

You Won’t Believe This Click: Content Rewriting for Agentic Choice

Tianyi Jin Zirui Wang David M. Chan
UC Berkeley

Abstract

Language models are increasingly used as agents to help humans decide what information is surfaced. This usage incentivizes content creators to optimize content in ways that appeal not only to humans, but also to agents that mediate access to them. In this paper, we study selection shifts induced by rewriting in agentic decision-making. Given a set of competing content snippets, we rewrite only one snippet while leaving the rest unchanged, and then measure how it impacts the agent’s choice. We operationalize this setup with **AgentBait**, an advisor-rewriter framework in which the advisor learns to propose rewriting strategies and the rewriter revises the snippet. While a rewriter with a fixed prompt improves target snippet selection from 17.1% to 34.8%, our **AgentBait** raises its selection to 98.5%. We further show that the advisor trained with **AgentBait** effectively transfers to setups with different agents, languages, and snippets in other domains (e.g., scientific papers). However, higher target selection can reflect unsupported rewrites rather than better content. Adding a reward for support from the original snippet redirects the advisor toward more supported rewriting strategies, revealing a trade-off between factuality and target selection. Together, our results show that once agents mediate access to information, content can be rewritten to be chosen by the agent, even when selection and usefulness diverge.

1 Introduction

Language models are increasingly used as active agents to resolve user queries for news recommendations (Li et al., 2023; Liu et al., 2024), web search (Zhang et al., 2026; Anupam et al., 2025) and paper discovery (Li et al., 2025; Kumar et al., 2026). Recent work frames this shift toward agent-mediated information access as part of an emerging “Agent Attention Economy” (Yang et al., 2025). Human-facing studies similarly show that headline

wording can causally affect click behavior (Robertson et al., 2023). Snippets (e.g., titles and abstracts) that are usually written for human readers or retrieval systems (Cohan et al., 2020; Karpukhin et al., 2020; Wu et al., 2019) are now also used by agents for decision-making (i.e., content selection). While Generative Engine Optimization (GEO) (Aggarwal et al., 2024) studies content adaptation for visibility, citation, or inclusion in synthesized LLM search responses, we isolate *post-retrieval selection*, where the candidate list is fixed and the agent



Figure 1: **Can rewriting a document make it more likely to be chosen over the same competitors?** A list of competing documents is shown to the target agent. We choose one target document from the list and generate a rewriting strategy for only that document. A separate rewriting model then revises the target document’s title and abstract, while all other documents in the list remain exactly the same. The target agent selects from this updated list, and whether the rewritten target document is selected is used to train the advisor.

must choose an item from it. We measure whether rewriting the content of a target item changes how often it is selected, and ground our experimental setup in news (e.g., MIND (Wu et al., 2020); EB-NeRD (Kruse et al., 2024)) and scientific document recommendations (e.g., SciRepEval (Singh et al., 2023)).

Inspired by Advisor Models (Asawa et al., 2026), we employ an advisor-rewriter framework in which the advisor learns to propose an explicit, high-level rewriting strategy for a target item and a rewriter model revises the content based on the strategy. Starting from a rewriter-only baseline with a fixed prompt, we show that pairing it with a fixed advisor improves its performance. Reinforcement Learning (i.e., GRPO (Shao et al., 2024)) with the rewriter-only baseline further improves target selection rate, but it does not generalize well in unseen setups compared to an advisor-rewriter setup with a learned advisor. We also find that simply learning rewriting strategies with RL comes at the cost of losing factuality on the rewritten snippet, and thus we augment the objective with an additional reward that encourages the rewritten snippet to remain supported by the original text. This factuality-aware variant redirects the advisor toward more grounded rewriting strategies, while exposing a trade-off between target selection and source support. Together, these results show that agent-mediated selection creates a direct optimization pressure on content presentation, where rewriting can substantially change what an agent selects, but selection success alone does not imply factuality or usefulness to human readers. Our contributions are summarized as the following:

1. We propose a controlled protocol for measuring how rewriting a single snippet changes agent-mediated content selection. By fixing the candidate list, competing snippets, and agent prompt, the protocol isolates the effect of the target item’s presentation on the agent’s final choice.
2. We introduce **AgentBait**, an advisor-rewriter framework for optimizing target selection. The advisor learns high-level rewriting strategies, while the rewriter revises the target snippet from those strategies. On news impressions, a fixed-prompt rewriter improves target selection from 17.1% to 34.8%, while **AgentBait** raises it to 98.5%.

3. We show that learned rewriting strategies transfer across agents, languages, and domains, but can also reduce factuality. Adding a reward for support from the original text redirects the advisor toward more grounded rewriting strategies, revealing a trade-off between target selection and source support.

2 Related Work

Generative engine optimization. Aggarwal et al. (2024) show that simple content edits, such as adding citations, quotations, statistics, or improving fluency, can substantially affect how prominently a document is used, surfaced, or cited in generated answers. More broadly, optimization for generative search can target several stages of the pipeline: retrieval into the candidate set, selection among retrieved candidates, utilization during generation, and presentation in the final answer. Work on retrieval-augmented generation and neural retrieval improves the first stage by increasing the likelihood that relevant content is retrieved or ranked highly before generation (Lewis et al., 2020; Khattab and Zaharia, 2020; Karpukhin et al., 2020; Sarthi et al., 2024). A second line of work studies how LLMs use retrieved evidence during answer generation, including retrieval-conditioned generation, self-retrieval, context ranking, and evidence filtering (Nakano et al., 2022; Asai et al., 2024; Izacard et al., 2023; Yu et al., 2024). A third line of work focuses on attribution and citation in generated answers, aiming to make generated claims verifiable by linking them to supporting sources (Liu et al., 2023; Gao et al., 2023; Patel et al., 2024; Qi et al., 2024; Xia et al., 2025). Other GEO work learns or diagnoses generative-engine preferences more directly, such as extracting rewriting rules for web content or repairing citation failures (Wu et al., 2026b; Tian et al., 2026). Recent query-grounded GEO methods either perform topic-level optimization over interpretable document features to improve citation visibility while maintaining content quality (Liu and Xu, 2026), or distill validated editing patterns into reusable, engine-specific optimization skills to improve visibility and citation fidelity (Wu et al., 2026a). In contrast, our work focuses on post-retrieval selection, where the candidates are already retrieved but the agent has not yet decided which item to act on.

Targeted content optimization for agents. Our setting is also related to LLM-based reranking,

where a language model selects or orders items from a fixed candidate set. Prior work studies LLMs as pointwise, pairwise, setwise, or listwise rankers for retrieved passages, including prompting-based rerankers, distilled open-source rankers, and reproducible reranking toolkits (Sun et al., 2023; Qin et al., 2024; Zhuang et al., 2024; Pradeep et al., 2023a,b; Sharifymoghaddam et al., 2025; Abdallah et al., 2025). Recent work further shows that LLM-mediated choices can be steered through agent-facing content, including product pages and web snippets (Kumar and Lakkaraju, 2024; Pfrommer et al., 2024; Tang et al., 2025; Jin et al., 2026), e-commerce listings and GEO benchmarks (Bagga et al., 2025), tool descriptions and metadata (Shi et al., 2026; Wang et al., 2026), webpages and environmental injections (Xu et al., 2025; Liao et al., 2025), visual distractions such as pop-ups (Zhang et al., 2025), retrieval-corpus poisoning (Zhong et al., 2023), and multimodal item-level preference or rank manipulation (Jiang et al., 2025; Du et al., 2026). In contrast, we study targeted rewriting of ordinary text snippets in non-adversarial selection tasks, measuring whether presentation alone can shift an LLM agent’s choice among otherwise fixed candidates.

3 AgentBait: Optimizing Documents for Agentic Selection

While real-world recommendation systems involve complex, interacting components, including retrieval, ranking, presentation, and user feedback, here we isolate the post-retrieval selection phase to rigorously evaluate our approach. Following the standard multi-stage formulation of recommender systems (Covington et al., 2016), we assume a candidate generation stage has already narrowed the item universe to a manageable list. Our focus is exclusively on the downstream selection stage.

Formally, we define an impression as a single information-access context (e.g., a recommendation feed or search results page) containing a fixed candidate list $L_i = (x_{i,1}, \dots, x_{i,k_i})$. Each candidate item $x_{i,j}$ is represented by its short-text rendering, comprising a title and an abstract. This representation aligns with standard practices in both neural news recommendation (Wu et al., 2019) and scientific document retrieval (Cohan et al., 2020; Wang and Blei, 2011).

We introduce a target agent A , which receives the full candidate list L_i and must select exactly




Figure 2: Overview of our advisor-rewriter setting.

one item, denoted as $\hat{x} = A(L_i)$. The candidate list remains fixed throughout the evaluation; the system cannot retrieve new items or remove competing candidates. The core objective of our methodology is to optimize the textual representation of a designated *target item*, $x_{i,target} \in L_i$, to maximize its probability of being selected by the target agent A .

3.1 The Advisor-rewriter Framework

To separate strategic planning from text generation, we propose a decoupled Advisor-rewriter framework, illustrated in Figure 2. In this setting, all non-target items in L_i remain fixed, and only $x_{i,target}$ is modified.

Instead of training a single model to rewrite the text directly, we separate the strategic reasoning from the text generation:

1. **The Advisor:** A trainable policy that takes the target item as input, outputting a discrete *rewriting strategy*.

2. **The Rewriter:** A frozen, high-capacity language model that executes the Advisor’s strategy to produce the rewritten target item, $x'_{i,target}$.

Table 1 details the ablations and variants of this framework explored in our study.

Method	Advisor trained	Rewriter trained	Factuality reward
Rewriter	—	✗	✗
Trained rewriter	—	✓	✗
Trained rewriter + MC	—	✓	✓
Untrained advisor	✗	✗	✗
Trained advisor	✓	✗	✗
Trained advisor + MC	✓	✗	✓

Table 1: **Rewriting settings.** MC: MiniCheck.

3.2 Selection-Optimized Reward and Training

We train the Advisor policy using a binary reward derived directly from the target agent’s final selection. For a given rewritten candidate list L'_i containing the rewritten target item $x'_{i,target}$, the selection reward is defined as:

$$r_{select}(L'_i, x'_{i,target}) = \mathbb{I}\{A(L'_i) = x'_{i,target}\} \quad (1)$$

We train the Advisor policy using Group-Relative Policy Optimization (GRPO) (Shao et al., 2024). For each training example, the policy samples a group of candidate strategies. Each strategy is realized by the frozen rewriter, inserted back into the candidate list, and evaluated by the target agent to obtain r_{select} . See Appendix D for background on GRPO and related optimization concepts.

3.3 Factuality-Constrained Optimization

Naively optimizing for selection rate risks reward hacking, where the rewriter might generate hyperbolic or unsupported claims to attract the target agent. To enforce fidelity to the source material, we introduce a factuality constraint using MiniCheck (Tang et al., 2024a), a lightweight textual entailment verifier.

We apply MiniCheck at the sentence level. Let the original target item $x_{i,target}$ serve as the grounding document, and let the rewritten target $x'_{i,target}$ be segmented into a set of claim sentences $\{c_1, \dots, c_m\}$. MiniCheck computes a support score $S_{MC}(c_j|x_{i,target})$ for each sentence. We define the overall support reward $r_{support}$ as the minimum sentence-level score to heavily penalize even a single unsupported claim:

$$r_{support} = \min_{j \in \{1, \dots, m\}} S_{MC}(c_j|x_{i,target}) \quad (2)$$

The final training objective combines the selection reward and the support reward:

$$R_{total} = r_{select} + \lambda \cdot r_{support} \quad (3)$$

We use $\lambda = 1$ in most experiments.

3.4 Experimental Design

Datasets We evaluate our framework across two distinct information-access domains: news recommendation and scientific-document selection. **(1) News Recommendation:** We utilize impression-level candidate lists from the MIND benchmark (Wu et al., 2020) and EB-NeRD (Kruse et al., 2024). Our primary training set consists of 8,000 MIND

impressions, capped at a maximum of 15 candidate items per list. To prevent overfitting to specific content, individual target documents are restricted to a maximum of two appearances. Held-out evaluation is conducted on 1,000 unseen MIND impressions.

(2) Scientific Document Selection: We derive candidate lists from the SciRepEval search benchmark (Singh et al., 2023). Our evaluation set contains 1,000 candidate lists with a maximum list size of 10. Further details regarding preprocessing, candidate filtering, and leakage controls are provided in Appendix B.

