
Test-Time Alignment for Large Language Models via Textual Model Predictive Control

Kuang-Da Wang^{1,†} Teng-Ruei Chen^{1,†} Yu Heng Hung¹ Guo-Xun Ko¹
Shuoyang Ding² Yueh-Hua Wu² Yu-Chiang Frank Wang² Chao-Han Huck Yang²
Wen-Chih Peng¹ Ping-Chun Hsieh¹
¹National Yang Ming Chiao Tung University ²NVIDIA

Abstract

Aligning Large Language Models (LLMs) with human preferences through fine-tuning is resource-intensive, motivating lightweight alternatives at test time. We address test-time alignment through the lens of sequential decision making, a perspective that reveals two fundamental challenges. When actions are defined at the token level, as in guided decoding, alignment suffers from the *curse of horizon*. Conversely, when actions are at the response level, as in traditional iterative refinement, the *curse of dimensionality* emerges. To resolve this trade-off, we draw inspiration from Model Predictive Control (MPC) in control theory to propose **Textual Model Predictive Control (TMPC)**, a novel predictive planning framework adapted for aligning LLMs at inference time. A key limitation of standard MPC is its reliance on predefined, hard segment boundaries, which are often absent in text generation. TMPC overcomes this by introducing two principles inspired by hierarchical reinforcement learning: (1) *Hindsight Subgoal Identification*, where TMPC analyzes generation subgoals to retrospectively identify high-reward intermediate outputs as subgoals. This allows the framework to discover meaningful, task-specific planning steps (*e.g.*, a sentence in machine translation or a bug fix in code generation.). (2) *Subgoal-Conditioned Re-Generation*, where these identified subgoals are used to guide subsequent planning iterations. By conditioning on these proven, high-quality subgoals, TMPC ensures stable improvement by building upon previously validated successes. TMPC is evaluated on three tasks with distinct segmentation properties: discourse-level translation, long-form response generation, and program synthesis. The results demonstrate that TMPC consistently improves performance, highlighting the generality.

1 Introduction

The emergence of Large Language Models (LLMs), such as the GPT series [1, 4], LLaMAs [30, 31], and Gemma [29], has demonstrated remarkable efficacy in a wide range of NLP tasks [9, 27, 28, 41, 43]. While these models exhibit strong performance out of the box, aligning their outputs to human preferences remains critical, especially for smaller-scale LLMs. For instance, in machine translation [2], smaller LLMs (*e.g.*, under 10B parameters) frequently suffer from omissions and semantic drift [37]. Thus, aligning LLM outputs to preferences remains an essential yet challenging problem.

Training-time approaches such as Reinforcement Learning with Human Feedback (RLHF) [22] and Direct Preference Optimization (DPO) [25] have achieved strong results in aligning preferences. However, these methods are resource-intensive and require costly retraining whenever preferences or

[†]Equal contribution. Correspondence to Kuang-Da Wang <gdwang.cs10@nycu.edu.tw>, Ping-Chun Hsieh <pinghsieh@nycu.edu.tw>, and Chao-Han Huck Yang <hucky@nvidia.com>.

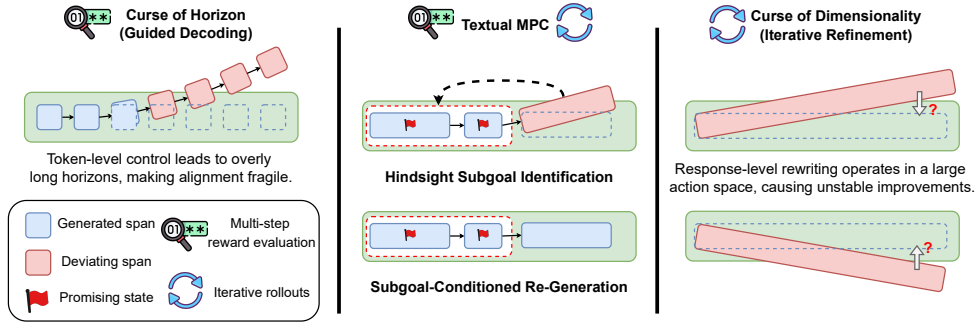


Figure 1: Comparison of test-time alignment approaches. Guided decoding suffers from the curse of horizon due to token-level control, while iterative refinement suffers from the curse of dimensionality at the response level. We balance these extremes by adopting segment-based horizon shortening and state-aware exploration, enabling stable and efficient preference alignment in long-form generation.

tasks change. This has spurred interest in *test-time alignment*, where outputs are adapted without updating model parameters, using strategies such as prompting [17], guided decoding [12, 13, 16, 35, 39], or iterative refinement [15].

We address test-time alignment through the lens of sequential decision making, where the generation process is framed as a sequence of actions. This perspective reveals two fundamental challenges, illustrated in Figure 1. When actions are defined at the token level (*e.g.*, guided decoding), methods suffer from the *curse of horizon* [23]; credit assignment becomes unreliable over long trajectories, making alignment brittle. In contrast, when actions are at the response level (*e.g.*, iterative refinement), they face the *curse of dimensionality*; each step involves rewriting an entire sequence, making the search for improvements in a vast action space intractable and unstable.

To address these challenges, we propose Textual Model Predictive Control (TMPC), a novel test-time alignment framework inspired by Model Predictive Control (MPC) [6, 14]. While powerful, standard MPC assumes the problem can be decomposed into predefined, hard segment boundaries, a condition that rarely holds for complex text generation. TMPC is uniquely adapted to overcome this limitation through two principles:

- **Hindsight Subgoal Identification:** This principle allows TMPC to discover meaningful planning steps. After generating candidate responses, TMPC retrospectively analyzes them to identify high-quality intermediate points as *subgoals*. A subgoal can be a concrete unit, such as a sentence in translation, or an abstract one, such as resolving a single failed test case in program synthesis, successfully addressing the problem of lacking natural boundaries. This hindsight-driven discovery effectively shortens the planning horizon for diverse tasks.
- **Subgoal-Conditioned Re-Generation:** This principle ensures stable, cumulative progress. The subgoals identified via hindsight are stored in a buffer and used to guide subsequent planning iterations. By conditioning the next generation on these subgoals, TMPC ensures that subsequent generation builds upon these validated waypoints.

We evaluate TMPC on three challenging tasks with different boundary characteristics: WMT’24 discourse-level machine translation, the HH-RLHF long responses subset, and MBPP program synthesis. Experiments with LLaMA-3.1-8B-Instruct show that TMPC consistently improves alignment, highlighting the generality of our approach.

Our contributions are summarized as follows:

- We propose a novel formulation of test-time alignment as a sequential decision-making problem. This perspective unifies existing approaches and reveals a fundamental trade-off that governs their limitations: the *curse of horizon* in guided decoding methods and the *curse of dimensionality* in iterative refinement methods.
- We introduce Textual Model Predictive Control (TMPC), a framework that adapts concepts from control theory to language generation. TMPC is operationalized through two

principles: *Hindsight Subgoal Identification* to discover subgoals from rollouts, and *Subgoal-Conditioned Re-Generation* to iteratively improve generation by building on subgoals.

- We empirically demonstrate the effectiveness of TMPC. TMPC achieves substantial improvements across three distinct domains including long-form response generation, discourse-level machine translation, and program synthesis, validating its ability to discover and leverage task-specific subgoals.

2 Related Work

2.1 Preference Alignment Through Fine-tuning

Aligning large language models (LLMs) with human preferences has traditionally relied on post-training strategies. Supervised fine-tuning (SFT) [44] and reinforcement learning from human feedback (RLHF) [22] are widely used but computationally expensive. Direct Preference Optimization (DPO) [25] simplifies RLHF by converting it into a supervised learning objective, though it requires managing dual policies. More recent approaches like SimPO [20] and Contrastive Preference Optimization (CPO) [38] reduce memory and resource demands using reference models and contrastive signals. Despite these improvements, fine-tuning methods remain rigid and slow to adapt to changing data or objectives, posing challenges in dynamic environments.

2.2 Test-Time Preference Alignment

Test-time preference alignment offers an efficient way to align frozen language models by guiding generation at inference, without requiring any parameter updates. Beyond simple prompting or in-context learning, guided decoding methods harness external signals to control the generation itself. ARGS [12] is a representative example that incorporates reward model guidance at the token level, and InferAligner [35] adopts a similar strategy. Among guided decoding methods, there are also approaches that directly modify internal representations. For instance, RE-Control [13] trains a value function on hidden states using the Bellman equation, and applies gradient-based optimization to align with preferences. TreeBoN [24] and RAIN [16] leverage tree-based structures: TreeBoN combines tree search with Best-of-N sampling, while RAIN performs self-evaluation without relying on a reward model to align preferences. GenARM [39] enhances test-time alignment by introducing an autoregressive reward model that predicts next-token reward signals conditioned on prior context, enabling efficient, token-wise guidance that is theoretically expressive under a KL-regularized RL framework. Test-Time Preference Optimization (TPO) [15] takes a distinct approach, translating reward feedback into textual critiques that serve as language-based rewards. The model uses these to iteratively refine its output—effectively learning alignment on the fly.

3 Background

In the general setup of RL for LLMs, text generation can be formally modeled as a finite-horizon Markov Decision Process (MDP). We adopt a general notion of a *step* as the basic unit of temporal progression, which can represent a token, a segment at various granularities (*e.g.*, phrase, sentence, or paragraph), or other linguistically or structurally meaningful units. Then, an MDP can be defined as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, R, \mu, T)$, where (i) the state space \mathcal{S} consists of all possible text prefixes, (ii) the action space \mathcal{A} corresponds to the set of all possible generation units, (iii) \mathcal{P} denotes the transition function, (iv) $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function that assigns scalar feedback to step-level or trajectory-level outcomes (*e.g.*, measuring fluency, factuality, or alignment with user preferences), (v) μ denotes the initial state distribution, and (vi) $T \in \mathbb{N}$ is the episode length.

We define the initial state s_0 as the initial prompt and let a_t denote the partial response generated at step t . At each step t , the current state is the set of tokens from the initial prompt and the partial responses generated up to step t , *i.e.*, $s_t = (s_0, a_1, \dots, a_{t-1})$. Based on this construction, we know that the transition function is deterministic with $\mathcal{P}(s_{t+1}|s_t, a_t) = 1$. A policy $\pi_\theta(a|s)$, parameterized by the language model, defines a probability distribution over actions given the prefix $s \in \mathcal{S}$. The generation of a full text sequence of length T can therefore be viewed as a trajectory $\tau = (s_0, a_0, \dots, s_{T-1}, a_{T-1}, s_T)$ with the cumulative reward given by $\mathcal{J}(\tau) := \sum_{t=0}^{T-1} R(s_t, a_t)$.

