
MetaSynth: Multi-Agent Metadata Generation from Implicit Feedback in Black-Box Systems

Shreeranjani Srirangamsridharan^{*†}

Ali Abavisani^{*}

Reza Yousefi Maragheh^{*‡}

Ramin Giahi

Kai Zhao

Jason Cho

Sushant Kumar

Abstract

Meta titles and descriptions strongly shape engagement in search and recommendation platforms, yet optimizing them remains challenging. Search engine ranking models are black boxes environments, explicit labels are unavailable, and feedback such as click-through rate (CTR) arrives only post-deployment. Existing template, LLM, and retrieval-augmented approaches either lack diversity, hallucinate attributes, or ignore whether candidate phrasing has historically succeeded in ranking. This leaves a gap in directly leveraging implicit signals from observable outcomes. We introduce MetaSynth, a multi-agent retrieval-augmented generation framework that learns from implicit search feedback. MetaSynth builds an exemplar library from top-ranked results, generates candidate snippets conditioned on both product content and exemplars, and iteratively refines outputs via evaluator-generator loops that enforce relevance, promotional strength, and compliance. On both proprietary e-commerce data and the Amazon Reviews corpus, MetaSynth outperforms strong baselines across NDCG, MRR, and rank metrics. Large-scale A/B tests further demonstrate +10.26% CTR and +7.51% clicks. Beyond metadata, this work contributes a general paradigm for optimizing content in black-box systems using implicit signals.

1 Introduction

Search and recommendation systems are central to online discovery, yet their internal ranking functions are typically opaque [11, 14]. These systems disclose little about how items are ordered, but their outputs such as search engine result pages (SERPs) consistently reflect stable preferences in the way information is phrased and structured. Among the most impactful elements are meta titles and descriptions, the short snippets displayed to users at decision time, which strongly influence click behavior and downstream traffic[10]. Optimizing these snippets represents a high-leverage intervention for organic acquisition. However, the optimization signal is fundamentally *black-box*: ranking models expose no gradients, and observable metrics such as impressions or click-through rate (CTR) are only available post-deployment, where exploration is expensive and feedback is biased by confounding factors such as position and popularity [25, 24, 40].

As illustrated in Fig. 1, even small stylistic changes to a snippet can significantly alter how users perceive and interact with results. In this example, the original snippet is accurate but lacks a strong promotional appeal, whereas the bottom snippet is a more engaging and policy-compliant description that highlights product attributes and use cases. This motivates the broader challenge: how can we

^{*}Equal contributions.

[†]All authors are from Walmart Global Tech, USA

[‡]Corresponding Author.

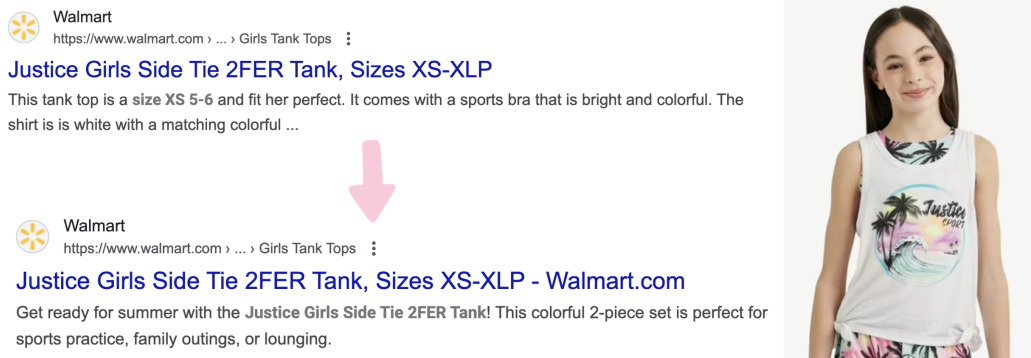


Figure 1: An example of how **MetaSynth** optimizes search engine meta descriptions. The top snippet (pre-optimization) is factual but generic, while the bottom snippet (MetaSynth) emphasizes promotional value, readability, and policy compliance. Such refinements directly impact user engagement and search-driven traffic by producing coherent and persuasive messages.

systematically design models that learn from observable outcomes to optimize text for engagement, despite the black-box nature of modern rankers.

Existing strategies have notable limitations. Template-based generation provides consistency but lacks expressiveness and generalization across domains [19, 5]. Prompt-only large language models (LLMs) generate fluent text but remain ungrounded, often hallucinating attributes or reverting to generic phrasing [29]. Retrieval-augmented generation (RAG) improves factual grounding, but retrieval is usually based solely on content similarity, disregarding whether candidate styles have historically been rewarded in ranking outcomes [37, 16]. As a result, current methods fail to fully exploit the implicit supervision embedded in black-box outputs [2].

We address this gap with **MetaSynth**, a model-based multi-agent, retrieval-augmented LM framework designed to “play the black-box search engine” by leveraging weak supervision from observable outcomes. MetaSynth constructs an exemplar success library by harvesting (query, metadata) pairs from top-ranked results, which we interpret as an implicit world model capturing ranking and click preferences. For a new webpage, an agentic retriever synthesizes plausible queries and retrieves relevant exemplars; when coverage is insufficient, it expands the library. A constrained LM generator proposes candidate snippets conditioned on content and exemplar styles. A panel of evaluator agents (critics) then scores relevance, coverage, promotional tone, and policy/brand compliance; a consensus coordinator integrates their feedback to plan targeted revisions. The result is an interpretable optimization loop guided by transparent objectives over a learned world model.

This framework reframes black-box optimization as a problem of “learning from weak supervision” while maintaining truthfulness and compliance. We evaluate MetaSynth through both offline experiments using proprietary and public data sets randomized online A/B tests. Offline studies enable principled ablation and iteration, while online deployment measures real-world impact on CTR and traffic.

Empirical results show that MetaSynth consistently outperforms prompt-only LLMs, and standard RAG, achieving state-of-the-art performance across NDCG, MRR, and average rank. Online A/B tests further demonstrate statistically significant improvements of **+10.26% CTR** and **+7.51% clicks**, validating both effectiveness and scalability. Beyond metadata optimization, our work contributes a general paradigm for leveraging implicit preference signals in black-box systems. We argue that this paradigm learning from weak but abundant observational cues opens new opportunities in ranking, recommendation, and personalization tasks where direct supervision is scarce but outcome-driven signals are observable.

Our main contributions are as follows:

- We propose **MetaSynth**, a novel multi-agent retrieval-augmented generation framework that exploits weak supervision from search and recommendation outcomes.

