

∇ -REASONER: LLM REASONING VIA TEST-TIME GRADIENT DESCENT IN TEXTUAL SPACE

Anonymous authors

Paper under double-blind review

ABSTRACT

Scaling inference-time compute for Large Language Models (LLMs) has unlocked unprecedented reasoning capabilities. However, existing inference-time scaling methods typically rely on inefficient and suboptimal discrete search algorithms or trial-and-error prompting to improve the online policy. In this paper, we propose ∇ -Reasoner, an iterative generation framework that integrates differentiable optimization over token logits into the decoding loop to refine the policy on the fly. Our core component, Differentiable Textual Optimization (DTO), leverages gradient signals from both the LLM’s likelihood and a reward model to refine textual representations. ∇ -Reasoner further incorporates rejection sampling and acceleration design to robustify and speed up decoding. **Theoretically, we show that performing inference-time gradient descent in the sample space to maximize reward is dual to aligning an LLM policy via KL-regularized reinforcement learning.** Empirically, ∇ -Reasoner achieves over 20% accuracy improvement on a challenging mathematical reasoning benchmark, while reducing **number of model calls** by approximately 10-40% compared to strong baselines. Overall, our work introduces a paradigm shift from zeroth-order search to first-order optimization at test time, offering a cost-effective path to amplify LLM reasoning.

1 INTRODUCTION

Large Language Models (LLMs) have unlocked remarkable reasoning capabilities (Radford et al., 2018; 2019; Brown et al., 2020), enabling machines to tackle challenges considered exclusive to human cognition, such as solving complex mathematical problems (Cobbe et al., 2021; Lewkowycz et al., 2022; Uesato et al., 2022; Lee et al., 2023; Yang et al., 2024b) and executing long-horizon planning (Liu et al., 2023a; Valmeeekam et al., 2023; Song et al., 2023). Such capabilities arise through large-scale pre-training on massive datasets, followed by careful post-training alignment (Wei et al., 2022a; Ouyang et al., 2022; Guo et al., 2025). A prevailing observation has indicated that scaling both model size and training data leads to continual improvements in LLM reasoning ability (Kaplan et al., 2020; Hoffmann et al., 2022).

Nevertheless, recent empirical findings increasingly suggest that scaling inference-time computation can be also crucial and perhaps more cost-effective than expanding pretraining to further enhance reasoning and problem-solving abilities (Snell et al., 2024). Chain-of-Thought (CoT) (Wei et al., 2022b) demonstrates that prompting LLMs at the test time to generate longer sequences with intermediate reasoning steps significantly improves their reasoning accuracy. Built on CoT, Wang et al. (2022) further scales the inference compute by sampling multiple reasoning chains and selecting the most consistent one, leading to enhanced performance. More recently, inference-time scaling has been augmented with reward models to refine reasoning quality. Notable approaches such as Tree-of-Thought (TOT) (Yao et al., 2024) and Reasoning-as-Planning (RAP) (Hao et al., 2023) cast LLM reasoning as a decision-making problem and employ strategic sampling algorithms to estimate the reward-to-go, thereby refining the sequential prediction policy at each decoding step. Underlying these approaches are extensive prompting-based search procedures that traverse the sequence space, with the LLM serving as a guiding heuristic. However, such approaches often struggle to adequately explore the sample space and thus become sensitive to sparse and noisy reward signals as reasoning chains grow longer and the search space expands exponentially. Consequently, their performance tends to saturate even when inference-time computation is substantially increased.

054 While existing methods fall into zeroth-order
 055 algorithms that rely solely on reward values, we
 056 note that first-order methods, providing direc-
 057 tional guidance for optimization, can be even
 058 more effective in searching for optimal solutions,
 059 overcoming the sparsity of the reward landscape
 060 (see an intuitive comparison in Fig. 1). In fact,
 061 gradient information is readily available during
 062 the LLM reasoning process, as both the LLM
 063 and reward function can be differentiable. In
 064 this paper, we introduce ∇ -Reasoner, a novel
 065 reasoning algorithm that applies inference-time
 066 gradient descent in the sample space to refine
 067 the outputs of a base policy prior to next-token
 068 prediction. The overall pipeline follows an iterative
 069 decoding process. At each step, the language model
 070 first generates a full completion together with its
 071 per-token logits, serving as the initial rollout. The core component, termed Differentiable Textual
 072 Optimization (DTO), then refines these token logits via gradient descent.

073 DTO formulates the reasoning process as a continuous optimization problem over the reward land-
 074 scape, directly leveraging gradients to refine textual representations. Specifically, DTO applies
 075 gradient descent to optimize the initial logit vectors from the base policy under an objective that
 076 combines the reward function with the sequence-level log-likelihood estimated by the language
 077 model. To enable end-to-end differentiability, we employ the straight-through estimator to map logit
 078 parameters into one-hot token vectors (Bengio et al., 2013). In this formulation, the reward function
 079 provides directional signals that guide tokens toward high-reward regions, while the log-likelihood
 080 term regularizes the tuned sequences to remain fluent and consistent with the pre-trained LLM
 081 distribution (Hoang et al., 2017; Qin et al., 2022; Kumar et al., 2022).

082 After refining the logits with DTO, ∇ -Reasoner treats the optimized logits of the first token as an
 083 improved policy and samples the next-token prediction from this updated distribution. We further
 084 integrate ∇ -Reasoner with rejection sampling, which accepts the token drawn from the refined policy
 085 only if it can yield continuation with higher reward; otherwise, the method reverts to the initial
 086 choice. Through iteratively interleaving decoding and refinement, ∇ -Reasoner scales inference-time
 087 reasoning by allocating additional computation to improve the LLM policy via gradient-based updates
 088 on the output space, efficiently backended by the parallel execution of transformer models. To further
 089 increase decoding throughput, we introduce a set of acceleration strategies that selectively skip tokens
 090 unlikely to benefit from DTO and reuse rollouts shared among decoding steps.

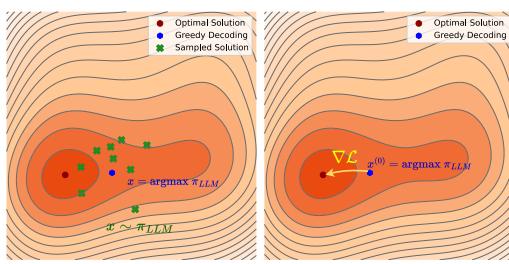
091 Theoretically, we show that DTO enables bidirectional gradient propagation along the sequence,
 092 facilitating global modifications that are crucial for effective reasoning (Bachmann & Nagarajan, 2024;
 093 Hao et al., 2023). Furthermore, we establish a close connection between DTO and RL algorithms
 094 (Schulman et al., 2017; Ouyang et al., 2022; Guo et al., 2025). We prove that sampling from an
 095 optimized LLM trained with RL is equivalent to directly drawing samples from the reference LLM
 096 and subsequently refining them through the gradient flow induced by DTO. This insight provides a
 097 new theoretical perspective for test-time approaches for reasoning.

098 Empirically, ∇ -Reasoner significantly enhances the mathematical reasoning capabilities by 10-40%
 099 across multiple models and benchmarks. It consistently outperforms strong inference-time baselines
 100 such as Best-of-N and RAP (Hao et al., 2023), while achieving accuracy on par with more costly
 101 training-based methods (e.g., GRPO). **We further show that ∇ -Reasoner scales compute more
 102 effectively: by leveraging parallelized execution of attention, it can utilize more compute per model
 103 forward pass. Henceforth, when comparing with sampling-only methods (e.g., BoN), ∇ -Reasoner
 104 achieves superior results while reducing the number of model calls by up to 40.2%.**

105 2 PRELIMINARIES

106 In this section, we formulate LLM reasoning as a decision-making problem, introducing the necessary
 107 notations and common approaches to address this problem along the way.

108 **Notations.** Let $\mathcal{V} = \{\delta_i \in \mathbb{R}^{|\mathcal{V}|} : i \in [|\mathcal{V}|]\}$ be the vocabulary set, where δ_i is the i -th canonical
 109 basis and $|\mathcal{V}|$ is the vocabulary size. We denote a sequence over this vocabulary as $x = [x_1, \dots, x_{|\mathcal{V}|}]$



110 Figure 1: Comparison between zeroth-order and
 111 first-order methods under the landscape of reward.
 112 The overall pipeline follows an iterative
 113 decoding process. At each step, the language model
 114 first generates a full completion together with its
 115 per-token logits, serving as the initial rollout. The core component, termed Differentiable Textual
 116 Optimization (DTO), then refines these token logits via gradient descent.

117 After refining the logits with DTO, ∇ -Reasoner treats the optimized logits of the first token as an
 118 improved policy and samples the next-token prediction from this updated distribution. We further
 119 integrate ∇ -Reasoner with rejection sampling, which accepts the token drawn from the refined policy
 120 only if it can yield continuation with higher reward; otherwise, the method reverts to the initial
 121 choice. Through iteratively interleaving decoding and refinement, ∇ -Reasoner scales inference-time
 122 reasoning by allocating additional computation to improve the LLM policy via gradient-based updates
 123 on the output space, efficiently backended by the parallel execution of transformer models. To further
 124 increase decoding throughput, we introduce a set of acceleration strategies that selectively skip tokens
 125 unlikely to benefit from DTO and reuse rollouts shared among decoding steps.

126 Theoretically, we show that DTO enables bidirectional gradient propagation along the sequence,
 127 facilitating global modifications that are crucial for effective reasoning (Bachmann & Nagarajan, 2024;
 128 Hao et al., 2023). Furthermore, we establish a close connection between DTO and RL algorithms
 129 (Schulman et al., 2017; Ouyang et al., 2022; Guo et al., 2025). We prove that sampling from an
 130 optimized LLM trained with RL is equivalent to directly drawing samples from the reference LLM
 131 and subsequently refining them through the gradient flow induced by DTO. This insight provides a
 132 new theoretical perspective for test-time approaches for reasoning.

133 Empirically, ∇ -Reasoner significantly enhances the mathematical reasoning capabilities by 10-40%
 134 across multiple models and benchmarks. It consistently outperforms strong inference-time baselines
 135 such as Best-of-N and RAP (Hao et al., 2023), while achieving accuracy on par with more costly
 136 training-based methods (e.g., GRPO). **We further show that ∇ -Reasoner scales compute more
 137 effectively: by leveraging parallelized execution of attention, it can utilize more compute per model
 138 forward pass. Henceforth, when comparing with sampling-only methods (e.g., BoN), ∇ -Reasoner
 139 achieves superior results while reducing the number of model calls by up to 40.2%.**

108 where $x_l \in \mathcal{V}$ is the l -th token for every $i \in [\lvert \mathbf{x} \rvert]$ and $\lvert \mathbf{x} \rvert$ represents the length of the sequence. We
109 also denote $\mathbf{x}_{\leq i}$ as the subsequence up to and including the i -th token, expressed as $[\mathbf{x}_1, \dots, \mathbf{x}_i]$.
110 The space of all sequences with finite length is given by $\mathcal{V}^* = \bigcup_{l \in \mathbb{N}} \mathcal{V}^l$.
111

112 **Language Models and Reward Models.** In this paper, we focus on autoregressive language
113 models (Radford et al., 2018; 2019; Brown et al., 2020), denoted as $\pi_{LLM} : \mathcal{V}^* \rightarrow [0, 1]$. The
114 language model can characterize the conditional probability of a question-answer pair. Given a pair of
115 questions and answers $\mathbf{x}, \mathbf{y} \in \mathcal{V}^*$, the model estimates their likelihood by the following factorization:
116 $\pi_{LLM}(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{|\mathbf{y}|} \pi_{LLM}(\mathbf{y}_i|\mathbf{y}_{\leq i-1}, \mathbf{x})$. We also denote $\text{Cat}(\pi_{LLM}(\cdot|\mathbf{y}_{\leq i-1}, \mathbf{x})) \in [0, 1]^{|\mathcal{V}|}$ as
117 the categorical distribution of \mathbf{y}_i given the prefix $\mathbf{y}_{\leq i-1}$ and \mathbf{x} . In addition, we define a reward model
118 as the function $r : \mathcal{V}^* \rightarrow \mathbb{R}$. $r(\mathbf{y}|\mathbf{x})$ evaluates the correctness of the response \mathbf{y} for question \mathbf{x} . In this
119 work, we mainly focus on *outcome reward* (Cobbe et al., 2021), which is often a sequence classifier,
120 offering an overall score for the entire response sequence. Our proposed method can also generalize
121 to process reward (Lightman et al., 2023).

