
GOLD PANNING: Iterative Bayesian Signal Anchoring for Many-Document Needle-in-Haystack Reasoning

Adam Byerly¹ Daniel Khashabi¹

Abstract

Large language models (LLMs) exhibit pronounced position bias in long-context needle-in-haystack problems, systematically prioritizing the location of information over its relevance. While current mitigations rely on *white-box* access, this is effectively impossible for many state-of-the-art models. We introduce GOLD PANNING, a *black-box* Bayesian framework that performs inference-time active search over long contexts by (i) reordering documents to concentrate high-belief items in highly diagnostic positions (signal anchoring) and (ii) updating beliefs over document relevance from model outputs. Unlike conventional active learning, which prioritizes uncertainty reduction, GOLD PANNING leverages anchoring—once flagged, keep it in sight—to preserve weak cues. We implement this using iterative assignment derived from the model’s diagnosticity profile, which provably identifies a target among N documents in $O(\log N)$ rounds, ensuring scalability to many-document settings. On needle-in-a-haystack retrieval and long-context QA, GOLD PANNING matches Permutation Self-Consistency’s target identification with 30–65% fewer queries and remains effective under calibration mismatch, suggesting coarse positional ordering drives performance gains. These results demonstrate that inherent model biases need not be failures, but can be used as tools for control.

1. Introduction

Many reasoning tasks require identifying relevant signals within large contexts, including scientific synthesis, legal analysis, and agentic fact retrieval. LLMs have emerged as practical tools for these tasks, enabling knowledge synthesis in retrieval-augmented generation (RAG) (Lewis et al., 2020; Gao et al., 2023) and agentic systems (Wang et al.,

2023a; Yao et al., 2023), but a critical bottleneck undermining reliability is *position bias*, where LLMs often prioritize information based on *where* it appears in context rather than its intrinsic relevance (Wang et al., 2023b; Zheng et al., 2024; Liu et al., 2024). Even state-of-the-art retrieval systems (Ke et al., 2025; Zhang et al., 2025) cannot guarantee optimal placement. When critical information falls into low-attention zones, reasoning failures increase substantially.

Current approaches to mitigate this bias fall into two categories. **White-box mitigations** modify model architectures to reduce bias (Peysakhovich & Lerer, 2023; Hsieh et al., 2024; Wang et al., 2024; Zhang et al., 2024), but require model access or extensive fine-tuning, rendering them impractical for proprietary API models. **Inference-time ensembling**, such as Permutation Self-Consistency (PSC; Tang et al. (2024)), offers a black-box alternative by randomly shuffling documents across multiple queries and aggregating results via majority vote. While PSC averages over random placements, it assumes independent error, an assumption violated in long-context settings where position bias induces strongly correlated failures (Byerly & Khashabi, 2025). Furthermore, PSC is inherently *stateless*, as each shuffle is drawn *independently*, discarding information from previous queries. Finally, PSC is bias-agnostic, treating all positions as exchangeable rather than exploiting known positional reliability differences.

Rather than treating position bias as noise to average away, we propose to *model it explicitly and learn from observations*. We introduce GOLD PANNING, a black-box, inference-time framework that maintains Bayesian beliefs over document relevance and iteratively reorders inputs to align promising documents with positions of high discriminative power. We observe that position bias has a consistent structure (Fig. 1, left pane), with some positions (e.g., the beginning and end of context) reliably surfacing relevant content, while others (the “lost in the middle” zone) tend to obscure it. We characterize this structure by estimating a position’s *diagnosticity* (§2.4), a measure of how reliably a position distinguishes relevant from irrelevant documents. We perform a lightweight, per-model calibration (§4.3, §B) using synthetic needle-in-haystack instances to estimate this positional reliability profile. Empirically, we find this

¹Johns Hopkins University. Correspondence to: Adam Byerly <abyerly2@jhu.edu>.

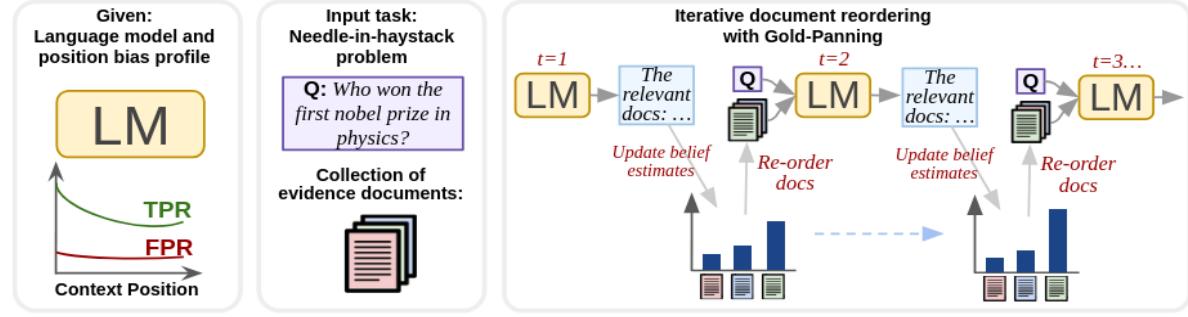


Figure 1. Overview of GOLD PANNING. **Given:** a language model with a measured positional diagnosticity profile (TPR/FPR by position; §2.4). **Task:** identify relevant information from a document collection (§4.1). **Method:** iteratively query the model, update Bayesian beliefs based on observed responses (§2.3), and reorder documents to place high-belief items in high-diagnosticity positions (§2.5). This accelerates belief separation, identifying relevant content with fewer queries than stateless ensembling.

profile transfers across multi-document question answering (MDQA) tasks when the protocol is fixed (§4.4).

Equipped with this profile, we re-imagine the context window not as a passive input buffer, but as an array of noisy detectors with known reliability. By adopting a Bayesian lens (Fig. 1, right pane), we maintain a running probabilistic belief about each document’s relevance, treating every LLM output as a noisy measurement rather than a final answer. This transforms the problem from stateless ensembling into *sequential Bayesian inference*, where each query extracts partial evidence that actively shapes the next input configuration, allowing the system to “pan for gold.”

As document retrieval is fundamentally a *search* problem, we do not need perfect belief calibration for all documents—we only need to efficiently identify the relevant ones. Consequently, reducing uncertainty about irrelevant documents yields only marginal gain towards this goal. In long-context LLM reasoning, the search objective dominates: placing high-belief documents in high-diagnosticity positions consistently outperforms probing uncertain ones.

The underlying mechanism behind this is *signal anchoring*. Beliefs in long-context reasoning are inherently fragile; anchoring increases the probability of receiving high-diagnostic evidence early, which accelerates belief separation and reduces the number of rounds to confidently select the needles. This principle, that **once identified, a needle must be kept in view**, maximizes recall with fewer queries.

GOLD PANNING operationalizes these insights via greedy belief-diagnosticity matching at each round. We evaluate on MDQA, a controlled setting where position bias is well-documented and retrieval success is directly measurable. Across configurations, GP-BELIEF outperforms GP-ENTROPY, confirming that signal anchoring dominates uncertainty reduction. Compared to PSC, GOLD PANNING identifies relevant documents using **30–65% fewer queries** while maintaining F_1 across model families and scales.

Contributions. (1) **Framework:** We formalize position

bias as a *diagnostic signal* and introduce GOLD PANNING, a black-box inference-time method for Bayesian active search over long contexts via calibrated position diagnosticity. (2) **Theory:** We prove that a greedy belief-based assignment identifies targets in $O(\log N)$ LLM queries, reducing the cost of inference. (3) **Finding:** GP-BELIEF outperforms GP-ENTROPY across all configurations, confirming that document retrieval is fundamentally a search problem, not a learning problem. (4) **Empirical:** On MDQA benchmarks, GOLD PANNING matches PSC’s F_1 with 30–65% fewer queries and remains effective under calibration mismatch, confirming that coarse positional ordering suffices.

2. The GOLD PANNING Framework

Our goal is to exploit position bias by repeatedly reordering documents across rounds to reliably surface relevant content with as few LLM queries as possible. We first present the algorithm (§2.1), then formalize each component: documents and positions (§2.2), belief dynamics (§2.3), positional diagnosticity (§2.4), and scoring objective (§2.5).

2.1. Overview and Algorithm

GOLD PANNING operates in rounds. Each round, we (1) score documents based on current beliefs, (2) construct an assignment matching high-scoring documents to high-diagnosticity positions, (3) query the LLM with this arrangement, and (4) update beliefs based on which documents were cited. Algorithm 1 summarizes this procedure: **Lines 3–4** compute positional diagnosticity and establish a fixed ordering of positions by discriminative power (§2.4); **Line 5** initializes uniform prior beliefs; the main loop (**Lines 6–13**) iterates: **Line 7** scores documents according to one of two objectives (§2.5), **Lines 8–9** construct the greedy assignment matching document and position orderings, **Lines 10–11** format and query the LLM, and **Line 12** performs the Bayesian belief updates (§2.3).

