

RADAR: Accelerate Large Language Model Inference with RL-Based Dynamic Draft Trees

Junjie Ma¹ Jinlong Li^{1*} Jiajun Luo²

¹University of Science and Technology of China

²Tsinghua University

mahiru@mail.ustc.edu.cn jlli@ustc.edu.cn luo-jj24@mails.tsinghua.edu.cn

Abstract

Inference with modern Large Language Models (LLMs) is expensive and slow, and speculative sampling has emerged as an effective solution to this problem. However, the number of calls to the draft model for generating candidate tokens in speculative sampling is a preset hyperparameter, lacking flexibility. To generate and utilize the candidate tokens more effectively, we propose RADAR, a novel speculative sampling method with RL-based dynamic draft trees. RADAR formulates the draft tree generation process as a Markov Decision Process (MDP) and employs offline reinforcement learning to train a prediction model, which enables real-time decisions on calls to the draft model, reducing redundant computations and further accelerating inference. Evaluations across three LLMs and four tasks show that RADAR achieves a speedup of 3.17x–4.82x over the auto-regressive decoding baseline. The code is available at <https://github.com/minaduki-sora/RADAR>.

1 Introduction

Modern Large Language Models (LLMs) exhibit remarkable capabilities and are widely applied across various tasks (Achiam et al., 2023; Guo et al., 2025). The parameter sizes of LLMs range from tens of billions to even hundreds of billions, and during auto-regressive generation, each generated token requires accessing entire model parameters, making LLM inference slow and costly. Here we focus on accelerating LLM inference.

To accelerate LLM inference, some methods have been proposed recently, such as prompt pruning (Zhou et al., 2023), quantization (Frantar et al., 2022), knowledge distillation (Agarwal et al., 2024), speculative sampling (Leviathan et al., 2023; Chen et al., 2023) and so on. Therein, as one of the most popular methods, speculative sampling

uses a small draft model to quickly generate candidate tokens, which are then verified by target LLM in parallel. Speculative sampling allows multiple tokens to be generated in a single LLM pass, significantly reducing inference latency while preserving the original output distribution of target LLM. For speculative sampling, the recent studies (Miao et al., 2024; Li et al., 2024b; Gao et al., 2025; Li et al., 2024a; Zhang et al., 2024) mostly focus on organizing and utilizing the tokens generated by the draft model more effectively.

In speculative sampling, the generated tokens are generally organized in a certain data structure, termed as *draft structure* here. And there are two types of draft structures: chain and tree-based. Speculative sampling (Leviathan et al., 2023; Chen et al., 2023) first employs an auto-regressive method to invoke the draft model and produce a sequence of candidate tokens, which is referred to as a chain-structured draft. To promote parallelism during LLM’s verifications, candidate tokens are organized as a verification tree (Miao et al., 2024), which is then serialized and verified via masked attention in a single LLM pass. However, most draft trees (Miao et al., 2024; Li et al., 2024b; Cai et al., 2024; Ankner et al., 2024) are static during token generation across all contexts or tasks. By reordering and pruning nodes of draft tree, methods like EAGLE-2 (Li et al., 2024a) and EAGLE-3 (Li et al., 2025) attempt to improve the acceptance length. However, the number of calls to the draft model remains a hyperparameter, lacking flexibility. For instance, we run the source code provided by EAGLE-3 on MT-bench (Zheng et al., 2023) and LLaMA-Instruct 3.1 8B (Grattafiori et al., 2024), results show that draft tokens are completely rejected at a frequency of about 31%, yet the draft model is still called 8 times (Li et al., 2025). While methods like SpecDec++ (Huang et al., 2024) and DISCO (Mamou et al., 2024) attempt to dynami-

*Corresponding author.

cally adjust the draft length by predicting the acceptance probability for each node, their methodologies are strictly designed for chain-based drafts and cannot be generalized to recent state-of-the-art tree-based methods such as EAGLE-3. Here, we attempted to dynamically decide the calls to the draft model during the draft stage, thereby generating draft trees with varying depths based on different contexts.

