# Algorithmic Thinking Theory

MohammadHossein Bateni
Google
bateni@google.com

Vincent Cohen-Addad
Google
cohenaddad@google.com

Yuzhou Gu
NYU
yuzhougu@nyu.edu

Silvio Lattanzi
Google
silviol@google.com

Simon Meierhans[*]
ETH Zurich
mesimon@inf.ethz.ch

Christopher Mohri[†]
Stanford, Google
xmohri@stanford.edu

## Abstract

Large language models (LLMs) have proven to be highly effective for solving complex reasoning tasks. Surprisingly, their capabilities can often be improved by iterating on previously generated solutions. In this context, a reasoning plan for generating and combining a set of solutions can be thought of as an algorithm for reasoning using a probabilistic oracle.

We introduce a theoretical framework for analyzing such reasoning algorithms. This framework formalizes the principles underlying popular techniques for iterative improvement and answer aggregation, providing a foundation for designing a new generation of more powerful reasoning methods.

Unlike approaches for understanding models that rely on architectural specifics, our model is grounded in experimental evidence. As a result, it offers a general perspective that may extend to a wide range of current and future reasoning oracles.

# 1 Introduction

The rapid advancement of Large Language Models (LLMs) has marked a paradigm shift in artificial intelligence, with models demonstrating superhuman performance on a wide array of language benchmarks and impressive capabilities on complex reasoning tasks [AVL+24, YSH+25]. Initial challenges, such as grade-school mathematics (GSM8K) and standard competition math (MATH dataset), have largely been surmounted, pushing the frontier of AI reasoning toward "grand challenge" problems, such as those found in the International Mathematical Olympiad (IMO).

These problems, renowned for their demand for deep insight, creativity, and rigorous proof, expose a fascinating weakness in modern LLMs. While a model's performance on a single attempt (termed $pass@1$) may be very low, its ability to produce a correct answer within $k$ attempts ($pass@k$) can be significantly higher. This $pass@1$ versus $pass@k$ gap, especially pronounced when sampling with high temperature to produce diverse outputs, suggests that models possess a vast, latent capability that is not accessible in a single, high-confidence generation.

Interestingly, to recover the full power of the model it is not sufficient to simply use multiple attempts. In fact, even the $pass@k$ metric fails to capture the full story. On the most difficult problems, simply sampling $k$ times and selecting the best answer (e.g., "best-of-32") still yields poor results. For instance, Huang and Yang (2025) report that a best-of-32 baseline on the IMO 2025 problems achieved an accuracy of only 31.6–38.1% for leading models [HY25]. This paradox lies at the heart of our work: the latent capability of LLMs is not merely a matter of *selection* (finding one correct needle in a haystack of $k$ attempts), but one of *synthesis*. While established techniques like Self-Consistency [WWS+23] leverage majority voting to improve reliability, they remain bound by the quality of individual samples. The "deep" reasoning ability of the model appears to be distributed across multiple, diverse, and often individually flawed chains of thought [WWS+22].

This shift highlights a growing focus on scaling inference-time compute—leveraging more computational resources *after* training to boost performance on hard tasks. These algorithms effectively simulate a "System 2" reasoning process atop the base model's "System 1" generations [Kah11], iteratively refining and synthesizing information to reach conclusions that are unreachable in a single forward pass. While empirical results show this works, we lack a formal theory governing the efficient allocation of this test-time budget. Existing theoretical work often focuses on the expressivity of standard Transformer forward passes, leaving the dynamics of these iterative, multi-call reasoning systems unexplored.

While the theoretical understanding of LLMs reasoning is still limited, the hypothesis that the ability to synthesize is a key ingredient in LLMs success is strongly backed by recent empirical breakthroughs. Work such as *Reflexion* [SCG+23] demonstrated the value of verbal reinforcement for iterative improvement, while Huang and Yang (2025) showed that a similar model-agnostic, multi-stage "verification and refinement pipeline" can achieve a stunning 85.7% accuracy (5 out of 6 problems) on the same IMO 2025 dataset [HY25]. Explicitly showing that a complex reasoning procedure—which iteratively generates, improves, verifies, and refines a solution—is able to harness the model's latent abilities to achieve a result far beyond what simple sampling could. Earlier structured approaches like *Tree of Thoughts* [YYZ+23] attempted to harness latent reasoning by framing it as a search over intermediate steps. While effective for moderate complexity, grand challenge problems have required even more involved, iterative pipelines.

A similar, complementary approach is proposed by Venkatraman et al. (2025) with their "Recursive Self-Aggregation" (RSA) algorithm [VJM+25]. Inspired by evolutionary methods, RSA

maintains a *population* of candidate solutions. In each step, it refines this population by "aggregating" subsets of solutions, prompting the model to combine their useful ideas and produce a new, improved generation. RSA is shown to "enable bootstrapping from partially correct intermediate steps within different chains of thought" [VJM$^+$25]. Again, this demonstrates a mechanism for *synthesis*, not selection.

These empirical successes present a clear theoretical gap. We have powerful, practical examples of "reasoning algorithms"—like the IMO pipeline or RSA—that unlock performance far exceeding a model's baseline. Yet, we lack a formal theory to understand *why* they work and *how* to design them. What properties of an LLM (as a "reasoning oracle") make aggregation and refinement effective? How do we model the trade-off between parallel branching (like RSA) and sequential depth (like the IMO pipeline)?

This paper initiates the theoretical study of the "algorithmic thinking" procedures. We formalize the components of these complex reasoning processes. We introduce the concept of a **reasoning oracle**, $\mathcal{A}$, which takes a *context* of previously generated solutions $C \subseteq \mathcal{S}$ and produces a new solution $s \in \mathcal{S}$. The probability of success is governed by a **transfer function**, $\mathcal{F}$, which models the quality of the output based on the quality of the solutions in the context.

This formalism allows us to precisely captures the empirical techniques described so far. For example, the classic baseline *pass*@1 performance is the success probability with an empty context, $\mathcal{F}(\emptyset)$, the *pass*@1 vs. *pass*@$k$ gap is captured by $\mathcal{F}(C) > \mathcal{F}(\emptyset)$ when $C$ contains $k$ solutions generated by $\mathcal{F}(\emptyset)$.

Finally, more complex techniques like Recursive Self-Aggregation [VJM$^+$25] can be described by an algorithm that iteratively applies $\mathcal{F}$ to contexts $C$ with $k > 1$ solutions.

Within this framework, we can now move from empirical observation to theoretical analysis. In particular, we first formalize the strategies seen in practice by defining the `Branching Algorithm` (Algorithm 1), which models tree-like synthesis, and the more efficient `Genetic Algorithm` (Algorithm 2), which re-uses solutions from a previous population, directly connecting to the evolutionary approach of RSA [VJM$^+$25]. We also study the `Random Sampling Algorithm` (Algorithm 3), a simple algorithm that reuses a random subset of previous solutions in each step.

Our central analysis hinges on a key property of the oracle: *monotonicity* (Definition 4.1), the formal assumption that a "better" context of solutions leads to a better output. We then introduce and study the class of `Decaying Models` (Definition 2.1), establishing results on algorithm behavior on such models. Then we will consider several special cases, including the `Uniform Model` (Definition 2.2) and two models where the oracle performance degrades as context size increases (Definitions 2.3 and 2.4). The former is the simplest non-trivial model under our framework, which already shows interesting behavior; the latter is closer to real-world behavior, capturing the experimental observation that performance can degrade as context size becomes too large.

## 1.1 Our Results

In Section 4, we characterize the *limits* of these algorithms and show that they obtain the maximum achievable success probability for Decaying Models (Proposition 4.6). We also analyze the *efficiency* of the algorithm by establishing the convergence rate of the algorithm to optimal success probability (Section 5). Then we make further discussions on the Uniform Model and specific Decaying Models in Section 6 and Section 7, respectively. Our goal is to move beyond empirical successes and develop a rigorous theory for designing and analyzing the next generation of reasoning algorithms.

In Appendix A, we provide experimental evidence to ground our modeling choices.

## 1.2 Technical Overview

We model a reasoning oracle as a probabilistic function $\mathcal{A}$ that samples from a solution space $\mathcal{S}$. The function can either be called with empty context $C = \emptyset$, or it can be provided with some previous solutions sampled from the solution space. If it is called with non-empty context, the properties of the provided solutions alter the distribution from which the solutions are sampled.

In this article, we focus on the following simple setting. Each solution is either correct or wrong, and the distribution from which $\mathcal{A}$ samples only depends on a) if there is a correct solution in the provided context and b) the number of solutions in the context. This allows us to capture two important properties observed in practice. Providing correct solutions helps, and hiding a correct solution alongside a lot of wrong solutions decreases its usefulness.

**Models.** Our framework leaves significant flexibility for modeling. We focus our initial study on the class of Decaying Models, where the success probability depends only on whether there is a correct solution in the context $C$, and the context size $|C|$.

We derive the maximum possible success probability one can boost to for such models, and derive the convergence rate to this success probability for the algorithms we consider.

Then we study the following two special cases of the Decaying Model.

1. The simplest possible model succeeds with probability $q > p$ whenever there is a correct solution in the context $C$ and $|C| \leq k$ (See Section 6).

2. To better capture the experiments, we replace the fixed probability $q$ with some function $f$ that decays with the context length $|C|$. In this article, we study exponential and polynomial decay functions (See Section 7).

