QD-REC: OPEN-ENDED EVOLUTION OF RECOMMENDATIONS VIA QUALITY–DIVERSITY

Anonymous authors
Paper under double-blind review

ABSTRACT

Open-ended search based on foundation models has produced faster algorithms, new results on open mathematical problems, and creative discoveries. We extend this paradigm to decision-making, where the set of possible choices is fixed but the reasons for choosing them are not. Examples are choosing which company to acquire, which paper to read, or which customer to target. In high-stakes decisions, comparing diverse, well-reasoned options can be more useful than a single ranked list. We formalize *open-ended recommendation* as building a diverse, high-quality set of recommendations, each an item together with a *hypothesis*, a natural-language reason why the item fits the needs of a decision maker. We propose QD-Rec, a quality–diversity method that selects items with a contextual bandit and generates hypotheses with a foundation model from earlier ones in an archive. A new recommendation is kept in the archive if it is novel or better than its nearest neighbor. Novel recommendations can thus serve as stepping stones for later hypotheses. Across M&A, customer targeting, and paper and book recommendation, QD-Rec gives the highest mean quality and builds the largest archive in most settings.

1 INTRODUCTION

In open-ended search with foundation models, the model repeatedly generates new candidates from earlier ones, and an archive keeps a selection of them for the model to build on in later rounds. This approach has found candidates that no single query to the model would produce, including new mathematical constructions (Romera-Paredes et al., 2024; Lange et al., 2026), faster algorithms (Novikov et al., 2025), agent designs (Hu et al., 2025), tasks that reveal a model's capabilities (Lu et al., 2025), adversarial prompts (Samvelyan et al., 2024), and creative text (Bradley et al., 2024).

We ask whether open-ended search can extend to recommendation. Unlike open-ended generation, where the model generates new candidates, recommendation works with a fixed pool of candidates: an acquirer chooses among the companies that exist, a reader among the papers already written, a marketer among the customers. At first glance, this leaves no room for open-endedness, since the pool is finite and known in advance. However, only the items are fixed. The same item can be recommended for many different reasons, so the search over reasons is still open-ended.

We therefore treat recommendation as an open-ended search over item–hypothesis pairs. A *hypothesis* is a natural-language explanation of why an item fits the decision maker's profile, such as the profile of one acquirer or one reader. The search returns a diverse, high-quality set of such pairs for a single profile, rather than one ranked list of items. We call this problem *open-ended recommendation*. In high-stakes decisions, comparing alternative reasons can be as important as comparing the items themselves (Aggarwal, 2016). We evaluate the set by its quality and diversity, following the quality–diversity (QD) view of *illuminating* a space of alternatives (Mouret & Clune, 2015).

Open-ended recommendation has two properties that make open-ended search a good fit. First, evaluation is costly: we learn a recommendation's value only by writing its hypothesis and judging it, so we cannot simply score every possible recommendation in advance and sort them. The system must decide which item to evaluate next and learn from each result. Second, the space is larger than the pool: each item allows unlimited hypotheses, so even evaluating every item once would leave most of the space unexplored. Good hypotheses often come from adapting or combining earlier ones,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

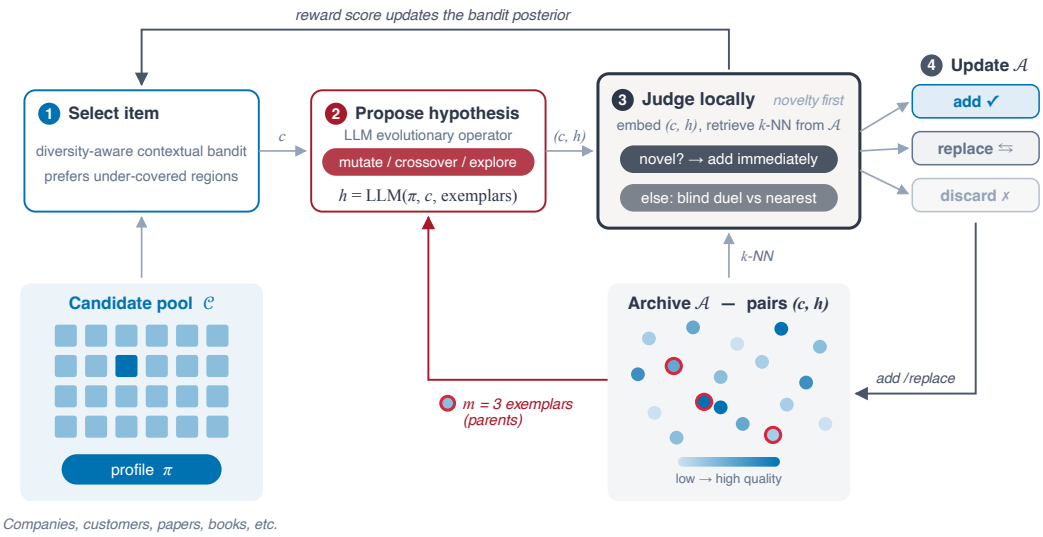

Figure 1: Overview of QD-Rec. Each iteration selects an item from the fixed pool $\mathcal{C}$ via a diversity-aware bandit, generates a hypothesis with an LLM conditioned on archive exemplars, which act as parents (an evolutionary operator), and judges the new pair against its k nearest neighbors: a novel pair is added immediately; a non-novel one survives only by winning a pairwise duel against its neighbor.

so keeping modest intermediate hypotheses as *stepping stones* can help (Lehman & Stanley, 2011; Stanley & Lehman, 2015). QD and novelty search were designed for such problems (Risi et al., 2025), but to our knowledge they have not been applied to recommendation.

We introduce **QD-Rec**, a quality–diversity search method for open-ended recommendation (Figure 1). In QD-Rec, an LLM called the *proposer* writes hypotheses, and an LLM judge or humans score them. Each iteration has three steps. A contextual bandit selects an item, balancing items that look promising against regions the archive has not yet covered. The proposer generates a hypothesis for that item, building on a few earlier hypotheses from the archive. The new recommendation is kept if it is novel, or if it is better than its nearest neighbor in the archive. Novelty and quality are thus judged *locally*, by comparison with nearby recommendations in an embedding space, instead of on a hand-designed grid of behavior descriptors.

Our contributions are: (i) we formalize *open-ended recommendation* as a quality–diversity search over item–hypothesis pairs, where the items come from a fixed pool; (ii) we propose QD-Rec, with two ablations isolating its components (Table 1); and (iii) across four domains and two proposers, QD-Rec gives the highest mean quality in most LLM-judged settings, usually builds the largest archive, and outperforms diversity-aware rerankers (MMR and DPP) on human-rated quality.

2 Background and Related Work

We briefly review quality–diversity methods and novelty search, evolutionary loops built around foundation models, and recommender systems.

Quality–diversity and novelty search. Quality–diversity (QD) methods aim to discover a *set* of high-performing yet diverse solutions rather than a single optimum. Multi-dimensional Archive of Phenotypic (MAP) Elites "illuminates" a search space by maintaining an archive and retaining elites across behavioral niches (Mouret & Clune, 2015). Novelty search emphasizes exploration by rewarding behavioral novelty, providing a mechanism for stepping-stone trajectories in which intermediate solutions enable later progress (Lehman & Stanley, 2011). Recent work continues to revisit local competition and novelty definitions in QD, including dominated novelty variants

(Bahlous-Boldi et al., 2025), and extensions that reduce dependence on hand-designed behavioral descriptors, such as dimensionless MAP-Elites (Dharna et al., 2025).

Foundation models in evolutionary loops. A growing family of systems places a frozen large model inside an evolutionary or QD loop as a proposal operator and/or evaluator, keeping an archive of artifacts that the model conditions on to propose the next: diverse creative text in QDAIF (Bradley et al., 2024), adversarial prompts in Rainbow Teaming (Samvelyan et al., 2024), agent skills in CycleQD (Kuroki et al., 2025), code and program search in ELM (Lehman et al., 2023) and FunSearch (Romera-Paredes et al., 2024), algorithm discovery in AlphaEvolve (Novikov et al., 2025), sample-efficient program evolution in ShinkaEvolve (Lange et al., 2026), open-ended evaluation tasks in Automated Capability Discovery (Lu et al., 2025), agent designs in ADAS (Hu et al., 2025), and image evolution in a Picbreeder replication whose vision–language models stand in for its human evaluators (Earle et al., 2026). A separate line of work is *self-referential*: STOP (Zelikman et al., 2024), the Darwin Gödel Machine (Zhang et al., 2026a), and Hyperagents (Zhang et al., 2026b) turn the search on the machinery itself, so the improved scaffold or agent performs the next round of improvement. QD-Rec belongs to the first group, with item–hypothesis pairs as the artifacts. However, these systems create new artifacts, while QD-Rec must choose its items from a fixed pool.

Recommender systems. Recommenders learn either *offline*, from a fixed dataset of past interactions (Aggarwal, 2016), or *online*, as a sequential decision problem with an exploration–exploitation trade-off, usually modeled as a (contextual) bandit (Lattimore & Szepesvári, 2020; Agarwal et al., 2009; Li et al., 2010; Agrawal & Goyal, 2013; Song et al., 2022). The evolutionary loops above are sequential as well, but they generate their candidates rather than choose them from a pool. QD-Rec does both, generating hypotheses for items selected from a fixed pool, so it selects items with a contextual bandit; unlike a recommender's bandit, however, it is judged, as in QD, by its final archive. Orthogonally, a long line of work treats diversity, novelty, coverage, and serendipity as first-class objectives (McNee et al., 2006; Ziegler et al., 2005; Vargas & Castells, 2011; Kaminskas & Bridge, 2016), pursued offline by re-ranking (Carbonell & Goldstein, 1998; Chen et al., 2018) and online by diversity-aware bandits (Radlinski et al., 2008; Liu et al., 2020). These methods diversify the list shown to a user; QD-Rec instead enforces diversity at archive admission, following recent QD systems (Lu et al., 2025; Dharna et al., 2025): an LLM judges whether each new pair is novel relative to its nearest archive neighbors. Finally, explainable recommendation adds a natural-language explanation to each recommended item (Zhang & Chen, 2020), and so do LLM recommenders that reason before choosing or justify afterward. QD-Rec also pairs each item with a reason, but it does so inside an evolutionary loop.

3 Open-Ended Recommendation

Let $\mathcal{C} = \{c_1, \ldots, c_N\}$ be a fixed candidate pool (e.g., companies for M&A, customer profiles for targeting, papers, or books), and let π be the *profile* of one decision maker. A *hypothesis* $h \in \mathcal{H}$ is a natural-language explanation of why an item fits π, and the hypothesis space $\mathcal{H}$ is unbounded. A *recommendation* is a pair $(c, h) \in \mathcal{C} \times \mathcal{H}$, so the same item can be recommended several times with different hypotheses. Its quality $q(c, h)$ is a score from a human or an LLM judge, for example on a 0 to 10 scale. The score is known only after the hypothesis is written and evaluated, and each evaluation is costly. Given π and $\mathcal{C}$, the task is to return a set $\mathcal{S} \subset \mathcal{C} \times \mathcal{H}$ of recommendations that have high quality and differ from one another, using at most T evaluations. Section 5.1 defines the metrics for quality and diversity.

4 QD-Rec

We present **QD-Rec**, an *archive-based quality–diversity (QD) method* for open-ended recommendation. QD-Rec keeps an archive $\mathcal{A}$ of retained recommendations and returns it as $\mathcal{S}$ at the end of the run. Each new entry is compared only with its nearest neighbors in the archive to decide whether it is kept. Each archive entry x stores a recommendation (c, h), its embedding $e = \mathrm{Embed}(c, h)$, and its quality score $q(x)$. The embedding is computed from a structured item profile followed by the hypothesis text. We use cosine similarity $\mathrm{sim}(x_i, x_j) = \langle e_i, e_j \rangle / (\|e_i\| \, \|e_j\|)$ to define these neighborhoods.

4.1 LOCAL NOVELTY, RANKING, AND ARCHIVE UPDATE

At each step we propose a new pair by selecting an item c and generating a hypothesis h (Section 4.2); we write $x_{\text{new}} = (c, h, e_{\text{new}})$ with $e_{\text{new}} = \text{Embed}(c, h)$ for the proposed entry (its quality estimate is not yet known at this point). We retrieve the k nearest neighbors from the archive,

$$\mathcal{N}_k(x_{\text{new}}) = \arg\max_{x \in \mathcal{A}}^{k} \text{sim}(x_{\text{new}}, x).$$

The archive update then proceeds in the following two steps:

- *Novelty check*: an LLM judge is prompted with the profile π, the new entry, and its k nearest neighbors, and returns a binary decision `is_novel` $\in \{0, 1\}$ (with brief reasoning). A novel entry is admitted immediately; no comparison is made.
- *Local quality comparison* (only when not novel): we use a blind pairwise duel: of the k neighbors only the single *nearest* one is contested. It and x_{new} are shuffled into two anonymous positions, pointwise scores are withheld, and the judge, never told which of the two is the newcomer, chooses the stronger under π (ties are not allowed). If x_{new} wins it replaces that neighbor; otherwise it is discarded. Algorithm 1 writes this comparison as Better_π. Separating the two questions keeps the retention comparison *blind* (the judge cannot favor or disfavor an entry for being new) and skips the comparison call entirely whenever the entry is novel.

This embedding-neighborhood admission scheme is inspired by recent novelty-search/QD systems that use learned representations for similarity and archive admission (Lu et al., 2025; Dharna et al., 2025).

The quality score $q(x)$ of each evaluated pair (Section 3) has two uses. The bandit reward $r(x)$ is the score divided by the maximum of its scale, so that rewards lie in $[0, 1]$. Reported Quality averages the score over the final archive, on its original scale.

4.2 ITEM SELECTION AND HYPOTHESIS GENERATION

The archive mechanics above leave two decisions open: which item $c \in \mathcal{C}$ to evaluate next, and how to generate a hypothesis h for it.

