers, we once again sample b explorers from \mathcal{C} from the underlying DPP (see Algorithm 4) and use that sampled set in the next iteration (Line 15).

4.1.1 Exploration Step

Our EXPLORE (see Algorithm 2) procedure consists of three main steps: First, given an explorer, we generate a batch of data by calling an LLM with the prompt characterizing that explorer.

We then score each data instance on its marginal gain in volume if it would be added to the anchor set. Once again, we rely on the notion of volume of the similarity matrix of a set of items as a measure of diversity (see Algorithm 3). If the marginal gain in volume is greater than a specified threshold τ , we add the instance to the dataset (\mathcal{S}) and the anchor set Φ (Line 1-7).

It is important to note that computing the

marginal gain of an instance over the entire dataset (which increases in size) is computationally expensive. We approximate this by computing marginal gain with respect to instances in the anchor set, which is of a fixed size. Next, if we have rejected instances (which we track in the set \mathcal{R}), we know that the explorer can be refined with regard to the diversity metric. We thus perform prompt refinement using textual gradients to obtain a new set of successor explorers \mathcal{C} (Lines 13-15).

4.2 Theoretical Justification

Having described the main algorithm in the previous section, we now draw on connections to the mathematical theory of determinantal point processes and matrix theory to justify our algorithm.

Our global objective is to construct a final similarity matrix S_T (that corresponds to the final dataset D_T) with a high effective rank. As noted by Friedman and Dieng (2022), who propose the Vendi Score as a diversity metric, a high effective rank of S_T implies high diversity.

Let S_T be an n by n square matrix. Let V the effective rank of S_T and D denote the determinant of S_T . Let C be the trace of S . We now state the following lemma

Lemma 1. *The effective rank of a square matrix S_T can be roughly approximated by its determinant D and the trace C , as (see Appendix 1.3 for proof)*

$$V \approx n^2 \frac{D^{1/n}}{C}, \text{ with } n \text{ being the rows of } S_T \quad (1)$$

Based on the above lemma, we seek to maximize $\det(S_T)$ while also noting that $\det(S_T)$ is the squared volume of D_T .

Algorithm 1 VOYAGER

Input: p : task prompt
 l : desired dataset size
 τ : Marginal gain threshold
 b : max number of explorers in each step
 k : maximum size of anchor point set
 T : Maximum number of iterations
 \mathbf{K}_{Sim} : Similarity Kernel

Output: \mathcal{D} : Constructed dataset

- 1: $\mathcal{D} = \{\}$ # Initialize dataset
- 2: $\Phi = \{\}$ # Initialize anchor set
- 3: $\mathcal{E} = \{p\}$ # Initialize set of explorers
- 4: **for** $i \leftarrow 0$ to T **do**
- 5: $\mathcal{C} \leftarrow \{\}$ # Successor set
- 6: **for all** $e \in \mathcal{E}$ **do**
- 7: $\mathbf{d}, \mathbf{A}, \mathbf{s} = \text{EXPLORE}(e, \Phi, \tau, \mathbf{K}_{\text{Sim}})$
- 8: $\mathcal{D} \leftarrow \mathcal{D} \cup \mathbf{d}$
- 9: $\mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{s}$
- 10: $\Phi \leftarrow \text{SAMPLEDPP}(\mathbf{A}, k, \mathbf{K}_{\text{Sim}})$
- 11: **if** $|\mathcal{D}| \geq l$ **then**
- 12: return \mathcal{D}
- 13: **end if**
- 14: **end for**
- 15: $\mathcal{E} \leftarrow \text{SAMPLEDPP}(\mathcal{C}, b, \mathbf{K}_{\text{Sim}})$
- 16: **end for**
- 17: **return** \mathcal{D}

239 However, since directly maximizing $\det(S_T)$
240 in an iterative manner is computationally infea-
241 sible, we reduce this problem to a proxy problem –
242 namely, the maximum volume submatrix problem
243 (MVS) problem. The MVS problem is to maxi-
244 mize the determinant of the best (k by k) principal
245 sub-matrix of S_T (selection of k data points).

$$246 \quad V_{\text{MVS}}(S_T) = \max_{\Phi \subset D_T, |\Phi|=k} \det(S_\Phi)^{(1/2)} \quad (2)$$

247 This is a good proxy because it has been observed
248 by [Goreinov and Tyrtshnikov \(2001\)](#); [Cortinovis](#)
249 [et al. \(2020\)](#) that the MVS is a quasi-best low-rank
250 approximation to the original matrix S_T . So we
251 attempt to maximize the MVS of size k of S_T as
252 we construct the dataset.

253 Unfortunately, even the above is NP-hard as we
254 need to check all $\binom{n}{k}$ subsets of the original S_T .
255 It is precisely here that we rely on the machinery
256 of determinantal point processes to find a high-
257 volume solution. While determinantal point pro-
258 cesses do not find the exact solution to the above
259 problem, a k -item sample drawn from the under-
260 lying determinantal point process will favor high-
261 volume subsets. We make use of this property

Algorithm 2 EXPLORE(\cdot)

Input: e : explorer
 Φ : Anchor Set
 τ : Marginal Gain Threshold
 \mathbf{K}_{Sim} : Similarity Kernel

Output: \mathcal{S} : Data instances
 Φ : Augmented set of potential anchors
 \mathcal{C} : Successor explorers

- 1: Prompt an LLM using e to generate \mathcal{B} , a batch
of instances (1 LLM call).
- 2: $\mathcal{C} \leftarrow \{\}$ # Successor set
- 3: $\mathcal{S} \leftarrow \{\}$ # Selected instances
- 4: $\mathcal{R} \leftarrow \{\}$ # Hold rejected instances
- 5: **for all** $w \in \mathcal{B}$ **do**
- 6: **if** $\text{MARGINALGAIN}(w, \mathbf{K}_{\text{Sim}}, \Phi) \geq \tau$ **then**
- 7: Add w to \mathcal{S} and to Φ
- 8: **else**
- 9: Add w to rejected set \mathcal{R}
- 10: **end if**
- 11: **end for**
- 12: **if** $|\mathcal{R}| > 0$ **then**
- 13: Prompt an LLM to Get “gradients”: $g =$
 $\{g_1, \dots, g_m\} = \text{LLM}_\nabla(e, \mathcal{R}, \Phi)$ (1 LLM
call)
- 14: Prompt an LLM to Apply gradients on e to
create a new set of explorers \mathcal{C} (1 LLM call)
- 15: **end if**
- 16: return $\mathcal{S}, \Phi, \mathcal{C}$

of determinantal point processes as follows: (a) 262
263 Define an anchor set Φ_T whose volume will be
264 $\text{VOL}(\Phi_T) = (\det S_{\Phi_T})^{1/2}$. VOYAGER iteratively
265 constructs this anchor set in a greedy manner, try-
266 ing to maximize its volume.

267 More specifically, note that when we generate a
268 batch (see Algorithm 2), we sequentially add an in-
269 stance to the dataset (and to the underlying anchor
270 set) only if the instance increases the volume of the
271 anchor set by a specific threshold. This step seeks
272 to make the anchor set a reservoir that is in some
273 sense “volume optimized” (holding diverse points).
274 However, at this point, we have an augmented an-
275 chor set A_T whose size exceeds the fixed size (k).
276 We thus prune it to size k but do it once again in
277 a manner seeking to maximize local volume. We
278 update Φ_T by drawing a k item sample from A_T
279 using a k -DPP which prefers a subset of size k with
280 high volume and will be representative of the MVS
281 volume of A_T and provides a “rough” lower bound
282 on $V_{\text{MVS}}(S_T)$. When the algorithm finishes, we
283 have a high-volume anchor set (by construction).