We give efficient algorithms for boosting the success probabilities in terms of the number of oracle calls. In many regimes, these are small multiples of $1/p$, the number of oracle calls needed for obtaining a single correct solution among all the solutions contained so far.

It may be unreasonable to hope to solve a very difficult problem without providing any context. We therefore propose various ideas for generalizing our model to partial solutions in Section 8, and hope that such a model can be both motivated by experiments and analyzed in the future.

## 1.3 Other Theoretical Models for LLMs

Early theoretical work focused on the expressivity of the architecture, proving that Transformers are universal approximators of sequence-to-sequence functions, meaning they can theoretically model any such function as then number of parameters approaches infinity [YBR+20]. Since then, there have been various proposals for understanding the experimentally observed benefit of providing context.

A leading theoretical approach explains in-context learning (ICL) as an algorithmic mechanism called "Induction Heads." These are 2-layer circuits that allow the model to copy and complete patterns from the context [OEN+22]. Another major theoretical angle formalizes ICL as an algorithm learning problem, where the Transformer implicitly constructs a hypothesis function at inference time [LIPO23]. Theoretical work has shown that Transformers trained by gradient flow can provably learn classes of linear models in-context [GTLV22].

Recent work has moved beyond showing CoT helps to proving it is theoretically required for certain classes of problems. A recent, pivotal study proved rigorous separation results, showing

that bounded-depth Transformers cannot directly solve basic arithmetic or linear equations unless their size grows super-polynomially. However, with CoT (autoregressive intermediate steps), they can solve these, effectively allowing them to handle problems up to P-complete (like dynamic programming) [FZG$^+$24]. Newer theoretical frameworks have derived scaling laws for an "optimal" CoT length, proving that while CoT is necessary, excessively long chains are susceptible to noise and error accumulation ("overthinking") [ZLM$^+$25]. In our model, we explain this behavior as the context becoming increasingly correlated and therefore providing less value.

Another line of complexity theoretic work established that both Transformers and SSMs generally fall into the TC0 computational complexity class [SVH24]. Despite being in the same broad class, they have distinct theoretical strengths. Transformers are provably better at state tracking and copying from long history (due to their ability to attend to any previous state), while SSMs excels at different types of structured state tracking without needing a full history cache.

Finally, recent work models grokking as a phase transition where the network shifts from a high-complexity "memorization" phase to a low-complexity "generalization" phase, effectively compressing the data [KN24]. More recent studies have begun linking this strictly to internal mechanisms, such as the evolution of routing pathways in Mixture of Experts (MoE) models during pretraining [LFZ25].

# 2   Preliminaries

**Notations.**   We let $\log(\cdot)$ refer to the natural logarithm and $[n] = \{1, \ldots, n\}$. $\mathbb{N} = \mathbb{Z}_{\geq 0}$ denotes the set of non-negative integers.

**Question.**   Throughout the article, we fix a question $Q$ which we seek to answer.

**Solution space and quality.**   Let $\mathcal{S}$ be the space of possible solutions of question $Q$ and $\mathcal{V} \subseteq \mathbb{R}$ denote the space of possible quality scores. We associate with $\mathcal{S}$ a scoring function score $: \mathcal{S} \to \mathcal{V}$. For every solution $s \in \mathcal{S}$, the value score$(s)$ assesses to what extent solution $s$ answers question $Q$. A reasoning algorithm does not have access to the function score. In this article, we will focus on $\mathcal{V} = \{0, 1\}$, i.e., the scoring function where every solution has a binary score.

**Reasoning oracle.**   To produce answers to question $Q$, we will exclusively use an oracle $\mathcal{A}$. This oracle is a randomized algorithm that given context $C \subseteq \mathcal{S}$ samples a solution $s \in \mathcal{S}$. The distribution from which the oracle samples depends only on the quality of the solutions in the context $C$.

We assume that distribution of the score of the output of the oracle depends only on the multiset of scores in the context $C$, in the following way. Let $\mathcal{M}(\mathcal{V})$ denote the set of all finite multisets with elements in $\mathcal{V}$, and $\mathcal{P}(\mathcal{V})$ denote the set of distributions over $\mathcal{V}$. We assume there exists a transfer function $\mathcal{F} : \mathcal{M}(\mathcal{V}) \to \mathcal{P}(\mathcal{V})$ such that the distribution of the score of the output of the oracle on context $C$ is given by $\mathcal{F}(\{\text{score}(s) : s \in C\}_{\text{multiset}})$. In particular, $\mathcal{F}(\emptyset)$ denotes the distribution of scores when the oracle is called with an empty context.

**Reasoning algorithm.**   An $(\mathcal{A}, n)$-reasoning algorithm is a procedure that calls the oracle $\mathcal{A}$ at most $n$ times to output a solution $s$. The goal is to maximize expected score $\mathbb{E}[\text{score}(s)]$ of the final output $s$. If $\mathcal{V} = \{0, 1\}$, then we say a reasoning algorithm succeeds if score$(s) = 1$.

In general, a reasoning algorithm has the form $(S_1, \ldots, S_n)$, where $S_k \subseteq [k-1]$ indicates the previous outputs used as context for the $k$-th oracle call. That is, for $k = 1, \ldots, n$, the algorithm calls the oracle $\mathcal{A}$ with context $\{s_i : i \in S_k\}_{\text{multiset}}$ to obtain $s_k$.

We can represent a reasoning algorithm as a DAG with $n$ nodes, where each node corresponds to an oracle call, and there is a directed edge from node $i$ to node $j$ if $i \in S_j$. The $n$-th node corresponds to the final output of the algorithm. We define the *depth* of an algorithm as the length (number of edges) of the longest paths in its DAG representation. The depth captures the inference time of the reasoning algorithm, assuming infinite parallelism.

**Answer population.** If the image of the score function is $\{0, 1\}$, then we call a subset $S \subset \mathcal{S}$ a $p$-correct answer population if at least a $p$ fraction of the elements in $S$ have score 1. For the rest of the article, we let $\mathcal{V} = \{0, 1\}$.[1]

**Decaying Model.** Most of our results focus on a class of models called the decaying model, defined as follows.

**Definition 2.1** (Decaying Model). *Let $\mathcal{V} = \{0, 1\}$. Let $f : \mathbb{N} \to [0, 1]$ and $g : \mathbb{N} \to [0, 1]$ be two functions satisfying $f(k) \geq g(k)$ for all $k \in \mathbb{N}$. The decaying model $\mathcal{A}^{(f,g)}$ is defined as the oracle satisfying for any context $C \subseteq \mathcal{S}$:*

$$\Pr[\text{score}(\mathcal{A}^{(f,g)}(C)) = 1] = \begin{cases} f(|C|), & \text{if } \sum_{s \in C} \text{score}(s) \geq 1, \\ g(|C|), & \text{if } \sum_{s \in C} \text{score}(s) = 0. \end{cases}$$

We note that the functions $f$ and $g$ are typically monotonically decreasing, which motivates the nomenclature. For this decaying model, we are able to give a precise characterization of the optimal success probability achievable by reasoning algorithms (see Section 4).

Next we will introduce several special cases of the decaying model.

**Uniform Model.** The simplest model is the following.

**Definition 2.2** (Uniform Model). *Let $0 \leq p \leq q \leq 1$ and integer $k \geq 1$. For the uniform model $\mathcal{A}_u^{(p,q,k)}$, each call has $\text{score}(\mathcal{A}_u^{(p,q,k)}(C)) = 1$ independently with probability*

$$p + (q - p)\mathbb{1}\left\{\sum_{s \in C} \text{score}(s) \geq 1 \text{ and } |C| \leq k\right\}$$

*and $\text{score}(\mathcal{A}_u^{(p,q,k)}) = 0$ otherwise.*

The uniform model is the special case of the decaying model where $f(k) = q$ and $g(k) = p$ for all $k \in \mathbb{N}$.

If we remove the constraint $|C| \leq k$ in the definition, then the optimal strategy is to always use all previous answers as context. We therefore limit the context size $|C|$ to be at most $k$ in the definition of the uniform model.

---

[1]Some of the general results presented in Section 4 apply to general sets $\mathcal{V}$.

**Exponentially and Polynomially Decaying Models.** When implementing the oracle with a large language model, we experimentally observe a slow decay in model accuracy as the context size increases. We therefore study models that capture this decay more explicitly.

We are particularly interested in two natural decay functions, exponential decay and polynomial decay.

**Definition 2.3** (Exponential Decay). *Let $0 < p \leq q \leq 1$. The exponential decaying model is the decaying model $\mathcal{A}^{(f,g)}$ where $f(k) = q^{k-1}$ and $g(k) = p\mathbb{1}_{\{k=0\}}$.*

**Definition 2.4** (Polynomial Decay). *Let $0 < p \leq q \leq 1$. The polynomial decaying model is the decaying model $\mathcal{A}^{(f,g)}$ where $f(k) = \frac{1}{k^q}$ and $g(k) = p\mathbb{1}_{\{k=0\}}$.*

# 3 Reasoning Algorithms

In this section, we introduce the reasoning algorithms we study in this article:

1. Branching Algorithm (Algorithm 1): A simple algorithm that merges independently generated solutions in a tree-like fashion. It achieves optimal success probability for a broad class of oracles.