- We introduce a *exemplar success library* and autonomous retrieval strategy that encode implicit ranking preferences and dynamically expand coverage through agentic query generation, and provides weak supervision for the generation process.
- We design an automated evaluator-generator refinement loop with consensus-driven feedback, ensuring outputs satisfy relevance, fluency, and brand/policy guardrails.
- We demonstrate MetaSynth’s effectiveness through both large-scale offline evaluations and online A/B tests, showing consistent improvements over strong baselines and significant real-world impact on CTR and clicks.

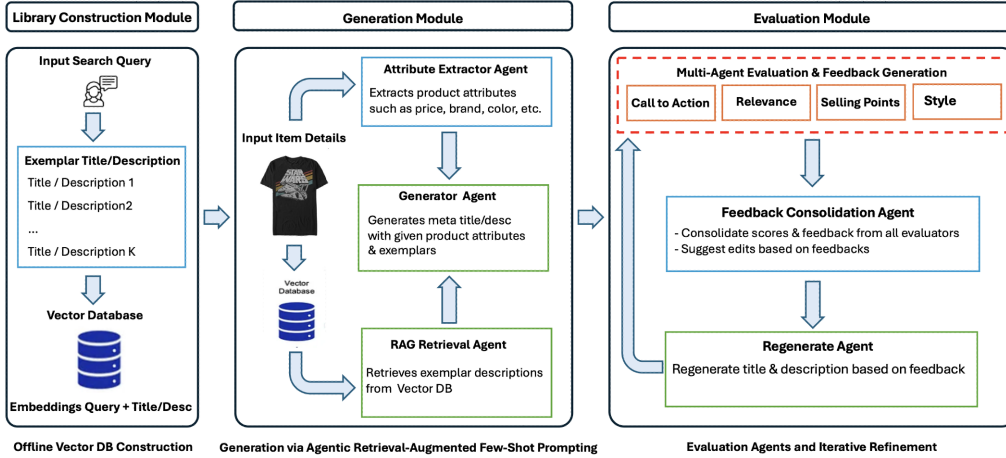


Figure 2: MetaSynth Framework: Three main modules to generate and optimize meta titles and descriptions for items, according to seller’s provided information, and constraints for better search engine ranking.

2 Related Work

Search and recommendation systems expose only their outputs, making it difficult to infer or optimize internal ranking preferences [32, 33, 20]. In this setting, user-facing snippets (meta titles and descriptions on SERPs) are known to shape attention and clicks, and thus downstream traffic [9, 23]. Prior work on query-biased summarization and snippet construction improves relevance and readability, but typically optimizes proxy text-quality metrics rather than outcome-driven objectives tied to ranking or clicks [3, 1, 8]. Moreover, exploration in production is expensive and biased by position and popularity, motivating methods that can learn from observational data without full access to the ranker [25, 15].

Template-based NLG offers control and consistency for product and listing metadata but struggles to generalize stylistically across domains [31, 12, 17]. Prompt-only LLMs increase fluency and diversity, yet can hallucinate attributes or regress to generic phrasing without grounding in historically successful styles [7, 29]. Retrieval-augmented generation (RAG) improves faithfulness by conditioning on retrieved evidence, but standard retrieval is primarily similarity-driven and agnostic to whether candidate styles have performed well under ranking [27, 36, 16]. Recent work examines leveraging observational signals for writing and recommendation [28, 2, 6], yet typically treats them as features for rerankers or classifiers rather than as priors for style-aware generation [13, 26].

Multi-agent LLM frameworks and self-refinement protocols use specialized roles (e.g., critics and verifiers) to improve reliability and adherence to constraints [38, 39, 7, 34, 21]. However, most evaluators optimize textual-quality proxies rather than outcome-aligned objectives and rarely close the loop with retrieval decisions [30].

In contrast to mentioned work MetaSynth is positioned at the intersection of RAG, weak supervision, and multi-agent self-improvement. It builds an *exemplar library* from top-ranked results and conditions generation jointly on product content and *outcome-informed* stylistic exemplars, thereby injecting an explicit, outcome-aligned prior into the generator rather than relying on content similarity

alone. An agentic retriever synthesizes queries to retrieve and expand coverage when library support is sparse, while a constrained generator produces candidates that are subsequently refined by a panel of evaluator agents scoring relevance, coverage, promotional strength, and brand/policy compliance; a consensus coordinator converts this feedback into targeted revisions. This closes the loop between retrieval, generation, and evaluation without gradient access to the ranker or curated preference labels, yielding an interpretable and controllable optimization process that aligns style with observed preferences. As such, MetaSynth operationalizes weak supervision for style-aware generation and offers a practical pathway to optimize text for engagement in black-box ranking environments.

3 Methodology

In this section we introduce MetaSynth, a multi agent generation and evaluation framework that takes exemplar meta data as input to generate meta snippets along with an evaluator-refinement loop. Our approach contains three main components i) Library Construction Module ii) Generation Module and iii) Evaluation Module. Fig 2 shows our entire framework design.

3.1 Problem Definition

Let \mathcal{X} denote the set of retailer product pages for a target eCommerce platform. For a page $x \in \mathcal{X}$, with its associated textual meta data like product name, brand, category, seller description $\mathbf{a}(x)$, the goal is to produce a meta-snippet $y = (\tau, \delta)$ comprising a meta title τ and meta description δ that maximizes downstream organic acquisition under a black-box search engine. We denote the (unknown) objective induced by the search engine and user behavior as

$$J(y | x) = \mathbb{E}[\text{traffic} | x, y], \quad (1)$$

which cannot be optimized directly because both ranking and user response are black-box. If the search engine was not a black box system, ideally we could have searched for optimal meta-snippet y^* by maximizing $J(\cdot)$:

$$y^* = \operatorname{argmax}_y J(y | x)$$

We therefore construct a multi-agent system that (i) learns from top-ranked search results as weak supervision for writing style, (ii) retrieves relevant exemplars, (iii) generates candidate snippets, and (iv) iteratively evaluates and refines them under brand guardrails.

