

Sample Efficient Alignment Learning With Episodic Control

Van Dai Do¹, Quan Hung Tran², Ahmed Kirmani², Lu Zhang², Hung Le¹

¹Applied AI Initiative, Deakin University, Australia ²Meta
{v.do, thai.le}@deakin.edu.au

Abstract

Aligning large language models (LLMs) with specific task objectives is challenging, especially when access to feedback signals for guiding the model is limited. While existing parametric methods perform reasonably, they rely heavily on large datasets and frequent feedback, making them impractical in scenarios with limited human feedback. We introduce Alignment Learning with Episodic Control (ALEC), a non-parametric framework that aligns LLM outputs during inference without fine-tuning. ALEC employs a key-value memory to store the associations between generated text and its corresponding values. It leverages a novel confidence-based writing scheme to update these stored values, maximizing the use of available data. During inference, ALEC utilizes a nearest-neighbor mechanism to estimate the values of generated texts, enabling the selection of the optimal text for decoding. Our method outperforms state-of-the-art baselines on harmless, helpful, and summarization tasks, demonstrating improved alignment with minimal interactions with the true reward model.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable performance across a wide range of natural language processing tasks, with their effectiveness improving as their parameter count grows (Brown et al., 2020; Wei et al., 2022; Chowdhery et al., 2023). However, despite these advancements, pre-trained LLMs often generate responses that are misaligned with human preferences and values, largely due to the vast and diverse nature of the data used to train them (Gehman et al., 2020; Deshpande et al., 2023). This misalignment underscores the critical challenge of ensuring that LLMs meet performance benchmarks and adhere to human objectives and safety considerations. Addressing this misalignment is particularly difficult, as aligning models to human values

typically requires fine-tuning based on expert feedback, which is both costly and scarce (Casper et al., 2023). The high expense and limited availability of expert input highlight an urgent need for more efficient alignment learning methods that reduce dependence on expert feedback while ensuring that LLM outputs remain consistent with human goals.

A key approach to addressing the alignment challenge is the Reinforcement Learning from Human Feedback (RLHF) framework (Christiano et al., 2017; Stiennon et al., 2020; Ouyang et al., 2022a), where a reward model (RM) evaluates outputs based on human preferences, guiding the LLM to adjust accordingly. Direct Policy Optimization (DPO) offers an alternative by aligning models directly with preference data, avoiding RLHF’s training instability (Rafailov et al., 2024). Despite their effectiveness, both RLHF and DPO rely heavily on frequent reward model interactions and large amounts of preference data, which may not always be available. Moreover, they require fine-tuning LLMs, which can be expensive and impractical for applications with limited resources.

Another approach in alignment learning focuses on inference-time methods. For instance, Yang (2021) introduces a lightweight prefix scorer applied during inference, while Mudgal (2024) and Khanov (2024) use RL-based frameworks that combine signals from a scorer and the base model. These methods allow output refinement without fine-tuning the LLMs, but *they heavily rely on training a parametric value model* to estimate candidate quality. This approach faces severe limitations when RM calls are restricted. Small models often fail to effectively map candidates to value scores (Yang and Klein, 2021), while large, complex models are prone to overfitting on limited preference data and suffer from slow execution times (Mudgal et al., 2024; Khanov et al., 2024), making them impractical in low-resource scenarios where frequent reward queries are not feasible.

In this paper, we address the research question: *How can alignment learning be effectively achieved under constraints on the number of queries to the expert reward model?* To this end, we propose Alignment Learning with Episodic Control (ALEC), a novel and efficient training-free approach designed to perform alignment learning with minimal reliance on expert reward model queries. Drawing inspiration from the rapid and instance-based learning observed in the hippocampus region of the human brain (Lengyel and Dayan, 2007), our method avoids the complexities and overfitting risks associated with training parametric estimators while maintaining reliability through performance-driven optimization. We argue that in scenarios where sample efficiency is critical, our non-parametric, memory-based learning framework can outperform traditional parametric methods. Concretely, we leverage Episodic Control (Blundell et al., 2016), treating each training prompt-response sequence as an episode and storing its value in memory. The traditional storing method, which only inserts the episode’s states, fails to revisit or refine stored states’ values, leading to inaccuracy in estimation. To address this, we propose a confidence-based memory writing that optimizes the sample order for memory writing and updates neighboring states to reduce bias toward familiar states, thereby enhancing overall performance. At test time, the memory functions as a non-parametric nearest-neighbor model, using stored values to approximate the expert reward model. It guides generation by selecting the next token candidates that maximize the estimated values. We show the overall framework in Figure 1.

In summary, the main contributions of this paper are: (1) We propose ALEC, a training-tuning-free framework for determining optimal sequences of tokens at inference time using a nearest-neighbor reading from episodic memory. (2) We introduce a confidence-based writing scheme that enhances value estimations stored in our memory. (3) We demonstrate empirically that ALEC outperforms strong baselines, including fine-tuning and inference-time decoding methods, across three benchmarks with limited RM calls.

2 Related Work

Generator Improvement Solutions For Alignment Learning. Fine-tuning LMs to reflect human preferences has gained attention recently due to its

flexible nature. In terms of generator improvement solutions, Reinforcement Learning from Human Feedback (RLHF) offers a direct route (Ouyang et al., 2022b). Within the RLHF framework, one notable solution is to use Proximal Policy Optimization and its variants (Askell et al., 2021; Ouyang et al., 2022b; Bai et al., 2022; Shao et al., 2024). However, due to the instability and data-intensive nature of RL training, researchers have sought alternative approaches. Rafailov (2024) derives a preference formula and alleviates RL training challenges by directly optimizing LMs based on preference data.

Inference Time Alignment Learning. This approach uses inference-time solutions for alignment learning without fine-tuning the base LM, where the model is sampled multiple times to select the best output. ST-BoN (2025) improves sampling efficiency with early truncation, but relies on heuristics rather than learned control. Mudgal (2024) introduces a method that trains a prefix scorer using a Deep-Q Network (DQN) framework (Mnih et al., 2015), providing additional signals for sampling. Similarly, Khanov (2024) employs a scorer to enhance sampling, but uses a standard reward model training approach based on the Bradley-Terry model (Bradley and Terry, 1952). However, training these scorers can be data-intensive, and their performance may degrade when only a limited number of queries to the true RM are available.

3 Method

3.1 Problem formulation

Problem Settings. We adopt the standard inference-time alignment setting for LLMs as described in (Mudgal et al., 2024). In this setup, we have access to the base language model $\mathcal{L}$, a RM, a training dataset $\mathcal{D}_{\text{train}} = \{x_i\}_{i=1}^{|\mathcal{D}_{\text{train}}|}$, and a separate hold-out set $\mathcal{D}_{\text{test}}$ for evaluation. Notably, we limit the number of calls to the RM during alignment to $E = 1000$, simulating real-world scenarios where feedback from the RM, such as output from a human expert or a user, is limited.