2. Genetic Algorithm (Algorithm 2): A more efficient version of the branching algorithm that re-uses solutions to reduce the number of oracle calls. It approaches the branching algorithm as the population sizes grow.

3. Random Sampling Algorithm (Algorithm 3): An algorithm that generates new solutions by randomly sampling from all previously generated solutions. It also achieves the optimal success probability, and has better convergence rate in certain settings.

**Branching algorithm.** The branching algorithm assembles a solution by initially generating a set of $pass@1$ solutions without providing any context. We call such solutions level 0 solutions. Then, it builds a tree obtaining a set of level $i$ solutions via combining groups of $k_i$ solutions from level $i-1$ until it creates a single level $L$ solution. See Algorithm 1 for pseudo-code of this algorithm.

---

**Algorithm 1:** BRANCHINGALGORITHM($\mathcal{A}, L, (k_1, \ldots, k_L)$))

**1 if** $L = 0$ **then**
**2** $\quad$ **return** $\mathcal{A}(\emptyset)$
**3** $C \leftarrow \emptyset$
**4 for** $j \in 1, \ldots, k_L$ **do**
**5** $\quad$ $C \leftarrow C \cup \{\text{BRANCHINGALGORITHM}(\mathcal{A}, L - 1, (k_1, \ldots, k_{L-1}))\}$
**6 return** $\mathcal{A}(C)$

---

From this construction, we observe that all solutions generated by this algorithm are completely independent, and that therefore the solutions at level $i$ also have to have the same guarantees. We conclude that all the solutions at level $i$ have the same success probability, and that any reasonable algorithm increases the success probability of the solutions over time (See Section 4).

**Genetic Algorithm.** While the branching algorithm achieves the best possible success probability for a broad class of oracles (Proposition 4.6), it exponentially grows the number of solutions with the depth $L$. This quickly blows up the necessary computational resources. Therefore, genetic algorithms are often used in practice [VJM+25]. These essentially follow the template from Algorithm 1, but re-use solutions from level $i-1$ multiple times to avoid branching as much. Concretely, they have a fixed population size for layer $i$, and each solution at layer $i > 0$ is created by sampling $k_i$ solutions from layer $i - 1$ uniformly at random with replacement. See Algorithm 2 for pseudo code of the algorithm and note that the population sizes $(s_1, \dots, s_L)$ are an additional input. We observe that this algorithm approaches the branching algorithm as the population sizes go towards infinity.

**Observation 3.1.** *The genetic algorithm (Algorithm 2) approaches the branching algorithm (Algorithm 1) when $s_1, \dots, s_L \to \infty$ and $\frac{s_i}{s_{i+1}} \to \infty$ for $i = 1, \dots L - 1$.*

---

**Algorithm 2:** GENETICALGORITHM$(\mathcal{A}, L, (s_1, \dots, s_L), (k_1, \dots, k_L))$

---

**1** $S_0 \leftarrow \emptyset$
**2 for** $j = 1, \dots, s_1$ **do**
**3**  $\quad$ $S_0 \leftarrow S_0 \cup \{\mathcal{A}(\emptyset)\}$
**4 for** $i = 1, \dots, L - 1$ **do**
**5**  $\quad$ $S_i \leftarrow \emptyset$
**6**  $\quad$ **for** $j = 1, \dots, s_{i+1}$ **do**
**7**  $\quad\quad$ Let $C$ consist of $k_i$ i.i.d. uniformly sampled items from $S_{i-1}$
**8**  $\quad\quad$ $S_i \leftarrow S_i \cup \{\mathcal{A}(C)\}$
**9** Let $C$ consist of $k_L$ i.i.d. uniformly sampled items from $S_{L-1}$
**10 return** $\mathcal{A}(C)$

---

In the following, we show that any branching algorithm can be turned into a genetic algorithm given access to a slightly stronger oracle in the case where every solution is either wrong or correct (i.e., $\mathcal{V} = \{0, 1\}$).

Recall the following standard Chernoff bound.

**Theorem 3.2** (Chernoff bound, see e.g. [MU05]). *Suppose that $X_1, \dots, X_n$ are i.i.d random variables taking values in $\{0, 1\}$. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Then*

$$\Pr[X \le (1 - \epsilon)\mu] \le e^{-\epsilon^2 \mu/2}.$$

**Definition 3.3.** *We say oracle $\mathcal{A}'(\cdot)$ $\alpha$-dominates oracle $\mathcal{A}(\cdot)$ if the success probability of $\mathcal{A}'(\cdot)$ is at least $\alpha$ times the success probability of oracle $\mathcal{A}$ for some $\alpha$ for all inputs to $\mathcal{A}$.*

**Lemma 3.4.** *Suppose an oracle $\mathcal{A}(\cdot)$, a parameter $L$ and a vector $(k_1, \dots, k_L)$ are given such that BRANCHINGALGORITHM$(\mathcal{A}, i, (k_1, \dots, k_i))$ returns a solution $s$ that is correct with probability $p_i$ for each $1 \le i \le L$. Then, given an oracle $\mathcal{A}'$ that $\frac{1}{1-\epsilon}$ dominates the oracle $\mathcal{A}$, GENETICALGORITHM$(\mathcal{A}', L, (s_1, \dots, s_L), (k_1, \dots, k_L))$ returns a solution $s'$ that is correct with probability at least $p_L$, provided that $s_i \ge -2\log(\epsilon/L)/(p_{i-1}\epsilon^2)$ for each $1 \le i \le L$.*

8

*Proof.* We prove by induction that for each $0 \leq i \leq L$, the set $S_i$ is a $p_{i-1}$-correct answer population with probability at least $1 - i \cdot \epsilon/L$.

For $S_0$, we sample directly from an oracle which guarantees that the solutions are correct with probability $\frac{1}{1-\epsilon}p_0$. Therefore, $s_1 = -2\log(\epsilon/L)/(p_0\epsilon^2)$ samples suffices with probability $1 - \epsilon/L$ by Theorem 3.2. This concludes the base case.

Then, in the step case, we have that if $S_{i-1}$ is a $p_{i-2}$-correct answer population with probability $1 - (i-1)\epsilon/L$ by induction. We assume that it is in fact a $p_{i-2}$-correct answer population, and then obtain that $S_i$ is a $p_{i-1}$-correct answer population with probability $1 - \epsilon/L$ by Theorem 3.2. The step case follows by union bound.

From this, we obtain that $S_L$ is a $p_{L-1}$-correct answer population with probability $1 - \epsilon$. If it is, then the output is correct with probability $\frac{1}{1-\epsilon}p_L$. The lemma follows. $\qquad\square$

**Random Sampling Algorithm.** The genetic algorithm samples solutions from the previous layer only. An alternative is to sample solutions from all previously generated solutions. This is the idea behind the random sampling algorithm, whose pseudo-code is given in Algorithm 3.

---

**Algorithm 3:** RANDOMSAMPLINGALGORITHM($\mathcal{A}, n, k$)

---

**1** $s_1 \leftarrow \mathcal{A}(\emptyset)$
**2** **for** $j = 2, \ldots, n$ **do**
**3** $\quad$ Let $C$ consist of $k$ i.i.d. uniformly sampled items from $\{s_1, \ldots, s_{j-1}\}$
**4** $\quad$ $s_j \leftarrow \mathcal{A}(C)$
**5** **return** $s_n$

---

We will see in Section 4 that this algorithm also achieves optimal success probability for decaying models.

# 4 Optimal Success Probability

In this section, we study the optimal success probability achievable by reasoning algorithms (without restrictions on the number of calls). We show that for decaying models, the algorithms introduced in Section 3 achieve the optimal success probability as the number of calls goes to $\infty$.

## 4.1 Monotonicity

In this section we establish some general properties of reasoning oracles and algorithms. We study two versions of monotonicity, and show that they have interesting consequences to the structure of an optimal algorithm.

**Definition 4.1** (Monotonicity). *An oracle $\mathcal{A}$ is called*

- weakly monotone *if for any two multisets $A = \{a_1, \ldots, a_k\}$, $B = \{b_1, \ldots, b_k\}$ satisfying $a_i \leq b_i$ for all $i \in [k]$, we have $\mathcal{F}(A) \preceq \mathcal{F}(B)$, where $\preceq$ denotes (first order) stochastic dominance (i.e., there is a monotone coupling between $\mathcal{F}(A)$ and $\mathcal{F}(B)$);*

- strongly monotone *if it is weakly monotone and for any two multisets $A$ and $B$ satisfying $A \subseteq B$ (as multisets), we have $\mathcal{F}(A) \preceq \mathcal{F}(B)$.*

The following result shows that weak monotonicity implies an FKG inequality for the joint distribution of scores of solutions [FGK71].