2.2. Documents and Positions

We consider N documents indexed by $i \in \{1, \dots, N\}$, each with an unknown binary relevance $Z_i \in \{0, 1\}$, where we wish to identify the subset of documents with $Z_i = 1$ (relevant) using a limited number of LLM calls. For each query, we also have N context positions indexed by $j \in \{1, \dots, N\}$, which we treat as binary *detectors* with true/false positive rates $\text{TPR}_j, \text{FPR}_j$. Letting $O_{ij} \in \{0, 1\}$ denote if the LLM cites document i at position j , we have:

$$\Pr(O_{ij} = 1 \mid Z_i = z) = z \cdot \text{TPR}_j + (1 - z) \cdot \text{FPR}_j, \quad (1)$$

with $\Pr(O_{ij} = 0 \mid \cdot) = 1 - \Pr(O_{ij} = 1 \mid \cdot)$. To ensure structured outputs, we enforce a JSON citation schema via constrained decoding when available, otherwise via format-constrained prompting and validation. The parameters $(\text{TPR}_j, \text{FPR}_j)$ are calibrated once (§B) and reused across tasks. In Algorithm 1, input documents D (**Line 1**) are arranged into a prompt P_t each round (**Line 10**).

2.3. Actions, Observations, and Beliefs

Observations. At each round t , we choose a *permutation* $\sigma_t \in S_N$ (the symmetric group on $\{1, \dots, N\}$) specifying which document is placed in which position; we interpret $\sigma_t(i) = j$ as “document i is placed in position j .” After querying the LLM, we observe $O_{t,i} \in \{0, 1\}$ indicating whether document i was cited at round t . Let $j = \sigma_t(i)$ denote the position assigned to i . The observation model (Eq. (1)) applies with $j = \sigma_t(i)$.

Concretely, the model returns a set of cited document IDs C_t . Given σ_t , we observe one outcome per document, $O_{t,i} = \mathbb{I}[i \in C_t]$, for $i \in \{1, \dots, N\}$.

Belief updates. We maintain posterior beliefs $\mathbf{b}_t = (b_{t,1}, \dots, b_{t,N})$, where $b_{t,i} = \Pr(Z_i = 1 \mid \mathcal{F}_t)$ is our belief that document i is relevant given all observations up to round t . Starting from a uniform prior, $b_{0,i} = 0.5$, we update these beliefs via Bayes’ theorem after each round using the observation model in Eq. (1):

$$b_{t,i} = \frac{b_{t-1,i} \mathcal{P}_1}{b_{t-1,i} \mathcal{P}_1 + (1 - b_{t-1,i}) \mathcal{P}_0}, \quad (2)$$

where $\mathcal{P}_z = \Pr(O_{t,i} \mid Z_i = z, \sigma_t(i) = j)$ (Eq. (1)) is the Bernoulli likelihood under label z for $j = \sigma_t(i)$. When upstream retrieval scores are available, these could inform a non-uniform prior; we adopt the uninformative case for generality and to isolate the contribution of positional inference. In Algorithm 1, beliefs are initialized uniformly (**Line 5**) and updated after each query via Eq. (2) (**Line 12**).

We define the log-odds $\lambda_{t,i} = \log \frac{b_{t,i}}{1 - b_{t,i}}$, which admits an additive Bayesian update (Eq. (2)) of the form $\lambda_{t,i} = \lambda_{t-1,i} + \ell_{\sigma_t(i)}(O_{t,i})$, where $\sigma_t(i)$ is the position assigned to

Algorithm 1 GOLD PANNING

```

1: Input: Query  $Q$ , Documents  $\mathcal{D} = \{d_1, \dots, d_N\}$ , Detector
   Profile  $(\text{TPR}_j, \text{FPR}_j)$ , Rounds  $T$ 
2: Output: Final beliefs  $\mathbf{b}_T$ 
3: Compute  $d_{\text{diag}}(j) \leftarrow |\text{TPR}_j - \text{FPR}_j|$ 
4:  $\pi_{\text{pos}} \leftarrow \text{argsort}(\mathbf{d}_{\text{diag}}, \text{descending})$ 
5: Initialize  $b_{0,i} \leftarrow 0.5$ 
6: for  $t = 1$  to  $T$  do
7:    $s_i \leftarrow \text{SCORE}(b_{t-1,i})$ 
8:    $\pi_{\text{doc}} \leftarrow \text{argsort}(\mathbf{s}, \text{descending})$ 
9:    $\sigma_t \leftarrow \text{zip}(\pi_{\text{doc}}, \pi_{\text{pos}})$ 
10:   $P_t \leftarrow \text{Format}(Q, \text{Permute}(\mathcal{D}, \sigma_t))$ 
11:   $C_t \leftarrow \text{QueryLLM}(P_t)$ 
12:  Update  $\mathbf{b}_t$  via Eq. (2) using citations  $C_t$ 
13: end for
14: return  $\mathbf{b}_T$ 

```

document i at round t , and $\ell_j(o)$ is the position-dependent log-likelihood ratio $\ell_j(o) = \log \frac{\Pr(O=o \mid Z=1, \text{pos}=j)}{\Pr(O=o \mid Z=0, \text{pos}=j)}$. Thus, the log-odds evolve as a random walk with increments determined by the assigned position.

We adopt a conditional-independence approximation where citations are treated as independent given relevance and position. While LLM decoding is autoregressive, we interpret $(\text{TPR}_j, \text{FPR}_j)$ as marginal rates under a fixed prompting protocol. Empirical results (§4) support this approximation.

2.4. Diagnosticity

We quantify how informative a position j is via Youden’s index (Peirce, 1884; Youden, 1950), a measure for assessing the discriminative power of diagnostic procedures. Youden’s J -statistic is defined as $J_j = \text{TPR}_j - \text{FPR}_j$. Under this measure, a position is maximally informative when $J_j = 1$, uninformative when $J_j = 0$, and maximally *anti*-informative when $J_j = -1$. Crucially, when $J_j < 0$, a citation from position j is more likely under irrelevance than relevance, and therefore provides strong *negative* evidence. Because we primarily need to decide *which positions yield the most information*, we rank positions by *magnitude*, $d_{\text{diag}}(j) = |J_j|$, while retaining the sign of J_j through the likelihoods $(\text{TPR}_j, \text{FPR}_j)$ in our Bayesian update (Eq. (2)).

For simplicity, we present the balanced case with N documents and N positions. Asymmetric cases where the number of retrieved documents differs from the number of usable positions can be reduced to this setting by adding *dummy* positions or items; we describe these details in Appendix C.

2.5. Scoring Objectives

GOLD PANNING is a greedy policy that uses calibrated positional diagnosticity and evolving document beliefs to construct assignments.

Document scoring. At each round, we assign a scalar score $s_{t,i}$ to each document i based on its current belief $b_{t-1,i}$, sort documents in descending order, and match the highest-scoring document to the most diagnostic position, second-highest to second-most diagnostic, and so on. The choice of scoring function reflects two natural objectives corresponding to active learning (Settles, 2009; 2011) and active search (Garnett et al., 2012). The first, **Uncertainty Reduction**, prioritizes exploring ambiguous documents to maximize information gain, assigning the most diagnostic positions to *high-entropy* documents. This follows standard intuition from active learning: probe the most uncertain cases to learn fastest. The second, **Target Identification**, prioritizes confirming likely positives to maximize recall, assigning diagnostic positions to *high-belief* documents. We implement these objectives via two distinct scoring variants:

GP-ENTROPY (uncertainty reduction). We set $s_{t,i}^{\text{Entropy}} = \mathcal{H}(b_{t-1,i})$, so high scores correspond to high uncertainty. This variant approximates the information-gain objective by repeatedly testing ambiguous documents.

GP-BELIEF (signal anchoring). We set $s_{t,i}^{\text{Belief}} = b_{t-1,i}$, so high scores correspond to documents believed most likely to be relevant. This variant approximates the active search objective by repeatedly anchoring likely-relevant documents in high-diagnosticity positions.

We formalize both objectives in §3, showing that uncertainty reduction is aligned with maximizing expected entropy decrease under our detector model, while signal anchoring maximizes a belief-weighted sum of position-wise detection rates. The algorithm returns the final posterior beliefs \mathbf{b}_T , which can be mapped to discrete predictions via standard decision rules, as detailed in the experimental setup (§4.1).

2.6. Computational Complexity

The per-round complexity of GOLD PANNING is $O(N \log N)$ for sorting documents by score, compared to $O(N)$ for PSC’s random shuffle. However, this is negligible CPU overhead, as the dominant cost in practice is LLM inference. With respect to the number of costly LLM calls, GOLD PANNING provides substantial savings, identifying targets in $O(\log N)$ rounds (proven in §3.1), compared to $O(T)$ rounds for PSC, where $T \gg \log N$ in practice.

3. Theoretical Properties

We analyze GOLD PANNING as a sequential Bayesian inference process. Here we characterize how the rounds required to identify a relevant document scale with collection size and assigned position quality. The central quantity governing convergence is the belief update *drift*, determined by the diagnostic strength of the assigned positions. All proofs are deferred to Appendix A.

3.1. Belief Dynamics and Sample Complexity

Recall that the log-odds evolve as a random walk with increments determined by the assigned position (§2.3). For a relevant document assigned to position j , the expected increment of this random walk is

$$\mu_j = \mathbb{E}[\ell_j(O) \mid Z = 1] = D_{\text{KL}}(P_{Z=1,j} \parallel P_{Z=0,j}), \quad (3)$$

the KL divergence between the observation distributions induced by position j . Positions with larger μ_j therefore induce faster belief concentration. The total rounds required to identify the relevant document thus depends on the cumulative drift achieved by the policy’s position assignments.