Models The trainable Advisor policy is initialized with Qwen3.5-9B (Qwen Team, 2026). The frozen rewriter and the default target agent are instantiated using GPT-5-mini (OpenAI, 2025). For agent-transfer evaluation, we replace the target agent with frontier models, including GPT-5.5 (OpenAI, 2026), Gemini 3 Flash, Gemini 3.1 Pro (Google, 2026a,b), Claude Sonnet 4.6, and Claude Opus 4.8 (Anthropic, 2026b,a). Fixed prompts for all models are detailed in Appendix C.

Reinforcement Learning Setup We implement GRPO using the SkyRL library (Griggs et al., 2025) and vLLM (Kwon et al., 2023) on 4 NVIDIA A100-SXM4-80GB GPUs. The primary Advisor policy is trained for 250 optimization steps (approximately 1 epoch) using a batch size of 32, a policy mini-batch size of 8, and 4 sampled rollouts per prompt. Optimization is performed with a learning rate of 3×10^{-7} , a sampling temperature of 1.0, a KL penalty coefficient of 0.0075, and 10 warm-up steps. Complete implementation details and settings are provided in Appendix D.

Factuality Verification For factuality-constrained optimization, we utilize MiniCheck-Flan-T5-Large (Tang et al., 2024b) to compute the dense factuality reward. Because incorporating this continuous auxiliary reward increases variance, the factuality-constrained variant is trained for 500 steps (two epochs) to ensure convergence.

4 Results

Our primary results are presented in Table 2. First, we observe that target selection shifts even before reinforcement learning is applied. The original target-selection rate is 17.1%, close to the row-wise random baseline for this held-out set (16.9%; Table B.1), indicating that designated targets are not systematically advantaged before in-

Method	Target agent	Held-out agents				
	GPT-5-mini	GPT-5.5	Gemini 3 Flash	Gemini 3.1 Pro	Sonnet 4.6	Opus 4.8
Original target	17.1	17.1	17.1	17.3	17.6	17.5
<i>Prompting-only</i>						
Rewriter-only	34.8 (+17.7)	24.1 (+7.0)	40.7 (+23.6)	20.2 (+2.9)	24.4 (+6.8)	41.0 (+23.5)
Advisor-rewriter	43.9 (+26.8)	39.4 (+22.3)	48.8 (+31.7)	27.0 (+9.7)	34.2 (+16.6)	51.5 (+34.0)
<i>RL-trained rewriter</i>						
Rewriter-only	95.9 (+78.8)	85.3 (+68.2)	97.4 (+80.3)	54.8 (+37.5)	49.2 (+31.6)	75.9 (+58.4)
+ MiniCheck	43.4 (+26.3)	26.3 (+9.2)	39.7 (+22.6)	18.6 (+1.3)	24.1 (+6.5)	43.9 (+26.4)
<i>RL-trained advisor</i>						
Advisor-rewriter	98.5 (+81.4)	93.3 (+76.2)	98.9 (+81.8)	78.7 (+61.4)	65.0 (+47.4)	94.1 (+76.6)
+ MiniCheck	68.6 (+51.5)	53.2 (+36.1)	65.8 (+48.7)	37.5 (+20.2)	41.0 (+23.4)	68.8 (+51.3)

Table 2: **Target selection rates (%) on 1,000 unseen English news impressions.** In advisor-rewriter settings, the trained advisor is Qwen3.5-9B and the frozen rewriter is GPT-5-mini. In trained standalone rewriter settings, Qwen3.5-9B directly generates the rewritten title and abstract. The advisor is trained against GPT-5-mini and evaluated on held-out impressions whose target documents do not appear in training. We also evaluate transfer to other target agents without additional training. Values in parentheses report absolute target selection rate gains in percentage points over the original text under the same target agent. The row-wise uniform random baseline is 16.9%. Darker cells indicate higher selection rates.

tervention. A prompt-only rewrite of the target title and abstract increases this rate to 34.8% (untrained rewriter-only) and 43.9% (untrained advisor), demonstrating that target agents are vulnerable to simple presentation changes. The trained rewriter-only model achieves a 95.9% selection rate, while the trained advisor reaches 98.5%, corresponding to an 81.4 percentage-point improvement over the original text.

Furthermore, these gains transfer to other target agents; although trained exclusively with GPT-5-mini, the trained advisor consistently achieves the highest selection rates across held-out agents from the GPT, Gemini, and Claude model families.

Finally, while introducing factuality constraints (via MiniCheck) predictably lowers absolute selection rates due to the additional factuality filter, the constrained advisor still maintains consistent target selection improvements over the original text baseline for every evaluated agent.

4.1 Cross-Lingual Transfer

To evaluate whether the learned policy transfers beyond the English training distribution, we test the English MIND-trained advisor across five further languages—Arabic, Spanish, Swahili, Turkish, and Chinese—translated from the original test set using the Google Cloud Translation API, without additional training (Table 3). See Appendix I for further translation details and examples.

To prevent language drift, we add a single instruction ensuring the rewrite matches the input language. We observe cumulative gains across all stages. The rewriter-only setting consistently outperforms the original text, and the untrained advisor provides further improvements. The strongest effects appear after training: despite being optimized exclusively on English data, the trained advisor significantly increases target selection across all evaluation languages. This demonstrates that advisor training captures transferable, selection-oriented rewriting strategies rather than language-specific memorization.

Transfer to Danish News. We further evaluate the English-trained advisor on EB-NeRD, a Danish news recommendation dataset constructed from Ekstra Bladet impression logs, to test generalization across a distinct national media context and topic distribution (Table 4). We use both the original Danish version and an English translation of the same test candidate lists by sampling 1,000 test impressions from records with at most 26 candidates. Treated targets are sampled from each candidate slate with a cap of two treated occurrences per document. The resulting Danish EB-NeRD set has 886 unique treated documents, 114 repeated treated documents, and a maximum treated-target repeat count of two. We also use an English machine-translated EB-NeRD variant. The Danish and English EB-NeRD versions preserve the same rows,

Language	Original target	Prompting-only		RL-trained		RL w/ MiniCheck	
		Rewriter-only	Advisor-rewriter	Rewriter-only	Advisor-rewriter	Rewriter-only	Advisor-rewriter
<i>Training language</i>							
English (en)	17.1	34.8 (+17.7)	43.9 (+26.8)	95.9 (+78.8)	98.5 (+81.4)	43.4 (+26.3)	68.6 (+51.5)
<i>Cross-lingual transfer</i>							
Arabic (ar)	16.1	28.9 (+12.8)	39.9 (+23.8)	81.0 (+64.9)	95.5 (+79.4)	33.6 (+17.5)	64.5 (+48.4)
Spanish (es)	17.8	31.2 (+13.4)	42.8 (+25.0)	85.8 (+68.0)	98.0 (+80.2)	35.7 (+17.9)	66.5 (+48.7)
Swahili (sw)	17.6	23.8 (+6.2)	39.4 (+21.8)	27.8 (+10.2)	93.7 (+76.1)	27.2 (+9.6)	62.2 (+44.6)
Turkish (tr)	17.5	30.3 (+12.8)	39.2 (+21.7)	86.7 (+69.2)	95.8 (+78.3)	31.1 (+13.6)	59.9 (+42.4)
Chinese (zh-CN)	17.3	33.2 (+15.9)	43.4 (+26.1)	87.1 (+69.8)	96.6 (+79.3)	36.8 (+19.5)	67.4 (+50.1)
Average	17.3	29.5 (+12.2)	40.9 (+23.7)	73.7 (+56.4)	95.9 (+78.7)	32.9 (+15.6)	64.1 (+46.8)

Table 3: **Target selection rates on MIND and cross-lingual transfer.** All values are percentages. The advisor is trained on English MIND. All other languages evaluate cross-lingual transfer. Values in parentheses report absolute target selection-rate gains in percentage points over the original text. All translated variants preserve the same candidate lists, so their row-wise uniform random baseline is 16.9%. Darker cells indicate higher selection rates.

Dataset	Language	Original target	Prompting-only		RL-trained		RL w/ MiniCheck	
			Rewriter-only	Advisor-rewriter	Rewriter-only	Advisor-rewriter	Rewriter-only	Advisor-rewriter
<i>Training dataset</i>								
MIND	English	17.1	34.8 (+17.7)	43.9 (+26.8)	95.9 (+78.8)	98.5 (+81.4)	43.4 (+26.3)	68.6 (+51.5)
<i>Transfer to other datasets and languages</i>								
MIND	Danish	16.5	31.1 (+14.6)	40.9 (+24.4)	95.0 (+78.5)	97.9 (+81.4)	35.6 (+19.1)	65.0 (+48.5)
EB-NeRD	English	12.8	43.8 (+31.0)	57.7 (+44.9)	93.1 (+80.3)	98.1 (+85.3)	50.2 (+37.4)	81.2 (+68.4)
EB-NeRD	Danish	11.6	32.2 (+20.6)	47.8 (+36.2)	86.1 (+74.5)	98.2 (+86.6)	40.4 (+28.8)	71.7 (+60.1)

Table 4: **Transfer across English and Danish news datasets.** The advisor is trained on English MIND. Values are target selection rates in percentages, with values in parentheses reporting absolute gains over the original text in percentage points. The row-wise uniform random baseline is 16.9% for MIND and 12.4% for EB-NeRD. Darker cells indicate higher selection rates.

impression IDs, treated article IDs, candidate slates, candidate order, and row alignment; only title and abstract text are translated.

The unconstrained advisor successfully overcomes both language and dataset shifts, achieving large selection gains on MIND-Da, EB-Da, and EB-En. This uniform performance indicates that the learned rewriting strategy is robust to both language translation and out-of-distribution content rather than being an artifact of the English MIND training distribution. Under the factuality-constrained setting, using an advisor has a clear advantage over the rewriter-only alternative across all four evaluation environments.

4.2 Cross-Domain Transfer

To evaluate whether the learned strategy transfers beyond news, we test the English MIND-trained advisor on academic document selection (Figure 3). To do this, we use search candidate lists derived

from Semantic Scholar clickthrough data, following the SciRepEval search setup (subsection 3.4). Rewriting alters the LLM agent’s decisions even before optimization, as both untrained rewriters substantially increase selection rates over the original text. Direct rewriter-only training on news fails to improve academic transfer, resulting in a slight performance decrease. In contrast, advisor training yields a large additional gain over its untrained baseline. This divergence suggests that advisor training learns transferable, selection-oriented guidance, whereas direct rewriter-only training primarily captures news-specific formatting and behavior.

4.3 Selection Hacking

High target selection does not imply that a rewrite is faithful to the original item, nor does it guarantee human usefulness, paralleling the known gap between clicks and post-click satisfaction in human-facing recommendation (Wang et al., 2021). To test




Figure 3: Selection rates on scientific-document impressions. MC: MiniCheck.

this, we evaluate whether rewritten titles and abstracts remain supported by the original target text using a post-hoc MiniCheck-Flan factuality check (Table 5). We can see from this experiment that unconstrained selection optimization frequently degrades source faithfulness, as models exploit unfaithful shortcuts and agent-legible presentation styles to maximize target selection. While applying factuality constraints mitigates this collapse at the cost of overall selection rates, these dynamics underscore the necessity of evaluating each axis as a distinct metric when ranking outcomes. Additional analyses of source-support reward weights and evaluations with independent source-support verifiers are reported in Appendix E.7 and Appendix F, respectively.

4.4 What Strategies Do Advisors Use?

One of the primary reasons for using the advisor model, rather than a black-box rewriter, is the interpretability afforded to the individual generated strategies. To explore this, we grouped advisor-generated strategies into the families shown in Figure 4. We then used StringSight-assisted inspection (Dunlap, 2025), followed by manual review, to identify recurring strategy types and summarize the distribution changes across training settings.

The results show that optimizing for a target agent does not teach the model a single, universal concept of good writing. Instead, it uncovers a wide spectrum of strategies. Some of these strategies act like traditional editing—focusing on clarity

	MIND-En	Academic	EB-En
<i>Prompting-only</i>			
Reewriter-only	60.7	68.2	53.4
Advisor-rewriter	42.2	53.0	37.9
<i>RL-trained rewriter</i>			
Reewriter-only	2.2	73.1	2.6
+ MiniCheck	66.7	83.4	63.8
<i>RL-trained advisor</i>			
Advisor-rewriter	2.0	24.1	1.8
+ MiniCheck	31.2	50.3	27.9

Table 5: MiniCheck sentence-support scores (%).

and detail, while others amount to structural shortcuts designed purely to exploit the evaluator’s blind spots. For example, when left completely unconstrained, the advisor trained against GPT-5-mini defaults to a shortcut strategy: it forces a technical topic into the text 96.2% of the time, regardless of whether that topic fits the source document. Rather than improving the writing itself, the model simply alters the content to match a specific pattern it knows the target agent prefers. Introducing a factuality constraint forces the model to take a different path. Here, the unfaithful technical shortcut drops to just 0.1%, replaced by more legitimate improvements like detail expansion (45.6%) and clearer framing of the story’s stakes (28.3%). Interestingly, different underlying models exhibit distinct optimization styles: the Gemini advisor favors adding a more formal, academic tone.