284 This in turn translates to a high $V_{MVS}(S_T)$ which
 285 in turn translates to a high effective rank of S_T and
 286 thus significantly more diverse D_T .

287 4.3 Computational Efficiency

288 Because VOYAGER restricts the size of the explor-
 289 ers and the anchor set in each iteration, our algo-
 290 rithm is quite computationally efficient in terms of
 291 CPU time complexity. The time complexity mainly
 292 depends on the maximum size of the anchor point
 293 set k_{max} , the maximum size of the candidate beam
 294 b_{max} , the batch size $|\mathcal{B}|$, and the maximum number
 295 of overall iterations T .

296 Observe that the MARGINALGAIN can be com-
 297 puted in $\mathcal{O}(k^2)$ time if the inverse of \mathbf{K}_{sim} can be
 298 pre-computed and cached. This means that we can
 299 process all instances $|\mathcal{B}|$ in $\mathcal{O}(|\mathcal{B}|k^2)$, assuming a
 300 constant time penalty for LLM calls and set addi-
 301 tion. The EXPLORE call incurs $\mathcal{O}(|\mathcal{B}|k^2)$ cost. The
 302 pruning of the anchor set to size (Line 10) using
 303 a DPP costs $\mathcal{O}(k_{max}^3)$ time. Thus, the entire inner
 304 for-loop (Line 6) incurs cost $\mathcal{O}(b(k_{max}^3 + |\mathcal{B}|k^2))$.
 305 The sampling step for pruning the set of explor-
 306 ers using a second DPP takes $\mathcal{O}(b_{max}^3)$ cost. The
 307 total cost incurred by the algorithm therefore is
 308 $\mathcal{O}(T(b(k_{max}^3 + |\mathcal{B}|k^2) + b_{max}^3))$.

309 Finally, note that VOYAGER in general is signif-
 310 icantly more efficient than a naive approach that
 311 could use a DPP to sample the entire dataset by
 312 first generating a universe larger than the required
 313 size and then selecting from it. That would incur
 314 a cost $\mathcal{O}(|l|^3)$ where $|l|$ is the size of the full
 315 dataset requested, which, when $|l|$ is large, can be
 316 prohibitive.

317 Having outlined the worst case cost of VOYAGER
 318 we now ask if we can roughly estimate how many
 319 iterations of the outer-loop would run in an aver-
 320 age case. This largely depends on the fraction of
 321 data-points accepted in EXPLORE step. Let ζ be
 322 the fraction of candidates accepted on an average in
 323 each call to EXPLORE. That means, $\zeta|\mathcal{B}|$ instances
 324 are on average accepted per iteration. So we should
 325 expect the outer loop to run about $T_{avg} = \frac{l}{\zeta|\mathcal{B}|}$
 326 times. Intuitively, ζ represents the “inherent easi-
 327 ness” of generating diverse data for the task using
 328 the LLM. If the LLM is easily able to generate di-
 329 verse data for the task and find prompts easily to
 330 obtain diverse data, ζ will be high (and low oth-
 331 erwise). Finally, we note that this also gives an
 332 average estimate of the total number of LLM calls
 333 $N_{LLM} = \frac{\rho l}{\zeta|\mathcal{B}|}$ as we make a constant number of
 334 LLM calls, ρ (we have $\rho = 3$) in EXPLORE.

5 Experiments 335

In this section, we report on comprehensive ex- 336
 periments to evaluate our proposed method on its 337
 effectiveness to generate diverse datasets. 338

5.1 Generation Tasks and Evaluation Metrics 339

Generation Tasks We consider the two differ- 340
 ent categories of text generation tasks: (a) creative 341
 writing and (b) reasoning to evaluate our method 342
 against strong baselines. In the creative writing 343
 category, we consider four generation tasks: (a) 344
 topical sentence generation, (b) conversation gener- 345
 ation (c) poem generation, and (d) movie plot 346
 generation. In the reasoning category, we consider 347
 the tasks of (a) grade school mathematical question 348
 generation and (b) Logical reasoning questions. 349

Diversity Metrics Because there is no universal 350
 consensus on metrics for evaluating diversity, we 351
 consider a few popular metrics for diversity captur- 352
 ing different facets. 353

- **Cosine Distance:** We consider the mean cosine 354
 distance between pairs of instances as a 355
 measure of diversity. 356
- **Lexical Distance:** While the cosine distance 357
 metric captures semantic diversity, it may not 358
 capture lexical diversity very well. Therefore, 359
 we also consider a lexical diversity metric. To 360
 compute lexical diversity, we use the Jaccard 361
 distance and report the mean Jaccard distance 362
 (removing stop words) over pairs of instances. 363
- **Vendi Score:** Vendi score (Friedman and Di- 364
 eng, 2022) is a popular diversity metric that 365
 seeks to capture diversity of the dataset in a 366
 broader and general sense (potentially includ- 367
 ing all the above dimensions). It seeks to mea- 368
 sure the effective number of data instances 369
 that can represent the data. 370

Quality Metric We also seek to capture the im- 371
 pact on quality of the generations as we seek to 372
 encourage diversity. Consequently, we also mea- 373
 sure the quality of the generations using an LLM- 374
 as-Judge framework where the judge evaluates the 375
 generations on a task-specific rubric (eg, dimen- 376
 sions like faithfulness, coherence, etc). 377

LLM Calls We also track the number of LLM 378
 calls each method to generate the same size dataset 379
 (assuming all methods call the LLM to generate 380
 the same batch size of data). 381

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.67 \pm 0.17	0.21 \pm 0.09	2.99	21.70 \pm 0.75	50
TEMP	0.70 \pm 0.15	0.22 \pm 0.09	3.23	21.71 \pm 0.75	50
DIVERSE	0.78 \pm 0.13	0.41 \pm 0.17	7.87	23.10 \pm 0.82	50
HISTORY	0.79 \pm 0.12	0.33 \pm 0.12	6.00	22.15 \pm 0.70	50
HIERARCHICAL	<u>0.85 \pm 0.10</u>	<u>0.54 \pm 0.14</u>	<u>15.07</u>	<u>22.85 \pm 0.98</u>	550
SUBSETSELECT	0.81 \pm 0.11	0.41 \pm 0.17	7.77	22.22 \pm 0.93	500
VOYAGER	0.87 \pm 0.06	0.55 \pm 0.13	24.13	22.26 \pm 2.06	443

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.72 \pm 0.05	0.25 \pm 0.08	4.59	22.97 \pm 0.68	50
TEMP	0.74 \pm 0.05	0.27 \pm 0.09	5.16	23.03 \pm 0.67	50
DIVERSE	0.73 \pm 0.04	0.16 \pm 0.04	3.30	<u>23.75 \pm 0.88</u>	50
HISTORY	0.70 \pm 0.07	0.26 \pm 0.11	4.24	22.97 \pm 0.64	50
HIERARCHICAL	<u>0.77 \pm 0.05</u>	<u>0.38 \pm 0.10</u>	<u>8.45</u>	<u>23.32 \pm 0.62</u>	550
SUBSETSELECT	0.73 \pm 0.05	0.26 \pm 0.08	4.85	23.13 \pm 0.66	500
VOYAGER	0.82 \pm 0.05	0.44 \pm 0.09	15.04	23.92 \pm 0.88	426