Theorem 3.1 (Sample Complexity). *Assume that log-likelihood ratios are bounded and that the policy assigns the relevant document positions with average drift at least $\underline{\mu} > 0$, i.e., $\sum_{t=1}^T \mu_{\sigma_t(i)} \geq \underline{\mu} \cdot T$. To identify a unique relevant document among N candidates with high probability, the number of rounds required satisfies*

$$T = O\left(\frac{\log N}{\underline{\mu}} + \frac{\log N}{\underline{\mu}^2}\right) = O\left(\frac{\log N}{\underline{\mu}^2}\right).$$

Interpretation. The $\log N$ factor reflects the intrinsic difficulty of distinguishing one relevant document from $N - 1$ distractors. The dominant term scales quadratically with $1/\underline{\mu}$, indicating that convergence speed is governed primarily by the diagnostic quality of the positions assigned to the true positive. Policies that concentrate high-drift positions on promising candidates therefore achieve substantially faster identification.

3.2. Information Rate Advantage of Strategic Placement

Thm 3.1 established that the convergence time is inversely proportional to the assigned drift $\underline{\mu}$. Consequently, the efficiency of the retrieval process depends entirely on the policy’s ability to maximize this quantity. We now quantify exactly how much GP-BELIEF improves the information rate compared to random assignment.

Let $\mu_{(1)} \geq \dots \geq \mu_{(N)}$ denote the position drifts sorted in descending order of quality. Let $\text{rank}_t(i) \in \{1, \dots, N\}$ be the rank of document i when sorting current beliefs $b_{t-1,i}$ in descending order.

Definition 3.2 (Information Rate). The *information rate* for a relevant document i under policy π at round t is the expected drift assigned to it, conditioned on the history:

$$\mathcal{R}_\pi(t) = \mathbb{E}[\mu_{\sigma_t(i)} \mid Z_i = 1, \mathcal{F}_{t-1}].$$

For the baseline **Permutation Self-Consistency (PSC)**, which assigns documents uniformly at random, the rate is fixed at the average drift: $\mathcal{R}_{\text{PSC}}(t) = \bar{\mu} = \frac{1}{N} \sum_{j=1}^N \mu_j$.

Proposition 3.3 (Exact Rate for GP-BELIEF). *Under GP-BELIEF, the information rate is:*

$$\mathcal{R}_{\text{GP-BELIEF}}(t) = \sum_{r=1}^N \mu_{(r)} \cdot \Pr(\text{rank}_t(i) = r \mid Z_i = 1, \mathcal{F}_{t-1}).$$

To interpret this, we look at the probability of the relevant document being “captured” in the high-quality positions.

Definition 3.4 (Top- m Capture Probability). For any cut-off $m \in \{1, \dots, N\}$, let

$$p_t(m) = \Pr(\text{rank}_t(i) \leq m \mid Z_i = 1, \mathcal{F}_{t-1}),$$

be the probability that the relevant document is ranked among the top m candidates.

Proposition 3.5 (Rate Advantage). *For any $m \in \{1, \dots, N\}$, the information rate of GP-BELIEF satisfies the lower bound:*

$$\mathcal{R}_{\text{GP-BELIEF}}(t) \geq p_t(m) \mu_{(m)} + (1 - p_t(m)) \mu_{(N)}.$$

Interpretation. This bound isolates the two drivers of GOLD PANNING’s performance: (1) **Ranking Quality** ($p_t(m)$): As the algorithm progresses, the probability of the relevant document being in the top m increases. Under random guessing (PSC), this probability is simply m/N . Whenever $p_t(m) > m/N$, GOLD PANNING is guaranteed to outperform the baseline. (2) **Position Heterogeneity** ($\mu_{(m)} - \mu_{(N)}$): The advantage is amplified when the position profile is “spiky.” If the top positions ($\mu_{(m)}$) are better than the worst ones ($\mu_{(N)}$), identifying the correct document accelerates in subsequent rounds.

This creates a positive feedback loop: a good ranking yields a high drift μ , which accumulates evidence faster, which further improves the ranking $p_{t+1}(m)$.

3.3. Why Anchoring Dominates Uncertainty Reduction

Prop. 3.3–3.5 establish that assigning documents to high-drift positions accelerates convergence. The central question remains: *which* documents should receive these positions?

GP-ENTROPY prioritizes high-uncertainty documents (maximizing information gain), while GP-BELIEF prioritizes high-belief documents (maximizing expected discoveries). In sparse search settings ($k \ll N$), these objectives diverge sharply. We formalize this divergence by analyzing the drift allocated to the true positive as it emerges from the haystack.

Proposition 3.6 (GP-ENTROPY Demotes Confident Positives). *Consider the $k = 1$ setting. As the belief in the true positive approaches certainty ($b_{t-1,i} \rightarrow 1$), its entropy $H(b_{t-1,i}) \rightarrow 0$. Consequently, if there exist at least m distractor documents with higher entropy (e.g., $b \approx 0.5$),*

GP-ENTROPY assigns the true positive a rank greater than m . In the limit, the diagnostic attention collapses to the minimum available drift:

$$\lim_{b_{t-1,i} \rightarrow 1} \mathcal{R}_{\text{GP-ENTROPY}}(t) = \mu_{(N)}.$$

Proposition 3.7 (GP-BELIEF Diagnostic Lock-in). *Consider the $k = 1$ setting. Under GP-BELIEF, if the true positive i achieves the highest belief rank at round τ , it receives maximal diagnostic attention $\mu_{(1)}$. Furthermore, there exists a strictly positive probability that i retains the lead in all subsequent rounds $t > \tau$, maintaining $\mathcal{R}_{\text{GP-BELIEF}}(t) = \mu_{(1)}$.*

Interpretation: Vicious vs. Virtuous Cycles.

The Vicious Cycle of Entropy: GP-ENTROPY penalizes success. As the model begins to identify the relevant document (belief rises), its entropy falls. The policy effectively says, “We are sure about this one, so let’s stop checking it.” It demotes the candidate to the noisiest positions ($\mu_{(N)}$), starving it of the signal required to cross the final confidence threshold $(1 - \delta)$. This explains the stagnation observed in Figure 2.

The Virtuous Cycle of Belief: GP-BELIEF rewards success. As the relevant document rises in rank, it is promoted to positions with higher drift ($\mu_{(1)}$). This strong signal reinforces the belief, further solidifying its top rank. This “Signal Anchoring” creates a lock-in effect that races the winner to the threshold with maximal velocity.

4. Empirical Validation

4.1. Setup

Benchmarks. We evaluate on two FLenQA benchmarks (Levy et al., 2024): **MonoRel** (monotonic transitive reasoning) and **PIR** (People in Rooms; compositional reasoning). We selected these tasks because they offer unambiguous, deterministic ground-truth labels. In contrast to open-domain QA, where implicit ambiguity can confound evaluation (Min et al., 2020) and necessitate complex semantic matching (Kamaloo et al., 2023), FLenQA allows for precise measurement of retrieval recall without the noise of partial relevance or answer granularity mismatches. While these benchmarks are natively single-fact QA, we convert each instance into a controlled MDQA task by augmenting the ground-truth fact with $N - 1$ sampled distractor, yielding exactly one relevant document.

Models. We test two model families at multiple scales: **Gemma-3** (Team et al., 2025) (12B, 27B) and **OLMo-3** (Olmo et al., 2025) (OLMo-3-7B, OLMo-3.1-32B). These models exhibit varying degrees of position bias as measured by their calibrated positional profiles (§B), allowing us to isolate when position-aware reordering helps.

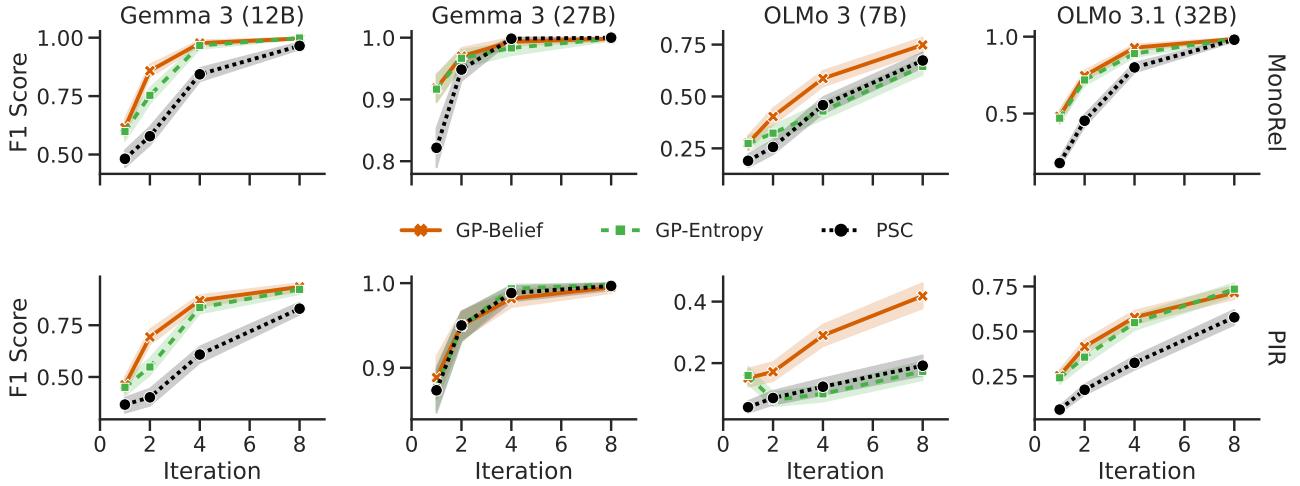


Figure 2. Main results across models and tasks. F_1 score over 8 iterations for $N=100$ documents. GP-BELIEF (orange) consistently outperforms GP-ENTROPY (green, dashed) and PSC (black, dotted), with advantages most pronounced on models exhibiting strong position bias. Gemma-3-27B shows no separation due to its lack of exploitable bias. Shaded regions indicate 95% confidence intervals.

