TEST-TIME ALIGNMENT FOR LARGE LANGUAGE MODELS VIA TEXTUAL MODEL PREDICTIVE CONTROL

Anonymous authors

000

001

002003004

010 011

012

013

014

016

017

018

019

021

024

025

026

027

028

029

031

032 033 034

035

037

038

040

041

042

043 044

046

047

048

050 051

052

Paper under double-blind review

ABSTRACT

Aligning Large Language Models (LLMs) with human preferences through finetuning 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 (Achiam et al., 2023; Brown et al., 2020), LLaMAs (Touvron et al., 2023a;b), and Gemma (Team et al., 2024), has demonstrated remarkable efficacy in a wide range of NLP tasks (Hendrycks et al., 2021; Srivastava et al., 2023; Stiennon et al., 2020; Yu et al., 2024; Zhong et al., 2024). 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 (Alves et al., 2024), smaller LLMs (*e.g.*, under 10B parameters) frequently suffer from omissions and semantic drift (Wu et al., 2024). Thus, aligning LLM outputs to preferences remains an essential yet challenging problem.

Training-time approaches such as Reinforcement Learning with Human Feedback (RLHF) (Ouyang et al., 2022) and Direct Preference Optimization (DPO) (Rafailov et al., 2023) have achieved strong results in aligning preferences. However, these methods are resource-intensive and require costly retraining whenever preferences or tasks change. This has spurred interest in *test-time alignment*, where outputs are adapted without updating model parameters, using strategies such as prompting (Lin et al., 2024), guided decoding (Khanov et al., 2024; Kong et al., 2024; Li et al., 2024; Wang et al., 2024b; Xu et al., 2025), or iterative refinement (Li et al., 2025).

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* (Park et al., 2025); credit assignment becomes unreliable over long

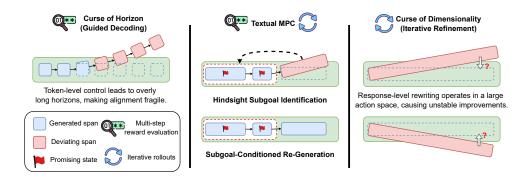


Figure 1: Textual Model Predictive Control (TMPC) balances the curse of horizon in guided decoding against the curse of dimensionality in naive iterative refinement. It employs Hindsight Subgoal Identification to dynamically discover promising states from rollouts and Subgoal-Conditioned Re-Generation to guide the search from these discovered subgoals, ensuring a stable alignment.

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) (Camacho & Bordons, 2007; Kouvaritakis & Cannon, 2016). While powerful, standard MPC assumes the problem can be decomposed into predefined, hard segments, 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. 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) (Ziegler et al., 2019) and reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022) are widely used but computationally expensive. Direct Preference Optimization (DPO) (Rafailov et al., 2023) simplifies RLHF by converting it into a supervised learning objective, though it requires managing dual policies. More recent approaches like SimPO (Meng et al., 2024) and Contrastive Preference Optimization (CPO) (Xu et al., 2024) 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 (Khanov et al., 2024) is a representative example that incorporates reward model guidance at the token level, and InferAligner (Wang et al., 2024b) adopts a similar strategy. Among guided decoding methods, there are also approaches that directly modify internal representations. For instance, RE-Control (Kong et al., 2024) trains a value function on hidden states using the Bellman equation, and applies gradient-based optimization to align with preferences. TreeBoN (Qiu et al., 2024) and RAIN (Li et al., 2024) 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 (Xu et al., 2025) 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) (Li et al., 2025) 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) $R: \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ 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., $s_t = (s_0, a_1, \cdots, a_{t-1})$. 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 $\tau = (s_0, a_0, \cdots, s_{T-1}, a_{T-1}, s_T)$ 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 (Chua et al., 2018; Lowrey et al., 2019), the goal of trajectory optimization is to find an optimal sequence of actions $\mathbf{a}^* = (a_0^*, \cdots, a_{T-1}^*)$ 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 $\tau = (s_0, a_0, \cdots, s_{T-1}, a_{T-1}, s_T)$ denote a trajectory generated under an action sequence $\mathbf{a}_{0:T-1} := (a_0, a_1, \ldots, a_{T-1})$. Given an initial prompt s_0 , the search for an optimal sequence $\mathbf{a}^*(s_0)$ can be formulated by the following optimization problem