122 **Reasoning as Decision Making.** Reasoning with LLMs can be framed as a search algorithm
123 that aims to identify a high-rewarding response: $\arg \min_{\mathbf{y} \in \mathcal{V}^*} -R(\mathbf{y}|\mathbf{x})$. Due to the combinatorial
124 nature of this optimization, directly finding the optimal solution is intractable, as the search space
125 grows exponentially with the sequence length, i.e., $|\mathcal{V}|^{|\mathbf{y}|}$. In autoregressive decoding, this challenge
126 reduces to making sequential decisions for the next token, each of which must ultimately contribute
127 to minimizing $-R(\mathbf{y}|\mathbf{x})$. This decision process is often formalized via a Bellman equation:
128

$$\pi_{LLM}^*(\cdot|\mathbf{y}_{\leq i-1}, \mathbf{x}) = \arg \max_{\mathbf{y}_i \in \mathcal{V}} \mathbb{E}_{\mathbf{y}_{\geq i+1} \sim \pi_{LLM}^*(\cdot|\mathbf{y}_i, \mathbf{y}_{\leq i-1}, \mathbf{x})} [r(\mathbf{y}_{\leq i-1}, \mathbf{y}_i, \mathbf{y}_{\geq i+1}|\mathbf{x})], \quad (1)$$

131 where π_{LLM}^* is a refined version of the original policy π_{LLM} . The expected reward-to-go in Eq.
132 1 is also known as the Q-function. The recursive structure of this formulation implies that greedy
133 decoding is not globally optimal, and identifying the optimal next-token prediction inherently requires
134 look-ahead rollouts and backtracking (Yao et al., 2024; Hao et al., 2023; Besta et al., 2024).
135

136 **Existing Approaches.** Current techniques tackling LLM reasoning via decision making can be
137 broadly categorized into training-time and inference-time methods. Training-time approaches include
138 supervised fine-tuning (SFT) as well as model-free policy optimization techniques, such as Schulman
139 et al. (2017); Guo et al. (2025); Rafailov et al. (2024). Our focus is on *inference-time methods*, which
140 aim to improve the decoding process of an LLM without additional training. These methods are
141 typically model-based and value-based, seeking to refine an existing policy by directly solving the
142 Bellman equation. For example, Best-of-N (BoN) (Stiennon et al., 2020) tackles Eq. 1 by sampling N
143 independent full trajectories from a base policy $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)} \sim \pi_{LLM}(\cdot|\mathbf{x})$, and selecting the one
144 with the highest reward $\mathbf{y}^* = \arg \max_{\mathbf{y} \in \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}\}} r(\mathbf{y}|\mathbf{x})$. More structured approaches, such
145 as Tree-of-Thoughts (ToT) (Yao et al., 2024) and Reasoning-as-Planning (RAP) (Hao et al., 2023),
146 explore the solution space and approximate Q-functions stochastically on the fly through rollouts and
147 recursive evaluation.

3 REASONING WITH GRADIENT-DRIVEN DECODING

148 **Overview.** In this section, we introduce ∇ -Reasoner, a novel reasoning algorithm that scales
149 inference-time computation by performing gradient descent in the sample space to refine the outputs
150 of a base policy. The overall pipeline, as illustrated in Fig. 2, is structured as an iterative decoding
151 process. Given a prefix \mathbf{x} , the model first generates an initial response $\mathbf{y}^{(0)}$. ∇ -Reasoner then
152 represents the generated sequence through its per-token pre-softmax logits $\mathbf{z}^{(0)}$ and optimizes
153 these logits via gradient descent to maximize the sequence-level reward $r(\mathbf{y}|\mathbf{x})$ (Sec. 3.1). After
154 optimization, ∇ -Reasoner resamples the *first token* of the generated sequence using the fine-tuned
155 logits $\tilde{\mathbf{z}}_1$. If the resampled token differs from the original, the subsequent tokens are regenerated, and
156 this candidate token is accepted only if its yielded response achieves a higher reward under $r(\cdot|\mathbf{x})$
157 (Sec. 3.2). The procedure then proceeds to the next token by incorporating the first generated token
158 into the prefix, and repeating this optimization-and-resampling loop. ∇ -Reasoner scales inference-
159 time reasoning by allocating additional computation to optimize the policy’s outputs via iterative
160 gradient descent. To further improve efficiency, we propose a series of system co-design strategies
161

162

163

164

165

166

167

168

169

170

171

172

173

174

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

Algorithm 1 ∇ -Reasoner: Decoding with DTO

Require: Prompt \mathbf{x} , language model π_{LLM} , reward model r , stop criteria $\text{StopCriteria}(\cdot)$.

 1: **repeat**

 2: $\mathbf{y}, \mathbf{z} \sim \pi_{LLM}(\cdot | \mathbf{x})$

 3: $\tilde{\mathbf{z}} \leftarrow \text{DTO}(\mathbf{x}, \mathbf{z}, \pi_{LLM}, r)$.

 4: $\tilde{\mathbf{y}}_1 \sim \text{softmax}(\tilde{\mathbf{z}}_1 / \tau)$.

 5: **if** $\tilde{\mathbf{y}}_1 \neq \mathbf{y}_1$ **then**

 6: $\tilde{\mathbf{y}}, \tilde{\mathbf{z}} \sim \pi_{LLM}(\cdot | \tilde{\mathbf{y}}_1, \mathbf{x})$

 7: **if** $r(\tilde{\mathbf{y}}, \tilde{\mathbf{y}}_1 | \mathbf{x}) > r(\mathbf{y} | \mathbf{x})$ **then**

 8: $\mathbf{x} \leftarrow \text{concat}[\mathbf{x}, \tilde{\mathbf{y}}_1]$

 9: **else**

 10: $\mathbf{x} \leftarrow \text{concat}[\mathbf{x}, \mathbf{y}_1]$

 11: **end if**

 12: **else**

 13: $\mathbf{x} \leftarrow \text{concat}[\mathbf{x}, \mathbf{y}_1]$

 14: **end if**

 15: **until** $\text{StopCriteria}(\mathbf{x})$

 16: **return** \mathbf{x}

Algorithm 2 Differentiable Textual Optimization (DTO)

Require: Prefix \mathbf{x} , initial logits \mathbf{z} , language model π_{LLM} , reward model r , and the number of training steps T .

 1: $\mathbf{z}^{(1)} \leftarrow \mathbf{z}$

 2: **for** $t = 1, \dots, T$ **do**

 3: **for** every $i = 1, \dots, |\mathbf{y}|$ **do**

 4: $j^* \leftarrow \arg \max_{j \in [|\mathcal{V}|]} \mathbf{z}_{ij}^{(t)}$

 5: $\mathbf{y}_i^{(t)} \leftarrow \delta_{j^*} + \text{softmax}(\mathbf{z}_i^{(t)} / \tau) - \text{StopGrad}(\text{softmax}(\mathbf{z}_i^{(t)} / \tau))$

 6: **end for**

 7: $\mathcal{L}_{\text{nll}} = -\sum_i \log \pi_{LLM}(\mathbf{y}^{(t)} | \mathbf{y}_{\leq i-1}^{(t)}, \mathbf{x})$

 8: $\mathcal{L}_{\text{reward}} = -r(\mathbf{y}^{(t)} | \mathbf{x})$.

 9: $\mathcal{L} = \mathcal{L}_{\text{nll}} + \lambda \mathcal{L}_{\text{reward}}$.

 10: $\mathbf{z}^{(t+1)} \leftarrow \mathbf{z}^{(t)} - \eta \nabla_{\mathbf{z}} \mathcal{L}$.

 11: **end for**

 12: **return** $\mathbf{z}^{(T)}$

Figure 2: Basic implementation of ∇ -Reasoner. ∇ -Reasoner is an iterative decoding algorithm driven by DTO. At each decoding step, DTO applies gradient descent on the logits initialized from the base model to optimize a reward-informed loss to refine the policy. The updated policy is then combined with rejection sampling, leading to high-reward responses. The pseudocode for the full implementation with acceleration techniques (Sec. 3.3) is deferred to Appendix B.

that selectively skip tokens unlikely to benefit from optimization and reuse model outputs and KV caches to accelerate decoding (Sec. 3.3).

3.1 DIFFERENTIABLE TEXTUAL OPTIMIZATION

The core step of our algorithm is leveraging gradient information to refine an initial response generated by the base policy. Existing reward-guided decoding methods (Wei et al., 2022b; Wang et al., 2022; Yao et al., 2024; Hao et al., 2023) can be regarded as zeroth-order approaches, as they rely solely on reward values. However, reward feedback is often sparse, and searching for improved solutions based only on scalar reward values can be sample-inefficient, particularly when the base policy is weak. We note that most reward models are inherently differentiable, as they are typically implemented with transformer-based sequence classifiers (Stiennon et al., 2020; Ouyang et al., 2022; Dong et al., 2024). This opens the door to exploiting not just reward values but also reward gradients, which provide richer directional information to guide samples toward high-reward regions. Motivated by this, we reformulate the search problem in Eq. 1 as a gradient-based differentiable optimization. We term this approach *Differentiable Textual Optimization (DTO)*, which differentiates reward over token space for progressive response improvement.

Objective. Our overall goal is to refine a given sequence of tokens $\mathbf{y}^{(0)}$ so as to maximize the reward. However, directly maximizing $r(\mathbf{y} | \mathbf{x})$ risks *reward hacking* (Pan et al., 2022), as the optimization trajectory of \mathbf{y} may drift away from the distribution under which $r(\mathbf{y} | \mathbf{x})$ is well-calibrated – typically the prior distribution induced by π_{LLM} . To mitigate this, we constrain \mathbf{y} to remain within the high-density region of π_{LLM} . Concretely, we regularize the log-likelihood of \mathbf{y} , thereby penalizing deviations from the distribution represented by the language model. The resulting objective function to be minimized is given by:

$$\mathcal{L}(\mathbf{y}) := -\lambda r(\mathbf{y} | \mathbf{x}) - \log \pi_{LLM}(\mathbf{y} | \mathbf{x}), \quad (2)$$

where $\lambda > 0$ is a hyper-parameter to balance the reward value and the regularization term. Intuitively, Eq. 2 seeks a response \mathbf{y} that not only achieves a high reward but also maintains fluency and faithfulness in natural language (Kumar et al., 2021; Qin et al., 2022; Yuan et al., 2025). To estimate the log-likelihood $\log \pi_{LLM}(\mathbf{y} | \mathbf{x})$, we decompose it sequentially from left to right, which results in the next-token prediction loss: $\log \pi_{LLM}(\mathbf{y} | \mathbf{x}) = \sum_{i=1}^{|\mathbf{y}|} \mathbf{y}_i^\top \log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq i-1}, \mathbf{x}))$.