Default configuration. Unless stated otherwise, we use $N = 100$ documents with $k = 1$ relevant document (Top-1 selection) and run $T = 8$ rounds. We follow model-recommended sampling parameters: temperature 1.0 with top- p 0.95 for Gemma, and temperature 0.6 with top- p 0.95 for OLMo. All methods enforce structured JSON outputs via constrained decoding to ensure parseable citations. Each round returns a set of cited document IDs, which we map to per-document indicators $O_{t,i}$ (Eq. (1)) for belief updates.

Baseline. We compare against **Permutation Self-Consistency (PSC)** (Tang et al., 2024), which queries the model on independent random permutations and aggregates per-document relevance predictions across rounds via per-document majority vote (ties broken deterministically). PSC is a black-box baseline that mitigates position bias by averaging over placements. All methods use the same number of LLM calls, enabling direct comparison of query efficiency.

Evaluation. GOLD PANNING maintains continuous beliefs $\mathbf{b}_t \in [0, 1]^N$, converted to binary predictions for evaluation. We assume the number of relevant documents k is known and given, and select the k highest-belief documents (Top- k). We report mean F_1 with 95% confidence interval.

Calibration of diagnosticity. For each model, we run a lightweight calibration to extract its diagnosticity profile. At position j , we construct calibration trials by placing a known-relevant “gold” document at position j , filling the remaining $N - 1$ positions with distractors, and query the model. Citing the gold at j contributes to TPR_j , while citing position j when gold is elsewhere contributes to FPR_j . Estimating diagnosticity at all N positions can be expensive for long contexts, even though this is a one-time cost. We therefore use sparse fixed-grid sampling with linear interpolation, assuming smooth diagnosticity trends. Specifically,

we evaluate a fixed grid of $K = 11$ positions with 50 trials each and linearly interpolate to obtain the full N -position profile (validated in §4.3; full details in Appendix B).

Fig. 3 shows the calibrated diagnosticity profile for Gemma-3-12B on both MonoRel and PIR with $N = 100$ documents. Both exhibit the characteristic “lost in the middle” pattern: positions near the beginning (primacy) and end (recency) of context show high diagnosticity, while middle positions are far less informative. This heterogeneity enables strategic placement—GP-BELIEF exploits the high-diagnosticity positions to accelerate belief separation.

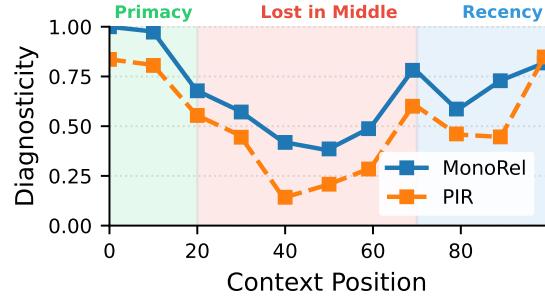


Figure 3. Diagnosticity profile for Gemma-3-12B on MonoRel and PIR ($N = 100$). Both tasks show primacy and recency effects, i.e., diminished diagnosticity in the middle.

4.2. Main Results

Fig. 2 compares GP-BELIEF, GP-ENTROPY, and PSC across four models. We analyze whether (i) our stateful Bayesian inference beats stateless ensembling, and (ii) signal anchoring outperforms uncertainty reduction.

Signal anchoring (GP-BELIEF) dominates uncertainty reduction (GP-ENTROPY). Whenever position bias is

exploitable, GP-BELIEF consistently outperforms GP-ENTROPY. On Gemma-3-12B (PIR), GP-BELIEF achieves $F_1 \approx 0.88$ at iteration 4 versus 0.72 for GP-ENTROPY, a gap that persists across model scales. This confirms that in needle-in-haystack retrieval ($k \ll N$), the optimal strategy is not to probe ambiguity, but to *protect likely signals*. GP-ENTROPY fails because it prioritizes documents near $b \approx 0.5$ —predominantly irrelevant distractors not yet ruled out. By contrast, GP-BELIEF *anchors likely signals* in high-diagnosticity positions to accelerate separation.

Bayesian statefulness beats stateless ensembling. Even with the suboptimal entropy objective, the Bayesian framework generally outperforms the stateless PSC baseline. On OLMO-3-7B (PIR), GP-ENTROPY attains $F_1 \approx 0.25$ at iteration 8 compared to PSC’s 0.17. The gains come from *adaptivity*: PSC treats iterations independently, often re-testing resolved negatives in high-value positions, whereas GOLD PANING accumulates evidence, deprioritizing low-belief documents to make room for promising ones. When combined with GP-BELIEF, this statefulness enables the theoretical information rate gains of Proposition 3.5.

Query efficiency. GP-BELIEF matches PSC using **30–65% fewer queries**. On Gemma-3-12B (MonoRel), GP-BELIEF reaches $F_1 \approx 0.97$ in just 4 iterations, matching PSC’s performance at iteration 8. On PIR, the gains are even larger: GP-BELIEF’s iteration-4 performance (0.88) strictly exceeds PSC’s final iteration-8 result (0.82), making it a viable strategy for latency-constrained applications.

No bias, no gains. Gemma-3-27B (Fig. 2, column 2) serves as a control. Because this model exhibits a fairly flat diagnosticity profile (Fig. 6), strategic reordering offers no leverage over random shuffling, and all methods converge indistinguishably. This negative result validates that GOLD PANING is not a “magic wand,” but a targeted intervention that turns position bias into a resource.

4.3. Validating Sparse Diagnosticity Calibration

As discussed in §4.1, we estimate diagnosticity via sparse position sampling. We validate this by comparing fixed-grid interpolation against uniform calibration over *all* positions. Tab. 1 reports the mean absolute residual Δ between fixed-grid interpolated values and uniformly-sampled ground truth ($N = 100$); **across models**, $\Delta < 0.071$. Thus, a grid of $K = 11$ positions recovers diagnosticity with minimal bias, *supporting lightweight calibration for full-scale evaluation*.

4.4. Cross-dataset Transferability

If position bias is a structural model property (as suggested by Fig. 3) rather than dataset-specific, diagnosticity profiles should transfer across datasets. We test this via cross-evaluation: *applying the diagnosticity profile calibrated*

Table 1. Mean absolute residual Δ (mean \pm SE) between uniformly-sampled ($N = 100$) and fixed-grid interpolated diagnosticity ($K = 11$ positions). All values < 0.071 confirm that sparse grid sampling accurately recovers the full profile.

Dataset	Model	Mean Δ Diagnosticity
MonoRel	Olmo-3-7B-Think	0.054 ± 0.002
	Olmo-3.1-32B-Think	0.051 ± 0.002
	Gemma-3-12B-IT	0.027 ± 0.002
	Gemma-3-27B-IT	0.008 ± 0.000
PIR	Olmo-3-7B-Think	0.041 ± 0.002
	Olmo-3.1-32B-Think	0.071 ± 0.003
	Gemma-3-12B-IT	0.028 ± 0.002
	Gemma-3-27B-IT	0.013 ± 0.001

on MonoRel to guide GP on PIR, and vice-versa. As shown in Fig. 4, **GP-BELIEF retains its advantage under cross-calibration**, with only modest degradation relative to matched calibration. This suggests most gains come from *coarse positional ordering*, i.e., with the ranking of positions being more important than precise TPR/FPR values.

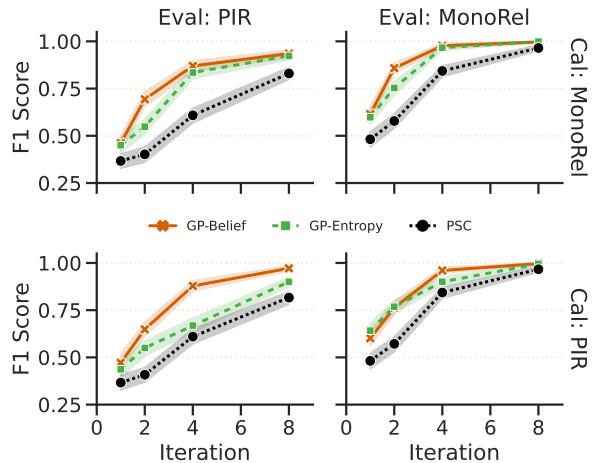


Figure 4. **Cross-task transferability for Gemma-3-12B** ($N = 100$). Rows indicate calibration source; columns indicate evaluation task. GP-BELIEF maintains its advantage even under cross-calibration (off-diagonal), with only modest degradation compared to matched calibration (diagonal). This robustness suggests that coarse positional ordering, the relative ranking of high- vs. low-diagnosticity positions, drives most of the gain.