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

Note that there is no need to take expectation in (1) as the state transitions are deterministic given $a_{0:T-1}$ 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 (Hansen et al., 2022), 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 $\boldsymbol{a}_{t:t+H}^*$ on a moving horizon H (usually H is smaller than T), given the state s_t , i.e.,

$$a^{\text{MPC}}(s_t) := \arg \max_{a_{t:t+H-1}} \sum_{i=t}^{t+H-1} R(s_t, a_t),$$
 (2)

and then select a subset of $a^{\text{MPC}}(s_t)$, denoted by $\widetilde{a}^{\text{MPC}}(s_t)$, for execution. In practice, $\widetilde{a}^{\text{MPC}}(s_t)$ can be selected as the first j contiguous actions ($1 \leq j \leq H$) or as a set of non-contiguous actions (Cagienard et al., 2007). 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 $\{\tau_t^{(i)} \equiv (s_{t:t+H-1}^{(i)}, a_{t:t+H-1}^{(i)})\}_{i=1}^K$ 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) (Williams et al., 2015) determines an approximate maximizer by performing a soft, utility-weighted aggregated selection as $a_t = \left(\sum_{i=1}^K \exp(\frac{1}{\lambda}\mathcal{J}(\tau_t^{(i)}))a_t^{(i)}\right)/\sum_{i=1}^K \exp(\frac{1}{\lambda}\mathcal{J}(\tau_t^{(i)}))$, 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.*,

$$\boldsymbol{a}^{\text{TMPC}}(s) \leftarrow \mathcal{G}\Big(\{\boldsymbol{\tau}^{(i)}\}_{i=1}^K, \{\mathcal{J}(\boldsymbol{\tau}^{(i)})\}_{i=1}^K; s\Big), \tag{3}$$

where $\{\tau^{(i)}\}_{i=1}^K$ are rollouts starting from s. Then, TMPC can leverage a sequence of non-contiguous actions, denoted by $\widetilde{a}^{\text{TMPC}}(s)$, 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

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 $\widetilde{\boldsymbol{\alpha}}^{\text{TMPC}}$. 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.

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. 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:

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

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

$$\widetilde{\boldsymbol{a}}_t^{\mathrm{TMPC}}(s) \leftarrow \mathcal{G}\Big(\{\boldsymbol{\tau}_t^{(i)}\}_{i=1}^K, R(\cdot) \mid s, \mathcal{B}\Big) := \left\{\boldsymbol{a} \mid R(s, a) \geq \alpha \text{ and } a \in \{\boldsymbol{\tau}_t^{(i)}\}_{i=1}^K\right\}, \tag{5}$$

where $\{\tau_t^{(i)}\}_{i=1}^K$ are the rollouts generated from subgoal-conditioned LLM $\pi(s,\mathcal{B})$. $\widetilde{a}_t^{\mathrm{TMPC}}(s)$ 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 $\widetilde a_t^{\rm TMPC}$ 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 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. Our framework addresses this by defining a segment abstractly through a functional milestone: the successful resolution of a single unit test.

5.1 Preference Dataset and Reward Model

Paragraph-Level MT Dataset. To construct a suitable preference dataset for long-text MT, we use the WMT'24 Discourse-Level Literary Translation benchmark (Wang et al., 2024a) for our experiments. The available language pairs include: Chinese \rightarrow English, Chinese \rightarrow German, and Chinese \rightarrow Russian. To fit within LLM context windows, each instance is segmented into up to 1024 tokens using GPT-4's tokenizer, ensuring paragraph-level MT remains within model limits.