Let $\mathbf{x}$ be a prompt and let $\mathbf{y} = y^T = [y_1, \dots, y_T]$ be the completed response consisting of T tokens, where $y_i \in \mathcal{V}$ and $\mathcal{V}$ represents the token vocabulary. The notation $\pi_{\mathcal{L}}$ denotes the pre-trained LLM $\mathcal{L}$ that is used to generate text in an auto-regressive manner. Specifically, $\pi_{\mathcal{L}}(\cdot | [\mathbf{x}, y^t])$ represents the probability over the vocabulary $\mathcal{V}$, where y^t is the response up to the

t -th token and $[\mathbf{x}, y^t]$ is the concatenation of the prompt and the (possibly unfinished) response.

Reinforcement Learning Formulation. We formulate inference-time alignment as an RL problem in which the LLM serves as the agent, with the following definitions:

State. Accurate and informative text embeddings are crucial for both memory storage and efficient retrieval in ALEC. We employ a pretrained encoder $\mathcal{E}$ from the SentenceTransformers library (Reimers, 2019) to generate these embeddings. At decoding step t , the agent’s state is $s^t = \mathcal{E}([\mathbf{x}, y^t])$.

Action. From state s^t , the LLM generate K actions and selects one as follows:

$$a^t = y^{t+1:t+M} \sim \pi_{\mathcal{L}}(\cdot \mid [\mathbf{x}, y^t]),$$

where each action is a continuation corresponding to up to M consecutive tokens. Consequently, each episode unfolds over $D = \left\lceil \frac{T}{M} \right\rceil$ steps, where T is the total number of tokens to be decoded. For simplicity, we use the terms *action* and *continuation* interchangeably from this point onward.

Reward. In our framework, reward is zero until a continuation that contains the end-of-sequence (EOS) token is reached. Let $s_1, \dots, s_D$ denote the sequence of states in an episode, with terminal state $s_D = \mathcal{E}([\mathbf{x}, \mathbf{y}])$. Assume $s_d, d \in [1, D]$ to be the state reached after generating action a^t . The reward for state s_d is formally written as:

$$R(s_d) = \begin{cases} 0, & \text{if EOS} \notin a^t, \\ \text{RM}(s_D), & \text{if EOS} \in a^t. \end{cases} \quad (1)$$

Return Derivation. Given the reward in Equation 1, the return for state s_d is defined as:

$$g(s_d) = \gamma^{D-d} \text{RM}(s_D) \quad (2)$$

where γ is the discount factor. These returns are written to memory as initial values g . To avoid confusion, we note that g may be updated later through our Memory Write method, which accounts for the distinct notation between g and v in Figure 1.

Value Function. We define the optimal value function for state $s^t = \mathcal{E}([\mathbf{x}, y^t])$ as:

$$v^*([\mathbf{x}, y^t]) = \mathbb{E}_{z_1, z_2, \dots \sim \pi_{\mathcal{L}}} \sum_{i \geq 0} R([\mathbf{x}, y^t, z^i]) \quad (3)$$

where z^i is the next possible chunk and z_i is the full response derived from policy $\pi_{\mathcal{L}}$. The objective is to estimate v^* to guide the decoding process,

ensuring that the output of $\mathcal{L}$ aligns with the RM. While previous approaches have relied on DQN (Mudgal et al., 2024) or the standard Bradley-Terry model (Khanov et al., 2024) to learn the value function, we propose a training-free approach based on Episodic Control (Blundell et al., 2016) to directly estimate the value. This alleviates the need for extensive training, making it particularly suitable for data-efficient frameworks.

Episodic Control (Blundell et al., 2016) is a special RL method that approximates action values by recording the highest returns observed when taking actions from specific states, using a growing memory table $Q^{EC}(s, a)$ rather than parametric learning. Each entry of the table contains the *highest return ever obtained* by taking action a from state s . At inference time, the policy estimates and selects the action with the highest stored return for the current state. If an exact (s_t, a_t) pair already appears in Q^{EC} , $Q^{EC}(s_t, a_t)$ is taken as the estimated value. Otherwise, it is estimated using k -nearest neighbors approximation:

$$\hat{Q}(s, a) = \frac{1}{k} \sum_{i=1}^k Q^{EC}(s^{(i)}, a) \quad (4)$$

where $s^{(1)}, \dots, s^{(k)}$ are the nearest neighbors of s under a defined distance metric. This enables rapid, gradient-free adaptation to recurring structures in near-deterministic environments by reusing prior successful trajectories. We will adapt the mechanism to the LLM alignment problem as follows.

3.2 Sample-Efficient Alignment Learning With Episodic Control

Memory Structure Our episodic memory $\mathcal{M}$ is structured as a dictionary, where states are stored as keys, and their corresponding value estimations are stored as values. In contrast to traditional model-free episodic memory (Blundell et al., 2016), where a memory element typically contains the state, a discrete action, and a reward, ALEC only stores the state as the key, along with its estimated value. The episodic memory can be represented by the following structure: $\mathcal{M} = \{s_i : v_i\}_{i=1}^L$ where s_i is a state collected from the training data, v_i is the associated estimated value of the state, and L denotes the maximum capacity of our episodic memory. We set memory capacity to $L = 10,000$, which is empirically sufficient—no overflow was observed in our sample-efficient setting.




Figure 1: **Overall framework of ALEC:** The framework consists of two phases: **Memory Construction** and **Inference**. In the memory construction phase, ALEC selectively writes data into Memory and interacts with the LLM to obtain state and value for Nearest-neighbor Memory Write. Note that the initial value g will be updated during Memory Write, and is thus denoted as v in the memory to reflect these updates. In the Inference phase, the memory serves as a non-parametric nearest-neighbor estimator for selecting the next continuation. If $is_certain$ is True, we select the continuation with the highest score estimated by the memory. Otherwise, we base the score on LLM probabilities. In this example, $is_certain$ is True, thus we get the continuation based on memory score.

Memory Operations Our framework consists of two main operations: *Memory Write* and *Memory Read*, corresponding to the memory construction and inference phases. During memory construction, we apply ϵ -greedy action selection at each generation step and introduce a strategy for determining the order in which samples are written. We also propose to update the nearest neighbor in memory for each write operation, refining the value approximation. During inference, the memory serves as a value approximator to select the highest-scoring continuation at each step. Further details are provided in the following sections.