**Lemma 4.2.** *If an oracle $\mathcal{A}$ is weakly monotone, then for any $(\mathcal{A}, n)$-reasoning algorithm $(S_1, \ldots, S_n)$ and any functions $f, g : \mathcal{V}^n \to \mathbb{R}$ that are both non-decreasing (or both non-increasing), we have*

$$\mathbb{E}[fg] \geq \mathbb{E}[f] \mathbb{E}[g].$$

*Proof.* Let $v_i = \text{score}(s_i)$ for $i \in [n]$. Let $(X_1, \ldots, X_n)$ be i.i.d. samples from $\text{Unif}([0, 1])$. For every $j \in [n]$ and every possible combination of values $(v_i : i \in S_j)$, we can couple $\text{Unif}([0, 1])$ with $P_{v_j | (v_i : i \in S_j)}$ in an order-preserving way. In this way, we can construct an order-preserving coupling between $(v_1, \ldots, v_n)$ and $(X_1, \ldots, X_n)$. So $f, g$ can be lifted to non-decreasing (or non-increasing) functions $\widetilde{f}, \widetilde{g} : [0, 1]^n \to \mathbb{R}$ such that $f(v_1, \ldots, v_n) = \widetilde{f}(X_1, \ldots, X_n)$ and $g(v_1, \ldots, v_n) = \widetilde{g}(X_1, \ldots, X_n)$. The FKG inequality [FGK71] implies that

$$\mathbb{E}[fg] = \mathbb{E}\left[\widetilde{f}\widetilde{g}\right] \geq \mathbb{E}\left[\widetilde{f}\right] \mathbb{E}[\widetilde{g}] = \mathbb{E}[f] \mathbb{E}[g].$$

$\square$

Strong monotonicity, in addition to the above, implies that an optimal algorithm must use all available context size.

**Lemma 4.3.** *If an oracle $\mathcal{A}$ is strongly monotone, then for any $(\mathcal{A}, n)$-reasoning algorithm, adding more solutions to the context of any oracle call never decreases the expected score of the final output.*

*Proof.* For two $(\mathcal{A}, n)$-reasoning algorithms operating on contexts $(S_1, \ldots, S_n)$ and $(S_1', \ldots, S_n')$ respectively satisfying $S_i \subseteq S_i'$, we can construct a monotone coupling step by step between the distributions of $(\text{score}(s_1), \ldots, \text{score}(s_n))$ and $(\text{score}(s_1'), \ldots, \text{score}(s_n'))$. So adding more solutions to the context never hurts. $\square$

In particular, if there are no limits on context size, then the optimal algorithm is to take $S_i$ to contain all previously generated solutions for all $i$.

For the binary quality model considered in this article (i.e., $\mathcal{V} = \{0, 1\}$), the transfer function $\mathcal{F}$ can be compactly represented by a function $F : \mathbb{N} \times \mathbb{N} \to [0, 1]$, where $F(a, b)$ is the probability that the oracle outputs a correct solution when the context contains $a$ correct solutions and $b$ incorrect solutions. With minor abuse of notation, we call $F$ the transfer function as well.

A binary quality model is weakly monotone if and only if $F$ satisfies $F(a + 1, b) \geq F(a, b + 1)$ for all $a, b \geq 0$. A binary quality model is strongly monotone if and only if $F$ satisfies $F(a + 1, b) \geq F(a, b + 1) \geq F(a, b)$ for all $a, b \geq 0$. So we have the following corollary.

**Corollary 4.4** (Monotonicity for binary quality models)**.** *Consider a binary quality model with transfer function $F$.*

*If $F(a + 1, b) \geq F(a, b + 1)$ for all $a, b \geq 0$, then for any $(\mathcal{A}, n)$-reasoning algorithm $(S_1, \ldots, S_n)$ and any functions $f, g : \{0, 1\}^n \to \mathbb{R}$ that are both non-decreasing (or both non-increasing), we have $\mathbb{E}[fg] \geq \mathbb{E}[f] \mathbb{E}[g]$.*

*If $F(a + 1, b) \geq F(a, b + 1) \geq F(a, b)$ for all $a, b \geq 0$, then for any $(\mathcal{A}, n)$-reasoning algorithm, adding more solutions to the context of any oracle call never decreases the expected score of the final output.*

*Proof.* Directly follows from Lemma 4.2 and Lemma 4.3. $\square$

## 4.2 Optimality of the branching algorithm

In the following, we show that for decaying models, the branching algorithm achieves the maximum achievable success probability.

The following lemma discusses solutions to a class of polynomial equations which arise in this study.

**Lemma 4.5.** *Let $p, q \in [0, 1]$ and $k \in \mathbb{N}$. Consider the equation $x = q - (1-x)^k(q-p)$ and its solutions in $[0, 1]$.*

- *When $k = 0$, the equation has a unique solution $x = p$.*

- *When $k = 1$ and $(p, q) \neq (0, 1)$, the equation has a unique solution $x = \frac{p}{1-q+p}$. If $k = 1$ and $(p, q) = (0, 1)$, every $x \in [0, 1]$ is a solution.*

- *When $k \geq 2$ and $p \geq q$, the equation has a unique solution in $[0, 1]$. When $k \geq 2$ and $0 < p < q$, the equation has a unique solution in $[0, 1]$. When $k \geq 2$, $p = 0$, and $kq \leq 1$, the equation has a unique solution $0$ in $[0, 1]$. When $k \geq 2$, $p = 0$, and $kq > 1$, the equation has two solutions in $[0, 1]$, one of which is $0$ and the other is in $(0, 1]$.*

*Proof.* Let $f(x) = x - q + (1-x)^k(q-p)$. Then $f(0) = -p \leq 0$, and $f(1) = 1 - q \geq 0$. The $k = 0, 1$ cases follow directly. We therefore assume $k \geq 2$.

Suppose $p \geq q$. Then $f$ is strictly increasing in $[0, 1]$, and there is a unique root in $[0, 1]$.

Suppose $p < q$. Then $f$ is strictly convex in $[0, 1]$. When $p > 0$, there is a unique root in $[0, 1]$. When $p = 0$, $0$ is a root, and $f'(0) = 1 - kq$. So when $kq \leq 1$, there is no root in $(0, 1]$; when $kq > 1$, there is a unique root in $(0, 1]$. $\qquad\square$

A decaying model is a binary quality model whose transfer function $F$ satisfies $F(a+1, b) = F(a, b+1)$ for all $a \geq 1, b \geq 0$. We have $f(k) = F(k, 0)$ and $g(k) = F(0, k)$ for all $k \in \mathbb{N}$.

**Proposition 4.6** (Optimality of branching algorithm). *Consider a decaying model.*

*Let $x_k^*$ be the largest solution in $[0, 1]$ to the equation $x = f(k) - (1-x)^k(f(k) - g(k))$ (c.f. Lemma 4.5).*

*Then the success probability of any $(\mathcal{A}, n)$-reasoning algorithm is at most $x^* = \sup_{k \geq 0} x_k^*$.*

*Furthermore, if $g(0) > 0$, then for any $\epsilon > 0$, a branching algorithm (Algorithm 1) achieves success probability at least $x^* - \epsilon$.*

*Proof.* We prove lower and upper bounds separately.

Upper bound. Let $(S_1, \ldots, S_n)$ be an $(\mathcal{A}, n)$-reasoning algorithm. Let $v_i = \text{score}(s_i)$ for $i \in [n]$. We prove by induction that $\Pr[v_i = 1] \leq x^*$ for all $i \in [n]$.

Fix any $m \in [n]$. Let $k_m = |S_m|$. By the induction hypothesis, for each $i \in S_m$, we have $\Pr[v_i = 0] \geq 1 - x^*$. By Corollary 4.4, $\Pr[v_i = 0 \text{ for all } j \in S_i] \geq (1-x^*)^{k_m}$.

Then $\Pr[v_i] \leq f(k_m) - (1-x^*)^{k_m}(f(k_m) - g(k_m)) \leq x^*$, where last step follows from the definition of $x^*$.

Lower bound. Choose $k$ such that $x_k^* > x^* - \epsilon$. Consider the branching algorithm Algorithm 1 with fan-in $k_i = k$ for all $i \in 1, \ldots, L$. Since all the generated solutions at a particular level are independent, the success probability evolves according to the function whose fixed point we are searching. Then, it is easy to observe that the success probability monotonically increases across levels, and goes to $x_k^*$ in the limit. $\qquad\square$

11

In fact, the branching algorithm achieves optimal success probability not only in the limit, but also among algorithms with a fixed depth.

**Proposition 4.7** (Optimality of branching algorithm for fixed-depth algorithms)**.** *Work under the setting of Proposition 4.6. Among algorithms with a fixed depth $L$, a branching algorithm with $L$ levels achieves the maximum success probability.*

*Proof.* Let $a_0 = g(0)$ and $a_i = \max_k \left( f(k) - (1 - a_{i-1})^k (f(k) - g(k)) \right)$ for $i \geq 1$. Clearly, a depth-$L$ branching algorithm can achieve success probability $a_L$. Furthermore, using the same induction as in the proof of Proposition 4.6, we can show that any depth-$L$ algorithm has success probability at most $a_L$. $\qquad\square$

These optimality results can be understood as that for these models, independence between solutions is favored. We will see in Section 6 that the natural algorithm that always passes the most recently generated solutions is not optimal because of said dependencies.

## 4.3 Optimality of the genetic algorithm

Via Lemma 3.4, Proposition 4.6 implies that the genetic algorithm also achieves optimal success probability.