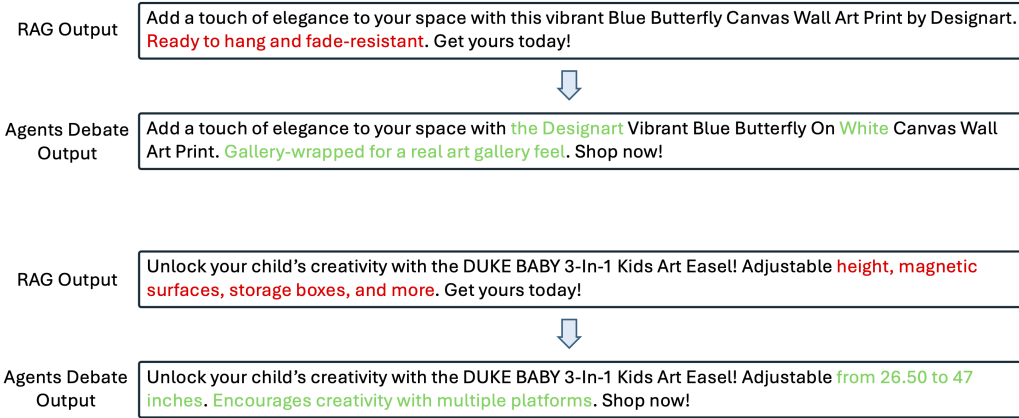


Figure 3: Two examples showcasing the edits done by evaluation and refinement agents on RAG outputs, to make the description more accurate and engaging, according to seller’s provided description.

Under this problem setting, one can model the search engine as a black-box function

$$\text{Search}(q, K) \rightarrow \{(u_i, \tau_i, \delta_i, r_i)\}_{i=1}^K, \quad (2)$$

which, given a query q , returns the top- K results with URL u_i , meta title τ_i , meta description δ_i , and rank $r_i = i$. We maintain a library \mathcal{L} of exemplars with entries $e = (q, u, \tau, \delta, r)$.

We embed all objects (including queries, product pages, and candidate snippets), we wish to compare into a shared vector space \mathbb{R}^d . Formally, $\mathbf{g}_q : \mathcal{Q} \rightarrow \mathbb{R}^d$ maps a textual query to a d -dimensional vector, $\mathbf{g}_x : \mathcal{X} \rightarrow \mathbb{R}^d$ maps a product page (aggregating its structured and unstructured attributes), and $\mathbf{g}_y : \mathcal{Y} \rightarrow \mathbb{R}^d$ maps a meta title–description pair. A common embedding space enables direct geometric comparisons across heterogeneous items: a query should lie near the products it meaningfully retrieves, and a high-quality snippet should lie near exemplars that reflect the desired style and content. While we do not assume a specific embedding model, we require that these maps are calibrated so that proximity corresponds to semantic relatedness. We measure proximity using cosine similarity.

Brand guardrails are denoted by set \mathcal{B} . In practice, these guardrails are stated by constraints, requirements, and acceptance thresholds denoted by \mathcal{H} , \mathcal{R} , and α respectively. The set \mathcal{H} captures *hard* prohibitions (e.g., legally sensitive claims, banned phrases), any of which yields immediate rejection. The set \mathcal{R} specifies *required* elements (e.g., presence of a call to action, inclusion of brand name) that must be verifiably present in the snippet. The vector α contains per-criterion thresholds for continuous quality scores (e.g., minimum relevance to the page, minimum promotional strength, minimum style compliance). During evaluation, the system computes a score vector and checks it against α while enforcing \mathcal{H} and \mathcal{R} ; only candidates satisfying all hard/required constraints and meeting or exceeding the thresholds are accepted, ensuring optimization for search performance remains aligned with brand and policy requirements.

3.2 Library Construction Module: Offline Vector DB Construction

Given a seed set of popular queries \mathcal{Q}_S , the Library Construction agent issues calls to the black-box engine and builds $\mathcal{L} = \bigcup_{q \in \mathcal{Q}_S} \{(q, u_i, \tau_i, \delta_i, r_i)\}_{i=1}^{K_{\text{lib}}}$. For saving each exemplar i , both its meta title τ_i and meta description δ are embedded using $\mathbf{g}_{y_i}(e_i) = \mathbf{g}_{y_i}(\tau_i \oplus \delta_i)$. We do not save the embedding vector if the cosine similarity with the most similar item in the library is more than a threshold ϵ_{dup} . In other words, e_i is considered a duplicate of existing e_j in the library if: $\text{sim}(\mathbf{g}_y(e_i), \mathbf{g}_y(e_j)) > \epsilon_{\text{dup}}$.

We also index a query-to-exemplar map $\mathcal{I}(q) = \{e \in \mathcal{L} : eq = q\}$ and a global ANN index over $\mathbf{g}_q(q)$ to support fast retrieval.

3.3 Generation Module: Generation via Agentic Retrieval-Augmented Few-Shot Prompting

3.3.1 Target Query Detection and Library Construction

To select a query for a target product page x , first its textual meta-data embedding vector is computed $\mathbf{z}_x = \mathbf{g}_x(a(x))$. Then, the most similar query q^* to \mathbf{z}_x is obtained and the similarity score between q^* and x is stored. In other words, we have

$$\begin{cases} q^* &= \arg \max_{q \in \text{dom}(\mathcal{I})} \text{sim}(\mathbf{z}_x, \mathbf{g}_q(q)), \\ s^* &= \max_q \text{sim}(\mathbf{z}_x, \mathbf{g}_q(q)). \end{cases} \quad (3)$$

Given a target similarity threshold $\tau_q \in (0, 1)$, if $s^* < \tau_q$ (i.e., no sufficiently similar query exists), the generation agent constructs candidate queries from attributes, by using Expand prompt template that generates new relevant queries $q_{\text{new},x}$ based on the product page’s associated textual data $a(x)$. In other words,

$$q_{\text{new},x} = \text{Expand}(a(x)), \quad (4)$$

and then invokes Search for each $q_{\text{new},x}$.

For each $q \in q_{\text{new},x}$, we augment \mathcal{L} with the top- K_{aug} results extracted by $\text{Search}(q, K)$ call and update the ANN index.

For the case where there exists a similar query in the library ($s^* \geq \tau_q$), we instead take the most relevant query q^* itself as well as all other similar queries passing the threshold τ_q . Then, we construct the Exemplar set as described below using obtained set.

3.3.2 Exemplar set construction and Generation

Let \mathcal{F}_x be the candidate exemplars set for webpage x . The m few-shot exemplars selected by a greedy algorithm using Maximal Marginal Relevance (MMR) [4] to balance relevance and diversity:

$$\begin{aligned} \text{MMR}(e \mid \mathcal{F}) &= \lambda \cdot \text{sim}(\mathbf{z}_x, \mathbf{g}_y(e)) \\ &\quad - (1 - \lambda) \cdot \max_{e' \in \mathcal{F}} \text{sim}(\mathbf{g}_y(e), \mathbf{g}_y(e')), \end{aligned} \quad (5)$$

where \mathcal{F} is the selected growing set and $\lambda \in [0, 1]$. We iterate m steps to obtain \mathcal{F}_x . This will ensure that we have a diverse enough set of successful examples to be passed to meta-snippet generation.