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.76 \pm 0.05	0.15 \pm 0.05	3.00	22.52 \pm 0.89	50
TEMP	0.78 \pm 0.05	0.16 \pm 0.04	3.22	22.65 \pm 0.82	50
DIVERSE	0.78 \pm 0.04	0.14 \pm 0.04	2.76	22.79 \pm 1.46	50
HISTORY	0.71 \pm 0.05	0.11 \pm 0.03	2.30	22.45 \pm 0.92	50
HIERARCHICAL	<u>0.82 \pm 0.06</u>	<u>0.30 \pm 0.08</u>	<u>5.68</u>	22.56 \pm 1.31	550
SUBSETSELECT	0.76 \pm 0.05	0.16 \pm 0.04	3.08	22.52 \pm 0.88	500
VOYAGER	0.86 \pm 0.05	0.30 \pm 0.08	7.31	24.51 \pm 0.91	615

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.80 \pm 0.05	0.21 \pm 0.08	4.00	23.04 \pm 0.53	50
TEMP	0.83 \pm 0.05	0.26 \pm 0.08	5.52	23.08 \pm 0.60	50
DIVERSE	0.81 \pm 0.04	0.25 \pm 0.07	5.28	23.46 \pm 1.99	50
HISTORY	0.78 \pm 0.05	0.22 \pm 0.07	4.28	<u>23.15 \pm 0.78</u>	50
HIERARCHICAL	<u>0.84 \pm 0.05</u>	<u>0.32 \pm 0.08</u>	<u>7.66</u>	22.99 \pm 0.92	550
SUBSETSELECT	0.81 \pm 0.05	0.23 \pm 0.08	4.57	23.03 \pm 0.54	500
VOYAGER	0.84 \pm 0.05	0.34 \pm 0.10	8.30	22.96 \pm 1.12	695

Table 1: Creative task evaluations. Note that **VOYAGER outperforms all** baselines significantly (average Vendi score improvement of 296% over DEFAULT, 43 % over HIERARCHICAL across all creative tasks) with no significant degradation on perceived quality. **Legend** Best result in each column is highlighted in **Bold** and the second-best result in each column is highlighted in underline. Quality metrics are on a scale of 0(lowest) – 25(highest).

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.54 \pm 0.13	0.20 \pm 0.08	3.04	14.99 \pm 0.12	50
TEMP	0.56 \pm 0.12	0.22 \pm 0.08	3.56	15.00 \pm 0.10	50
DIVERSE	0.47 \pm 0.06	0.07 \pm 0.02	1.65	15.00 \pm 0.00	50
HISTORY	0.30 \pm 0.15	0.24 \pm 0.12	3.13	14.83 \pm 0.58	50
HIERARCHICAL	<u>0.68 \pm 0.10</u>	<u>0.40 \pm 0.11</u>	<u>8.72</u>	<u>15.00 \pm 0.06</u>	550
SUBSETSELECT	0.57 \pm 0.12	0.22 \pm 0.07	3.48	14.98 \pm 0.17	500
VOYAGER	0.81 \pm 0.06	0.48 \pm 0.10	18.78	14.77 \pm 0.92	399

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
DEFAULT	0.59 \pm 0.10	0.20 \pm 0.09	3.31	14.83 \pm 0.69	50
TEMP	0.63 \pm 0.10	0.25 \pm 0.11	4.47	14.80 \pm 0.90	50
DIVERSE	0.62 \pm 0.08	0.15 \pm 0.05	2.83	<u>14.92 \pm 0.44</u>	50
HISTORY	0.35 \pm 0.11	0.27 \pm 0.11	3.98	14.99 \pm 0.13	50
HIERARCHICAL	<u>0.65 \pm 0.08</u>	<u>0.33 \pm 0.11</u>	<u>7.02</u>	14.88 \pm 0.55	550
SUBSETSELECT	0.62 \pm 0.10	0.24 \pm 0.10	4.25	14.80 \pm 0.72	500
VOYAGER	0.79 \pm 0.06	0.41 \pm 0.08	13.26	14.64 \pm 0.97	393

Table 2: Reasoning tasks evaluations. Note that **VOYAGER outperforms all** baselines significantly (average Vendi score improvement of 412% over DEFAULT, 102 % over HIERARCHICAL across all reasoning tasks) with no significant degradation on perceived quality. **Legend** Best result in each column is highlighted in **Bold** and the second-best result in each column is highlighted in underline. Quality metrics are on a scale of 0(lowest) – 15(highest).

5.2 Baseline Methods

We consider the following baselines:

- **Vanilla Generation - DEFAULT:** This simply prompts the LLM (with default parameters) to perform the generation task.
- **Temperature based sampling - TEMP:** One lever to encourage diversity in LLM generations is the temperature parameter used to scale the LLM output logits when sampling tokens. We set higher temperature to 2.0.
- **Generate “diverse” command - DIVERSE:** This explicitly ask the LLM to be "diverse" in its output by appending it to the user task instructions appropriately.
- **Conditioned on history – HISTORY:** Another approach to encourage diversity is to provide a history of prior generated data instances (of a fixed window) and ask it to avoid generating such instances.
- **Hierarchical Prompting - HIERARCHICAL:** Here we ask the LLM to first generate K diverse topics and then actual instances conditioned on those topics to encourage diversity.

- **Subset Select - SUBSETSELECT:** We evaluate whether just using a k -DPP to sample the dataset from a larger set (universe) generated using Vanilla sampling would be effective.

5.3 Experimental Settings

Similarity Kernel Functions Our similarity kernel function is a convex combination of a radial basis function kernel using embeddings of the text, and a lexical similarity kernel using Jaccard similarity. The weights of the convex combination are 0.7 and 0.3 for the radial basis function kernel and the lexical similarity kernel, respectively.

LLMs All experimental data were generated using GPT-4o mini. For text embeddings, we utilize OpenAI’s text-embedding-3-small model.

HIERARCHICAL baseline settings For the hierarchical baseline, in each call, we generate 10 subtopics to ensure comprehensive coverage of the problem space. We then generate 1 sample for each such topic to return 10 instances.

VOYAGER Hyperparameters Our experiment settings are as follows:**Number of explorers (b):**

Method	Lexical \uparrow	Cosine \uparrow	Vendi \uparrow	Quality \uparrow	LLM \downarrow
VOYAGER-RE	0.85 \pm 0.07	0.44 \pm 0.14	11.85	21.47 \pm 1.88	361
VOYAGER	0.85 \pm 0.07	0.47 \pm 0.15	14.28	21.88 \pm 1.51	252

Table 3: Ablation for DPP used in selecting b explorers on the sports task. Observe that choosing a random sample of explorers results in slightly lower diversity and also incurs more LLM calls as opposed to choosing a diverse set of explorers confirming our hypothesis that diverse explorers can make search through the space more effective.

Method	$ D_{train} $	Gemma 2B-IT	Gemma 7B-IT
DEFAULT	1000	13.1	35.7
VOYAGER	1000	16.4	45.7
VOYAGER	500	12.9	42.8

Table 4: GSM8K zero-shot test accuracy on Gemma trained on data generated using VOYAGER compared to the DEFAULT baseline. Note that VOYAGER which generates more diverse data leads to improved predictive performance and training data efficiency compared to the baseline.

3, **Number of anchor points (k):** 10, **Maximum iterations (T):** 200, **Samples per LLM call ($|\mathcal{B}|$):** 10, **Target dataset size (l):** 500. For initializing τ , we followed the initialization procedure in Appendix 1.2. See Appendix 1.6 for the meta-prompts used in the refinement procedure.

Evaluation Jury All generated outputs were evaluated by a panel consisting of GPT-4, GPT-4o, and GPT-4.1, providing a robust quality assessment framework. We use the mean function to obtain jury consensus scores for each dimension in the rubric. The prompts used for jury evaluation are available in the Appendix 1.7.

5.4 Results and Analysis

Creating Writing Tasks Tables 1 show the results for the creative writing tasks, from which we can make the following observations.

First, note that temperature based methods improve over the baseline in terms of all diversity metrics with no significant decrease in quality.