Item selection. Each candidate item c has a *precomputed candidate embedding* $\psi(c) \in \mathbb{R}^d$ (e.g., title+abstract for papers, or a structured company/customer-profile representation), distinct from the archive-entry embedding e and taken to be L2-normalized. We model a scalar reward $r \in [0, 1]$ by Bayesian linear regression, $r = \psi(c)^\top \theta + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$ and prior $\theta \sim \mathcal{N}(0, \lambda^{-1} I)$. For the set $\mathcal{D}$ of evaluated pairs we maintain the sufficient statistics

$$A \;=\; \lambda I + \sigma^{-2} \sum_{(\psi, r) \in \mathcal{D}} \psi \psi^\top, \qquad b \;=\; \sigma^{-2} \sum_{(\psi, r) \in \mathcal{D}} r\, \psi, \qquad \theta \mid \mathcal{D} \;\sim\; \mathcal{N}(A^{-1} b,\; A^{-1}),$$

and a new observation $(\psi(c), r)$ updates them in closed form:

$$A \;\leftarrow\; A + \sigma^{-2}\, \psi(c) \psi(c)^\top, \qquad b \;\leftarrow\; b + \sigma^{-2}\, r\, \psi(c).$$

Linear Thompson sampling (Agrawal & Goyal, 2013) then draws $\tilde{\theta} \sim \mathcal{N}(A^{-1} b, A^{-1})$ at each step, and we select the candidate that maximizes the *diversity-aware selection score*

$$s(c) \;=\; \psi(c)^\top \tilde{\theta} \;+\; \beta\, d(c, \mathcal{R}), \qquad d(c, \mathcal{R}) \;=\; 1 - \max_{c' \in \mathcal{R}} \psi(c)^\top \psi(c'), \tag{1}$$

where $\mathcal{R}$ collects the candidate embeddings of the current archive. The bonus $d(c, \mathcal{R})$ is the cosine distance to the *nearest* covered item (the standard nearest-neighbor novelty measure), so it is large exactly when nothing similar has been archived yet, and $\beta \geq 0$ trades this coverage of under-explored regions against the sampled quality estimate $\psi(c)^\top \tilde{\theta}$. The bandit *reward* is unchanged: after evaluation we update the posterior using the normalized reward $r(x)$ defined above, so β affects only *which* items are chosen for evaluation, steering the archive to spread rather than cluster. Setting $\beta = 0$ recovers the plain quality-only bandit; we use $\beta = 1$ throughout. We warm-start the posterior

Algorithm 1 QD-Rec. The Novelty+Evo ablation uses uniform-random selection from $\mathcal{C}$; Novelty additionally uses no exemplars ($m = 0$).

Require: Candidate pool $\mathcal{C}$, profile π, seed archive $\mathcal{A}$, iterations T, neighborhood size k, exemplars m, diversity weight $\beta \geq 0$, prior precision λ, noise variance σ^2; warm-start $A = \lambda I$, $b = 0$ with one observation $(\psi(c), r(x))$ per seed entry

for $t = 1$ **to** T **do**
  Sample $\tilde{\theta} \sim \mathcal{N}(A^{-1}b,\ A^{-1})$ *// diversity-aware bandit item selection*
  $c \leftarrow \arg\max_{c \in \mathcal{C}} \left[\psi(c)^\top \tilde{\theta} + \beta\, d(c, \mathcal{R})\right]$
  $h \leftarrow \mathrm{LLM}_{\text{mut}}(\pi, c, \mathrm{SampleParents}(\mathcal{A}, m))$ *// few-shot hypothesis generation*
  $x_{\text{new}} \leftarrow (c, h, \mathrm{Embed}(c, h))$ *// local novelty + ranking archive update*
  Obtain $q(x_{\text{new}})$; set $r \leftarrow r(x_{\text{new}})$ *// score scaled to* $[0, 1]$
  $\mathcal{N}_k \leftarrow k\mathrm{NN}(x_{\text{new}}, \mathcal{A})$
  `novel` $\leftarrow \mathrm{LLM}_{\text{novelty}}(\pi, x_{\text{new}}, \mathcal{N}_k)$ *// novelty first*
  if `novel` **then**
    Add x_{new} to $\mathcal{A}$ *// admitted; no comparison call*
  else if $\mathrm{Better}_\pi(x_{\text{new}},$ nearest neighbor$)$ **then**
    Replace nearest neighbor with x_{new} *// blind pairwise duel*
  end if
  $A \leftarrow A + \sigma^{-2}\psi(c)\psi(c)^\top,\ b \leftarrow b + \sigma^{-2}\, r\, \psi(c)$ *// update even if discarded*
end for
return $\mathcal{A}$

from the seed archive: each seed entry contributes one observation $(\psi(c), r(x))$, using the same normalization.

An item can be selected more than once, because a new hypothesis for it can receive a different quality score. The same item can also appear in the archive several times. If a new hypothesis for an item differs enough from the entries near it, the novelty check admits it as a new entry.

Hypothesis generation (evolutionary operator). Given the selected item c, we prompt the proposer LLM with the profile π, the item c, and m exemplars sampled uniformly from the current archive. The prompt presents the exemplars as existing hypotheses and instructs the model to apply one of three operators, *mutate* (adapt one exemplar's angle to c), *crossover* (combine angles from several exemplars), or *explore* (propose an angle the exemplars miss), so the archive acts as an evolving population (Guo et al., 2024; Lehman et al., 2023).

Simpler variants. We also define two simpler variants of QD-Rec, which serve as ablations of the full method. *Novelty+Evo* replaces the bandit with uniform-random item selection. *Novelty* removes both the bandit and the exemplars, so the proposer generates hypotheses zero-shot. All three methods share the archive update of Section 4.1, and Table 1 compares them. Algorithm 1 gives QD-Rec in full.

4.3 Implementation details

To speed up experiments, we run the loop in parallel batches where workers propose and evaluate candidates against a snapshot of the current archive. To select a batch of b items, we extend the single-draw rule of Section 4.2: we take b independent posterior draws $\tilde{\theta}_1, \ldots, \tilde{\theta}_b$, and for each $\tilde{\theta}_i$ select the candidate maximizing $s(c)$ while excluding items already chosen in the batch and adding them to $\mathcal{R}$, so the diversity bonus discourages near-duplicates within a batch as well as against the archive. After each batch, we reconcile results by filtering near-duplicate additions (above a cosine-similarity threshold) and resolving replacement collisions (multiple workers targeting the same neighbor). The bandit uses prior precision $\lambda = 1$ and noise standard deviation $\sigma = 0.1$, the novelty check uses $k = 3$ nearest neighbors, and each run starts from 10 seed entries.

5 Experiments

We evaluate the three variants (Novelty, Novelty+Evo, QD-Rec; Table 1) across four domains.

	Local novelty filter	Evolutionary operator	Bandit item selection
Novelty	✓	×	×
Novelty+Evo	✓	✓	×
QD-Rec (full)	✓	✓	✓

Table 1: QD-Rec and its two ablations. All three use the same archive update (Section 4.1).

5.1 SETUP

Common protocol. Our main experiment uses `claude-sonnet-5` as the hypothesis proposer; we additionally report a weaker-proposer ablation that swaps in `claude-haiku-4-5`. The judge is `gpt-5.6-terra` in *every* configuration and domain, a different model family from both proposers, to prevent self-preference bias (Wataoka et al., 2024). Proposer and judge both run with reasoning effort set to `high`. The weak-proposer ablation model does not support the effort parameter and runs without it or extended thinking. We leave the temperature setting to the provider default for `claude-haiku-4-5`. Each configuration runs $1{,}000$ evaluations (proposal–judge cycles) per method in all four domains, executed in generations of up to 20 parallel evaluations, with the archive reconciled between generations (Section 4.3). Novelty+Evo and QD-Rec condition each proposal on $m = 3$ archive exemplars. Runs are repeated for five random seeds. Each run begins from a seed archive under the same profile π; the seed controls the pipeline's own sampling randomness (parent selection, item ordering) while LLM sampling stays at the provider default, so across-seed spread measures the stochasticity of the full pipeline.

We embed archive entries and candidate items with OpenAI's `text-embedding-3-small` at 512 output dimensions, L2-normalized, and use cosine similarity for k-nearest-neighbor retrieval.

Methods. We evaluate the proposed method QD-Rec and its two ablations (Novelty, Novelty+Evo) under matched evaluation budgets and shared hyperparameters within each domain. Table 1 summarizes the three variants as an "ingredients" view: all share the archive mechanics; Novelty+Evo adds the evolutionary few-shot operator, and QD-Rec additionally adds diversity-aware bandit item selection.

Domains. We instantiate the framework on four fixed-pool domains, each realizing the hypothesis h in its native genre. The first, Goodreads, is a *grounding check* whose quality signal is a real human rating rather than an LLM judge; the other three are heterogeneous applied settings. Full dataset construction (sources, filtering, pool sizes, and the specific profiles used) is in Appendix C.

- **Goodreads.** The pool $\mathcal{C}$ is the children's books a reader has rated (Wan & McAuley, 2018); the profile π is that reader (summarized from past reviews); h is a reading memo on whether they would enjoy a book. Uniquely, quality is *not* an LLM judge but the reader's own 1–5 rating, held out of all prompts and used only as the reward and replacement signal (Appendix C), testing whether QD still yields meaningful quality–diversity tradeoffs under a real human signal. This Quality measure evaluates the book for the reader, not the generated memo: changing the memo cannot change the rating, although it can affect novelty and archive geometry.
- **M&A target discovery.** $\mathcal{C}$ is a universe of 6,332 US companies built from Refinitiv/LSEG Workspace data (business descriptions, industry taxonomy, financials), π an acquirer (Walt Disney or JPMorgan Chase), and h an acquisition thesis (e.g. capability acquisition vs. market access vs. vertical integration). Deal sourcing motivates it: teams want a *diverse set of strategic theses* rather than a single shortlist.
- **Marketing.** $\mathcal{C}$ is 10,000 synthetic customer profiles from California, sampled from NVIDIA Nemotron-Personas-USA (Meyer & Corneil, 2025), π a product (a desktop 3D printer or an online upskilling subscription), and h a targeting hypothesis explaining why a profile is an attractive target for π, yielding a map of distinct targeting angles rather than a ranked list.
- **Paper-abstract exploration.** $\mathcal{C}$ is the titles and abstracts of 17,511 papers accepted at ICLR 2026, ICML 2026 and NeurIPS 2025 (Paper Copilot (Jing Yang), 2026), π a researcher profile (an open-endedness researcher or an applied researcher at an AI startup), and h a research memo linking a paper to π. The scale of the AI literature motivates it: researchers need *diverse*, well-

Domain	Profile	Proposer	Judge	Novelty				Novelty+Evo				QD-Rec			
				Qual.	Div.	Cov.	n	Qual.	Div.	Cov.	n	Qual.	Div.	Cov.	n
Goodreads†	series/YA	Sonnet 5	human	**4.17**$\pm$**0.01**	.463$\pm$.001	.26$\pm$.00	222$\pm$5	**4.15**$\pm$**0.01**	**.480**$\pm$**.001**	**.27**$\pm$**.00**	**226**$\pm$5	**4.15**$\pm$**0.02**	**.480**$\pm$**.001**	**.28**$\pm$**.01**	**244**$\pm$3
		Haiku 4.5	human	**4.09**$\pm$**0.01**	.480$\pm$.001	.26$\pm$.00	233$\pm$2	4.08$\pm$0.01	**.493**$\pm$**.001**	**.29**$\pm$**.00**	254$\pm$3	**4.12**$\pm$**0.01**	**.492**$\pm$**.002**	**.28**$\pm$**.01**	**257**$\pm$3
	picture-book	Sonnet 5	human	**4.19**$\pm$**0.01**	.475$\pm$.002	.32$\pm$.00	297$\pm$6	**4.17**$\pm$**0.02**	**.481**$\pm$**.002**	.33$\pm$.01	318$\pm$6	4.12$\pm$0.02	**.485**$\pm$**.001**	**.34**$\pm$**.01**	**336**$\pm$7
		Haiku 4.5	human	**4.14**$\pm$**0.01**	.489$\pm$.001	.32$\pm$.01	320$\pm$2	**4.15**$\pm$**0.02**	.497$\pm$.003	.33$\pm$.01	337$\pm$4	**4.11**$\pm$**0.02**	**.506**$\pm$**.001**	**.35**$\pm$**.01**	**367**$\pm$4
M&A	Walt Disney Co	Sonnet 5	GPT-5.6	2.40$\pm$0.12	.330$\pm$.001	.12$\pm$.00	85$\pm$1	2.51$\pm$0.07	**.353**$\pm$**.001**	.18$\pm$.01	135$\pm$5	**2.91**$\pm$**0.05**	**.352**$\pm$**.001**	**.26**$\pm$**.00**	**214**$\pm$2
		Haiku 4.5	GPT-5.6	2.12$\pm$0.05	**.326**$\pm$**.002**	.19$\pm$.01	143$\pm$4	2.11$\pm$0.04	**.322**$\pm$**.002**	.23$\pm$.01	188$\pm$4	**2.55**$\pm$**0.03**	.318$\pm$.001	**.27**$\pm$**.00**	**262**$\pm$5
	JPMorgan	Sonnet 5	GPT-5.6	2.88$\pm$0.04	.332$\pm$.001	.13$\pm$.01	95$\pm$3	2.54$\pm$0.08	**.363**$\pm$**.003**	.15$\pm$.01	108$\pm$4	**3.04**$\pm$**0.04**	.353$\pm$.001	**.23**$\pm$**.00**	**183**$\pm$5
		Haiku 4.5	GPT-5.6	2.25$\pm$0.04	**.331**$\pm$**.001**	.20$\pm$.00	150$\pm$3	2.17$\pm$0.04	.324$\pm$.003	.22$\pm$.01	166$\pm$9	**2.71**$\pm$**0.02**	.316$\pm$.001	**.28**$\pm$**.00**	**238**$\pm$7
ML Papers	OpenEndedness	Sonnet 5	GPT-5.6	4.96$\pm$0.02	.469$\pm$.001	.43$\pm$.01	578$\pm$5	5.26$\pm$0.02	.479$\pm$.002	**.47**$\pm$**.01**	690$\pm$8	**5.32**$\pm$**0.02**	**.501**$\pm$**.001**	.44$\pm$.01	**730**$\pm$8
		Haiku 4.5	GPT-5.6	4.26$\pm$0.02	.489$\pm$.001	**.39**$\pm$**.01**	**468**$\pm$9	**5.03**$\pm$**0.07**	.504$\pm$.002	.26$\pm$.02	247$\pm$31	**4.87**$\pm$**0.02**	**.523**$\pm$**.002**	.30$\pm$.01	353$\pm$34
	AI Startup	Sonnet 5	GPT-5.6	**6.21**$\pm$**0.01**	.495$\pm$.001	.46$\pm$.00	658$\pm$7	**6.24**$\pm$**0.03**	.502$\pm$.001	**.48**$\pm$**.00**	737$\pm$8	**6.24**$\pm$**0.02**	**.521**$\pm$**.001**	.45$\pm$.01	**789**$\pm$8
		Haiku 4.5	GPT-5.6	**5.81**$\pm$**0.02**	.516$\pm$.001	**.46**$\pm$**.01**	616$\pm$4	**5.87**$\pm$**0.03**	.507$\pm$.001	**.47**$\pm$**.01**	674$\pm$12	**5.83**$\pm$**0.02**	**.526**$\pm$**.002**	.43$\pm$.01	**720**$\pm$6
Marketing	3D printer	Sonnet 5	GPT-5.6	6.81$\pm$0.04	.219$\pm$.001	.13$\pm$.01	96$\pm$4	6.73$\pm$0.01	.237$\pm$.001	**.18**$\pm$**.01**	**156**$\pm$5	**7.05**$\pm$**0.05**	**.243**$\pm$**.001**	**.18**$\pm$**.00**	**164**$\pm$2
		Haiku 4.5	GPT-5.6	**6.65**$\pm$**0.05**	.190$\pm$.003	.11$\pm$.00	80$\pm$3	6.29$\pm$0.02	**.202**$\pm$**.002**	**.14**$\pm$**.00**	**110**$\pm$3	**6.83**$\pm$**0.03**	**.199**$\pm$**.001**	**.13**$\pm$**.00**	**104**$\pm$3
	Upskilling	Sonnet 5	GPT-5.6	6.47$\pm$0.08	.225$\pm$.001	.09$\pm$.00	61$\pm$1	**6.70**$\pm$**0.06**	**.230**$\pm$**.002**	**.10**$\pm$**.00**	**72**$\pm$3	**6.89**$\pm$**0.06**	**.235**$\pm$**.001**	**.11**$\pm$**.00**	**82**$\pm$3
		Haiku 4.5	GPT-5.6	**6.47**$\pm$**0.03**	.221$\pm$.001	**.08**$\pm$**.01**	**69**$\pm$4	**6.35**$\pm$**0.12**	**.238**$\pm$**.005**	**.08**$\pm$**.00**	58$\pm$3	**6.62**$\pm$**0.10**	**.245**$\pm$**.002**	**.07**$\pm$**.00**	**54**$\pm$3

Table 2: Main results: mean $\pm$ standard error over five seeds (metrics in Section 5.1; †: Goodreads Quality is a human star rating from 1 to 5). Per metric, bold text with blue shading marks the best mean and every cell not significantly different from it (paired two-sided t-test across seeds, $p \geq 0.05$, uncorrected; the three variants share each seed's initial archive). A single highlighted cell is therefore significantly better than both other variants.

justified reading directions under limited time. Candidate embeddings for bandit item selection are precomputed from title+abstract; archive embeddings are computed from title+abstract+memo.