216 **Parameterization.** The token space of \mathbf{y} is a discrete where gradients cannot directly operate. Therefore, we propose to parameterize the tokens via the underlying logit vectors used to sample them. At the initialization stage, we use the LLM-generated logits to initialize $\mathbf{z}^{(0)} \in \mathbb{R}^{|\mathbf{y}^{(0)}| \times |\mathcal{V}|}$. During optimization, we use [Gumbel-softmax straight-through](#) trick to parameterize $\mathbf{y}_i^{(t)} = \delta_{\arg \max_{j \in |\mathcal{V}|} \mathbf{z}_{ij}^{(t)}} + \text{softmax}(\mathbf{z}_i^{(t)} / \tau) - \text{StopGrad}(\text{softmax}(\mathbf{z}_i^{(t)} / \tau))$ (Bengio et al., 2013; Jang et al., 2016), where δ_i denotes the i -th canonical basis and $\tau > 0$ is the temperature coefficient. By this means, gradient descent can be equivalently performed on the space of $\mathbf{z}^{(t)}$ as: $\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} - \eta \nabla_{\mathbf{z}} \mathcal{L}(\mathbf{z}^{(t)})$.

225 As we will demonstrate in Sec. 4, the gradient of \mathcal{L} propagates information bidirectionally. Preceding
 226 tokens act as a regularizer on successive tokens, enforcing consistency with the autoregressive
 227 generation process, while trailing tokens propagate outcome-level reward signals and full-sequence
 228 alignment back to earlier tokens through attention. This implements a closed-loop control of earlier
 229 predictions influencing subsequent decoding steps, thereby capturing the key recursive structure of
 230 reasoning characterized in Eq. 1. Furthermore, in Sec. 4, we establish a close connection between
 231 DTO, which optimizes directly in the sample space, and policy-optimization (e.g., PPO) that operate
 232 in the parameter space. We show that DTO provably shifts the drawn samples toward the reward-
 233 maximizing distribution induced by the original policy.

235 3.2 ITERATIVE DECODING WITH DTO

237 In this section, we elaborate on the detailed iterative generation process with DTO integrated for
 238 policy improvement. Akin to autoregressive decoding, ∇ -Reasoner generates the full response token
 239 by token. The sampling of each token consists of the following two steps:

241 **Policy Improvement via DTO.** Starting from a prefix $\mathbf{x} \in \mathcal{V}^*$, we let the LLM π_{LLM} generate a
 242 continuation sequence $\mathbf{y}^{(0)}$ along with its pre-softmax logits $\mathbf{z}^{(0)}$. We then apply the DTO algorithm
 243 to optimize $\mathbf{z}^{(0)}$ for T steps, yielding refined logits $\tilde{\mathbf{z}}$. The logits corresponding to the first token are
 244 treated as the improved policy for predicting the immediate next token, intentionally adjusted to yield
 245 higher reward when used to generate the continuing responses. Accordingly, we resample the next
 246 token from this updated policy: $\tilde{\mathbf{y}}_1 \sim \text{softmax}(\tilde{\mathbf{z}}_1 / \tau)$.

248 **Rejection Sampling.** Once a new next-token candidate $\tilde{\mathbf{y}}_1$ is obtained, we first compare it with
 249 the initial prediction \mathbf{y}_1 . If $\tilde{\mathbf{y}}_1 = \mathbf{y}_1$, no effective policy update occurs, and we proceed directly to
 250 the next token for policy refinement. If $\tilde{\mathbf{y}}_1 \neq \mathbf{y}_1$, we perform an additional rollout conditioned on
 251 $\tilde{\mathbf{y}}_1$ as the next token, yielding a new response $\tilde{\mathbf{y}}$. Both \mathbf{y} and $\tilde{\mathbf{y}}$ are then evaluated under the reward
 252 function, and the token that yields a full response with the higher reward is retained.

254 **Test-Time Scaling.** We scale computation in ∇ -Reasoner along two axes: (1) increasing the number
 255 of gradient update steps used by DTO to refine the policy, and (2) performing rejection sampling
 256 among rollouts yielded by the original and updated policy. Comparatively, allocating additional
 257 compute to gradient-based optimization is not only more effective in incorporating reward signals
 258 into the policy, but also more efficient than purely autoregressive decoding. This efficiency arises
 259 because computing the full-sequence gradient $\nabla_{\mathbf{z}} \mathcal{L}$ leverages the parallel execution of transformers:
 260 a single gradient step propagates updates across all tokens within one model call, whereas a standard
 261 autoregression generates only a single token per model call. As we will show in Sec. 5.4, sampling
 262 from the policy refined by DTO yields a significantly higher chance of reward improvement.

264 3.3 ACCELERATING ∇ -REASONER

266 The naive implementation of ∇ -Reasoner is inefficient due to two primary bottlenecks: (1) decoding
 267 each token requires a full optimization procedure, where each step involves backpropagation through
 268 two large models; and (2) generating a single token requires an additional full rollout. In this section,
 269 we demonstrate that ∇ -Reasoner is amenable to several strategies that significantly accelerate both
 optimization and generation, while adaptively allocating compute to the tokens that matter most.

270 **Gradient Caching.** The gradient backpropagated to the logits \mathbf{z} can be decomposed via the chain
 271 rule as: $\nabla_{\mathbf{z}} \mathcal{L} = \frac{\partial \mathbf{z}}{\partial \mathbf{y}} \frac{\partial \mathcal{L}}{\partial \mathbf{y}}$, wherein the term $\frac{\partial \mathcal{L}}{\partial \mathbf{y}}$ dominates the computational cost, since it requires a
 272 full forward and backward pass through both the language model and the reward model. However,
 273 we observe that \mathbf{y} – the one-hot vectors indicating the maximal entries in the soft logits \mathbf{z} – changes
 274 infrequently as optimization proceeds. Exploiting this property, we cache the gradient $\frac{\partial \mathcal{L}}{\partial \mathbf{y}}$ once
 275 computed, and reuse it until the maximal entries of \mathbf{z} flip. In our implementation, we retain the
 276 cached gradient $\mathbf{g}_i = \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i}$ for every $i \in [|\mathbf{y}|]$ whenever \mathbf{y} is updated, and otherwise recover it
 277 efficiently using the surrogate loss $\mathcal{L}_{cache} = \sum_{i=1}^{|\mathbf{y}|} \mathbf{y}_i^\top \mathbf{g}_i$ to recover the saved gradients $\{\mathbf{g}_i\}_{i \in [|\mathbf{y}|]}$
 278 whenever \mathbf{y} remains unchanged from the previous iteration. See Algorithm 4 in Appendix B.
 279

280 **Rollout Trajectory Reusing.** We further note that the rollout strategy can be improved to reduce
 281 unnecessary computation and better leverage the KV cache. In the naive implementation (Sec. 3.2 or
 282 Algorithm 1), ∇ -Reasoner generates a sequential rollout to optimize the next-token prediction policy
 283 for every decoding step. However, the rollout trajectory, including both tokens and logits, continuing
 284 from the previously accepted token can be directly reused as the rollout for the subsequent token. In
 285 Algorithm 1, we skip the rollout at the beginning of each step and reuse $\mathbf{y}_{\geq 2}$ and $\mathbf{z}_{\geq 2}$ as the rollout
 286 for the next token if the resampled token $\tilde{\mathbf{y}}$ is rejected; otherwise, we continue with $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{z}}$ for the
 287 next step. We also limit the total number of rollouts as N_{max} . Once the total number of rollouts
 288 exceeds the N_{max} , we terminate the iterative policy refinement and generate the remaining tokens
 289 using standard autoregressive decoding. See more details in Algorithm 3 in Appendix B.

290 **Confidence- and Gradient-Guided Token Selection.** Running DTO to optimize the policy at
 291 every decoding step can result in redundant computation. We observe that token logits with either
 292 high confidence (see Appendix C.3) or small gradients are unlikely to be modified under DTO. To
 293 address this, we introduce two selection criteria, *entropy-based* and *gradient-based* to determine
 294 which tokens should undergo policy refinement. Specifically, we define two hyperparameters, ϵ_{ent}
 295 and ϵ_{grad} . DTO is applied only when the entropy of the token logits satisfies $H(\mathbf{z}_1) > \epsilon_{ent}$ and
 296 the gradient magnitude exceeds $\|\nabla_{\mathbf{z}_1} \mathcal{L}\|_2 > \epsilon_{grad}$, where $H(\cdot)$ denotes the entropy of a categorical
 297 distribution. We refer readers to Algorithm 4 in Appendix B for more details.

4 THEORETICAL ANALYSIS

301 **Interpretation of Gradient Updates.** We analyze the gradient of \mathcal{L} to reveal how DTO updates
 302 the response. The derivatives in terms of the l -th token $\partial \mathcal{L} / \partial \mathbf{x}_l$ under loss Eq. 2 can be decomposed
 303 as follows:

$$304 \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i} = - \underbrace{\log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq i-1}, \mathbf{x}))}_{\delta_{prefix}} - \underbrace{\sum_{j=i+1}^{|\mathbf{y}|} \frac{\partial \log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq j-1}, \mathbf{x}))}{\partial \mathbf{y}_i} \mathbf{x}}_{\delta_{postfix}} - \lambda \underbrace{\frac{\partial r(\mathbf{y} | \mathbf{x})}{\partial \mathbf{y}_i}}_{\delta_{reward}}.$$

309 We defer the derivation to Appendix C.1. The first term, δ_{prefix} , updates the token \mathbf{x}_i based on its
 310 preceding context, aligning the next-token policy with the autoregressive prediction probabilities
 311 produced by the language model. The second term, $\delta_{postfix}$, propagates information from subsequent
 312 tokens through the attention mechanism, encouraging global consistency with respect to its future
 313 context. Finally, the reward gradient δ_{reward} provides a sequence-level signal, transmitting information
 314 from later tokens to \mathbf{y}_i via attention. As highlighted by Bachmann & Nagarajan (2024), the order of
 315 generation plays a crucial role in complex reasoning or algorithmic tasks. Pure left-to-right generation
 316 can fall short of *error accumulation*, making it insufficient for intricate logical reasoning processes.
 317 An ideal decoding method for reasoning should allow for iterative refinement of the reasoning chain
 318 in both forward and backward directions (Yao et al., 2024; Hao et al., 2023).