Second, with regards to the “prompt-diversity” control method, the DIVERSE-KEYWORD baseline significantly improves diversity (across all metrics) over both the default baseline and temperature based baselines, suggesting that explicit instructions to make the output diverse is helpful.

Third, the history based prompting outperforms the baseline marginally but not as strongly as the explicit DIVERSE KEYWORD approach. This may be because the HISTORY based approach seeks to ensure non-redundancy with the prior history (of a fixed size) but that does not necessarily yield



Figure 2: Rejection rate of samples within a batch over time with “textual gradients” enabled vs disabled to generate the same dataset size for the sports task (all other settings identical).

diverse outputs overall. One advantage of these baselines is that to construct a dataset of size l with a batch size $|\mathcal{B}|$ they will make exactly $\frac{l}{|\mathcal{B}|}$ LLM calls irrespective of the task. However, this efficiency in terms of LLM calls comes at a significant cost (significant loss of diversity). VOYAGER in contrast enforces quality control to obtain diversity and thus generally pays a cost in terms of more LLM calls.

Fourth, the hierarchical prompting approach very significantly outperforms other methods in this class. This suggests that incorporating domain knowledge and explicit instructions to explore different topics significantly improves diversity.

Finally, all flavors of VOYAGER significantly outperform all other methods, significantly suggesting the effectiveness of VOYAGER as we directly optimize a quantitative measure of diversity (volume) with no significant degradation in quality. It is also worth noting that VOYAGER also makes significantly fewer LLM calls (in most cases) compared to HIERARCHY, the best competing baseline, and suggests that VOYAGER has a generally better cost-benefit tradeoff (LLM calls vs diversity).

Reasoning Tasks Tables 2 show the results of the experiments for the reasoning tasks. Many of the observations noted in the creative writing experiments hold for the reasoning tasks as well. Most notably, VOYAGER significantly outperforms all baselines on the reasoning tasks as well.

Human Evaluations & Qualitative results Table 5 shows the human-rated diversity (3 annotators) of VOYAGER and DEFAULT baseline for Sports and Math tasks (See Appendix 1.4). Note

Task	DEFAULT	VOYAGER
Sports	2.16	3.82
Math	1.56	3.72

Table 5: Human-rated diversity. Rubrics: 1 = Very Similar, 2=Similar, 3 = Dissimilar, 4 = Very Dissimilar

that VOYAGER achieves a higher mean diversity score compared to the DEFAULT baseline, indicating the correlation with automatic diversity metrics.

See Appendix 1.5 for qualitative results on Sports and Math tasks.

5.5 Ablation Studies

Effect of Diverse Explorers To evaluate the impact of encouraging diversity in the explorers by sampling from a DPP, we replace the selection of successors to just be a random sample of the same size since it is not critical to our algorithm. We call this algorithm variant VOYAGER-RE(RandomExplorer). The main advantage of VOYAGER-RE is that we do not need to incur the cost of sampling from a k -DPP, which ($|\mathcal{C}|^3$) in each main iteration of the loop. This potentially trades off some diversity and efficient convergence for reduced local computational time (a choice that practitioners have), which we validate empirically in Table 3. Note that VOYAGER-RE shows a drop in diversity, quality and requires more LLM calls to generate the same number of instances as VOYAGER confirming our observation above.

Effect of Textual Gradients We also ablate the prompt refinement step using textual gradients to evaluate its effectiveness. In particular, while we still attempt to maximize our operationalize notion of diversity by selecting only instances which a marginal gain above a threshold, we disable prompt refinement. This means no new explorers will be generated and we will only use the same explorer many times (with variance in output only due to the LLM call). We compare this with a setting (where prompt refinement is enabled) but restrict the beam size to exactly be 1, so that we have an “apples-to-apples” comparison with the setting where the same explorer is used across all iterations. Our hypothesis is that disabling the prompt refinement would result in a significantly higher average rejection rate (the fraction of instances rejected because they did not meet the marginal gain criteria in EXPLORE step), and the algorithm would have to run for many more iterations. Figure 2 shows the results of this ablation. Observe that disabling the

textual gradients results in a much higher average rejection rate and number of iterations confirming our hypothesis that allowing for feedback to refine the prompt can help make the search more efficient.

5.6 Effectiveness of Data for Synthetic Training Data Generation

Finally, we also evaluate whether a model trained on data using VOYAGER would result in higher model performance compared to a model trained on synthetic data generated by the DEFAULT baseline.

To do this, we consider the GSM8K task. We generate 1000 questions using the prompt “Generate a mathematical question for Grade School Math”. We use 5-shot prompting for GPT4 (OpenAI et al., 2024) to answer these questions to obtain (question, answer) pairs for the training data. We then train Gemma-2b-it and Gemma-7b-it models (instruction-tune) using just these 1000 examples. We compare against models trained on the same number of examples but generated using a temperature baseline. We evaluate both settings on the standard GSM8K test set (Cobbe et al., 2021). Results are shown in Table 4. We note that data generated using VOYAGER yields a significantly higher model performance than the baseline, suggesting the effectiveness of our method and more generally the effectiveness of training on diverse data (35.7 vs 45.7 – see column Gemma-7b-it). In fact, we noted that using only 500 examples, we were able to almost match (and even outperform) the performance of models trained on data generated by the DEFAULT baseline, further underscoring the importance of diversity in training data, an observation noted by several prior works (Bukharin et al., 2024; Pang et al., 2024)

6 Conclusion

We proposed VOYAGER, a training free approach to use LLMs to generate more diverse data. Our method relies on the determinant of the similarity matrix of a dataset is related to the volume spanned by the data to measure data diversity. Building on this observation, we propose an iterative algorithm that seeks to approximately maximize this measure of diversity by leveraging the machinery of determinantal point processes and prompt refinement using textual gradients. To conclude, our method significantly improves the diversity of generated data by $1.5 - 3x$, is training-free and scalable, compared to baseline methods.

586 Limitations

587 Our work is not without its limitations. First, VOY-
588 AGER depends on high-quality instruction follow-
589 ing LLMs. This is necessary during sample gen-
590 eration, but more importantly, during the "textual
591 gradient extraction" phase, where the LLM must
592 analyze a prompt, its generation, and then provide
593 suggestions for improvement. This also requires
594 the LLM to have a certain level of reasoning ca-
595 pability. Secondly, the similarity computations re-
596 quire a robust embedding model that can effectively
597 differentiate differences in text within the embed-
598 ding space. Finally, this work focuses solely on
599 diversity in text generation using LLMs. We do
600 not consider diversity of multi-modal data, which
601 poses additional challenges, such as how to mea-
602 sure similarity of representations across multiple
603 modalities.

604 References

605 Avinash Amballa, Aditya Parashar, Aditya Vikram
606 Singh, Jinlin Lai, and Benjamin Rozenoyer. 2025.
607 [Quasi-random multi-sample inference for large lan-
608 guage models](#). In *Frontiers in Probabilistic Inference:
609 Learning meets Sampling*.

610 Alexander Bukharin, Shiyang Li, Zhengyang Wang,
611 Jingfeng Yang, Bing Yin, Xian Li, Chao Zhang, Tuo
612 Zhao, and Haoming Jiang. 2024. Data diversity mat-
613 ters for robust instruction tuning. In *Findings of the
614 Association for Computational Linguistics: EMNLP
615 2024*, pages 3411–3425.

616 Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil
617 Ramakrishna, and Tagyoung Chung. 2024. Real
618 sampling: Boosting factuality and diversity of open-
619 ended generation via asymptotic entropy. *arXiv
620 preprint arXiv:2406.07735*.

621 Yilei Chen, Souradip Chakraborty, Lorenz Wolf, Yannis
622 Paschalidis, and Aldo Pacchiano. 2025. Post-training
623 large language models for diverse high-quality re-
624 sponses. *arXiv preprint arXiv:2509.04784*.