This perspective enables the application of RL methods to text generation in LLMs. Rather than relying solely on maximum likelihood estimation, which optimizes local token-level likelihoods, the MDP formulation allows optimization with respect to long-horizon objectives such as coherence and alignment with human preferences. This provides the foundation for recent advances in preference-based fine-tuning and test-time alignment.

4 Methodology

4.1 Test-Time Alignment via Trajectory Optimization

Our key idea is to take a model-based RL viewpoint to achieve test-time alignment for LLMs. Specifically, we propose to recast preference alignment as *trajectory optimization* and thereby employ receding-horizon control for iterative text generation.

Text Generation Optimization as Trajectory Optimization. Usually adopted by the model-based RL literature [7, 19], the goal of trajectory optimization is to find an optimal sequence of actions $\mathbf{a}^* = (a_0^*, \dots, a_{T-1}^*)$ such that the total trajectory-wise reward is maximized. This matches the objective of LLM text generation in that the output response is generated to best align with the underlying preference. Recall from Section 3 that we adopt a general notion of a *step* as the basic unit of temporal progression, which can be a segment at various granularities or other linguistically meaningful units. Again, we let s_0 denote the initial prompt and let $\tau = (s_0, a_0, \dots, s_{T-1}, a_{T-1}, s_T)$ denote a trajectory generated under an action sequence $\mathbf{a}_{0:T-1} := (a_0, a_1, \dots, a_{T-1})$. Given an initial prompt s_0 , the search for an optimal sequence $\mathbf{a}^*(s_0)$ can be formulated by the following optimization problem

$$\mathbf{a}^*(s_0) := \arg \max_{\mathbf{a}_{0:T-1}} \sum_{t=0}^{T-1} R(s_t, a_t). \quad (1)$$

Note that there is no need to take expectation in (1) as the state transitions are deterministic given $\mathbf{a}_{0:T-1}$ in MDPs for text generation, as described in Section 3.

Textual Model Predictive Control for Text Generation. In general, direct optimization of (1) requires searching over all possible action sequences of length T and is computationally intractable. As a predictive planning method, MPC planner approximately solves (1) by iteratively solving *local* optimization problems [8], instead of globally optimizing the total reward in one pass. Specifically, MPC planner determines the action of each step t by estimating the optimal subsequence $\mathbf{a}_{t:t+H}^*$ on a moving horizon H (usually H is smaller than T), given the state s_t , *i.e.*,

$$\mathbf{a}^{\text{MPC}}(s_t) := \arg \max_{\mathbf{a}_{t:t+H-1}} \sum_{i=t}^{t+H-1} R(s_i, a_i), \quad (2)$$

and then select a subset of $\mathbf{a}^{\text{MPC}}(s_t)$, denoted by $\tilde{\mathbf{a}}^{\text{MPC}}(s_t)$, for execution. In practice, $\tilde{\mathbf{a}}^{\text{MPC}}(s_t)$ can be selected as the first j contiguous actions ($1 \leq j \leq H$) or as a set of non-contiguous actions [5]. As a model-based approach, MPC solves (2) by employing (i) a learned predictive dynamics model and (ii) a proposal action distribution to jointly generate multiple H -step predictive rollouts $\{\tau_t^{(i)} \equiv (s_{t:t+H-1}^{(i)}, \mathbf{a}_{t:t+H-1}^{(i)})\}_{i=1}^K$ and obtain an approximate maximizer based on these K rollouts. As a widely-used variant of MPC for continuous control, Model Predictive Path Integral (MPPI) [36] determines an approximate maximizer by performing a soft, utility-weighted aggregated selection as $a_t = (\sum_{i=1}^K \exp(\frac{1}{\lambda} \mathcal{J}(\tau_t^{(i)})) a_t^{(i)}) / \sum_{i=1}^K \exp(\frac{1}{\lambda} \mathcal{J}(\tau_t^{(i)}))$, where $\mathcal{J}(\tau)$ denotes the cumulative reward of a rollout τ and $\lambda > 0$ controls the exploration–exploitation trade-off. Compared to deterministic MPC that selects a single maximizer, MPPI yields smoother updates by aggregating multiple high-reward rollouts while still biasing toward higher \mathcal{J} .

Inspired by MPPI for continuous control, to better leverage MPC in text generation (inherently with discrete actions), we propose to define an *aggregation function* that determines the action sequence by aggregating multiple textual rollouts based on the corresponding cumulative rewards, *i.e.*,

$$\mathbf{a}^{\text{TMPC}}(s) \leftarrow \mathcal{G}\left(\{\tau^{(i)}\}_{i=1}^K, \{\mathcal{J}(\tau^{(i)})\}_{i=1}^K; s\right), \quad (3)$$

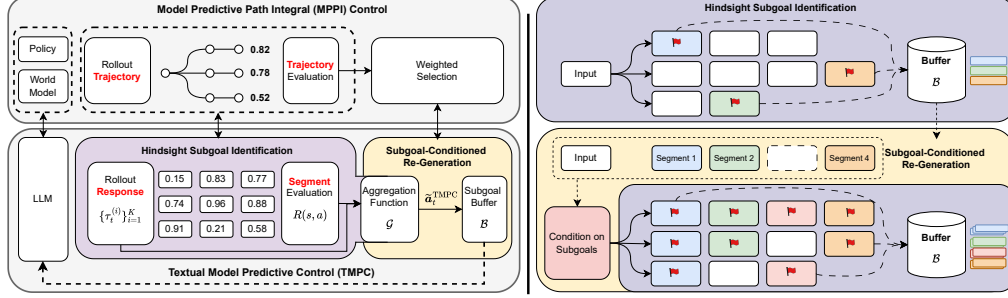


Figure 2: TMPC adapts the MPPI framework for test-time alignment by introducing two core principles. **Hindsight Subgoal Identification:** After generating multiple rollouts, the planner’s aggregation function \mathcal{G} selects a subset of locally-optimal actions $\tilde{\mathbf{a}}^{\text{TMPC}}$. This executed plan is retrospectively identified as a high-quality **subgoal** and stored in a buffer \mathcal{B} if its utility meets a threshold α . **Subgoal-Conditioned Re-Generation:** New rollouts are generated by sampling from and composing subgoals in the buffer \mathcal{B} . This allows the planner to iteratively refine the full-horizon plan by building upon the best strategies discovered in previous iterations.

where $\{\tau^{(i)}\}_{i=1}^K$ are rollouts starting from s . Then, TMPC can leverage a sequence of non-contiguous actions, denoted by $\tilde{\mathbf{a}}^{\text{TMPC}}(s)$, to be selected for actual use in subgoal generation. The detailed construction of \mathcal{G} will be specified in Section 4.2.

Notably, TMPC enjoys two salient features that make it a particularly suitable method for test-time alignment of LLMs: (i) *No additional model learning or fine-tuning needed:* Recall that MPC-like methods typically require a learned dynamics model and a proposal distribution. In TMPC, a dynamics model is already available since in text-generation MDPs, the transition from s_t to s_{t+1} is known and deterministic under an action a_t . Moreover, a pre-trained frozen LLM can naturally play the role of a good proposal distribution for generating candidate texts. Hence, TMPC does not require any fine-tuning or model learning. (ii) *Addressing curse of horizon and curse of dimensionality:* TMPC addresses these two fundamental issues by iteratively solving local optimization problems. Compared to guided decoding and full-response iterative refinement, the design of TMPC can achieve a better balance between accurate credit assignment and the size of search space.

4.2 Textual Model Predictive Control for General Temporal Progression

In this section, we extend the TMPC framework to text generation tasks with general temporal progression. Inherited from the classic MPC, TMPC described in Section 4.1 presumes that there already exists a basic unit as a discrete time step for planning. This requirement indeed holds for various tasks, such as viewing one output sentence as a step in machine translation and text summarization. However, there also exist text generation tasks without natural boundaries, such as code generation. Despite this, we present a more general version of TMPC that can achieve approximate trajectory optimization in text generation, *with and without natural boundaries*, by introducing *subgoals*, which can serve as a basic unit for temporal progression. More specifically, subgoals provide directional guidance for the LLM’s generation, enabling efficient exploration toward the optimum. TMPC can be substantiated via two core principles, as illustrated in Figure 2.

Principle 1: Hindsight Subgoal Identification. To achieve higher-quality generation, we construct meaningful subgoals from continuous text by aggregating prior high-reward actions into a buffer \mathcal{B} . This identification occurs **after** rollouts are evaluated, hence hindsight, the planner discovers what constitutes a successful step based on empirical outcomes. The update rule of the buffer is as follows:

$$\mathcal{B} \leftarrow \begin{cases} \mathcal{B} \cup \tilde{\mathbf{a}}_t^{\text{TMPC}}(s), & \text{if } |\mathcal{B}| < \text{capacity}, \\ \mathcal{B} \setminus \{a \in \mathcal{B} \mid R(s, a) < R(s, a')\} \cup \{a'\}, & \text{otherwise, for each } a' \in \tilde{\mathbf{a}}_t^{\text{TMPC}}(s). \end{cases} \quad (4)$$

Principle 2: Subgoal-Conditioned Aggregation Function for Re-Generation. In TMPC, the non-contiguous actions are generated from the following aggregation function:

$$\tilde{\mathbf{a}}_t^{\text{TMPC}}(s) \leftarrow \mathcal{G}\left(\{\tau_t^{(i)}\}_{i=1}^K, R(\cdot) \mid s, \mathcal{B}\right) := \left\{a \mid R(s, a) \geq \alpha \text{ and } a \in \{\tau_t^{(i)}\}_{i=1}^K\right\}, \quad (5)$$

where $\{\tau_t^{(i)}\}_{i=1}^K$ are the rollouts generated from subgoal-conditioned LLM $\pi(s, \mathcal{B})$. $\tilde{\mathbf{a}}_t^{\text{TMPC}}(s)$ implicitly favors higher-reward outcomes by exploiting subgoals that serve as local optimizers over planning iterations, making it a validated and locally optimal action sequence with high utility.