Let G be a generation function conditioned on (i) product content, (ii) \mathcal{F}_x , and (iii) brand guardrails \mathcal{B} . The initial meta-snippet is obtained by

$$y^{(0)} = G(x, \mathcal{F}_x, \mathcal{B}), \quad (6)$$

where the prompt includes structured slotting of the associated text with webpage x , $a(x)$, and the selected exemplars \mathcal{F}_x , and guardrails \mathcal{B} .

Please note that The relevance filter requiring u_x to appear in the top- K_{hit} enforces that only queries demonstrably leading to the target page are retained; the few-shot pool \mathcal{F}_x is then assembled from the *top results of those queries*, serving as weakly supervised exemplars of effective writing styles.

3.4 Evaluation Module: Evaluation Agents and Iterative Refinement

The Evaluation agent E_ϕ scores a candidate y for page x along K criteria and returns a score vector $\mathbf{s}(y, x) \in [0, 1]^K$ and textual feedback $\mathbf{c}(y, x)$. We instantiate the following four primary criteria: (i) $s_{\text{rel}}(y, x) \in [0, 1]$ which evaluates if the generated meta-snippet, y , is relevant to the target item page, x , (ii) $s_{\text{promo}}(y) \in [0, 1]$ which evaluates if y has a promotional tone, (iii) $s_{\text{cta}}(y) \in \{0, 1\}$ which evaluates if y has a call-to-action (phrases like “buy now”), and (iv) $s_{\text{brand}}(y; \mathcal{H}) \in [0, 1]$ which evaluates if y is abiding the brand/style guidelines.

The generated feedback at each round t by evaluators is stored $\mathbf{c}^{(t)} = \mathbf{c}(y^{(t)}, x)$ and given this feedback the generator produces a revised meta-snippet:

$$y^{(t+1)} = G(x, \mathcal{F}_x, \mathcal{B}, y^{(t)}, \mathbf{c}^{(t)}). \quad (7)$$

In this case, $\mathbf{c}^{(t)}$ is injected as structured constraints (e.g., “increase promotional strength,” “insert CTA,” “remove forbidden term h ”). The cycle stops at iteration t^* if either (i) if the Evaluator accepts the generated text on all criteria or (ii) iteration hits the max iteration budget K_{max} . Optionally, a stagnation rule halts the iterations if enough improvement does not happen at all (or enough) for two consecutive steps. Fig. 3 shows an example of how evaluation module modifies the initial generated output of 3.3.2. Please refer to Appendix B for entire algorithm.

4 Experiments and Results

4.1 Data and Experiment Setting

We conduct experiments on two datasets to comprehensively evaluate MetaSynth.

- **Proprietary dataset:** A large-scale e-commerce catalog containing 40,000 items. We sample an equal proportion of products from four diverse categories—Clothing, Electronics, Toys, and Home & Garden. Each item is associated with rich metadata, including product titles, specifications, brand information, and customer reviews.
- **Amazon Review dataset [22]:** A widely used public benchmark curated by McAuley Lab. We extract 30,000 items across three domains, Home, Toys, and Fashion, with 10,000 items in each category. This dataset provides structured metadata such as item descriptions, prices, and review text, making it complementary to our proprietary corpus.

For preprocessing, we apply standard NLP techniques to remove non-ASCII characters, special symbols, and noise. For each item, we employ an LLM to generate the most likely user query that

could lead to that item. We then identify the closest matching query from our VectorDB (as described in section 3.3.1) and retrieve the top- k exemplar meta titles and descriptions. These exemplars serve as weakly supervised demonstrations for the **Generation Module** (section 3.3.2), ensuring that outputs are conditioned on successful historical styles. We further regenerate meta-snippets based on feedback and use them as outputs to compare again these three benchmarks (section 3.4). This setup allows us to test MetaSynth on both controlled proprietary environments and an open, publicly reproducible benchmark.

4.2 Evaluation Metrics

We compare MetaSynth against three representative baselines:

1. **Vanilla (Baseline):** A single LLM call that generates meta titles and descriptions directly from item metadata, without retrieval or reasoning.
2. **DRE (Direct Retrieval Enhancement) [18]:** An approach that extracts keywords from items and incorporates them into generation prompts, thereby enriching outputs with content-derived terms.
3. **CoT (Chain-of-Thought prompting) [35]:** A reasoning-based method that guides LLMs through structured intermediate steps, improving factual alignment between metadata and generated snippets.

To ensure consistent evaluation, we use **GPT-4.1-mini** as a judgment model, applying the LLM-as-a-judge paradigm [41] to rank snippets generated by each of these baseline models. We report three widely accepted ranking metrics:

- **NDCG (Normalized Discounted Cumulative Gain):** Measures the overall quality of ranked outputs, giving higher weight to relevant outputs that appear at the top of the list. Scores are normalized between 0 and 1, with higher values indicating stronger alignment with ideal rankings.
- **MRR (Mean Reciprocal Rank):** Captures how quickly the first highly relevant output appears in the ranking. An MRR close to 1 implies that the correct snippet is almost always ranked first.
- **Average Rank:** Records the average position of generated outputs across all items. Lower values indicate better performance, as strong methods consistently place outputs near the top.

Together, these metrics provide complementary perspectives: NDCG highlights ranking quality, MRR emphasizes efficiency in surfacing relevant snippets, and Average Rank evaluates overall placement robustness.

4.3 Experiment Results

On the Amazon Review dataset (See Table 1), Baseline establishes a modest benchmark (e.g., NDCG 0.5970, MRR 0.4617 on Fashion titles), but its lack of explicit reasoning often results in suboptimal ranking. DRE provides incremental gains in some settings, such as improved meta description retrieval on Home (NDCG 0.5885 vs. 0.5609 for Baseline), but overall lags behind other methods, as its heuristic adjustments fail to capture deeper semantic structure. COT delivers more consistent improvements by incorporating structured reasoning, achieving higher retrieval quality (e.g., Fashion meta description MRR of 0.5282 vs. 0.4013 for Baseline). However, while COT narrows the gap, its reasoning alone cannot fully address domain-specific nuances in product data.