Evaluation. Four metrics summarize a run's final active archive.

- Quality is the mean pointwise score of its entries: the judge score (0–10), or in Goodreads the reader's real star rating (1–5).
- Diversity is $1-$mean pairwise cosine similarity over the entries' embeddings, computed in the original high-dimensional space after L2 normalization (the complement of the mean pairwise similarity used to quantify homogenization by Jiang et al., 2025).
- Coverage is the fraction of occupied cells on a 25×25 grid over a 2D PCA projection (the maps of Figure 2(a)). For each profile, proposer, and trial seed, we fit PCA by full SVD on the union of all three variants' active archives and use shared axis bounds padded by 1% on each side. Coverage is comparable across methods within that grid; separate profiles, proposers, and seeds use separate grids.
- Archive size n is the number of entries that survived admission.

Online evaluation protocol. We treat the pointwise evaluator as the reward oracle of an online problem, not a held-out benchmark judge: each oracle call reveals a score that updates the bandit posterior, and the goal is a high-quality, diverse archive within a fixed budget of calls. Reporting final Quality under the same oracle is then standard: bandit and QD methods report reward and fitness under the objective they optimize.

5.2 Results

Table 2 reports our main results: mean Quality and semantic Diversity of the three variants (Novelty, Novelty+Evo, QD-Rec) on both profiles of every domain, under both proposers.

Main results. QD-Rec gives the best mean Quality in most LLM-judged settings (Table 2), often significantly better than both ablations, and it is never significantly worse than the best variant. The gains are largest on M&A. QD-Rec also builds the largest archive in most settings and often has the best MAP-Elites coverage, so its higher Quality does not come from a smaller, more selective archive. Its semantic Diversity is the best or tied for best in every domain except M&A, where it is slightly lower. On Goodreads, where Quality is the reader's own star rating, the three variants reach similar mean ratings. There, QD-Rec keeps this rating while matching or beating both ablations on archive size, Diversity and coverage.

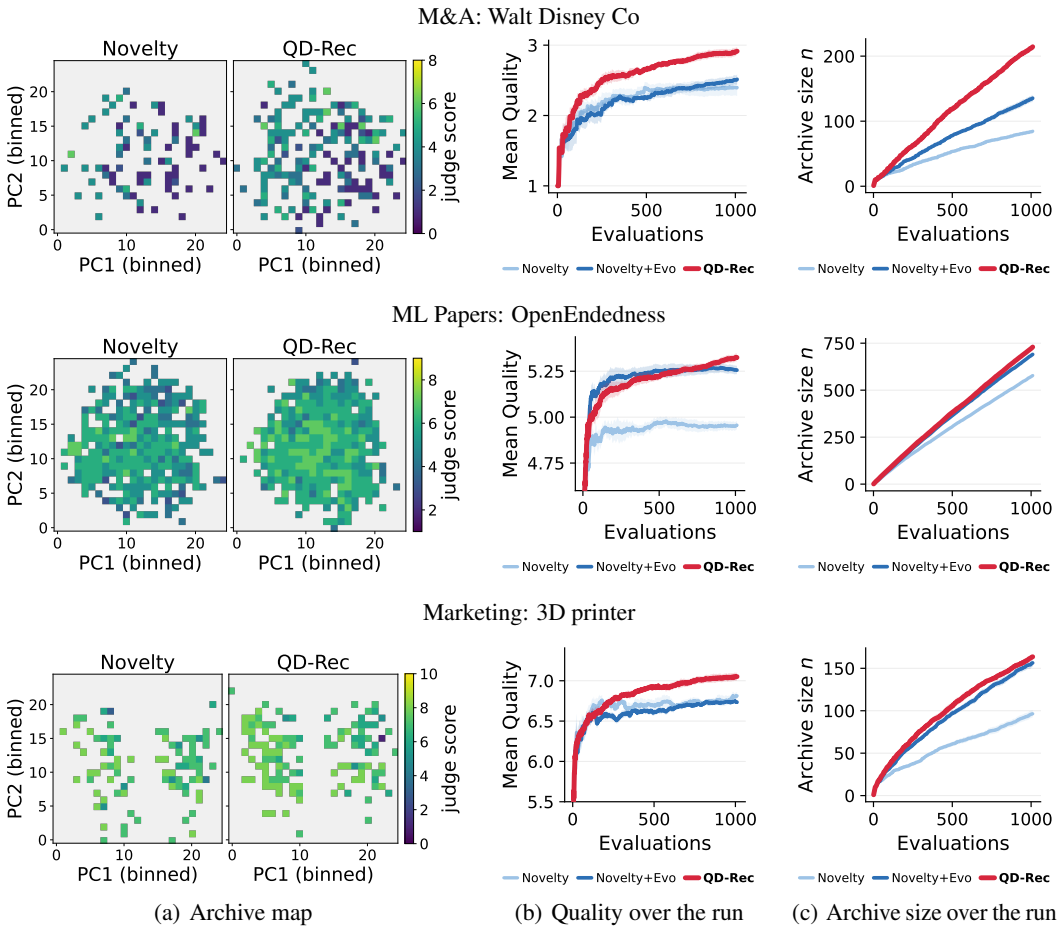

(a) Archive map (b) Quality over the run (c) Archive size over the run

Figure 2: Three profiles, one per row (M&A: Walt Disney Co; ML Papers: OpenEndedness; Marketing: 3D printer). **(a)** MAP-Elites maps of Novelty and QD-Rec (best quality per cell of the Coverage grid, Section 5.1; seed 2). **(b,c)** Mean Quality and size n of the current active archive after every evaluation (`claude-sonnet-5` proposer, mean $\pm$ standard error over five seeds). All eight scenarios are in the appendix: maps in Figure 3 and curves over the run in Figure 5.

Cost. We measure cost as the API spend of a run in US dollars at list prices. It covers every model call: the proposer, the quality judge, the novelty check, the duel, and the embeddings. The three variants make the same calls for each of the $1{,}000$ evaluations. They differ only in the proposer prompt, which includes three archive exemplars in Novelty+Evo and QD-Rec, and in how often a duel is needed. The bandit runs locally and adds no API cost. The runs did not log token usage, so we estimate the cost of each call type by re-sending a sample of 840 archived calls with the same models and settings. We then multiply by the exact number of calls in each run. QD-Rec costs the same as Novelty+Evo, and both cost about $1.3\times$ as much as Novelty. More than 90% of the difference comes from the proposer, because the exemplars roughly triple its input. A single run costs between about \$13 (Novelty on Goodreads with the `claude-haiku-4-5` proposer) and \$62 (QD-Rec on Marketing with the `claude-sonnet-5` proposer). Goodreads is the cheapest domain because its ratings replace the LLM judge and the duel.

Qualitative maps. The left column of Figure 2 shows MAP-Elites-style heatmaps for Novelty and QD-Rec on one profile per domain. In each, QD-Rec covers more of the space than Novelty, most visibly for Disney and the 3D printer. Maps for all eight scenarios, including Goodreads and Novelty+Evo, are in Appendix A, and example recommendations are in Appendix E.

Domain	Profile	Novelty				Novelty+Evo				QD-Rec			
		Qual.	Div.	Cov.	n	Qual.	Div.	Cov.	n	Qual.	Div.	Cov.	n
Goodreads†	series/YA	4.12	.496	.27	282	4.10	.523	**.28**	280	**4.14**	**.533**	**.28**	**287**
M&A	Walt Disney Co	2.05	**.520**	.18	141	1.86	.512	.26	209	**2.30**	.518	**.29**	**270**
ML Papers	OpenEndedness	4.25	.541	**.44**	**538**	**5.05**	.567	.39	402	4.86	**.623**	.34	425
Marketing	3D printer	6.72	.444	.14	99	6.44	**.474**	**.19**	**144**	**6.89**	.470	.17	129

Table 3: Embedding-model sensitivity: one profile per domain re-run with `Qwen3-Embedding-0.6B` embeddings (`claude-haiku-4-5` proposer, seed 0, one run). Bold with blue shading marks the best cell per metric; with a single run, no significance test is applied.

Comparison to a no-learning baseline. To ask whether the online learning and archive machinery earns its cost, we compare against two training-free re-rankers, relevance-MMR (Carbonell & Goldstein, 1998) and a quality-weighted DPP with greedy MAP inference (Chen et al., 2018), on Goodreads, where Quality is a real human star rating. Sweeping each baseline's relevance–diversity weight γ,no setting of either reaches any QD variant on Quality, and they approach QD-Rec's semantic Diversity only at their most diversity-favoring settings. The baselines see no ratings while the QD variants query up to $1{,}000$, so the gap reflects rating feedback and the archive together, not item selection alone. Full setup and the per-setting table are in Appendix B.

Quality and archive size over the run. Figures 2(b) and 2(c) plot the current archive's mean Quality and size n after every evaluation, replayed from each run's add/replace/discard history (all eight scenarios in Figure 5). QD-Rec's Quality keeps rising through the run and ends highest, so its advantage builds up during the run rather than coming from the seeds. Archive size n grows throughout the run as novel hypotheses are added.

Embedding-model sensitivity. The embedding model defines the novelty neighborhoods, the Diversity metric, the MAP-Elites grid and the bandit's item features. We therefore re-run one profile per domain with a different embedding model, `Qwen3-Embedding-0.6B` truncated to the same 512 dimensions (`claude-haiku-4-5` proposer, seed 0, one run with batches of 16 instead of 20, due to the compute limits of the local machine that runs the embedding model; Table 3). The best variant on Quality and on archive size is the same as with the original embeddings on all four profiles. Diversity is higher for every variant, because the Qwen embeddings are spread more evenly. Diversity and Coverage can therefore be compared only within one embedding space, while the Quality and archive-size orderings are unchanged.

6 Conclusion

We introduced QD-Rec, an archive-based quality–diversity method for open-ended recommendation, where the items are chosen with contextual bandits, and the reasons for recommending them are searched open-endedly. Across four domains, QD-Rec gave the best mean Quality in most settings and usually built the largest archive. As a first work along this line, we kept it relatively simple, and several directions follow.

First, our bandit models reward by Bayesian linear regression over fixed embeddings; neural contextual bandits (Zhou et al., 2020; Zhang et al., 2021) could capture nonlinear structure. More radically, the bandit could be discarded altogether in favor of a coding agent that reads the archive directly, reasons over which regions look exhausted and which look promising, and writes its own selection rule, replacing a fixed parametric reward model with a policy that can be inspected and revised in natural language. Second, exemplars are sampled uniformly from the archive, and could instead be selected by quality, recency, or how often they have already served as exemplars (Earle et al., 2026; Zhang et al., 2026a). Third, this work focused on the *outer loop*: which item to evaluate next, and which pairs to retain. The inner hypothesis generator was a single prompted LLM call, and could instead be a full research agent (Shang et al., 2025); we kept it lightweight to show that the outer loop is itself a meaningful lever, and expect the two to compose.

AI USE STATEMENT

We used generative AI tools to help develop and refine research ideas and methodology, implement the experimental and analysis code, interpret results, and draft and revise the manuscript. The authors take responsibility for the final content of this work, including all AI-assisted text, code, analyses, and interpretations.

ETHICS STATEMENT

Generated hypotheses can contain unsupported claims, so any practical use, for example in acquisitions or marketing, needs factual verification and human oversight. The marketing experiments use synthetic personas rather than real customer records.

REPRODUCIBILITY STATEMENT

Following the ICLR 2027 author guidelines,[1] we will share our code with the reviewers and area chairs through a link to an anonymous repository, which we will post as a comment once the discussion forums open. Section 5.1 gives the models and settings, and the appendices describe the data sources, how each candidate pool is built, and the prompts.

REFERENCES

Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. Explore/exploit schemes for web content optimization. In *ICDM*, 2009.

Charu C. Aggarwal. *Recommender Systems: The Textbook*. Springer, 2016.

Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *ICML*, 2013.

Ryan Bahlous-Boldi, Maxence Faldor, Luca Grillotti, Hannah Janmohamed, Lisa Coiffard, Lee Spector, and Antoine Cully. Dominated novelty search: Rethinking local competition in quality-diversity. In *GECCO*, 2025.

Herbie Bradley, Andrew Dai, Hannah Teufel, Jenny Zhang, Koen Oostermeijer, Marco Bellagente, Jeff Clune, Kenneth Stanley, Grégory Schott, and Joel Lehman. Quality-diversity through AI feedback. In *ICLR*, 2024.

Jaime Carbonell and Jade Goldstein. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *SIGIR*, 1998.

Laming Chen, Guoxin Zhang, and Eric Zhou. Fast greedy MAP inference for determinantal point process to improve recommendation diversity. In *NeurIPS*, 2018.

Aaron Dharna, Cong Lu, and Jeff Clune. Foundation model self-play: Open-ended strategy innovation via foundation models. In *RLC*, 2025.

Sam Earle, Kai Arulkumaran, Akarsh Kumar, Andrew Dai, Julian Togelius, and Sebastian Risi. In search of the ingredients of open-endedness: Replicating Picbreeder with large vision-language models. In *GECCO*, 2026.

Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. In *ICLR*, 2024.

Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. In *ICLR*, 2025.

Liwei Jiang, Yuanjun Chai, Margaret Li, Mickel Liu, Raymond Fok, Nouha Dziri, Yulia Tsvetkov, Maarten Sap, and Yejin Choi. Artificial hivemind: The open-ended homogeneity of language models (and beyond). In *NeurIPS*, 2025.

[1] `https://iclr.cc/Conferences/2027/AuthorGuidelines`

Marius Kaminskas and Derek Bridge. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(1):1–42, 2016.

So Kuroki, Taishi Nakamura, Takuya Akiba, and Yujin Tang. Agent skill acquisition for large language models via CycleQD. In *ICLR*, 2025.