3.2.1 Memory Construction

Memory Write. We first describe how to write into our memory. Given a current state s and a set of K candidate actions $\{a_i\}_{i=1}^K$, we define the next state, resulting from taking action a_i as $s' = s'(a_i | s)$. For a next state s' , we estimate its value $\hat{v}(s')$ as:

$$\hat{v}(s') = \begin{cases} \frac{\sum_{s_j \in \mathcal{N}(s')} cs(s', s_j) \cdot \mathcal{M}[s_j]}{\sum_{s_j \in \mathcal{N}(s')} cs(s', s_j)}, & \text{if } s' \notin \mathcal{M} \\ \mathcal{M}[s'], & \text{otherwise} \end{cases} \quad (5)$$

where $\mathcal{M}[s']$ represents the current memory value for s' , $\mathcal{N}(s')$ denotes the set of nearest neighbors of s' in the memory $\mathcal{M}$, and $cs(s', s_j)$ is the cosine similarity between the two states, defined as: $cs(s', s_j) = \frac{s' \cdot s_j}{\|s'\| \|s_j\|}$. The weighted sum ensures that states with higher semantic similarity contribute more to the final estimation. During mem-

ory construction, we employ an ϵ -greedy policy to select one action from K possible actions, which helps explore diverse and high-quality actions. This is repeated using Equation 5 until completion. To avoid confusion, we note that $\hat{v}(\cdot)$ in Equation 5 is used solely for estimating candidate values during decoding, distinct from the return $g(\cdot)$ in Equation 2, which is used to assign initial discounted returns to states in memory.

Once generation concludes, we compute $g(s)$ and write the pair $(s, g(s))$ to memory. Importantly, this process also updates the values of neighboring states. Since large state spaces with limited samples may contain sparsely distributed neighbors, we propose that only those satisfying $cs(s, s_i) \geq cs_{\text{threshold}}$ are considered for updates. This ensures only sufficiently similar states contribute to the memory estimation. To update neighboring states, we follow a weighted rule inspired by Le et al. (2021), applying an update rate α :

$$\mathcal{M}[s_i] \leftarrow \mathcal{M}[s_i] + \alpha (g(s) - \mathcal{M}[s_i]) \frac{cs(s, s_i)}{\sum_j cs(s, s_j)} \quad (6)$$

Here, j indexes the neighbor set $\mathcal{N}(s_i)$. This guarantees that neighbors with greater similarity receive larger updates in proportion.

Order Of Construction. A key feature of our approach is the incremental construction of memory during the construction phase. Recall that we propose the Memory Write operation to store new state-value pairs in memory while also updating

the values of existing states, rather than simply appending new pairs. As data is progressively written, the partially constructed memory is used to approximate values. Consequently, the order in which samples are introduced into memory becomes critical. Without careful control, early memory writes can disproportionately shape the value estimates of candidate continuations.

Sample Selection For Memory Construction. To address the aforementioned issue, we introduce a novel sample selection strategy for the Memory Write operation. Similar to the Upper Confidence Bound approach (Sutton, 2018), in our framework, a sample is preferred for storage in memory if its estimated value is high and substantially different from those already stored. Specifically, during memory construction, each sample $\mathbf{x}_i$ with its corresponding state s_i is ranked using the following insert score:

$$\text{insert_score}(\mathbf{x}_i) = \hat{v}(s_i) + \beta\sigma_i, \quad (7)$$

where σ_i is the variance of the values in $\mathcal{N}(s_i)$ used to estimate $\hat{v}(s_i)$ using Equation 5, β is the parameter controlling the influence of the uncertainty of the memory when estimating an insert score for state s_i . We iteratively select the sample with the highest insert score for Memory Write to enhance memory adaptability and improve value estimation. A full description of the Memory Write procedure is provided in Algorithm 1, Appendix A.5.

3.2.2 Inference With Memory Read

In this section, we describe the Memory Read process, which plays a crucial role during decoding. The goal at each inference step is to estimate the values of K candidate actions and select the one with the highest estimated value, provided that the memory is sufficiently certain.

In practice, the model $\mathcal{L}$ frequently encounters novel states—those not previously stored in memory. To handle this, we treat the memory as a *nearest-neighbor value estimator*. Given a current state s , we estimate the value of taking an action a by evaluating the resulting next state $s' = s'(a|s)$ for K candidate actions. The estimated value $\hat{v}(s')$ is then calculated using Equation 5.

When executing the reading process, we also evaluate how certain the memory is in approximating the scores for the candidate actions. If the score gap between the top two actions exceeds a threshold, the memory prediction is considered certain,

and the action is accepted. Mathematically, we define a boolean variable, *is_certain*, as follows:

$$\begin{aligned} \text{is_certain} = & \left| \max_{a_i} \hat{v}(s'(a_i|s)) - \max_{a_j \neq a_i} \hat{v}(s'(a_j|s)) \right| \\ & > \zeta \cdot \text{std}(\hat{v}) \end{aligned} \quad (8)$$

where $\text{std}(\hat{v})$ is the standard deviation of the values given by the memory for all possible continuations, and ζ denotes the weight that controls the threshold for acceptance. When *is_certain* = False, the decision is made based on the original probability provided by the base LM $\pi_{\mathcal{L}}$.

Based on the aforementioned procedure, during testing, the optimal action is selected using a greedy strategy as follows:

$$a^* = \begin{cases} \text{argmax}_{a_i} \pi_{\mathcal{L}}(a_i|s), & \text{if } \text{is_certain} = \text{False} \\ \text{argmax}_{a_i} \hat{v}(s'(a_i|s)), & \text{otherwise} \end{cases} \quad (9)$$

where $\pi_{\mathcal{L}}(a_i|s)$ is the probability the LLM assigns to continuation a_i given state s . This encourages each step to favor tokens with high certainty and estimated value, guiding the LLM toward more optimal outputs.

4 Experiments

We aim to demonstrate how ALEC can guide LLMs to decode in alignment with human preferences while minimizing calls to the expert reward model. The experiments are conducted using 3 base open-source LMs: **Llama-2-7b-chat-hf**; **Vicuna-7b-v1.5**; **Mistral-7B-Instruct-v0.2** (Touvron et al., 2023; Zheng et al., 2023b; Jiang et al., 2023).

4.1 Datasets

Anthropic HH. (Bai et al., 2022) We use the standard benchmark for alignment learning problems. This dataset contains conversations between a human and an agent, where the goal is to complete the next turn in the conversation. To represent distinct alignment goals, we create 2 different tasks using 2 subsets: *helpful-base* and *harmless-base*, respectively. Each subset focuses on aligning the agent’s responses to either be maximally helpful or harmless, ensuring targeted performance evaluation for both objectives. **TL;DR.** (Stiennon et al., 2022) We also evaluate on the TL;DR summarization task, which includes Reddit posts paired with two summaries and human preference annotations.