319 **Inference-Time Gradient Descent is “Deamortized” PPO.** We theoretically establish the connec-
 320 tion between the test-time textual optimization and parametric RL-based training. RLHF (Schulman
 321 et al., 2017; Ouyang et al., 2022) and RLVR (Guo et al., 2025; Shao et al., 2024) have been demon-
 322 strated to be particularly effective for mathematical reasoning tasks (Wang et al., 2023a; Zhao et al.,
 323 2023; Dong et al., 2024; Shao et al., 2024). The primary objective of RL is to fine-tune a pre-trained
 LLM using RL algorithms, ensuring that its responses to specific prompts maximize a given reward

function. Among various RL algorithms, KL-regularized approaches, such as Proximal Policy Optimization (PPO) (Schulman et al., 2017), have been widely adopted in practice. In this section, we uncover the hidden connection between PPO and our proposed DTO. Formally, let $\rho : \mathcal{V}^* \rightarrow \mathbb{R}$ represent an LLM to be aligned with the reward function r , initialized from the pre-trained policy π_{LLM} . PPO optimizes for ρ by minimizing the following functional objective defined over the space of distributions:

$$\mathcal{L}_{PPO}(\rho) := -\mathbb{E}_{\mathbf{y} \sim \rho}[\lambda r(\mathbf{y})] + D_{KL}(\rho \parallel \pi_{LLM}), \quad (3)$$

where the first expectation term estimates the expected reward, while in the second term, the KL-divergence regularizes the distributional discrepancy between ρ and π_{LLM} . Assuming LLMs’ input domain can be extended to the ambient space beyond discrete vocabularies, then we can show the relation between (stochastic) gradient flow of Eq. 2 and functional solution to PPO:

Theorem 4.1. Suppose $\{\rho^t\}_{t \geq 0}$ denotes the Wasserstein gradient flow minimizing Eq. 3 in the distribution space with boundary conditions $\rho^0 = \pi_{LLM}$ and $\rho^\infty = \rho^* = \arg \min_\rho \mathcal{L}_{PPO}(\rho)$. Then we can draw samples from ρ^* by first initializing $\mathbf{x}^0 \sim \pi_{LLM}$ and simulating a trajectory $\{\mathbf{x}^t\}_{t \geq 0}$ following the stochastic gradient flow of Eq. 2: $\frac{d\mathbf{x}^t}{dt} = -\nabla \mathcal{L}(\mathbf{x}^t) + \sqrt{2}\epsilon_t$, where $\{\epsilon_t \in \mathcal{N}(\mathbf{0}, \mathbf{I})\}_{t \geq 0}$ are Brownian motions.

Theorem 4.1 is proved in Appendix C.2. Theorem 4.1 shows that instead of optimizing the entire policy to satisfy the reward function w.r.t. Eq. 3, there exists a trajectory driven by the gradients of Eq. 2 on the sample space that can directly generate samples from the optimal distribution minimizing Eq. 3. Pre-training scaling and test-time scaling can be unified and interpreted through Theorem 4.1 as two complementary forms of statistical inference: parametric and non-parametric (particle-based) inference (Liu & Wang, 2016; Chen et al., 2018). The pre-training stage corresponds to parametric inference: a global parameter is optimized to minimize the overall loss across a dataset, amortizing the cost of individual samples into a shared parameter. Increasing the size of this parameter space enhances the model’s representational capacity, thereby reducing the average cost per sample. In contrast, test-time scaling via DTO is analogous to non-parametric inference, which performs optimization in the sample space, treating each sample as an independent “particle” that minimizes its own cost. This allows for fine-grained adaptation at the individual sample level. The Wasserstein gradient flow provides a mathematical framework to describe the relationship between the dynamics of measures (global distributions) and individual samples, thereby bridging the conceptual gap between pre-training scaling and test-time scaling.

5 EXPERIMENTS

5.1 RESULTS ON MATH REASONING

Experiment Details. Tab. 1 compares our test-time method, ∇ -Reasoner, against a variety of baselines. We benchmark its performance against other test-time approaches, including greedy decoding, Self-Consistency (SC) (Xie et al., 2024), Best-of-N (BoN) (Stiennon et al., 2020), tree-search based methods: Tree-of-Thought (ToT) (Yao et al., 2024) and Reasoning via Planning (RAP) (Hao et al., 2023), and the iterative refinement approach: TPO (Li et al., 2025). Additionally, we include training-based methods such as Supervised Fine-Tuning (SFT) and GRPO (Guo et al., 2025) for a comprehensive comparison. We evaluate two model families, Qwen-2.5-math (Yang et al., 2024a) and Llama-3.1 (Grattafiori et al., 2024), on four representative mathematical reasoning benchmarks: MATH-500 (Hendrycks et al., 2021), AIME24, AIME25, and AMC (Li et al., 2024). **We leverage reward models from the Skywork-V2 family (Liu et al., 2025): for Skywork-V2-Qwen-4B for Qwen-based models and Skywork-V2-Llama-8B for Llama family models.** For BoN and SC, we let $N = 8$ to match $N_{max} = 8$ used in our methods. For TPO, we set the number of samples per step as $N_{samples} = 2$ and the number of refinement steps as $N_{refine} = 2$. For ToT and RAP, we adopt the default hyperparameters in Hao et al. (2024) to yield meaningful results. Please refer to Appendix D for more experimental details.

Performance Comparison. Our method shows superior performance across all models and benchmarks on test-time methods. We even reach comparable performance with training-based methods. Specifically, with the Qwen-2.5-7B base model, our approach achieves the highest scores among

378
379
380
381
382
Table 1: Accuracy (%) on math reasoning datasets compared with baseline methods, including
both test-time and training-time approaches. We skip results on AIME datasets for Llama-3.1-8B
as it is incapable of generating reasonable performance. We mark the best performer in **bold** and
the runner-up with underline. Our method outperforms all test-time baselines and even achieves
performance on par with the training-based methods (SFT and GRPO), respectively.

Models	Methods	MATH-500	AMC	AIME24	AIME25
Qwen-2.5-7B	Greedy decoding	43.8	33.0	6.7	6.7
	SC (Xie et al., 2024) ($N = 8$)	69.8	49.4	22.5	20.0
	BoN (Stiennon et al., 2020) ($N = 8$)	70.2	50.1	22.5	13.3
	ToT (Yao et al., 2024)	57.8	42.4	6.7	10.0
	RAP (Hao et al., 2023)	68.6	50.1	18.3	14.2
	SFT (Ouyang et al., 2022)	65.8	36.4	6.3	11.7
	GRPO (Guo et al., 2025)	70.8	52.8	20.8	16.7
Qwen-2.5-7B-Instruct	∇ -Reasoner ($N_{max} = 8$)	71.0	<u>51.5</u>	23.3	<u>15.0</u>
	Greedy decoding	71.2	43.0	5.3	7.5
	SC (Xie et al., 2024) ($N = 8$)	76.6	55.5	25.0	22.5
	BoN (Stiennon et al., 2020) ($N = 8$)	77.8	55.9	22.5	18.3
	ToT (Yao et al., 2024)	75.4	48.2	20.0	18.3
	RAP (Hao et al., 2023)	80.2	54.6	1.6	12.5
	TPO (Li et al., 2025)	77.6	55.9	6.7	11.1
Llama-3.1-8B-Instruct	∇ -Reasoner ($N_{max} = 8$)	80.4	56.8	26.6	<u>20.0</u>
	Greedy decoding	40.6	19.3	-	-
	SC (Xie et al., 2024) ($N = 8$)	54.8	25.7	-	-
	BoN (Stiennon et al., 2020) ($N = 8$)	52.2	26.1	-	-
	ToT (Yao et al., 2024)	50.2	25.6	-	-
	RAP (Hao et al., 2023)	55.4	25.8	-	-
	SFT (Ouyang et al., 2022)	46.6	20.2	-	-
∇ -Reasoner ($N_{max} = 8$)	55.8	28.9	-	-	

404
405
406
407
408
409
410
test-time methods on MATH-500 (71.0%) and AIME24 (23.3%), and remains highly competitive
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
989
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1287
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1297
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1939
1940
1941
19

complexity under idealized system-level and hardware optimizations, while mitigating discrepancies arising from implementation details (e.g., compatibility with serving engines).

Fig. 3 compares the computational cost of our method with several baselines. For fair comparison, we set the number of generated samples to 8 for all approaches ($N = 8$ for SC and BoN, and $N_{max} = 8$ for our ∇ -Reasoner). Our method delivers superior performance at a significantly lower cost than SC and BoN. For instruction-tuned models, it reduces the number of model calls by up to 40.2%, while for base models, it outperforms all baselines using only about 90% of this metric. The reason for the reduced cost is twofold: (1) with confidence- and gradient-guided selection, rollouts usually start from the middle of the sequence, instead of from the beginning as BoN and SC do; (2) the optimization cost with gradient caching remains lightweight while our DTO enables efficient parallelizable execution of transformers and revision of tokens. These results imply that ∇ -Reasoner has stronger prospects for achieving inference efficiency. Unlike trial-and-error-based test-time scaling methods (e.g., BoN), which repeatedly resample outputs without guidance, our method updates the decoding policy in a targeted and strategic fashion.

In the meantime, we measure and compare the wallclock running time of our methods with others on AMC and AIME dataset in Tab. 2. The experiment was conducted on the AIME-25 dataset using the Qwen-2.5-math-Instruct-7B

model running on eight 80GB NVIDIA A100 GPUs. Note that to mitigate the implementation discrepancies, we do not exploit dynamic serving engines for LLM inference. Even though our approach currently demonstrates a similar running time to BoN, we emphasize that ∇ -Reasoner has greater potential to leverage compute more effectively by executing transformers in parallel mode. While BoN can potentially benefit from highly optimized serving engines, we speculate that integrating the optimization procedure of ∇ -Reasoner better with the generation pipeline would create a much larger efficiency gap between our method and inference-only approaches.

5.3 TEST-TIME SCALING LAW

Table 2: Wall-clock time measurements.

Method	BoN ($N = 8$)	SC ($N = 8$)	Ours ($N_{max} = 8$)
AMC	21.9 s	23.5 s	23.6 s
AIME	40.3 s	39.5 s	41.1 s

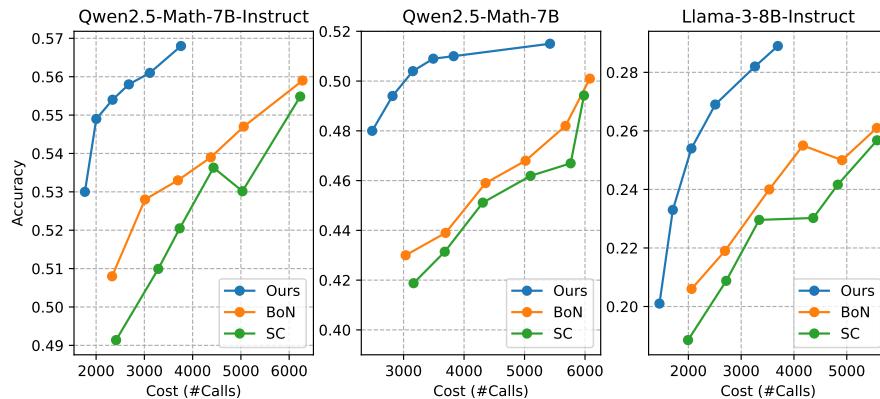


Figure 4: Test-time scaling curves comparing our method with BoN and SC. We change the number of samples N for BoN and SC and number of rollouts N_{max} for our method. The results show ∇ -Reasoner achieves superior performance with reduced cost across multiple models.

We present our test-time scaling curves in Fig. 4, comparing our method against Best-of-N (BoN) and Self-Consistency (SC). The figure plots accuracy against computational cost, showing how each method’s performance scales as more resources are used. As is evident across all models, our method’s curve consistently lies above the baselines. This indicates that for any given computational budget (number of calls), our approach achieves a higher accuracy. These results demonstrate that ∇ -Reasoner offers a superior trade-off between performance and computational cost, establishing a more efficient frontier compared to these sample-heavy techniques.

5.4 ALGORITHM ANALYSIS

486 **Dependencies on Reward Models.** Our
 487 approach relies on the gradient signal from
 488 the reward model to optimize policy at test
 489 time. To study the dependency on the qual-
 490 ity of reward models, we further evaluate

491 our approach on Qwen2.5-Math-7B-Instruct paired with the larger Skywork-Reward-V2-Qwen3-8B
 492 reward model. We note that our original choice Skywork-Reward-V2-Qwen3-4B is a smaller and
 493 weaker reward model according to the RewardBench (Malik et al., 2025). According to the Tab. 3,
 494 the performance gap between the 4B and 8B variants remains consistently below 1 point across both
 495 MATH-500 and AMC. This indicates that using a smaller reward model does not lead to significant
 496 performance degradation compared with the larger, stronger version. This justifies our original choice
 497 (in Tab. 1) and further suggests that smaller reward models may even be preferable for improving
 498 efficiency.