MetaSynth achieves the strongest results across all datasets, significantly outperforming prior methods. On Fashion titles, MetaSynth reaches NDCG of 0.8190 and MRR of 0.7601, compared to 0.6280 and 0.5046 for COT. A similar trend holds for meta descriptions, where MetaSynth (NDCG 0.7911, MRR 0.7213) consistently outperforms all baselines. The improvements extend to Toys (title MRR 0.7551 vs. 0.4562 for COT) and Home (description NDCG 0.7996 vs. 0.6048 for COT), with MetaSynth reducing average rank to 1.7 across domains. These results highlight MetaSynth’s ability to generalize across categories while preserving fine-grained detail in retrieval.

Table 1: Offline evaluation results: LLM-as-a-judge metrics across Amazon and proprietary datasets for meta titles and meta descriptions.

Dataset	Approach	Meta Title			Meta Description		
		NDCG	MRR	Avg Rank	NDCG	MRR	Avg Rank
Amazon Fashion	Baseline	0.5970	0.4617	2.5284	0.5504	0.4013	2.8355
	DRE	0.5095	0.3505	3.2731	0.5673	0.4261	2.9335
	COT	0.6280	0.5046	2.4911	0.6443	0.5282	2.5221
	MetaSynth	0.8190	0.7601	1.7025	0.7911	0.7213	1.6987
Amazon Toys	Baseline	0.6302	0.5051	2.3450	0.5873	0.4492	2.5834
	DRE	0.5174	0.3606	3.2197	0.5972	0.4657	2.7714
	COT	0.5911	0.4562	2.7260	0.5804	0.4450	2.9459
	MetaSynth	0.8149	0.7551	1.7015	0.7881	0.7169	1.6921
Amazon Home	Baseline	0.6648	0.5510	2.2187	0.5609	0.4147	2.7284
	DRE	0.5486	0.4005	3.0045	0.5885	0.4543	2.8140
	COT	0.5846	0.4478	2.8052	0.6048	0.4764	2.7941
	MetaSynth	0.7598	0.6809	1.9694	0.7996	0.7316	1.6627
Proprietary	Baseline	0.5762	0.4352	2.6834	0.4771	0.3160	3.2655
	DRE	0.5553	0.4102	2.9622	0.5066	0.3552	3.1309
	COT	0.6527	0.5381	2.4191	0.7117	0.6243	1.9640
	MetaSynth	0.7631	0.6882	1.9716	0.7835	0.7204	1.6416

On the e-commerce Proprietary dataset, the differences become even clearer. On this dataset, Vanilla and DRE show relatively weak performance, with DRE only marginally improving over Vanilla in description generation. COT achieves notable gains, particularly in meta descriptions (NDCG of 0.7117 and MRR of 0.6243), demonstrating the utility of structured reasoning in this domain. However, MetaSynth delivers the strongest results across all metrics, achieving the highest NDCG (0.7631 for titles, 0.7835 for descriptions) and the lowest average rank (1.9716 for titles, 1.6416 for descriptions). Still, MetaSynth delivers the largest gains, achieving the highest NDCG and MRR.

The offline evaluation metrics clearly show that our proposed approach MetaSynth consistently performs in all of them beating all benchmark models indicating the importance of exemplar library and evaluator feedback loop. Although benchmarks like COT and DRE fare better than Vanilla, MetaSynth yields state-of-the-art retrieval performance across all datasets and metrics. Please refer to Appendix A for real examples from our experiments showing the results of each benchmark model.

4.4 Ablation Study

Table 2: Ablation Studies

Approach	Average rank	NDCG	MRR
MetaSynth wo RAG	2.4830	0.6245	0.5018
Meta Synth wo Evaluation	1.9769	0.7267	0.6353
MetaSynth	1.6416	0.7835	0.7204

To systematically analyze the contribution of individual components within our pipeline, we conducted ablation studies across 3 variants i) MetaSynth without Library Construction Module ii) MetaSynth without Evaluation module iii) Complete MetaSynth pipeline. For this analysis, we compared the performance of each these variants with benchmark models on our Proprietary dataset.

The results summarized in (see Table 2) show that exclusion of Library Construction Module leads to a substantial degradation in performance there is a 33% drop in average rank compared to the complete Meta Synth framework which highlights the role of high quality exemplar titles and descriptions. Additionally, when we remove the Evaluation and Feedback module, there is again a considerable decline in all metrics indicating that the loop of evaluation, feedback consolidation and regeneration is essential for good performance. In ranking effectiveness, MetaSynth yields a 25.5% gain in MRR and a 25.4% gain in NDCG compared to MetaSynth w/o RAG, while achieving a further 13.4% (MRR) and 7.8% (NDCG) improvement over MetaSynth w/o Evaluation. Notably, the observed performance drop is more pronounced when Library Module is omitted compared to ii) W/o Evaluation suggesting that Library Module exerts a greater influence on the performance, Overall, having all three components surpasses the other methods demonstrating the significance of each of the components. In particular, evaluation modules enhances ranking precision, while Library Construction Module enriches the VectorDB enabling Meta Synth to perform the best in all offline evaluation metrics.

4.5 A/B Test Results

To evaluate the impact of meta-titles and snippets on user engagement, we conducted 4-weeks long A/B test to compare a proprietary control search engine meta generator with the ones generated through MetaSynth. MetaSynth leads to a +10.26% improvement in clicks and 7.51% in CTR when compared to the control model (Table 3). These statistically significant lifts confirm that MetaSynth’s offline improvements translate directly into real-world user engagement. Importantly, the gains substantially outweigh the additional inference costs introduced by the multi-agent pipeline, demonstrating that MetaSynth is both effective and scalable for large catalogs.

Table 3: A/B test performance

Metric	Lift
CTR	+10.26%
Overall Clicks	+7.51%

5 Conclusion

We presented **MetaSynth**, a multi-agent retrieval-augmented generation framework for optimizing metadata in search and recommendation settings. Unlike prior template-based, prompt-only, or standard RAG approaches, MetaSynth directly leverages implicit signals from observable outcomes, treating top-ranked results as weak supervision for learning style and content preferences. Our design integrates three key components: an exemplar library, a constrained generator conditioned on product content and exemplars, and an evaluator–generator refinement loop that enforces relevance, promotional strength, and compliance.

Extensive experiments across both proprietary e-commerce data and the Amazon Reviews corpus demonstrate consistent gains in NDCG, MRR, and ranking quality compared to strong baselines. Large-scale online A/B tests further confirm the practical impact, yielding **+10.26% CTR** and **+7.51% clicks**. These results highlight MetaSynth’s ability to bridge the gap between black-box ranking systems and generative optimization, offering a reproducible methodology that is both effective and deployable.