Figure 5 provides a complementary lexical diagnostic: unconstrained training shifts advisor strategies toward technical and research-oriented terms, whereas factuality-constrained training yields more news-compatible framing terms. Figure 6 gives a qualitative example of this shift: the unconstrained advisor wins selection by introducing unsupported content, whereas the MiniCheck-constrained advisor preserves the original factual core while still making the item more attractive to the target agent. Additional examples and taxonomy details are provided in Appendices G and H.

4.5 Fixed Global Strategy Baselines

The distribution of strategy families above raises a natural question: whether the gains from advisor training can be explained by a single reusable rewriting strategy. We therefore compare the learned advisor with five hand-written global strategies and a global strategy optimized by GEPA (Agrawal et al., 2026). GEPA is a reflective prompt




Figure 4: **Strategy-family distribution across advisor variants.** Each bar shows the percentage distribution of strategy families over 1,000 advisor outputs for a given condition. Remaining low-frequency strategy families are grouped into “Other strategies.”



Figure 5: **Quantitative lexical diagnostic of advisor strategy outputs on the English evaluation set.** Each panel shows frequent non-stopword terms from 1,000 advisor strategies, with font size proportional to term frequency. From left to right: (a) the untrained Qwen3.5-9B advisor, (b) the unconstrained trained advisor, and (c) the factuality-constrained trained advisor. Unconstrained training shifts strategy language toward technical and research registers, while factuality-constrained strategies use more news-compatible framing terms. Rewriter-only models have no advisor strategies.

Method	GPT-5-mini	Gemini 3 Flash	Sonnet 4.6
Original	17.1	17.1	17.6
Untrained advisor	43.9	48.8	34.2
Best hand-written	72.4	71.1	36.2
GEPA	77.5	51.7	38.4
RL advisor	98.5	98.9	65.0

Table 6: **Advisor training versus fixed global strategies.** Target-selection rates (%) on held-out English MIND impressions. “Best hand-written” reports the strongest of five fixed global strategies for each target agent. GEPA optimizes a single global strategy, whereas the RL advisor generates input-dependent strategies. Darker cells indicate higher selection rates.

optimizer that searches over instructions using downstream task feedback while keeping model weights fixed. In our setting, we use GEPA to search for a single global strategy instruction that is then reused for every target item. As shown in Table 6, fixed global strategies produce substantial selection gains, but remain below the RL-trained advisor across all three evaluated target agents. On GPT-5-mini, the GEPA-optimized global strategy reaches 77.5% target selection, compared with

98.5% for the RL-trained advisor. The gap also persists under Gemini 3 Flash and Sonnet 4.6. These results indicate that the gains from advisor training cannot be reduced to a single reusable global strategy instruction. Full results for all hand-written strategies and the selected GEPA instruction are reported in Appendix E.3.

5 Conclusion

In this paper we present **AgentBait**, an advisor-rewriter framework that optimizes the presentation of content to measure and mitigate presentation-induced selection shifts in language-model-mediated selection. Our results demonstrate that target agents are vulnerable to localized textual rewrites, with rewritten texts achieving large selection gains that transfer across model families, languages, and content domains. While unconstrained optimization frequently leads to unfaithful selection hacking, we further show that incorporating a factuality constraint redirects the learned rewriting policy toward grounded strategies while maintaining a substantial selection advantage. Overall, our findings show that language-model-mediated




Figure 6: **Qualitative comparison between unconstrained and factuality-constrained advisor rewrites.** The candidate list is fixed, and only the target item’s title and abstract are rewritten. Abstracts are lightly abbreviated for space. Both rewrites cause the GPT-5-mini target agent to select the target. The unconstrained TRAINED ADVISOR attributes the airport’s punctuality to an unsupported AI scheduling system, whereas TRAINED ADVISOR + MINICHECK preserves the original factual core while reframing the item around the operational puzzle.

selection can be shaped by presentation alone, underscoring the need to evaluate AI agents across factuality, robustness, and user-aligned utility before deployment.

Limitations

First, real interfaces contain layout, ranking position, images, thumbnails, metadata, and interaction affordances that can affect both human and agent choices. We focus only on the textual presentation of an item, namely its title and abstract, and therefore do not model the full multimodal and interface-dependent nature of selection.

Second, and more importantly, our experiments assume that the candidate list is already fixed before rewriting. The target agent receives only the candidate titles and abstracts and no explicit user query. This isolates a query-free post-retrieval selection problem, but it does not test whether the rewritten item would still be retrieved or ranked into the same list by an upstream retrieval system. In a full pipeline, rewriting could change lexical or semantic matching and may therefore help or hurt retrieval under the original query or recommendation context. Our findings should therefore be interpreted as evidence of query-free selection effects conditional on exposure, not as evidence of query relevance or end-to-end optimization over retrieval, ranking, and selection.

Third, although the impressions come from real logs, we treat each impression as a fixed snapshot.

Real information-access systems are dynamic and non-stationary: articles enter and leave the corpus, competing stories evolve, user interests drift, social trends change, and platform feedback affects future retrieval and ranking. The selection pressures observed in a fixed candidate list may therefore change over time. Studying how rewriting incentives evolve under changing candidate pools and long-run feedback loops is important future work.

We do not collect human judgments of relevance, quality, or usefulness, so higher target selection should not be interpreted as evidence of user benefit. Dataset access conditions, release restrictions, privacy safeguards, and artifact release policies are documented in Appendix J.

Acknowledgments

As part of their affiliation with UC Berkeley, the authors were supported in part by the U.S. Department of War, the Berkeley Artificial Intelligence Research (BAIR) Industrial Alliance program, and gifts from Accenture, AMD, Anyscale, Broadcom, Cisco, Google, IBM, Intel, Intesa Sanpaolo, Lambda, Lightspeed, Mibura, Microsoft, NVIDIA, Qualcomm, Samsung SDS, SAP, and Veridian Technologies. This material is based upon work supported by the Defense Advanced Research Projects Agency and the Air Force Research Laboratory under contract FA8650-23-C-7316. The authors also acknowledge VESSL AI for providing compute resources. We thank Sewon Min and

Alane Suhr for helpful feedback on an earlier version of this work as part of UC Berkeley CS 288. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of AFRL or DARPA.

The authors used AI-based tools to assist with code generation, editing, and writing during the preparation of this paper. Specifically, AI assistance was used to help draft and revise portions of the manuscript for clarity, grammar, and organization, and to support the development, debugging, and refinement of code used in the research workflow. All AI-generated or AI-assisted content, code, analyses, and interpretations were reviewed, verified, and, where necessary, modified by the authors. The authors take full responsibility for the accuracy, integrity, originality, and final content of the paper, including any code or text developed with AI assistance.

References

- Abdelrahman Abdallah, Bhawna Piryani, Jamshid Mozafari, Mohammed Ali, and Adam Jatowt. 2025. [How good are LLM-based rerankers? an empirical analysis of state-of-the-art reranking models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5693–5709, Suzhou, China. Association for Computational Linguistics.
- Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet Deshpande. 2024. [GEO: Generative engine optimization](#). In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, pages 5–16. ACM.
- Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alex Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. 2026. [GEPA: Reflective prompt evolution can outperform reinforcement learning](#). In *International Conference on Learning Representations*, volume 2026, pages 8479–8565.
- Anthropic. 2026a. Introducing claude opus 4.8. <https://www.anthropic.com/news/claude-opus-4-8>. Accessed: 2026-07-21.
- Anthropic. 2026b. Introducing Claude Sonnet 4.6. <https://www.anthropic.com/news/claude-sonnet-4-6>. Anthropic model announcement.
- Sagnik Anupam, Davis Brown, Shuo Li, Eric Wong, Hamed Hassani, and Osbert Bastani. 2025. [Browserarena: Evaluating llm agents on real-world web navigation tasks](#). *ArXiv preprint*, abs/2510.02418.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Parth Asawa, Alan Zhu, Abigail O'Neill, Matei Zaharia, Alexandros G. Dimakis, and Joseph E. Gonzalez. 2026. [How to train your advisor: Steering black-box LLMs with advisor models](#). In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306, Seoul, South Korea.
- Puneet S Bagga, Vivek F Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. 2025. [E-geo: A testbed for generative engine optimization in e-commerce](#). *ArXiv preprint*, abs/2511.20867.
- Bespoke Labs. 2024. Bespoke-MiniCheck-7B. <https://huggingface.co/bespokelabs/Bespoke-MiniCheck-7B>. Hugging Face model card. Accessed: 2026-05-24.
- Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. [SPECTER: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282, Online. Association for Computational Linguistics.
- Paul Covington, Jay Adams, and Emre Sargin. 2016. [Deep neural networks for youtube recommendations](#). In *Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15-19, 2016*, pages 191–198. ACM.
- Yixuan Du, Chenxiao Yu, Haoyan Xu, Ziyi Wang, Yue Zhao, and Xiyang Hu. 2026. [Multimodal generative engine optimization: Rank manipulation for vision-language model rankers](#). In *Proceedings of the 4th Workshop on Towards Knowledgeable Foundation Models (KnowFM 2026)*, pages 115–128, San Diego, California, United States. Association for Computational Linguistics.
- Lisa Dunlap. 2025. Stringsight. <https://github.com/lisadunlap/StringSight>. Software. Accessed: 2026-08-28.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.
- Google. 2026a. Gemini 3 Flash Preview. <https://ai.google.dev/gemini-api/docs/models/gemini-3-flash-preview>. Google AI for Developers model documentation; model code: gemini-3-flash-preview. Accessed: 2026-08-28.
- Google. 2026b. Gemini 3.1 Pro Preview. <https://ai.google.dev/gemini-api/docs/models/gemini-3.1-pro>.

- [ini-3.1-pro-preview](#). Google AI for Developers model documentation; model code: [gemini-3.1-pro-preview](#). Accessed: 2026-08-28.
- Tyler Griggs, Sumanth Hegde, Eric Tang, Shu Liu, Shiyi Cao, Dacheng Li, Charlie Ruan, Philipp Moritz, Kourosh Hakhamaneshi, Richard Liaw, Akshay Malik, Shishir G. Patil, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025. [Evolving skylr into a highly-modular rl framework](#). NovaSky blog, posted June 26, 2025.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43.
- Tanqiu Jiang, Min Bai, Nikolaos Pappas, Yanjun Qi, and Sandesh Swamy. 2025. [Cross-modal content optimization for steering web agent preferences](#). *ArXiv preprint*, abs/2510.03612.
- Haibo Jin, Ruoxi Chen, Peiyan Zhang, Yifeng Luo, Huimin Zeng, Man Luo, and Haohan Wang. 2026. [Controlling output rankings in generative engines for llm-based search](#). *ArXiv preprint*, abs/2602.03608.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. [Colbert: Efficient and effective passage search via contextualized late interaction over BERT](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25–30, 2020*, pages 39–48. ACM.
- Johannes Kruse, Kasper Lindskow, Saikishore Kalloori, Marco Polignano, Claudio Pomo, Abhishek Srivastava, Anshuk Uppal, Michael Riis Andersen, and Jes Frellsen. 2024. [EB-NeRD: A large-scale dataset for news recommendation](#). In *Proceedings of the Recommender Systems Challenge 2024*, pages 1–11. Association for Computing Machinery.
- Aounon Kumar and Himabindu Lakkaraju. 2024. Manipulating large language models to increase product visibility. *arXiv preprint arXiv:2404.07981*.
- Komal Kumar, Aman Chadha, Salman Khan, Fahad Shahbaz Khan, and Hisham Cholakkal. 2026. [Paper circle: An open-source multi-agent research discovery and analysis framework](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 42170–42188, San Diego, California, United States. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Long Li, Weiwen Xu, Jiayan Guo, Ruochen Zhao, Xingxuan Li, Yuqian Yuan, Boqiang Zhang, Yuming Jiang, Yifei Xin, Ronghao Dang, Yu Rong, Deli Zhao, Tian Feng, and Lidong Bing. 2025. [Chain of ideas: Revolutionizing research via novel idea development with LLM agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 8971–9004, Suzhou, China. Association for Computational Linguistics.
- Xinyi Li, Yongfeng Zhang, and Edward C. Malthouse. 2023. [A preliminary study of ChatGPT on news recommendation: Personalization, provider fairness, and fake news](#). In *Proceedings of the International Workshop on News Recommendation and Analytics*, volume 3561 of *CEUR Workshop Proceedings*, Singapore. CEUR-WS.org.
- Zeyi Liao, Lingbo Mo, Chejian Xu, Mintong Kang, Jiawei Zhang, Chaowei Xiao, Yuan Tian, Bo Li, and Huan Sun. 2025. [EIA: Environmental injection attack on generalist web agents for privacy leakage](#). In *International Conference on Learning Representations*, volume 2025, pages 66972–67003.
- Dairui Liu, Boming Yang, Honghui Du, Derek Greene, Neil Hurley, Aonghus Lawlor, Ruihai Dong, and Irene Li. 2024. [RecPrompt: A self-tuning prompting framework for news recommendation using large language models](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, pages 3902–3906. ACM.
- Nelson Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7001–7025, Singapore. Association for Computational Linguistics.
- Zikang Liu and Peilan Xu. 2026. [Think before writing: Feature-level multi-objective optimization for generative citation visibility](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20290–20303, San Diego, California, United States. Association for Computational Linguistics.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse,

- Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2022. [Webgpt: Browser-assisted question-answering with human feedback](#). *Preprint*, arXiv:2112.09332.
- OpenAI. 2025. GPT-5 mini Model. <https://developers.openai.com/api/docs/models/gpt-5-mini>. OpenAI API model documentation.
- OpenAI. 2026. GPT-5.5 Model. <https://developers.openai.com/api/docs/models/gpt-5.5>. OpenAI API model documentation.
- Nilay Patel, Shivashankar Subramanian, Siddhant Garg, Pratyay Banerjee, and Amita Misra. 2024. [Towards improved multi-source attribution for long-form answer generation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3906–3919, Mexico City, Mexico. Association for Computational Linguistics.
- Samuel Pfommer, Yatong Bai, Tanmay Gautam, and Somayeh Sojoudi. 2024. [Ranking manipulation for conversational search engines](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9523–9552, Miami, Florida, USA. Association for Computational Linguistics.
- Ronak Pradeep, Sahel Sharifmoghadam, and Jimmy Lin. 2023a. [Rankvicuna: Zero-shot listwise document reranking with open-source large language models](#). *ArXiv preprint*, abs/2309.15088.
- Ronak Pradeep, Sahel Sharifmoghadam, and Jimmy Lin. 2023b. [Rankzephyr: Effective and robust zero-shot listwise reranking is a breeze!](#) *ArXiv preprint*, abs/2312.02724.
- Jirui Qi, Gabriele Sarti, Raquel Fernández, and Arianna Bisazza. 2024. [Model internals-based answer attribution for trustworthy retrieval-augmented generation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6037–6053, Miami, Florida, USA. Association for Computational Linguistics.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. 2024. [Large language models are effective text rankers with pairwise ranking prompting](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1504–1518, Mexico City, Mexico. Association for Computational Linguistics.
- Qwen Team. 2026. [Qwen3.5: Towards native multi-modal agents](#).
- Claire E. Robertson, Nicolas Pröllochs, Kaoru Schwarzenegger, Philip Pärnamets, Jay J. Van Bavel, and Stefan Feuerriegel. 2023. [Negativity drives online news consumption](#). *Nature Human Behaviour*, 7:812–822.
- Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. 2024. [RAPTOR: recursive abstractive processing for tree-organized retrieval](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Alessandro Scirè, Karim Ghonim, and Roberto Navigli. 2024. [FENICE: Factuality evaluation of summarization based on natural language inference and claim extraction](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14148–14161, Bangkok, Thailand. Association for Computational Linguistics.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *ArXiv preprint*, abs/2402.03300.
- Sahel Sharifmoghadam, Ronak Pradeep, Andre Slavescu, Ryan Nguyen, Andrew Xu, Zijian Chen, Yilin Zhang, Yidi Chen, Jasper Xian, and Jimmy Lin. 2025. [RankLLM: A python package for reranking with LLMs](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’25*, pages 3681–3690. ACM.
- Jiawen Shi, Zenghui Yuan, Guiyao Tie, Pan Zhou, Neil Zhenqiang Gong, and Lichao Sun. 2026. [Prompt injection attack to tool selection in LLM agents](#). In *Proceedings 2026 Network and Distributed System Security Symposium, NDSS 2026, San Diego, CA, USA*. Internet Society.
- Amanpreet Singh, Mike D’Arcy, Arman Cohan, Doug Downey, and Sergey Feldman. 2023. [SciRepEval: A multi-format benchmark for scientific document representations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5548–5566, Singapore. Association for Computational Linguistics.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? investigating large language models as re-ranking agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937, Singapore. Association for Computational Linguistics.
- Liyan Tang, Philippe Laban, and Greg Durrett. 2024a. [MiniCheck: Efficient fact-checking of LLMs on grounding documents](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8818–8847, Miami, Florida, USA. Association for Computational Linguistics.

- Liyan Tang, Philippe Laban, and Greg Durrett. 2024b. MiniCheck-Flan-T5-Large. <https://huggingface.co/lytang/MiniCheck-Flan-T5-Large>. Hugging Face model card. Accessed: 2026-05-24.
- Yiming Tang, Yi Fan, Chenxiao Yu, Tiankai Yang, Yue Zhao, and Xiyang Hu. 2025. Stealthrank: Llm ranking manipulation via stealthy prompt optimization. *ArXiv preprint*, abs/2504.05804.
- Zhihua Tian, Yuhan Chen, Yao Tang, Jian Liu, and Ruoxi Jia. 2026. Diagnosing and repairing citation failures in generative engine optimization. *ArXiv preprint*, abs/2603.09296.
- Chong Wang and David M. Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, August 21-24, 2011*, pages 448–456. ACM.
- Wenjie Wang, Fuli Feng, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2021. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, pages 1288–1297. ACM.
- Zihan Wang, Rui Zhang, Yu Liu, Wenshu Fan, Wenbo Jiang, Qingchuan Zhao, Hongwei Li, and Guowen Xu. 2026. MPMA: Preference manipulation attack against model context protocol. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(42):35838–35846.
- Beining Wu, Fuyou Mao, Jiong Lin, Cheng Yang, Jiaxuan Lu, Yifu Guo, Siyu Zhang, Yifan Wu, Ying Huang, and Fu Li. 2026a. From experience to skill: Multi-agent generative engine optimization via reusable strategy learning. In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 43305–43315, San Diego, California, United States. Association for Computational Linguistics.
- Chuhan Wu, Fangzhao Wu, Suyu Ge, Tao Qi, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with multi-head self-attention. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6389–6394, Hong Kong, China. Association for Computational Linguistics.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020. MIND: A large-scale dataset for news recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606, Online. Association for Computational Linguistics.
- Yujia Wu, Shanshan Zhong, Yubin Kim, and Chenyan Xiong. 2026b. What generative search engines like and how to optimize web content cooperatively. In *International Conference on Learning Representations*, volume 2026, pages 136858–136888.
- Sirui Xia, Xintao Wang, Jiaqing Liang, Yifei Zhang, Weikang Zhou, Jiaji Deng, Fei Yu, and Yanghua Xiao. 2025. Ground every sentence: Improving retrieval-augmented LLMs with interleaved reference-claim generation. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 969–988, Albuquerque, New Mexico. Association for Computational Linguistics.
- Chejian Xu, Mintong Kang, Jiawei Zhang, Zeyi Liao, Lingbo Mo, Mengqi Yuan, Huan Sun, and Bo Li. 2025. AdvAgent: Controllable blackbox red-teaming on web agents. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 69318–69330. PMLR.
- Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen, Meng Fang, Muhao Chen, Shangding Gu, Ming Jin, Costas Spanos, Yang Yang, Pieter Abbeel, Dawn Song, Weinan Zhang, and Jun Wang. 2025. Agentic web: Weaving the next web with ai agents. *Preprint*, arXiv:2507.21206.
- Yue Yu, Wei Ping, Zihan Liu, Boxin Wang, Jiaxuan You, Chao Zhang, Mohammad Shoeybi, and Bryan Catanzaro. 2024. RankRAG: Unifying context ranking with retrieval-augmented generation in LLMs. In *Advances in Neural Information Processing Systems*, volume 37, pages 121156–121184. Curran Associates, Inc.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. AlignScore: Evaluating factual consistency with a unified alignment function. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.
- Weizhi Zhang, Yangning Li, Wei-Chieh Huang, Junyu Luo, Guancheng Wan, Liangwei Yang, Yuyao Yang, Chenxuan Xie, Chunyu Miao, Henry Peng Zou, Xiao Luo, Yusheng Zhao, Yankai Chen, Chunkit Chan, Peilin Zhou, Xinyang Zhang, Chenwei Zhang, Jingbo Shang, Ming Zhang, and 3 others. 2026. From web search towards agentic deep research: Incentivizing search with reasoning agents. In *4th Workshop on Towards Knowledgeable Foundation Models at ACL 2026*.
- Yanzhe Zhang, Tao Yu, and Diyi Yang. 2025. Attacking vision-language computer agents via pop-ups. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8387–8401, Vienna, Austria. Association for Computational Linguistics.
- Zexuan Zhong, Ziqing Huang, Alexander Wettig, and Danqi Chen. 2023. Poisoning retrieval corpora by

injecting adversarial passages. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13764–13775, Singapore. Association for Computational Linguistics.

Shengyao Zhuang, Honglei Zhuang, Bevan Koopman, and Guido Zuccon. 2024. [A setwise approach for effective and highly efficient zero-shot ranking with large language models](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 38–47. ACM.

Appendix

The appendix provides additional details organized as follows:

- **Rewriting variants and intervention settings.** [Appendix A](#) describes the rewriting variants considered in our experiments, including standalone and advisor-guided rewriting, and the fixed-list intervention setting used for conditional exposure evaluation.
- **Data construction and dataset statistics.** [Appendix B](#) provides detailed data construction workflows, filtering criteria, candidate-list statistics, and dataset split statistics for MIND, EB-NeRD, and SciRepEval.
- **Prompts and output formats.** [Appendix C](#) presents the complete evaluation prompts, candidate rendering formats, and output constraints.
- **Reproducibility details and optimization rationale.** [Appendix D](#) reports the implementation and training details needed to reproduce our experiments, including policy optimization hyperparameters, GEPA budget choices, sampling and reward-evaluation budgets, validation/model-selection procedures, and hardware and software settings.
- **Additional Baselines and Robustness Analyses.** [Appendix E](#) reports chooser-prompt and inference-setting robustness, original-target and random-choice diagnostics, fixed-strategy baselines, support-reward sensitivity, and additional advisor and rewriter ablations.
- **Verifier robustness and factuality diagnostics.** [Appendix F](#) evaluates verifier coupling with MiniCheck, Bespoke-MiniCheck-7B, AlignScore, and FENICE and reports additional source-support diagnostics.
- **Strategy taxonomy and advisor behavior.** [Appendix G](#) describes the qualitative strategy taxonomy, reports family distributions across optimized advisor variants, and analyzes differences between unconstrained and factuality-constrained strategies.
- **Qualitative case studies.** [Appendix H](#) presents additional case studies covering unconstrained rewrites, factuality-aware constrained rewrites, cross-domain examples, and typical failure modes.
- **Translation protocol and multilingual examples.** [Appendix I](#) reports the machine translation protocols used for cross-lingual evaluation and provides representative multilingual examples.
- **Data sources, licenses, and release policy.** [Appendix J](#) documents dataset sources, data-use alignment, user privacy safeguards, licenses, and asset release policies.

A Rewriting Variants and Intervention Settings

This section defines the rewriting settings compared in [section 3](#). All settings operate on the same fixed candidate list: only the target item may be rewritten, and all non-target items remain unchanged.

A.1 Original Text

The target agent receives the original candidate list without rewriting. This setting provides the reference target selection rate.

A.2 Rewriter-Only

In the standalone rewriter setting, the rewriter directly generates the rewritten title and abstract for the target item. We compare an untrained standalone rewriter, which uses a frozen rewriter model with a fixed prompt, and a trained standalone rewriter, which directly optimizes the rewriter using the target agent reward.

A.3 Untrained Advisor

An untrained advisor first produces a rewriting strategy for the target item. The same frozen rewriter then uses this strategy to generate the rewritten title and abstract. This baseline controls for the advisor–rewriter pipeline without selection-based learning.

A.4 Advisor-Guided Rewriter

In the advisor-guided setting, an advisor first produces a rewriting strategy for the target item, and a frozen rewriter then uses this strategy to generate the rewritten title and abstract.

We compare an untrained advisor, which controls for the advisor–rewriter pipeline without selection-based learning, and a trained advisor, which optimizes only the advisor using the target agent reward.