Robert Tjarko Lange, Yuki Imajuku, and Edoardo Cetin. ShinkaEvolve: Towards open-ended and sample-efficient program evolution. In *ICLR*, 2026.

Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.

Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.

Joel Lehman, Jonathan Gordon, Shawn Jain, Kamal Ndousse, Cathy Yeh, and Kenneth O Stanley. Evolution through large models. In *Handbook of evolutionary machine learning*, pp. 331–366. Springer, 2023.

Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, 2010.

Yong Liu, Yingtai Xiao, Qiong Wu, Chunyan Miao, Juyong Zhang, Binqiang Zhao, and Haihong Tang. Diversified interactive recommendation with implicit feedback. In *AAAI*, 2020.

Cong Lu, Shengran Hu, and Jeff Clune. Automated capability discovery via foundation model self-exploration. *arXiv preprint arXiv:2502.07577*, 2025.

Sean M McNee, John Riedl, and Joseph A Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI Extended Abstracts*, 2006.

Yev Meyer and Dane Corneil. Nemotron-Personas-USA: Synthetic personas aligned to real-world distributions, June 2025. URL `https://huggingface.co/datasets/nvidia/Nemotron-Personas-USA`.

Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J R Ruiz, Abbas Mehrabian, et al. AlphaEvolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.

Paper Copilot (Jing Yang). paperlists: Processed / cleaned data for Paper Copilot. `https://github.com/papercopilot/paperlists`, 2026. GitHub repository.

Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In *ICML*, 2008.

Sebastian Risi, Yujin Tang, David Ha, and Risto Miikkulainen. *Neuroevolution: Harnessing Creativity in AI Agent Design*. MIT Press, Cambridge, MA, 2025.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco J R Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475, 2024.

Mikayel Samvelyan, Sharath Chandra Raparthy, Andrei Lupu, Eric Hambro, Aram H Markosyan, Manish Bhatt, Yuning Mao, Minqi Jiang, Jack Parker-Holder, Jakob Foerster, et al. Rainbow teaming: Open-ended generation of diverse adversarial prompts. In *NeurIPS*, 2024.

Yu Shang, Peijie Liu, Yuwei Yan, Zijing Wu, Leheng Sheng, Yuanqing Yu, Chumeng Jiang, An Zhang, Fengli Xu, Yu Wang, Min Zhang, and Yong Li. AgentRecBench: Benchmarking LLM agent-based personalized recommender systems. In *NeurIPS Datasets and Benchmarks Track*, 2025.

Yu Song, Shuai Sun, Jianxun Lian, Hong Huang, Yu Li, Hai Jin, and Xing Xie. Show me the whole world: Towards entire item space exploration for interactive personalized recommendations. In *WSDM*, 2022.

Kenneth O. Stanley and Joel Lehman. *Why Greatness Cannot Be Planned: The Myth of the Objective*. Springer, 2015.

Saúl Vargas and Pablo Castells. Rank and relevance in novelty and diversity metrics for recommender systems. In *RecSys*, 2011.

Mengting Wan and Julian McAuley. Item recommendation on monotonic behavior chains. In *RecSys*, 2018.

Koki Wataoka, Tsubasa Takahashi, and Ryokan Ri. Self-preference bias in LLM-as-a-Judge. In *NeurIPS Safe Generative AI Workshop*, 2024.

Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai. Self-taught optimizer (STOP): Recursively self-improving code generation. In *COLM*, 2024.

Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin Gödel Machine: Open-ended evolution of self-improving agents. In *ICLR*, 2026a.

Jenny Zhang, Bingchen Zhao, Wannan Yang, Jakob Foerster, Jeff Clune, Minqi Jiang, Sam Devlin, and Tatiana Shavrina. HyperAgents. *arXiv preprint arXiv:2603.19461*, 2026b.

Weitong Zhang, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural Thompson sampling. In *ICLR*, 2021.

Yongfeng Zhang and Xu Chen. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14(1):1–101, 2020.

Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with UCB-based exploration. In *ICML*, 2020.

Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In *WWW*, 2005.

A Additional results: maps and curves over the run for all scenarios

This appendix collects the qualitative maps for all eight scenarios (Figure 3) and the curves over the run, including semantic Diversity, for all eight scenarios (Figure 5), with weaker-proposer (`claude-haiku-4-5`) counterparts in Figures 4 and 6.

A.1 Qualitative maps: all scenarios

Figure 3 shows the MAP-Elites-style heatmaps of Figure 2(a) for *all* eight scenarios (both profiles of every domain), each with all three variants. Figure 4 repeats the construction for the weaker-proposer ablation (`claude-haiku-4-5`).

A.2 Metrics over the run: all scenarios

Figure 5 plots mean Quality, archive size n, and semantic Diversity against the number of evaluations for all eight scenarios. Figure 6 shows the same curves under the weaker proposer `claude-haiku-4-5`.

Unlike Quality, Diversity changes little after the first $\sim$100 evaluations, while archive size n keeps growing. A likely reason is that admission is decided by the LLM novelty check against nearest neighbors, not by an embedding-distance threshold. New entries then add distinct hypotheses within the region the archive already spans, which raises n but not the mean pairwise distance.

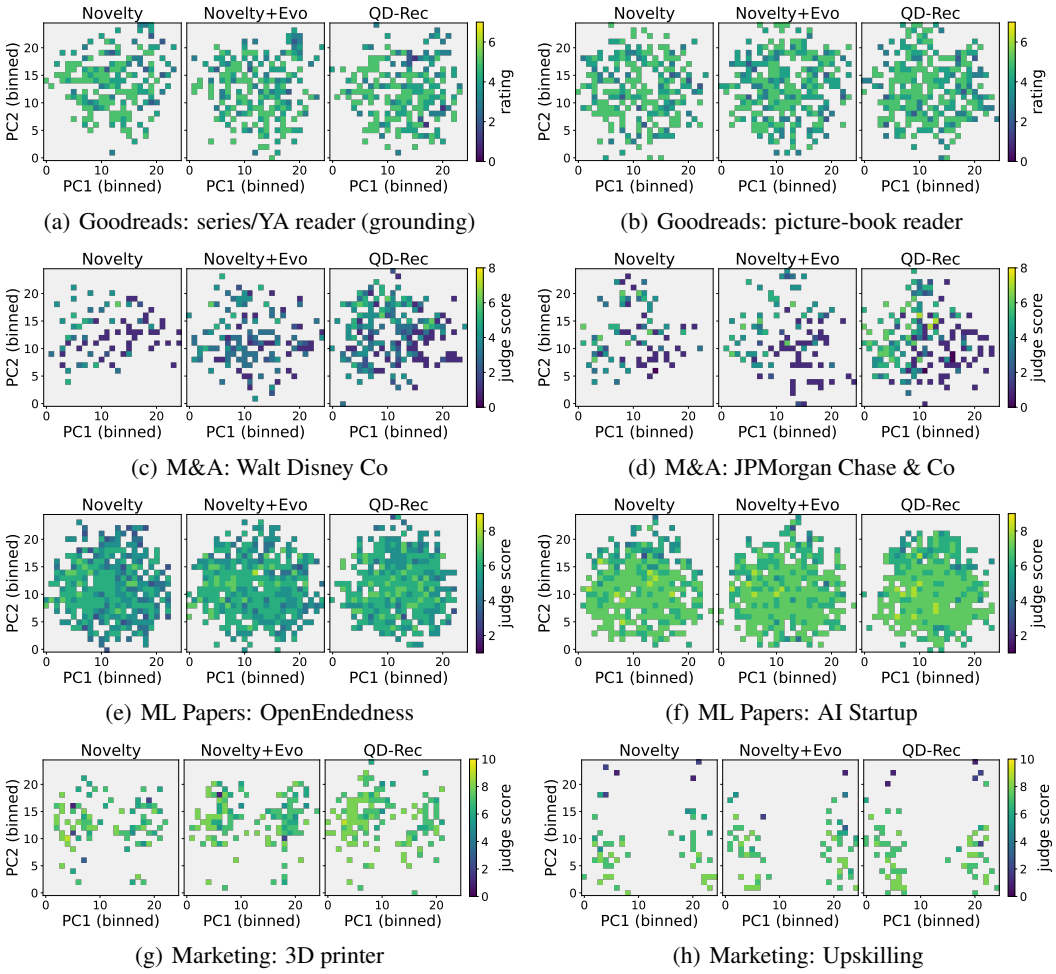

(a) Goodreads: series/YA reader (grounding)

(b) Goodreads: picture-book reader

(c) M&A: Walt Disney Co

(d) M&A: JPMorgan Chase & Co

(e) ML Papers: OpenEndedness

(f) ML Papers: AI Startup

(g) Marketing: 3D printer

(h) Marketing: Upskilling

Figure 3: Qualitative maps for all eight scenarios (both profiles of every domain), complementing Figure 2(a). Same construction: 25×25 MAP-Elites-style heatmaps (best quality per PCA-binned cell), here showing all three variants (Novelty, Novelty+Evo, QD-Rec); color ranges are auto-scaled per domain. All panels use the seed-0 `claude-sonnet-5`-proposer runs. Goodreads uses human star ratings rather than an LLM judge.

B NO-LEARNING BASELINES: RELEVANCE-MMR AND A QUALITY-WEIGHTED DPP

To ask whether the online learning and archive machinery earns its cost, we compare against two training-free re-rankers on Goodreads, where Quality is a real human star rating that neither the learning methods nor the baselines can game. *Relevance-MMR* (Carbonell & Goldstein, 1998) greedily selects items by $\gamma\cos(\text{item},\pi) - (1-\gamma)\max_{j\in S}\cos(\text{item},j)$ over the same content embeddings, where S is the set of items selected so far. *Quality-weighted DPP* (Chen et al., 2018) is the principled counterpart: greedy MAP inference for a determinantal point process maximizing $\gamma\sum_{i\in S}\text{rel}_i + (1-\gamma)\log\det(L_S)$. Here $\text{rel}_i = \cos(\text{item}_i,\pi)$ is the relevance of item i, the same relevance MMR uses. The matrix L_S holds the pairwise cosine similarities of the items in S, so $\log\det(L_S)$ measures the diversity of the set as its log-volume. In both baselines, the weight $\gamma\in[0,1]$ trades relevance against diversity, and $\gamma{=}1$ recovers plain top-K retrieval exactly.

Neither baseline uses an archive, novelty filter, or feedback, and no LLM touches anything that could affect the headline result: selection is pure geometry over precomputed content embeddings, with relevance $\cos(\text{item},\pi)$ rather than an LLM score (scoring the full pool with an LLM is precisely

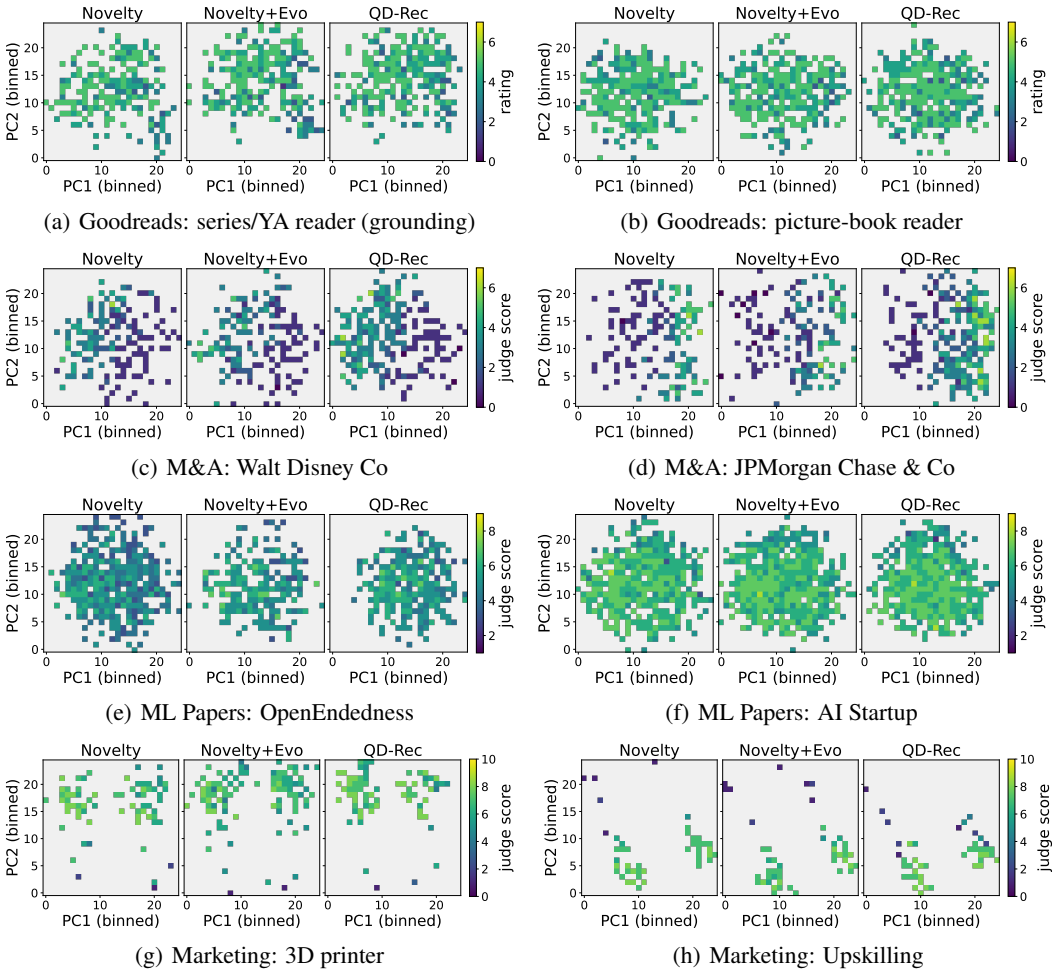

(a) Goodreads: series/YA reader (grounding)

(b) Goodreads: picture-book reader

(c) M&A: Walt Disney Co

(d) M&A: JPMorgan Chase & Co

(e) ML Papers: OpenEndedness

(f) ML Papers: AI Startup

(g) Marketing: 3D printer

(h) Marketing: Upskilling

Figure 4: Qualitative maps for all eight scenarios under the weaker proposer `claude-haiku-4-5`, complementing Figure 3. Same construction: seed-0 runs, all three variants, 25×25 best-quality-per-cell heatmaps, with color ranges auto-scaled per domain within this proposer's runs.

the cost our setting excludes), and Quality is the human star rating, fixed the moment an item is selected. The only LLM calls are measurement plumbing: after selection, each chosen item receives a zero-shot memo (proposer `claude-sonnet-5`) and the same joint item+memo embedding the QD entries carry, so that Diversity and Coverage are computed in the same embedding space as Table 2 and every row emits the same item–hypothesis artifact. These memos cannot influence which items are selected or their ratings.

For each reader, the Coverage grid is fitted jointly to all QD, MMR, and DPP rows using full SVD and the same 1% padding convention. This larger reference pool differs from Table 2, so absolute Coverage need not match across tables. The comparison is not a controlled ablation of item selection, however: the QD variants observe the rating of every item they evaluate (up to 1,000 per run) and use it for replacement and, in QD-Rec, for the bandit, whereas the baselines see none. They are no-feedback geometric references, and the gap measures what rating feedback and the archive machinery buy together. Every run is matched in archive size n to the QD-Rec archive of its reader; selection is deterministic given the embeddings, so rows are single runs (Table 4).