Beyond metadata generation, this work contributes a broader paradigm for learning from implicit feedback in black-box environments. By showing how multi-agent generation can integrate weak supervision, retrieval, and iterative critique, we open new directions for applying similar techniques to recommendation, personalization, and ranking-adjacent tasks. Future research may extend this framework to richer modalities (e.g., images, video), integrate counterfactual debiasing of implicit signals, and explore theoretical guarantees for convergence under noisy supervision.

References

- [1] T. Baumel, M. Eyal, and M. Elhadad. Query focused abstractive summarization: Incorporating query relevance, multi-document coverage, and summary length constraints into seq2seq models. *arXiv preprint arXiv:1801.07704*, 2018.
- [2] H. Cao. Writing style matters: An examination of bias and fairness in information retrieval systems. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*, pages 336–344, 2025.
- [3] Z. Cao, W. Li, S. Li, F. Wei, and Y. Li. Attsum: Joint learning of focusing and summarization with neural attention. *arXiv preprint arXiv:1604.00125*, 2016.
- [4] J. Carbonell and J. Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 335–336, 1998.
- [5] H. Chen, X. Chen, S. Shi, and Y. Zhang. Generate natural language explanations for recommendation. *arXiv preprint arXiv:2101.03392*, 2021.
- [6] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3):1–39, 2023.
- [7] J. Chen, K. Yao, R. Yousefi Maragheh, K. Zhao, J. Xu, J. Cho, E. Korpeoglu, S. Kumar, and K. Achan. Carts: Collaborative agents for recommendation textual summarization. *arXiv preprint arXiv:2506.17765*, 2025.
- [8] Z. Chen, X. Li, X. Ren, and J. Yu. Abstractive snippet generation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1301–1310. ACM, 2020.
- [9] A. Chuklin and M. de Rijke. The anatomy of relevance: Topical, snippet and perceived relevance in search result evaluation. *arXiv preprint arXiv:1501.06412*, 2015.
- [10] C. L. Clarke, E. Agichtein, S. Dumais, and R. W. White. The influence of caption features on clickthrough patterns in web search. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 135–142, 2007.
- [11] P. Covington, J. Adams, and E. Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198, 2016.
- [12] K. V. Deemter, M. Theune, and E. Krahmer. Real versus template-based natural language generation: A false opposition? *Computational linguistics*, 31(1):15–24, 2005.
- [13] Y. Deldjoo, Z. He, J. McAuley, A. Korikov, S. Sanner, A. Ramisa, R. Vidal, M. Sathiamoorthy, A. Kasrizadeh, S. Milano, et al. Recommendation with generative models. *arXiv preprint arXiv:2409.15173*, 2024.
- [14] F. Doshi-Velez and B. Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- [15] B. Ermiş, P. Ernst, Y. Stein, and G. Zappella. Learning to rank in the position based model with bandit feedback. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2405–2412, 2020.
- [16] N. Forouzandehmehr, R. Y. Maragheh, S. Kollipara, K. Zhao, T. Biswas, E. Korpeoglu, and K. Achan. Cal-rag: Retrieval-augmented multi-agent generation for content-aware layout design. *arXiv preprint arXiv:2506.21934*, 2025.
- [17] S. Gajbhiye and M. Lopes. Template-based nlg for tabular data using bert. In *2021 Grace Hopper Celebration India (GHCI)*, pages 1–5. IEEE, 2021.
- [18] S. Gao, Y. Wang, J. Fang, L. Chen, P. Han, and S. Shang. Dre: Generating recommendation explanations by aligning large language models at data-level. *arXiv preprint arXiv:2404.06311*, 2024.
- [19] A. Gatt and E. Krahmer. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170, 2018.
- [20] R. Giahi, J. Xu, R. Y. Maragheh, E. Korpeoglu, and K. Achan. Systems and methods for siamese wide and deep neural network ranking, July 31 2025. US Patent App. 18/429,128.

- [21] R. Giahi, K. Yao, S. Kollipara, K. Zhao, V. Mirjalili, J. Xu, T. Biswas, E. Korpeoglu, and K. Achan. VI-clip: Enhancing multimodal recommendations via visual grounding and llm-augmented clip embeddings. *arXiv preprint arXiv:2507.17080*, 2025.
- [22] Y. Hou, J. Li, Z. He, A. Yan, X. Chen, and J. McAuley. Bridging language and items for retrieval and recommendation. *arXiv preprint arXiv:2403.03952*, 2024.
- [23] M. A. Islam, R. Srikant, and S. Basu. Micro-browsing models for search snippets. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1904–1909. IEEE, 2019.
- [24] M. A. Islam, K. Vasilaky, and E. Zheleva. Correcting for position bias in learning to rank: A control function approach. *arXiv preprint arXiv:2506.06989*, 2025.
- [25] T. Joachims, A. Swaminathan, and T. Schnabel. Unbiased learning-to-rank with biased feedback. In *WSDM*, 2017.
- [26] A. Khan, A. Wang, S. Hager, and N. Andrews. Learning to generate text in arbitrary writing styles. *arXiv preprint arXiv:2312.17242*, 2023.
- [27] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474, 2020.
- [28] E. Loghmani. Aligning language models with observational data: Opportunities and risks from a causal perspective. *arXiv preprint arXiv:2506.00152*, 2025.
- [29] R. Y. Maragheh and Y. Deldjoo. The future is agentic: Definitions, perspectives, and open challenges of multi-agent recommender systems. *arXiv preprint arXiv:2507.02097*, 2025.
- [30] S. R. Motwani, C. Smith, R. J. Das, R. Rafailov, I. Laptev, P. H. Torr, F. Pizzati, R. Clark, and C. S. de Witt. Malt: Improving reasoning with multi-agent llm training. *arXiv preprint arXiv:2412.01928*, 2024.
- [31] E. Reiter. Nlg vs. templates. *arXiv preprint cmp-lg/9504013*, 1995.
- [32] M. Rolínek, V. Musil, A. Paulus, M. Vlastelica, C. Michaelis, and G. Martius. Optimizing rank-based metrics with blackbox differentiation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7620–7630, 2020.
- [33] M. Sazanovich, A. Nikolskaya, Y. Belousov, and A. Shpilman. Solving black-box optimization challenge via learning search space partition for local bayesian optimization. In *NeurIPS 2020 Competition and Demonstration Track*, pages 77–85. PMLR, 2021.
- [34] D. Wan, J. C.-Y. Chen, E. Stengel-Eskin, and M. Bansal. Mamm-refine: A recipe for improving faithfulness in generation with multi-agent collaboration. *arXiv preprint arXiv:2503.15272*, 2025.
- [35] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [36] R. Yousefi Maragheh, P. Vadla, P. Gupta, K. Zhao, A. Inan, K. Yao, J. Xu, P. Kanumala, J. Cho, and S. Kumar. Arag: Agentic retrieval augmented generation for personalized recommendation. *arXiv preprint arXiv:2506.21931*, June 2025.
- [37] H. Yu, A. Gan, K. Zhang, S. Tong, Q. Liu, and Z. Liu. Evaluation of retrieval-augmented generation: A survey. In *CCF Conference on Big Data*, pages 102–120. Springer, 2024.
- [38] P. Yu, G. Chen, and J. Wang. Table-critic: A multi-agent framework for collaborative criticism and refinement in table reasoning. *arXiv preprint arXiv:2502.11799*, 2025.
- [39] Y. Yuan and T. Xie. Reinforce llm reasoning through multi-agent reflection. *arXiv preprint arXiv:2506.08379*, 2025.
- [40] Y. Zhao, Z. Liu, T. Cai, H. Zhang, C. Zhuang, and J. Gu. Mitigate position bias with coupled ranking bias on ctr prediction. *arXiv preprint arXiv:2405.18971*, 2024.
- [41] L. Zheng, W. L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, volume 36, 2024.