**Proposition 4.8.** *Work under the setting of Proposition 4.6. Fix $k \geq 1$ and $\epsilon > 0$. The Genetic Algorithm (Algorithm 2) with $k$-way branching achieves success probability at least $x_k^* - \epsilon$ for large enough $L$, where $x_k^*$ is the unique solution of $x = f(k) - (f(k) - g(k))(1 - x)^k$ in $[0, 1]$.*

*Proof.* Let $\mathcal{A}$ denote the given oracle. Choose $\epsilon_1 > 0$ small enough such that the unique solution to $x(1 - \epsilon_1) = f(k) - (f(k) - g(k))(1 - x)^k$ in $[0, 1]$ is at least $x_k^* - \epsilon$. This means that there is a decaying model that is $1/(1 - \epsilon_1)$ dominated by $\mathcal{A}$ such that the optimal success probability is at least $x_k^* - \epsilon$. Then Lemma 3.4 shows that there is a genetic algorithm that achieves success probability at least $x_k^* - \epsilon$. $\qquad\square$

## 4.4 Optimality of the random sampling algorithm

The analysis of the random sampling algorithm relies on the theory of stochastic approximation [Duf13]. We use the following classic result.

**Theorem 4.9** ([Duf13, Theorem 2.2.12])**.** *Suppose $(X_n)_{n \geq 1}$ and $(Y_n)_{n \geq 1}$ are two sequences of random variables in $\mathbb{R}$ that are adapted to a filtration $(\mathcal{F}_n)_{n \geq 1}$ and linked by the equation $X_{n+1} = X_n + \frac{1}{n} Y_{n+1}$. Suppose*

$$\mathbb{E}\left[Y_{n+1} | \mathcal{F}_n\right] = f(X_n),$$
$$\mathbb{E}\left[(Y_{n+1} - f(X_n))^2 | \mathcal{F}_n\right] = \Gamma(X_n).$$

*In addition, suppose that the following holds for a finite constant $K$:*

1. *$f$ is a function of class $C^2$ such that*

$$f(x)^2 \leq K(1 + x^2),$$
$$f(x^*) = 0 \text{ and for } x \neq x^*, f(x)(x - x^*) < 0;$$

12

2. $\Gamma$ *is continuous in a neighborhood of* $x^*$ *and* $\Gamma(x) \leq K(1 + x^2)$;

3. *There exists an* $a > 0$ *such that* $\sup_n \mathbb{E}\left[|Y_{n+1} - f(X_n)|^{2+a}|\mathcal{F}_n\right] < \infty$.

*Then we have the following.*

1. $(X_n)_{n \geq 1}$ *converges to* $x^*$ *almost surely.*

2. *Let* $\tau = -f'(x^*)$ *and* $\Gamma^* = \Gamma(x^*)$.

   (a) *If* $\tau > 1/2$, *then* $\sqrt{n}(X_n - x^*)$ *converges in distribution to* $\mathcal{N}(0, \Gamma^*/(2\tau - 1))$.

   (b) *If* $\tau = 1/2$, *then* $\sqrt{n/\log n}(X_n - x^*)$ *converges in distribution to* $\mathcal{N}(0, \Gamma^*)$.

   (c) *If* $0 < \tau < 1/2$, *then* $n^\tau(X_n - x^*)$ *converges almost surely to a finite random variable.*

**Proposition 4.10.** *Work under the setting of Proposition 4.6. Fix* $k \geq 1$ *and* $\epsilon > 0$. *If* $g(k) > 0$, *then the Random Sampling Algorithm (Algorithm 3) with context size* $k$ *achieves success probability at least* $x_k^* - \epsilon$ *for large enough* $n$, *where* $x_k^*$ *is the unique solution of* $x = f(k) - (f(k) - g(k))(1 - x)^k$ *in* $[0, 1]$.

*Proof.* Define $X_t$ as the fraction of correct solutions in $\{s_1, \ldots, s_t\}$. Then we have

$$X_{t+1} = X_t + \frac{1}{t+1}(-X_t + Y_t)$$

where $Y_t = 1$ with probability $f(k) + (f(k) - g(k))(1 - X_t)^k$ and $Y_t = 0$ otherwise. Define $h(x) = f(k) - (f(k) - g(k))(1 - x)^k - x$. Then we have

$$X_{t+1} = X_t + \frac{1}{t+1}(h(X_t) + Z_t)$$

where $Z_t = Y_t - \mathbb{E}[Y_t|X_t]$ is a martingale difference sequence with bounded variance.

To apply Theorem 4.9, we need to verify that $h(x)(x - x_k^*) < 0$ for $x \neq x_k^*$. Note that $h(x)$ is concave on $[0, 1]$ and $h(0) = g(k) > 0$. So $h(x) > 0$ for $0 \leq x < x_k^*$ and $h(x) < 0$ for $x_k^* < x \leq 1$. Thus the condition holds.

By Theorem 4.9, $X_t$ converges almost surely to the unique root of $h(x) = 0$, which is exactly $x_k^*$. $\square$

## 5  Convergence Analysis

In this section we analyze the convergence rate of different reasoning algorithms. For simplicity, we focus on uniform models, but our analysis can be extended to general decaying models as well.

We fix a uniform model $\mathcal{A}_u^{(p,q,k)}$ as defined in Definition 2.2, where $0 < p < q < 1$ and integer $k \geq 1$. Let $x^*$ be the unique solution in $[0, 1]$ of the equation $x = q - (q - p)(1 - x)^k$. By Proposition 4.6, the optimal success probability of any reasoning algorithm using the oracle $\mathcal{A}_u^{(p,q,k)}$ converges to $x^*$ as the number of steps goes to infinity.

## 5.1 Convergence analysis of the branching algorithm

**Proposition 5.1.** *The success probability of the branching algorithm (Algorithm 1) with $L$ levels of $k$-way branching is $x^* - \left( k\frac{q-x^*}{1-x^*} \pm o(1) \right)^L$.*

*In terms of the number of calls $n = k^L$, the success probability is $x^* - \epsilon_n$ where*

$$\lim_{n \to \infty} \frac{\log(1/\epsilon_n)}{\log n} = \frac{-\log\left( k\frac{q-x^*}{1-x^*} \right)}{\log k}.$$

*Proof.* Define $f(x) = q - (q-p)(1-x)^k$. Let $x_j$ be the success probability after $j$ levels of branching. We have $x_0 = p$ and $x_{j+1} = f(x_j)$ for all $j \geq 0$.

We have shown in Proposition 4.6 that $x_j \to x^*$ as $j \to \infty$. The convergence rate is governed by $f'(x^*) = k(q-p)(1-x^*)^{k-1} = k\frac{q-x^*}{1-x^*}$. By the proof of Proposition 4.10, we have $0 < f'(x^*) < 1$. So $\lim_{j \to \infty} \frac{|x_{j+1}-x^*|}{|x_j-x^*|} = f'(x^*)$. Therefore, $|x_L - x^*| = (f'(x^*) \pm o(1))^L$. □

## 5.2 Convergence analysis of the genetic algorithm

**Proposition 5.2.** *There exists a Genetic Algorithm (Algorithm 2) with $k$-way branching whose success probability $x^* - \epsilon_n$ (where $n$ is the number of calls to the oracle) satisfies*

$$\lim_{n \to \infty} \frac{\log(1/\epsilon_n)}{\log n} = \frac{1}{2}.$$

*Proof.* In the proof of Proposition 4.8, we can take $\epsilon_1 = \Theta(\epsilon)$. Then by Lemma 3.4, the number of calls in the genetic algorithm (which achieves error probability $x^* - \epsilon$) is $\epsilon_1^{-2\pm o(1)} = \epsilon^{-2\pm o(1)}$. □

## 5.3 Convergence analysis of the random sampling algorithm

**Proposition 5.3.** *The success probability of the Random Sampling Algorithm (Algorithm 3) with context size $k$ after $n$ steps is $x^* - \epsilon_n$, where*

$$\lim_{n \to \infty} \frac{\log(1/\epsilon_n)}{\log n} = \min\left\{ 1/2, 1 - k\frac{q-x^*}{1-x^*} \right\}.$$

*Proof.* The result essentially follows from Theorem 4.9.

Let $h(x) = q - (q-p)(1-x)^k - x$ as in the proof of Proposition 4.10. Then $\tau = -h'(x^*) = 1 - k\frac{q-x^*}{1-x^*}$. We have $0 < \tau < 1$ by the proof of Proposition 4.10. By Theorem 4.9, the error decays polynomially in $n$ with exponent $-\min\{1/2, \tau\}$. □

# 6 Uniform Model

In this section, we consider the uniform model. We first recall its definition.

**Definition 2.2** (Uniform Model)**.** *Let $0 \leq p \leq q \leq 1$ and integer $k \geq 1$. For the uniform model $\mathcal{A}_u^{(p,q,k)}$, each call has $\mathrm{score}(\mathcal{A}_u^{(p,q,k)}(C)) = 1$ independently with probability*

$$p + (q-p)\mathbb{1}\left\{ \sum_{s \in C} \mathrm{score}(s) \geq 1 \text{ and } |C| \leq k \right\}$$

*and* $\operatorname{score}(\mathcal{A}_u^{(p,q,k)}) = 0$ *otherwise.*

We remark that this corresponds to a very simple special case of the decaying model where $f(k)$ is a step function. Clearly, it is best to pass exactly $k$ solutions to the context whenever $k$ solutions are available. We first observe that the uniform model is strongly monotone, and therefore Lemma 4.3 applies.

**Observation 6.1.** *In the uniform model with $q \geq p$, it is always beneficial to add $k$ solutions to the context.*

*Proof.* Follows directly from Lemma 4.3 since the uniform oracle is strongly monotone Definition 4.1. $\square$

From this, we immediately obtain that the maximum achievable success probability is the fixed point achieved when always adding $k$ independently generated solutions.