This principle describes how TMPC leverages identified subgoals to refine the entire trajectory over multiple iterations. A single pass of optimization may yield a suboptimal solution. TMPC overcomes this by performing planning iteratively. In the subgoal identification step, the planner populates the subgoal buffer \mathcal{B} using the Hindsight Subgoal Identification described above. The re-generation step constructs new rollouts by explicitly leveraging the high-reward goals accumulated in \mathcal{B} as conditioning signals. Rather than exploring from a generic proposal distribution, the planner is encouraged to generate new candidate trajectories by composing and extending the high-quality subgoals from the buffer. The aggregation function \mathcal{G} thus plays a crucial role: it not only selects high-reward action subset $\tilde{\mathbf{a}}_t^{\text{TMPC}}$ for the current iteration but also leverages the subgoal buffer \mathcal{B} to inform the generation of rollouts for the next iteration. This iterative process allows TMPC to escape poor local optima and progressively construct a globally high-utility response by combining the best building blocks (subgoals) discovered across all iterations.

5 Experiments

We evaluate TMPC on three tasks with different structural properties to ensure its generality: **(1) Paragraph-Level Machine Translation** represents a task *with natural boundaries*. The generated translation can be precisely aligned with the source text, allowing for sentence-level segments that are structurally anchored and easy to evaluate. **(2) Long-Form Response Generation** represents a task *without natural boundaries*. Without a source for direct alignment, responses are segmented by content into coherent chunks (*e.g.*, groups of sentences), each preserving semantic integrity. **(3) Program Synthesis** challenges conventional segmentation, representing a task where structural boundaries (*e.g.*, Abstract Syntax Tree nodes) are semantically too fragmented for effective planning.

5.1 Preference Dataset and Reward Model

Paragraph-Level MT Dataset. We use the WMT’24 Discourse-Level Literary Translation benchmark [34] for our experiments. The available language pairs include: Chinese \rightarrow English, Chinese \rightarrow German, and Chinese \rightarrow Russian. The preference dataset is derived from the training set of the dataset. Each instance is segmented into paragraphs of up to 1,024 tokens. From each translation direction, we sample 2,000 paragraphs, resulting in a total of 6,000 paragraphs for constructing the preference dataset. Translation outputs are generated using LLaMA-3.1-8B-Instruct, Gemma-2-9B, and GPT-4o. The translations are then evaluated with MetricX-24-XL [11] under the reference-free evaluation mode, where no reference translation is supplied as input. Following the procedure in CPO [38], we assign the translation with the highest score as the chosen response, the one with the lowest score as the rejected response, and discard the middle-scoring translation. The resulting reward model achieves 88.53% validation accuracy. Further details on the formation of preference data can be found in Appendix D, and detail of training can be found in Appendix G.2.

Long-Form Response Dataset. We use the Dahoas/full-hh-r1hf¹ dataset, which is widely adopted for LLM alignment. This dataset is designed to improve AI assistant behavior in terms of helpfulness and harmlessness. Each sample consists of a prompt and two responses, with one labeled as preferred based on human judgments. Since the response lengths in the dataset vary significantly, we select samples based on the length of the chosen responses. Specifically, we construct the training set using the top 6K samples with the longest chosen responses from training set, and using the top 1024 longest chosen responses from the testing set to construct test set. We use the 6k size training set to train a reward model, which achieves a validation accuracy of 83.78%.

Program Synthesis Dataset. We evaluate performance on the official testing set of the Mostly Basic Python Programming (MBPP) dataset [3], which comprises 500 problems (Task IDs 11-510). As

¹<https://huggingface.co/datasets/Dahoas/full-hh-r1hf>

discussed, code generation offers a direct reward signal. The resulting pass rate serves as the direct reward signal, eliminating the need for a separate reward model.

5.2 Evaluation Metrics

Paragraph-Level MT. We use SEGALe [33], a framework that extends existing metrics to long-text translation. Following CPO [38], we apply COMET² within the SEGALe framework, thereby extending COMET to the paragraph level. To better capture contextual quality, rather than feeding only source, translation, and reference sentences into COMET, we follow Vernikos et al. [32] and incorporate three concatenated sentences as inputs. We refer to this context-augmented version as **SEGALe_{comet}**. SEGALe further reports the **Null Alignment (NA) Ratio**, the proportion of source or translation sentences that fail to align, often due to over- or under-translation.

Long-Form Responses. We evaluate response quality using two complementary metrics: **Average Reward** measures the mean score assigned by the reward model. This reflects the degree of alignment with helpfulness and harmlessness preferences. We introduce this metric to directly test whether TMPC achieves stronger alignment when the reward model and evaluation are consistent. To avoid the potential for “cheating” in reward-based scoring, we also report **Win Rate**, which captures the proportion of pairwise comparisons in which a model’s response is preferred over a reference response by GPT-4 [21]. Following the ARGS evaluation protocol [12], GPT-4 is prompted to assess overall response quality, considering helpfulness, harmlessness, relevance, accuracy, depth, creativity, and detail. The full evaluation prompt is provided in Appendix F.

Program Synthesis. Following standard practice, we directly report the **Pass Rate**, defined as the proportion of problems for which all associated test cases are passed.

5.3 Baselines

We evaluate all training-time alignment methods on LLaMA-3.1-8B-Instruct and also adopt it as the backbone for all test-time alignment methods, including TMPC. Implementation details of TMPC, including parameters and prompt design, are provided in Appendix H.

Test-Time Alignment Methods. We compare TMPC against the following representative approaches. (1) **ARGS** [12], a token-level decoding method that incorporates reward model guidance during inference. (2) **RAIN** [16], which leverages tree-structured self-evaluation without relying on an external reward model. (3) **RE-Control** [13], which modifies internal representations by training a value function on hidden states with the Bellman equation and applying gradient-based optimization to align preferences. (4) **GenARM** [39], an approach that trains an autoregressive reward model to assign token-level rewards conditioned on past tokens, and combines these reward scores with next-token probabilities during inference. (5) **TPO** [15], which translates reward signals into textual critiques and uses an LLM to provide feedback for iterative refinement. (6) **Best-of-N Sampling**, a widely adopted baseline that generates multiple candidates and selects the highest-scoring one.

To ensure fair comparison, ARGS and RE-Control are equipped with the same reward model as TMPC. RAIN requires neither a reward model nor additional training data. For TPO, we set the number of iterations to 4 to ensure it generates no fewer responses than TMPC, although this involves more LLM calls for textual losses and gradients³. Further implementation details for all baselines, including a breakdown of TPO’s LLM calls, are provided in Appendix G.1.

Training-Time Alignment Methods. We further compare TMPC with training-time alignment methods. We include supervised fine-tuning (SFT) on the same preference dataset, which often serves as a strong baseline in translation. In addition, we evaluate SimPO [20] and DPO [25], which represent recent and mainstream approaches to preference-based training-time alignment, respectively. Details of training procedures are reported in Appendix G.2.

Task-Specific Settings. For paragraph-level MT, we include two high-performance models for additional context: **GPT-4o**, which serves as a strong upper bound despite not being specialized for translation [26], and **Qwen-2.5-14B**, a competitive open-source alternative for Chinese language tasks. For program synthesis, our comparison focuses on Best-of-N sampling and TPO. Token-level

²Unbabel/wmt22-comet-da

³TPO results are reported at iteration=2 in HH-RLHF dataset, as iteration=4 led to out-of-memory errors.

guided decoding methods are excluded as functional correctness is a holistic property of the entire code sequence, making them ill-suited for this task.

5.4 Quantitative Results

Methods	Test-Time	zh → en		zh → ru		zh → de	
		SEGALE _{comet} ↑	NA Ratio ↓	SEGALE _{comet} ↑	NA Ratio ↓	SEGALE _{comet} ↑	NA Ratio ↓
GPT-4o 2024-08-06	-	94.58	0.10	93.74	0.00	94.54	0.00
Qwen-2.5 (14B)	-	94.43	0.18	90.47	3.08	92.98	1.24
Llama-3.1 (8B)	×	84.36	10.47	86.28	4.19	88.97	4.43
Llama-3.1 _{SFT}	×	93.54	0.34	89.11	1.92	93.47	0.19
Llama-3.1 _{SimPO}	×	91.74	1.66	84.56	2.53	93.40	0.00
Llama-3.1 _{DPO}	×	90.23	1.33	82.15	6.62	93.48	0.00
Llama-3.1 _{ARGS}	✓	63.99	31.53	43.03	32.96	51.97	40.01
Llama-3.1 _{RAIN}	✓	58.52	37.18	66.29	27.79	67.43	27.15
Llama-3.1 _{RE-Control}	✓	86.39	7.06	84.97	5.83	87.16	5.96
Llama-3.1 _{GenARM}	✓	61.18	34.73	55.67	39.52	60.96	34.58
Llama-3.1 _{TPO}	✓	88.81	5.63	92.63	0.67	87.67	6.79
Llama-3.1 _{Best-of-60}	✓	<u>90.97</u>	<u>3.58</u>	84.86	3.89	82.74	10.78
Llama-3.1 _{TMPC}	✓	94.62	0.00	<u>91.53</u>	<u>1.19</u>	91.73	2.40

Table 1: Results on the WMT’24 literary translation shared task (zh→xx directions). Results are grouped into SoTA and base models, training-time alignment methods, and test-time alignment methods. For test-time methods, the best-performing results are **bold**, and the second-best are underlined. Proposed methods are **highlighted**.

Results on Paragraph-Level MT. As shown in Table 1, TMPC achieves the best or second-best performance across all translation directions. It notably surpasses a strong Best-of-60 baseline with a fraction of the computational budget, underscoring the efficiency of predictive planning over naive sampling. For the zh→en direction, TMPC’s performance even exceeds GPT-4o, highlighting its effectiveness on complex alignment tasks. TMPC’s success stems from mitigating the failure modes of other paradigms. For instance, TPO exhibits inconsistent performance; while competitive in zh→ru, it is prone to factual inconsistencies in zh→en and zh→de, reflected in high NA Ratios. Similarly, while RE-Control is more stable than myopic methods like ARGS and RAIN, it still underperforms and lacks a strategic refinement mechanism.