B Data Construction and Dataset Statistics

We use MINDlarge ([Wu et al., 2020](#)), EB-NeRD ([Kruse et al., 2024](#)), and SciRepEval-derived scientific-document candidate sets ([Singh et al., 2023](#)). We also construct translated evaluation-only variants of held-out MIND and EB-NeRD slates using Google Cloud Translation Advanced

v3. This appendix summarizes the data sources, derived evaluation settings, intended use, and dataset statistics used in our experiments.

Random baseline. For each candidate list, exactly one item is designated as the treated target. The random baseline reports the expected target-selection rate of a uniformly random chooser that selects one item from each fixed candidate list. For candidate list L_i , the probability of selecting the treated target under uniform random choice is $1/|L_i|$. Since candidate-list sizes vary across rows, we compute the random baseline as the row-wise average:

$$\text{Random} = \frac{1}{N} \sum_{i=1}^N \frac{1}{|L_i|}.$$

This differs from the inverse of the mean candidate-list size. The baseline provides a chance-level reference for how often the designated target would be selected in the absence of any model preference or presentation effect. An original target-selection rate close to this baseline indicates that the designated targets are not systematically advantaged under their unmodified presentations.

Dataset	Rows	Mean	Max	Random
MIND test	1,000	8.42	15	16.9%
EB-NeRD	1,000	9.93	26	12.4%
SciRepEval	1,000	9.29	10	11.2%

Table B.1: **Evaluation dataset statistics.** Rows are impression-level candidate lists. Mean and max report candidate-list size. Random reports the row-wise uniform target-selection baseline defined above. For MIND, the English test set and its Arabic, Danish, Spanish, Swahili, Turkish, and Chinese translations use the same candidate lists. For EB-NeRD, the Danish set and its English machine-translated version use the same candidate lists.

B.1 MIND

We use MINDlarge as the source dataset, but construct our own train, validation, and held-out evaluation splits rather than using the original release split directly. MINDlarge is collected from recommendation logs and has a long-tailed document-exposure distribution: a small number of news articles appear in many impressions, while many articles appear only rarely. This is appropriate for studying recommendation logs, but it is not ideal for our fixed-slate rewrite intervention setting,

where we want to measure presentation-induced selection changes rather than repeated exposure to a few high-frequency articles.

We therefore construct the primary news training and evaluation settings from MINDlarge impression logs in two stages. First, we build a doc-capped impression pool from the raw impression logs. Second, we construct fixed-slate rewrite examples by sampling one treated target from each retained impression.

Doc-capped impression-pool construction. To reduce domination by highly repeated documents, we first construct a more diverse impression pool using a capacity-constrained bipartite matching procedure.

Let $\mathcal{I}$ denote the set of impression rows and let $\mathcal{D}$ denote the set of news documents. Each impression $i \in \mathcal{I}$ has a candidate set $C_i \subseteq \mathcal{D}$. We build a bipartite graph

$$G = (\mathcal{D}, \mathcal{I}, E), \quad E = \{(d, i) : d \in C_i\}.$$

The goal is to retain as many impression rows as possible while assigning each retained row to one representative document and limiting how often any document can be assigned. With document usage cap K , we solve

$$\max_z \sum_{i \in \mathcal{I}} \sum_{d \in C_i} z_{d,i}$$

subject to

$$\begin{aligned} \sum_{d \in C_i} z_{d,i} &\leq 1 & \forall i \in \mathcal{I}, \\ \sum_{i: d \in C_i} z_{d,i} &\leq K & \forall d \in \mathcal{D}, \\ z_{d,i} &\in \{0, 1\} & \forall (d, i) \in E. \end{aligned}$$

Here $z_{d,i} = 1$ means that impression i is retained and assigned to document d . This is a bipartite b -matching problem. We implement it by duplicating each document node K times and then applying ordinary Hopcroft–Karp maximum bipartite matching on the expanded graph:

$$d \mapsto \{d^{(1)}, d^{(2)}, \dots, d^{(K)}\}.$$

Mapping matched duplicate nodes back to their original document IDs gives an assignment

$$i \mapsto a_i, \quad a_i \in C_i,$$

where each document appears as an assigned representative document at most K times.

In our run, we set the upstream document usage cap to $K = 10$. This cap controls representative-document reuse during impression-pool construction. It is separate from the treated-target caps used when constructing the final rewrite examples. The matched-impression files preserve the full original candidate list for each retained impression, so this preprocessing step should not be interpreted as constructing candidate-pool-disjoint splits.

Final rewrite-example construction. The matched-impression pool is then used as input to the rewrite dataset builder. The downstream builder does not directly use the matched assigned document as the treated target. Instead, for each retained impression, it keeps the full candidate slate and samples one treated candidate from that slate.

For an impression i with candidate list

$$L_i = (d_{i1}, d_{i2}, \dots, d_{im_i}),$$

we retain the row only if $m_i \leq 15$ for the MIND setting. We do not truncate candidate lists. All retained candidates remain in the slate in their original order. A treated document $t_i \in L_i$ is selected uniformly at random subject to a downstream treated-target reuse cap:

$$\begin{aligned} t_i &\sim \text{Uniform}(L_i), \\ \sum_i \mathbf{1}[t_i = d] &\leq R & \forall d \in \mathcal{D}. \end{aligned}$$

The advisor-training split uses 8,000 MIND impressions and the validation split uses 500 impressions. For these splits, we use $R = 2$, so a document can appear as the treated target at most twice. The final held-out MIND evaluation set uses 1,000 test impressions with the stricter cap $R = 1$, so no evaluation target repeats.

Each constructed row contains one impression, one treated target, the treated target’s title and abstract, and the full candidate slate in the original order. The top-level title, abstract, and candidate_id fields correspond to the treated candidate. The full fixed chooser slate is stored in reward_impression_candidates. During evaluation, the rewritten title and abstract replace only the treated candidate’s original text. All other candidates, the candidate order, and the slate size are

held fixed:

$$L_i = (d_{i1}, \dots, t_i, \dots, d_{im_i}),$$

$$L'_i = (d_{i1}, \dots, \tilde{t}_i, \dots, d_{im_i}).$$

where $\tilde{t}_i$ denotes the rewritten version of the treated candidate. The outcome is whether the chooser selects the treated target from the fixed slate:

$$Y_i = \mathbf{1}\{\text{Chooser}(L'_i) = t_i\}.$$

The original no-rewrite condition is evaluated analogously using L_i instead of L'_i .

Held-out overlap filtering. For leakage control, we check overlap using both document identifiers and normalized title–abstract text. Identifier checks remove exact document reuse, while text checks additionally catch duplicated or re-indexed articles whose IDs differ but whose displayed content is the same. In the strict paper-test split, we apply this check at the full-slate level rather than only to the treated target: if any candidate in a test impression overlaps with any training or validation candidate, the entire test impression is skipped.

Let

$$\kappa(d) = \text{normalize}(\text{title}(d), \text{abstract}(d))$$

be the normalized text key for document d . Let

$$\mathcal{T}_{\text{train/val}} = \{\kappa(d) : d \in L_i, i \in \mathcal{I}_{\text{train/val}}\}$$

be the set of candidate text keys observed anywhere in the training and validation slates. A test impression j is kept only if

$$\{\kappa(d) : d \in L_j\} \cap \mathcal{T}_{\text{train/val}} = \emptyset.$$

Thus, if any candidate in a test impression overlaps textually with any training or validation candidate, the entire test impression is skipped. The checked MIND paper-test split has zero overlap with the MIND train/validation splits for target IDs, target text, candidate IDs, and candidate text.

B.2 Translated MIND

We construct translated MIND variants as evaluation-only versions of the same 1,000 held-out MIND slates. The translated variants preserve the same rows, treated targets, candidate IDs, candidate order, candidate sizes, and reward/evaluation metadata; only the displayed title and abstract text are replaced by translations. These variants test whether selection gains transfer across display languages while holding the candidate-slate structure fixed. Details of the translation procedure and preserved fields are provided in Section I.1.

Split	Rows	Mean size	Max size	Random
Train	8,000	9.09	15	15.6%
Validation	500	9.36	15	15.0%
Test	1,000	8.42	15	16.9%

Table B.2: **MIND split statistics.** Rows denote impression-level candidate lists. Candidate-list sizes are computed from the fixed chooser candidate lists. Random reports the row-wise uniform target-selection baseline defined above.

B.3 EB-NeRD

We construct an EB-NeRD evaluation set from Danish news impressions. We sample 1,000 test impressions from records with at most 26 candidates and preserve full candidate lists rather than truncating them. Treated targets are sampled from each candidate slate with a cap of two treated occurrences per document. The resulting Danish EB-NeRD set has 886 unique treated documents, 114 repeated treated documents, and a maximum treated-target repeat count of two.

We also use an English machine-translated EB-NeRD variant. The Danish and English EB-NeRD versions preserve the same rows, impression IDs, treated article IDs, candidate slates, candidate order, and row alignment; only title and abstract text are translated.

B.4 Category Composition of News Evaluation Impressions

Figure B.1 reports the category composition of the final news evaluation impressions for MIND and EB-NeRD. For each dataset, we compare the distribution over all candidate articles appearing in the fixed chooser impressions with the distribution over treated targets, i.e., the documents whose title and abstract are rewritten. Both datasets retain the long-tailed category structure typical of news recommendation data. Importantly, the treated-target distributions closely track the corresponding candidate distributions, suggesting that the intervention targets are not concentrated in a substantially different category subset from the overall candidate pool. This provides a dataset-construction sanity check that the reported selection shifts are not driven by a heavily skewed target-category mix.

Category Composition of Final News Evaluation Impressions




Figure B.1: **Category composition of final news evaluation impressions.** Candidate panels summarize all articles appearing in the fixed chooser impressions, while treated-target panels summarize the documents whose title and abstract are rewritten. For both MIND and EB-NeRD, treated-target category distributions closely track the corresponding candidate distributions, suggesting that intervention targets are not concentrated in a substantially different category subset from the overall candidate pool. EB-NeRD labels show original Danish categories with English annotations.

B.5 SciRepEval-Derived Scientific-Document Candidate Lists

The scientific-document setting is derived from SciRepEval search-style candidate lists. Each example contains a fixed list of scientific-document candidates represented by titles and abstracts. The target document is selected from the candidate list and rewritten under the same fixed-slate intervention as the news settings.

C Prompts and Output Formats

Table C.1 reports the two prompts used for advisor-guided rewriting, and Table C.2 reports the single prompt used for the rewriter-only.

Table C.3 reports the target-agent prompts used to obtain target-selection decisions.

Table C.4 reports the shared rewrite objective and candidate rendering format.

Generation prompts require valid JSON outputs.

Target-agent prompts return only the selected candidate label. Empty title or abstract fields are rendered as None provided.

ADVISOR PROMPT

System message:
You are a rewriting strategist.
Your job is to propose one concrete rewrite strategy.
Do not write the final rewrite.
Return valid JSON only.

User message:
Objective: {rewrite_objective}

Current candidate title: {title}

Current candidate abstract: {abstract}

Return exactly one JSON object with this schema:
{"strategy": "<one concrete open-ended instruction>"}

ADVISOR-GUIDED REWRITER PROMPT

System message:
You are an execution model for rewriting.
Your job is to execute the provided advisor strategy as
↪ directly as possible
and rewrite the title and abstract.
Return valid JSON only with rewritten_title and
↪ rewritten_abstract.

User message:
Objective: {rewrite_objective}

Current candidate title: {title}

Current candidate abstract: {abstract}

Advisor strategy: {strategy}

Return exactly one JSON object with this schema:
{"rewritten_title": "...", "rewritten_abstract": "..."}"

Table C.1: **Advisor-guided generation prompts.** Each prompt box shows the concatenated system and user messages for a single model call. Advisor-guided rewriting uses two model calls: the advisor first proposes a rewrite strategy, and the rewriter then executes that strategy. All generation prompts require valid JSON outputs. Candidate identifiers, candidate order, and non-target candidate texts are fixed across original and rewritten conditions.

REWRITER-ONLY PROMPT

System message:
You are a strategic text rewriter.
Your job is to rewrite the title and abstract to
↪ maximize selection likelihood.
Return valid JSON only with rewritten_title and
↪ rewritten_abstract.

User message:
Objective: {rewrite_objective}

Current candidate title: {title}

Current candidate abstract: {abstract}

Return exactly one JSON object with this schema:
{"rewritten_title": "...", "rewritten_abstract": "..."}"

Table C.2: **Rewriter-only generation prompt.** The rewriter-only baseline uses a single model call that directly rewrites the title and abstract without an advisor-provided strategy. The prompt requires a valid JSON output. Candidate identifiers, candidate order, and non-target candidate texts are fixed across original and rewritten conditions.