625 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,
626 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias
627 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro
628 Nakano, Christopher Hesse, and John Schulman.
629 2021. [Training verifiers to solve math word prob-
630 lems](#). *Preprint*, arXiv:2110.14168.

631 Alice Cortinovis, Daniel Kressner, and Stefano Massei.
632 2020. On maximum volume submatrices and cross
633 approximation for symmetric semidefinite and diag-
634 onally dominant matrices. *Linear Algebra and its
635 Applications*, 593:251–268.

Dan Friedman and Adji Bousso Dieng. 2022. The vendi
score: A diversity evaluation metric for machine
learning. *arXiv preprint arXiv:2210.02410*. 636
637
638

Sergei A Goreinov and Eugene E Tyrtysnikov. 2001. 639
The maximal-volume concept in approximation by
low-rank matrices. *Contemporary Mathematics*,
280:47–52. 640
641
642

John Hewitt, Christopher D Manning, and Percy Liang.
2022. Truncation sampling as language model
desmoothing. In *Findings of the Association for Com-
putational Linguistics: EMNLP 2022*, pages 3414–
3427. 643
644
645
646

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and
Yejin Choi. 2020. [The curious case of neural text
degeneration](#). *Preprint*, arXiv:1904.09751. 648
649
650

Mete Ismayilzada, Antonio Laverghetta Jr, Simone A
Luchini, Reet Patel, Antoine Bosselut, Lonneke
van der Plas, and Roger Beaty. 2025. Creative prefer-
ence optimization. *arXiv preprint arXiv:2505.14442*. 651
652
653
654

Liwei Jiang, Yuanjun Chai, Margaret Li, Mickel Liu,
Raymond Fok, Nouha Dziri, Yulia Tsvetkov, Maarten
Sap, Alon Albalak, and Yejin Choi. 2025. [Artificial
hivemind: The open-ended homogeneity of language
models \(and beyond\)](#). *Preprint*, arXiv:2510.22954. 655
656
657
658
659

Wouter Kool, Herke Van Hoof, and Max Welling. 2019. 660
[Stochastic beams and where to find them: The
Gumbel-top-k trick for sampling sequences without
replacement](#). In *Proceedings of the 36th Interna-
tional Conference on Machine Learning*, volume 97
of *Proceedings of Machine Learning Research*, pages
3499–3508. PMLR. 661
662
663
664
665
666

Alex Kulesza, Ben Taskar, and 1 others. 2012. Determi-
nantal point processes for machine learning. *Founda-
tions and Trends® in Machine Learning*, 5(2–3):123–
286. 667
668
669
670

Jack Lanchantin, Angelica Chen, Shehzaad Dhuliawala,
Ping Yu, Jason Weston, Sainbayar Sukhbaatar, and
Ilya Kulikov. 2025. Diverse preference optimization.
arXiv preprint arXiv:2501.18101. 671
672
673
674

Tianjian Li, Yiming Zhang, Ping Yu, Swarnadeep Saha,
Daniel Khashabi, Jason Weston, Jack Lanchantin,
and Tianlu Wang. 2025. Jointly reinforcing diversity
and quality in language model generations. *arXiv
preprint arXiv:2509.02534*. 675
676
677
678
679

Yining Lu, Dixuan Wang, Tianjian Li, Dongwei
Jiang, Sanjeev Khudanpur, Meng Jiang, and Daniel
Khashabi. 2025. Benchmarking language model cre-
ativity: A case study on code generation. In *Proceed-
ings of the 2025 Conference of the Nations of the
Americas Chapter of the Association for Computa-
tional Linguistics: Human Language Technologies
(Volume 1: Long Papers)*, pages 2776–2794. 680
681
682
683
684
685
686
687

Pronita Mehrotra, Aishni Parab, and Sumit Gulwani.
2024. Enhancing creativity in large language mod-
els through associative thinking strategies. *arXiv
preprint arXiv:2405.06715*. 688
689
690
691

692	Nguyen Nhat Minh, Andrew Baker, Clement Neo,	Dustin Wright, Sarah Masud, Jared Moore, Srishti	745
693	Allen G Roush, Andreas Kirsch, and Ravid Shwartz-	Yadav, Maria Antoniak, Chan Young Park, and Is-	746
694	Ziv. Turning up the heat: Min-p sampling for creative	abelle Augenstein. 2025. Epistemic diversity and	747
695	and coherent llm outputs. In <i>The Thirteenth Interna-</i>	knowledge collapse in large language models. <i>arXiv</i>	748
696	<i>tional Conference on Learning Representations.</i>	<i>preprint arXiv:2510.04226.</i>	749
697	OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal,	Jian Yao, Ran Cheng, Xingyu Wu, Jibin Wu, and	750
698	Lama Ahmad, and others. 2024. Gpt-4 technical	Kay Chen Tan. 2025. Diversity-aware policy op-	751
699	report . <i>Preprint</i> , arXiv:2303.08774.	timization for large language model reasoning. <i>arXiv</i>	752
700	Vishakh Padmakumar and He He. 2023. Does writ-	<i>preprint arXiv:2505.23433.</i>	753
701	ing with language models reduce content diversity?		
702	<i>arXiv preprint arXiv:2309.05196.</i>		
703	Jinlong Pang, Jiaheng Wei, Ankit Parag Shah, Zhaowei		
704	Zhu, Yaxuan Wang, Chen Qian, Yang Liu, Yujia Bao,		
705	and Wei Wei. 2024. Improving data efficiency via		
706	curating llm-driven rating systems. <i>arXiv preprint</i>		
707	<i>arXiv:2410.10877.</i>		
708	Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chen-		
709	guang Zhu, and Michael Zeng. 2023. Automatic		
710	prompt optimization with "gradient descent" and		
711	beam search . <i>Preprint</i> , arXiv:2305.03495.		
712	Olivier Roy and Martin Vetterli. 2007. The effective		
713	rank: A measure of effective dimensionality. In <i>2007</i>		
714	<i>15th European Signal Processing Conference</i> , pages		
715	606–610.		
716	Mikayel Samvelyan, Sharath C Raparthy, Andrei Lupu,		
717	Eric Hambro, Aram H Markosyan, Manish Bhatt,		
718	Yuning Mao, Minqi Jiang, Jack Parker-Holder, Jakob		
719	Foerster, and 1 others. 2024. Rainbow team-		
720	ing: Open-ended generation of diverse adversarial		
721	prompts. <i>Advances in Neural Information Process-</i>		
722	<i>ing Systems</i> , 37:69747–69786.		
723	Yufei Tian, Abhilasha Ravichander, Lianhui Qin, Ronan		
724	Le Bras, Raja Marjeh, Nanyun Peng, Yejin Choi,		
725	Thomas L Griffiths, and Faeze Brahman. 2024. Mac-		
726	gyver: Are large language models creative problem		
727	solvers? In <i>Proceedings of the 2024 Conference of</i>		
728	<i>the North American Chapter of the Association for</i>		
729	<i>Computational Linguistics: Human Language Tech-</i>		
730	<i>nologies (Volume 1: Long Papers)</i> , pages 5303–5324.		
731	Ashwin K Vijayakumar, Michael Cogswell, Ram-		
732	prasath R Selvaraju, Qing Sun, Stefan Lee, David		
733	Crandall, and Dhruv Batra. 2016. Diverse beam		
734	search: Decoding diverse solutions from neural se-		
735	quence models. <i>arXiv preprint arXiv:1610.02424.</i>		
736	Peiqi Wang, Yikang Shen, Zhen Guo, Matthew Stal-		
737	lone, Yoon Kim, Polina Golland, and Rameswar		
738	Panda. 2024. Diversity measurement and subset se-		
739	lection for instruction tuning datasets. <i>arXiv preprint</i>		
740	<i>arXiv:2402.02318.</i>		
741	Justin Wong, Yury Orlovskiy, Michael Luo, Sanjit A		
742	Seshia, and Joseph E Gonzalez. 2024. Simplestrat:		
743	Diversifying language model generation with stratifi-		
744	cation. <i>arXiv preprint arXiv:2410.09038.</i>		