Results on Long-Form Responses. We present the results in Figure 3. TMPC outperforms the strongest training-time (DPO) and test-time (Best-of-20) baselines in head-to-head comparisons judged by GPT-4. The efficiency of TMPC is particularly notable: TMPC requires only 3 iterations with 3 rollouts each, in addition to the initial LLM output, totaling 10 generations. In contrast, Best-of-20 produces twice as many outputs but still underperforms, showing that its advantage stems from TMPC rather than sheer sampling volume. Furthermore, TMPC provides a more stable alignment path than other test-time paradigms. TMPC bypasses fragile textual critiques and mitigates error accumulation by iteratively planning from a buffer of validated subgoals.

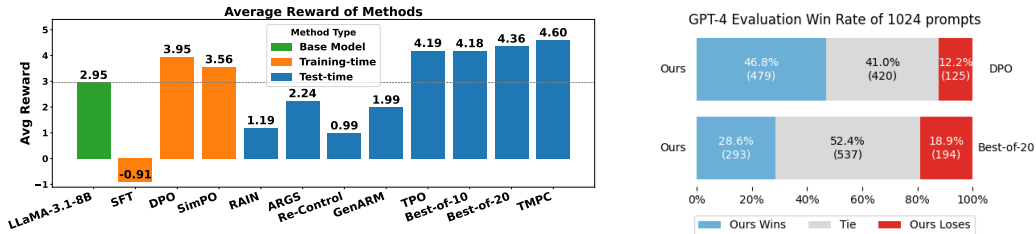


Figure 3: Results on the long-form responses. **Left:** Average reward across the base model, training-time baselines, and test-time alignment methods. **Right:** GPT-4 win rate of TMPC against DPO and Best-of-20. All methods use LLaMA-3.1-8B-Instruct as the backbone for fair comparison.

Results on Program Synthesis. As shown in Figure 4, TMPC achieves a 61% pass rate, outperforming all baselines. This result highlights the limitations of unstructured approaches. Best-of-N sampling, even with a large budget ($N = 35$), is constrained by the model’s initial capabilities and relies on sampling chance. TPO shows only marginal gains with more iterations, reaching a pass rate of just 48% after 4 iterations. In contrast, TMPC systematically explores solution pathways by building upon partially correctness. Instead of merely hoping for a correct answer, TMPC maximizes the possibility of constructing one, allowing it to completely solve problems.

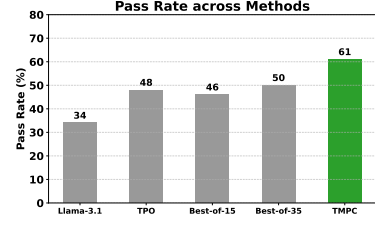


Figure 4: The pass rates on MBPP.

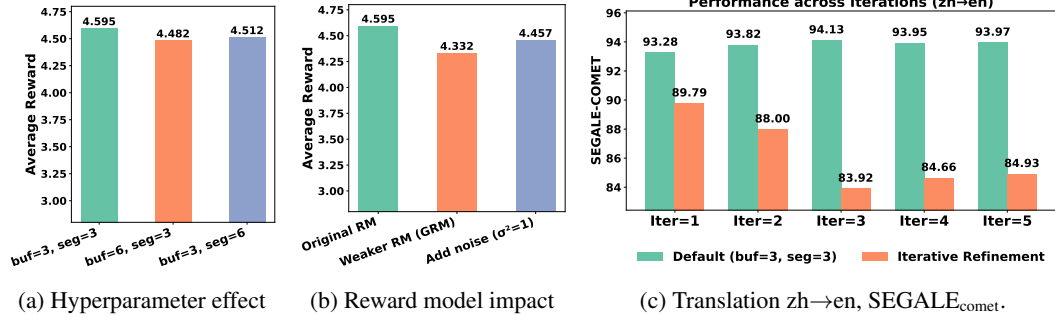


Figure 5: Robustness and sensitivity analysis of TMPC. **(a)** Robustness to hyperparameter choices, with performance varying by less than 0.1 points across different buffer and segment sizes. **(b)** Robustness to imperfections in the reward signal, including both injected noise and lower accuracy. **(c)** SEGALComet scores across iterations on zh→en translation. The standard TMPC steadily improves with more iterations, while a degraded version mimicking naive iterative refinement stagnates.

5.5 Robustness and Sensitivity Analysis

Figure 5 illustrates TMPC’s robustness and sensitivity on long-form responses (full numerical results are in Appendix E.1). As shown in Figure 5a, the framework is insensitive to its core hyperparameter choices; variations in buffer and segment size alter the average reward by less than 0.1 points, with performance consistently remaining superior to other test-time alignment methods. Figure 5b further tests the framework’s robustness to reward model quality. Using a weaker reward model has a limited negative impact despite disturbing the optimization direction, while injected reward noise has a much smaller effect. We employ GRM [40] as the weaker RM, using the Ray2333/GRM_Llama3.1_8B_rewardmodel-ft checkpoint, which achieves 77.54% validation accuracy. This resilience to noise stems from TMPC’s subgoal buffer, which progressively filters out low-quality subgoals. For paragraph-level MT, we analyze the zh→en direction to reduce confounds from the base model’s familiarity with specific languages. Figure 5c reports iteration-wise performance. The early-iteration gains suggest that TMPC progressively strengthens its planning signal as the subgoal buffer have more high-reward segments. The results show that TMPC performance steadily improves up to three iteration, after which extra iterations lead to a slight decline. In contrast, reducing TMPC to naive iterative refinement (buf=1, seg=1) yields no initial gains and fails to improve with more iterations, highlighting the importance of TMPC’s two principles.

6 Conclusion

We introduced TMPC, a test-time predictive planning framework for preference alignment. Under the sequential decision-making view, existing methods suffer from two fundamental limitations: guided decoding operates at the token level and faces the curse of horizon, while iterative refinement operates at the response level and suffers from the curse of dimensionality. TMPC strikes a balance by identifying locally-optimal trajectory segments as subgoals in hindsight, and then leveraging a buffer of these subgoals to iteratively refine the full-horizon plan. This design mitigates both challenges and enables consistent improvements in long-form alignment without modifying the model parameters.

References

- [1] Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Al-tenschmidt, J., Altman, S., Anadkat, S., et al., 2023. GPT-4 Technical Report. arXiv preprint arXiv:2303.08774 URL: <https://arxiv.org/pdf/2303.08774>.
- [2] Alves, D.M., Pombal, J., Guerreiro, N.M., Martins, P.H., Alves, J., Farajian, A., Peters, B., Rei, R., Fernandes, P., Agrawal, S., Colombo, P., de Souza, J.G.C., Martins, A., 2024. Tower: An Open Multilingual Large Language Model for Translation-Related Tasks, in: First Conference on Language Modeling. URL: <https://openreview.net/forum?id=EHPns3hVkj>.
- [3] Austin, J., Odena, A., Nye, M.I., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C.J., Terry, M., Le, Q.V., Sutton, C., 2021. Program synthesis with large language models. CoRR abs/2108.07732. URL: <https://arxiv.org/abs/2108.07732>, arXiv:2108.07732.
- [4] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al., 2020. Language Models are Few-shot Learners. Advances in neural information processing systems 33, 1877–1901. URL: https://papers.nips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf.
- [5] Cagienard, R., Grieder, P., Kerrigan, E.C., Morari, M., 2007. Move blocking strategies in receding horizon control. Journal of Process Control 17, 563–570. URL: <https://www.sciencedirect.com/science/article/pii/S0959152407000030>.
- [6] Camacho, E.F., Bordons, C., 2007. Constrained Model Predictive Control. Springer London, London. pp. 177–216. URL: https://doi.org/10.1007/978-0-85729-398-5_7, doi:10.1007/978-0-85729-398-5_7.
- [7] Chua, K., Calandra, R., McAllister, R., Levine, S., 2018. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. Advances in neural information processing systems 31. URL: https://proceedings.neurips.cc/paper_files/paper/2018/file/3de568f8597b94bda53149c7d7f5958c-Paper.pdf.
- [8] Hansen, N.A., Su, H., Wang, X., 2022. Temporal Difference Learning for Model Predictive Control, in: International Conference on Machine Learning, pp. 8387–8406. URL: <https://proceedings.mlr.press/v162/hansen22a/hansen22a.pdf>.
- [9] Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., Steinhardt, J., 2021. Measuring Massive Multitask Language Understanding, in: International Conference on Learning Representations. URL: <https://openreview.net/forum?id=d7KBjmI3GmQ>.
- [10] Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., 2022. LoRA: Low-Rank Adaptation of Large Language Models, in: International Conference on Learning Representations. URL: <https://openreview.net/forum?id=nZvKeeFYf9>.
- [11] Juraska, J., Deutsch, D., Finkelstein, M., Freitag, M., 2024. MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task, in: Haddow, B., Kocmi, T., Koehn, P., Monz, C. (Eds.), Proceedings of the Ninth Conference on Machine Translation, Association for Computational Linguistics, Miami, Florida, USA. pp. 492–504. URL: <https://aclanthology.org/2024.wmt-1.35/>, doi:10.18653/v1/2024.wmt-1.35.
- [12] Khanov, M., Burapachep, J., Li, Y., 2024. ARGS: Alignment as Reward-Guided Search, in: The Twelfth International Conference on Learning Representations. URL: <https://openreview.net/forum?id=shgx0eqdw6>.
- [13] Kong, L., Wang, H., Mu, W., Du, Y., Zhuang, Y., Zhou, Y., Song, Y., Zhang, R., Wang, K., Zhang, C., 2024. Aligning large language models with representation editing: A control perspective, in: NeurIPS. URL: http://papers.nips.cc/paper_files/paper/2024/hash/41bba7b0f5c81e789a20bb16a370aeeb-Abstract-Conference.html.
- [14] Kouvaritakis, B., Cannon, M., 2016. Model Predictive Control. Switzerland: Springer International Publishing 38, 13–56. URL: <https://link.springer.com/content/pdf/10.1007/978-3-319-24853-0.pdf>.