We propose to decide whether to invoke the draft model at each step of draft generation process, rather than predicting the required depth of the draft tree. This is because the depth of the draft tree, equivalent to the acceptance length, is stochastic due to the rejection sampling in the speculative sampling (Leviathan et al., 2023; Chen et al., 2023), making it impossible to obtain labeled data. However, we observed that the draft tree generation process can be modeled as a Markov Decision Process (MDP), which eliminates the need for labeled data by utilizing intrinsic and extrinsic rewards. We propose to calculate these rewards based on acceptance length distributions of draft trees, which are obtained by running speculative sampling algorithm on shareGPT¹ dataset, and then all these distributions are collected to construct a dataset for training the prediction model. During training, the prediction model samples policy trajectories in real time based on the real MDP, which enables offline reinforcement learning without extrapolation error (Fujimoto et al., 2019).

Upon these insights, we propose a Reinforcement learning Adjusted Draft-generation Algorithm for speculative sampling (RADAR), leveraging the confidence scores from the draft model to dynamically decide the calls to the draft model. In summary, our key contributions include:

- We introduce a speculative sampling framework with dynamic draft trees, namely RADAR, which employs a lightweight prediction model to adaptively decide the calls to the draft model during the draft stage.
- To handle the lack of labeled data for training the prediction model, we model the draft tree generation process as a Markov Decision Process (MDP), and address it through a newly introduced dataset of acceptance length distribution, from which an acceptance length is

sampled and then used for the calculation of action reward, eliminating the need for expensive real-time interactions with the LLM.

2 Related Work

2.1 Speculative decoding

Since the proposal of speculative decoding, researchers have been improving the algorithm from different perspectives. One line of research focuses on enhancing alignment between draft and target models via distillation (Zhou et al., 2024) or feature-level integration (Li et al., 2024b). Another direction involves optimizing draft structures, shifting from chain-based to tree-based decoding (Miao et al., 2024) and utilizing pruning and reranking (Li et al., 2024a) to increase the acceptance rate. The third targets system parallelism to minimize drafting latency, such as Medusa’s (Cai et al., 2024) parallel decoding heads and PEARL’s (Liu et al., 2025) framework that overlaps the drafting and verification phases asynchronously. Our work advances the latter by introducing a dynamic draft tree that adaptively adjusts the number of calls to the draft model, resulting in minimizing drafting latency. Crucially, our approach is orthogonal to PEARL: while PEARL achieves dynamic candidate length, its primary contribution and motivation lie in maximizing temporal parallelism between draft and verification via asynchronous execution. In contrast, RADAR focuses on the redundant calls to the draft model and make real-time decisions to adjust the draft tree. These two paradigms can be integrated to achieve even better acceleration.

2.2 Dynamic draft structure

A major efficiency bottleneck in conventional speculative decoding methods stems from their inherent dependency on a fixed hyperparameter K , representing the number of calls to the draft model. Vanilla Speculative Sampling (Leviathan et al., 2023) derives an optimal K based on the i.i.d. assumption of token acceptance probabilities. To introduce context-awareness, Kangaroo (Liu et al., 2024) adopts a simple heuristic that terminates speculation if the current draft token’s confidence falls below a threshold. Going a step further, SpecDec++ (Huang et al., 2024) employs a trained prediction head to estimate individual acceptance probabilities and stops speculation based on their calculated cumulative product. In a similar vein, the primary distinction in DISCO (Mamou et al.,

¹https://huggingface.co/datasets/anon8231489123/ShareGPT_Vicuna_unfiltered

2024) is the use of a specialized head to directly predict this cumulative product, bypassing the need for manual multiplication of individual probabilities. Despite these advancements, previous methods (Liu et al., 2024; Huang et al., 2024; Mamou et al., 2024) are inherently designed for chain-based draft, making them difficult to generalize to more complex tree-based draft structures that involve pruning and reranking (Li et al., 2024a). Furthermore, these approaches typically rely on supervised learning to estimate acceptance probabilities and use a predefined threshold to implement termination, which remains a rigid hyperparameter. In contrast, RADAR addresses these limitations by modeling the speedup ratio end-to-end using reinforcement learning. This framework enables context-aware dynamic drafting that is compatible with state-of-the-art tree-based speculative decoding methods, effectively accelerating inference across diverse contexts.