**Corollary 6.2.** *The maximum achievable success probability of a $(\mathcal{A}_u^{(p,q)}, n)$-reasoning algorithm with $q \geq p$ and maximum context length $k$ is the unique solution of*

$$x = q + (p - q)(1 - x)^k$$

*as $n \to \infty$.*

*Proof.* Directly follows from Observation 6.1 and Proposition 4.6 by reformulating. $\square$

The algorithm for achieving this probability may be quite inefficient in terms of the number of oracle calls for achieving a fixed probability. In the following we show that introducing dependencies to reduce the number of oracle calls in fact reduces the achieved success probability. First, we show that this is the case for $k = 2$ as a warm-up in Claim 6.3. This proof can be seamlessly generalized to larger $k$, which we show in Proposition 6.4.

**Claim 6.3** (Warm-up). *The sliding window approach where the oracle is always called on the $k$ most recently generated solutions achieves sub-optimal performance for $k = 2$.*

*Proof.* We can imagine this process as generating a sequence of solutions $\{s_i\}_{i=1,\dots,n}$, and we let $v_i = \operatorname{score}(s_i)$. Then imagine being at some stage $n$. The only quantities that matter for the transfer function describing the distribution of $v_{n+1}$ are $v_n$ and $v_{n-1}$. We distinguish three states. The state is the smallest number $i$ such that $v_{n-i} = 1$. If no such $i$ exists, or $i \geq k$, the state is $k$. We then obtain the following transition matrix from the definition of the oracle.

$$W = \begin{pmatrix} q & q & p \\ 1 - q & 0 & 0 \\ 0 & 1 - q & 1 - p \end{pmatrix}^{\top}$$

We then compute the stationary distribution of this matrix. Solving the system $W^T \pi = \pi$ yields

$$\pi = \frac{1}{N} \begin{pmatrix} p \\ (1 - q)p \\ (1 - q)^2 \end{pmatrix}$$

where $N = p + (1-q)p + (1-q)^2$. We output the correct solution whenever we are in state 0, so the probability of returning a correct solution is given by $p/N$.

Now, given Corollary 6.2 we aim to show that $p/N < q + (p-q)(1-p/N)^2$. Reformulating yields

$$(p-q)^2(1-q)^2 > 0$$

which is true whenever $q < 1$ and $p \neq q$. □

Next, we generalize the previous claim to all $k$.

**Proposition 6.4.** *The sliding window approach where the oracle is always called on the $k$ most recently generated solutions achieves sub-optimal performance.*

*Proof.* We again consider a sequence of solutions $\{s_i\}_{i=1,\ldots,n}$ and let $v_i = \text{score}(s_i)$ for $i \in 1, \ldots, n$.

Then, for $0 \leq i \leq n$, $0 \leq j \leq n$, we let $x_{i,j}$ denote the probability that

$$j = \min\{k, \max\{a : v_i = \cdots = v_{i-a+1} = 0\}\}.$$

Then, we have $x_{0,j} = \mathbb{1}_{j=k}$ and $x_{i+1,*} = x_{i,*}W$ where the transition matrix $W$ is given by $W(k,0) = p$, $W(k,k) = 1-p$, $W(i,0) = q$, $W(i,i+1) = 1-q$ for $0 \leq i \leq k-1$ and $W$ contains 0 entries everywhere else. So $x_{n,*} = x_{0,*}W^n$. We remark that this transition matrix is completely analogous to the warm up for $k = 2$ presented in Claim 6.3.

The stationary distribution $\pi$ is given by $\pi(j) = C(1-q)^j$ for $0 \leq j \leq k-1$, $\pi(k) = C(1-q)^k/p$, where $C^{-1} = \sum_{0 \leq j \leq k-1}(1-q)^j + (1-q)^k/p = (1-(1-q)^k)/q + (1-q)^k/p$. Since the algorithm succeeds whenever it is in state 0, the limiting success probability of the sliding window algorithm is equal to

$$\pi(0) = \left( (1-(1-q)^k)/q + (1-q)^k/p \right)^{-1}.$$

In general, this is not a solution to the equation $x = q + (p-q)(1-x)^k$. So the sliding window algorithm is not optimal for general $k$. □

Therefore, even in this extremely simple model correlations between the solutions are an issue, which motivates considering the branching and geometric algorithm. We remark that for certain decaying models (such that the exponential decaying model and the polynomial decaying model) the success probability of a sliding window algorithm goes to 0 as the number of steps goes to infinity because the context will eventually only contain wrong solutions. After this point, the oracle never generates a correct solution again. This can be seen as modeling the "overthinking" phenomenon observed in practice [ZLM+25].

## 7 Exponentially and Polynomially Decaying Models

Next, we analyze the exponentially and polynomially decaying models $\mathcal{A}_d^{(p,f)}$, which are a more realistic model that introduces a non-trivial dependency on the context size (See Appendix A for experimental results). We first recall the general model definition and then the exponentially and polynomially decaying special cases.

**Definition 2.1** (Decaying Model). *Let $\mathcal{V} = \{0, 1\}$. Let $f : \mathbb{N} \to [0, 1]$ and $g : \mathbb{N} \to [0, 1]$ be two functions satisfying $f(k) \geq g(k)$ for all $k \in \mathbb{N}$. The decaying model $\mathcal{A}^{(f,g)}$ is defined as the oracle satisfying for any context $C \subseteq \mathcal{S}$:*

$$\Pr[\text{score}(\mathcal{A}^{(f,g)}(C)) = 1] = \begin{cases} f(|C|), & \text{if } \sum_{s \in C} \text{score}(s) \geq 1, \\ g(|C|), & \text{if } \sum_{s \in C} \text{score}(s) = 0. \end{cases}$$

**Definition 2.3** (Exponential Decay). *Let $0 < p \leq q \leq 1$. The exponential decaying model is the decaying model $\mathcal{A}^{(f,g)}$ where $f(k) = q^{k-1}$ and $g(k) = p\mathbb{1}_{\{k=0\}}$.*

**Definition 2.4** (Polynomial Decay). *Let $0 < p \leq q \leq 1$. The polynomial decaying model is the decaying model $\mathcal{A}^{(f,g)}$ where $f(k) = \frac{1}{k^q}$ and $g(k) = p\mathbb{1}_{\{k=0\}}$.*

Unlike the uniform setting, where it is always beneficial to pass as many solutions as possible in the context, a non-constant decay function $f(k)$ realizes a non-trivial trade off. In the following, we explore the two decay functions that are arguably most natural, exponential and polynomial decay (See Definition 2.3 and Definition 2.4). We sometimes write $\mathcal{A}^{(f,p)}$ when referring to the exponential and polynomial decay model with $g(k) = p\mathbb{1}_{\{k=0\}}$.

## 7.1 Exponential Decay

In this section, we examine exponential decay functions $e_q(k) = q^{k-1}$ (See Definition 2.3). These already lead to a non-trivial trade off when selecting the context size. Due to the quick reduction in success probability, they encourage only combining a small amount of solutions.

**A two stage process.**     We first show how this trade-off is realized by examining the simple two stage process *pass@k*. In the first stage, we assume to be given an algorithm that samples an arbitrarily sized collection of solutions, each of which is correct independently with some probability $x$. Then, we aim to select an optimal number of such solutions to put into the context for a single oracle call to maximize the probability of correctness of the generated solution. This corresponds to maximizing the success probability that a *pass@k* algorithm can obtain.

**Claim 7.1.** *The optimal amount of independent solutions that are correct with probability $x$ to pass to the oracle $\mathcal{A}_d^{(e_q,p)}$ for $x \geq p$ is*

$$\arg\max_{k \in \mathbb{N}_{\geq 0}} q^{k-1}(1 - (1-x)^k) = \max\left(1, \left\lfloor 1 + \frac{\log\left(\frac{1-q}{1-(1-x)q}\right)}{\log(1-x)} \right\rfloor\right).$$

*Proof.* Since passing a single solution yields success probability $x \geq p$ it is never beneficial to pass 0 solutions. We now analyze the function $g(k) \stackrel{\text{def}}{=} q^{k-1}(1 - (1-x)^k)$ that we aim to maximize. We examine how this function changes as we increase $k$.

We have $f(k+1)/g(k) = q\frac{1-(1-x)^{k+1}}{1-(1-x)^k}$. Since the term $\frac{1-(1-x)^{k+1}}{1-(1-x)^k}$ is monotonically decreasing with $k$, we have to find the largest $k$ for which $q\frac{1-(1-x)^{k+1}}{1-(1-x)^k} \geq 1$ which is equivalent to

$$(1-x)^k \geq \frac{1-q}{1-q(1-x)}. \tag{1}$$

We then solve for the continuous $k'$ for which (1) holds with equality by taking the logarithm on each side. We obtain that $(1-x)^{k'} = \frac{1-q}{1-q(1-x)}$ for

$$k' = \frac{\log\left(\frac{1-q}{1-(1-x)q}\right)}{\log(1-x)}.$$

Since this shows there is still an increase when going from $k'$ to $k'+1$ in this case, we obtain that the optimal discrete solution is as claimed. $\qquad\square$

We observe that whenever progress is possible for some $k$, setting $k = 2$ also makes progress. This property of the exponential decay model enables us to derive a closed form solution of the maximum achievable success probability.