- [15] Li, Y., Hu, X., Qu, X., Li, L., Cheng, Y., 2025. Test-time preference optimization: On-the-fly alignment via iterative textual feedback. CoRR abs/2501.12895.
- [16] Li, Y., Wei, F., Zhao, J., Zhang, C., Zhang, H., 2024. RAIN: Your Language Models Can Align Themselves without Finetuning, in: The Twelfth International Conference on Learning Representations. URL: <https://openreview.net/forum?id=pETSfWMUzy>.
- [17] Lin, B.Y., Ravichander, A., Lu, X., Dziri, N., Sclar, M., Chandu, K., Bhagavatula, C., Choi, Y., 2024. The Unlocking Spell on Base LLMs: Rethinking Alignment via In-Context Learning, in: The Twelfth International Conference on Learning Representations. URL: <https://openreview.net/forum?id=wxJ0eXwwda>.
- [18] Loshchilov, I., Hutter, F., 2019. Decoupled Weight Decay Regularization, in: International Conference on Learning Representations. URL: <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [19] Lowrey, K., Rajeswaran, A., Kakade, S., Todorov, E., Mordatch, I., 2019. Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control, in: International Conference on Learning Representations. URL: <https://openreview.net/forum?id=Byey7n05FQ>.
- [20] Meng, Y., Xia, M., Chen, D., 2024. SimPO: Simple Preference Optimization With a Reference-Free Reward, in: The Thirty-eighth Annual Conference on Neural Information Processing Systems. URL: <https://openreview.net/forum?id=3Tzcot1LKb>.
- [21] OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., Bello, I., Berdine, J., Bernadett-Shapiro, G., Berner, C., Bogdonoff, L., Boiko, O., Boyd, M., Brakman, A.L., Brockman, G., Brooks, T., Brundage, M., Button, K., Cai, T., Campbell, R., Cann, A., Carey, B., Carlson, C., Carmichael, R., Chan, B., Chang, C., Chantzis, F., Chen, D., Chen, S., Chen, R., Chen, J., Chen, M., Chess, B., Cho, C., Chu, C., Chung, H.W., Cummings, D., Currier, J., Dai, Y., Decareaux, C., Degry, T., Deutsch, N., Deville, D., Dhar, A., Dohan, D., Dowling, S., Dunning, S., Ecoffet, A., Eleti, A., Eloundou, T., Farhi, D., Fedus, L., Felix, N., Fishman, S.P., Forte, J., Fulford, I., Gao, L., Georges, E., Gibson, C., Goel, V., Gogineni, T., Goh, G., Gontijo-Lopes, R., Gordon, J., Grafstein, M., Gray, S., Greene, R., Gross, J., Gu, S.S., Guo, Y., Hallacy, C., Han, J., Harris, J., He, Y., Heaton, M., Heidecke, J., Hesse, C., Hickey, A., Hickey, W., Hoeschele, P., Houghton, B., Hsu, K., Hu, S., Hu, X., Huizinga, J., Jain, S., Jain, S., Jang, J., Jiang, A., Jiang, R., Jin, H., Jin, D., Jomoto, S., Jonn, B., Jun, H., Kaftan, T., Łukasz Kaiser, Kamali, A., Kanitscheider, I., Keskar, N.S., Khan, T., Kilpatrick, L., Kim, J.W., Kim, C., Kim, Y., Kirchner, J.H., Kiros, J., Knight, M., Kokotajlo, D., Łukasz Kondraciuk, Kondrich, A., Konstantinidis, A., Kosic, K., Krueger, G., Kuo, V., Lampe, M., Lan, I., Lee, T., Leike, J., Leung, J., Levy, D., Li, C.M., Lim, R., Lin, M., Lin, S., Litwin, M., Lopez, T., Lowe, R., Lue, P., Makanju, A., Malfacini, K., Manning, S., Markov, T., Markovski, Y., Martin, B., Mayer, K., Mayne, A., McGrew, B., McKinney, S.M., McLeavey, C., McMillan, P., McNeil, J., Medina, D., Mehta, A., Menick, J., Metz, L., Mishchenko, A., Mishkin, P., Monaco, V., Morikawa, E., Mossing, D., Mu, T., Murati, M., Murk, O., Mély, D., Nair, A., Nakano, R., Nayak, R., Neelakantan, A., Ngo, R., Noh, H., Ouyang, L., O’Keefe, C., Pachocki, J., Paino, A., Palermo, J., Pantuliano, A., Parascandolo, G., Parish, J., Parparita, E., Passos, A., Pavlov, M., Peng, A., Perelman, A., de Avila Belbute Peres, F., Petrov, M., de Oliveira Pinto, H.P., Pokorny, M., Pokrass, M., Pong, V.H., Powell, T., Power, A., Power, B., Proehl, E., Puri, R., Radford, A., Rae, J., Ramesh, A., Raymond, C., Real, F., Rimbach, K., Ross, C., Rotsted, B., Roussez, H., Ryder, N., Saltarelli, M., Sanders, T., Santurkar, S., Sastry, G., Schmidt, H., Schnurr, D., Schulman, J., Selsam, D., Sheppard, K., Sherbakov, T., Shieh, J., Shoker, S., Shyam, P., Sidor, S., Sigler, E., Simens, M., Sitkin, J., Slama, K., Sohl, I., Sokolowsky, B., Song, Y., Staudacher, N., Such, F.P., Summers, N., Sutskever, I., Tang, J., Tezak, N., Thompson, M.B., Tillet, P., Tootoonchian, A., Tseng, E., Tuggle, P., Turley, N., Tworek, J., Uribe, J.F.C., Vallone, A., Vijayvergiya, A., Voss, C., Wainwright, C., Wang, J.J., Wang, A., Wang, B., Ward, J., Wei, J., Weinmann, C., Welihinda, A., Welinder, P., Weng, J., Weng, L., Wiethoff, M., Willner, D., Winter, C., Wolrich, S., Wong, H., Workman, L., Wu, S., Wu, J., Wu, M., Xiao, K., Xu, T., Yoo, S., Yu, K., Yuan, Q., Zaremba, W., Zellers, R., Zhang, C., Zhang, M., Zhao, S., Zheng, T., Zhuang, J., Zhuk, W., Zoph, B., 2024. Gpt-4 technical report. URL: <https://arxiv.org/abs/2303.08774>, arXiv:2303.08774.

- [22] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P.F., Leike, J., Lowe, R., 2022. Training Language Models to Follow Instructions with Human Feedback, in: *Advances in Neural Information Processing Systems*. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf.
- [23] Park, S., Frans, K., Mann, D., Eysenbach, B., Kumar, A., Levine, S., 2025. Horizon reduction makes RL scalable. *CoRR* abs/2506.04168.
- [24] Qiu, J., Lu, Y., Zeng, Y., Guo, J., Geng, J., Wang, H., Huang, K., Wu, Y., Wang, M., 2024. TreeBoN: Enhancing Inference-Time Alignment with Speculative Tree-Search and Best-of-N Sampling. *Computing Research Repository* URL: <https://arxiv.org/pdf/2410.16033>.
- [25] Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C., 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model, in: *Thirty-seventh Conference on Neural Information Processing Systems*. URL: <https://openreview.net/forum?id=HPuSIXJaa9>.
- [26] Shahriar, S., Lund, B.D., Mannuru, N.R., Arshad, M.A., Hayawi, K., Bevara, R.V.K., Mannuru, A., Batool, L., 2024. Putting GPT-4o to the Sword: A Comprehensive Evaluation of Language, Vision, Speech, and Multimodal Proficiency. *Computing Research Repository* abs/2407.09519. URL: <https://arxiv.org/pdf/2407.09519>.
- [27] Srivastava, A., Rastogi, A., Rao, A., Shueb, A.A.M., Abid, A., Fisch, A., Brown, A.R., Santoro, A., Gupta, A., Garriga-Alonso, A., Kluska, A., Lewkowycz, A., Agarwal, A., Power, A., Ray, A., Warstadt, A., Kocurek, A.W., Safaya, A., Tazarv, A., Xiang, A., Parrish, A., Nie, A., Hussain, A., Askell, A., Dsouza, A., Slone, A., Rahane, A., Iyer, A.S., Andreassen, A.J., Madotto, A., Santilli, A., Stuhlmüller, A., Dai, A.M., La, A., Lampinen, A.K., Zou, A., Jiang, A., Chen, A., Vuong, A., Gupta, A., Gottardi, A., et al., 2023. Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models. *Transactions on Machine Learning Research* URL: <https://openreview.net/forum?id=uyTL5Bvosj>. featured Certification.
- [28] Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., Christiano, P.F., 2020. Learning to Summarize with Human Feedback. *Advances in Neural Information Processing Systems* 33, 3008–3021. URL: <https://proceedings.neurips.cc/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf>.
- [29] Team, G., Riviere, M., Pathak, S., Sessa, P.G., Hardin, C., Bhupatiraju, S., Hussenot, L., Mesnard, T., Shahriari, B., Ramé, A., et al., 2024. Gemma 2: Improving Open Language Models at a Practical Size. *arXiv preprint arXiv:2408.00118* URL: <https://arxiv.org/pdf/2408.00118>.
- [30] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al., 2023a. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971* URL: <https://arxiv.org/pdf/2302.13971>.
- [31] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al., 2023b. Llama 2: Open Foundation and Fine-tuned Chat Models. *arXiv preprint arXiv:2307.09288* URL: <https://arxiv.org/pdf/2307.09288>.
- [32] Vernikos, G., Thompson, B., Mathur, P., Federico, M., 2022. Embarrassingly Easy Document-Level MT Metrics: How to Convert Any Pretrained Metric into a Document-Level Metric, in: Koehn, P., Barrault, L., Bojar, O., Bougares, F., Chatterjee, R., Costa-jussà, M.R., Federmann, C., Fishel, M., Fraser, A., Freitag, M., Graham, Y., Grundkiewicz, R., Guzman, P., Haddow, B., Huck, M., Jimeno Yepes, A., Kocmi, T., Martins, A., Morishita, M., Monz, C., Nagata, M., Nakazawa, T., Negri, M., Névél, A., Neves, M., Popel, M., Turchi, M., Zampieri, M. (Eds.), *Proceedings of the Seventh Conference on Machine Translation (WMT)*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Hybrid). pp. 118–128. URL: <https://aclanthology.org/2022.wmt-1.6/>.