499
 500 Table 4: Analysis of rejection rate (%) in rejection
 501 sampling. We set $N = 8$ for the BoN baseline
 502 and also set $N_{max} = 8$ for our ∇ -Reasoner. The
 503 theoretical rejection rate of the baseline is 66.0%.

Model	Baseline	∇ -Reasoner
Qwen-2.5	65.9	32.8
Qwen-2.5-Instruct	66.5	28.9
Llama-3-Instruct	66.9	40.1

504 To this end, we measure the *rejection rate*, defined as the percentage of candidates produced by
 505 DTO that are rejected for failing to improve the reward. For comparison, we also evaluate this
 506 metric on a baseline that performs rejection sampling without DTO, which is equivalent to BoN.
 507 Theoretically, performing rejection sampling N times over an identical distribution yields a rejection
 508 rate of $1 - (\sum_{k=1}^N 1/k)/N$ that converges to one as $N \rightarrow \infty$. For $N = 8$, the expected rejection rate
 509 is approximately 66.0%. We report our measured rejection rate in Tab. 4. We report the empirical
 510 rejection rates in Tab. 4. The results show that rejection sampling without DTO closely matches the
 511 theoretical prediction, while rejection sampling with DTO significantly reduces the rejection rate (by
 512 up to 30%). This confirms that DTO is effective in improving the next-token policy, producing tokens
 513 that lead to continuations with higher rewards.

521 6 CONCLUSION AND LIMITATIONS

522 We presented ∇ -Reasoner, an inference-time reasoning framework that introduces Differentiable
 523 Textual Optimization (DTO) to refine token logits via gradient-based optimization. By combining
 524 gradient signals from both the LLM likelihood and a reward model, ∇ -Reasoner enables more effec-
 525 tive policy improvement than zeroth-order search methods, while incorporating rejection sampling
 526 and speedup techniques to boost effectiveness and efficiency. Theoretically, we show that aligning
 527 with a reward function is equivalent to gradient-based optimization in the sample space. ∇ -Reasoner
 528 delivers substantial performance gains over base models while consistently reducing computation
 529 cost, illustrating a sharper and more efficient scaling paradigm for LLM reasoning.

530 **Limitations.** In line with previous observations (Yue et al.), the performance of ∇ -Reasoner appears
 531 to remain bounded by the capabilities of the underlying base model and reward model, particularly
 532 under limited computation. **The base and reward models are required to share the same vocabulary**
 533 **to allow for end-to-end logit optimization.** Moreover, integrating ∇ -Reasoner into efficient LLM
 534 serving pipelines requires more careful system co-design to incorporate test-time gradient descent.

535 ETHICS STATEMENT

536 This research primarily concentrates on developing inference algorithms to enhance reasoning in
 537 large language models (LLMs), with a particular emphasis on mathematical reasoning. It relies solely

538 Table 3: Ablation study on reward model choice.

Model	Skywork-Qwen3-4B	Skywork-Qwen3-8B
MATH-500	80.4	80.8 (+0.4)
AMC	56.8	57.1 (+0.3)

539 **Rejection Rate Analysis.** As described in
 540 Sec. 3.2 and Algorithm 1, ∇ -Reasoner first ap-
 541 plies DTO to directly optimize the policy in the
 542 logit space, and then compares a continuation \tilde{y}
 543 generated from a resampled token with the origi-
 544 nal rollout sequence y . The token sampled from
 545 the optimized policy is adopted only if it yields
 546 a continuation with a higher reward. Henceforth,
 547 it becomes essential to quantify the acceptance
 548 rate of tokens drawn from the optimized policy
 549 in order to justify the effectiveness of DTO.

540 on extant LLMs and does not involve training, fine-tuning model weights, or creating new LLMs. As
 541 a result, the work does not raise any novel domain-specific ethical considerations or societal impacts
 542 beyond those already well-documented in relation to large-scale language models more broadly.
 543

544 **REPRODUCIBILITY STATEMENT**

545 We include pseudocode plus a detailed version in both the main text and Appendix B. Complete
 546 derivations and proofs are provided in Appendix C. Additionally, Appendix D contains the full list of
 547 hyperparameters, datasets, and model checkpoints required to reproduce our experimental results.
 548

549 **THE USE OF LARGE LANGUAGE MODELS**

550 Large language models are used solely for sentence-level proofreading. All research ideation and
 551 paper writing were originally carried out by the authors.
 552

553 **REFERENCES**

554 Gregor Bachmann and Vaishnav Nagarajan. The pitfalls of next-token prediction. *arXiv preprint*
 555 *arXiv:2403.06963*, 2024.

556 Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through
 557 stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

558 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawska, Lukas Gianinazzi,
 559 Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, et al. Graph of thoughts:
 560 Solving elaborate problems with large language models. In *Proceedings of the AAAI Conference*
 561 *on Artificial Intelligence*, volume 38, pp. 17682–17690, 2024.

562 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
 563 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
 564 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

565 Changyou Chen, Ruiyi Zhang, Wenlin Wang, Bai Li, and Liqun Chen. A unified particle-optimization
 566 framework for scalable bayesian sampling. *arXiv preprint arXiv:1805.11659*, 2018.

567 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
 568 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve
 569 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

570 Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen
 571 Sahoo, Caiming Xiong, and Tong Zhang. Rlhf workflow: From reward modeling to online rlhf.
 572 *arXiv preprint arXiv:2405.07863*, 2024.

573 Yilun Du, Shuang Li, Joshua Tenenbaum, and Igor Mordatch. Learning iterative reasoning through
 574 energy minimization. In *International Conference on Machine Learning*, pp. 5570–5582. PMLR,
 575 2022.

576 Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R Bartoldson,
 577 Bhavya Kaikhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent
 578 reasoning: A recurrent depth approach. *arXiv preprint arXiv:2502.05171*, 2025.

579 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 580 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
 581 models. *arXiv preprint arXiv:2407.21783*, 2024.

582 Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna
 583 Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, Ashima Suvarna, Benjamin Feuer, Liangyu
 584 Chen, Zaid Khan, Eric Frankel, Sachin Grover, Caroline Choi, Niklas Muennighoff, Shiye Su,
 585 Wanja Zhao, John Yang, Shreyas Pimpalaonkar, Kartik Sharma, Charlie Cheng-Jie Ji, Yichuan
 586 Deng, Sarah Pratt, Vivek Ramanujan, Jon Saad-Falcon, Jeffrey Li, Achal Dave, Alon Albalak,
 587 Kushal Arora, Blake Wulfe, Chinmay Hegde, Greg Durrett, Sewoong Oh, Mohit Bansal, Saadia
 588 Gabriel, Aditya Grover, Kai-Wei Chang, Vaishaal Shankar, Aaron Gokaslan, Mike A. Merrill,
 589

594 Tatsunori Hashimoto, Yejin Choi, Jenia Jitsev, Reinhard Heckel, Maheswaran Sathiamoorthy,
 595 Alexandros G. Dimakis, and Ludwig Schmidt. Openthoughts: Data recipes for reasoning models,
 596 2025. URL <https://arxiv.org/abs/2506.04178>.

597 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 598 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 599 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

600 Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu.
 601 Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*,
 602 2023.

603 Shibo Hao, Yi Gu, Haotian Luo, Tianyang Liu, Xiyan Shao, Xinyuan Wang, Shuhua Xie, Haodi Ma,
 604 Adithya Samavedhi, Qiyue Gao, et al. Llm reasoners: New evaluation, library, and analysis of
 605 step-by-step reasoning with large language models. *arXiv preprint arXiv:2404.05221*, 2024.

606 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 607 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv
 608 preprint arXiv:2103.03874*, 2021.

609 Cong Duy Vu Hoang, Gholamreza Haffari, and Trevor Cohn. Towards decoding as continuous
 610 optimization in neural machine translation. *arXiv preprint arXiv:1701.02854*, 2017.

611 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
 612 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al.
 613 Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

614 James Y Huang, Sailik Sengupta, Daniele Bonadiman, Yi-an Lai, Arshit Gupta, Nikolaos Pappas,
 615 Saab Mansour, Katrin Kirchhoff, and Dan Roth. Deal: Decoding-time alignment for large language
 616 models. *arXiv preprint arXiv:2402.06147*, 2024.

617 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv
 618 preprint arXiv:1611.01144*, 2016.

619 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child,
 620 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models.
 621 *arXiv preprint arXiv:2001.08361*, 2020.

622 Maxim Khanov, Jirayu Burapachaphee, and Yixuan Li. Args: Alignment as reward-guided search.
 623 *arXiv preprint arXiv:2402.01694*, 2024.

624 Sachin Kumar, Eric Malmi, Aliaksei Severyn, and Yulia Tsvetkov. Controlled text generation as
 625 continuous optimization with multiple constraints. *Advances in Neural Information Processing
 626 Systems*, 34:14542–14554, 2021.

627 Sachin Kumar, Biswajit Paria, and Yulia Tsvetkov. Gradient-based constrained sampling from
 628 language models. *arXiv preprint arXiv:2205.12558*, 2022.

629 Nayoung Lee, Kartik Sreenivasan, Jason D Lee, Kangwook Lee, and Dimitris Papailiopoulos.
 630 Teaching arithmetic to small transformers. *arXiv preprint arXiv:2307.03381*, 2023.

631 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt
 632 tuning. *arXiv preprint arXiv:2104.08691*, 2021.

633 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-
 634 masesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative
 635 reasoning problems with language models. *Advances in Neural Information Processing Systems*,
 636 35:3843–3857, 2022.

637 Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang,
 638 Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau,
 639 Guillaume Lample, and Stanislas Polu. Numinamath. [<https://huggingface.co/AI-MO/NuminaMath-CoT>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.

648 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv*
 649 *preprint arXiv:2101.00190*, 2021.

650

651 Yafu Li, Xuyang Hu, Xiaoye Qu, Linjie Li, and Yu Cheng. Test-time preference optimization:
 652 On-the-fly alignment via iterative textual feedback. *arXiv preprint arXiv:2501.12895*, 2025.

653 Yuhui Li, Fangyun Wei, Jinjing Zhao, Chao Zhang, and Hongyang Zhang. Rain: Your language
 654 models can align themselves without finetuning. *arXiv preprint arXiv:2309.07124*, 2023.

655

656 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 657 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint*
 658 *arXiv:2305.20050*, 2023.

659 Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone.
 660 Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint*
 661 *arXiv:2304.11477*, 2023a.

662

663 Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiacai Liu, Chaojie Wang, Rui Yan, Wei Shen,
 664 Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. Skywork-reward-v2: Scaling preference
 665 data curation via human-ai synergy. *arXiv preprint arXiv:2507.01352*, 2025.

666

667 Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference
 668 algorithm. *Advances in neural information processing systems*, 29, 2016.

669

670 Xin Liu, Muhammad Khalifa, and Lu Wang. Bolt: Fast energy-based controlled text generation with
 671 tunable biases. *arXiv preprint arXiv:2305.12018*, 2023b.

672

673 Saumya Malik, Valentina Pyatkin, Sander Land, Jacob Morrison, Noah A. Smith, Hannaneh Ha-
 jishirzi, and Nathan Lambert. Rewardbench 2: Advancing reward model evaluation. <https://huggingface.co/spaces/allenai/reward-bench>, 2025.

674

675 Dimitra Maoutsas, Sebastian Reich, and Manfred Opper. Interacting particle solutions of fokker-
 676 planck equations through gradient-log-density estimation. *Entropy*, 22(8):802, 2020.