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 (Juraska et al., 2024) under the reference-free evaluation mode, where no reference translation is supplied as input. Following the procedure in CPO Xu et al. (2024), 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-rlhf¹ 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 (Austin et al., 2021), which comprises 500 problems (Task IDs 11-510). As 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.

https://huggingface.co/datasets/Dahoas/full-hh-rlhf

5.2 EVALUATION METRICS

Paragraph-Level MT. We use SEGALE (Wang et al., 2025), a framework that extends existing metrics to long-text translation. Following CPO (Xu et al., 2024), 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. (2022) 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 (OpenAI et al., 2024). Following the ARGS evaluation protocol (Khanov et al., 2024), 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** (Khanov et al., 2024), a token-level decoding method that incorporates reward model guidance during inference. (2) **RAIN** (Li et al., 2024), which leverages tree-structured self-evaluation without relying on an external reward model. (3) **RE-Control** (Kong et al., 2024), 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** (Xu et al., 2025), 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** Li et al. (2025), 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. GenARM trains its own autoregressive reward model using the same training data employed for TMPC's reward model. 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 gradients. 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 (Meng et al., 2024) and DPO (Rafailov et al., 2023), 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-40**, which serves as a strong upper bound despite not being specialized for translation (Shahriar et al., 2024), 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 guided decoding methods are excluded as functional correctness is a holistic property of the entire code sequence, making them ill-suited for this task.

²Unbabel/wmt22-comet-da

Methods	Test-Time	$zh \to en $		$zh \to ru$		$zh \to de $	
		SEGALE _{comet} ↑	NA Ratio ↓	SEGALE _{comet} ↑	NA Ratio ↓	SEGALE _{comet} ↑	NA Ratio ↓
GPT-40 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}	\checkmark	58.52	37.18	66.29	27.79	67.43	27.15
Llama-3.1 _{RE-Control}	\checkmark	86.39	7.06	84.97	5.83	87.16	<u>5.96</u>
Llama-3.1 _{GenARM}	\checkmark	61.18	34.73	55.67	39.52	60.96	34.58
Llama-3.1 _{TPO}	\checkmark	88.81	5.63	92.63	0.67	<u>87.67</u>	6.79
Llama-3.1 _{Best-of-60}	\checkmark	90.97	3.58	84.86	3.89	82.74	10.78
Llama-3.1 _{TMPC}	\checkmark	94.62	0.00	91.53	<u>1.19</u>	91.73	2.40

Table 1: Results on the WMT'24 literary translation shared task (zh \rightarrow 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**.

5.4 QUANTITATIVE RESULTS

Results on Paragraph-Level MT. As shown in Table 1, TMPC consistently outperforms all test-time alignment baselines. 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-40, 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. TMPC inherits the stability of response-level refinement while avoiding the compounding errors of token-level guidance, striking a more effective balance.



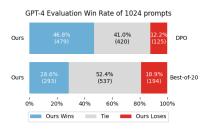


Figure 3: Results on the long-form responses. **Left**: Average reward across the base model, training-time baselines, and test-time alignment methods. **Right**: GPT-4 win rate of TMPC against DPO and Best-of-20. All methods use LLaMA-3.1-8B-Instruct as the backbone for fair comparison.

Results on 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 paradigms³. TMPC bypasses fragile textual critiques and mitigates error accumulation by iteratively planning from a buffer of validated subgoals.

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

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.

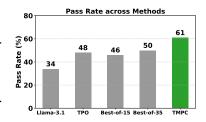


Figure 4: The pass rates on MBPP.

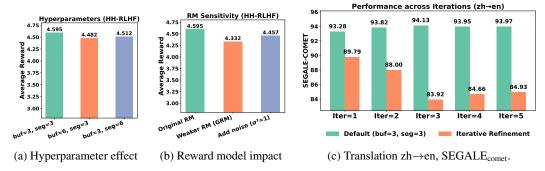


Figure 5: Robustness and sensitivity analysis of TMPC. (a) Robustness to hyperparameter choices, with performance varying by less than 0.1 points across different buffer and segment sizes. (b) Robustness to imperfections in the reward signal, including both injected noise and lower accuracy. (c) SEGALE_{comet} 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 5.5). 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 (Yang et al., 2024) 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 to illustrate the trend of improvement over time 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.