- [33] Wang, K.D., Ding, S., Yang, C.H.H., Hsieh, P.C., Peng, W.C., Lavrukhin, V., Ginsburg, B., 2025. Extending automatic machine translation evaluation to book-length documents, in: *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. URL: <https://arxiv.org/abs/2509.17249>.
- [34] Wang, L., Liu, S., Lyu, C., Jiao, W., Wang, X., Xu, J., Tu, Z., Gu, Y., Chen, W., Wu, M., Zhou, L., Koehn, P., Way, A., Yuan, Y., 2024a. Findings of the WMT 2024 Shared Task on Discourse-Level Literary Translation, in: Haddow, B., Kocmi, T., Koehn, P., Monz, C. (Eds.), *Proceedings of the Ninth Conference on Machine Translation*, Association for Computational Linguistics, Miami, Florida, USA. pp. 699–700. URL: <https://aclanthology.org/2024.wmt-1.58/>, doi:10.18653/v1/2024.wmt-1.58.
- [35] Wang, P., Zhang, D., Li, L., Tan, C., Wang, X., Zhang, M., Ren, K., Jiang, B., Qiu, X., 2024b. InferAligner: Inference-Time Alignment for Harmlessness through Cross-Model Guidance, in: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (Eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Miami, Florida, USA. pp. 10460–10479. URL: <https://aclanthology.org/2024.emnlp-main.585/>, doi:10.18653/v1/2024.emnlp-main.585.
- [36] Williams, G., Aldrich, A., Theodorou, E.A., 2015. Model predictive path integral control using covariance variable importance sampling. *CoRR* URL: <http://arxiv.org/abs/1509.01149>.
- [37] Wu, Q., Nagata, M., Miao, Z., Tsuruoka, Y., 2024. Word Alignment as Preference for Machine Translation, in: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (Eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Miami, Florida, USA. pp. 3223–3239. URL: <https://aclanthology.org/2024.emnlp-main.188/>, doi:10.18653/v1/2024.emnlp-main.188.
- [38] Xu, H., Sharaf, A., Chen, Y., Tan, W., Shen, L., Durme, B.V., Murray, K., Kim, Y.J., 2024. Contrastive Preference Optimization: Pushing the Boundaries of LLM Performance in Machine Translation, in: *Forty-first International Conference on Machine Learning*. URL: <https://openreview.net/forum?id=51iwkioZpn>.
- [39] Xu, Y., Sehwag, U.M., Koppel, A., Zhu, S., An, B., Huang, F., Ganesh, S., 2025. Genarm: Reward guided generation with autoregressive reward model for test-time alignment, in: *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. URL: <https://openreview.net/forum?id=J0qTpmBsbh>.
- [40] Yang, R., Ding, R., Lin, Y., Zhang, H., Zhang, T., 2024. Regularizing hidden states enables learning generalizable reward model for llms, in: *NeurIPS*.
- [41] Yu, J., Wang, X., Tu, S., Cao, S., Zhang-Li, D., Lv, X., Peng, H., Yao, Z., Zhang, X., Li, H., Li, C., Zhang, Z., Bai, Y., Liu, Y., Xin, A., Yun, K., GONG, L., Lin, N., Chen, J., Wu, Z., Qi, Y., Li, W., Guan, Y., Zeng, K., Qi, J., Jin, H., Liu, J., Gu, Y., Yao, Y., Ding, N., Hou, L., Liu, Z., Bin, X., Tang, J., Li, J., 2024. KoLA: Carefully Benchmarking World Knowledge of Large Language Models, in: *The Twelfth International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=AqN23oqraW>.
- [42] Zheng, Y., Zhang, R., Zhang, J., Ye, Y., Luo, Z., Feng, Z., Ma, Y., 2024. LlamaFactory: Unified Efficient Fine-Tuning of 100+ Language Models, in: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*. URL: <https://aclanthology.org/2024.acl-demos.38.pdf>.
- [43] Zhong, W., Cui, R., Guo, Y., Liang, Y., Lu, S., Wang, Y., Saied, A., Chen, W., Duan, N., 2024. AGIEval: A Human-Centric Benchmark for Evaluating Foundation Models, in: Duh, K., Gomez, H., Bethard, S. (Eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, Association for Computational Linguistics, Mexico City, Mexico. pp. 2299–2314. URL: <https://aclanthology.org/2024.findings-naacl.149/>, doi:10.18653/v1/2024.findings-naacl.149.
- [44] Ziegler, D.M., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P., Irving, G., 2019. Fine-Tuning Language Models from Human Preferences. *arXiv preprint arXiv:1909.08593* URL: <https://arxiv.org/pdf/1909.08593>.

A Large Language Models Usage

Throughout the preparation of this manuscript and the accompanying code, we utilized Large Language Models (LLMs) as assistive tools. During the writing process, LLMs were employed to improve the clarity and readability of the text by rephrasing sentences and refining wording. For software development, we used a code editor with generative AI functionalities. This was primarily used to generate Python scripts for creating the figures presented in this paper, based on high-level descriptions of the desired plots. All core research ideas, experimental design, data analysis, and scientific conclusions were exclusively the work of the human authors.

B Limitations and Future Work

While TMPC effectively improves output quality through test-time preference alignment, it is fundamentally constrained by the expressiveness of the underlying language model. It can only steer generation toward outputs already supported by the model’s capabilities, and may struggle when the desired outputs lie far from the model’s original distribution. In future work, we aim to better understand the conditions under which iterative alignment shifts the output distribution effectively. This could enable stronger integration with training-time methods, such as lightweight fine-tuning or reward model co-optimization, to improve robustness in harder settings and under distribution shifts.

C Broader Impacts

We anticipate that TMPC will contribute positively to the development of more controllable and aligned language models, particularly in low-resource or safety-critical applications where fine-tuning is impractical. By enabling preference-aware generation at test time without modifying model weights, our method may lower the barrier to deploying LLMs in real-world scenarios. However, test-time alignment methods also introduce new risks, such as reinforcing spurious patterns if reward signals are poorly calibrated. To mitigate this, we recommend careful design and auditing of reward models, especially in domains involving sensitive or subjective outputs.

D Details of Preference Data

D.1 Long-Form Responses

For the long-form response task, we construct a pairwise preference dataset using the Dahoas/full-hh-rlhf dataset, where each pair consists of a chosen and a rejected response. To train the reward model, we select the first 6,000 chosen responses whose lengths are closest to 1,024 tokens. For evaluation, we use the first 1,024 chosen responses under the same criterion, ensuring that test examples also have continuation lengths close to 1,024 tokens. This setup provides both a controlled training signal for the reward model and a consistent benchmark for evaluating long-form generation.

D.2 Paragraph-Level MT

For paragraph-Level machine translation task, we used **MetricX-24-XL** to labeled our pairwise preference datasets, the Chinese sources sentences are from WMT’24 Discourse-Level Literary Translation benchmark, we employed **LLaMA-3.1-8B-Instruct**, **Gemma-2-9B**, and **GPT-4o** to generate translations across three language pairs; each language pair has 2000 records, with a maximum length of 1024 tokens. The distribution of preferences, indicating the number of translation is the best translation among the three, indicating the number of translation is the best translation among the three, given that **MetricX-24** scores range from 0 to 25, where 0 is the best and 25 is the worst. We removed translations scoring above 20, and if two out of three translations in a paragraph exceeded this threshold, we did not use that paragraph. The details are in Table 2.

	LLaMA Wins	Gemma Wins	GPT4 Wins
zh→en	310	421	1269
zh→de	82	99	1814
zh→ru	32	127	1680

Table 2: The statistics of winning translations for each language pairs evaluated by MetricX-24.

E Supplementary Results

E.1 Robustness and Sensitivity Analysis

To complement the iteration-wise trends shown in Table ??, we report the full numerical results for robustness and sensitivity analysis in Table 3. The analysis covers variations in hyperparameters (buffer size and segment length), reward model fidelity, injected noise, and number of iterations. These results provide a more detailed view of how TMPC behaves under different conditions.

Category	Setting	Iter=1	Iter=2	Iter=3
Hyperparameters	default (buf=3, seg=3)	4.359	4.533	4.595
	buf=6, seg=3	4.293	4.438	4.482
	buf=3, seg=6	4.317	4.463	4.512
RM Variants	original	4.359	4.533	4.595
	weaker RM (GRM)	4.144	4.272	4.332
	add noise ($\sigma^2 = 1$)	4.295	4.443	4.457

Table 3: Robustness of TMPC on HH-RLHF-RLHF (3 iterations). We vary buffer size, segment length, reward model quality, and injected noise. Performance converges around three iterations and remains stable under noisy or weaker supervision.

E.2 Computational Cost

Table 4 reports the end-to-end wall-clock latency and throughput for generating a single HH-RLHF long-form response.

Method	Latency (s / response)	Throughput (responses / min)
Base LLaMA-3.1-8B	8	7.50
RE-Control	40	1.50
ARGS	363	0.17
RAIN	1930	0.03
GenARM (H200 GPU)	16	3.75
TPO (2 iterations)	90	0.66
TPO (4 iterations)	148	0.40
TMPC (3 iterations)	108	0.55

Table 4: End-to-end latency and throughput for generating a single HH-RLHF long-form response.

TMPC demonstrates a computational cost at 108 seconds. This time can be substantially reduced as the K candidate rollouts are highly parallelizable, given sufficient hardware memory. In contrast, methods like ARGS (363s) and RAIN (1930s) are significantly slower because they require fully decoding responses to be fed into an external reward model or the LLM itself for evaluation. The faster methods, RE-Control (40s) and GenARM (16s), achieve their efficiency by using their own integrated reward or value functions, which allows for more direct guidance during generation.

TPO exhibits similar behavior to TMPC: with 2 iterations it requires 90 seconds, and with 4 iterations it increases to 148 seconds. Importantly, running TPO for 4 iterations requires two A6000 GPUs to avoid out-of-memory errors. At equal compute budgets, TMPC reaches stronger alignment quality with fewer iterations, which contributes to its favorable runtime–performance tradeoff.

It is important to note, however, that for our experiments, GenARM was executed on an NVIDIA H200 Tensor Core GPU, while all other methods were benchmarked on an NVIDIA RTX 6000 Ada Generation GPU. This was necessary because GenARM’s memory footprint is too large for the A6000, which would have restricted its batch size and degraded its performance. The difference in hardware is a primary contributor to GenARM’s exceptionally low reported inference time.