677

678 Fatemehsadat Mireshghallah, Kartik Goyal, and Taylor Berg-Kirkpatrick. Mix and match: Learning-
 679 free controllable text generation using energy language models. In *Proceedings of the 60th Annual*
 680 *Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 401–415,
 681 2022.

682

683 Reichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher
 684 Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted
 685 question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.

686

687 Bernt Øksendal. *Stochastic differential equations*. Springer, 2003.

688

689 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 690 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 691 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
 692 27744, 2022.

693

694 Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping
 695 and mitigating misaligned models. *arXiv preprint arXiv:2201.03544*, 2022.

696

697 Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. Automatic prompt
 698 optimization with “gradient descent” and beam search. *arXiv preprint arXiv:2305.03495*, 2023.

699

700 Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. Cold decoding: Energy-based
 701 constrained text generation with langevin dynamics. *Advances in Neural Information Processing*
 702 *Systems*, 35:9538–9551, 2022.

703

704 Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language under-
 705 standing by generative pre-training. 2018.

706

707 Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language
 708 models are unsupervised multitask learners. 2019.

702 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 703 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
 704 *in Neural Information Processing Systems*, 36, 2024.

705 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 706 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

708 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 709 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 710 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

711 Weijia Shi, Xiaochuang Han, Hila Gonen, Ari Holtzman, Yulia Tsvetkov, and Luke Zettlemoyer.
 712 Toward human readable prompt tuning: Kubrick’s the shining is a good movie, and a good prompt
 713 too? *arXiv preprint arXiv:2212.10539*, 2022.

715 Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt:
 716 Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint*
 717 *arXiv:2010.15980*, 2020.

718 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 719 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

721 Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su.
 722 Llm-planner: Few-shot grounded planning for embodied agents with large language models. In
 723 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2998–3009, 2023.

724 Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,
 725 Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in*
 726 *Neural Information Processing Systems*, 33:3008–3021, 2020.

728 Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia
 729 Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and
 730 outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

731 Karthik Valmeeakam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao Kambham-
 732 pati. On the planning abilities of large language models (a critical investigation with a proposed
 733 benchmark). *arXiv preprint arXiv:2302.06706*, 2023.

734 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Y Wu, and Zhifang
 735 Sui. Math-shepherd: A label-free step-by-step verifier for llms in mathematical reasoning. *arXiv*
 736 *preprint arXiv:2312.08935*, 2023a.

738 Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Haotian Luo, Jiayou Zhang, Nebojsa Jojic, Eric P
 739 Xing, and Zhiting Hu. Promptagent: Strategic planning with language models enables expert-level
 740 prompt optimization. *arXiv preprint arXiv:2310.16427*, 2023b.

741 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 742 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 743 *arXiv preprint arXiv:2203.11171*, 2022.

745 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama,
 746 Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models.
 747 *arXiv preprint arXiv:2206.07682*, 2022a.

748 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 749 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 750 *neural information processing systems*, 35:24824–24837, 2022b.

751 Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein.
 752 Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery.
 753 *Advances in Neural Information Processing Systems*, 36, 2024.

755 Haoyi Wu, Zhihao Teng, and Kewei Tu. Parallel continuous chain-of-thought with jacobi iteration.
 756 *arXiv preprint arXiv:2506.18582*, 2025.

756 Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael
 757 Xie. Self-evaluation guided beam search for reasoning. *Advances in Neural Information Processing*
 758 *Systems*, 36, 2024.

759 Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang.
 760 Iterative preference learning from human feedback: Bridging theory and practice for rlhf under
 761 kl-constraint. In *Forty-first International Conference on Machine Learning*, 2024.

762 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-
 763 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical
 764 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024a.

765 Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J
 766 Prenger, and Animashree Anandkumar. Leandojo: Theorem proving with retrieval-augmented
 767 language models. *Advances in Neural Information Processing Systems*, 36, 2024b.

768 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan.
 769 Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural*
 770 *Information Processing Systems*, 36, 2024.

771 Yige Yuan, Teng Xiao, Li Yunfan, Bingbing Xu, Shuchang Tao, Yunqi Qiu, Huawei Shen, and Xueqi
 772 Cheng. Inference-time alignment in continuous space. *arXiv preprint arXiv:2505.20081*, 2025.

773 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang.
 774 Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model?
 775 2025. URL <https://arxiv.org/abs/2504.13837>.

776 Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and
 777 James Zou. Textgrad: Automatic” differentiation” via text. *arXiv preprint arXiv:2406.07496*,
 778 2024.

779 Boyi Zeng, Shixiang Song, Siyuan Huang, Yixuan Wang, He Li, Ziwei He, Xinbing Wang, Zhiyu
 780 Li, and Zhouhan Lin. Pretraining language models to ponder in continuous space. *arXiv preprint*
 781 *arXiv:2505.20674*, 2025.

782 Siyan Zhao, John Dang, and Aditya Grover. Group preference optimization: Few-shot alignment of
 783 large language models. *arXiv preprint arXiv:2310.11523*, 2023.

784 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan,
 785 and Jimmy Ba. Large language models are human-level prompt engineers. *arXiv preprint*
 786 *arXiv:2211.01910*, 2022.

787 Hanlin Zhu, Shibo Hao, Zhiting Hu, Jiantao Jiao, Stuart Russell, and Yuandong Tian. Reasoning
 788 by superposition: A theoretical perspective on chain of continuous thought. *arXiv preprint*
 789 *arXiv:2505.12514*, 2025.

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A OTHER RELATED WORK 811

812 **Scaling LLM Reasoning at Inference Time.** The paradigm of scaling inference-time compute
813 has emerged as a powerful strategy to amplify the reasoning capabilities of LLMs. Chain-of-thought
814 (CoT) prompting Wei et al. (2022b) pioneered this direction by eliciting multi-step reasoning explicitly.
815 Subsequent works scale inference-time compute by sampling more reasoning chains or exploring
816 larger reasoning spaces. For instance, Best-of-N (BoN) (Stiennon et al., 2020; Nakano et al., 2021)
817 and Self-consistency (SC-CoT) Wang et al. (2022) repeatedly sample multiple chains and select the
818 best one using heuristics such as predefined rewards or the consistency. More advanced approaches
819 have introduced various search algorithms to efficiently explore the reasoning space Yao et al. (2024);
820 Besta et al. (2024); Hao et al. (2023; 2024). Tree-of-thought (ToT) Yao et al. (2024) formulates
821 reasoning as a tree search problem, where the LLM generates thoughts while heuristic rewards
822 guide tree traversal. RAP Hao et al. (2023; 2024) employs Monte Carlo Tree Search (MCTS) to
823 strategically balance exploration and exploitation in the search space. Recent work by Snell et al.
824 (2024) further formalizes the empirical benefits of test-time compute scaling. In contrast to these
825 sampling and search-based methods, our work proposes gradient-based optimization to traverse the
826 reward landscape more efficiently, bypassing the trial-and-error inefficiencies of sampling approaches.
827 Notably, our method is orthogonal to recent RL-trained methods for test-time scaling, such as
828 OpenAI’s o1 and DeepSeek’s R1 – our method can be directly applied to these models to refine their
829 long CoT inference, though we leave this exploration to future work.
830

831 **Inference-time Constrained Decoding.** Constrained decoding is a traditional problem to study
832 in text generation, with popular applications including controllable text generation and preference
833 alignment. Recently, inference-time alignment has been heavily studied to align LLMs with human
834 values as alignment rewards Li et al. (2023); Huang et al. (2024); Khanov et al. (2024). Most of them
835 formulate the problem as a reward-guided search process. More relevantly, controllable text generation
836 has leveraged energy-based models (EBMs) (Kumar et al., 2021; Qin et al., 2022; Kumar et al.,
837 2022; Mireshghallah et al., 2022; Liu et al., 2023b; Yuan et al., 2025) to relax discrete text sequences
838 into continuous spaces, which enables gradient-based optimization (e.g., Langevin dynamics) to
839 steer generation toward objectives defined by arbitrary energy functions. While sharing similar
840 reward-constrained decoding principles and gradient-based methodologies, end-to-end sequence
841 optimization is notoriously unstable, often producing broken sequences or failing to converge, which
842 limits its applicability to reasoning tasks. In contrast, our approach introduces an efficient and novel
843 gradient-based framework integrated with the iterative decoding, tailored for reasoning tasks.
844

845 **Prompt Optimization.** A distinct but related line of research focuses on optimizing prompts
846 rather than decoding sequences. Established principles from the above sections apply here as well.
847 Search-based methods Zhou et al. (2022); Wang et al. (2023b); Pryzant et al. (2023); Yuksekgonul
848 et al. (2024) iteratively refine the prompts through automated trial-and-error, while gradient-based
849 approaches include directly optimizing soft prompts Li & Liang (2021); Lester et al. (2021), or
850 searching for discrete prompts via gradients Shin et al. (2020); Shi et al. (2022); Wen et al. (2024).
851 Although these techniques share conceptual similarities with our method, esp. the gradient-based
852 optimization methods, they operate on the input (prompt) space rather than the output (reasoning
853 chain) space. Our work focuses on optimizing reasoning trajectories directly, complementing rather
854 than competing with prompt optimization methods – future work could explore synergies between
855 the two paradigms.
856

857 **Continuous Latent Space Reasoning.** Another line of research explores reasoning in a continuous
858 latent space, bypassing discrete token-level operations. These approaches typically perform iterative
859 refinement on the model’s hidden states to solve reasoning problems. For example, some methods
860 frame reasoning as an energy minimization process (Du et al., 2022) or a fixed-point iteration problem
861 within the latent space, which can be parallelized for efficiency (Wu et al., 2025). Others propose
862 increasing test-time compute by applying recurrent updates to latent representations, effectively
863 deepening the model’s computation on-the-fly (Geiping et al., 2025). This paradigm has also been
864 supported by specialized pre-training objectives that encourage models to “ponder” in a continuous
865 space (Zeng et al., 2025) and has received theoretical analysis under the lens of continuous
866 chain-of-thought (Zhu et al., 2025). While conceptually related in their use of iterative refinement,
867 these methods are fundamentally different from ours. They operate within the LLM’s latent space,
868 modifying internal hidden representations. In contrast, our work, ∇ -Reasoner, performs optimization
869

directly in the output space. We manipulate the token logits, a continuous relaxation of the discrete vocabulary, guided by reward gradients. This direct textual optimization allows us to refine the reasoning chain itself at inference time without altering the base model’s internal forward pass or requiring any specialized training, uniquely positioning our method as a post-hoc, gradient-based search algorithm over the sequence space.

B IMPLEMENTATION

B.1 PSEUDOCODE

In Algorithms 1 and 2, we present a basic implementation for ∇ -Reasoner. Below we give a detailed pseudocode for a full version with all acceleration techniques integrated (Sec. 3.3). In Alg. 3, we list the complete version of iterative decoding with confidence- and gradient-guided token selection, rollout reusing, and early stop techniques. In Alg. 4, we present the full DTO algorithm with gradient caching. We note that these techniques significantly accelerate the decoding speed of ∇ -Reasoner, as demonstrated in Sec. 5.

B.2 GENERALIZATION TO PROGRESS REWARD.