No setting of either baseline reaches any QD variant on Quality. The baselines approach QD-Rec's Diversity only at their most diversity-favoring settings, and a few settings exceed its Coverage. The DPP gives almost the same results as MMR, so the gap is not specific to MMR's heuristic.

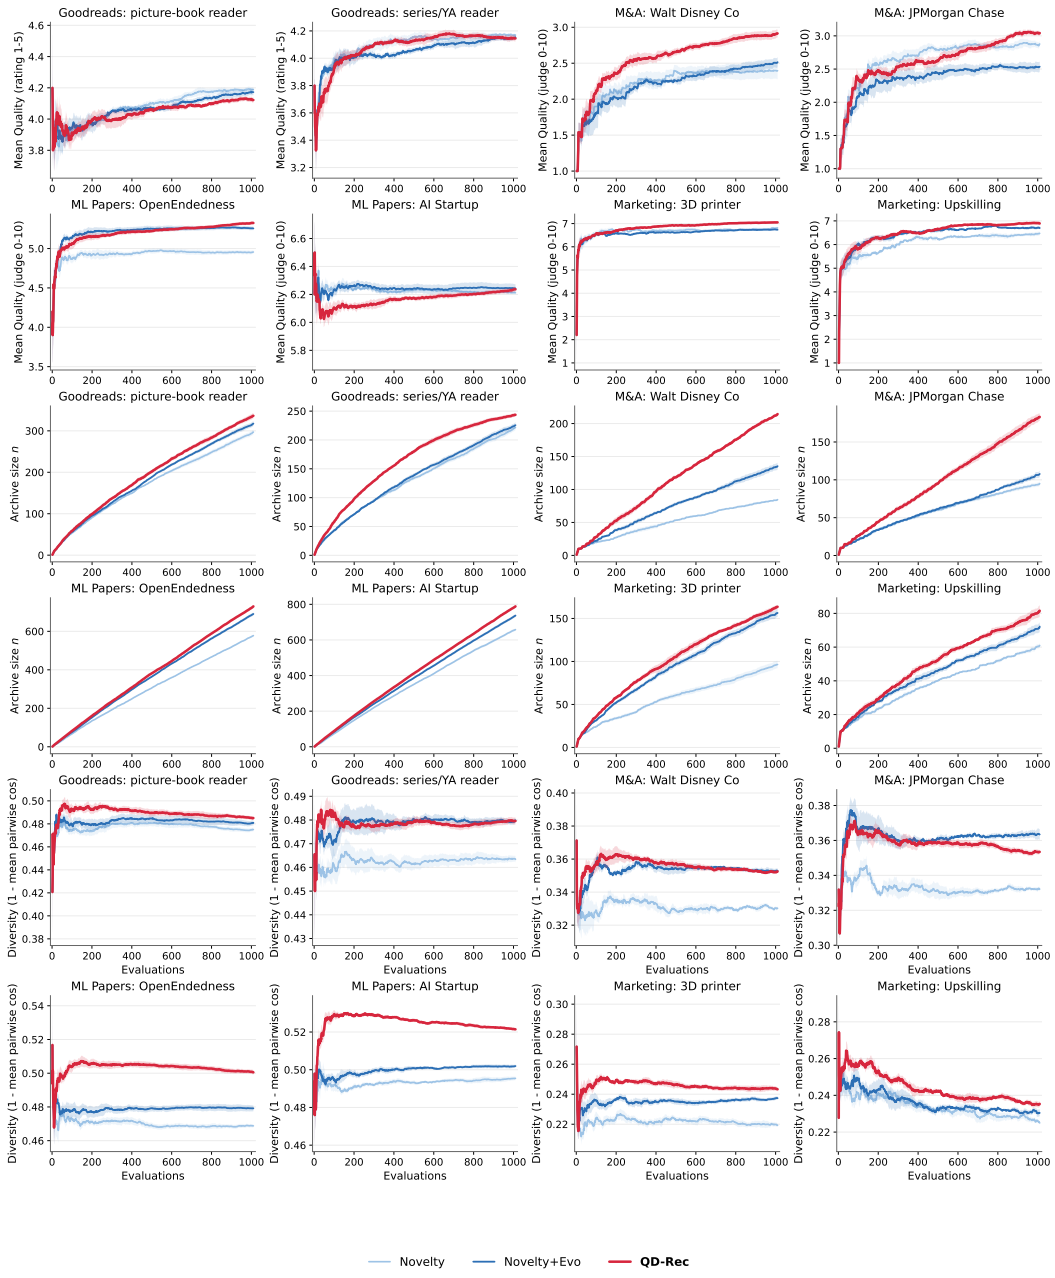

Figure 5: Metrics over the run for all eight scenarios (same construction and `claude-sonnet-5`-proposer runs as Figure 2). Rows 1–2 show mean Quality for the eight scenarios, rows 3–4 the corresponding archive size n, and rows 5–6 semantic Diversity; panel titles identify the scenario. Goodreads Quality is near-flat because it is the reader's fixed 1–5 rating.

C Dataset details

This appendix gives the full dataset construction for the four domains of Section 5.1 (sources, filtering, pool sizes, and the specific profiles used).

Goodreads (real human-feedback grounding). We use the Children subset of the UCSD Goodreads Book Graph (Wan & McAuley, 2018). For each reader, the candidate pool is their

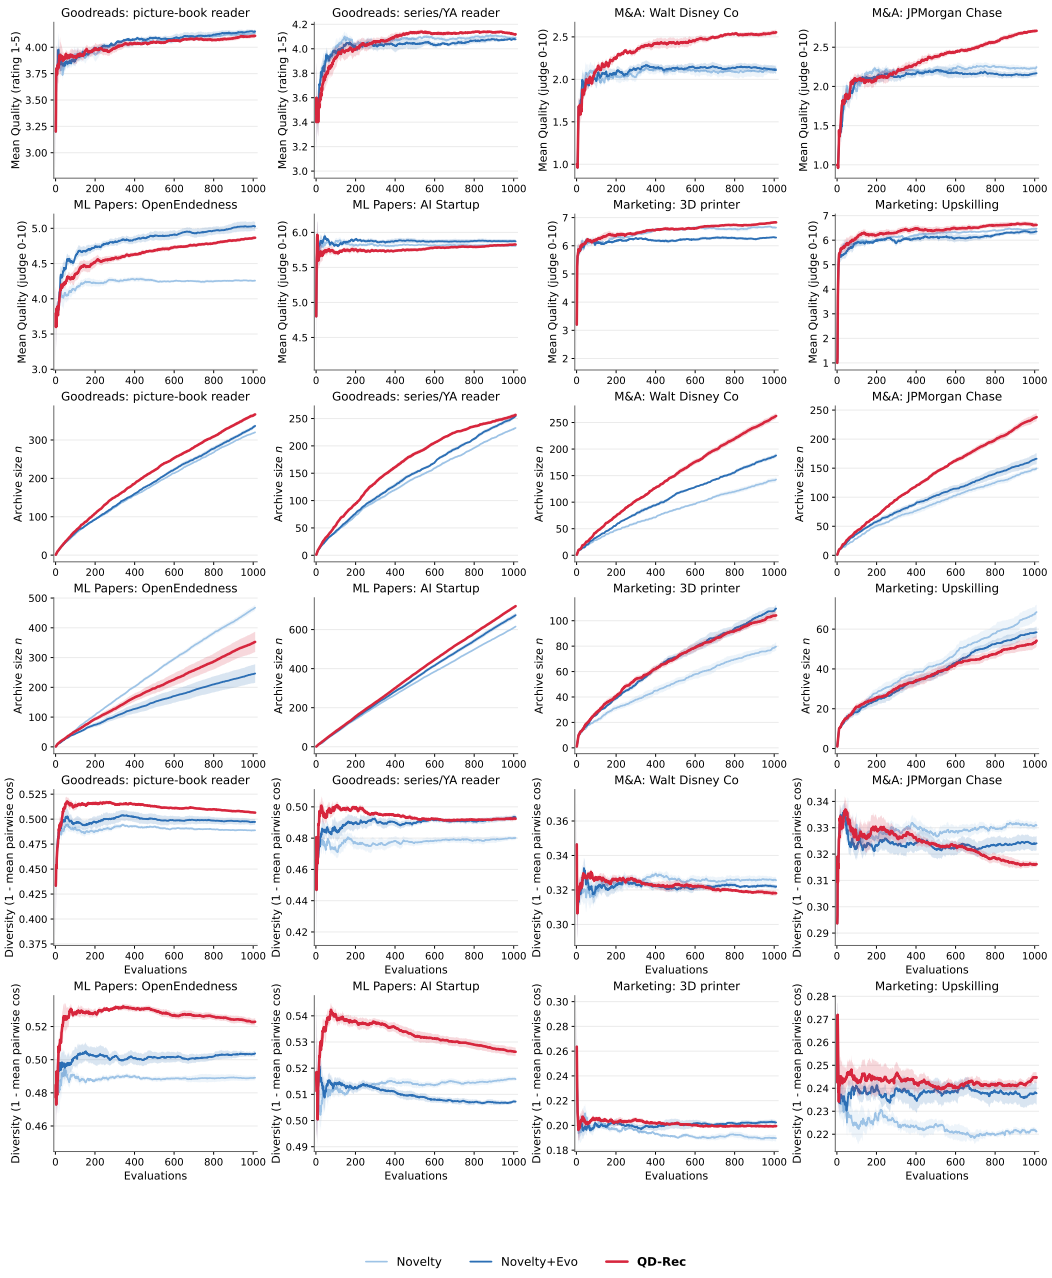

Figure 6: Metrics over the run for all eight scenarios under the weaker proposer (`claude-haiku-4-5`; same construction as Figure 5). Rows 1–2 show mean Quality, rows 3–4 archive size n, and rows 5–6 semantic Diversity; panel titles identify the scenario.

own rated books (minus the few used to construct the profile), so every item carries a ground-truth rating. We select two readers by a reproducible procedure (high rating volume, many written reviews, full use of the 1–5 scale, and broad reading across book types), yielding two distinct reader types: a picture-book reader (1,761 books) and a series/young-adult reader (1,117 books). Goodreads uses the reader's 1–5 rating in place of the LLM judge, so the bandit reward is $q(x)/5$. In the local comparison, a new book replaces its nearest neighbor only if its rating is strictly higher; ties keep the incumbent. Books are selected without replacement in all three variants, including seed books and discarded proposals, because a fixed rating means a new memo for the same book would get the same score.

Method	Qual.	Div.	Cov.	n
Goodreads: picture-book reader				
Novelty	**4.22**	.474	.285	281
Novelty+Evo	4.13	.477	.318	322
QD-Rec	4.12	.484	.326	**344**
MMR γ=1.0	3.69	.408	.312	**344**
MMR γ=0.7	3.70	.420	.326	**344**
MMR γ=0.5	3.69	.442	.328	**344**
MMR γ=0.3	3.68	.478	.318	**344**
DPP γ=1.0	3.69	.408	.323	**344**
DPP γ=0.7	3.71	.453	.314	**344**
DPP γ=0.5	3.71	.477	**.330**	**344**
DPP γ=0.3	3.68	**.484**	.328	**344**
Goodreads: series/YA reader				
Novelty	**4.18**	.467	.234	206
Novelty+Evo	4.16	**.481**	.246	231
QD-Rec	4.12	.481	.248	**234**
MMR γ=1.0	3.31	.401	.229	**234**
MMR γ=0.7	3.48	.420	**.259**	**234**
MMR γ=0.5	3.61	.428	.250	**234**
MMR γ=0.3	3.68	.443	.230	**234**
DPP γ=1.0	3.31	.397	.234	**234**
DPP γ=0.7	3.59	.433	.248	**234**
DPP γ=0.5	3.68	.451	.248	**234**
DPP γ=0.3	3.67	.454	.242	**234**

Table 4: Learning vs. no-learning on the human-rated domain. The three QD variants vs. the training-free relevance-MMR and quality-weighted DPP (γ sweep) re-rankers on Goodreads (real 1–5 star ratings). Every baseline run selects as many items as the QD-Rec archive of its reader holds and writes a zero-shot memo for each (proposer `claude-sonnet-5`); QD rows are the seed-0 runs of the same generation (`claude-sonnet-5` proposer). The baselines see no ratings, so they are no-feedback geometric references rather than a controlled ablation of item selection. Bold marks the best per metric within each reader block.

M&A target discovery. We construct country-specific company universes from Refinitiv/LSEG Workspace exports that combine (i) tabular financial snapshot features (as of 2024-08-31) and (ii) textual business descriptions. We filter to a target country, retain feature columns with at least 0.80 fill rate, and merge description fields into a single company profile. This yields universes of 6,332 USA companies. The profile π is an acquirer selected from the same universe; in our runs we use two USA acquirers, Walt Disney Co and JPMorgan Chase & Co. For QD-Rec, each candidate company has a precomputed embedding used for bandit item selection.

Marketing. We use a subset of the NVIDIA Nemotron-Personas-USA dataset, a large-scale *synthetic customer-profile* corpus aligned to real-world US demographic distributions (Meyer & Corneil, 2025). We use this synthetic dataset as a proxy in lieu of sensitive real-world customer data. Each customer profile includes demographic attributes (e.g., age, sex, marital status, education, occupation, city/state) and multiple profile facets (professional, sports, arts, travel, culinary). To control for regional heterogeneity we restrict to a single state (California) and sample 10,000 customer profiles for experiments. We use two distinct product profiles (a desktop 3D printer and an online upskilling subscription), spanning different purchase contexts (a one-time tool purchase vs. career reskilling). The seed archive is initialized from sampled customer profiles, and the three variants iteratively generate and evaluate targeting hypotheses for additional profiles.

Paper-abstract exploration. We use a fixed collection of paper metadata (title and abstract only) drawn from three recent machine-learning conferences: ICLR 2026, ICML 2026, and NeurIPS 2025. We filter to accepted statuses (Oral, Spotlight, and Poster, plus ICLR's conditional accepts) and retain entries with non-empty identifier, title, and abstract, yielding 17,511 papers (5,358 / 6,341 / 5,812 for ICLR / ICML / NeurIPS respectively). The metadata is obtained from the `papercopilot/paperlists` repository (Paper Copilot (Jing Yang), 2026). We precompute candidate

embeddings from title+abstract for bandit item selection. Archive embeddings are computed from the concatenated title, abstract, and memo (`item_hypothesis` mode) and are used for novelty-neighbor retrieval, semantic Diversity, and Coverage. The three variants generate memos for selected papers and apply the same local novelty-first archive update procedure.

D PROMPT TEMPLATES

This appendix lists the prompts used in our experiments. The three variants use the same generation and judging prompts. They differ only in how the next item is selected and whether the generator is shown archive exemplars. For each domain we show the generator prompt, the judge prompt (0–10 score), the novelty-check prompt, and the pairwise-duel prompt, which runs only when an entry is not novel. Goodreads has only the generator and novelty-check prompts, because it uses human ratings for Quality and replacement. Placeholders such as `{acquirer_profile}` are filled in at run time. A novelty-check or duel response that cannot be parsed never adds an entry to the archive.