NEWS TARGET-AGENT PROMPT

System message:
In this step, you are viewing a content feed.
Each content card is clickable and has a unique label.
You will be shown labeled items from the same impression.
Each has: label, title, and abstract.

User message:
Items: {candidates}

Choose exactly one item to click next, based only on the
↪ title and abstract.
Return the selected label only.

ACADEMIC TARGET-AGENT PROMPT

System message:
In this step, you are viewing a set of candidate papers.
Each paper is clickable and has a unique label.
You will be shown labeled items from the same search
↪ session.
Each has: label, title, and abstract.

User message:
Items: {candidates}

Choose exactly one item to click next, based only on the
↪ title and abstract.
Return the selected label only.

Table C.3: **Target-agent prompts.** Each prompt box shows the concatenated system and user messages for a single target-agent call. The target agent receives the fixed candidate impression and returns exactly one selected label. The news and academic settings differ only in the framing of the candidate set.

SHARED REWRITE OBJECTIVE

Rewrite the title and abstract to make this item more
 $\hookrightarrow$ likely to be selected by
the target agent based on its title and abstract.

You may revise the title and abstract by changing their
 $\hookrightarrow$ wording, emphasis,
ordering, framing, clarity, structure, and tone when
 $\hookrightarrow$ helpful.

MULTILINGUAL LANGUAGE CONSTRAINT

Original text is in {language}.

The rewritten_title and rewritten_abstract must be
 $\hookrightarrow$ entirely in {language}.

CANDIDATE RENDERING

A.
Title: {title}
Abstract: {abstract}

B.
Title: {title}
Abstract: {abstract}

...

Empty title or abstract fields are rendered as None
 $\hookrightarrow$ provided.

Table C.4: **Shared rewrite objective, multilingual language constraint, and candidate rendering.** The same rewrite objective is used for advisor-guided and standalone rewriting. In multilingual runs, a language constraint is appended to the rewriter system message. The same candidate rendering format is used in original and rewritten target-agent conditions.

D Reproducibility Details and Optimization Rationale

This appendix reports the implementation details needed to reproduce the optimization and evaluation pipeline. We focus on the evaluation protocol, policy-optimization hyperparameters, GEPA prompt-optimization budget, reward-evaluation accounting, validation and model-selection procedure, and hardware/software environment.

D.1 Evaluation Protocol

Our evaluation is conditional on exposure: for each held-out impression, the candidate list is fixed, and only the treated target item’s title and abstract are changed. Original and rewritten conditions use the same candidate order, same target-agent prompt, and same target-selection metric.

D.2 Policy Optimization Hyperparameters

Training uses GRPO through SkyRL with a Qwen3.5-9B policy. In the standalone setting, the

policy directly generates the rewritten title and abstract. In the advisor-guided setting, the policy generates an advisor strategy that is passed to GPT-5-mini for the final rewrite. The reward is based on final target selection by the target agent. Source-support variants add a factuality-oriented support constraint.

D.3 Optimization Budgets and Reward Accounting

The primary unconstrained run uses a nominal generation budget of $250 \times 32 \times 4 = 32,000$ sampled outputs. The factuality-constrained run uses $500 \times 32 \times 4 = 64,000$ sampled outputs. One downstream target-agent evaluation is requested for each sampled output.

The GEPA comparison uses at most 32,000 sample-level metric calls, matching the nominal

Category	Item	Setting
Hardware	GPUs	4 NVIDIA A100-SXM4-80GB GPUs
Training	Framework	SkyRL
	Runtime	PyTorch, vLLM, Ray
Software	Python	3.12.3
	PyTorch	2.10.0
	Transformers	5.6.0.dev0
	vLLM	0.19.1
	SkyRL	0.3.0
	SkyRL-Train	0.3.1
	FlashAttention	2.8.3
	Ray	2.51.1

Table D.1: **Hardware and software environment.** Environment used for advisor training and evaluation.

Setting	Value
Policy initialization	Qwen3.5-9B
Optimization objective	GRPO
Optimization steps	250 (unconstrained); 500 (factuality-constrained)
Examples per step	32
Policy mini-batch size	8
Sampled rollouts per example	4
Learning rate	3×10^{-7}
Sampling temperature	1.0
KL coefficient	0.0075
Warm-up steps	10
Source-support weight	$\lambda = 1$ unless otherwise stated

Table D.2: **GRPO training configuration.** Settings shared by the primary advisor and rewriter-only runs; rows explicitly distinguish the unconstrained and factuality-constrained variants where their budgets differ.

32,000-output budget of the primary unconstrained GRPO run. Reflection mini-batches contain 32 examples. Because GRPO updates model parameters whereas GEPA searches over a global instruction, metric calls and optimizer steps are reported separately rather than treated as interchangeable units. Further GEPA details are provided in [subsection E.3](#).

D.4 Validation and Model Selection

Validation examples are used for checkpoint or prompt-variant selection. The held-out test split is reserved for final reporting. Unless otherwise stated, cross-evaluator results are used to measure transfer rather than to select the optimized system.

D.5 Hardware and Software Environment

Training was run with SkyRL, PyTorch, and vLLM on four NVIDIA A100-SXM4-80GB GPUs. The primary 9B unconstrained advisor run used 250 optimization steps.

D.6 Practical Considerations

Because optimized prompts and policy checkpoints can vary in target-selection performance, validation examples are used for selection before final test evaluation. The test split is not used for iterative system selection. GRPO and GEPA budgets are reported in their native units, as detailed in [subsection D.3](#).

E Additional Baselines and Robustness Analyses

E.1 Robustness to Target Agent Prompts and Inference Settings

We re-evaluate the same frozen rewriters used in the main experiments on 1,000 held-out English MIND impressions under one same-objective paraphrase and three criterion stress tests. As shown in [Table E.1](#), every rewriting method remains above its matched original-text baseline, although the size of the gain varies across chooser criteria.

Chooser-prompt variants. Every user prompt begins with Items: {candidates} and requires the chooser to return only the selected label. The decision instructions are:

Default. Choose exactly one item to click next, based only on the title and abstract.

Paraphrased. Using only the title and abstract of each item, select exactly one item to click next.

Reader usefulness. Choose the item that seems most informative and substantively useful to a general reader. Base the decision only on the title and abstract.

Evidence-focused. Choose the item that appears most specific, concrete, and evidence-oriented. Do not reward vague, sensational, or generic wording. Base the decision only on the title and abstract.

Style-deemphasized. Choose exactly one item to click next, based only on the title and abstract. Do not choose an item only because its wording sounds more technical, academic, or polished.

Inference settings. We additionally vary GPT-5-mini reasoning effort across minimal, medium, and high settings while holding the inference configuration fixed across methods within each comparison. All rewriting methods remain above their matched baselines, and within-method target-selection rates vary by at most 5.1 percentage points.

E.2 Meaning of the Original-Target Baseline and Random-Choice Diagnostics

The no-rewrite condition reports the target agent’s selection rate for the designated target item before any rewriting. This condition is not a retrieval baseline. The 1,000 held-out English MIND impressions contain 8.42 candidate items on average. As shown in [Table E.3](#), the original-target selection rate is 17.10%, close to the row-wise uniform baseline of 16.92%. This aggregate similarity does not mean that GPT-5-mini behaves like a uniform chooser: when each impression is evaluated under 10 candidate-order permutations, its pairwise selection agreement is 51.49%, compared with 16.92% under uniformly random choice. Thus, the designated targets are not systematically advantaged on average, while the target agent still exhibits stable preferences for some items.

E.3 Fixed-Strategy Baselines

We test five hand-written global strategy prompts and a GEPA-style optimized strategy to determine whether advisor-training gains can be explained by a single global instruction rather than by learned input-dependent strategy generation. GEPA is a reflective prompt optimizer that revises instructions using downstream task feedback while keeping model parameters fixed ([Agrawal et al., 2026](#)). At a high level, it maintains a set of candidate prompts $\mathcal{P}$, evaluates each prompt $p \in \mathcal{P}$ with a task score

Method	Default	Paraphrased	Reader usefulness	Evidence-focused	Style-deemphasized
Original	17.1	16.4	17.8	19.9	17.7
<i>Prompting-only</i>					
Rewriter-only	34.8 (+17.7)	34.0 (+17.6)	28.4 (+10.6)	27.2 (+7.3)	31.4 (+13.7)
Advisor-rewriter	43.9 (+26.8)	45.1 (+28.7)	40.4 (+22.6)	39.4 (+19.5)	43.0 (+25.3)
<i>RL-trained rewriter</i>					
Rewriter-only	95.9 (+78.8)	84.6 (+68.2)	80.8 (+63.0)	97.7 (+77.8)	83.6 (+65.9)
+ MiniCheck	43.4 (+26.3)	35.8 (+19.4)	35.3 (+17.5)	43.4 (+23.5)	37.9 (+20.2)
<i>RL-trained advisor</i>					
Advisor-rewriter	98.5 (+81.4)	91.9 (+75.5)	82.9 (+65.1)	96.6 (+76.7)	91.0 (+73.3)
+ MiniCheck	68.6 (+51.5)	66.5 (+50.1)	58.9 (+41.1)	57.5 (+37.6)	65.1 (+47.4)

Table E.1: **Robustness to chooser-prompt variants.** Target-selection rates (%) on 1,000 held-out English MIND impressions using the same frozen rewriters as in Table 2. Parentheses report absolute gains over the original-text baseline under the same chooser prompt. “Original” denotes the same unmodified target condition labeled “Original target” in the main tables. Darker cells indicate higher selection rates.

Method	Minimal	Medium	High
<i>Reference</i>			
Original target	17.1	18.0	19.1
<i>Prompting-only</i>			
Rewriter-only	34.8 (+17.7)	32.8 (+14.8)	31.7 (+12.6)
Advisor-rewriter	43.9 (+26.8)	41.0 (+23.0)	40.2 (+21.1)
<i>RL-trained rewriter</i>			
Rewriter-only	95.9 (+78.8)	96.7 (+78.7)	95.5 (+76.4)
+ MiniCheck	43.4 (+26.3)	45.1 (+27.1)	40.0 (+20.9)
<i>RL-trained advisor</i>			
Advisor-rewriter	98.5 (+81.4)	99.3 (+81.3)	98.5 (+79.4)
+ MiniCheck	68.6 (+51.5)	67.6 (+49.6)	65.6 (+46.5)

Table E.2: **Sensitivity to target-agent reasoning effort.** Target-selection rates (%) under minimal, medium, and high GPT-5-mini reasoning-effort settings. Parenthesized values report absolute gains in percentage points over the original target under the same reasoning-effort setting.

Diagnostic	Observed	Uniform baseline
Original-target selection rate	17.10%	16.92%
Pairwise agreement (10 permutations)	51.49%	16.92%

Table E.3: **Original-target and random-choice diagnostics.** Results are computed on 1,000 held-out English MIND impressions. Pairwise agreement is measured across 10 candidate-order permutations of each impression. Because candidate-list sizes vary, the uniform baseline is the row-wise average of $1/|L_i|$.

Method	GPT-5-mini	Gemini 3 Flash	Sonnet 4.6
<i>Reference baselines</i>			
Original	17.1	17.1	17.6
Untrained advisor	43.9 (+26.8)	48.8 (+31.7)	34.2 (+16.6)
<i>Fixed global strategies</i>			
Academic/empirical	72.4 (+55.3)	71.1 (+54.0)	30.6 (+13.0)
Technical	66.5 (+49.4)	67.1 (+50.0)	29.1 (+11.5)
Authority/evidence	67.2 (+50.1)	61.7 (+44.6)	32.8 (+15.2)
Urgency/stakes	52.1 (+35.0)	50.9 (+33.8)	30.0 (+12.4)
Grounded detailing	48.5 (+31.4)	50.5 (+33.4)	36.2 (+18.6)
GEPA-optimized	77.5 (+60.4)	51.7 (+34.6)	38.4 (+20.8)
<i>Learned advisor policy</i>			
RL advisor	98.5 (+81.4)	98.9 (+81.8)	65.0 (+47.4)

Table E.4: **Fixed global strategy baselines.** Target-selection rates (%) on held-out English MIND impressions. Fixed-strategy rows reuse one global instruction across target items. Values in parentheses report absolute gains over the Original baseline. The RL-trained advisor outperforms every tested fixed global strategy across all three target agents.

$S(p)$, and searches for a higher-scoring prompt:

$$p^* = \arg \max_{p \in \mathcal{P}} S(p).$$

Here, $S(p)$ is the downstream target-selection score obtained when using prompt p in the rewriting pipeline.