The reward function can take different forms: it may provide an *outcome reward* (Cobbe et al., 2021), offering an overall score for the entire response sequence, or a *process reward* (Lightman et al., 2023), which assesses individual intermediate steps and assigns a series of scores accordingly. Thus, beyond a single reward defined over the whole sequence, we denote the total reward as the sum of rewards obtained from different subsequences: $R(\mathbf{y}|\mathbf{x}) = \sum_{l=1}^{|\mathbf{y}|} r(\mathbf{y}_{\leq l}|\mathbf{x})$. In the case of an outcome reward, the reward is only assigned at the end of the response sequence, meaning $r(\mathbf{y}_{\leq l}|\mathbf{x}) = 0$ if $l < |\mathbf{y}|$. Conversely, when using a process reward, rewards are assigned incrementally, with $r(\mathbf{y}_{\leq l}|\mathbf{x}) \neq 0$ only if \mathbf{y}_l is an end token of a thought (Xiong et al., 2024). Our framework can seamlessly incorporate a progress reward by replacing $r(\mathbf{y}|\mathbf{x})$ with this generalized version $R(\mathbf{y}|\mathbf{x})$.

918 **Algorithm 4** Differentiable Textual Optimization (DTO)

919

920 **Require:** Prefix \mathbf{x} , initial logits \mathbf{z} , language model π_{LLM} , reward model r , and the number of
921 training steps T .

922 1: $\hat{\mathbf{y}} \leftarrow \text{None}$

923 2: $\mathbf{g}_1, \mathbf{g}_2, \dots \leftarrow \text{None}$

924 3: **while** $t < T$ **do**

925 4: **for** every $i = 1, \dots, |\mathbf{y}|$ **do**

926 5: $j^* \leftarrow \arg \max_{j \in [\mathcal{V}]} \mathbf{z}_{ij}^{(t)}$

927 6: $\mathbf{y}_i^{(t)} \leftarrow \delta_{j^*} + \text{softmax}(\mathbf{z}_i^{(t)} / \tau) - \text{StopGrad}(\text{softmax}(\mathbf{z}_i^{(t)} / \tau))$

928 7: **end for**

929 8: **if** $\mathbf{y} \neq \hat{\mathbf{y}}$ **then**

930 9: $\mathcal{L}_{nll} = - \sum_{i=1}^{|\mathbf{y}^{(t)}|} \log \pi_{LLM}(\mathbf{y}^{(t)} | \mathbf{y}_{\leq i-1}^{(t)}, \mathbf{x})$

931 10: $\mathcal{L}_{reward} = -r(\mathbf{y}^{(t)} | \mathbf{x})$.

932 11: $\mathcal{L} = \mathcal{L}_{nll} + \lambda \mathcal{L}_{reward}$. ▷ Eq. 2

933 12: $\hat{\mathbf{y}} \leftarrow \mathbf{y}$

934 13: $\mathbf{g}_i \leftarrow \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i^{(t)}}$ for every $i \in [|\mathbf{y}^{(t)}|]$ ▷ Gradient caching (Sec. 3.3)

935 14: **else**

936 15: $\mathcal{L} = \sum_{i=1}^{|\mathbf{y}^{(t)}|} \mathbf{g}_i^\top \mathbf{y}_i^{(t)}$ ▷ Surrogate loss with cached gradient (Sec. 3.3)

937 16: **end if**

938 17: $\mathbf{z}^{(t+1)} \leftarrow \mathbf{z}^{(t)} - \eta \nabla_{\mathbf{z}} \mathcal{L}$.

939 18: $t \leftarrow t + 1$.

940 19: **end while**

941 20: **return** $\mathbf{z}^{(T)}$

942

943 **C DEFERRED THEORY**

944

945

946 **C.1 GRADIENT DERIVATION OF \mathcal{L}**

947

948 We derive the gradients of Eq. 2 summarized as the following proposition. We consider the
949 generalized reward function which is written as a summation over rewards defined over different
950 subsequences (Sec. B.2).

951 **Proposition C.1.** *The gradient of loss function $\mathcal{L}(\mathbf{y}) = -\lambda \sum_{i=1}^{|\mathbf{y}|} r(\mathbf{y}_{\leq i} | \mathbf{x}) - \log \pi_{LLM}(\mathbf{y} | \mathbf{x})$ takes
952 the form of $\frac{\partial \mathcal{L}(\mathbf{y})}{\partial \mathbf{y}_l} = \delta_{prefix} + \delta_{postfix} + \lambda \delta_{reward}$ where:*

953

954
$$\delta_{prefix} = -\log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq l-1}, \mathbf{x})), \quad (4)$$

955

956
$$\delta_{postfix} = - \sum_{i=l+1}^{|\mathbf{y}|} \frac{\partial \log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq i-1}, \mathbf{x}))}{\partial \mathbf{y}_l} \mathbf{y}, \quad (5)$$

957

958
$$\delta_{reward} = - \sum_{i=l}^{|\mathbf{y}|} \frac{\partial r(\mathbf{y}_{\leq i} | \mathbf{x})}{\partial \mathbf{y}_l}. \quad (6)$$

959

960 *Proof.* The proof is done by elementary derivative calculation. First of all, we write down the
961 expanded expression of the loss function:

962

963
$$\mathcal{L}(\mathbf{y}) = -\lambda \sum_{i=1}^{|\mathbf{y}|} r(\mathbf{y}_{\leq i} | \mathbf{x}) - \sum_{i=1}^{|\mathbf{y}|} \log \pi_{LLM}(\mathbf{y}_i | \mathbf{y}_{\leq i-1}, \mathbf{x}) \quad (7)$$

964

965
$$= -\lambda \sum_{i=1}^{|\mathbf{y}|} r(\mathbf{y}_{\leq i} | \mathbf{x}) - \sum_{i=1}^{|\mathbf{y}|} \sum_{v \in [\mathcal{V}]} \mathbf{y}_{i,v} \log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq i-1}, \mathbf{x}) \quad (8)$$

966

972 For a specific token index $l \in [|\mathbf{y}|]$, we decompose the loss into five components:
973

$$974 \quad \mathcal{L}(\mathbf{y}) = -\lambda \underbrace{\sum_{i=1}^{l-1} r(\mathbf{y}_{\leq i} | \mathbf{x})}_{\Phi_1} - \lambda \underbrace{\sum_{i=l}^{|\mathbf{y}|} r(\mathbf{y}_{\leq i} | \mathbf{x})}_{\Phi_2} - \underbrace{\sum_{i=1}^{l-1} \sum_{v \in [|\mathcal{V}|]} \mathbf{y}_{i,v} \log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq i-1}, \mathbf{x})}_{\Pi_1} \quad (9)$$

$$979 \quad - \underbrace{\sum_{v \in [|\mathcal{V}|]} \mathbf{y}_{l,v} \log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq l-1}, \mathbf{x})}_{\Pi_2} - \underbrace{\sum_{i=l+1}^{|\mathbf{y}|} \sum_{v \in [|\mathcal{V}|]} \mathbf{y}_{i,v} \log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq i-1}, \mathbf{x})}_{\Pi_3}, \quad (10)$$

983 where $\frac{\partial \Phi_1}{\partial \mathbf{y}_l} = 0$ and $\frac{\partial \Pi_1}{\partial \mathbf{y}_l} = 0$ because they do not involve \mathbf{y}_l . Π_2 only depends on \mathbf{y}_l through the
984 term $\mathbf{y}_{l,v}$ while Π_3 only depends on \mathbf{y}_l via $\log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq i-1}, \mathbf{x})$. Next, we compute the gradients
985 for Φ_2 , Π_2 , and Π_3 , respectively.
986

$$987 \quad \delta_{reward} = \frac{\partial \Phi_2}{\partial \mathbf{y}_l} = \sum_{i=l}^{|\mathbf{y}|} \frac{\partial r(\mathbf{y}_{\leq i} | \mathbf{x})}{\partial \mathbf{y}_l}, \quad (11)$$

$$990 \quad \delta_{prefix} = \frac{\partial \Pi_2}{\partial \mathbf{y}_l} = [-\log \pi_{LLM}(\mathbf{e}_1 | \mathbf{y}_{\leq l-1}, \mathbf{x}), \dots, -\log \pi_{LLM}(\mathbf{e}_N | \mathbf{y}_{\leq l-1}, \mathbf{x})]^\top \quad (12)$$

$$992 \quad = -\log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq l-1}, \mathbf{x})), \quad (13)$$

$$994 \quad \delta_{postfix} = \frac{\partial \Pi_3}{\partial \mathbf{y}_l} = -\sum_{i=l+1}^{|\mathbf{y}|} \sum_{v \in [|\mathcal{V}|]} \frac{\partial \log \pi_{LLM}(\mathbf{e}_v | \mathbf{y}_{\leq i-1}, \mathbf{x})}{\partial \mathbf{y}_l} \mathbf{y}_{l,v} \quad (14)$$

$$997 \quad = -\sum_{i=l+1}^{|\mathbf{y}|} \frac{\partial \log \text{Cat}(\pi_{LLM}(\cdot | \mathbf{y}_{\leq i-1}, \mathbf{x}))}{\partial \mathbf{y}_l} \mathbf{y}_l, \quad (15)$$

1000 as desired. \square

1002 *Remark C.2.* Our proposed DTO fundamentally differs from previous works that utilize gradients for
1003 controlled generation (Qin et al., 2022; Kumar et al., 2021; 2022; Mireshghallah et al., 2022; Liu
1004 et al., 2023b), where $\delta_{postfix}$ is often detached from the computational graph, and only prior context
1005 is used to guide subsequent token prediction.

1007 C.2 PROOF OF THEOREM 4.1

1008 **Theorem C.3** (Restatement of Theorem 4.1). *Suppose $\{\rho^t\}_{t \geq 0}$ denotes the Wasserstein gradient
1009 flow minimizing Eq. 3 in the distribution space with boundary conditions $\rho^0 = \pi_{LLM}$ and $\rho^\infty =$
1010 $\rho^* = \arg \min_\rho \mathcal{L}_{PPO}(\rho)$. Then we can draw samples from ρ^* by first initializing $\mathbf{x}^0 \sim \pi_{LLM}$ and
1011 simulating a trajectory $\{\mathbf{x}^t\}_{t \geq 0}$ following the stochastic gradient flow of Eq. 2: $\frac{d\mathbf{x}^t}{dt} = -\nabla \mathcal{L}(\mathbf{x}^t) +$
1012 $\sqrt{2}\epsilon_t$, where $\{\epsilon_t \in \mathcal{N}(\mathbf{0}, \mathbf{I})\}_{t \geq 0}$ are Brownian motions.*

1015 *Proof.* First of all, we derive the Wasserstein gradient flow for ρ^t on $\mathbb{W}_2(\mathbb{R}^{L_x \times N})$ under the functional
1016 cost \mathcal{L}_{PPO} (Chen et al., 2018):

$$1018 \quad \partial_t \rho^t = -\nabla_{\mathbb{W}} \mathcal{L}_{PPO}(\rho^t) \quad \Rightarrow \quad \partial_t \rho^t + \nabla_{\mathbf{x}} \cdot \left(\rho^t \nabla_{\mathbf{x}} \frac{\delta \mathcal{L}_{PPO}}{\delta \rho^t} \right) = 0, \quad (16)$$

1020 where $\frac{\delta \mathcal{L}_{PPO}}{\delta \rho^t}$ denotes the first variation of \mathcal{L}_{PPO} in terms of ρ_t , and the $\nabla \cdot$ is the divergence operator.
1021 We derive the first variation as below:

$$1023 \quad \frac{\delta \mathcal{L}_{PPO}}{\delta \rho^t} = \frac{\delta}{\delta \rho^t} \left[\int -\lambda r(\mathbf{x}) + \log \frac{\rho^t(\mathbf{x})}{\pi_{LLM}(\mathbf{x})} \rho^t(\mathbf{x}) d\mathbf{x} \right] \quad (17)$$

$$1025 \quad = -\lambda r(\mathbf{x}) + \log \rho^t(\mathbf{x}) - \log \pi_{LLM}(\mathbf{x}) + 1. \quad (18)$$