4.5. Controlled Ablations with Simulated Environment

Real LLM experiments confound multiple factors: calibration error, intrinsic model variability, and task-specific noise. To isolate the conditions under which GP-BELIEF succeeds, we turn to simulation, where we can precisely control detector heterogeneity, calibration accuracy, and context scale independently. We characterize GP-BELIEF’s operating conditions through simulation (Fig. 5; Top-1 selection, $N = 100$, $k = 1,500$ trials unless noted).

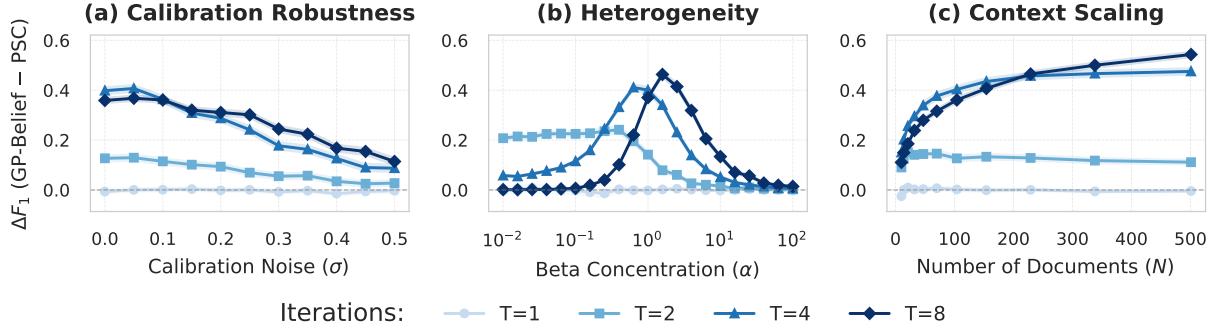


Figure 5. Ablation studies characterizing GP-BELIEF’s operating conditions (simulation; Top-1 selection, $N=100$, $k=1$, 5,000 trials unless noted). (a) Calibration robustness: advantage degrades gracefully as Gaussian noise (σ) is injected into TPR/FPR estimates, confirming that relative position rankings matter more than precise calibration. (b) Heterogeneity: advantage peaks at moderate Beta concentration ($\alpha \approx 1$) and vanishes as positions become homogeneous ($\alpha \rightarrow \infty$), since uniform detectors offer no strategic leverage. (c) Context scaling: advantage grows with N , reaching $\Delta F_1 \approx 0.5$ at $N = 500$, confirming that strategic placement becomes critical as random placement increasingly misses high-diagnosticity positions.

Robustness to calibration error (a). We test GP-BELIEF under miscalibration by adding Gaussian noise ($\sigma \in [0, 0.5]$) to TPR/FPR estimates (Fig. 5a). Performance degrades smoothly: even at $\sigma = 0.4$, GP-BELIEF retains a clear advantage over PSC ($\Delta F_1 \approx 0.15\text{--}0.25$). This indicates the method relies on relative ranking of positional diagnosticities rather than their precise absolute values.

Heterogeneity (b). GP-BELIEF’s advantage depends on profile variation. When detector quality is drawn from $\text{BETA}(\alpha, \alpha)$, advantage peaks at moderate concentration ($\alpha \approx 1$) and vanishes as positions become homogeneous ($\alpha \rightarrow \infty$), since uniform detectors are indistinguishable and therefore offer no opportunity for strategic leverage.

Context scaling (c). The advantage grows with N , reaching $\Delta F_1 \approx 0.5$ at $N = 500$. As context expands, random placement rarely assigns relevant documents to diagnostic positions, amplifying the value of strategic placement.

5. Related Work

Exploiting systematic biases. Systematic biases are typically viewed as obstacles, but prior work in human-computer interaction has leveraged cognitive biases to optimize interfaces (Gajos & Weld, 2005; Gajos et al., 2008) and guide user behavior (Cialdini & Cialdini, 2007; Mathur et al., 2019). In LLMs, predictable regularities enable similar efficiency gains, as in speculative decoding and early-exit methods (Leviathan et al., 2023; Schuster et al., 2022; Chen et al., 2024). In a similar spirit, we leverage position bias as a *diagnostic signal* to improve in-context NIAH retrieval.

Position bias in LLMs. Position bias in long-context LLMs is well-studied (Wang et al., 2023b; Zheng et al., 2023; Liu et al., 2024, *inter alia*), with mitigations broadly divided into white-box and black-box approaches. *White-box methods* suppress bias by modifying model internals, includ-

ing alternative positional encodings (Zhang et al., 2024) and attention manipulations (Peysakhovich & Lerer, 2023; Hsieh et al., 2024; Wang et al., 2024). *Black-box methods* instead treat the model as fixed, such as Permutation Self-Consistency (PSC) (Tang et al., 2024), which averages predictions over random permutations. PSC, however, is stateless and does not exploit information accumulated across queries. While Liu et al. (2025) explores relevance via internal model states, we are the first to formalize inference-time reordering as *sequential Bayesian active search*, enabling stateful inference using only API-level outputs.

Active search and iterative retrieval. We frame context reordering as a *Bayesian active search* problem (Garnett et al., 2012; Jiang et al., 2017), distinct from conventional active learning (Settles, 2009; 2011), which prioritizes uncertainty reduction. Bandit-based methods (Slivkins, 2019) have been applied to optimize retrieval in RAG systems (Duan et al., 2025), but these primarily operate at training time. At inference time, GOLD PANNING is most closely related to iterative reranking methods such as RankGPT (Sun et al., 2023), but differs fundamentally: rather than prompting the model to explicitly sort candidates (itself subject to position bias), we infer relevance from the model’s natural responses under structured layouts, requiring no instruction tuning.

6. Conclusion

We introduce GOLD PANNING, a Bayesian framework that exploits position bias as a diagnostic signal for long-context retrieval. By treating positions as noisy detectors and actively anchoring high-belief documents in high-diagnosticity positions, signal anchoring consistently outperforms uncertainty-based baselines, achieving large query reductions with matched asymptotic accuracy. More broadly, our results suggest that predictable inference-time biases can be leveraged as information rather than treated as noise.

Impact Statement

This paper presents a method for improving the efficiency of long-context retrieval in large language models. By reducing the number of LLM queries required to identify relevant documents, GOLD PANNING may lower computational costs and latency in retrieval-augmented generation systems, with associated reductions in energy consumption.

Our method exploits position bias, a known limitation of current LLMs. While we frame this bias as a resource rather than a flaw, we do not believe this work discourages efforts to reduce position bias at the model level. We see no direct ethical concerns specific to this work beyond those inherent to improving information retrieval systems generally.

References

Byerly, A. and Khashabi, D. Self-consistency falls short! the adverse effects of positional bias on long-context problems. 2025. URL <https://arxiv.org/abs/2411.01101>.

Chen, Y., Pan, X., Li, Y., Ding, B., and Zhou, J. EE-LLM: Large-scale training and inference of early-exit large language models with 3D parallelism. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 7163–7189. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/chen24ae.html>.

Cialdini, R. B. and Cialdini, R. B. *Influence: The psychology of persuasion*, volume 55. Collins New York, 2007.

Duan, S., Li, X., Liu, Z., Yi, X., Yan, Y., Wang, S., Gu, Y., Yu, G., and Sun, M. Chunks as arms: Multi-armed bandit-guided sampling for long-context llm preference optimization. *arXiv preprint arXiv:2508.13993*, 2025.

Gajos, K. and Weld, D. S. Preference elicitation for interface optimization. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pp. 173–182, 2005.

Gajos, K. Z., Wobbrock, J. O., and Weld, D. S. Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 1257–1266, 2008.

Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, H., and Wang, H. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2(1), 2023.

Garnett, R., Krishnamurthy, Y., Xiong, X., Schneider, J., and Mann, R. Bayesian optimal active search and surveying. In *International conference on machine learning*, pp. 843–850. PMLR, 2012.

Hsieh, C.-Y., Chuang, Y.-S., Li, C.-L., Wang, Z., Le, L., Kumar, A., Glass, J., Ratner, A., Lee, C.-Y., Krishna, R., et al. Found in the middle: Calibrating positional attention bias improves long context utilization. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 14982–14995, 2024.

Jiang, S., Malkomes, G., Converse, G., Shofner, A., Moseley, B., and Garnett, R. Efficient nonmyopic active search. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1714–1723. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/jiang17d.html>.

Kamalloo, E., Dziri, N., Clarke, C., and Rafiei, D. Evaluating open-domain question answering in the era of large language models. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5591–5606, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.307. URL <https://aclanthology.org/2023.acl-long.307/>.

Ke, W., Zheng, Y., Li, Y., Xu, H., Nie, D., Wang, P., and He, Y. Large language models in document intelligence: A comprehensive survey, recent advances, challenges and future trends. *ACM Transactions on Information Systems*, 2025.