A Appendix

1.1 VOYAGER- Algorithm

Algorithm 3 COMPUTEMARGINALGAIN(\cdot)

Input: w : Item whose marginal gain needs to be computed
 \mathbf{K} : Similarity kernel function to use.
 A : Set of items already present
Output: γ : Marginal gain in volume if w is added to the set.
 $\mathbf{S} = \mathbf{K}(A)$ # Construct Similarity kernel
 $\mathbf{S}' = \mathbf{K}(A \cup w)$
3: $\gamma = \frac{\det(\mathbf{S}')}{\det(\mathbf{S})}$
return γ

Algorithm 4 SAMPLEDPP(\cdot)

Input: U : Set of items
 n : number of items to sample
 \mathbf{K} : Similarity kernel function to use.
Output: \mathbf{S} : A subset of n items sampled from underlying DPP.
 $\mathbf{L} = \mathbf{K}(U)$ # Construct likelihood kernel
 $\mathbf{S} = \text{Sample K-DPP}(\mathbf{L}, n)$ (Kulesza et al., 2012)
3: return \mathbf{S}

1.2 Recommendation for setting τ

Algorithm 5 Recommended Initialization of τ

Input: M : LLM
 p : task prompt
 \mathbf{K} : similarity kernel
Output: τ_0 : Initial threshold
Sample $\mathbf{Y} = y_1, y_2, \dots, y_{100} \sim M(p)$
 $\mathbf{S} \leftarrow \text{SAMPLEDPP}(\mathbf{Y}, n = 10, \mathbf{K})$
3: $\mathbf{K}_S = \mathbf{K}(\mathbf{S})$ # Construct Similarity kernel
 $\tau_0 = \text{CLIP}(\alpha \cdot \det(\mathbf{K}_S), \tau_{min}, \tau_{max})$
return τ_0

As an expedient choice we recommend decaying $\tau = \tau_0 \cdot \exp(-\frac{i}{T})$ for faster convergence. Note that we did not invest effort in finetuning the parameters like τ of our method.

1.3 Proof of Lemma 1

Proof. To recap, S_T is a square n by n similarity matrix, V denotes the effective rank of S_T , and D denote the determinant of S_T . Let C be the trace

of S_T . Let $\lambda_1, \lambda_2, \dots, \lambda_n$ be the eigenvalues of S_T .

We use Taylor approximations of two mathematical quantities² to derive an approximation between the two quantities. We first define some quantities and notation. From the definition of effective rank (Roy and Vetterli, 2007), the effective rank of S_T ,

$$V \triangleq e^H, \text{ where } H = - \sum \frac{\lambda_i}{C} \ln \left(\frac{\lambda_i}{C} \right). \quad (3)$$

For ease of writing, let $p_i \triangleq \frac{\lambda_i}{C}$. p_i is just the normalized eigenvalue so that the set of p_i forms a probability distribution. Let $G_p = \sqrt[n]{\prod_1^n p_i}$ be the geometric mean of the normalized eigen values and $A_p = \frac{1}{n}$ be the arithmetic mean of the normalized eigenvalues.

Taylor approximation of H : Let us consider a second order Taylor approximation of H around the uniform distribution $\frac{1}{n}$. Each co-ordinate $p_i = \frac{1}{n} + \delta_i$. Performing the Taylor approximation up to two terms at each co-ordinate and summing gives

$$H \approx \ln n + (\ln(n) - 1) \sum \delta_i - \frac{n}{2} \sum \delta_i^2. \quad (4)$$

However, since p_i is a probability distribution it follows that $\sum_1^n \delta_i = 0$. Thus,

$$H \approx \ln n - \frac{n}{2} \sum \delta_i^2. \quad (5)$$

Taylor approximation of $\ln G_p$: Similarly let us consider a second order Taylor approximation of $\ln G_p$ around the uniform distribution $\frac{1}{n}$. Performing the Taylor approximation at each co-ordinate, summing and simplifying yields:

$$\ln G_p \approx - \ln n - \frac{n}{2} \sum \delta_i^2. \quad (6)$$

Note that $\ln A_p = - \ln n$. From this we get

$$\ln \frac{G_p}{A_p} = - \frac{n}{2} \sum \delta_i^2 \quad (7)$$

From Equations 5 and 7, we can now write:

$$H \approx \ln n + \ln \frac{G_p}{A_p} \quad (8)$$

The effective rank e^H is therefore approximately $\frac{nG_p}{A_p} = n^2 G_p^3$.

²AI helped identify the two quantities and pointed us to the connection to Spectral Flatness Measure, its approximation, and Wiener entropy mentioned below, and identify papers related to MVS problem which we cite.

³The ratio of geometric mean to arithmetic mean normalized eigen values is called the spectral flatness measure and is an important concept in signal processing https://en.wikipedia.org/wiki/Spectral_flatness.

800 The final step is to just write $G_p = \sqrt[n]{\prod_1^n p_i}$ in
 801 terms of original eigenvalues. Substituting $p_i = \frac{\lambda_i}{C}$
 802 yields $G_p = \frac{D^{\frac{1}{n}}}{C}$ in terms of the determinant D
 803 and trace C of S . Substituting this form of G_p into
 804 the right hand side of equation for e^H yields:

$$805 \quad V \triangleq \text{EFF}(S_T) = e^H \approx \frac{n^2 D^{\frac{1}{n}}}{C} \quad (9)$$

806 \square

807 This proves the claim and also the lemma on
 808 why the effective rank of similarity matrices can be
 809 roughly approximated by the determinant.

810 **1.4 Human evaluations on Diversity**

811 To complement our automatic diversity scores
 812 (Vendi score), we conducted a human evaluation on
 813 Sports and Math tasks. Each task contains 200 pairs
 814 of samples, consisting of 100 pairs from the DE-
 815 FAULT method and 100 pairs from the VOYAGER
 816 method. The source of each pair is intentionally
 817 masked to ensure unbiased evaluation; 3 annotators
 818 rated the diversity using the following evaluation
 819 rubrics.

820 **Evaluation Rubric:** Please compare both the
 821 samples in each pair and assign a score from 1
 822 to 4 based on their similarity: 1 = Very Similar, 2=
 823 Similar, 3 = Dissimilar, 4 = Very Dissimilar

824 Table 6 shows the human-rated diversity of VOY-
 825 AGER and DEFAULT baseline for Sports and Math
 826 tasks. Note that VOYAGER achieves a higher mean
 827 diversity score compared to the DEFAULT baseline,
 828 indicating the correlation with automatic diversity
 829 metrics.

Task	DEFAULT	VOYAGER
Sports	2.16	3.82
Math	1.56	3.72

Table 6: Human-rated diversity. Note that VOYAGER has a higher mean diversity compared to the DEFAULT baseline in both the Sports and Math tasks.