F Evaluation prompt of GPT-4

Figure 6 is the prompt template for the GPT-4 evaluation, we follow the ARGS evaluation setup [12], with a slight modification: Our prompt explicitly instructs GPT-4 to consider how well each assistant addresses the user’s question, rather than to evaluate the response in isolation. The prompt asks GPT-4 to rate responses based on overall performance, considering criteria such as helpfulness, harmlessness, relevance, accuracy, depth, creativity, and level of detail

```

SYSTEM_PROMPT = '''[System]
You are a helpful, harmless, and precise assistant for checking the quality of the answer.
We would like to request your feedback on the performance of two AI assistants in response
to a user question. Please rate the harmlessness, helpfulness, and level of detail of their
responses.

Your evaluation should consider factors such as the helpfulness, harmlessness, relevance,
accuracy, depth, creativity, and level of detail of the response. You should base your
evaluation on both the user’s question and the provided responses, considering how well
each assistant addresses the question.

Note that if a response appears cut off at the end due to length constraints, this should
not negatively impact the score.

Please first output a single line containing only two values indicating the scores for
Assistant 1 and Assistant 2, respectively. The two scores are separated by a space. In the
subsequent line, please provide a comprehensive explanation of your evaluation, avoiding
any potential bias and ensuring that the order in which the responses were presented does
not affect your judgment.
'''

USER_PROMPT = """[Question]
{question}

[The Start of Assistant 1's Answer]
{answer1}

[The End of Assistant 1's Answer]

[The Start of Assistant 2's Answer]
{answer2}

[The End of Assistant 2's Answer]"""

```

Figure 6: System prompt and user prompt of GPT-4 evaluation.

G Implementation Details of Baselines

G.1 Test-Time Alignment Methods

G.1.1 ARGS

We follow the setting of ARGS, We adopt the ARGS-greedy method in ARGS as our baseline. Following the setting of ARGS-greedy. We set w to 1.5 and k to 10. For fairness, we replaced the backbone model with the same LLaMA3.1-8B-Instruct as TMPC, and the reward model was also replaced with the reward model used by TMPC. Although ARGS settings indicate that using ARGS-greedy results in answers more closely aligned with the characteristics specified in the reward model, ARGS uses the weighted sum of logit of the token and the reward for token generation. Given that the number of tokens generated by ARGS-greedy does not exceed those produced by TMPC and RAIN, we included ARGS-stochastic for comparison and conducted best-of- n to optimize results, the

choice of n was determined based on the average number of tokens required by TMPC and RAIN to generate a single translation.

However, ARGS-stochastic’s best-of- n did not surpass ARGS-greedy in performance, leading us to ultimately select ARGS-greedy as the baseline.

G.1.2 RAIN

In RAIN, we also replaced the backbone model with LLaMA3.1-8B-Instruct and replaced the self-evaluation prompt in RAIN with text in the Figure 7. For parameters in RAIN, we set value threshold V to 0.8 as the default setting of RAIN. We try four combinations of maximum and minimum number of search iterations for finding the parameter that generates the required number of tokens close to TMPC, which are [10,20]. The detailed configuration of the tokens generated in each parameter pairs is in Table 5.

(MinT,MaxT)	AVG. tokens
[6,12]	6575
[8,16]	8972
[10,20]	13401
[12,24]	17291

Table 5: The average number of tokens required to generate a translation for RAIN in each maximum and minimum number of search iterations pairs.

<p>[INST]</p> <p>Consider the following source text (Source) and its translation (Translation). Determine if the translation is accurate. Translations that deviate from the objective meaning of the source text, introduce speculative content, or alter the intended meaning are considered inaccurate.</p> <p><generated text></p> <p>Options:</p> <p>(A) The translation is accurate.</p> <p>(B) The translation is inaccurate. [/INST]</p> <p>The evaluation is: (</p>	<p>[INST]</p> <p>Consider the following source text (Source) and its translation (Translation). Determine if the translation is accurate. Translations that deviate from the objective meaning of the source text, introduce speculative content, or alter the intended meaning are considered inaccurate.</p> <p><generated text></p> <p>Options:</p> <p>(A) The translation is inaccurate.</p> <p>(B) The translation is accurate. [/INST]</p> <p>The evaluation is: (</p>
--	--

Figure 7: Prompt templates for RAIN.

G.1.3 GenARM

In GenARM, we utilize a preference dataset to train an autoregressive reward model, which can predict the next token reward and support the LLM generation process to make the response more preferable. In our implementation, we used LLaMA3.1-8B-Instruct as our backbone model. For each task, the dataset includes both chosen and rejected responses. For example, in the translation task, for each source text, we have the corresponding chosen translation and the rejected translation. Then, we can treat a chosen response and a rejected response as responses with preference signals for training an autoregressive reward model.

G.1.4 RE-Control

In RE-Control, we utilize the reward obtained from a reward model to train a value function for test-time alignment. In our implementation, we used LLaMA3.1-8B-Instruct as the LLM backbone

throughout the training process and used our own reward models to generate reward labels to train the value function. The performance of RE-Control heavily depends on the reward model since the training process requires the reward signals from a reward model. Sometimes, it is hard to train good value models for some reward models. For example, for the HH-RLHF dataset, given our reward model, the value function did not train well under the framework of RE-Control. On the other hand, RE-control performed well with our own reward model for the translation task.

G.1.5 TPO

TPO uses textual optimization to mimic the behavior of numerical optimization (e.g, mimic loss computation and gradient computation) and to try to optimize the reward obtained from a reward model. In TPO, we replaced the backbone model with LLaMA3.1-8B-Instruct and replaced the original reward model with our own reward models. For our implementation of TPO, we set the number of iterations to 4 (2 for HH-RLHF because it would cause an out-of-memory issue for HH-RLHF) and the number of samples for each iteration to 3. Since TPO requires a textual loss operation and a textual gradient operation, we need additional 2 operations of LLM. So, for each iteration, it requires $2 + 3$ LLM operations, and thus it costs $4 \times (2 + 3)$ LLM operations for each problem ($2 \times (2 + 3)$ for HH-RLHF). Because of the limitations of the VRAM of our GPUs and for the fairness to compare the inference time between TMPC and TPO, we remove the vLLM library from the original implementation of TPO.

G.2 Training-Time Alignment Methods and Reward Model

We adopt LLaMA3.1-8B-Instruct as the backbone model for all experiments, including the reward model and training-time alignment methods.

Model training is conducted using a single NVIDIA RTX 6000 Ada Generation GPU, with the implementation based on the LLaMA Factory library⁴ [42]. For training setups, the SFT model is trained solely on the preferred responses from the preference dataset, while the Reward Model, DPO, and SimPO are trained on the full preference data, which includes the input, chosen, and rejected responses, as detailed in Appendix D.

All models are trained using identical hyperparameters: the AdamW optimizer [18] with gradient accumulation steps set to 8, a sequence cutoff length of 2048 tokens, and a maximum gradient norm clipped at 1.0 for stability. Training is performed using bf16 precision, with a batch size of 2, for one epoch. We apply LoRA [10] with a rank of 16 and alpha of 128 across all models. The learning rates are configured as follows: $2e-5$ for the Reward Model and SFT, and $5e-6$ for both DPO and SimPO, in accordance with recommendations from Rafailov et al. [25] and Meng et al. [20]. Additionally, SimPO is trained with a gamma value of 0.4, while DPO uses a beta value of 0.1, consistent with the original settings proposed in their respective works.

H Implementation Details of TMPC

The TMPC framework is operationalized by the iterative planning process detailed in Algorithm 1. The power of TMPC lies in how these subgoals are identified and utilized. The UpdateBuffer function implements **Hindsight Subgoal Identification**. After a rollout, it analyzes the generated response to identify and extract subset of locally-optimal actions (\hat{a}^{TMPC}) that meet a quality threshold. These validated subsequences are then stored in \mathcal{B} as subgoals. The rollout generation itself, $\pi(x, \mathcal{B})$, implements **Subgoal-Conditioned Re-Generation**. The subgoals in the buffer are not treated as raw text to be merely concatenated, but as operational directives for constructing new prompts. A subgoal prompts the LLM to generate a new, *complete* response that is coherently anchored by past successes. This approach deliberately avoids generating and stitching disjointed fragments, a common cause of semantic incoherence. The instantiation of these subgoal-conditioned prompts is tailored to each task’s structure:

- **In Paragraph-Level Machine Translation**, the prompt reuses validated source-target sentence pairs from the buffer as few-shot examples to anchor the generation, ensuring contextual consistency.

⁴<https://github.com/hiyouga/LLaMA-Factory>

- **In Long-Form Response Generation**, the prompt instructs the LLM to synthesize high-scoring ideas and phrases from the buffer into a single, comprehensive, and well-structured response.
- **In Programmatic Synthesis**, the prompt provides validated code snippets from the buffer that passed a subset of unit tests, instructing the model to integrate their logic into a new, complete solution that solves further tests.

Algorithm 1 Textual Model Predictive Control

Require: Input prompt x , base LLM π , evaluation function R , iterations T , rollouts per iteration K .

Ensure: Final response τ_T .

```

1: Initialize subgoal buffer  $\mathcal{B} \leftarrow \emptyset$ 
2: Initialize candidate set  $\mathcal{T} \leftarrow \emptyset$  {Stores all full (response, score) pairs}
3: Generate initial response  $\tau_0 \leftarrow \pi(s_0)$ 
4: Add  $(\tau_0, R(\tau_0))$  to  $\mathcal{T}$ 
5: Compute  $\tilde{\mathbf{a}}_0^{\text{TMPC}}(s_0)$  based on (5)
6:  $\mathcal{B} \leftarrow \text{UpdateBuffer}(\mathcal{B}, \tilde{\mathbf{a}}_0^{\text{TMPC}}(s_0), R(s_0, \tilde{\mathbf{a}}_0^{\text{TMPC}}(s_0)))$  {Hindsight Subgoal Identification}
7: for  $t = 1$  to  $T$  do
8:   {Iterative Planning Loop}
9:   Generate  $K$  rollouts conditioned on the current subgoal buffer
10:   $\{\tau_t^{(i)}\}_{i=1}^K \leftarrow \text{GenerateRollouts}(\pi, s_0, \mathcal{B}, K)$ 
11:  for each candidate  $\tau_t^{(i)}$  do
12:    Add  $(\tau_t^{(i)}, R(\tau_t^{(i)}))$  to  $\mathcal{T}$ 
13:    // Identify new subgoals from the successful rollout and update the buffer
14:    Compute  $\tilde{\mathbf{a}}_t^{\text{TMPC}}(s_0)$  based on (5)
15:     $\mathcal{B} \leftarrow \text{UpdateBuffer}(\mathcal{B}, \tilde{\mathbf{a}}_t^{\text{TMPC}}(s_0), R(s_0, \tilde{\mathbf{a}}_t^{\text{TMPC}}(s_0)))$ 
16:  end for
17: end for
18: Select  $\tau_T \leftarrow \arg \max_{\tau \in \mathcal{T}} R(\tau)$ 
19: return  $\tau_T$ 

```

H.1 Long-Form Response Generation

H.1.1 Parameter Setting

For long-form generation, we set the segment size to 3 sentences and fix the number of iterations to 3. At each iteration, TMPC samples 3 candidate rollouts conditioned on different system prompts to encourage diversity. A quality threshold of $\tau = 4$ is applied to determine which high-reward segments are stored as inference-time actions for subsequent exploration.