In our setting, GEPA operates at the strategy layer and can be viewed as a fixed-strategy variant of the advisor-rewriter framework. Instead of learning an input-dependent advisor policy, GEPA searches for one global strategy-producing instruction that is reused for all target items while all model weights remain fixed. This provides both a fixed-strategy baseline and an alternative strategy-

optimization route within the same advisor-rewriter setting. If a single optimized strategy instruction were sufficient to account for the gains from advisor training, this baseline would be expected to approach the RL-trained advisor. In contrast, the RL-trained advisor learns an input-dependent strategy policy with GRPO and can generate a different strategy for each target item. This comparison therefore separates global prompt-level strategy search from learned per-example strategy adaptation.

We run GEPA with a maximum budget of 32,000 sample-level metric calls, approximately matching the nominal rollout budget of our GRPO advisor training setup: 8,000 training examples with 4 sampled advisor outputs per example. Reflection minibatches contain 32 examples. The selected GEPA strategy is evaluated under the same fixed-list English MIND protocol used in Table 2.

Hand-written fixed-strategy prompts. The five global instructions are reused unchanged for every target item:

Academic/empirical reframing. Rewrite the item as an empirical study or analytical report, emphasizing the research question, evidence, measured outcomes, and broader implications.

Technical framing. Reframe the item using a technical or analytical angle that makes the topic appear systematic, mechanism-driven, and consequential.

Authority/evidence framing. Rewrite the item to sound credible, investigative, and evidence-oriented, emphasizing concrete evidence, institutional context, expert relevance, legal or policy implications, or data-backed framing.

Urgency/stakes framing. Foreground why the item matters now, emphasizing risk, rarity, consequences, human cost, or time sensitivity.

Grounded detailing. Expand concrete details to clarify the actors, event, context, and consequences.

The results in Table E.4 suggest that the advisor gains cannot be reduced to a single global strategy instruction. Several hand-written strategies and GEPA produce substantial selection shifts, but all remain below the RL-trained advisor. GEPA is a meaningful strategy-layer variant of our advisor-rewriter framework rather than a trivial baseline: it directly searches for a fixed global strategy instruction while keeping all model weights fixed. In contrast, the RL-trained advisor learns a per-example strategy policy and can adapt its guidance

SELECTED GLOBAL STRATEGY INSTRUCTION

Adopt a **“Priority-First Verification”** strategy. **Titles** must lead with the primary event, action, time, or metric, such as a number, date, outcome, or other news hook, rather than the reporting organization or actor. Append publication or agency attribution at the beginning or end only when it clarifies the authority of the claim or is itself the subject, such as a press release, to maximize signal clarity and avoid redundancy.

Abstracts must open with a **Stakes** line that reframes the outcome as a direct reader consequence or a specific named-group impact, ensuring the impact is concrete and immediate while removing passive voice or hypothetical phrasing.

Facts lines must strictly list only discrete, verifiable data points, such as names, figures, dates, and counts, that substantiate the title’s claims; omit background context or inferred motives.

Action sections must provide exactly three steps, each instructing the user to locate and read the specific sentence or phrase, including quotes where necessary, in the source that verbatim validates a distinct element of the title, ensuring full textual grounding.

Table E.5: **Selected global strategy instruction from GEPA.** We show the final global strategy instruction selected by the reported GEPA reflection run. The instruction is reused for all target items rather than generated per example. The selected prompt is meaningful and interpretable, but it also illustrates the fixed-template nature of global strategy optimization.

to each target item. This supports the view that advisor training improves target selection through adaptive strategy generation, not simply through a reusable global rewriting prompt.

E.4 Cross Target Agent Transfer

It is reasonable to ask whether a particular target agent used during training performs better for a particular target model at evaluation. In Table E.6, we show the selection-rate improvements over baseline for GPT-5-mini and Gemini 3 Flash at different stages of training, while keeping the data, rewriter, policy model, and training steps fixed. We can see that two models show strong mutual transfer.

Training target	GPT-5-mini	Gemini 3 Flash
Original target	17.1	17.1
Prompting-only	43.9 (+26.8)	48.8 (+31.7)
<i>Advisor-rewriter target agent</i>		
GPT-5-mini	98.5 (+81.4)	98.9 (+81.8)
Gemini 3 Flash	84.3 (+67.2)	90.6 (+73.5)

Table E.6: **Matched fixed-step cross-target-agent selection rates.** For the trained advisor-rewriter rows, row labels indicate the target agent used for scoring during training; columns indicate the target agent used during evaluation. Values are target selection rates in percentages; values in parentheses report absolute gains over the corresponding no-rewrite selection rate in percentage points. The row-wise uniform random baseline is 16.9%. Darker cells indicate higher selection rates.

E.5 Training Advisor Variants

We evaluate whether advisor-training gains depend on the choice of trainable advisor backbone. All variants keep the frozen GPT-5-mini rewriter, target agent, reward, training data, and fixed-list English MIND evaluation protocol unchanged. We vary the advisor along two axes: model family and model scale.

E.6 Training Rewriter Variants

To test whether learned advisor strategies are independent of the rewriting model, we decouple the rewriter used during advisor training from the rewriter used at evaluation time, while holding fixed the target agent, advisor backbone, training data, reward, optimization budget, and evaluation split.

The strongest results occur when the training-

Advisor backbone	Untrained advisor	Trained advisor	Gain over original
Qwen3.5-9B (default)	43.9	98.5	+81.4
Qwen3.5-4B	32.5	94.5	+77.4
SmolLM3	27.1	73.7	+56.6
Llama-3.1-8B-Instruct	30.4	68.4	+51.3

Table E.7: **Training advisor variants.** We vary the trainable advisor backbone while keeping the frozen GPT-5-mini rewriter, target agent, reward, training data, and fixed-list English MIND evaluation protocol unchanged. Qwen3.5-9B is the default advisor backbone used in the main experiments; the remaining rows are backbone ablations. Each cell reports target-selection rate in percentages. The original no-rewrite target-selection rate is 17.1%; gains are absolute percentage-point improvements over this original baseline.

	Evaluation-time rewriter	
	GPT-5-mini	Gemini 3.1 Flash-Lite
Advisor training		
<i>Untrained-advisor baseline</i>		
Untrained advisor	43.9	29.7
<i>Trained advisor</i>		
GPT-5-mini rewriter	98.5	28.7
Gemini 3.1 Flash Lite rewriter	56.1	78.2

Table E.8: **Train-time versus evaluation-time rewriter ablation.** Rows indicate the rewriter used during advisor training to realize sampled strategies and compute target-selection rewards. Columns indicate the rewriter used at evaluation time to execute the advisor’s strategies on held-out English MIND impressions. All cells report target selection rates in percentages. The untrained-advisor baseline is evaluated separately for each evaluation-time rewriter. For the trained rows, the target agent, advisor backbone, training data, reward, optimization budget, and evaluation split are held fixed.

λ	Selection rate (%)	MiniCheck (%) $\uparrow$	AlignScore (%) $\uparrow$	FENICE (native) $\uparrow$
0.0	98.5	2.0	1.1	-0.784
0.5	66.9	23.8	15.8	-0.219
1.0	68.6	31.2	16.1	-0.382

Table E.9: **Sensitivity to the source-support reward weight.** Final-checkpoint results for the RL advisor-rewriter on held-out English MIND impressions. The $\lambda = 0$ setting optimizes only target selection; positive values add the MiniCheck source-support reward. Higher source-support scores are better.

time and evaluation-time rewriters match. Cross-rewriter transfer is asymmetric: the GPT-5-mini-trained advisor transfers poorly to Gemini 3.1 Flash Lite despite GPT-5-mini having the stronger untrained-advisor baseline, whereas the Gemini-trained advisor transfers partially to GPT-5-mini while performing best with the matched Gemini rewriter. This pattern suggests that advisor strategies are not fully model-agnostic; they partly depend on the execution style and capability profile of the rewriter used during training.

E.7 Support-Reward Sensitivity

We vary the source-support reward weight λ for the RL advisor-rewriter. Both positive-weight final checkpoints have lower target selection and higher scores on MiniCheck, AlignScore, and FENICE than the $\lambda = 0$ checkpoint. This shows that the support constraint changes the learned rewriting behavior rather than acting as a cost-free regularizer.

Method	MiniCheck (%) ↑	Bespoke-MiniCheck-7B (%) ↑	AlignScore (%) ↑	FENICE (native) ↑
Prompting-only rewriter	60.7	51.7	44.8	0.123
RL rewriter	2.2	0.3	0.9	−0.814
RL rewriter + MiniCheck	66.7	38.6	48.6	0.313
Prompting-only advisor-rewriter	42.2	35.7	28.7	−0.193
RL advisor-rewriter	2.0	0.9	1.1	−0.784
RL advisor-rewriter + MiniCheck	31.2	26.8	16.1	−0.382

Table E.10: **Verifier coupling under independent source-support metrics.** We train the constrained variants using MiniCheck and additionally evaluate all methods with Bespoke-MiniCheck-7B, AlignScore, and FENICE. Bespoke-MiniCheck-7B scores are aggregated over rewritten title and abstract claims on the English MIND evaluation split ($N = 1,000$ per setting). Higher values indicate greater source support. MiniCheck-constrained methods score higher than their corresponding unconstrained RL variants under all four metrics.

F Verifier Robustness and Factuality Diagnostics

We use automatic source-support scores as diagnostics rather than definitive factuality labels. In addition to MiniCheck-Flan-T5-Large (Tang et al., 2024b), we evaluate the same outputs using Bespoke-MiniCheck-7B (Bespoke Labs, 2024), AlignScore (Zha et al., 2023), and FENICE (Scirè et al., 2024).

Table E.10 shows that both MiniCheck-constrained methods score higher than their corresponding unconstrained RL variants under all four metrics. Bespoke-MiniCheck-7B yields the same qualitative comparison on English MIND, although absolute score levels differ across verifiers.

These differences suggest that the metrics have distinct calibration and should not be interpreted as directly comparable measures of factuality across settings. We therefore use them as complementary diagnostics of source support. Selection remains the primary outcome in the main experiments, while the factuality metrics provide an auxiliary check on the factual-consistency risk associated with the generated rewrites.

G Strategy Taxonomy and Advisor Behavior

G.1 Strategy analysis with StringSight.

To characterize the rewriting behavior learned by the advisor, we conduct a post-hoc qualitative analysis using StringSight (Dunlap, 2025). We import per-example advisor rewrite outputs and associated strategy descriptions into StringSight, which provides an initial clustering of recurring rewriting behaviors. We treat these clusters as an exploratory taxonomy rather than as automatic ground-truth

labels. We then manually merge and normalize the raw clusters into paper-facing strategy families, focusing on behavioral rewriting moves such as article-specific reframing, urgency or engagement framing, actionability, grounding, credibility, and factuality risk. All quantitative task results are computed independently of StringSight; StringSight is used only to support interpretation of the advisor’s learned strategies.

G.2 Advisor Strategy Taxonomy

Table G.1 summarizes the qualitative strategy taxonomy used to analyze advisor outputs. We code the advisor instruction rather than the final rewritten title or abstract, since the instruction is the object directly optimized by the advisor policy. The taxonomy is derived from rule-normalized StringSight clusters and merged into paper-facing families that capture the main behavioral rewriting move, such as topic substitution, grounded expansion, stakes framing, evidential framing, or genre reframing. These categories describe mechanisms by which advisor instructions may increase target selection rate; they should not be interpreted as judgments that the resulting rewrites are more accurate, useful, or socially desirable.

All counts are computed on the English MIND evaluation set with one advisor instruction per example. The no-support-constraint advisor is dominated by unsupported technical substitution, while the support-constrained advisor distributes more broadly across grounded expansion, urgency and stakes framing, mixed investigative reframing, and academic or empirical genre reframing. Evaluator-specific advisors also show distinct policies: the Gemini 3 Flash evaluator advisor mostly academicizes items into empirical-study formats.

Strategy family	Definition
Unsupported technical substitution	Replaces the original topic with a technical, scientific, or AI-centered framing, often introducing unsupported mechanisms or claims.
Academic or empirical genre reframing	Recasts the item as an academic study, policy analysis, or methodological contribution using metrics, methods, data, or results.
Grounded article expansion and detailing	Preserves the original topic while adding source-relevant details, concrete actors, causes, examples, or consequences.
Urgency and stakes framing	Increases salience by foregrounding risk, rarity, human cost, consequences, or time sensitivity while generally preserving the topic.
Authority and evidential framing	Adds investigative, legal, forensic, institutional, expert, or evidence-oriented framing to increase perceived credibility.
Human and emotional impact framing	Foregrounds vulnerability, community impact, grief, resilience, identity, or lived experience to make the event more personally consequential.
Positive benefit and agency reframing	Emphasizes resolution, safety, usefulness, agency, or practical benefit, often reducing alarm while preserving selection appeal.
Causal and structural clarity	Highlights mechanisms, causes, systems, or structural implications, emphasizing explanation rather than evidential authority.
Other or mixed strategy	Combines multiple rewriting moves or does not cleanly fit a single dominant strategy family.