1026 The gradient of $\frac{\delta \mathcal{L}_{PPO}}{\delta \rho^t}$ can be expressed as:
 1027

$$1028 \nabla_{\mathbf{x}} \frac{\delta \mathcal{L}_{PPO}}{\delta \pi_t} = \nabla_{\mathbf{x}} (-\lambda r(\mathbf{x}) + \log \rho^t(\mathbf{x}) - \log \pi_{LLM}(\mathbf{x})). \quad (19)$$

1030 Now we substitute the above equations into Eq. 16 and find the following partial differential equation
 1031 of ρ^t :
 1032

$$1033 \partial_t \rho^t + \nabla_{\mathbf{x}} \cdot [\rho^t \nabla_{\mathbf{x}} (-\lambda r(\mathbf{x}) + \log \rho^t(\mathbf{x}) - \log \pi_{LLM}(\mathbf{x}))] = 0 \quad (20)$$

$$1034 \partial_t \rho^t + \nabla_{\mathbf{x}} \cdot [\rho^t (-\lambda \nabla_{\mathbf{x}} r(\mathbf{x}) - \nabla_{\mathbf{x}} \log \pi_{LLM}(\mathbf{x}))] + \sum_{i,j=1}^{L_x, N} \frac{\partial}{\partial \mathbf{x}_{i,j}} (\nabla_{\mathbf{x}} \log \rho^t(\mathbf{x})) = 0 \quad (21)$$

$$1037 \partial_t \rho^t + \nabla_{\mathbf{x}} \cdot [\rho^t (-\lambda \nabla_{\mathbf{x}} r(\mathbf{x}) - \nabla_{\mathbf{x}} \log \pi_{LLM}(\mathbf{x}))] + \Delta \rho^t(\mathbf{x}) = 0, \quad (22)$$

1039 where Δ means the Laplacian operator $\sum_{ij} \frac{\partial^2}{\partial \mathbf{x}_{ij}^2}$. By Fokker-Plank equation (Maoutsa et al., 2020),
 1040 we obtain a velocity field for particles $\mathbf{x}^t \in \mathbb{R}^{L_x \times N}$:
 1041

$$1042 \frac{d\mathbf{x}^t}{dt} = -\lambda \nabla_{\mathbf{x}} r(\mathbf{x}) - \nabla_{\mathbf{x}} \log \pi_{LLM}(\mathbf{x}) + \sqrt{2} \epsilon_t = -\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^t) + \sqrt{2} \epsilon_t, \quad (23)$$

1044 where $\{\epsilon_t\}_{t \geq 0}$ denotes a Brownian motion (Øksendal, 2003)).
 1045

1046 Conversely, if \mathbf{x}^t follows the above dynamics starting from the initial distribution $\mathbf{x}^t \sim \rho^0$, then the
 1047 density function ρ^t of \mathbf{x}^t follows the time evolution in Eq. 16 (see Maoutsa et al. (2020)). Since we
 1048 assume the limiting condition $\rho^t \rightarrow \rho^*$, we conclude that $\mathbf{x}^t \sim \rho^*$ in distribution as $t \rightarrow \infty$. \square
 1049

1050 C.3 JUSTIFICATION OF CONFIDENCE-BASED TOKEN SELECTION

1051 We re-parameterize the policy as a categorical distribution through a softmax function to ensure
 1052 differentiability. In Sec. 3.3, we mention that we can skip optimizing tokens whose logits is over-
 1053 confident as gradient descent is unlikely to significantly change its resultant distribution. We provide
 1054 theoretical evidence for this argument.

1055 Without loss of generality, we consider the scenario where we only optimize a single token $\mathbf{x} \in \mathbb{R}^{|\mathcal{V}|}$
 1056 within a vocabulary \mathcal{V} . We initialize a logit vector $\mathbf{z} \in \mathbb{R}^N$, then apply the softmax transformation to
 1057 obtain the corresponding categorical distribution:
 1058

$$1059 \mathbf{x}_i = \frac{\exp(\mathbf{z}_i)}{\sum_{j=1}^{|\mathcal{V}|} \exp(\mathbf{z}_j)}, \quad \forall i \in [|\mathcal{V}|]. \quad (24)$$

1062 The loss function is defined over \mathbf{x} and by chain rule, we derive the gradient w.r.t \mathbf{z}_i for $i \in [|\mathcal{V}|]$:
 1063

$$1064 \frac{\partial \mathcal{L}}{\partial \mathbf{z}_i} = \frac{\partial \mathcal{L}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{z}_i} = (\text{diag}(\mathbf{x}) - \mathbf{x} \mathbf{x}^\top)_i \frac{\partial \mathcal{L}}{\partial \mathbf{x}} = \mathbf{x}_i \left(\left[\frac{\partial \mathcal{L}}{\partial \mathbf{x}} \right]_i - \mathbf{x}^\top \frac{\partial \mathcal{L}}{\partial \mathbf{x}} \right), \quad (25)$$

1066 where we use the fact that the Jacobian matrix of softmax function is $\text{diag}(\mathbf{x}) - \mathbf{x} \mathbf{x}^\top$. The derivation
 1067 above indicates that the gradient magnitude for the i -th logit is proportional to its corresponding
 1068 post-softmax probability. When \mathbf{x}_i is small at the initialization, then its underlying representation \mathbf{z}_i
 1069 cannot be updated effectively. This limitation underscores the necessity of skipping tuning \mathbf{x} with
 1070 high confidence throughout the decoding stage.
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080 **D MORE ON EXPERIMENT**
10811082 **D.1 EXPERIMENT DETAILS**
1083

1084 We used a temperature of 0.5 and a top-p of 0.95, and set the maximum generation length to 1024 for
 1085 AMC and MATH-500, and 3072 for AIME, for all baselines and our methods. We report the average
 1086 performance across 4 independent runs on smaller datasets such as AMC, AIME-24, and AIME-25.
 1087 For SFT experiments, we randomly sampled a 10k subset from the Open-thoughts dataset (Guha
 1088 et al., 2025). For GRPO experiments, we used a random 35k subset from the Numina math dataset (LI
 1089 et al., 2024). Exclusively for ∇ -Reasoner, we set $\epsilon_{ent} = 0.25$, $\epsilon_{grad} = 8$, learning rate to 0.01, and
 1090 number of iterations to 20 in all experiments. We use Skywork-Reward-V2-Qwen3-4B (Liu et al.,
 1091 2025) as the reward model for Qwen family and Skywork-Reward-V2-Llama-3.1-8B (Liu et al.,
 1092 2025) as the reward model for Llama model. Below, we list a more comprehensive summary of the
 1093 experiment setups.

1094
1095 **Table 5: Summary of Experimental Settings.**

1096 Setting	1097 Value
Generation Hyperparameters	
1098 Temperature	0.5
1099 Top-p	0.95
1100 Max Generation Length (AMC, MATH-500)	1024
1101 Max Generation Length (AIME)	3072
Evaluation	
1103 Independent Runs (MATH-500)	1
1104 Independent Runs (AMC, AIME-24, AIME-25)	4
Training Data Subsets	
1105 SFT	10k random sample from Open-thoughts (Guha et al., 2025)
1106 GRPO	35k random sample from Numina math (LI et al., 2024)
∇-Reasoner Hyperparameters	
1108 ϵ_{ent}	0.25
1109 ϵ_{grad}	8
1110 Learning Rate	0.01
1111 Number of Iterations	20
1112 Optimizer	Adam-W
1113 LR Scheduler	Cosine
1114 Min LR	0.001
Reward Models	
1115 For Qwen Family	Skywork-Reward-V2-Qwen3-4B (Liu et al., 2025)
1116 For Llama Model	Skywork-Reward-V2-Llama-3.1-8B (Liu et al., 2025)

1119 **D.2 EFFECTIVENESS OF ACCELERATION TECHNIQUES**
1120

1121 In Sec. 3.3, we proposed three major techniques to improve the decoding efficiency of ∇ -Reasoner.
 1122 In this section, we quantify the contribution of each technique to the overall speedup. Instead of
 1123 ablating them in running time (since removing any one of these components causes the algorithm to
 1124 run in an prohibitive amount of time, often more than several hours per sample), we monitor how
 1125 each technique reduces the cost of the specific stage to which it is applied. As a result, we find our
 1126 gradient-caching mechanism bypasses more than 63.8% of the model calls in parallel forms required
 1127 for gradient acquisition, and trajectory reuse eliminates more than 74.1% of autoregressive model
 1128 calls. Both techniques exploit the sparsity of token updates. The token-selection strategy further
 1129 complements these components by skipping the full optimization procedure when appropriate; we
 1130 observe that it effectively avoids 89.2% of token-optimization steps across the sequence.

1132 **D.3 A SIMPLIFIED EXAMPLE ON GSM8K**
1133

The problem asks:

1134 Josh buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased
 1135 the value of the house by 150%. How much profit did he make?
 1136

1137 **Greedy Decoding Baseline (Failure Analysis)**

1138 The original output reads:

1140 Josh buys a house for \$80,000 and then puts in \$50,000 in repairs.
 1141 The value of the house increased by 150%, so the new value of the
 1142 house is $80,000 \times 1.5 = \$120,000$. Josh sold the house for 120,000,
 1143 so he made a profit of $120,000 - 80,000 - 50,000 = -10,000$. The
 1144 answer is -10,000.

1145 This greedy decoding baseline fails by interpreting "increased by" as simple multiplication $(80,000 \times$
 1146 1.5), resulting in:

$$\text{Incorrect New Value} = \$80,000 \times 1.5 = \$120,000$$

1147 and an incorrect profit of:

$$\text{Incorrect Profit} = \$120,000 - \$80,000 - \$50,000 = -\$10,000$$

1152 **Iterative Token Revision Process (Correction)**

1153 Our method corrects this error through an iterative token revision process.

1155 Initial Input and Refinement Round 1 (Operator Correction)

- 1157 • Initial Output Segment: "..., so the new value of the house is $80,000 \times 1.5 = \$120,000$."
- 1158 • The algorithm performs refinement at the multiplication token "×".
- 1159 • Through gradient descent, the model updates the multiplication to-
 1160 ken "×" to the addition symbol "+", revising the logical structure to:
 1161 "..., so the new value of the house is $80,000 + \dots$ "

1163 Refinement Round 2 (Calculation Correction)

- 1165 • Subsequently, when calculating the increment, the model initially outputs $80,000 \times 1.5 =$
 1166 \$120,000. The calculation is incorrect because the model ignores the preceding $80,000 +$.
- 1167 • The algorithm flags the token **120** (representing \$120,000) due to a low score.
- 1168 • Gradient optimization identifies **200** as the top replacement, correcting the calculation for
 1169 the new value:

$$\text{Correct New Value} = \$80,000 + (\$80,000 \times 1.5) = \$80,000 + \$120,000 = \$200,000$$

1172 Final Correct Derivation

1174 The model completes the string to form "\$200,000", allowing it to derive the correct profit:

$$\text{Correct Profit} = \$200,000 - \$80,000 - \$50,000 = \$70,000$$

1177 **Conclusion**

1178 In conclusion, after two rounds of iterative refinement, the final output becomes:

1180 Josh buys a house for \$80,000 and then puts in \$50,000 in repairs.
 1181 The value of the house increased by 150%, so the new value of the
 1182 house is $80,000 + 80,000 \times 1.5 = \$200,000$. Josh sold the house for
 1183 200,000, so he made a profit of $200,000 - 80,000 - 50,000 = 70,000$.
 1184 The answer is 70,000.