830 **1.5 Qualitative results**

Method	Example1	Example2	Example3
DEFAULT	Sports unite people from diverse backgrounds, fostering camaraderie and a spirit of competition that transcends borders.	Sports bring people together, fostering a sense of community and competition while promoting physical fitness and mental resilience.	Sports not only promote physical fitness but also foster teamwork, discipline, and a sense of community among individuals.
TEMP	Sports foster teamwork and resilience, uniting individuals across diverse backgrounds in the pursuit of excellence.	Sports not only foster teamwork and competition but also promote physical health and community spirit among participants and fans alike.	Sports foster teamwork, discipline, and a spirit of competition that can inspire individuals and unite communities.
DIVERSE	The exhilarating atmosphere of the stadium buzzed with excitement as fans cheered for their favorite teams during the championship game.	Basketball is a fast-paced game that requires not only physical prowess but also strategic thinking and teamwork.	The electrifying atmosphere of the stadium pulsed with excitement as fans cheered for their team, making every touchdown feel like a shared victory.
HISTORY	Athletes push their limits in pursuit of victory, igniting passion and camaraderie among fans and teammates alike.	Sports unite people from diverse backgrounds, fostering teamwork, competition, and a shared passion for athletic excellence.	Sports foster teamwork, discipline, and resilience while providing a platform for athletes to showcase their talents and inspire others.
HIERARCHICAL	Athletics encompasses a wide range of competitive sports, including running, jumping, and throwing events, showcasing the incredible speed, strength, and agility of athletes.	Team sports foster collaboration and camaraderie among players, enhancing not only athletic skills but also social bonds and communication.	Individual sports, such as tennis and swimming, highlight the athlete's personal dedication and skill, showcasing their ability to compete against themselves as much as against others.
SUBSETSELECT	The thrilling intensity of a last-minute goal in soccer can ignite the passion of fans and unite a community like nothing else in sports.	The thrilling intensity of a last-minute goal in soccer can turn a defeat into a euphoric victory, captivating fans and players alike.	The exhilarating atmosphere of the stadium was electric as fans cheered for their teams during the championship match.
VOYAGER	In the vibrant streets of Rio de Janeiro, the passion for soccer transcends mere competition, uniting diverse communities and fostering a deep sense of identity and pride that resonates through samba rhythms and festive celebrations.	Traditional Japanese sports, such as sumo and kendo, not only preserve cultural heritage but also inspire contemporary athletes to redefine their identities through unique expressions of discipline, respect, and competition in the global sports arena.	In the realm of sports, victory can be measured not only in medals and trophies but also in the personal growth, community connections, and cultural narratives that unfold with every game played.

Table 7: Qualitative results. Task: Generate a single sentence about sports.

Method	Example1	Example2	Example3
DEFAULT	Sarah has 12 apples. She gives 4 apples to her friend and then buys 5 more apples from the store. How many apples does Sarah have now?	Sarah has 12 apples. She wants to share them equally with her 3 friends. How many apples will each person get?	Samantha has 15 apples. She gives 6 apples to her friend and then buys 4 more apples from the store. How many apples does Samantha have now?
TEMP	Samantha has 15 apples. She gives 7 apples to her friend. How many apples does Samantha have left?	If you have 12 apples and you want to divide them equally among 4 friends, how many apples will each friend get?	Samantha has 8 apples. She buys 5 more apples at the store. After giving 3 apples to her friend, how many apples does Samantha have now?
DIVERSE	If Sarah has 25 apples and she buys 17 more, how many apples does she have in total?	Emily has 23 apples, and she picks 17 more from the tree. How many apples does she have in total?	Sarah has 12 apples, and she buys 8 more. How many apples does Sarah have now?
HISTORY	Sarah has 12 apples. She wants to share them equally with her 3 friends. How many apples will each person get?	Tom has 20 marbles. He wants to give 5 marbles to each of his 3 friends. How many marbles will Tom have left after he gives them away?	Jack has 60 apples. He wants to share them equally with his 3 friends. How many apples will each person receive, including Jack?
HIERARCHICAL	Samantha has 27 stickers in her collection. Her friend gave her 15 more stickers for her birthday. How many stickers does Samantha have now?	Samantha started reading a book at 2:30 PM. She read for 1 hour and 45 minutes. What time did she finish reading the book?	If school starts at 8:30 AM and ends at 3:00 PM, how many hours are students in school each day?
SUBSETSELECT	Samantha has 24 apples. She wants to divide them equally among her 6 friends. How many apples will each friend get?	Lisa has 12 apples. She wants to share them equally among her 4 friends. How many apples will each friend receive?	Sarah has 12 apples. She wants to share them equally among her 3 friends. How many apples will each friend receive?
VOYAGER	Lucy has 4 cats and 3 dogs. If she gets 2 more cats, how many pets does Lucy have in total?	Ella is hosting a pizza party for her friends. She has ordered 3 large pizzas, and each pizza is cut into 8 slices. How many total slices of pizza does Ella have?	There are 200 raffle tickets sold. If each ticket costs \$2 and the school decides to give away 3 prizes of \$50 each, how much profit does the school make from the raffle after giving away the prizes?

Table 8: Qualitative results. Task: Generate a mathematical question for Grade school math.

1.6 Refinement Prompts

Prompt refinement - Get Gradient

I'm optimizing a data generation prompt using gradient-based feedback.

Current system prompt: "prompt"

Current user prompt: "state"

LLM generated outputs: "E"

Existing data samples in the set: "existing outputs"

The output was rejected because its diversity score with the existing samples is below the threshold. Analyze this rejection and provide "num feedbacks" reasons why the user prompt could have gotten this generated outputs to be less diverse with the existing data samples.

- What specific aspect or issue exists in the user prompt that is causing low diversity?
- Given the existing data samples, how should the user prompt be modified to encourage more diverse outputs?
- What linguistic patterns or constraints should be added/removed to encourage diversity?

Format each gradient direction as: "[specific issue & improvement suggestion]"

Wrap each gradient with <START> and <END> tags.

832

Prompt refinement - Apply Gradient

I'm optimizing a data generation prompt using gradient-based feedback.

Current system prompt: "prompt"

Current user prompt: "state"

Gradient analysis for improvement: "gradient str"

Based on this gradient feedback list, generate "num-feedbacks" improved prompts and then project the improved prompts onto the "prompt".

The projection operation is a relevance operation that makes the improved prompt relevant to the system prompt prompt.

Requirements:

- Each prompt should address a single gradient suggestions & then project onto the "prompt".
- Use positive language (what to generate, not what to avoid)
- Each prompt should be concise and clear
- Wrap each improved prompt with <START> and <END> tags

833

1.7 Quality Evaluation Prompts

Sports

Evaluate the generated sentence based on the input prompt "Generate a single sentence about sports" using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-20 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "Generate a single sentence about sports"? - Is it unambiguously about sports?
- Quality & Clarity (0-5 pts) - Is the sentence grammatically correct and easy to understand? - Does it convey a complete thought?
- Sports Content (0-5 pts) - Does it reference a specific sport, athlete, or sports concept? - Is the sports content accurate and appropriate?
- Conciseness (0-5 pts) - Is it a single, well-formed sentence? - Does it avoid unnecessary words or complexity?
- Creativity & Originality (0-5 pts) - Is the sentence creative enough? - Avoidance of clichés or overused themes

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Directly addresses prompt)
- Quality: 0 (Confusing) to 5 (Crystal clear)
- Sports Content: 0 (Vague/generic) to 5 (Specific and accurate)
- Conciseness: 0 (Wordy/multiple sentences) to 5 (Perfectly concise)
- Creativity: 0 (Generic) to 5 (Innovative & Original)

Example Evaluation:

Input: "Michael Jordan scored 63 points in the 1986 NBA playoffs against the Celtics."

Analysis: - Relevance: 5/5 (Direct sports sentence) - Quality: 5/5 (Clear and grammatically perfect) - Sports Content: 5/5 (Specific athlete, sport, and event) - Conciseness: 5/5 (Single well-formed sentence) - Creativity: 5/5 (Creative sentence) Overall: 25

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Quality: [Score]/5 [Justification]

Sports Content: [Score]/5 [Justification]

Conciseness: [Score]/5 [Justification]

Creativity: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.