Leviathan, Y., Kalman, M., and Matias, Y. Fast inference from transformers via speculative decoding. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 19274–19286. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/leviathan23a.html>.

Levy, M., Jacoby, A., and Goldberg, Y. Same task, more tokens: the impact of input length on the reasoning performance of large language models. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15339–15353, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.818. URL <https://aclanthology.org/2024.acl-long.818/>.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttrler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474, 2020.

Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., and Liang, P. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12, 2024.

Liu, Z., Amjad, R. A., Adkathimar, R., Wei, T., and Tong, H. Selfelicit: Your language model secretly knows where is the relevant evidence. *arXiv preprint arXiv:2502.08767*, 2025.

Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., and Narayanan, A. Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on human-computer interaction*, 3 (CSCW):1–32, 2019.

Min, S., Michael, J., Hajishirzi, H., and Zettlemoyer, L. AmbigQA: Answering ambiguous open-domain questions. In Webber, B., Cohn, T., He, Y., and Liu, Y. (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5783–5797, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.466. URL <https://aclanthology.org/2020.emnlp-main.466/>.

Olmo, T., Ettinger, A., Bertsch, A., Kuehl, B., Graham, D., Heineman, D., Groeneveld, D., Brahman, F., Timbers, F., Ivison, H., et al. Olmo 3. *arXiv preprint arXiv:2512.13961*, 2025.

Peirce, C. S. The numerical measure of the success of predictions. *Science*, (93):453–454, 1884.

Peysakhovich, A. and Lerer, A. Attention sorting combats recency bias in long context language models. *arXiv preprint arXiv:2310.01427*, 2023.

Schuster, T., Fisch, A., Gupta, J., Dehghani, M., Bahri, D., Tran, V., Tay, Y., and Metzler, D. Confident adaptive language modeling. *Advances in Neural Information Processing Systems* (NeurIPS), 35:17456–17472, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/6fac9e316a4ae75ea244ddcef1982c71-Paper.pdf.

Settles, B. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.

Settles, B. From theories to queries: Active learning in practice. In *Active learning and experimental design workshop in conjunction with AISTATS 2010*, pp. 1–18. JMLR Workshop and Conference Proceedings, 2011.

Slivkins, A. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.

Sun, W., Yan, L., Ma, X., Wang, S., Ren, P., Chen, Z., Yin, D., and Ren, Z. Is ChatGPT good at search? investigating large language models as re-ranking agents. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 14918–14937, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.923. URL <https://aclanthology.org/2023.emnlp-main.923/>.

Tang, R., Zhang, C., Ma, X., Lin, J., and Türe, F. Found in the middle: Permutation self-consistency improves listwise ranking in large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 2327–2340, 2024.

Team, G., Kamath, A., Ferret, J., Pathak, S., Vieillard, N., Merhej, R., Perrin, S., Matejovicova, T., Ramé, A., Rivière, M., et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L., and Anandkumar, A. Voyager: An open-ended embodied agent with large language models. *Transactions on Machine Learning Research*, 2023a.

Wang, Y., Cai, Y., Chen, M., Liang, Y., and Hooi, B. Primacy effect of chatgpt. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 108–115, 2023b.

Wang, Z., Zhang, H., Li, X., Huang, K.-H., Han, C., Ji, S., Kakade, S. M., Peng, H., and Ji, H. Eliminating position bias of language models: A mechanistic approach. *arXiv preprint arXiv:2407.01100*, 2024.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023.

Youden, W. J. Index for rating diagnostic tests. *Cancer*, 3 (1):32–35, 1950.

Zhang, Q., Chen, S., Bei, Y., Yuan, Z., Zhou, H., Hong, Z., Chen, H., Xiao, Y., Zhou, C., Chang, Y., et al. A survey of graph retrieval-augmented generation for customized large language models. *arXiv preprint arXiv:2501.13958*, 2025.

Zhang, Z., Chen, R., Liu, S., Yao, Z., Ruwase, O., Chen, B., Wu, X., Wang, Z., et al. Found in the middle: How language models use long contexts better via plug-and-play positional encoding. *Advances in Neural Information Processing Systems*, 37:60755–60775, 2024.

Zheng, C., Zhou, H., Meng, F., Zhou, J., and Huang, M. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*, 2023.

Zheng, C., Zhou, H., Meng, F., Zhou, J., and Huang, M. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=shr9PXz7T0>.

A. Proofs

In this section, we provide proofs for the theoretical properties of GOLD PANNING. We begin by defining the filtration and the martingale structure of the belief updates.

Setup and notation. Let (Ω, \mathcal{F}, P) be a probability space carrying the sequence of document relevance variables Z and observation sequence O . Let $\mathcal{F}_t = \sigma(O_1, \dots, O_t, \sigma_1, \dots, \sigma_{t+1})$ be the filtration representing the information available after round t (including the policy’s choice for round $t+1$). The log-odds for a fixed document i update as:

$$\lambda_{t,i} = \lambda_{t-1,i} + \ell_{\sigma_t(i)}(O_{t,i}),$$

where $\ell_j(o)$ is the log-likelihood ratio at position j .

A.1. Proof of Consistency

Proof. Consider a single document i . Assume without loss of generality that i is relevant ($Z_i = 1$); the case for $Z_i = 0$ follows by symmetry with negative drift. Let $X_t = \ell_{\sigma_t(i)}(O_{t,i})$ denote the update increment at round t . We can decompose the accumulated log-odds $\lambda_{t,i} = \sum_{s=1}^t X_s$ into a drift term and a martingale difference sequence.

Define the expected drift at step s conditioned on the history:

$$\mu_s = \mathbb{E}[X_s \mid \mathcal{F}_{s-1}, Z_i = 1].$$

By definition of the log-likelihood ratio, this expectation is the KL divergence:

$$\mu_s = D_{\text{KL}}(P_{Z=1, \sigma_s(i)} \parallel P_{Z=0, \sigma_s(i)}).$$

Let M_t be the centered process:

$$M_t = \sum_{s=1}^t (X_s - \mu_s).$$

M_t is a martingale with respect to \mathcal{F}_t because $\mathbb{E}[M_t - M_{t-1} \mid \mathcal{F}_{t-1}] = \mathbb{E}[X_t - \mu_t \mid \mathcal{F}_{t-1}] = 0$.

As in Thm. 3.1, we assume the drift is bounded, and so $|X_s| \leq L$. Consequently, the martingale increments are also bounded: $|X_s - \mu_s| \leq 2L$. By the Strong Law of Large Numbers for Martingales, we have:

$$\lim_{t \rightarrow \infty} \frac{M_t}{t} = 0 \quad \text{almost surely.}$$

Now consider the log-odds normalized by time:

$$\frac{\lambda_{t,i}}{t} = \frac{1}{t} \sum_{s=1}^t \mu_s + \frac{M_t}{t}.$$

As $t \rightarrow \infty$, the second term vanishes. For the first term, the boundedness guarantees that the average drift is bounded away from zero:

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mu_s \geq \underline{\mu} > 0.$$

Therefore, $\liminf_{t \rightarrow \infty} \frac{\lambda_{t,i}}{t} \geq \underline{\mu} > 0$. This implies $\lim_{t \rightarrow \infty} \lambda_{t,i} = +\infty$ almost surely. Since the belief $b_{t,i} = \sigma(\lambda_{t,i})$ is the sigmoid of the log-odds, $b_{t,i} \xrightarrow{a.s.} 1$, recovering the ground truth $Z_i = 1$. \square

A.2. Proof of Thm. 3.1 (Sample Complexity)

Proof. We analyze the probability that the log-odds $\lambda_{T,i}$ fail to reach the threshold Λ by time T ; assume $Z_i = 1$. We are given that for all $t \leq T$, the assigned positions satisfy $\mu_{\sigma_t(i)} \geq \underline{\mu}$. Decomposing the log-odds as before:

$$\lambda_{T,i} = \sum_{t=1}^T \mu_t + M_T \geq T\underline{\mu} + M_T.$$

A failure occurs if $\lambda_{T,i} < \Lambda$. This implies:

$$T\underline{\mu} + M_T < \Lambda \implies M_T < \Lambda - T\underline{\mu}.$$

We apply the Azuma-Hoeffding inequality to the martingale M_T . The increments $Y_t = X_t - \mu_t$ are bounded within an interval of size at most $2L$ (since $X_t \in [-L, L]$). The inequality states that for any $K > 0$:

$$\Pr(M_T \leq -K) \leq \exp\left(-\frac{K^2}{2TL^2}\right).$$

Let us choose T large enough such that the expected drift $T\underline{\mu}$ exceeds the threshold Λ by a safety margin K . Specifically, we set the safety margin to be half the total drift:

$$\text{Let } T\underline{\mu} = 2\Lambda \implies \Lambda - T\underline{\mu} = -\Lambda.$$

However, we also need to satisfy the confidence requirement δ' . We set T according to the bound in the proposition:

$$T = \frac{2\Lambda}{\underline{\mu}} + \frac{8L^2}{\underline{\mu}^2} \log\left(\frac{1}{\delta'}\right).$$