Concurrent with our work, LTD (Zhang et al., 2026) also addresses the efficiency issue caused by redundant draft calls in speculative decoding, and likewise applies reinforcement learning to dynamic draft tree control. Different from RADAR, LTD relies on online reinforcement learning and requires costly interaction with the target LLM during training, whereas RADAR adopts an offline reinforcement learning framework that avoids such online interaction. We include an empirical comparison with LTD in the main experiments, and further discuss differences in Appendix B.1.

2.3 EAGLE Series

Distinct from vanilla speculative decoding that employs a separate small language model for token-level drafting, EAGLE (Li et al., 2024b) introduces auto-regression at the feature level. It predicts the second-to-top layer hidden states of the target model M_p and reuses M_p 's original LM head for sampling, which mitigates the distribution shift between the draft and target models.

The framework has evolved through two major iterations. EAGLE-2 (Li et al., 2024a) identifies that acceptance rates are context-dependent rather than purely position-dependent. It leverages the draft model's confidence scores as a proxy for acceptance probability to construct a context-aware draft tree, reordering and pruning nodes at runtime to optimize the verification efficiency. More recently, EAGLE-3 (Li et al., 2025) addresses the limitations of feature prediction and scaling. By shifting from

pure feature extrapolation to direct token prediction and incorporating *multi-layer feature fusion*, EAGLE-3 significantly enhances drafting accuracy and benefits from large-scale data training.

Despite these advancements, the drafting process in the EAGLE family remains constrained by a fixed number of draft model calls. As observed in our preliminary experiments, this static hyperparameter often leads to redundant drafting efforts, highlighting a lack of flexibility in adapting to real-time verification.

3 Preliminaries

Speculative Sampling (Leviathan et al., 2023; Chen et al., 2023) has emerged as a standard technique to accelerate the inference of LLMs without compromising output quality. The core mechanism involves a collaboration between a computationally efficient draft model M_q and a powerful target model M_p . Given an input prefix $\mathbf{x}$, the process begins with a drafting phase, where M_q auto-regressively predicts t future tokens. Specifically, at each drafting step $i \in \{1, \dots, t\}$, a candidate token $\hat{x}_i$ is sampled from the distribution $q_i(\cdot) = M_q(\cdot \mid \mathbf{x}, \hat{x}_1, \dots, \hat{x}_{i-1})$. This stage results in a draft sequence $\hat{\mathbf{X}} = \{\hat{x}_1, \dots, \hat{x}_t\}$ along with its corresponding draft probabilities.

Following the drafting phase, the target model M_p performs a single parallel forward pass on the concatenated sequence $[\mathbf{x}, \hat{x}_1, \dots, \hat{x}_t]$ to compute the target distributions $p_1, \dots, p_{t+1}$, where $p_i = M_p(\cdot \mid \mathbf{x}, \hat{x}_1, \dots, \hat{x}_{i-1})$. To ensure the output follows the target distribution p exactly, a speculative sampling criterion is applied sequentially. For each candidate $\hat{x}_i$, it is accepted with a probability as Eq. (1):

$$\alpha_i = \min \left(1, \frac{p_i(\hat{x}_i)}{q_i(\hat{x}_i)} \right). \quad (1)$$

If the token is accepted, the verification proceeds to $\hat{x}_{i+1}$; otherwise, the process is truncated at index i , and a new token is resampled from the rectified distribution as Eq. (2):

$$p'_i(x) = \frac{\max(0, p_i(x) - q_i(x))}{\sum_{x'} \max(0, p_i(x') - q_i(x'))}. \quad (2)$$

If all t tokens are accepted, an additional token is sampled from p_{t+1} to maximize efficiency. This ensures that in each iteration, the system generates between 1 and $t + 1$ tokens, while the final

output remains statistically identical to the distribution produced by the target model alone, thereby maintaining the quality of the generation.