Table G.1: **Qualitative taxonomy of advisor strategies.** Each row defines one paper-facing strategy family used in the post-hoc analysis.

Strategy family	Base	No support	Support-constrained	Gemini 3 Flash
Unsupported technical substitution	6 (0.6%)	962 (96.2%)	1 (0.1%)	31 (3.1%)
Academic or empirical genre reframing	28 (2.8%)	13 (1.3%)	67 (6.7%)	951 (95.1%)
Grounded article expansion and detailing	553 (55.3%)	1 (0.1%)	456 (45.6%)	4 (0.4%)
Urgency and stakes framing	300 (30.0%)	13 (1.3%)	283 (28.3%)	6 (0.6%)
Authority and evidential framing	44 (4.4%)	1 (0.1%)	41 (4.1%)	1 (0.1%)
Human and emotional impact framing	18 (1.8%)	1 (0.1%)	27 (2.7%)	0 (0.0%)
Positive benefit and agency reframing	21 (2.1%)	5 (0.5%)	8 (0.8%)	2 (0.2%)
Causal and structural clarity	5 (0.5%)	2 (0.2%)	9 (0.9%)	3 (0.3%)
Other or mixed strategy	25 (2.5%)	2 (0.2%)	108 (10.8%)	2 (0.2%)
Total	1000 (100.0%)	1000 (100.0%)	1000 (100.0%)	1000 (100.0%)

Table G.2: **Post-hoc distribution of advisor strategy families.** Each column reports the number and percentage of advisor instructions assigned to each strategy family on the English MIND evaluation set ($N = 1,000$ per column). Strategy families are assigned automatically from the generated advisor instruction. The distributions show how factual-support constraints and the transfer target agent change the dominant rewriting strategy.

H Additional Qualitative Examples

Case	Original target title	Rewritten target title	Rewriter variant	Selection	Mini-Check
Grounded selection gain	Cold, wet weather rolls into western Washington, could lower snow levels to 1,000 feet by Tuesday	Holiday Travel Alert: Snow Levels Could Drop to 1,000 Feet; Expect Freezing, Hazardous Mountain Driving This Week	Untrained advisor	0 → 1	0.868
Unsupported selection gain	Brave sea turtle shows up to visit scuba divers at shark dive	DeepSeaDetect: A Machine Learning Pipeline for Reliable Hawksbill Turtle Detection in Shark-Associated Midwater Environments	Trained advisor	0 → 1	0.014
Support-aware selection gain	Alaskan city sees heat and snowfall records in single day	A Rare Weather Paradox: Anchorage Records a 45°F High and Over a Foot of Snow in One Day	Trained advisor + MiniCheck	0 → 1	0.643
Faithful non-selection	Duke climbs to No. 1 in AP Top 25 following Kentucky's loss	Duke Surges to No. 1 in AP Top 25 Poll Following Kentucky's Loss	Trained standalone + MiniCheck	0 → 0	0.960

Table H.1: **Qualitative examples under fixed candidate lists.** Each row compares the original target with its rewritten version while holding the candidate identities, candidate order, non-target text, and chooser conditions fixed. Selection reports the target outcome before and after rewriting. MiniCheck reports the source-support score of the rewritten target. The examples show that selection gains can arise from both source-supported and unsupported reframing, while high source support alone does not guarantee selection.

I Translation Protocol and Multilingual Examples

I.1 Translation Protocol

For MIND, we export unique document title–abstract pairs from the 1,000 held-out slates and translate them with Google Cloud Translation API (Advanced v3) using the default neural machine translation model. The target languages are Arabic, Danish, Spanish, Swahili, Turkish, and Simplified Chinese. For EB-NeRD, we translate Danish title–abstract pairs into English.

Across translation-based evaluations, only the displayed candidate title and abstract are translated. Row order, impression IDs, treated candidate IDs, candidate counts, candidate order, and nested candidate-ID sequences are preserved. Candidate identifiers, labels, and prompt structure are not translated. Empty title or abstract fields are rendered as None provided. No additional training, fine-tuning, or translation-specific rewriting is performed.

This protocol tests language transfer under structurally aligned candidate slates. It does not claim to reproduce native language-specific presentation conventions, such as natural headline style, emphasis, or framing.

I.2 Representative Translation Examples

Tables I.1 and I.2 show two representative MIND target items after machine translation. The first is a longer informational item with numeric content, and the second is a shorter food and culture item. Across all examples, the candidate list, target item, and prompt structure are preserved; only the displayed title and abstract fields are translated.

Language	Target item displayed to the chooser
English source	Title: Why Tokyo's Haneda is one of the world's most punctual airports Abstract: Haneda, officially called Tokyo International Airport, is the world's fifth busiest airport. Yet it delivered 85.6% of its flights on time in 2018, making it the most punctual mega airport in the world.
Arabic	Title: لماذا يُعد مطار هانيدا في طوكيو أحد أكثر المطارات التزاماً بالمواعيد في العالم؟ Abstract: يُعد مطار هانيدا، المعروف رسميًا باسم مطار طوكيو الدولي، خامس أكثر مطارات العالم ازدحامًا. ومع ذلك، فقد أُنجز 85,6 بالمئة من رحلاته في الموعد المحدد عام 2018، مما يجعله أكثر المطارات الضخمة التزامًا بالمواعيد في العالم.
Danish	Title: Hvorfor Tokyos Haneda er en af verdens mest punktlige lufthavne Abstract: Haneda, officielt kaldet Tokyo International Airport, er verdens femte travleste lufthavn. Alligevel afviklede den 85,6% af sine flyvninger til tiden i 2018, hvilket gør den til den mest punktlige megalufthavn i verden.
Spanish	Title: ¿Por qué el aeropuerto de Haneda, en Tokio, es uno de los aeropuertos más puntuales del mundo? Abstract: Haneda, cuyo nombre oficial es Aeropuerto Internacional de Tokio, es el quinto aeropuerto con mayor tráfico del mundo. Sin embargo, en 2018 logró que el 85,6% de sus vuelos llegaran a tiempo, lo que lo convierte en el megaaeropuerto más puntual del mundo.
Swahili	Title: Kwa nini Haneda ya Tokyo ni mojawapo ya viwanja vya ndege vinavyofanya kazi kwa wakati zaidi duniani Abstract: Haneda, ambayo inaitwa rasmi Uwanja wa Ndege wa Kimataifa wa Tokyo, ni uwanja wa ndege wa tano wenye shughuli nyingi zaidi duniani. Hata hivyo, ilitoa 85.6% ya safari zake za ndege kwa wakati mwaka wa 2018, na kuifanya kuwa uwanja mkubwa wa ndege unaofanya kazi kwa wakati zaidi duniani.
Turkish	Title: Tokyo'nun Haneda Havalimanı neden dünyanın en dakik havalimanlarından biri? Abstract: Resmi adı Tokyo Uluslararası Havalimanı olan Haneda, dünyanın en yoğun beşinci havalimanıdır. Buna rağmen, 2018 yılında uçuşlarının %85,6'sını zamanında gerçekleştirerek dünyanın en dakik mega havalimanı olmuştur.
Chinese (zh-CN)	Title: 为什么东京羽田机场是世界上最准时的机场之一？ Abstract: 羽田机场,正式名称为东京国际机场,是世界第五繁忙的机场。然而,2018年其航班准点率高达85.6%,使其成为世界上最准点的巨型机场。

Table I.1: **Representative MIND multilingual translation example 1: airport punctuality.** Target candidate ID: N84539; fixed candidate list size: 9.

Language	Target item displayed to the chooser
English source	Title: 5 Things You Didn't Know About Campbell's Iconic Green Bean Casserole Abstract: We even got our hands on the original 1955 recipe card from the test kitchen.
Arabic	Title: خمسة أشياء لم تكن تعرفها عن طبق الفاصوليا الخضراء الشهير من كامبل Abstract: بل إننا حصلنا على بطاقة الوصفة الأصلية لعام 1955 من مطبخ الاختبار.
Danish	Title: 5 ting du ikke vidste om Campbells ikoniske grønne bønnegryde Abstract: Vi fik endda fat i det originale opskriftskort fra 1955 fra testkøkkenet.
Spanish	Title: 5 cosas que no sabías sobre la icónica cazuela de judías verdes de Campbell's Abstract: Incluso conseguimos la tarjeta de recetas original de 1955 de la cocina de pruebas.
Swahili	Title: Mambo 5 Uliyokuwa Huyajui Kuhusu Kitoweo cha Maharagwe Kijani cha Campbell Abstract: Tulipata hata kadi ya mapishi ya awali ya 1955 kutoka jikoni ya majaribio.
Turkish	Title: Campbell's'ın İkonik Yeşil Fasulye Güveci Hakkında Bilmediğiniz 5 Şey Abstract: Hatta test mutfağından 1955 yılına ait orijinal tarif kartını bile ele geçirdik.
Chinese (zh-CN)	Title: 关于坎贝尔经典青豆砂锅,你不知道的5件事 Abstract: 我们甚至还找到了1955年测试厨房的原始食谱卡。

Table I.2: **Representative MIND multilingual translation example 2: food and culture.** Target candidate ID: N66002; fixed candidate list size: 2.

J Data Sources, Licenses, and Release Policy

We use MINDlarge (Wu et al., 2020), EB-NeRD (Kruse et al., 2024), and SciRepEval-derived scientific-document candidate sets (Singh et al., 2023). We also construct translated evaluation-only variants of held-out MIND and EB-NeRD slates using Google Cloud Translation API.

Data sources and derived selection format. MINDlarge provides English news recommendation impressions and is used for advisor training and in-domain evaluation. EB-NeRD provides Danish news recommendation impressions from Ekstra Bladet and is used only for out-of-distribution news evaluation. For scientific-document evaluation, we use candidate-paper sets derived from the SciRepEval search setting. Across all settings, each example is converted into a fixed candidate-list selection instance: one target item may have its displayed title and abstract rewritten, while all other candidates remain unchanged. Translated MIND and EB-NeRD variants preserve the original candidate lists and target assignments, but replace the displayed candidate text with machine-translated text. These translated variants are used only for controlled evaluation.

Intended use. We use these datasets for controlled research on presentation-induced selection shifts under fixed exposure. Our experiments do not train or deploy a production recommendation system, do not model long-term platform feedback, and do not attempt to infer private user attributes. Although MIND and EB-NeRD contain interaction logs, we use them only to construct candidate-list selection contexts. The advisor is optimized using chooser-model feedback, not by predicting individual users' clicks.

Licenses, access conditions, and release policy. We use all datasets and derived evaluation sets under their respective license and access terms. MINDlarge is made available for research purposes under the Microsoft Research License Terms. For MIND-derived settings, we release code, preprocessing scripts, prompts, target assignments, model outputs, and aggregate evaluation results, but do not redistribute raw MIND impression logs or article metadata; users must obtain MIND from the original provider and run our preprocessing scripts.

EB-NeRD is made available under the dataset's General License Terms for research-only, non-

commercial use. Because the EB-NeRD terms restrict redistribution of exact dataset contents and other derivative uses, including training language models, we use EB-NeRD only for controlled research evaluation. We do not train or fine-tune language-model, advisor, or rewriter parameters on EB-NeRD text, logs, or article metadata, and we do not release EB-NeRD text, logs, article metadata, translated EB-NeRD slates, or reconstructed EB-NeRD candidate-list corpora.

For the scientific-document experiments, we use candidate sets derived from SciRepEval. The SciRepEval aggregate benchmark is released under the Open Data Commons Attribution License (ODC-BY), and we follow the applicable licensing requirements of its constituent datasets. Where permitted by the SciRepEval license and constituent dataset licenses, we release the processed scientific-document evaluation files used in our experiments, with attribution to the original sources.

Translations preserve the access restrictions of the underlying source datasets. Accordingly, we release translated or processed artifacts only where permitted by the source dataset terms.

Privacy and content considerations. We do not collect new user data or conduct new human-subject data collection. MIND and EB-NeRD are existing datasets with anonymized user identifiers. We do not attempt to identify users, recover private attributes, link anonymized identifiers to external accounts, or reconstruct user-level histories beyond the candidate-list context needed for our experiments. When reporting qualitative examples, we avoid releasing raw user logs and exact EB-NeRD dataset contents.

Potential misuse. Our method studies how rewriting displayed text can increase selection by an LLM chooser, which could be misused to make content more appealing to automated selection systems without improving its underlying quality. We therefore frame the work as a controlled measurement of selection vulnerability, report factuality variants, and discuss limitations and risks in the limitations section.