D.1 EXAMPLE PROFILES USED IN EXPERIMENTS

Listing 1: Researcher profile used in ML Papers experiments: OpenEndedness Researcher.

```
An ML researcher working on open-ended and quality-diversity methods. Their
    through-line is divergent search, reaching ambitious capabilities through diverse
    intermediate stepping stones rather than direct climbs toward a fixed objective,
    but in practice they bring a handful of separable questions to a paper. Sometimes
    a paper is interesting for how it measures novelty or behavioral diversity without
    hand-designed descriptors, through learned representations and embeddings that
    define what a good behavior space even is; sometimes for its archive and selection
    mechanics, such as what to retain and how local competition and illumination shape
    search; sometimes for its operators, such as mutation, crossover, or using large
    models as variation operators; sometimes for auto-curricula and co-evolution,
    where agents and the environments that challenge them generate their own next
    problems; sometimes for what it says about exploration and the deception of
    objectives, including intrinsic motivation, when direct optimization stalls, and
    what transfers from reinforcement-learning exploration; and sometimes for
    recursive self-improvement, where systems rewrite their own code, operators, or
    objectives so that gains compound, as in self-referential and self-modifying
    agents. They have a skeptical streak toward pure scaling and benchmark-chasing,
    and they actively look to import techniques from adjacent areas into open-ended
    search. They do not expect every paper at a major ML conference to be about
    open-endedness; they look for the specific angle (a method, a measurement, a
    failure mode, an assumption) that connects to their interest.
```

Listing 2: Researcher profile used in ML Papers experiments: AI Startup Researcher.

```
An applied ML researcher at an early-stage AI startup who reads research through the
    lens of practical leverage. They are not only asking, "Could this become a
    product?" but also, "Could this make a product cheaper, safer, faster, more
    reliable, more defensible, easier to evaluate, or easier to specialize?" Their
    interest is in papers that could plausibly affect a 3-to-12-month startup roadmap,
    whether by enabling a new product surface, improving unit economics, reducing
    inference cost, increasing reliability, strengthening safety guarantees, opening a
    vertical market, or creating a data or evaluation moat. When reading papers from
    major ML conferences in particular, they do not expect every paper to be
    immediately productizable. They recognize that much of this work is theoretical,
    exploratory, or several steps removed from company use, so they look for the
    subset of work that hints at future engineering leverage, changes an important
    assumption, improves the efficiency/safety/reliability envelope, or suggests a
    direction that could compound into practical advantage over time.
```

Listing 3: Acquirer profile used in M&A experiments: Walt Disney Co.

```
Company: Walt Disney Co
Location: Burbank, Americas (United States of America)
```

```
Classification:
  - Sector: Consumer Cyclicals / Cyclical Consumer Services
  - Industry: Media & Publishing / Broadcasting
  - Activity: Broadcasting (NEC)

Business Description:
The Walt Disney Company is a diversified worldwide entertainment company. The Company's
    segments include Entertainment, Sports and Experiences. The Entertainment segment
    generally encompasses its non-sports focused global film and episodic content
    production and distribution activities. The lines of business within the
    Entertainment segment along with their business activities include Linear
    Networks, Direct-to-Consumer, and Content Sales/Licensing. The Sports segment
    encompasses its sports-focused global television and direct-to-consumer (DTC)
    video streaming content production and distribution activities. The lines of
    business within the Sports segment include ESPN and Star. The Experiences segment
    includes Parks and Experiences and Consumer Products. Parks and Experiences
    consists of Walt Disney World Resort in Florida, Disneyland Resort in California,
    Disney Cruise Line, and others. Consumer Products includes licensing of its trade
    names, characters, visual, literary and other IP.

Financial Profile (Snapshot: 2024-08-31, Local Currency):
  - Enterprise Value: 210223027404.4 (Current)
  - Historical EV (1y/2y/3y/4y ago): 198446140333.96 / 257149691645.81 /
      387047387650.601 / 296707312395.75
  - Performance: 0.562294155788656% YTD, 4.84709850044702% 2-Year
```

Listing 4: Acquirer profile used in M&A experiments: JPMorgan Chase & Co.

```
Company: JPMorgan Chase & Co
Location: New York City, Americas (United States of America)

Classification:
  - Sector: Financials / Banking & Investment Services
  - Industry: Banking Services / Banks
  - Activity: Banks (NEC)

Business Description:
JPMorgan Chase & Co. is a financial holding company. The Company is engaged in
    investment banking, financial services for consumers and small businesses,
    commercial banking, financial transaction processing and asset management. The
    Company operates through three segments: Consumer & Community Banking (CCB),
    Commercial & Investment Bank (CIB), and Asset & Wealth Management (AWM). Its CCB
    segment offers products and services to consumers and small businesses through
    bank branches, ATMs, digital and telephone banking. Its CIB segment consists of
    banking and payments and markets and securities services, and offers a suite of
    investment banking, lending, payments, market-making, financing, custody and
    securities products and services to a global base of corporate and institutional
    clients. AWM segment offers investment and wealth management solutions. It offers
    multi-asset investment management solutions, retirement products and services,
    brokerage, custody, estate planning, and others.

Financial Profile (Snapshot: 2024-08-31, Local Currency):
  - Enterprise Value: 962614030629.599 (Current)
  - Historical EV (1y/2y/3y/4y ago): 619830439781.86 / 289049648861.661 /
      420986173300.6 / 442185113328.18
  - Performance: 34.4997336773436% YTD, 75.6934494012104% 2-Year
```

Listing 5: Product profile used in Marketing experiments: 3D printer.

```
LAYERCRAFT is a fictional desktop 3D-printer brand in the $500-$900 price range. It is
    aimed at people who already make, fix, or teach things and want to do so with
    their own parts and designs, and it positions itself as "the printer that just
    prints": the machine arrives assembled and calibrated, and the first print starts
    within fifteen minutes of unboxing.
```

```
The lineup comprises three models. "Layercraft Studio" is a compact, fully enclosed
    printer, quiet enough for an apartment or a classroom. "Layercraft Workshop" has a
    larger build volume and prints tougher engineering materials for functional parts.
    "Layercraft Color" holds four spools and switches between them automatically for
    multi-color prints.

Every model has automatic bed leveling, a companion app with a library of
    ready-to-print designs, and a 2-year warranty with US-based support. A design
    marketplace lets owners sell or share their own models, and educators get free
    lesson plans and a classroom management mode. Filament is sold in recyclable
    spools at $20-$30 per kilogram.

Owners use it for different things: replacing a broken part on an appliance or a car
    instead of ordering one; prototyping a product for a small business or a side
    venture; STEM projects with children or students; props, cosplay and
    tabletop-gaming pieces; custom organizers, tools and gifts. It is not for
    everyone: it needs a permanent desk-sized spot, a few evenings to learn basic 3D
    design and slicing, filament at $20-$30 per kilogram, and patience, since a
    typical print takes several hours.

LAYERCRAFT frames a 3D printer as a tool for people who make things, like a sewing
    machine or a table saw: something that pays for itself in the things it makes and
    the things it fixes. By taking the setup and the failed prints out of the
    experience, it aims to make owning one ordinary in that kind of household.
```

Listing 6: Product profile used in Marketing experiments: Upskilling.

```
SkillForge is a fictional online learning platform offered at $89-$249 per month.
    Centered on business skills, technology, and creative fields, it provides content
    aimed at practical, job-ready skill acquisition. Its defining feature is that it
    is designed not as passive video-watching but around the goal of "being able to
    use what you learn at work."

Courses are organized into three learning tracks. The "Business Track" delivers skills
    that make you immediately useful in white-collar roles -- data analysis, project
    management, financial modeling, presentation. The "Tech Track" offers hands-on
    courses from the basics to applications of Python, SQL, cloud, and AI,
    understandable even for non-engineers. The "Creative Track" teaches expressive
    skills usable for side work and career changes, such as UI/UX design, video
    editing, and writing.

Each course is taught by a working expert and includes a weekly live Q&A session, a
    Slack community of fellow learners, and assignment submissions with feedback
    modeled on real work. Progress is visualized on a dashboard, and completion
    certificates can be posted on LinkedIn or a portfolio. A corporate plan adds
    team-level enrollment management and progress reporting.

The intended users are diverse: late-20s-to-30s professionals aiming to upskill in
    their current jobs, women considering reskilling as they plan a return from
    childcare, seniors preparing for a second career after retirement, and employees
    with a side-gig mindset who want a source of income beyond their main job.

With the concept of "making re-learning unremarkable," SkillForge is designed so that
    even busy working adults can steadily build skills with 3-4 hours of study a week.
    With no enrollment fee and flexible, cancel-anytime plans, it lowers the barrier
    to staying with learning and aims to support the career autonomy of working people
    in the US.
```

D.2 M&A TARGET DISCOVERY

Listing 7: M&A few-shot hypothesis generation (system).

```
You are an M&A strategist who generates acquisition hypotheses.
```

Listing 8: M&A few-shot hypothesis generation (user).

```
## Acquiring Company:
{acquirer_profile}

---

## Existing Acquisition Theses (for inspiration only):

{existing_theses}

---

## Target Company to Evaluate:
{target_profile}

---

Generate an acquisition hypothesis explaining why the acquirer should consider
    acquiring the target company.

Focus on:
- Strategic fit with the acquirer's existing business
- Specific synergies (revenue growth, cost savings, capability acquisition)
- Market positioning and competitive advantages
- Financial attractiveness (valuation, growth trajectory, deal sizing relative to
    acquirer)
- Timing considerations (is the target's current valuation/performance favorable?)

Be concrete and specific to the actual companies involved. Reference financial metrics
    when relevant (e.g., enterprise value, growth rates, returns).

Make the strongest plausible case for why the acquirer should acquire this company;
    recommend against it only if the strategic fit is genuinely poor.

Use the existing theses as inspiration. You may:
- **Mutate**: Take a strategic angle from ONE thesis and adapt it to this new target
- **Crossover**: Combine elements from TWO OR MORE theses into a novel rationale
- **Explore**: Propose a completely different strategic direction not yet covered by
    any thesis

Also report which operator you used and the thesis number(s) you built on, using the
"Thesis N" numbers shown above (leave the source list empty for Explore).

Do not copy phrases verbatim from the existing theses.

Provide your response in this EXACT JSON format:

{
  "operation": "mutate" | "crossover" | "explore",
  "source_examples": [thesis numbers you built on, e.g. [2] for mutate or [1, 3] for
      crossover; [] for explore],
  "evolution_note": "ONE sentence on how the example thesis/theses shaped this
      hypothesis, e.g. which angle came from which Thesis N. This is the ONLY field
      that may reference the examples; leave empty for explore.",
  "hypothesis": "2-6 sentences explaining WHY acquiring this company would benefit the
      acquirer. Be specific about strategic fit, synergies, and value creation
      potential. Write it as a STANDALONE rationale, in the same style as one written
      with no examples in front of you: do NOT reference, number, or allude to the
      example theses, and do NOT narrate the operator or how it was derived (put any
      such mention only in evolution_note)."
}
```

Listing 9: M&A pointwise judge (system).

```
You are a critical M&A advisor scoring acquisition hypotheses.
```

Listing 10: M&A pointwise judge (user).

```
Evaluate the following acquisition hypothesis.

## Acquiring Company:
{acquirer_profile}

## Target Company:
{target_profile}

## Acquisition Hypothesis:
{hypothesis}

---

Score how compelling and plausible the strategic case for the acquisition is,
    considering:
- Strategic fit with the acquirer's existing business
- Synergy potential (revenue and cost synergies)
- Feasibility (is this acquisition realistic and executable?)
- Risk factors (integration challenges, market risks, regulatory concerns)
- Value creation potential

A hypothesis that argues against the acquisition, concludes the target is a poor fit,
    or offers only a hollow/implausible rationale is a weak recommendation and must
    score low, even if its reasoning is accurate. Reserve high scores for genuinely
    compelling, specific, plausible theses.

Provide your evaluation in this EXACT JSON format:

{
  "reasoning": "Few sentences summarizing your overall assessment",
  "score": <integer from 0-10>
}

Scoring guide (strength of the recommendation to acquire):
- 0-2: No real case (argues against, no strategic fit, or hollow/implausible rationale)
- 3-4: Weak case (strained or generic rationale; major concerns)
- 5-6: Moderate case (some plausible merit but notable gaps)
- 7-8: Strong case (specific, plausible fit with clear value creation)
- 9-10: Excellent case (highly compelling, specific, plausible thesis)
```

Listing 11: M&A novelty check (system).

```
You are a senior M&A advisor who assesses the novelty of acquisition theses.
```

Listing 12: M&A novelty check (user).

```
## Acquiring Company:
{acquirer_profile}

---

## Existing Acquisition Theses:

{existing_text}

---

## New Thesis to Assess:

{new_text}
```

```
---

Assess whether the NEW thesis above represents a meaningfully different strategic
    direction or value-creation mechanism compared to the existing ones:
- Consider target industry, strategic rationale (vertical/horizontal integration,
    capability acquisition, market access), and synergy type (revenue, cost,
    technology, talent).
- Evaluate strictly: minor variations or superficially different targets with the same
    thesis are NOT novel. Only truly distinct strategic directions count.

Respond in the following JSON format:

{
  "novelty_reasoning": "1-2 sentences explaining whether the new thesis is a genuinely
      novel strategic direction compared to the existing ones.",
  "is_novel": true or false
}
```

Listing 13: M&A blind pairwise duel (system).

```
You are a senior M&A advisor who compares acquisition theses and decides which is
    stronger.
```

Listing 14: M&A blind pairwise duel (user).

```
## Acquiring Company:
{acquirer_profile}

---

## Acquisition Theses to Compare:

{duel_text}

---

Which thesis is stronger overall for the acquirer? Judge objectively, on the strategic
    case alone, by the following criteria:
- Strategic fit with the acquirer's existing business
- Synergy potential (revenue and cost synergies)
- Feasibility and execution risk
- Value creation potential

You must pick exactly one winner; ties are not allowed.

Respond in the following JSON format:

{
  "comparison_reasoning": "2-3 sentences explaining the key differentiators and your
      decision.",
  "winner": "A" or "B"
}
```

D.3 US CUSTOMER TARGETING (MARKETING)

Listing 15: Marketing few-shot hypothesis generation (system).

```
You are a market research expert who generates hypotheses about users' interest in
    products and services.
```