Substituting this T into the failure condition is algebraically complex, so we verify the bound using the safety margin approach directly. Let us define the target margin $K = T\underline{\mu} - \Lambda$. We require $\Pr(M_T < -K) \leq \delta'$. Using Azuma-Hoeffding, it suffices to have:

$$\begin{aligned} \exp\left(-\frac{K^2}{2TL^2}\right) \leq \delta' &\implies \frac{K^2}{2TL^2} \geq \log(1/\delta') \\ &\implies K \geq L\sqrt{2T \log(1/\delta')}. \end{aligned}$$

Substitute $K = T\underline{\mu} - \Lambda$:

$$T\underline{\mu} - \Lambda \geq L\sqrt{2T \log(1/\delta')}.$$

We now check if the proposed T satisfies this inequality. The proposed T has two terms. Let $T_1 = \frac{2\Lambda}{\underline{\mu}}$ and $T_2 = \frac{8L^2}{\underline{\mu}^2} \log(1/\delta')$. So $T = T_1 + T_2$. The LHS of the inequality becomes:

$$(T_1 + T_2)\underline{\mu} - \Lambda = (2\Lambda + T_2\underline{\mu}) - \Lambda = \Lambda + T_2\underline{\mu} > T_2\underline{\mu}.$$

The RHS requires us to bound $L\sqrt{2(T_1 + T_2) \log(1/\delta')}$. Notice that for δ' small enough (or Λ small relative to the noise), the stochastic term dominates. To strictly prove the bound holds, we can simply ensure $T_2\underline{\mu}$ alone is sufficient to cover the noise from the *entire* duration T . Observe that $T_2\underline{\mu} = \frac{8L^2}{\underline{\mu}} \log(1/\delta')$. Squaring the LHS requirement (conservatively using just $T_2\underline{\mu}$):

$$(T_2\underline{\mu})^2 = \frac{64L^4}{\underline{\mu}^2} (\log(1/\delta'))^2.$$

Squaring the RHS (substituting $T \approx T_2$ for the worst case noise regime):

$$2T_2L^2 \log(1/\delta') = \frac{16L^4}{\underline{\mu}^2} (\log(1/\delta'))^2.$$

Since $64 > 16$, the drift from T_2 alone is more than sufficient to cover the variance from T_2 . The addition of T_1 adds more drift (Λ) than variance (since variance grows as \sqrt{T} while drift grows as T), so the inequality holds comfortably. Thus, with probability at least $1 - \delta'$, the threshold is crossed by time T .

We note that the derived bound for T is an upper bound and is not tight. The constants 2 and 8 are conservative estimates chosen to simplify the application of the Azuma-Hoeffding inequality and ensure the safety margin K is robust against worst-case noise. While a tighter constant might be achievable through more complex analysis (e.g., using exact characteristic functions), the current bound is sufficient to establish the fundamental logarithmic scaling, $T \propto \log N$, required for efficient search. \square

A.3. Proofs of Information Rate Properties (Section 3.2)

Setup. Recall that $\mu_{(1)} \geq \mu_{(2)} \geq \dots \geq \mu_{(N)}$ denotes the fixed position drifts sorted in descending order. Let $\text{rank}_t(i)$ denote the rank of document i according to the beliefs b_{t-1} at the start of round t . The GP-BELIEF policy assigns the document with rank r to the position with drift $\mu_{(r)}$. Thus, if $\text{rank}_t(i) = r$, the drift assigned to document i is $\mu_{\sigma_t(i)} = \mu_{(r)}$.

A.3.1. PROOF OF PROPOSITION 3.3 (EXACT RATE)

Proof. The information rate is defined as the expected drift assigned to a relevant document i ($Z_i = 1$) conditioned on the history \mathcal{F}_{t-1} :

$$\mathcal{R}_{\text{GP-BELIEF}}(t) = \mathbb{E} [\mu_{\sigma_t(i)} \mid Z_i = 1, \mathcal{F}_{t-1}] .$$

Since the policy is deterministic given the beliefs (which are fixed in \mathcal{F}_{t-1}), the only randomness comes from the uncertainty of the rank itself relative to the ground truth. Conditioning on the specific rank r that the document achieves:

$$\mathbb{E} [\mu_{\sigma_t(i)}] = \sum_{r=1}^N \mathbb{E} [\mu_{\sigma_t(i)} \mid \text{rank}_t(i) = r] \cdot \Pr(\text{rank}_t(i) = r) .$$

Under GP-BELIEF, the assignment is strictly coupled: the event $\{\text{rank}_t(i) = r\}$ implies that the document is assigned to the position with drift $\mu_{(r)}$. Therefore, $\mathbb{E} [\mu_{\sigma_t(i)} \mid \text{rank}_t(i) = r] = \mu_{(r)}$. Substituting this back into the summation yields the exact rate:

$$\mathcal{R}_{\text{GP-BELIEF}}(t) = \sum_{r=1}^N \mu_{(r)} \cdot \Pr(\text{rank}_t(i) = r \mid Z_i = 1, \mathcal{F}_{t-1}) .$$

□

A.3.2. PROOF OF PROPOSITION 3.5 (RATE ADVANTAGE)

Proof. We seek a lower bound on $\mathcal{R}_{\text{GP-BELIEF}}(t)$ based on the Top- m capture probability $p_t(m)$. We split the summation from Proposition 3.3 into two parts: the top m ranks (where the document is "captured") and the remaining $N - m$ ranks.

$$\begin{aligned} \mathcal{R}_{\text{GP-BELIEF}}(t) &= \underbrace{\sum_{r=1}^m \mu_{(r)} \Pr(\text{rank}_t(i) = r)}_{\text{Term 1}} \\ &\quad + \underbrace{\sum_{r=m+1}^N \mu_{(r)} \Pr(\text{rank}_t(i) = r)}_{\text{Term 2}} . \end{aligned}$$

Bounding Term 1: The drifts are sorted descending, so for any rank $r \leq m$, we have $\mu_{(r)} \geq \mu_{(m)}$.

$$\begin{aligned} \text{Term 1} &\geq \sum_{r=1}^m \mu_{(m)} \Pr(\text{rank}_t(i) = r) \\ &= \mu_{(m)} \sum_{r=1}^m \Pr(\text{rank}_t(i) = r) . \end{aligned}$$

By definition, the sum of probabilities for the top m ranks is the capture probability $p_t(m)$. Thus, $\text{Term 1} \geq \mu_{(m)} p_t(m)$.

Bounding Term 2: For any rank $r > m$, the drift is at least the minimum drift $\mu_{(N)}$.

$$\begin{aligned} \text{Term 2} &\geq \sum_{r=m+1}^N \mu_{(N)} \Pr(\text{rank}_t(i) = r) \\ &= \mu_{(N)} \sum_{r=m+1}^N \Pr(\text{rank}_t(i) = r) . \end{aligned}$$

The sum of probabilities for the remaining ranks is simply the complement of the capture probability, $1 - p_t(m)$. Thus, $\text{Term 2} \geq \mu_{(N)} (1 - p_t(m))$.

Combining: Adding the lower bounds for Term 1 and Term 2 yields the result:

$$\mathcal{R}_{\text{GP-BELIEF}}(t) \geq \mu_{(m)} p_t(m) + \mu_{(N)} (1 - p_t(m)) .$$

□

A.4. PROOFS OF SIGNAL ANCHORING PROPERTIES (SECTION 3.3)

A.4.1. PROOF OF PROPOSITION 3.6 (ENTROPY DEMOTION)

Proof. The binary entropy function $H(p) = -p \log p - (1-p) \log(1-p)$ is concave and maximized at $p = 0.5$. Consider the case where the true positive i has high belief $b_{t-1,i} > 0.5$, and there are $N - 1$ distractors. In sparse search ($k \ll N$), the vast majority of distractors are true negatives. However, early in the process, many distractors will have beliefs near the prior 0.5 (high entropy) simply due to noise or lack of observation.

Let $\mathcal{U} = \{j \neq i : H(b_{t-1,j}) > H(b_{t-1,i})\}$ be the set of documents with higher entropy than the true positive. GP-ENTROPY assigns positions by descending order of entropy. Thus, the rank of document i is $|\mathcal{U}| + 1$. As $b_{t-1,i} \rightarrow 1$, $H(b_{t-1,i}) \rightarrow 0$. Since distractor beliefs fluctuate around 0.5 (or decay to 0 slower than the winner rises to 1), the size of the set $|\mathcal{U}|$ approaches $N - 1$ (all other documents have more uncertainty than the almost-certain winner).

Consequently, the rank of i approaches N . Under the assignment policy, rank N corresponds to the position with drift $\mu_{(N)}$. Thus, $\mathcal{A}_{\text{GP-ENTROPY}}(t) \rightarrow \mu_{(N)}$. □

A.4.2. PROOF OF PROPOSITION 3.7 (BELIEF LOCK-IN)

Proof. Let i be the relevant document ($Z_i = 1$) and let $\mathcal{K} = \{k : k \neq i\}$ be the set of irrelevant distractors ($Z_k = 0$). At round τ , we are given that document i achieves the highest belief rank. In terms of log-odds, this implies that $\lambda_{\tau,i} > \max_{k \in \mathcal{K}} \lambda_{\tau,k}$.