**Corollary 7.2.** *Whenever* $\arg\max_{k\in\mathbb{N}_{\geq 0}} q^{k-1}(1 - (1-x)^k) > k$, *then* $q(1-x)^2 \geq x$.

*Proof.* As observed in the proof of Claim 7.1, the function $g(k) = q^{k-1}(1-x)^k$ first monotonically increases until it reaches the optimal step $k$, and then monotonically decreases. Since $g(1) = x$, we are guaranteed that $g(2) > x$ whenever the optimal $k$ is at least 2. This proves the corollary. $\qquad\square$

**The maximum achievable success probability.**     Similar to the case of the uniform model, we study the maximum achievable success probability with the exponential decay function.

**Lemma 7.3.** *The maximum achievable success probability of a* $(\mathcal{A}_d^{(e_q,p)}, n)$-*reasoning algorithm with* $q \geq p$ *is*

$$\max\left(p, 2 - \frac{1}{q}\right).$$

*Proof.* We start with the same observations as in the proof of Section 4, namely that all solutions generated by a reasoning algorithm are either positively correlated or independent, and that it is always optimal to pass independent solutions with as high correctness probability as possible. We therefore obtain from Proposition 4.6 that the solution has to be a fixed point of the function $f(x) = q(1 - (1-x)^2)$ because the largest fixed point is achieved at $k = 2$ by Corollary 7.2. This function has only one fixed point $x = 2 - \frac{1}{q}$. $\qquad\square$

We remark that it is never useful to pass more than one solution as context to $\mathcal{A}_d^{(e_q,p)}$ when $q \leq \frac{1}{2}$, as in this case the model performs worse than randomly sampling a solution from the context. Therefore, we have $q \geq 1/2$ for all boostable models.

**Convergence analysis.**     In the following, we analyze the speed of convergence. In contrast to Section 5, we focus on the behavior when the initial success probability approaches 0. This corresponds to the setting where boosting is the most useful.

Here, we assume $k = 2$ throughout to achieve a lower bound. Furthermore, we will artificially weaken the oracle by a factor $(1 - \epsilon)$. This allows us to immediately apply Lemma 3.4 to obtain a genetic algorithm and will not significantly change the convergence analysis.

We consider the branching algorithm that starts by generating a first generation of solutions with success probability $p_0 = p$, and then iteratively generates the next generation of solutions with success probability $p_i$ by combining two fresh solutions from the previous generation. Formally, this is equivalent to the branching algorithm (See Algorithm 1) with $k_i = 2$ for all $i$.

18

**Lemma 7.4.** *Let $p_{\max} = 2 - \frac{1}{q} > p$. Then we obtain a solution with success probability $p_{\max} - \epsilon$ after $L = O(\frac{1}{\epsilon}\log(1/p))$ generations. This also holds if the success probability of the oracle is decreased by a multiplicative factor $1 - \epsilon/8$.*

*Proof.* We have $p_{i+1} = q(1 - (1 - p_i)^2)$. We will show that $p_{i+1} \geq p_i(1 + \frac{\epsilon}{4})$, or equivalently

$$p_i(1 + \frac{\epsilon}{4}) \leq (1 - \epsilon/8) \cdot q(1 - (1 - p_i)^2)$$

which is implied by

$$p_i(1 + \frac{\epsilon}{2}) \leq q(1 - (1 - p_i)^2) \tag{2}$$

as long as $p_i \leq p_{\max} - \epsilon$. Reformulating (2) yields

$$p_i \leq 2 - \frac{1 + \frac{\epsilon}{2}}{q}.$$

Since $q \in [\frac{1}{2}, 1]$, this condition always holds when $p_i \leq p_{\max} - \epsilon$. Therefore, the success probability increases by at least a multiplicative factor $(1 + \frac{\epsilon}{4})$ in each iteration. $\square$

**Corollary 7.5.** *There is a genetic algorithm (See Algorithm 2) with $O(\log(1/\epsilon)\log(1/p)/(p\epsilon^3))$ oracle calls that achieves success probability $p_{\max} - \epsilon$. The sequential depth of the algorithm is $O(\log(1/p)/\epsilon)$.*

*Proof.* Directly follows from Lemma 7.4 and Lemma 3.4. $\square$

In our proof, we merely focus on context sizes equal to 2. For very large $q$, there can however be substantial gains in the number of rounds when combining more answers at a time. In particular, for $q = (1 - p)$ we can get rid of the dependence on $\log(1/p)$ in the number of rounds entirely by initially combining $\approx 1/p$ solutions. For most practical settings where $q$ should be thought of as a constant close to 1, the studied setting captures the behavior.

## 7.2 Polynomial Decay

In this section, we examine polynomial decay functions (See Definition 2.4). Since these functions drop off much slower than the exponential decay we examined in the previous section, they allow much faster boosting. In particular, we can show that approximately $\approx \log\log 1/p$ iterations suffice to boost to constant success probability with a simple algorithm.

**A simple algorithm.** We describe a simple algorithm that automatically scales $k$ and quickly increases the success probability to a small constant. Given access to independent solutions with correctness probability at least $p$, the algorithm samples $k = \lfloor 1/p \rfloor$ such solutions as context $C$ and passes them to the oracle $\mathcal{A}_d^{(p_q,p)}(C)$. Then, the oracle generates a solution that is correct with probability $p'$ bounded by

$$p' \geq \frac{1}{k^q}(1 - (1 - p)^k) \geq p^q(1 - (1 - p)^{1/p - 1}).$$

For $p \leq 1/2$, we obtain

$$p' \geq p^q(1 - 2/e).$$

For a constant $q < 1$, this series will recover a squaring type behavior. We next show that this series increases very quickly for a general constant $c$. This later allows us to replace $(1 - 2/e)$ with $(1 - 2/e)(1 - \epsilon')$ to weaken the oracle and obtain convergence guarantees for the genetic algorithm via Lemma 3.4.

**Claim 7.6.** *Consider $p_0 = p$, and $p_i = p_{i-1}^q \cdot c$ for $i \geq 1$. Then, $p_n = p^{q^n} \cdot c^{\frac{1-q^n}{1-q}}$ for $n \geq 1$.*

*Proof.* The proof is by induction. The base case $n = 1$ directly follows from the definition of $p_1$.
   Now, we assume the recurrence holds for $n$ and aim to show it for $n + 1$. We have

$$
\begin{aligned}
p_{n+1} &= p_n^q \cdot c \\
&= \left( p^{q^n} \cdot c^{\frac{1-q^n}{1-q}} \right)^q \cdot c \\
&= p^{q^{n+1}} \cdot c^{1 + q\frac{1-q^n}{1-q}} = p^{q^{n+1}} c^{\frac{1-q^{n+1}}{1-q}}.
\end{aligned}
$$

$\square$

   From Claim 7.6, we observe that this simple algorithm converges to a success probability of $c^{\frac{1}{1-q}}$. For constant $q$ and $c$, this corresponds to a constant success probability.

**Claim 7.7.** *We have $p_n \geq (1 - \epsilon)c^{\frac{1}{1-q}}$ for some $n = O(\log\log(1/p)/q + \log(\epsilon^{-1})/q)$.*

*Proof.* We have $p_n = p^{q^n} c^{\frac{1-q^n}{1-q}} = p^{q^n} c^{\frac{1}{1-q}} c^{\frac{-q^n}{1-q}}$. We first show that $c^{\frac{-q^n}{1-q}} \geq (1 - \epsilon/2)$ for $n = O(\log\log(1/p) + \log(\epsilon^{-1}/q)$. We take the logarithm on both sides

$$\frac{-q^n}{1-q} \log(c) \leq \log(1 - \epsilon/2)$$

$$q^n \geq -(1 - q)\frac{\log(1 - \epsilon/2)}{\log(c)}$$

We then approximate $\log(1 - \epsilon) \approx \epsilon$ and collect the constant terms into $C$ obtaining the condition

$$n \geq C' \log(1/\epsilon)/q$$

for some constant $C'$.
   Next, we show that $p^{q^n} \geq (1 - \epsilon/2)$. We have

$$
\begin{aligned}
p^{q^n} &\geq (1 - \epsilon/2) \\
q^n \log(p) &\geq \log(1 - \epsilon/2) \\
n &\geq C(\log(\epsilon^{-1}) + \log\log(1/p))/q
\end{aligned}
$$

for some large constant $C$. This concludes the proof. $\square$

From Claim 7.7, we obtain that this algorithm indeed converges very quickly. Furthermore, it does so even when $c$ is slightly smaller. This allows it to pair very well with the genetic algorithm described in Algorithm 2. We therefore derive a genetic algorithm.

In the following, we focus on the behavior when the initial success probability approaches 0. Once a reasonably high (e.g. constant) success probability is achieved, the type of convergence analysis from Section 5 could be applied to obtain the speed of convergence for the final bit.

**Lemma 7.8.** *There is a genetic algorithm that using the model $\mathcal{A}_d^{(p_q,p)}$ finds a solution that is correct with probability at least $(1/16)^{1/(1-q)}$ in $O(\frac{1}{pq} \log\log 1/p)$ oracle calls.*

*Proof.* Directly follows from Claim 7.7 and Lemma 3.4 and $(1 - (1 - p)^{1/p-1} \geq (1 - 2/e)$ for $p \geq 1/2$. $\qquad\square$

**The maximum achievable success probability.** To characterize the maximum achievable success probability we resort to Proposition 4.6 since there does not appear to be an obvious simpler characterization.