Listing 16: Marketing few-shot hypothesis generation (user).

```
## Target product/service:
{target_product}

---

## Existing example hypotheses (for inspiration only):

{existing_theses}

---

## Persona to evaluate:
{persona_profile}

---

Generate a compelling hypothesis explaining why this persona would be interested in
    this product.

Focus on the following:
- The person's specific needs or pain points that the product can address
- How the product fits into that person's lifestyle and daily life
- The emotional and practical benefits for this specific persona
- The specific features or attributes that resonate with that person

Refer specifically to the actual persona, citing their background, interests, and
    situation.

Make the strongest plausible case for why this persona is worth targeting; conclude
    they are a poor target only if the fit is genuinely poor.

You can use the existing hypotheses as inspiration:
- **Mutate**: Take the reasoning angle of ONE existing hypothesis and adapt it to this
    persona
- **Crossover**: Combine elements of TWO OR MORE hypotheses to create new logic
- **Explore**: Propose a completely different motivation or use case not covered by any
    hypothesis

Also report which operation you used, and which of the "Persona N" numbers above you
    based it on (source_examples). For Explore, set source_examples to an empty list.

Do not copy phrases verbatim from the existing hypotheses.

Respond in the following JSON format:

{
  "operation": "mutate" | "crossover" | "explore",
  "source_examples": [The persona numbers you based it on. E.g., for mutate [2], for
      crossover [1, 3], for explore []],
  "evolution_note": "One sentence explaining how this hypothesis was influenced by the
      existing example hypotheses (which part of which 'Persona N' you incorporated,
      etc.). This is the only field where you may reference the examples. Empty string
      for explore.",
  "hypothesis": "2-6 sentences explaining why this persona would be interested in this
      product (or, only if the fit is genuinely poor, why they would not). Write it as
      a standalone text, as if you had written it without the example hypotheses at
      hand: do not reference, number, or hint at the example hypotheses, and do not
      explain in the body which operation you used or how it was derived (those
      references go only in evolution_note)."
}
```

Listing 17: Marketing pointwise judge (system).

```
You are a critical market research analyst scoring persona-targeting recommendations.
```

Listing 18: Marketing pointwise judge (user).

```
Evaluate the following persona-targeting recommendation.

## Target product/service:
{target_product}

## Persona profile:
{persona_profile}

## Hypothesis (why this persona would or would not be interested):
{hypothesis}

---

Score how compelling and plausible the case for targeting this persona is, considering:
- Relevance: Does the product actually address this persona's needs?
- Specificity: Is the case specific to this persona, or is it generic?
- Plausibility: Is it believable that this persona would actually buy/use the product?
- Depth of insight: Does it demonstrate a genuine understanding of the persona?
- Value alignment: Does the product align with this persona's values, lifestyle, and
    means?

A hypothesis that argues the persona would NOT be interested, concludes the fit is
    poor, or offers only a hollow/implausible rationale is a weak recommendation and
    must score low, even if its reasoning is accurate. Otherwise evaluate strictly and
    honestly; not every recommendation is good.

Provide your evaluation in the following JSON format:

{
  "reasoning": "A concise one-sentence summary of the evaluation",
  "score": <integer from 0 to 10>
}

Scoring guide (strength of the recommendation to target this persona):
- 0-2: No real case (argues against, no fit, generic reasoning, or not believable)
- 3-4: Weak case (weak fit, serious concerns)
- 5-6: Moderate case (some relevance, but lacking specificity or depth)
- 7-8: Strong case (clear fit, manageable concerns)
- 9-10: Excellent case (compelling, specific, and insightful connection)
```

Listing 19: Marketing novelty check (system).

```
You are a senior market research expert who assesses the novelty of persona-product
    hypotheses.
```

Listing 20: Marketing novelty check (user).

```
## Target product/service:
{target_product}

---

## Existing hypotheses:

{existing_text}

---

## New hypothesis to assess:

{new_text}

---
```

```
Assess whether the NEW hypothesis above represents a meaningfully different hypothesis
    compared to the existing ones:
- Does the new hypothesis explore a different motivation or use case for the product?
- Focus on the novelty of the hypothesis, not the difference in users.
- Evaluate strictly: trivial variations are not novel.

Respond in the following JSON format:

{
  "novelty_reasoning": "1-2 sentences explaining whether the new hypothesis is novel
      compared to the existing ones.",
  "is_novel": true or false
}
```

Listing 21: Marketing blind pairwise duel (system).

```
You are a senior market research expert who compares persona-product hypotheses.
```

Listing 22: Marketing blind pairwise duel (user).

```
## Target product/service:
{target_product}

---

## Hypotheses to compare:

{duel_text}

---

Which hypothesis is better overall? Judge objectively, on content quality alone, by the
    following criteria:
- Relevance: How well does the product fit this persona's needs?
- Specificity: Is the hypothesis specific and insightful?
- Plausibility: Would this persona realistically use the product?
- Value alignment: Does the product align with this persona's values, lifestyle, and
    means?

You must pick exactly one winner; ties are not allowed.

Respond in the following JSON format:

{
  "comparison_reasoning": "2-3 sentences explaining the key differentiators and your
      decision.",
  "winner": "A" or "B"
}
```

D.4 PAPER-ABSTRACT EXPLORATION

Listing 23: ML Papers few-shot memo generation (system).

```
You are a research advisor who writes memos on whether papers are worth reading.
```

Listing 24: ML Papers few-shot memo generation (user).

```
## Researcher Profile:
{researcher_profile}

---
```

```
## Existing Example Memos (for inspiration only):

{existing_recommendations}

---

## New Paper:
Title: {title}

Abstract:
{abstract}

---

Write a concise research memo for this researcher.

The abstract is already a summary, so DO NOT re-summarize the paper. Instead, focus on:
- Why this paper could be interesting for THIS researcher (your "hypothesis" of
    relevance)
- One concrete next step (experiment, ablation, or extension) if they read it

Use the existing memos as evolutionary building blocks:
- Mutate: adapt a useful reasoning pattern from ONE example to the new paper and this
    researcher
- Crossover: combine complementary angles from TWO OR MORE examples into a stronger memo
- Explore: propose a genuinely different angle not covered by any example

Also REPORT which operator you used and the example number(s) you built on, using the
"Example N" numbers shown above (leave the source list empty for Explore).

IMPORTANT:
- Do not copy phrases verbatim from the examples.
- Be grounded in the abstract; do not invent unrelated claims.
- Keep it short: 2-6 sentences (one short paragraph).
- Write in plain text; do NOT include LaTeX or math notation.

Format your response as JSON:
{
  "operation": "mutate" | "crossover" | "explore",
  "source_examples": [example numbers you built on, e.g. [2] for mutate or [1, 3] for
      crossover; [] for explore],
  "evolution_note": "ONE sentence on how the example memo(s) shaped this --- e.g. which
      angle came from which Example N. This is the ONLY field that may reference the
      examples; leave empty for explore.",
  "research_notes": "2-6 sentences: start with your relevance hypothesis for this
      researcher (no paper summary), and end with one concrete next step (experiment,
      ablation, or extension) grounded in the abstract. Write the memo as STANDALONE:
      do NOT reference, number, or allude to the example memos, and do NOT narrate the
      operator or how it was derived (put any such mention only in evolution_note)."
}
```

Listing 25: ML Papers pointwise judge (system).

```
You are a critical research advisor who evaluates paper recommendations.
```

Listing 26: ML Papers pointwise judge (user).

```
## Researcher Profile:
{researcher_profile}

---

## Paper Abstract:
{paper_abstract}
```

```
---

## Research Notes:
{research_memo}

---

Evaluate this paper recommendation for the researcher.

Assess how valuable this paper + research memo is for this researcher, considering:
- Relevance to the researcher's stated interests
- Technical quality and novelty
- Quality of insights in the research memo
- Potential for inspiring new research directions

Be rigorous and honest. Most recommendations are NOT outstanding. It might be written
    as though there are deep connections with the researcher's interests, but don't be
    fooled easily.

Calibration (important):
- Use the full 0-10 scale. Avoid score inflation.
- If the memo is generic (could fit many papers) or offers vague directions, score it
    lower even if it "sounds reasonable".
- Prefer concrete, falsifiable, technically grounded next steps over broad
    brainstorming.

Provide your evaluation in this EXACT JSON format:

{
  "reasoning": "2-3 sentences summarizing your assessment",
  "score": <integer from 0-10>
}

Scoring guide:
- 0-2: Poor recommendation (irrelevant/low-signal; likely a waste of time)
- 3-4: Weak recommendation (thin connection; generic memo; few concrete next steps)
- 5-6: OK / mixed (some relevance, but limited novelty or actionable direction;
    unsurprising)
- 7: Good (clear relevance + shows specific, plausible research directions)
- 8: Very good (strong relevance + multiple specific directions; minimal hand-waving)
- 9-10: Exceptional (highly relevant, genuinely novel insight, very interesting)
```

Listing 27: ML Papers novelty check (system).

```
You are a senior research advisor who assesses the novelty of paper recommendations.
```

Listing 28: ML Papers novelty check (user).

```
## Researcher Profile:
{researcher_profile}

---

## Existing Recommendations:

{existing_text}

---

## New Recommendation to Assess:

{new_text}
```

```
---

Assess whether the NEW recommendation above offers a meaningfully different reading
    angle or memo perspective compared to the existing ones.

Consider:
- Does the new memo explore a different relevance hypothesis or research direction for
    this researcher?
- Focus on the angle of the memo, not just the topic of the paper.
- Evaluate strictly: superficial differences in paper topic with the same reading angle
    are NOT novel. Only a genuinely different memo framing or research direction
    counts.

Respond in the following JSON format:

{
  "novelty_reasoning": "1-2 sentences explaining whether the new recommendation is a
      genuinely novel reading angle or memo framing compared to the existing ones.",
  "is_novel": true or false
}
```

Listing 29: ML Papers blind pairwise duel (system).

```
You are a senior research advisor who compares paper recommendations and decides which
    is more valuable.
```

Listing 30: ML Papers blind pairwise duel (user).

```
## Researcher Profile:
{researcher_profile}

---

## Recommendations to Compare:

{duel_text}

---

Which recommendation is more valuable for this researcher? Judge objectively, on memo
    content alone, by the following criteria:
- Relevance to the researcher's stated interests
- Specificity and depth of the reading hypothesis
- Quality and concreteness of the proposed research directions
- Potential to inspire new work

You must pick exactly one winner; ties are not allowed.

Respond in the following JSON format:

{
  "comparison_reasoning": "2-3 sentences explaining the key differentiators and your
      decision.",
  "winner": "A" or "B"
}
```

D.5 GOODREADS-CHILDREN (GROUNDING)

Goodreads has no LLM quality judge (Quality is the reader's held-out star rating), so its prompt set is smaller: memo generation and the novelty check. For a non-novel book, replacement compares its human rating directly with that of its nearest neighbor: a strictly higher rating replaces the neighbor, while a tie or lower rating retains the incumbent. The bandit uses rating/5 for both seed initialization

and subsequent updates, while reported Quality remains on the 1–5 scale. Previously tried books, including seeds, are excluded from subsequent selection.

Listing 31: Goodreads few-shot memo generation (system).

```
You are a reading advisor who writes memos on whether a reader would enjoy a book.
```

Listing 32: Goodreads few-shot memo generation (user).

```
## Reader Profile:
{profile}

---

## Existing Example Memos (for inspiration only):

{existing_recommendations}

---

## New Candidate Book:
Title: {title}

Description:
{description}

---

Write a concise memo on why this reader might (or might not) enjoy this book.

The description already summarizes the book, so DO NOT re-summarize it. Instead, focus
    on:
- Why THIS reader in particular might enjoy (or not enjoy) this book, grounded in the
    specific tastes, themes, and pet peeves visible in their profile.
- What about the book connects to or clashes with what they have rewarded or panned
    before.

Use the existing memos as evolutionary building blocks:
- Mutate: adapt a useful reasoning angle from ONE example to the new book and this
    reader
- Crossover: combine complementary angles from TWO OR MORE examples
- Explore: propose a genuinely different angle not covered by any example

Also REPORT which operator you used and the example number(s) you built on, so the
reasoning can be traced. Use the "Example N" numbers shown above; for Explore, leave
the source list empty.

IMPORTANT:
- Do not copy phrases verbatim from the examples.
- Be grounded in the book description and the reader's profile; do not invent plot
    details not in the description.
- Keep it short: 2-6 sentences (one short paragraph).
- Do not mention or guess a star rating.

Format your response as JSON:
{
  "operation": "mutate" | "crossover" | "explore",
  "source_examples": [example numbers you built on, e.g. [2] for mutate or [1, 3] for
      crossover; [] for explore],
  "evolution_note": "ONE sentence on how the example memo(s) shaped this --- e.g. which
      angle came from which Example N. This is the ONLY field that may reference the
      examples; leave empty for explore.",
  "notes": "2-6 sentences: start directly with your relevance hypothesis grounded in
      this reader's demonstrated tastes; do not re-summarize the plot. Write it as
      STANDALONE: do NOT reference, number, or allude to the example memos, and do NOT
```

```
      narrate the operator or how it was derived (put any such mention only in
      evolution_note)."
}
```

Listing 33: Goodreads novelty check (system).

```
You are a reading advisor who assesses whether a new book recommendation adds genuine
    variety to a reader's shelf.
```

Listing 34: Goodreads novelty check (user).

```
## Reader Profile:
{profile}

---

## Recommendations on the shelf:

{recommendations_text}

---

The NEW recommendation is: **{new_recommendation_label}**

Assess whether the NEW recommendation is meaningfully different from the existing ones
    --- i.e. whether it adds genuine variety to a diverse shelf for this reader.

Important:
- Judge by the kind of book and the kind of appeal, not by quality.
- Be strict: a superficially different title that offers the same kind of reading
    experience as an existing one is NOT novel.
- Do not use or guess any ratings/scores; judge by content and appeal only.

Provide your response in this EXACT JSON format:

{
  "novelty_reasoning": "1-2 sentences explaining why the new recommendation is or isn't
      novel.",
  "is_novel": true or false
}
```

E EXAMPLE QD-REC RECOMMENDATIONS

The single highest-scoring active QD-Rec entry per applied scenario (proposer `claude-sonnet-5`, judge `gpt-5.6-terra`, seed 0; ties broken by earliest evaluation): the chosen item's profile and the generated hypothesis, verbatim apart from the replacement of non-ASCII characters.