Appendix

A Case studies





	<p>Trademark Fine Art</p> <p>Trademark Fine Art 'Us Capitol Building And Red Hot Air Balloons' Canvas Art by Fab Funky</p>		<p>Visit the Walker Edison Store</p> <p>Walker Edison Modern Farmhouse Sliding Door Vertical Cabinet, Barnwood</p>
Baseline	Elevate your space with this 24"x32" modern canvas art of the US Capitol and red balloons. Ready to hang gallery wrap. Buy now and brighten your walls!	Baseline	Add rustic charm with the Walker Edison Farmhouse Sliding Door Cabinet. Features adjustable shelves & durable MDF wood. Get yours today!
DRE	Discover the vibrant 'Us Capitol Building And Red Hot Air Balloons' canvas art by Fab Funky. Modern, gallery-wrapped, 24x32 inches. Buy yours today!	DRE	Add rustic charm with the Walker Edison Farmhouse Sliding Door Cabinet. Features adjustable shelves & durable MDF wood. Organize in style—Shop Now!
CoT	Trademark Fine Art US Capitol Building And Red Hot Air Balloons Canvas Art, 24"x32", Modern Gallery-Wrapped Giclee Print. Vibrant, Ready to Hang. Buy Now!	CoT	Walker Edison Modern Farmhouse Sliding Door Vertical Cabinet, 68"H x 36"L, features rustic barnwood finish, grooved sliding door, and adjustable shelves. Shop Now!
MetaSynth	Discover the stunning US Capitol Building & Red Hot Air Balloons canvas art by Fab Funky . Add a touch of modern & contemporary style to your space with this gallery-wrapped masterpiece. Get yours today!	MetaSynth	Organize in style with the Walker Edison Modern Farmhouse Sliding Door Vertical Cabinet. Rustic design, ample storage space , and durable materials . Get yours today!
	<p>Avia</p> <p>Avia Women's Transition V-Neck Short Sleeve T-Shirt</p>		<p>Trademark Global</p> <p>Buffalo Sabres Portable Metal Bar Table with Carrying Case</p>
Baseline	Stay cool and comfortable in the Avia Women's Transition V-Neck T-Shirt with moisture-wicking fabric and soft peached jersey. Shop now and grab yours today!	Baseline	Elevate your game day with the NHL Portable Bar featuring Buffalo Sabres wrap, metal build, two shelves, and a carrying case. Get yours today!
DRE	Stay cool and comfortable in the Avia Women's Transition V-Neck T-Shirt with moisture-wicking fabric and soft peached jersey. Shop now!	DRE	Officially licensed Buffalo Sabres portable bar with 2 shelves, metal build, collapsible design, and carrying case. Perfect for tailgates! Get yours today!
CoT	Avia Women's Transition V-Neck Short Sleeve T-Shirt features Moisture Wicking Fabric & Soft Peached Jersey. Relaxed Fit for All Activities. Shop Now!	CoT	Official NHL Portable Bar with Carrying Case, 39" x 15" Metal Top, Collapsible Design, Buffalo Sabres Wrap. Perfect for BBQs & Picnics. Get Yours Today!
MetaSynth	Upgrade your activewear collection with the Avia Women's Transition V-Neck T-Shirt! Stay cool and comfortable during workouts with moisture-wicking fabric and a soft hand feel . Shop Now and get yours today!	MetaSynth	Get ready to elevate your next picnic or BBQ with the NHL Portable Bar featuring the Buffalo Sabres logo. Collapsible, spacious , and officially licensed. Shop Now!

Figure 4: Case studies comparing MetaSynth method with three other studies (baseline, DRE, CoT), in which MetaSynth outcome ranked first among others in the search engine. Some influential words that might be affected the ranking are demonstrated with bold font.

Below we compare the four variants for the “US Capitol Building and Red Balloons” canvas (See Figure 4, top-left case). Relative to the Baseline, DRE, and CoT texts, the MetaSynth description exhibits stronger lexical economy and discourse naturalness: it foregrounds the core entity and creator (“... by Fab Funky”), compresses categorical qualifiers (“modern & contemporary style”), and uses a single, semantically rich construction (“gallery-wrapped masterpiece”) rather than a list of loosely coupled attributes. In contrast, Baseline omits brand/creator and leans on generic language (“brighten your walls”), DRE reads as a keyword list with rigid slot filling (“Modern, gallery-wrapped, 24x32 inches”), and CoT over-enumerates proper nouns and media terms (“Trademark Fine Art ... Canvas Art ... Giclee Print”) in a way that resembles inventory metadata rather than user-centric copy. This shift from enumeration to fluent phrasing improves readability while preserving key tokens that matter for retrieval (e.g., “gallery-wrapped”, style cues), thus aligning with relevance, readability, and technical-compliance components of the surrogate objective. Second, the MetaSynth variant better harmonizes intent coverage with precision by balancing audience framing and attribute salience. The phrase “modern & contemporary style” broadens matchability across adjacent intents without resorting to repetitive keyword stuffing, while mentioning the creator “Fab Funky” supplies a trusted brand cue that can disambiguate entity searches. Compared with DRE’s mechanical cadence and CoT’s concatenated title-like string, MetaSynth’s clause structure distributes salient tokens across sentences (lead with identity → situate style → state mounting/finish) to satisfy evaluator checks on relevance and detail without triggering duplication penalties.