	Llama-2-7b-chat-hf				Vicuna-7b-v1.5				Mistral-7B-Instruct-v0.2			
	harm.	help.	summ.	aver.	harm.	help.	summ.	aver.	harm.	help.	summ.	aver.
DPO $_{\beta=0.01}$	51.21	49.28	49.95	50.15	49.31	51.88	50.05	50.41	47.27	53.63	50.77	50.56
DPO $_{\beta=100}$	50.61	49.32	49.74	49.89	50.52	<u>53.58</u>	48.50	50.87	46.02	<u>52.09</u>	52.12	50.07
M=256												
random	45.54	46.22	51.08	45.71	49.70	48.51	49.43	49.21	49.74	48.15	48.29	48.73
FUDGE	49.31	46.60	49.02	48.31	49.95	50.93	48.40	49.76	49.52	47.49	49.33	48.78
CD	50.69	50.55	50.26	50.50	48.31	41.84	59.85	50.00	48.39	33.00	50.57	43.99
CD + $\pi_{\mathcal{L}}$	50.12	50.29	44.89	48.43	60.75	41.08	43.55	48.46	55.56	33.01	<u>57.17</u>	48.58
ARGS	53.20	49.91	53.35	52.15	51.08	37.60	52.53	47.07	56.48	31.18	49.53	45.73
ARGS + $\pi_{\mathcal{L}}$	53.38	50.80	43.75	49.31	61.70	35.25	44.48	47.14	<u>56.69</u>	31.48	53.56	47.24
ALEC (ours)	<u>60.90</u>	<u>61.64</u>	<u>52.60</u>	58.38	73.44	54.37	51.60	59.80	56.89	51.91	55.80	54.86
M=40												
random	44.00	42.31	44.17	43.49	50.69	45.24	46.44	47.46	49.22	48.35	48.19	48.59
FUDGE	46.54	44.70	44.99	45.41	50.00	45.79	46.13	47.31	49.52	46.60	50.26	48.79
CD	50.77	45.60	47.68	48.02	51.51	41.80	50.77	48.03	54.58	33.43	51.08	46.36
CD + $\pi_{\mathcal{L}}$	51.21	50.08	50.46	50.58	56.92	43.46	53.46	51.28	54.88	33.26	54.28	47.47
ARGS	48.05	47.15	44.69	46.63	49.87	39.59	50.36	46.61	55.32	34.79	48.19	46.10
ARGS + $\pi_{\mathcal{L}}$	48.87	49.65	47.88	48.80	53.67	37.68	50.98	47.44	54.28	35.47	53.97	47.91
ALEC (ours)	60.94	61.81	51.08	<u>57.94</u>	<u>72.88</u>	50.51	<u>53.77</u>	<u>59.05</u>	53.70	52.03	58.70	<u>54.81</u>

Table 1: Win rates over base model $\mathcal{L}$ using three backbones across three benchmarks (harmless, helpful, summarize) and their average. Results use $E = 1000$ calls to reward model $\mathcal{R}$, with inference-time methods evaluated at chunk lengths $M = 256$ and $M = 40$, and $K = 15$ continuations per step. **Bold** and underline indicate highest and second-highest scores, respectively. Our method is shown in gray.




Figure 2: Average rewards across multiple datasets for ALEC and other methods. For all inference-time baselines, we use $M = 256$ and $K = 15$. The dashed lines represent the average reward of the base model generations.

4.2 Reward Models (RM) and Evaluation Metrics

RM. We utilize pre-trained RM for the 3 tasks above. For *harmless-base* and *helpful-base*, we use GPT-2 large models fine-tuned on their respective datasets. These models achieve 73.7% and 72.6% accuracy on the corresponding test sets. For *summarize*, we use a fine-tuned DeBERTa-large model, which achieves 72.23% accuracy on its test set. Further details can be seen in Appendix A.13

Win-rate against the base policy. Following previous work, we evaluate the effectiveness of ALEC in different datasets by measuring the win rate against the base LM $\mathcal{L}$. We define π_1 as winning against π_2 on prompt $\mathbf{x}$ if $r(\mathbf{x}, \mathbf{y}_1) > r(\mathbf{x}, \mathbf{y}_2)$, where $\mathbf{y}_1 \sim \pi_1$

and $\mathbf{y}_2 \sim \pi_2$.

Average reward. This metric measures the average reward across all benchmark samples.

4.3 Baselines

We compare our work with state-of-the-art training-based and decoding-based methods. For a fair comparison, we only use $E = 1000$ training samples provided with true rewards for optimization across all baselines. Following prior work (Mudgal et al., 2024), we evaluate the decoding-based methods with different chunk lengths M .

Random serves as a dummy baseline without optimization. Given K candidate continuations, one is randomly selected. **FUDGE** (Yang and

ALEC vs.	FUDGE			$CD + \pi_{\mathcal{L}}$			ARGS		
	W	D	L	W	D	L	W	D	L
<i>harm.</i>	12.0	79.0	9.0	20.0	62.0	18.0	18.0	70.0	12.0
<i>help.</i>	53.0	36.0	11.0	32.5	53.0	14.5	28.5	51.5	20.0
<i>summ.</i>	37.5	51.0	11.5	44.0	48.5	7.5	23.5	66.0	10.5

Table 2: Head-to-head of Win-Draw-Lose (W-D-L) rates (%) between ALEC and other methods across different datasets, using 200 samples per dataset and *Llama-2-7b-chat-hf* ($M = 40$ and $K = 15$). Evaluation was performed using the Command-R+ model.

Klein, 2021) trains the scorer using explicit rewards from the decoding path. **Controlled Decoding** (Mudgal et al., 2024) proposes a DQN-based prefix scorer v^{CD} to guide generation by estimating the value of future continuations. We implement two variants: $CD + \pi_{\mathcal{L}}$ and CD , which differ in their decoding strategy. The CD variant uses the trained scorer to greedily select the next token, while $CD + \pi_{\mathcal{L}}$ follows the original formulation, sampling from a modified distribution π_{λ}^* , defined as: $\pi_{\lambda}^*(a|[x, y^t]) \propto \pi_{\mathcal{L}}(a|[x, y^t]) \cdot \exp(\lambda v^{CD}([x, y^t, a]))$, which reweights the base LM distribution with parameter λ . **ARGS** (Khanov et al., 2024) similarly employs a learned scorer v^{ARGS} , trained using the Bradley-Terry model (Bradley and Terry, 1952) to capture pairwise preferences. It modifies the decoding distribution to π_w^* , defined as: $\pi_w^*(a|[x, y^t]) \propto \pi_{\mathcal{L}}(a|[x, y^t]) + w v^{ARGS}([x, y^t, a])$, which linearly combines base LM probabilities with the scorer output via parameter w . We refer to the full ARGS method as $ARGS + \pi_{\mathcal{L}}$, and the variant using only the scorer as $ARGS$. **DPO** (Rafailov et al., 2024) improves the generator by further training the base model on a preference dataset. With DPO, for a fair comparison, for each training sample we roll out 2 responses from the base LM, label them using the RM, and train based on those data.