M&A TARGET DISCOVERY: WALT DISNEY CO (ACQUIRER)

*Acquirer profile: Walt Disney Co. Hypothesis = acquisition thesis (judge 0–10). Selected entry: judge score **6/10**, generated by* `explore`*.*

Listing 35: Chosen item.

```
Company: Magnite Inc
Location: New York City, United States of America
Sector / Industry: Consumer Cyclicals / Advertising & Marketing
Activity: Advertising & Marketing (NEC)
Enterprise Value (2024-08-31): 2172629799.04
YTD / 2Y total return: 47.6445396134457 / 30.2171860206116

Business Description:
```

```
Magnite, Inc. is an independent sell-side advertising company. The Company provides
    technology solutions to automate the purchase and sale of digital advertising
    inventory. Its platform features applications and services for sellers of digital
    advertising inventory, or publishers, that own and operate connected television
    (CTV) channels, applications, websites and other digital media properties, to
    manage and monetize their inventory; applications and services for buyers,
    including advertisers, agencies, agency trading desks, and demand side platforms,
    to buy digital advertising inventory; and a transparent, independent marketplace
    that brings buyers and sellers together and facilitates intelligent decision
    making and automated transaction execution at scale. Its streaming sell-side
    advertising platform (SSP) and ad server offers CTV sellers a holistic solution to
    manage and monetize their entire portfolio of CTV ad inventory, across both
    programmatic and direct-sold video inventory.
```

Listing 36: QD-Rec generated hypothesis (score 6/10).

```
Disney's DTC ambitions (Disney+, Hulu, ESPN app) increasingly depend on monetizing
    ad-supported tiers efficiently, and Magnite is the leading independent sell-side
    platform for CTV advertising, giving Disney direct control over the programmatic
    infrastructure that currently sits with third-party intermediaries taking a cut of
    its ad revenue. At a $2.17B enterprise value -- up sharply from $1.55B a year ago
    but still far below its $4.47B level three years prior -- Magnite is available at
    a discount to its historical peak while riding the secular growth of CTV ad spend,
    making this a timely, financially digestible acquisition at roughly 1% of Disney's
    $210B EV. Owning the sell-side stack would let Disney internalize yield
    optimization, unify ad inventory management across its Entertainment, Sports
    (ESPN), and Experiences-linked digital properties, and reduce reliance on external
    SSPs like Google's ad exchange or The Trade Desk integrations, directly improving
    ad-tier margins as Disney+ and Hulu scale their advertising businesses. Magnite's
    existing publisher relationships across the broader CTV ecosystem also offer
    Disney a data and demand-aggregation advantage that could be monetized beyond its
    own properties, turning a cost-center into a new B2B revenue line. The main risk
    is channel conflict -- other publishers may be wary of routing inventory through a
    platform owned by a competing content owner -- but Disney could preserve
    independence via a separately operated subsidiary while still capturing the
    strategic control and margin upside internally.
```

M&A TARGET DISCOVERY: JPMORGAN CHASE & CO (ACQUIRER)

*Acquirer profile: JPMorgan Chase & Co. Hypothesis = acquisition thesis (judge 0–10). Selected entry: judge score **7/10**, generated by* `mutate`.

Listing 37: Chosen item.

```
Company: Hamilton Lane Inc
Location: Conshohocken, United States of America
Sector / Industry: Financials / Investment Management & Fund Operators
Activity: Investment Management
Enterprise Value (2024-08-31): 8483299524.44
YTD / 2Y total return: 35.8664681857622 / 146.341522708497

Business Description:
Hamilton Lane Incorporated is a private markets investment company providing solutions
    to institutional and private wealth investors around the world. The Company offers
    a variety of investment solutions to address its clients' needs across a range of
    private markets, including private equity, private credit, real estate,
    infrastructure, real assets, growth equity, venture capital and impact. These
    solutions are constructed from a range of investment types, including primary
    investments in funds managed by third-party managers, direct investments alongside
    such funds and acquisitions of secondary stakes in such funds, with a number of
    its clients utilizing multiple investment types. These solutions are offered in a
    variety of formats, covering some or all phases of private markets investment
    programs, which include customized separate accounts, specialized funds, advisory
    services, distribution management, and reporting, monitoring, data and analytics.
```

Listing 38: QD-Rec generated hypothesis (score 7/10).

```
Hamilton Lane's $8.48B enterprise value private markets platform -- spanning primary
    fund commitments, direct co-investments, and secondaries across private equity,
    credit, infrastructure and real assets -- fills a structural capability gap in
    JPMorgan's Asset & Wealth Management segment, which currently lacks the scaled,
    institutional-grade alternatives manufacturing that rivals like Blackstone and
    Apollo have used to capture both institutional and, increasingly, private wealth
    flows. JPMorgan's global private bank and AWM distribution network, serving
    ultra-high-net-worth and institutional clients across dozens of countries, is
    exactly the channel Hamilton Lane needs to scale its 'private wealth investor'
    ambitions beyond its current advisory and fund-of-funds base, creating an
    immediate cross-sell opportunity to embed Hamilton Lane's data, analytics, and
    portfolio construction tools into Chase Private Client and J.P. Morgan Wealth
    Management offerings. At roughly 0.88% of JPMorgan's $962.6B enterprise value, the
    deal is sizeable but absorbable, and importantly the target's EV has grown from
    $5.3B a year ago to $8.48B now (+60% YTD-adjacent, +146% over two years),
    reflecting genuine secular tailwinds in private markets democratization rather
    than one-off volatility. Acquiring now lets JPMorgan capture Hamilton Lane's
    technology-enabled secondaries and analytics franchise before further multiple
    expansion, while adding recurring, capital-light management and advisory fee
    revenue that diversifies AWM away from pure asset-based fee compression risk.
```

PAPER-ABSTRACT EXPLORATION: OPEN-ENDEDNESS RESEARCHER

*Researcher profile: open-endedness. Hypothesis = research memo (judge 0–10). Selected entry: judge score **7/10**, generated by* `mutate`.

Listing 39: Chosen item.

```
Title: Dropping Just a Handful of Preferences Can Change Top Large Language Model
    Rankings
Year: 2026

Abstract:
We propose a method for evaluating the robustness of widely used LLM ranking
    systems---variants of a Bradley--Terry model---to dropping a worst-case very small
    fraction of preference data. Our approach is computationally fast and easy to
    adopt. When we apply our method to matchups from popular LLM ranking platforms,
    including Chatbot Arena and derivatives, we find that the rankings of
    top-performing models can be remarkably sensitive to the removal of a small
    fraction of preferences; for instance, dropping just 0.003% of human preferences
    can change the top-ranked model on Chatbot Arena. Our robustness check identifies
    the specific preferences most responsible for such ranking flips, allowing for
    inspection of these influential preferences. We observe that the rankings derived
    from MT-bench preferences are notably more robust than those from Chatbot Arena,
    likely due to MT-bench's use of expert annotators and carefully constructed
    prompts. Finally, we find that neither rankings based on crowdsourced human
    evaluations nor those based on LLM-as-a-judge preferences are systematically more
    sensitive than the other.
```

Listing 40: QD-Rec generated hypothesis (score 7/10).

```
This is relevant because quality-diversity and open-ended systems that rank policies or
    behaviors via pairwise comparisons (tournament selection, LLM-as-judge fitness,
    human-in-the-loop archives) often implicitly assume Bradley-Terry-style
    aggregation is stable, but this paper shows that top rankings can flip from
    removing a vanishingly small fraction of comparisons, which is a serious warning
    for any illumination or local-competition scheme that relies on such rankings to
    decide what enters or dominates an archive niche. The influence-diagnostic method
    for identifying which specific preferences drive ranking flips is also a
    transferable tool: it could be repurposed to audit which pairwise comparisons are
    secretly determining elite status in a MAP-Elites-style archive using
    preference-based fitness. A concrete next step would be to take an archive built
```

```
    with pairwise preference-based selection (e.g., LLM-judge tournaments deciding
    niche occupants) and run this robustness check to see how many comparisons would
    need to flip to change which solutions are retained as elites, then compare that
    sensitivity across different aggregation schemes (raw win-rate, Bradley-Terry,
    Elo) to see if some are structurally more robust for QD selection.
```

PAPER-ABSTRACT EXPLORATION: AI-STARTUP RESEARCHER

*Researcher profile: AI startup. Hypothesis = research memo (judge 0–10). Selected entry: judge score **8/10**, generated by* `crossover`.

Listing 41: Chosen item.

```
Title: Multilingual Unlearning in LLMs: Transfer, Dynamics, and Reversibility
Year: 2026

Abstract:
Large language models (LLMs) can memorize sensitive facts, motivating *unlearning*
    methods that remove targeted knowledge without costly retraining. However,
    unlearning research remains heavily English-centric. We study multilingual
    unlearning by extending the TOFU benchmark to five languages, and fine-tune,
    unlearn and query our models with different permutations of languages. We find
    that unlearning transfer -- the ability of an unlearned model to "forget" facts in
    languages other than the unlearning language -- is highly variable: e.g., it is
    strongest between languages sharing scripts and families, and we show that the
    *unlearning language* predicts which *query languages* are most likely to yield
    the strongest transfer. Layer-wise analysis reveals that unlearning leaves the
    shared cross-lingual latent space largely intact in early layers, instead
    operating primarily in later decoding layers. This suggests that unlearning does
    not truly erase knowledge, but rather induces superficial suppression. Exploiting
    this structure, a single inference-time steering direction reverses much of this
    suppression across languages, recovering 50% (Qwen) and 90% (Gemma) of the
    unlearned knowledge.
```

Listing 42: QD-Rec generated hypothesis (score 8/10).

```
If your product touches GDPR/CCPA-style right-to-be-forgotten requirements or
    memorized-PII removal in a multilingual deployment (support bots, international
    user data, etc.), this paper is a serious red flag for a compliance feature you
    might assume is 'solved': it shows current unlearning methods only suppress
    knowledge in later decoding layers rather than erasing it from the shared
    cross-lingual representation, meaning a user querying in a different language than
    the one used for unlearning can often still extract the 'forgotten' fact, and
    worse, a single steering vector can undo most of the suppression entirely. This
    means any unlearning-as-a-service or data-deletion pipeline built on current
    methods could be legally and reputationally fragile, especially for multilingual
    products. Concrete next step: run their extended TOFU benchmark (or a similar
    synthetic sensitive-fact set) against whatever unlearning or deletion method you
    currently use or plan to use, testing cross-lingual query leakage and
    susceptibility to the same steering-based reversal, to decide whether you need a
    stronger erasure method (e.g. weight editing at earlier layers) before advertising
    deletion guarantees to customers.
```

MARKETING: 3D PRINTER (LAYERCRAFT)

*Product: LAYERCRAFT desktop 3D printer. Hypothesis = targeting hypothesis (judge 0–10). Selected entry: judge score **9/10**, generated by* `crossover`.

Listing 43: Chosen item.

```
Clyde Wasserman is a pragmatic, competitively driven aerospace veteran who blends
    meticulous engineering with spontaneous community gatherings, a lingering worry
    about standards, and a love for hands-on hobbies, and they keep life both
    structured and delightfully unpredictable.
```

```
Age / Sex: 58 / Male Marital status: married_present
Education: high_school
Occupation: aerospace_engineer Location: Wilton, CA

Career goals & ambitions:
Clyde aims to transition from day-to-day engineering tasks to a senior technical
    consultant role where he can codify his practical knowledge into standardized
    procedures for the aerospace industry. He wants to lead a small, high-performance
    team that focuses on developing reliable, low-cost satellite components for
    commercial launch providers. As he approaches retirement, Clyde is also looking to
    secure a legacy by publishing a guide on pragmatic aerospace maintenance for
    technicians without advanced degrees. While he values stability and
    predictability, his competitive nature pushes him to take on challenging projects
    that can demonstrate measurable performance improvements and position him as a
    go-to expert in the Sacramento aerospace community.

Professional persona:
Clyde Wasserman is an aerospace engineer with three decades of hands-on experience in
    structural design, propulsion integration, and launch-vehicle maintenance, whose
    methodical approach to risk assessment and FAA compliance is tempered by a relaxed
    attitude toward paperwork, a competitive drive that pushes teammates toward higher
    performance, and they naturally mentor technicians through practical,
    no-degree-required solutions.
```

Listing 44: QD-Rec generated hypothesis (score 9/10).

```
Clyde already designs in SolidWorks and CATIA and spends weekends building and
    launching model rockets with his hobby club, so the Layercraft Workshop's larger
    build volume and engineering-grade materials give him a way to print custom nose
    cones, fin brackets, payload bays, and jig fixtures instead of hand-shaping wood
    or waiting on machined parts, and the marketplace lets him share or sell those
    designs with fellow rocketry enthusiasts at regional maker fairs where he already
    showcases his builds. As a volunteer in the county's STEM outreach program
    mentoring high-schoolers, the free educator lesson plans and classroom management
    mode would let him turn 3D printing into a hands-on teaching tool, echoing the
    practical, no-degree-required approach he prides himself on with junior
    technicians. His three decades of hands-on troubleshooting mean he won't be
    intimidated by slicing or calibration, but the 'assembled and printing in fifteen
    minutes' promise still appeals to his relaxed attitude toward fussy setup and
    paperwork, letting him get straight to the fabrication he actually enjoys. It also
    slots naturally into his woodshop-adjacent maker space, complementing his custom
    furniture projects with precision plastic parts he can't easily cut or carve.
    Given his competitive drive to demonstrate measurable performance improvements,
    he'd likely enjoy iterating on rocket components to squeak out better flight
    results, making him a strong, high-engagement fit for the Workshop model in
    particular.
```

MARKETING: UPSKILLING (SKILLFORGE)

*Product: SkillForge learning platform. Hypothesis = targeting hypothesis (judge 0–10). Selected entry: judge score **8/10**, generated by* `crossover`.

Listing 45: Chosen item.

```
Tha Nguyen, a methodical yet sociable designer whose anxiety drives meticulous
    spreadsheets, can be found juggling design drafts, a simmering pot of broth, and a
    karaoke mic in the same evening, while their habit of double-checking emails late
    into the night sometimes bleeds into sleep.

Age / Sex: 31 / Female Marital status: never_married
Education: some_college
Occupation: graphic_designer Location: Long Beach, CA
```

```
Career goals & ambitions:
Tha aspires to progress from senior graphic designer to art director within the next
    five years, eventually taking on a creative director role where she can shape
    brand narratives and lead multidisciplinary design teams. She aims to build a
    portfolio that showcases culturally resonant visual storytelling, especially for
    Asian-American brands and social impact projects. She also plans to mentor junior
    designers, expand her expertise in motion graphics and interactive media, and
    possibly launch her own boutique design studio that blends contemporary design
    with Vietnamese cultural motifs.

Professional persona:
Tha Nguyen, a 31-year-old senior graphic designer in Los Angeles, blends expert mastery
    of the Adobe Creative Suite, digital illustration, branding, and UI/UX prototyping
    with a methodical project-management style and a naturally talkative, gregarious
    demeanor that makes them the catalyst for design meetups and client workshops,
    driving their swift ascent toward an art-director role.
```

Listing 46: QD-Rec generated hypothesis (score 8/10).

```
Tha's explicit five-year goal of moving from senior designer to art director, and
    eventually launching her own boutique studio blending Vietnamese motifs with
    contemporary design, maps directly onto SkillForge's Creative Track for expanding
    into motion graphics and interactive media, and its Business Track for the
    project-management, client communication, and financial-modeling skills she'll
    need to run a studio and lead teams. Her already-methodical, spreadsheet-driven
    anxiety-management style suggests she'd thrive with SkillForge's progress
    dashboard and structured weekly cadence, turning her nervous energy into visible,
    trackable growth. As someone who is gregarious and already organizes design
    meetups, she'd likely become an active voice in the Slack community and live Q&As,
    mentoring less experienced designers much as she hopes to mentor junior staff in
    her career plan. The 3-4 hour weekly commitment fits neatly around her busy but
    disciplined schedule of hiking, cooking pho, and karaoke nights, since she's shown
    she can juggle multiple demanding pursuits simultaneously. Given her clear
    ambition, existing design fluency, and appetite for structured self-improvement,
    Tha represents a strong-fit user who would use SkillForge specifically to
    formalize UI/UX and motion-graphics expertise while building the business acumen
    to eventually run her own creative venture.
```