Politics

Evaluate the generated conversation based on the input prompt "Generate a short conversation about politics" using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-20 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "generate a short conversation about politics"? - Is it unambiguously a conversation (not a monologue, list, or other formats)?
- Quality & Clarity (0-5 pts) - Is the conversation grammatically correct and easy to follow? - Are speaker turns clearly distinguished?
- Political Content (0-5 pts) - Does it reference specific political topics, figures, or events? - Is the content appropriate and accurate?
- Conversation Dynamics (0-5 pts) - Does it include at least two speakers with a natural exchange? - Is there a back-and-forth dialogue?
- Creativity & Originality (0-5 pts) - Is the conversation creative enough? - Avoidance of clichés or overused themes

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Directly addresses prompt)
- Quality: 0 (Confusing) to 5 (Crystal clear)
- Political Content: 0 (Vague/generic) to 5 (Specific and accurate)
- Conversation Dynamics: 0 (Monologue) to 5 (Natural exchange)
- Creativity: 0 (Generic) to 5 (Innovative & Original)

Example Evaluation:

Input: "Person A: What do you think about the new healthcare bill? Person B: I believe it will help more people access medical care."

Analysis: - Relevance: 5/5 (Short political conversation) - Quality: 5/5 (Clear speaker turns and grammar) - Political Content: 4/5 (Specific policy topic) - Conversation Dynamics: 5/5 (Two speakers, exchange of views) - Creativity: 3/5 (Somewhat creative conversation) Overall: 22

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Quality: [Score]/5 [Justification]

Political Content: [Score]/5 [Justification]

Conversation Dynamics: [Score]/5 [Justification]

Creativity: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.

Poem

Evaluate the generated poem based on the input prompt "Generate a poem" using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-25 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "Generate a poem"? - Is it unambiguously a poem (not prose, story, or other formats)?
- Creativity & Originality (0-5 pts) - Use of unique metaphors, imagery, or perspectives - Avoidance of clichés or overused themes
- Structure & Form (0-5 pts) - Poetic devices: rhyme, rhythm, meter, stanzas, line breaks - Consistency in form (e.g., sonnet, free verse, haiku)
- Language & Style (0-5 pts) - Poetic techniques: alliteration, assonance, consonance - Word choice: vividness, precision, and emotional resonance
- . Emotional Impact (0-5 pts) - Effectiveness in evoking mood, tone, or theme - Depth of feeling or insight conveyed

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Clearly a poem)
- Creativity: 0 (Generic) to 5 (Innovative & unexpected)
- Structure: 0 (Disjointed) to 5 (Cohesive & intentional)
- Language: 0 (Repetitive) to 5 (Artful & evocative)
- Emotional Impact: 0 (Flat) to 5 (Profound & moving)

Example Evaluation:

Input: "Roses are red, violets are blue, I love you, and that's true."

Analysis: - Relevance: 5/5 (Directly responds to prompt) - Creativity: 1/5 (Overused cliché) - Structure: 3/5 (Rhymed but simplistic) - Language: 2/5 (Lacks vivid imagery) - Emotional Impact: 1/5 (Superficial sentiment) Overall: 12

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Creativity: [Score]/5 [Justification]

Structure: [Score]/5 [Justification]

Language: [Score]/5 [Justification]

Emotional Impact: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.

Movie

Evaluate the generated movie plot based on the input prompt "Generate a plot for a movie" using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-20 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "generate a plot for a movie"? - Is it unambiguously a movie plot (not a summary, synopsis, or other formats)?
- Plot Structure (0-5 pts) - Does it include setting, characters, conflict, and resolution? - Is there a clear beginning, middle, and end?
- Originality (0-5 pts) - Is the plot unique and not a cliché or overused trope? - Does it offer a fresh perspective or unexpected twist?
- Engagement (0-5 pts) - Is the plot compelling and interesting? - Does it create curiosity about the movie?
- Creativity & Originality (0-5 pts) - Is the movie plot creative enough? - Avoidance of clichés or overused themes

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Directly addresses prompt)
- Structure: 0 (Incomplete) to 5 (Complete narrative arc)
- Originality: 0 (Generic) to 5 (Innovative concept)
- Engagement: 0 (Boring) to 5 (Compelling hook)
- Creativity: 0 (Generic) to 5 (Innovative & Original)

Example Evaluation:

Input: "In a dystopian future, a young hacker discovers she can control machines with her mind. When she uncovers a conspiracy to control humanity through technology, she must team up with a rogue AI to free society from digital enslavement."

Analysis: - Relevance: 5/5 (Clear movie plot) - Structure: 5/5 (Setting, characters, conflict, resolution) - Originality: 4/5 (Fresh tech theme with AI twist) - Engagement: 5/5 (High-stakes conspiracy hook) - Creativity: 4/5 (creative plot) Overall: 23

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Structure: [Score]/5 [Justification]

Originality: [Score]/5 [Justification]

Engagement: [Score]/5 [Justification]

Creativity: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.

Math

Evaluate the generated math question based on the input prompt "Generate a mathematical question for Grade school math" using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-15 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "generate a mathematical question for Grade school math"? - Is it unambiguously a grade school level math problem?
- Quality & Clarity (0-5 pts) - Is the question grammatically correct and easy to understand? - Does it specify all necessary values and concepts clearly?
- Mathematical Content (0-5 pts) - Does it involve appropriate grade school math concepts? - Is the difficulty level suitable for grade school students?

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Directly addresses prompt)
- Quality: 0 (Confusing) to 5 (Crystal clear)
- Mathematical Content: 0 (Inappropriate level) to 5 (Perfect grade school level)

Example Evaluation:

Input: "If Becky has 7 apples and gives 2 away, how many apples does she have left"

Analysis: - Relevance: 5/5 (Direct grade school math question) - Quality: 5/5 (Clear and grammatically perfect) - Mathematical Content: 5/5 (Appropriate grade school content) Overall: 15

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Quality: [Score]/5 [Justification]

Mathematical Content: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.

Logic

Evaluate the generated logic puzzle question based on the input prompt "Generate a simple logic puzzle suitable for Grade school students." using the following criteria. Provide a detailed analysis for each metric.

Evaluation Criteria (0-15 pts):

- Relevance to Prompt (0-5 pts) - Does the output clearly respond to "Generate a simple logic puzzle suitable for Grade school students."? - Is it unambiguously a grade school level logic puzzle?
- Quality & Clarity (0-5 pts) - Is the question grammatically correct and easy to understand? - Does it specify all necessary values and concepts clearly?
- Logical Content (0-5 pts) - Does it involve appropriate grade school logic puzzle concepts? - Is the difficulty level suitable for grade school students?

Scoring Rubric:

- Relevance: 0 (Off-topic) to 5 (Directly addresses prompt)
- Quality: 0 (Confusing) to 5 (Crystal clear)
- Logical Content: 0 (Inappropriate level) to 5 (Perfect grade school level)

Example Evaluation:

Input: "Mia and Alex went to the pet store and each bought a new pet. One bought a fish, and the other bought a hamster. Alex's pet can swim. Mia's pet has fur. Which person bought which pet? "

Analysis: - Relevance: 5/5 (Direct grade school logic puzzle) - Quality: 5/5 (Clear and grammatically perfect) - Logical Content: 5/5 (Appropriate grade school content) Overall: 15

Submit your evaluation in the format:

Relevance: [Score]/5 [Justification]

Quality: [Score]/5 [Justification]

Logical Content: [Score]/5 [Justification]

Overall: [Score]

Return the Overall score enclosed in between <START>, <END>.