4.4 Benchmarking Using Expert RM

ALEC consistently achieves the highest win rates within 1000 calls to the RM. The main results are shown in Table 1. In most settings, ALEC outperforms previous baselines by a large margin across all backbones. On the *harmless* dataset, with *Vicuna-7b-v1.5* and $M = 256$, we achieve an 11.74% gain over the second-best method ($ARGS + \pi_{\mathcal{L}}$). ALEC also ranks first for the *Llama-2-7b-chat-hf* and *Mistral-7B-Instruct-v0.2* models, with win rates of 60.90% and 56.89%, respectively. For the *helpful* dataset, ALEC achieves the high-

est performance for both *Llama-2-7b-chat-hf* and *Mistral-7B-Instruct-v0.2*, with win rates of 61.81% ($M = 40$) and 54.37% ($M = 256$), respectively. It only trails DPO for *Mistral-7B-Instruct-v0.2*, which is fine-tuned, while all other inference-time baselines struggle to surpass the base model performance. For *summarize*, ALEC remains among the best performers across all models, achieving the highest win rate of 58.70% with $M = 40$ on *Mistral-7B-Instruct-v0.2*. Overall, ALEC delivers the best average performance across all three backbone models, highlighting the effectiveness of our method in limited RM call settings.

ALEC improves the average reward across the benchmarks. As shown in Figure 2, ALEC demonstrates superior performance across the *harmless*, *helpful*, and *summarize* datasets. On the *harmless* dataset with *Llama-2-7b-chat-hf*, ALEC increases the average reward by 28.57% (from 0.7 to 0.9) and achieves similarly strong results with other models, notably outperforming baselines such as $ARGS + \pi_{\mathcal{L}}$. For *helpful* dataset, ALEC maintains its competitive edge, with *Mistral-7B-Instruct-v0.2* being the only approach to surpass the base model’s mean reward. Finally, for *summarize* dataset, while the improvements are less pronounced than in the previous datasets, ALEC still ranks among the highest-performing baselines, highlighting the robustness and effectiveness of our method across different backbone models.

4.5 LLM Evaluation

To assess nuanced language quality beyond standard metrics, we employ Command-R+, a 104B-parameter LLM from CohereAI (Cohere For AI, 2024), as a proxy for human evaluation. It scores outputs from ALEC and baselines (FUDGE, ARGS, and $CD + \pi_{\mathcal{L}}$) on 200 prompts across three datasets. To reduce positional bias (Zheng et al., 2023a), we randomize response order in the evaluation prompt (see Appendix A.15).

Table 2 presents the results of the Win-Draw-Lose evaluation, which show that ALEC **outperforms all baselines by a significant margin**. This indicates that ALEC not only generates more accurate responses but also better meets the specific requirements of each dataset, whether it is harmless, helpful, or summarization. Notably, ALEC’s performance gap is particularly pronounced in challenging datasets, where its ability to address nuanced task requirements sets it apart. For instance, ALEC achieves a 44.0% win rate with only 7.5%

Embedding type	harm.	help.	summ.
Llama-2-7b-chat-hf	54.28	56.07	50.36
SentenceTransformers	60.90	61.64	52.60

Table 3: Ablation on generation LLM embeddings: Win rates of ALEC using *Llama-2-7b-chat-hf* with $M = 256$ and $K = 15$ across different embedding types. The best win rates are highlighted in **bold**.

Method	256	40
Insert-only	53.45	54.43
ϵ -greedy	43.58	42.95
Confidence-based (Ours)	60.90	60.94

Table 4: Ablation on different Memory Write methods: Win rates of ALEC for the *harmless* dataset under various memory write schedules with $K \in [40, 256]$ using *Llama-2-7b-chat-hf*. Best win rates are **bolded**.

losses against $CD + \pi_{\mathcal{L}}$ on *summarize*, and 53.0% wins vs. 11.0% losses against FUDGE on *harmless*. These results align with those in Table 1, further confirming the robustness of ALEC across all tasks.

5 Model Analysis and Ablation Study

Why Not Using Generation LLM Embeddings?

We include results from ALEC using generation LLM embeddings to validate our choice of sentence embeddings. The model we use is *Llama-2-7b-chat-hf*. For embedding type *Llama-2-7b-chat-hf*, we take the mean pooling of embeddings across all layers across all non-padding token of the last hidden state. As seen in Table 3, the results using SentenceTransformers are better across all datasets. We hypothesize that, since our method relies on semantic retrieval in the Memory Read process, embeddings that can better distinguish sentences based on their semantic meaning are more suitable.

Effect Of Memory Write Method. We analyze the impact of our writing method by comparing it with two standard scheduling approaches: **insert-only** and **ϵ -greedy**. The insert-only method adds new state-value pairs without updates, while ϵ -greedy selects new examples based solely on their estimated value $\hat{v}_i$, neglecting uncertainty in updating neighbors. Table 4 shows that the insert-only method becomes inefficient when reducing chunk length from 256 $\rightarrow$ 40. Our writing schedule shows improved performance, whereas the insert-only and ϵ -greedy methods do not show the same trend, with ϵ -greedy results significantly lower than




Figure 3: Ablation on memory size: Win rates (bars) and inference time (lines) of ALEC versus number of examples on three datasets using *Llama-2-7b-chat-hf*.

expected. This may stem from over-smoothing of memory values, where bias from the initial sample set results in inaccurate subsequent value estimations, negatively impacting performance.

Analysis On Memory Size We evaluate scalability by testing memory sizes of 100, 1000, and 10,000 samples on *Llama-2-7b-chat-hf* across all three datasets. As shown in Figure 3, adding more examples generally improves final performance. For *harmless* dataset, using 10,000 samples increases performance by approximately 7% compared to using 100 samples. For *helpful* dataset, the performance hits the highest at 1000 samples. Notably, increasing the number of elements in memory has only a marginal impact on generation time, highlighting the time efficiency of our method.

Other Ablation Studies. Additional ablations—covering alignment trade-offs, scaling to larger LMs, embedding visualization, different number of K , farthest neighbors, memory usage, different number of neighbors—are provided in Figures 4, 5, 6, 7 and Tables 5, 9, 10 (see Appendix A.1–A.11). These results confirm that ALEC significantly improves alignment performance while remaining effective in sample-efficient scenarios.

6 Conclusion

We introduce Alignment Learning with Episodic Control (ALEC), a non-parametric framework that aligns LLM outputs during inference without fine-tuning. ALEC stores key-value associations and uses a confidence-based writing scheme for data efficiency. Evaluations on harmless, helpful, and summarization tasks show its effectiveness, with detailed model analyses highlighting its performance advantages with minimal RM interactions.

Limitations

Our work introduces a data-efficient approach to aligning LLMs through Episodic Control, which we term ALEC. While our initial experiments demonstrate encouraging results—showing that ALEC can improve alignment with relatively few examples—these evaluations are currently limited to medium-scale datasets. As such, further empirical evaluations are required to assess the scalability of ALEC in more complex, large-scale datasets.

Acknowledgments

This work was supported by compute credits from a Cohere Labs Research Grant, these grants are designed to support academic partners conducting research with the goal of releasing scientific artifacts and data for good projects.