ETHICS STATEMENT

This work introduces Textual Model Predictive Control (TMPC), a general framework for the test-time alignment of large language models. Our primary goal is to develop more stable and efficient methods for aligning models with beneficial human preferences, such as helpfulness and harmlessness. All experiments were conducted on publicly available and widely used academic benchmarks (HH-RLHF, WMT'24, and MBPP), and no new data involving human subjects was collected.

The primary ethical consideration of our work is that the alignment outcome is determined by the provided reward signal. While we have used it for positive alignment, a malicious or biased reward signal could steer a model toward generating harmful, unfair, or toxic content. TMPC, like other alignment techniques, could potentially amplify biases present in the preference data used to train the reward model. We therefore stress the importance of careful design, auditing, and red-teaming of reward models before deploying systems using this technology in sensitive, real-world applications. We believe that by providing a more transparent and controllable test-time alignment mechanism, our work can contribute positively to the development of safer AI systems.

REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our research. Our implementation of TMPC, along with all experimental scripts, will be made publicly available in a permissively licensed open-source repository upon publication.

The core methodology is described in Section 4, with a detailed, task-agnostic algorithm provided in Algorithm 1. All datasets used in our experiments—HH-RLHF, WMT'24, and MBPP—are public benchmarks, with details on their specific versions and preprocessing steps provided in Section 5 and Appendix G. All hyperparameters, prompt templates, and task-specific implementation details necessary to replicate our results for long-form response generation, machine translation, and programmatic synthesis are documented in Appendix H. We believe these resources provide a clear and sufficient basis for the community to reproduce and build upon our findings.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023. URL https://arxiv.org/pdf/2303.08774.

Duarte Miguel Alves, José Pombal, Nuno M Guerreiro, Pedro Henrique Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and Andre Martins. Tower: An Open Multilingual Large Language Model for Translation-Related Tasks. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=EHPns3hVkj.

Jacob Austin, Augustus Odena, Maxwell I. Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. Program synthesis with large language models. *CoRR*, abs/2108.07732, 2021. URL https://arxiv.org/abs/2108.07732.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models are Few-shot Learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. URL https://papers.nips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

Raphael Cagienard, Pascal Grieder, Eric C Kerrigan, and Manfred Morari. Move blocking strategies in receding horizon control. *Journal of Process Control*, 17(6):563–570, 2007. URL https://www.sciencedirect.com/science/article/pii/S0959152407000030.

E. F. Camacho and C. Bordons. *Constrained Model Predictive Control*, pp. 177–216. Springer London, London, 2007. ISBN 978-0-85729-398-5. doi: 10.1007/978-0-85729-398-5_7. URL https://doi.org/10.1007/978-0-85729-398-5_7.

- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. *Advances in neural information processing systems*, 31, 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/3de568f8597b94bda53149c7d7f5958c-Paper.pdf.
 - Nicklas A Hansen, Hao Su, and Xiaolong Wang. Temporal Difference Learning for Model Predictive Control. In *International Conference on Machine Learning*, pp. 8387–8406, 2022. URL https://proceedings.mlr.press/v162/hansen22a/hansen22a.pdf.
 - Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring Massive Multitask Language Understanding. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=d7KBjmI3GmQ.
 - Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
 - Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task. In Barry Haddow, Tom Kocmi, Philipp Koehn, and Christof Monz (eds.), *Proceedings of the Ninth Conference on Machine Translation*, pp. 492–504, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.wmt-1.35. URL https://aclanthology.org/2024.wmt-1.35/.
 - Maxim Khanov, Jirayu Burapacheep, and Yixuan Li. ARGS: Alignment as Reward-Guided Search. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=shgx0eqdw6.
 - Lingkai Kong, Haorui Wang, Wenhao Mu, Yuanqi Du, Yuchen Zhuang, Yifei Zhou, Yue Song, Rongzhi Zhang, Kai Wang, and Chao Zhang. Aligning large language models with representation editing: A control perspective. In *NeurIPS*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/41bba7b0f5c81e789a20bb16a370aeeb-Abstract-Conference.html.
 - Basil Kouvaritakis and Mark Cannon. Model Predictive Control. *Switzerland: Springer International Publishing*, 38:13–56, 2016. URL https://link.springer.com/content/pdf/10.1007/978-3-319-24853-0.pdf.
 - Yafu Li, Xuyang Hu, Xiaoye Qu, Linjie Li, and Yu Cheng. Test-time preference optimization: On-the-fly alignment via iterative textual feedback. *CoRR*, abs/2501.12895, 2025.
 - Yuhui Li, Fangyun Wei, Jinjing Zhao, Chao Zhang, and Hongyang Zhang. RAIN: Your Language Models Can Align Themselves without Finetuning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=pETSfWMUzy.
 - Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. The Unlocking Spell on Base LLMs: Rethinking Alignment via In-Context Learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=wxJ@eXwwda.
 - Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
 - Kendall Lowrey, Aravind Rajeswaran, Sham Kakade, Emanuel Todorov, and Igor Mordatch. Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Byey7n05FQ.
 - Yu Meng, Mengzhou Xia, and Danqi Chen. SimPO: Simple Preference Optimization With a Reference-Free Reward. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=3Tzcot1LKb.

595

596

597

598

600

601

602

603

604

605

606

607

608

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

629

630

631

632

633

634

635

636

637 638

639

640

641

642

644

645 646

647

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training Language Models to Follow Instructions with Human Feedback. In *Advances in Neural Information Processing Systems*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf.

Seohong Park, Kevin Frans, Deepinder Mann, Benjamin Eysenbach, Aviral Kumar, and Sergey Levine. Horizon reduction makes RL scalable. *CoRR*, abs/2506.04168, 2025.

Jiahao Qiu, Yifu Lu, Yifan Zeng, Jiacheng Guo, Jiayi Geng, Huazheng Wang, Kaixuan Huang, Yue Wu, and Mengdi Wang. TreeBoN: Enhancing Inference-Time Alignment with Speculative Tree-Search and Best-of-N Sampling. *Computing Research Repository*, 2024. URL https://arxiv.org/pdf/2410.16033.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=HPuSIXJaa9.

Sakib Shahriar, Brady D. Lund, Nishith Reddy Mannuru, Muhammad Arbab Arshad, Kadhim Hayawi, Ravi Varma Kumar Bevara, Aashrith Mannuru, and Laiba Batool. Putting GPT-40 to the Sword: A Comprehensive Evaluation of Language, Vision, Speech, and Multimodal Proficiency. *Computing Research Repository*, abs/2407.09519, 2024. URL https://arxiv.org/pdf/2407.09519.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ambrose Slone, Ameet Rahane, Anantharaman S. Iyer, Anders Johan Andreassen, Andrea Madotto, Andrea Santilli, Andreas Stuhlmüller, Andrew M. Dai, Andrew La, Andrew Kyle Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, et al. Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=uyTL5Bvosj. Featured Certification.

Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to Summarize with Human Feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020. URL https://proceedings.neurips.cc/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving Open Language Models at a Practical Size. *arXiv preprint arXiv:2408.00118*, 2024. URL https://arxiv.org/pdf/2408.00118.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*, 2023a. URL https://arxiv.org/pdf/2302.13971.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open Foundation and Fine-tuned Chat Models. *arXiv preprint arXiv:2307.09288*, 2023b. URL https://arxiv.org/pdf/2307.09288.