Figure 1: Pipeline of RADAR. Figure (a) illustrates the steps of the prediction model controlling the draft model generation, while the Figure (b) demonstrates the corresponding generation results.

4 RADAR

RADAR consists of three components: the target LLM, the draft model, and the prediction model, as illustrated in Figure 1 (a). Following the description of the inference pipeline, we explain the *prediction model*, the core component, in detail.

4.1 Inference Pipeline

Compared with previous popular speculative sampling methods such as EAGLE-3 (Li et al., 2025), we introduce a prediction model that dynamically decides to *continue/stop* the draft model’s autoregressive generation at each generation step. Here, we illustrate the inference pipeline of our approach by example, as shown in Figure 1. A detailed step-by-step walkthrough of this inference pipeline is provided in Appendix A.1.

4.2 Prediction Model

The primary goal of the prediction model is to dynamically decide the calls to the draft model during the draft phase, enabling the construction of a dynamic draft tree. Considering the sequential nature of the generation process, our prediction model is designed as an LSTM-based network, which takes a hidden vector initialized to zero and the top- k confidence scores from the draft model as inputs. While the MLP policy uses only the confidence scores as its state, the LSTM policy additionally maintains a recurrent hidden state that summarizes previous drafting steps. We use the confidence scores as inputs because there is a strong positive correlation between the draft model’s confidence score and the acceptance rate of the token (Li et al., 2024a). The model then outputs a binary control signal *continue/stop*, recorded as *logits*, and an updated hidden vector h .

We elaborate our prediction model on two aspects: problem modeling and model training.

Problem modeling. We formulate the decision of calls to the draft model as a Markov Decision Process (MDP), defined by a tuple $\mathcal{M} = (S, A, P, r, \gamma)$. The detailed definitions of these elements are as Eqs. (3-9):

- **State Space (S):** The state $s_t \in S$ is defined as the confidence scores of the top- k candidate tokens generated by the draft model at time step t , as Eq. (3):

$$s_t = [c_1, c_2, \dots, c_k], \quad 1 \leq t \leq t_{\max}, \quad (3)$$

where $c_i \in [0, 1]$, k denotes the branching factor of the draft tree, and $t_{\max}$ is the maximum number of calls to the draft model.

- **Action Space (A):** The action space consists of two actions, as defined in Eq. (4):

$$A = \{0, 1\}. \quad (4)$$

Action "0" indicates the termination of draft generation, while action "1" indicates the continuation to generate candidate tokens.

- **Transition Probability (P):** The state transition function $P(s_{t+1}|s_t, a_t)$ is defined as Eq. (5):

$$P(s_{t+1}|s_t, a_t) = \begin{cases} \delta(s_{t+1} = s_{\perp}) & \text{if } a_t = 0 \\ P_{\text{draft}}(s_{t+1}|s_t) & \text{if } a_t = 1 \end{cases}, \quad (5)$$

where $s_{\perp}$ is an absorbing terminal state reached when $a_t = 0$, and $P_{\text{draft}}(\cdot)$ represents the transition probability distribution based on the draft model's generation when $a_t = 1$.

- **Reward Function (r):** The reward r_t imposes a constant penalty $-\alpha$ at each step to discourage excessive drafting, and provides a final reward estimating the inference speed upon termination. A detailed interpretation of α is provided in Appendix A.4. The reward at time step t is defined as Eq. (6):

$$r_t = \begin{cases} -\alpha & \text{if } t < T \\ \frac{\ell_{\text{acc}}}{T_{\text{gen}}(T)} & \text{if } t = T \end{cases}, \quad (6)$$

where ℓ_{acc} is the acceptance length. $T_{\text{gen}}(t)$ estimates the total time cost, defined as Eq. (7):