References

- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, and 1 others. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Charles Blundell, Benigno Uria, Alexander Pritzel, Yazhe Li, Avraham Ruderman, Joel Z Leibo, Jack Rae, Daan Wierstra, and Demis Hassabis. 2016. Model-free episodic control. *arXiv preprint arXiv:1606.04460*.
- Ralph Allan Bradley and Milton E. Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, and 1 others. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *Transactions on Machine Learning Research*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and 1 others. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 30.
- Cohere For AI. 2024. [c4ai-command-r-plus-08-2024](#).
- Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. Toxicity in chatgpt: Analyzing persona-assigned language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1236–1270.
- Van Dai Do, Quan Hung Tran, Svetha Venkatesh, and Hung Le. 2025. [Dynamic steering with episodic memory for large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 13731–13749, Vienna, Austria. Association for Computational Linguistics.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. 2020. Realltoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369.
- James Y. Huang, Sailik Sengupta, Daniele Bonadiman, Yi-An Lai, Arshit Gupta, Nikolaos Pappas, Saab Mansour, Katrin Kirchhoff, and Dan Roth. 2025. [DeAL: Decoding-time alignment for large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26280–26300, Vienna, Austria. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Maxim Khanov, Jirayu Burapachee, and Yixuan Li. 2024. Args: Alignment as reward-guided search. In *The Twelfth International Conference on Learning Representations*.
- Hung Le, Thommen Karimpanal George, Majid Abolshah, Truyen Tran, and Svetha Venkatesh. 2021. Model-based episodic memory induces dynamic hybrid controls. *Advances in Neural Information Processing Systems*, 34:30313–30325.

- Máté Lengyel and Peter Dayan. 2007. Hippocampal contributions to control: the third way. *Advances in neural information processing systems*, 20.
- Sheng Liu, Haotian Ye, Lei Xing, and James Zou. 2024. [In-context vectors: Making in context learning more effective and controllable through latent space steering](#). *Preprint*, arXiv:2311.06668.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, and 1 others. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, and 1 others. 2024. Controlled decoding from language models. In *Forty-first International Conference on Machine Learning*.
- Long Ouyang, Jeff 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 Christiano, Jan Leike, and Ryan Lowe. 2022a. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022b. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- N Reimers. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2022. [Learning to summarize from human feedback](#). *Preprint*, arXiv:2009.01325.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Richard S Sutton. 2018. Reinforcement learning: An introduction. *A Bradford Book*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui Wang. 2025. Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early decoding. *arXiv preprint arXiv:2503.01422*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, and 1 others. 2022. Emergent abilities of large language models. *Transactions on Machine Learning Research*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Hugging-face’s transformers: State-of-the-art natural language processing](#). *Preprint*, arXiv:1910.03771.
- Kevin Yang and Dan Klein. 2021. [FUDGE: Controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3511–3535, Online. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023a. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023b. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the*

A Appendix

A.1 Alignment Trade-off Results




Figure 4: Ablation on performance versus KL divergence: *harmless-base* win rate versus KL divergence across various baselines using *Llama-2-7b-chat-hf*. ALEC outperforms both training-based and inference-based techniques, showcasing effective guided optimization with a favorable trade-off curve, while other baselines fail to surpass the base model’s performance.

We analyze the alignment trade-off by presenting the trade-off curve between the aligned model and the base model using *harmless* dataset. A good answer should achieve a high reward with a low KL divergence compared to the base model, $KL(\pi||\pi_{ref})$, where π_{ref} is the aligned policy used for sampling answers. For inference-time based methods, with L_i representing the number of tokens in the response of $\mathbf{x}_i$, we use the bound of KL divergence $KL(\pi||\pi_{ref}) \leq \mathbb{E}_{\mathbf{x} \sim \mu} (\log(K) - \frac{K-1}{K}) \lceil \frac{L_x}{M} \rceil$ (Stiennon et al., 2022; Mudgal et al., 2024). For DPO, we adjust the β parameter to explore various KL divergence values.

The win-rate and KL divergence trade-off are illustrated in Figure 4. As the outputs generated by ALEC deviate from the base model, they consistently achieve higher rewards, indicating effective optimization. In contrast, the other baselines exhibit less coherent learning patterns, underscoring the challenges these methods face with limited samples. This observation further emphasizes the efficacy of ALEC as a non-parametric approach, demonstrating its superior ability to leverage available data for enhanced performance. We note that the KL value is calculated from the responses over the dataset, making it difficult to obtain an exact value. This explains the longer lines in our figure.

Algorithm 1 Memory Write

Require: Memory $\mathcal{M}$, Reward Model RM , Number of generation candidates K , Train dataset $\mathcal{D}_{train}$, Encoder $\mathcal{E}$, LM policy $\pi_{\mathcal{L}}$

- 1: Select a sample $\mathbf{x} \in \mathcal{D}_{train}$ using Equation 7
- 2: Get initial context $c = [\mathbf{x}]$, set $y^{L_1} \leftarrow \emptyset$
- 3: Get initial state $s^1 \leftarrow \mathcal{E}(c)$
- 4: **for** $t = 1, \dots, T$ **do**
- 5: Get $a_i = \pi_{\mathcal{L}}(\cdot|c)$
- 6: Estimate $\hat{v}(s^{t+1}(a_i|s^t)), i = [1, \dots, K]$ using Equation 5
- 7: Select action $a = y^{L_t:L_t+M}$ with ϵ -greedy policy
- 8: $y^{L_{t+1}} \leftarrow [y^{L_t}, y^{L_t:L_t+M}]$
- 9: Update context $c \leftarrow [\mathbf{x}, y^{L_{t+1}}]$
- 10: Move to next state $s^{t+1} \leftarrow \mathcal{E}(c)$
- 11: **end for**
- 12: Retrieve reward for full response $R_T = RM(s_D) = RM(\mathbf{x}, \mathbf{y})$
- 13: **for** $d = D, D-1, \dots, 1$ **do**
- 14: $g(s_d) \leftarrow \gamma^{D-d} R_T$ (Equation 2)
- 15: Insert $(s_d, g(s_d))$ into memory
- 16: **for each** $s_i \in \mathcal{N}(s_d)$ **do**
- 17: **if** $cs(s_d, s_i) \geq cs_{threshold}$ **then**
- 18: Update $\mathcal{M}[s_i]$ using Equation 6
- 19: **end if**
- 20: **end for**
- 21: **end for**

A.2 Additional Related Work

Our work operates strictly under a budget of 1000 RM queries. While Best-of-K is a strong and widely adopted baseline, it requires K RM evaluations per sample, resulting in a total of $K \times N$ RM calls for N test examples. This makes it impractical under our budget constraints.

For instance, even the smallest dataset (TL;DR) contains approximately 1,000 examples. Running Best-of-2 would already require at least 2,000 RM queries, and larger K values or datasets would only increase this cost further. This violates our sample-efficient evaluation setting, which is why Best-of-K is not included in our main experimental comparison.