# 8 Conclusion

We introduce a theoretical model that formalizes reasoning oracles that are given an array of solutions as context. These formalize the experimentally observed behavior that adding correct solutions to the context helps, but the accuracy decreases with the context size. The derived theory provides a framework for developing and analyzing reasoning algorithms that have been shown to perform well in practice [VJM+25].

**Future work.** The algorithms we develop are naturally parallel. It is interesting to further explore the trade off between solution quality, parallelism and the total number of model calls. We believe that less parallelism could lead to savings in logarithmic factors that are highly relevant given the cost of executing such algorithms in practice.

**Limitations and future model extensions.** We limit ourselves to the binary model, where a solution is either correct or not. While this enables us to devise simple experiments, this is a significant drawback when considering more difficult questions for which we cannot hope to obtain correct solutions directly.

It therefore makes sense to extend our models to more fine grained score functions in the future. Such an extension should go hand in hand with the introduction of a measure of answer diversity. That is, because combining two partially correct solutions should only be beneficial if they complement each other. This could be modeled by additionally associating a unit vector with each solution where orthogonal vectors correspond to solutions that perfectly complement each other. These vectors can either be hidden alongside the score function or exposed through some oracle, where the former is simpler than the latter but yields less flexibility when developing algorithm.

While our model offers exciting opportunities for extensions, we caution that it is important to ground them in experiments to ensure that the theory captures something interesting about the world we live in.

# 9  References

[AVL+24]  Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. Large language models for mathematical reasoning: Progresses and challenges. In Neele Falk, Sara Papi, and Mike Zhang, editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 225–237, St. Julian's, Malta, March 2024. Association for Computational Linguistics. 2

[Duf13]  Marie Duflo. *Random iterative models*, volume 34. Springer Science & Business Media, 2013. 12

[FGK71]  C. M. Fortuin, J. Ginibre, and P. W. Kasteleyn. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89 − 103, 1971. 10

[FZG+24]  Guhao Feng, Bohang Zhang, Yuxian Gu, Haotian Ye, Di He, and Liwei Wang. Towards revealing the mystery behind chain of thought: A theoretical perspective. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2024. 5

[GTLV22]  Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes. In *Advances in Neural Information Processing Systems*, volume 35, 2022. 4

[HY25]  Yichen Huang and Lin F. Yang. Winning gold at imo 2025 with a model-agnostic verification-and-refinement pipeline, 2025. 2

[Kah11]  Daniel Kahneman. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York, 2011. 2

[KN24]  T. Kumar and Neel Nanda. The complexity dynamics of grokking. *arXiv preprint arXiv:2412.09810*, 2024. 5

[LFZ25]  Ziyue Li, Chenrui Fan, and Tianyi Zhou. Grokking in llm pretraining? monitor memorization-to-generalization without test, 2025. 5

[LIPO23]  Yingcong Li, M. Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. Transformers as algorithms: Generalization and stability in in-context learning. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. 4

[MU05]  Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, USA, 2005. 8

[OEN+22]  Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Nelson, Stephen Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. 4

[SCG+23] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 2

[SVH24] Yash Sarrof, Yana Veitsman, and Michael Hahn. The expressive capacity of state space models: A formal language perspective. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 5

[VJM+25] Siddarth Venkatraman, Vineet Jain, Sarthak Mittal, Vedant Shah, Johan Obando-Ceron, Yoshua Bengio, Brian R. Bartoldson, Bhavya Kailkhura, Guillaume Lajoie, Glen Berseth, Nikolay Malkin, and Moksh Jain. Recursive self-aggregation unlocks deep thinking in large language models, 2025. 2, 3, 8, 21

[WWS+22] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc. 2

[WWS+23] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 2

[YBR+20] Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank J Reddi, and Sanjiv Kumar. Are transformers universal approximators of sequence-to-sequence functions? In *International Conference on Learning Representations*, 2020. 4

[YSH+25] Yibo Yan, Jiamin Su, Jianxiang He, Fangteng Fu, Xu Zheng, Yuanhuiyi Lyu, Kun Wang, Shen Wang, Qingsong Wen, and Xuming Hu. A survey of mathematical reasoning in the era of multimodal large language model: Benchmark, method & challenges. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11798–11827, Vienna, Austria, July 2025. Association for Computational Linguistics. 2

[YYZ+23] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 2

[ZLM+25]   Chentian Zhang, Xubo Liu, Marcus Mak, Jackie Chi Kit Cheung, and Yanshuai Cao. When more is less: Understanding chain-of-thought length in llms. In *ICLR 2025 Workshop on Reasoning and Planning for LLMs*, 2025. 5, 16

# A    Experiments

In this section, we verify some of our theoretical assumptions with the AIME 2025 math dataset and the Gemini 2.5 Pro model. The dataset became available after the model, which helps us avoid test set contamination. AIME 2025 has 30 questions, and we plot the average accuracy for each question across independent model calls in Figure 1. The prompt consists only of the question. As expected, Gemini 2.5 Pro has a strong performance on many of the questions.
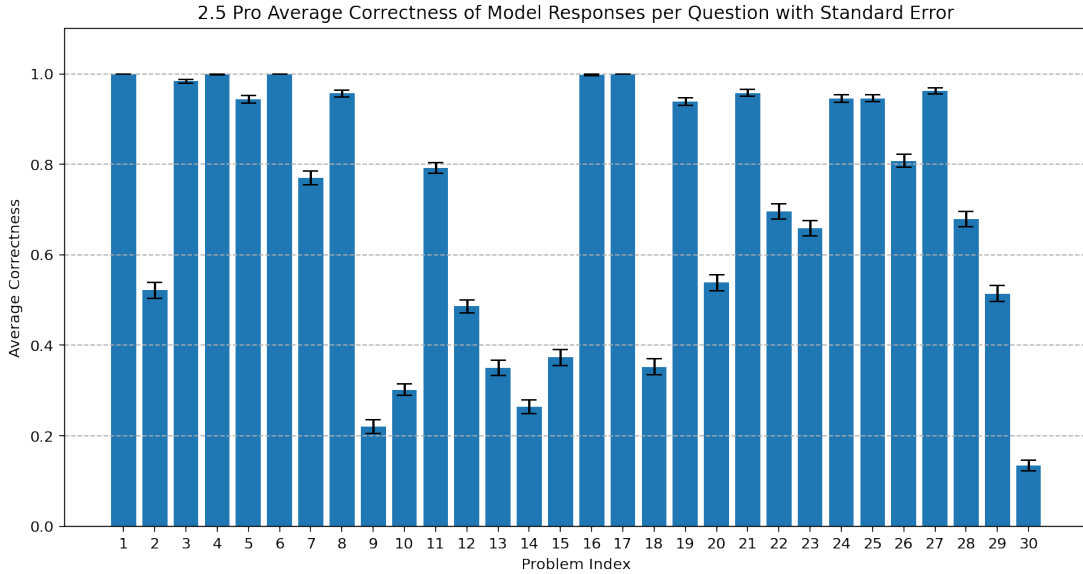


Figure 1: Average AIME 2025 accuracy per question for Gemini 2.5 Pro. We use 780 model calls per question. The error bars represent standard error.

Next, we saved the outputs and picked questions 10 and 12 to examine model accuracy when solutions are provided in context. These are questions for which the model struggled but occasionally generated correct answers. We study two settings below.

## A.1    Fixing one correct solution and increasing incorrect solutions

We start by investigating model accuracy when only one of the solutions is correct, and the number of incorrect solutions grows from 0 to 12. For each of the configurations, we make 30 calls with Gemini 2.5 Pro and make sure none of the solutions are repeated across any of the calls. The solutions are also shuffled before placing them in context. Our results are in Figure 2, and we observe a decaying accuracy as the number of incorrect questions increases. This motivates our theoretical models with decay.

## A.2    Fixing total number of solutions and varying correct solutions

Finally, we move to a setting where instead of fixing one correct answer, we fix the total number of solutions and vary the number of correct versus incorrect answers. We show results in Figure 3. We

Figure 2: Gemini 2.5 Pro accuracy when providing 1 correct solution answer and $k = 0$ to 12 incorrect solutions. The red line represents base accuracy of Pro on the question without any solutions. The green lines represent $1/(k+1)$, the accuracy when returning a solution in the context uniformly at random. The orange bar is the most likely configuration according to the base accuracy (the $k$ for which $1/(k+1)$ is closest to base accuracy).

again used 30 calls per configuration, did not reuse solutions across any model calls, and shuffled the order of the solutions.

We observe a very smoothly increasing accuracy as the number of correct solutions increases. We also note that the performance of this setup—sampling 5 solution answers and providing them as context for a final model call—always improves over the base accuracy of the model on average. On the two questions we examined for Pro, the gap is about 40%, suggesting that it is a particularly strong verifier.

Figure 3: Gemini 2.5 Pro accuracy when providing 5 total solutions and varying the number of correct versus incorrect answers. The red line represents base accuracy of Pro on the question without any solutions. The numbers on top of each bar represent the probability of that configuration according to the average correctness. The orange bar highlights the most likely configuration. The green line uses these weights to give the accuracy when sampling 5 solutions and using them as context for a final model call.