Giorgos Vernikos, Brian Thompson, Prashant Mathur, and Marcello Federico. Embarrassingly Easy Document-Level MT Metrics: How to Convert Any Pretrained Metric into a Document-Level Metric. In Philipp Koehn, Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Tom Kocmi, André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, Matteo Negri, Aurélie Névéol, Mariana Neves, Martin Popel, Marco Turchi, and Marcos Zampieri (eds.), *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pp. 118–128, Abu Dhabi, United Arab Emirates (Hybrid), December 2022. Association for Computational Linguistics. URL https://aclanthology.org/2022.wmt-1.6/.

Kuang-Da Wang, Shuoyang Ding, Chao-Han Huck Yang, Ping-Chun Hsieh, Wen-Chih Peng, Vitaly Lavrukhin, and Boris Ginsburg. Extending automatic machine translation evaluation to book-length documents. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, November 2025. URL https://arxiv.org/abs/2509.17249.

Longyue Wang, Siyou Liu, Chenyang Lyu, Wenxiang Jiao, Xing Wang, Jiahao Xu, Zhaopeng Tu, Yan Gu, Weiyu Chen, Minghao Wu, Liting Zhou, Philipp Koehn, Andy Way, and Yulin Yuan. Findings of the WMT 2024 Shared Task on Discourse-Level Literary Translation. In Barry Haddow, Tom Kocmi, Philipp Koehn, and Christof Monz (eds.), *Proceedings of the Ninth Conference on Machine Translation*, pp. 699–700, Miami, Florida, USA, November 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.wmt-1.58. URL https://aclanthology.org/2024.wmt-1.58/.

- Pengyu Wang, Dong Zhang, Linyang Li, Chenkun Tan, Xinghao Wang, Mozhi Zhang, Ke Ren, Botian Jiang, and Xipeng Qiu. InferAligner: Inference-Time Alignment for Harmlessness through Cross-Model Guidance. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 10460–10479, Miami, Florida, USA, November 2024b. Association for Computational Linguistics. doi: 10. 18653/v1/2024.emnlp-main.585. URL https://aclanthology.org/2024.emnlp-main.585/.
- Grady Williams, Andrew Aldrich, and Evangelos A. Theodorou. Model predictive path integral control using covariance variable importance sampling. *CoRR*, 2015. URL http://arxiv.org/abs/1509.01149.
- Qiyu Wu, Masaaki Nagata, Zhongtao Miao, and Yoshimasa Tsuruoka. Word Alignment as Preference for Machine Translation. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 3223–3239, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.188. URL https://aclanthology.org/2024.emnlp-main.188/.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. Contrastive Preference Optimization: Pushing the Boundaries of LLM Performance in Machine Translation. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=51iwkioZpn.
- Yuancheng Xu, Udari Madhushani Sehwag, Alec Koppel, Sicheng Zhu, Bang An, Furong Huang, and Sumitra Ganesh. 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*, 2025. URL https://openreview.net/forum?id=J0qTpmbSbh.
- Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. Regularizing hidden states enables learning generalizable reward model for llms. In *NeurIPS*, 2024.
- Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, Chunyang Li, Zheyuan Zhang, Yushi Bai, Yantao Liu, Amy Xin, Kaifeng Yun, Linlu GONG, Nianyi Lin, Jianhui Chen, Zhili Wu, Yunjia Qi, Weikai Li, Yong Guan, Kaisheng Zeng, Ji Qi, Hailong Jin, Jinxin Liu, Yu Gu, Yuan Yao, Ning Ding, Lei Hou, Zhiyuan Liu, Xu Bin, Jie Tang, and Juanzi Li. KoLA: Carefully Benchmarking World Knowledge of Large Language Models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=AqN23oqraW.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. 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)*, 2024. URL https://aclanthology.org/2024.acl-demos.38.pdf.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. AGIEval: A Human-Centric Benchmark for Evaluating Foundation Models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 2299–2314, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.149. URL https://aclanthology.org/2024.findings-naacl.149/.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-Tuning Language Models from Human Preferences. arXiv preprint arXiv:1909.08593, 2019. URL https://arxiv.org/pdf/1909.08593.