Nonetheless, for reference, we report Best-of-K results using Vicuna-7B v1.5 in Table 6. These results are intended purely for context and lie outside the efficiency regime that our method targets. Notably, our win rates on the Harmless (73.44%) and Helpful (54.37%) datasets achieve performance

Mechanism	M	harmless	helpful	summarize
Farthest	40	56.83	56.79	47.47
Nearest (ours)	40	60.94	61.81	51.08
Farthest	256	52.25	56.03	49.22
Nearest (ours)	256	60.90	61.64	52.60

Table 5: Ablation Study on Using farthest neighbors: Win rates of ALEC using *Llama-2-7b-chat-hf* with $K = 15$, $M = 256$ and $M = 40$ across different neighbor mechanisms. The best win rates are highlighted in **bold**.

closest to this baseline, demonstrating that our method is highly effective while remaining significantly more sample-efficient.

While our work focuses on reward-model-based alignment during inference via ALEC, other methods such as DeAL (Huang et al., 2025) take a fundamentally different approach. Instead of relying on a trained reward model, DeAL enforces alignment directly at decoding time using user-specified reward functions. Since ALEC is designed to assess and leverage reward-model predictions with episodic memory, a direct empirical comparison with DeAL is not meaningful. Another line of work explores alignment through activation steering, which modifies the latent representations of the language model to guide its behavior (Liu et al., 2024; Do et al., 2025). However, activation steering requires direct intervention in the model’s internal activations, whereas ALEC operates entirely at the inference level without modifying the model parameters or hidden states. As a result, these methods pursue orthogonal objectives, making direct empirical comparison similarly uninformative.

Dataset	Win Rate (%)
Harmless	86.28
Helpful	87.98
Summarize	86.58

Table 6: Win rates of Best-of-K vs. base model.

A.3 F1 and ROC-AUC for Reward Models

We provide additional F1 and ROC-AUC metrics to better contextualize the reported accuracy. For our pre-trained reward model evaluated on the balanced test sets, we report in Table 7.

A.4 Ablation on is_certain

We conduct additional experiments to study the effect of the `is_certain` mechanism. Since

Dataset	Accuracy (%)	F1 Score (%)	ROC-AUC (%)
Harmless	73.71	84.05	69.87
Helpful	72.59	83.12	69.19

Table 7: Performance of the pre-trained reward model on balanced test sets.

`is_certain` is determined by a threshold parameter ζ , we evaluate the impact of varying ζ values. All experiments are conducted using Vicuna-7B-v1.5 with $M = 256$ on the Harmless dataset. We provide the results in Table 8.

ζ	Win Rate vs. Base (%)
0.5	60.70
0.3	60.73
0.1	60.90

Table 8: Ablation on `is_certain` with different ζ thresholds.

A.5 Memory Write Algorithm

We provide the detailed information on the Memory Write algorithm in Algorithm 1.

A.6 Confirming The Effect Of Nearest Neighbors

We include the results of ALEC using *Llama-2-7b-chat-hf* across all datasets when retrieving farthest neighbors instead of nearest neighbors. This ablation aims to verify the importance of retrieving semantically similar examples in our setting. Intuitively, using farthest neighbors is expected to degrade performance, as these examples are less relevant and likely to yield inaccurate value estimates. As shown in Table 5, substituting nearest neighbors with farthest neighbors leads to a substantial drop in performance across datasets. This outcome is consistent with our expectations, as the retrieved examples no longer reflect the characteristics of the test query, thereby providing little to no useful information for estimating its value and ultimately leading to reduced effectiveness.

A.7 Memory Usage

We report the memory usage of our method in megabytes (MB) to assess its efficiency in terms of storage overhead. As expected, ALEC exhibits minimal memory consumption, as it maintains only the essential components needed for value estimation—specifically, the embeddings of each state

# samples	100	1000	10000
harmless	0.75	8.52	85.31
helpful	1.01	9.43	85.36
summarize	0.95	9.02	86.03

Table 9: Memory Usage: Memory used (in MB) with different settings of ALEC. ALEC stores only state embeddings and scalar values, resulting in minimal memory usage.

and a single scalar value representing the estimated value of that state. This lightweight design avoids the need for storing large auxiliary structures or complex model parameters during inference. The results presented in Table 9 empirically validate the memory efficiency of ALEC, highlighting its scalability and suitability for deployment in environments where memory is a limiting factor. Compared to methods with more elaborate memory footprints, ALEC offers a clear advantage in terms of compactness and efficiency.

A.8 Scaling To Larger Language Models.

To further investigate the effectiveness and scalability of ALEC on larger LMs, we conduct additional evaluations using two 13-billion parameter models: *Llama-2-13b-chat-hf* and *Vicuna-13b-v1.5*. Specifically, we compare the win rates of ALEC and the $CD + \pi_{\mathcal{L}}$ baseline against the respective base policies for each model. The results, summarized in Figure 5, reveal that while the $CD + \pi_{\mathcal{L}}$ baseline continues to suffer from the inefficiency of its reward model and limited generalization, ALEC consistently demonstrates strong performance improvements across all datasets and model configurations.

For instance, on the *harmless* dataset with *Vicuna-13b-v1.5*, ALEC achieves a dominant lead with a 17.65% higher win rate compared to the *DPO* baseline. Similarly, on the *helpful* dataset using *Llama-2-13b-chat-hf*, ALEC surpasses $CD + \pi_{\mathcal{L}}$ by a margin of 9.59%. These results not only confirm the robustness of ALEC in larger model settings but also highlight its ability to deliver high-




Figure 5: Ablation on larger models: Paired heatmaps showing win rates (in %) of three decoding-time methods of ALEC, Controlled Decoding with π_L ($CD + \pi_L$) and Direct Preference Optimization (DPO, $\beta = 0.1$) against the base policy on two 13-Billion backbones (Llama-2-13b-chat-hf and Vicuna-13b-v1.5) with $M = 256$ and $K = 15$. Cell color intensity encodes the win rate, and each cell is annotated with its exact value.

$\mathcal{K}_{read}$	1	3	5	10	15
harmless	58.91	61.76	73.44	65.61	65.10
helpful	43.37	44.41	54.37	53.86	53.95
summarize	53.25	51.59	51.60	51.59	50.12

Table 10: Ablation on different numbers of neighbors: Performance of ALEC using *Vicuna-7b-v1.5* with different numbers of neighbors ($\mathcal{K}_{read}$). Best win rates are **bolded**.

quality alignment while requiring only a limited number of RM queries, making it a practical and efficient solution for large-scale policy improvement.

A.9 Ablation On Number Of Neighbors.

We vary the number of neighbors $\mathcal{K}_{read}$ used to estimate the values to observe the performance of ALEC on different datasets using *Vicuna-7b-v1.5* with $M = 256$, and report the results in Table 10. For the *helpful* and *summarize* datasets, the performance remains stable regardless of the number of neighbors. For *harmless*, increasing the number of neighbors shows a nuanced effect: while a larger neighborhood provides more context for value estimation, it may also introduce additional noise, which can impact performance. This highlights the importance of balancing context richness and noise sensitivity when choosing $\mathcal{K}_{read}$.