Additionally, we compare the four variants for the “Walker Edison Farmhouse Sliding Door Cabinet” product description (See Figure 4, top-right case). Relative to the Baseline, DRE, and CoT texts, the MetaSynth description demonstrates superior lexical efficiency: it leads with action-oriented framing

("Organize in style"), emphasizes core value propositions (ample storage space, durable materials), and maintains brand recognition while avoiding specification overload. In contrast, Baseline and DRE rely on identical phrasing, while CoT overwhelms with technical specifications (68"H x 36"L, "grooved sliding door") that read as catalog entries rather than persuasive copy. MetaSynth's selective emphasis highlights functionality and longevity without dimensional clutter, while "Organize in style" efficiently combines utility with aesthetic appeal. This creates broader search matching than CoT's specification-heavy approach while maintaining more substance than Baseline's generic "rustic charm" language, improving both readability and conversion potential.

Third, we compare the four variants for the "Avia Women's Transition V-Neck T-Shirt" product description (See Figure 4, bottom-left). Relative to the Baseline, DRE, and CoT texts, the MetaSynth description demonstrates superior market positioning: it leads with aspirational framing ("Upgrade your activewear collection"), strategically emphasizes category context (activewear) and tactile appeal (soft hand feel), while maintaining technical benefits within engaging copy. In contrast, Baseline and DRE open with generic comfort claims, while CoT defaults to feature enumeration ("Moisture Wicking Fabric & Soft Peached Jersey. Relaxed Fit") that prioritizes specifications over lifestyle integration. MetaSynth's selective bolding of activewear establishes clear category positioning for broader intent matching, while soft hand feel translates technical "peached jersey" into consumer-friendly sensory language. The phrase "Upgrade your activewear collection" positions the product as enhancement rather than necessity, creating stronger purchase motivation than CoT's functional listing or Baseline's passive "stay cool" messaging, improving both category relevance and conversion appeal.

Finally, we compare the four variants for the "NHL Portable Bar featuring Buffalo Sabres" product description (See Figure 4, bottom-right). Relative to the Baseline, DRE, and CoT texts, the MetaSynth description demonstrates superior contextual targeting: it leads with experiential framing ("Get ready to elevate your next picnic"), emphasizes specific use scenarios (your next picnic) and key functional benefits (spacious), while streamlining technical details into essential selling points. In contrast, Baseline opens with generic "game day" positioning, DRE focuses heavily on structural specifications, while CoT overwhelms with dimensional data (39" x 15") and feature lists that prioritize inventory details over lifestyle appeal. MetaSynth's selective bolding of your next picnic creates direct personal relevance and expands beyond traditional sports contexts, while spacious translates technical shelf configurations into practical consumer benefit language. The phrase "Get ready to elevate your next picnic" positions the product within broader outdoor entertainment rather than limiting to sports events, creating wider intent matching than Baseline's "game day" restriction or CoT's specification-heavy approach, improving both contextual relevance and purchase motivation.

B Algorithm

In this section of the appendix, we present a pseudocode for the proposed methodology of MetaSynth.

Algorithm 1 Concise Multi-Agent Meta-Snippet Generation

Require: Seed queries \mathcal{Q}_S ; cutoffs $K_{\text{lib}}, K_{\text{hit}}, K_{\text{aug}}$; thresholds $\epsilon_{\text{dup}}, \tau_q$; few-shot size m ; MMR weight λ ; max iters K_{max} ; guardrails $\mathcal{B} = (\mathcal{H}, \mathcal{R}, \alpha)$

Offline: Library

```

1: for  $q \in \mathcal{Q}_S$  do
2:    $R \leftarrow \text{Search}(q, K_{\text{lib}})$ 
3:   for  $(u, \tau, \delta, r) \in R$  do
4:     if  $\max_{e \in \mathcal{L}} \text{sim}(\mathbf{g}_y(\tau \oplus \delta), \mathbf{g}_y(e)) < \epsilon_{\text{dup}}$  then
5:       add  $e = (q, u, \tau, \delta, r)$  to  $\mathcal{L}$ ;  $\mathcal{I}(q) \leftarrow \mathcal{I}(q) \cup \{e\}$ 
6:     end if
7:   end for
8: end for

Online: Page  $x$  (URL  $u_x$ )
9:  $\mathbf{z}_x \leftarrow \mathbf{g}_x(\mathbf{a}(x)); \quad s^* \leftarrow \max_q \text{sim}(\mathbf{z}_x, \mathbf{g}_q(q))$ 
10: if  $s^* < \tau_q$  then  $\triangleright$  no similar repo query  $\Rightarrow$  agentic search
11:    $\mathcal{Q}_{\text{new}} \leftarrow \text{Expand}(\mathbf{a}(x))$ 
12:    $\mathcal{Q}_S(x) \leftarrow \{q \in \mathcal{Q}_{\text{new}} : u_x \in \text{top-}K_{\text{hit}} \text{ of } \text{Search}(q, K_{\text{hit}})\}$ 
13:   augment  $\mathcal{L}$  with top- $K_{\text{aug}}$  items (excluding  $u_x$ ) from those searches; update  $\mathcal{I}(\cdot)$ 
14: else
15:    $\mathcal{Q}_S(x) \leftarrow \{q : \text{sim}(\mathbf{z}_x, \mathbf{g}_q(q)) \geq \tau_q\}$ 
16: end if
17:  $\mathcal{E}(x) \leftarrow \bigcup_{q \in \mathcal{Q}_S(x)} \mathcal{I}(q);$ 
18:  $\mathcal{F}_x \leftarrow \text{MMR\_Select}(\mathcal{E}(x), \mathbf{z}_x, m, \lambda)$ 
19:  $y^{(0)} \leftarrow G(x, \mathcal{F}_x, \mathcal{B})$ 
20: for  $t = 0$  to  $K_{\text{max}} - 1$  do
21:    $(\mathbf{s}^{(t)}, \mathbf{c}^{(t)}) \leftarrow E_\phi(y^{(t)}, x, \mathcal{B})$ 
22:   if  $(\forall k, s_k^{(t)} \geq \alpha_k)$  then
23:     break  $\triangleright$  accepted
24:   else
25:      $y^{(t+1)} \leftarrow G(x, \mathcal{F}_x, \mathcal{B}, y^{(t)}, \mathbf{c}^{(t)})$ 
26:   end if
27: end for
28: return  $y^{(t)}$ 

```