Define the "margin" of the relevant document as the gap between its log-odds and that of its strongest competitor:

$$\Delta_t = \lambda_{t,i} - \max_{k \in \mathcal{K}} \lambda_{t,k} \quad (4)$$

We are given the initial condition $\Delta_\tau > 0$. We analyze the evolution of this gap for $t > \tau$ conditioned on the event that i retains rank 1.

1. **Drift of the relevant document:** Since i holds rank 1, the GP-BELIEF policy assigns it to the position with maximal diagnosticity. The expected increment for i is strictly positive:

$$\mathbb{E}[\lambda_{t,i} - \lambda_{t-1,i}] = \mu_{(1)} > 0 \quad (5)$$

2. **Drift of distractors:** Any distractor $k \in \mathcal{K}$ is assigned to a position $j = \sigma_t(k)$. Since $Z_k = 0$, the expected update is the negative KL divergence between the likelihoods (negative drift):

$$\begin{aligned} \mathbb{E}[\lambda_{t,k} - \lambda_{t-1,k}] &= -D_{KL}(P_{Z=0,j} || P_{Z=1,j}) \\ &\leq 0 \end{aligned}$$

Thus, the gap Δ_t evolves as a random walk with a strictly positive expected drift (since the relevant document accumulates positive signal while distractors accumulate negative or neutral signal).

From the theory of random walks with positive drift, starting at $\Delta_\tau > 0$, there exists a strictly positive probability that the walk never crosses zero (i.e., $\Delta_t > 0$ for all $t > \tau$). On this event, document i never loses the top rank. Consequently, it remains permanently assigned to the optimal position, maintaining the maximal information rate $\mathcal{R}_{\text{GP-BELIEF}}(t) = \mu_{(1)}$ for all subsequent rounds. \square

B. Calibration

GOLD PANNING requires position-wise detector parameters $(\text{TPR}_j, \text{FPR}_j)$ estimated once per model. Here we describe the calibration procedure and validate that simple fixed-grid sampling suffices.

Calibration procedure. For each dataset (MonoRel, PIR), we construct calibration instances to estimate position-wise detector parameters. Each instance is formed by selecting one sample as the known-relevant “gold” document and drawing $N - 1$ distractors from other samples in the same dataset. We place the gold at a target position j , fill remaining positions with distractors, and query the model once, yielding a set of cited document IDs. This induces binary observations for every document–position pair: a true positive when the gold is cited, a false positive when a distractor is cited, and analogous negatives otherwise.

For each position, we repeat this procedure 50 times with different gold documents and distractor sets. We average over 10 independent trials to obtain stable estimates of TPR_j and FPR_j . Diagnosticity $d_{\text{diag}}(j) = |\text{TPR}_j - \text{FPR}_j|$ is then computed from these estimates.

Conditional independence assumption. We adopt a conditional-independence approximation: given the true labels \mathbf{Z} and the assignment σ_t , citation outcomes are independent across documents within a round. In practice, citation events may exhibit within-round dependencies due to limited citation budgets or answer structure constraints. We therefore interpret $(\text{TPR}_j, \text{FPR}_j)$ as position-wise *marginal* detection rates under a fixed protocol. Under this interpretation, within-round dependencies are absorbed into the calibrated parameters rather than modeled explicitly. This approximation is conservative: if citations are negatively correlated (e.g., due to a citation budget), our independence assumption overestimates false positive rates, leading to more cautious belief updates.

C. Asymmetric Document-Position Cases

The main text presents the balanced case where the number of documents N equals the number of context positions. In practice, these quantities may differ: retrieval may return more documents than fit in context, or the context window may have unused capacity. We describe how to reduce asymmetric cases to the balanced setting via dummy items.

More documents than positions ($N_{\text{doc}} > N_{\text{pos}}$). When the retrieved set exceeds context capacity, we cannot place all documents in a single query. We introduce $N_{\text{doc}} - N_{\text{pos}}$ *dummy positions* with parameters $\text{TPR} = \text{FPR} = 0$. Documents assigned to dummy positions are excluded from the current query and receive no observation, leaving their beliefs unchanged. The greedy matching then operates over N_{doc} documents and N_{doc} positions (real + dummy). Under GP-BELIEF, the lowest-belief documents are assigned to dummy positions and effectively “wait” until higher-belief documents are resolved. Under GP-ENTROPY, the lowest-uncertainty documents wait instead.

This reduction preserves the algorithm’s convergence guarantee as long as every document is eventually placed at an informative real position. In practice, with T rounds and N_{pos} real positions per round, we can test up to $T \cdot N_{\text{pos}}$ document-position pairs, which suffices for convergence when T is moderate.

More positions than documents ($N_{\text{pos}} > N_{\text{doc}}$). When context capacity exceeds the document set, some positions remain unfilled. We introduce $N_{\text{pos}} - N_{\text{doc}}$ *dummy documents* with fixed belief $b = 0$ (known irrelevant). These are assigned to the lowest-diagnosticity positions, ensuring that real documents occupy the most informative slots. Dummy documents do not participate in belief updates; their presence simply pads the assignment to achieve a balanced matching.

Implementation note. In our experiments, we use $N = N_{\text{doc}} = N_{\text{pos}} = 100$ throughout, so these reductions are

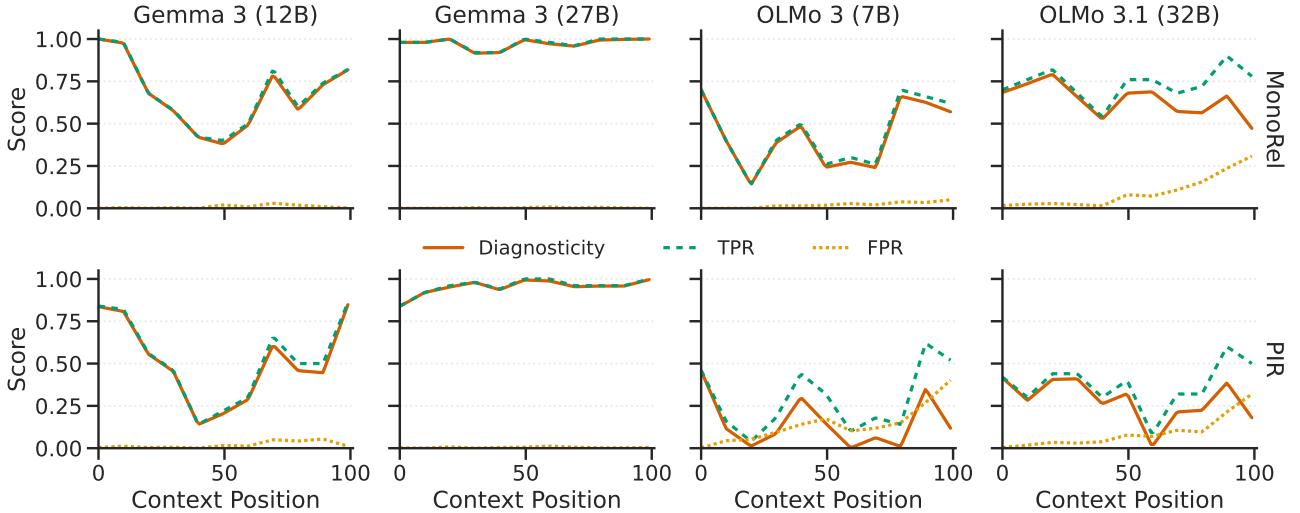


Figure 6. Calibrated position profiles across models and tasks ($N = 100$). Each panel shows TPR (solid), FPR (dashed), and diagnosticity $|TPR - FPR|$ (dotted) by context position. Gemma-3-12B and OLMo models exhibit pronounced primacy and recency effects, while Gemma-3-27B shows a flat profile with minimal position bias, explaining why strategic reordering offers no advantage for this model.

not exercised. However, the framework naturally extends to production settings where retrieval set sizes vary across queries. The key insight is that dummy items with degenerate parameters (TPR = FPR for positions, $b = 0$ or $b = 1$ for documents) do not affect the information-theoretic properties of the algorithm.

D. Profile Stability Across Tasks

§4.4 demonstrated that GOLD PANNING remains effective under cross-calibration, suggesting that diagnosticity profiles are structural properties of the model rather than task-specific artifacts. Here we quantify this stability directly by measuring the rank correlation between profiles calibrated on MonoRel versus PIR.

Table 2. Spearman’s ρ and Kendall’s τ between diagnosticity profiles calibrated on MonoRel vs. PIR. Red entries are not statistically significant ($p > 0.05$); these correspond to degenerate cases where the model exhibits minimal position bias.

Model	ρ	τ
OLMo-3-7B-Think	0.595	0.455
OLMo-3.1-32B-Think	0.488	0.315
Gemma-3-12B-IT	0.959	0.842
Gemma-3-27B-IT	0.102	0.112

Tab. 2 reports Spearman’s ρ and Kendall’s τ for each model at $N = 100$ positions. Models with exploitable position bias show strong rank correlations ($\rho > 0.49$), indicating that the relative ordering of high- vs. low-diagnosticity positions is preserved across tasks. This stability explains why cross-calibration incurs only modest performance degradation

(Fig. 4): the algorithm primarily relies on *which* positions are reliable, not their precise TPR/FPR values.