A.10 Visualization Of Memory Embeddings.

We visualize 50 randomly selected test embeddings (orange) alongside memory embeddings (blue). As




Figure 6: t-SNE visualization of 50 random test samples (orange) with the memory embeddings (blue) with samples from *harmless* dataset

shown in Figure 6, the t-SNE plot demonstrates that memory embeddings sufficiently cover the 2D space for the test embedding set, supporting the effectiveness of a nearest neighbor retrieval mechanism. The distribution indicates that test queries tend to fall within or near dense regions of the memory embedding space, suggesting that relevant memory entries are likely to be found through local similarity search. This overlap implies that the embedding model successfully maps semantically similar inputs close together, enabling reliable reuse of prior experience via nearest neighbor lookups. This strongly supports the success in decoding of our method.

Parameter	harmless	helpful	summarize
E		1000	
K		15	
γ		0.9	
ϵ		0.3	
$cS_{threshold}$		0.7	
α		0.5	
β		0.1	
ζ		0.1	
temperature		0.7	
top_p		0.9	
$\mathcal{K}_{read}$		5	

Table 11: Hyperparameters used in ALEC




Figure 7: Results of ALEC using different number of K on *harmless* dataset using *Vicuna-7b-v1.5* and $M = 256$.

A.11 Ablation on the number of generations per step K .

We further investigate the effect of the number of generations per step, denoted by K , on alignment performance. Specifically, we vary K for the *harmless* dataset using the *Vicuna-7b-v1.5* model and present the results in Figure 7. When $K = 1$, the win rate is 50%, corresponding to the performance of the base model without any additional candidates—no choice is provided, so alignment cannot occur through selection. As K increases, the model is allowed to generate and evaluate a larger set of candidate completions at each step. Consequently, the win rate improves steadily with higher values of K , reaching 73.44% at $K = 15$. This trend is consistent with the intuition that a larger candidate pool increases the likelihood of including a more aligned or preferable response, effectively expanding the search space over which the

Method	Parameters	Value
DPO	Number of epochs	1
	Learning rate	$5e^{-6}$
	Learning rate scheduler type	cosine
	Gradient accumulation step	8
	Warm up ratio	0.1
	Precision	bf16
ARGS	Number of epochs	1
	Cutoff length	1024
	Learning rate	$1e^{-4}$
	Learning rate scheduler type	cosine
	Gradient accumulation step	8
	Warm up ratio	0.1
FUDGE	Number of epochs (max)	10000
	Input size	1024
	Hidden size	128
	Dropout probability	0.5
	Learning rate	$1e^{-6}$
	Batch size	64
Controlled Decoding	Precision	bf16
	Number of epochs	1
	LoRa rank	64
	LoRa alpha	16
	LoRa dropout	0.1
	Learning rate	$1e^{-6}$
	Batch size	1
	Precision	bf16

Table 12: Baselines training details

selection policy can operate. Notably, the gain is non-linear, suggesting diminishing returns at higher values of K , which may reflect a saturation point in diversity or utility among sampled candidates.

A.12 ALEC implementation details

We implement ALEC with all base models on Huggingface Library (Wolf et al., 2020). We also provide the hyper-parameters used for ALEC across 3 datasets and 3 models in Table 11.

In this section, we also note the reason for using ϵ -greedy strategy during decoding in memory construction phase. The reason is that we want to avoid local optima and balance exploration and exploitation, following common RL practice. By default, we choose $\epsilon = 0.3$, which helps occasionally select less likely continuations from top- K candidates, enabling exploration within high-quality continuations. Due to resource constraints, we did not tune this parameter. The performance, however, can still be improved with excessive tune of ϵ .

A.13 Reward Model details

We provide information of the reward models used. All of the reward models below are public available.

For harmless-base, we use https://huggingface.co/Ray2333/gpt2-large-harmless-reward_model

For helpful-base, we use https://huggingface.co/Ray2333/gpt2-large-helpful-reward_model

For summarize, we use <https://huggingface.co/OpenAssistant/reward-model-deberta-v3-large>

A.14 Baseline details

We provide details on how the baselines are implemented in table 12. We note that for DPO (Rafailov et al., 2024), the parameters shown are for training the generator LM, while for the other methods, the parameters are for the scorer of the candidates for continuation. We utilize LlamaFactory framework (Zheng et al., 2024) to train DPO baseline, the prefix scorers of Controlled Decoding paper (Mudgal et al., 2024) and ARGS paper (Khanov et al., 2024). As mentioned, the dataset used for training is the same with our method, ensuring a fair comparison.

A.15 Evaluation prompts

We provide detailed evaluation prompts used for *harmless-base*, *helpful-base* and *summarize* in Fig. 9, Fig. 10 and Fig. 11. We note that the order of responses has been shuffled to avoid any possible order bias from LLM.




Figure 8: Comparative example between ALEC, base generation, ARGS and CD on *helpful-base*




Figure 9: Evaluation prompt for *harmless-base*

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 four AI assistants in response to the user question displayed above.
 Please rate the harmfulness and level of detail of their responses.
 Your evaluation should consider the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response.
 Each assistant receives an overall score on a scale of 1 to 10, where a higher score indicates better overall performance.
 Please first output a single line containing four values indicating the scores for Assistant 1, 2, 3 and 4, respectively. The four 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.

[Question] {dialogue}

[The Start of Assistant 1's Answer]
 {output_1}
 [The End of Assistant 1's Answer]

[The Start of Assistant 2's Answer]
 {output_2}
 [The End of Assistant 2's Answer]

[The Start of Assistant 3's Answer]
 {output_3}
 [The End of Assistant 3's Answer]

[The Start of Assistant 4's Answer]
 {output_4}
 [The End of Assistant 4's Answer]

Figure 10: Evaluation prompt for *helpful-base*

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 four AI assistants in response to the user question displayed above.
 Please rate the summarize ability of the answers, given the original text.
 Your evaluation should consider the relevance, conciseness, accuracy, depth, creativity, and level of detail of the response.
 Each assistant receives an overall score on a scale of 1 to 10, where a higher score indicates better overall performance.
 Please first output a single line containing four values indicating the scores for Assistant 1, 2, 3 and 4, respectively. The four 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.

[Question] {dialogue}

[The Start of Assistant 1's Answer]
 {output_1}
 [The End of Assistant 1's Answer]

[The Start of Assistant 2's Answer]
 {output_2}
 [The End of Assistant 2's Answer]

[The Start of Assistant 3's Answer]
 {output_3}
 [The End of Assistant 3's Answer]

[The Start of Assistant 4's Answer]
 {output_4}
 [The End of Assistant 4's Answer]

Figure 11: Evaluation prompt for *summarize*