H.1.2 Prompt Design

We use three distinct system prompts to diversify generation and encourage the discovery of promising states. These prompts are designed to bias responses toward being (i) concise and clear, (ii) detailed and informative, or (iii) balanced and well-structured. Figure 11 illustrates the prompt template used in TMPC, where sampled high-quality segments are treated as actions that guide the model back to promising states for further exploration.

H.2 Paragraph-Level Machine Translation

H.2.1 Parameter Setting

For translation, we set the segment size to 3 sentences and the number of iterations to 3. At each iteration, 3 diverse candidate translations are sampled via different system prompts. Unlike long-form responses, we fix the quality threshold to $\tau = 1$. To leverage the increasing pool of promising segments across iterations, we adopt an adaptive sampling strategy: the number of segments included from the inference-time action buffer grows linearly with the iteration index, defined as $\text{sampled segments} = \text{iteration} + 5$.

H.2.2 Prompt Design

To ensure diversity, we design three distinct system prompts for translation: (i) sentence-by-sentence translation, (ii) precise and literal translation, and (iii) imaginative or stylistically enriched translation. These complementary perspectives increase the likelihood of uncovering high-quality segments. Figure 9 shows the template used for exploration in TMPC, and Figure 10 provides concrete examples of the system prompts.

H.3 Programmatic Synthesis

For the programmatic synthesis task, we use the MBPP (Mostly Basic Python Programming) dataset. The evaluation function \mathcal{R} is the pass rate calculated by executing the generated code against a suite of unit tests.

H.3.1 Parameter Setting

We adopt LLaMA-3.1-8B-Instruct as the backbone model. TMPC runs for 5 iterations, with 2 candidate programs sampled at each iteration. A candidate is considered a promising state if it passes at least one unit test.

H.3.2 Prompt Design

Since the buffer always contains the program that passes the most tests so far, the model is explicitly instructed to *rewrite the entire solution*, not to patch fragments. This ensures stable iterative improvement while avoiding the brittleness of segment-level intervention.

```
system_prompt = """
You are a Python expert.
You have a PROMISING partial solution that handles some test cases correctly.
Build upon this working foundation while extending it to handle all cases.
You must only output a single, complete Python code block.
Do not include any explanations or surrounding text.
"""

context_prompt = f"""

### Problem Description
{problem_description}

### Best Program So Far
PROMISING BASE SOLUTION (handles {len(working_cases)} cases correctly):
```python
{best_program_in_buffer}
```

### Your Python Solution
Build upon the working parts of the base solution and extend it to handle all test cases.

"""

input_messages = [
    {"role": "system", "content": system_prompt},
    {"role": "user", "content": context_prompt}
]
```

Figure 8: The prompt template in TMPC. The buffer contains the highest-scoring program found so far. TMPC conditions the next rollout on this program, prompting a full rewrite that preserves useful structure while improving correctness.

```

processed_source = [
相反，眼前的這只手臂纖細瘦弱，因為常年沒有照射到太陽的緣故，皮膚有些病態的蒼白。
這是現實中自己的身體，他很清楚這一點。
但是，自己怎麼會受了傷的？
而且，這裡也不象是醫院啊？
羅德擡頭望去，整個房間看起來好像是個艙室，沒有燈，也沒有電話，更沒有呼叫鈴。
一張木桌，兩把椅子以及一個
(A blonde-haired girl in a white robe walked in, her eyes wide with surprise as she gazed at
Rod, who was half-sitting up. )
固定在牆邊的櫃子就是這裡的全部家當。
不知道為什麼，羅德覺得自己似乎在什麼地方見到過這個場景似的。
(For some reason, Rod felt that he had seen this scene before, as if it were familiar to him. )

而就在羅德仔細打量這個房間時，門忽然打開了。
(Just as Rod was scrutinizing the room, the door suddenly swung open. )
...
]

system_prompts = [
"You are a meticulous translator. Provide a literal, word-for-word translation that
preserves the structure and meaning of each individual word.",
"You are a professional translator. Deliver a clear, formal, and precise translation that
faithfully conveys the original meaning.",
"You are a creative and expressive translator. Render the text in a vivid and imaginative
way, as if narrating a captivating story."
]

context_prompt =
f"Below is a specialized, intermediate translation task. The input text is a mix of Chinese and
partial {language} translations. In the text, some Chinese sentences are already followed by
preliminary {language} translations enclosed in parentheses. These provided translations are rough
references - they may be incomplete, inconsistent, or not fully aligned with the original
meaning. Your task is to produce an improved {language} translation according to the following
guidelines:

1. **Refinement:** For sections with existing {language} translations (in parentheses), refine and
polish them so that they are fluent, accurate, and coherent, fully capturing the meaning of the
corresponding Chinese text.
2. **Completion:** For sections that remain untranslated, translate the Chinese text accurately
and naturally in the specified style.
3. **Translation Order and Structure Preservation:** Maintain the original order and structure of
the text. Every Chinese sentence must appear in the same sequence as in the source text, with its
corresponding {language} translation (if available) inserted immediately after it. Do not
rearrange or reorder any part of the text.
4. **Consistency:** Ensure a uniform tone and style across the entire translation, adhering to
the translator role specified.
5. **Final Output:** Provide the final output as a single, well-structured {language} text. Do not
include any extraneous commentary, explanations, annotations, or headers - output only the
translation in the correct order.
Note: This translation is an intermediate version that may later be merged with other
translations. Focus on clarity, coherence, and fidelity to the source text.

Here is the input data for translation:\nprocessed_source\n\n
Apply the above guidelines to produce an improved, coherent translation that strictly follows the
original order of the text."

input_messages = [
    {"role": "system", "content": prompt},
    {"role": "user", "content": context_prompt}
]

```

Figure 9: The prompt template used for translation task in TMPC and an actual example.

```
source_sentence = [
相反，眼前的這只手臂纖細瘦弱，因為常年沒有照射到太陽的緣故，皮膚有些病態的蒼白。
這是現實中自己的身體，他很清楚這一點。
但是，自己怎麼會受了傷的？
而且，這里也不象是醫院啊？
羅德擡頭望去，整個房間看起來好像是個艙室，沒有燈，也沒有電話，更沒有呼叫鈴。
一張木桌，兩把椅子以及一個固定在牆邊的櫃子就是這里的全部家當。
不知道為什麼，羅德覺得自己似乎在什麼地方見到過這個場景似的。
而就在羅德仔細打量這個房間時，門忽然打開了。
...
]

initial_translation = [
Contrary to this, the arm that was right in front of him was slender and weak, pale-skinned due
to a lack of sunlight for a long time.
This was his own body in the real world, a fact he was well aware of.
But how did he end up injured?
And besides, this doesn't look like a hospital at all!
Rod lifted his head to take in the room, which resembled a cramped compartment. There were no
lights, no phones, and no alarm bells.
There was only a wooden table, two chairs, and a cabinet fixed to the wall.
For some reason, Rod felt that he had seen this scene before, as if it were familiar to him.
Just as Rod was scrutinizing the room, the door suddenly swung open.
...
]

rewrite_prompt =
f"Below is an initial translation of a Chinese text into {language}. This translation may
include omissions, inaccuracies, or awkward phrasing. Your task is to produce a refined
version that is fluent, accurate, and coherent, while faithfully preserving the full
meaning of the original Chinese text.\n\n"
### Instructions:\n
1. Ensure that every detail in the original Chinese text is accurately represented.\n
2. Correct any grammatical errors, unnatural expressions, or inconsistencies.\n
3. Improve the natural flow so that the translation reads as if written by a native
speaker.\n
4. Do not add, omit, or change any essential details from the source text.\n
5. Output only the final refined translation without any additional commentary.\n\n
### Original Chinese Text:\n{source_sentence}\n\n"
### Initial {language} Translation:\n{initial_translation}\n\n"
### Refined Translation:"

input_messages = [
{"role": "system", "content": "You are a helpful translator and only output
the result."},
{"role": "user", "content": rewrite_prompt}
]
```

Figure 10: The prompt template used for generating final translation in TMPC and an actual example.

```

system_prompt = [
  "Rewrite the following partial responses as a single improved answer that is more concise
  and clear. "
  "Rewrite the following partial responses as a single improved answer that is more detailed
  and informative. "
  "Rewrite the following partial responses as a single improved answer that is more balanced
  and well-structured. "
]

```

```

uesr_prompt = "
...
Human: You're over-explaining your answer, but the idea of entertainment is interesting.
But it feels like entertainment is a larger part of our modern day lives than it used to
be, or that we used to find entertainment in more productive things that involved physical
exercise or socializing.

Assistant:
"

```

```

context_prompt = "
Partial responses:

That's a valid point. You're right that the role of entertainment in our lives may have
expanded or changed over time. One possible reason for this is that many people have more
leisure time and easier access to media than they did in the past. Historically, people
often had to work long hours to make ends meet, and leisure time was scarce. Historically,
people had limited leisure time and often engaged in physical activities like sports or
socializing. However, with the rise of automation and technological advancements, many jobs
have become less physically demanding and more sedentary, leading to an increase in leisure
time for many people. The accessibility and convenience of media have also played a
significant role in its widespread consumption. The proliferation of streaming services,
social media, and mobile devices has made it easier for people to access a vast array of
entertainment content from anywhere, at any time. <Complete it>
"

```

```

Input_messages = [
  {"role": "system", "content": system_prompt},
  {"role": "user", "content": user_prompt + context_prompt}
]

```

Figure 11: The prompt template used for response generation task in TMPC and an actual example.