$$T_{\text{gen}}(t) = T_o + T_d \cdot t + T_{\text{pred}}(t), \quad (7)$$

$$T_{\text{pred}}(t) = \begin{cases} T_p \cdot (t + 1) & \text{if } t < t_{\text{max}} \\ T_p \cdot t & \text{if } t = t_{\text{max}} \end{cases}, \quad (8)$$

where T_o denotes a fixed overhead, T_d denotes the latency of a single forward pass of the draft model, T_p denotes the latency of a single forward pass of the prediction model, and T_{pred} is the inference cost of the prediction model, which is different at the maximum step t_{max} , as the draft tree generation terminates, thus eliminating the need for a subsequent call to the prediction model. A breakdown of latency is provided in Appendix A.2.

- **Discount Factor (γ):** We set the discount factor to balance immediate and future rewards, following standard reinforcement learning practices (Mnih et al., 2015), as Eq. (9):

$$\gamma = 0.99. \quad (9)$$

Model training. The training loss is the negative expected reward, as Eq. (10):

$$\mathcal{L}(\theta) = -E_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \left(\sum_{t'=t}^T \gamma^{t'-t} r_{t'} \right) \log \pi_{\theta}(a_t | s_t) \right], \quad (10)$$

where $\tau = (s_1, a_1, \dots, s_T)$ is a trajectory sampled from policy π_{θ} .

4.3 Offline Dataset

We construct an offline dataset to train the prediction model to avoid expensive real-time interactions with the LLM during training. Following the description of the dataset and its construction, we explain the use of our dataset.

For each data point x in the dataset, $x = \{\mathcal{S}, \mathcal{D}\}$, where $\mathcal{S} = \{s_1, s_2, \dots, s_{t_{\text{max}}}\}$ denotes the sequence of states, where s_t and t_{max} are defined by Eq. (3); $\mathcal{D} = \{d_1, d_2, \dots, d_{t_{\text{max}}}\}$ denotes the set of distributions for acceptance lengths corresponding to different numbers of calls to the draft model. Specifically, for each i , $d_i = \{p_0, p_1, \dots, p_{t_{\text{max}}}\}$ where p_j denotes the probability that the acceptance length is j given that i calls were made to the draft model.

To construct the dataset, we run EAGLE-3 on ShareGPT¹ dataset, and for each prefix in the corpus, we repeat the following procedure. First, during the draft stage, we enumerate the number of calls to the draft model from 1 to t_{max} , to generate multiple draft trees. The confidence scores of the top- k tokens at each generation step are collected to form the state sequence $\mathcal{S}$. Next, draft trees are serialized and verified in parallel by the LLM, and during the verification phase, for each i , the distribution d_i of acceptance lengths for the draft tree generated with i calls is derived. These distributions are then aggregated into the set $\mathcal{D}$ and we can obtain one data point $x = \{\mathcal{S}, \mathcal{D}\}$.

Here, we detail the derivation of the acceptance length distribution, defined by $d_i = \{p_0, p_1, \dots, p_{t_{\text{max}}}\}$. For a draft tree node v_1 with parent v_0 and children $l_1, l_2, \dots, l_n$, as shown in Figure 2, the probability p_i used in the dataset is defined by Eqs. (11-13).

$$p_i = \sum_{v, \text{depth}=i} P_{\text{stop}}(v), \quad (11)$$

$$P_{\text{stop}}(v_1) = P_{\text{acc}}(v_0) \cdot P_{\text{acc}}(v_1 | v_0) \cdot P_{\text{rej}}(l | v_1), \quad (12)$$

$$P_{\text{rej}}(l | v_1) = 1 - \sum_{j=1}^n P_{\text{acc}}(l_j | v_1), \quad (13)$$

where $P_{\text{stop}}(v_1)$ denotes the probability that v_1 is accepted while all of its children are rejected by LLM during verification stage; $P_{\text{acc}}(v_0)$ denotes the probability that node v_0 is accepted, which can be calculated based on speculative sampling algorithm (Leviathan et al., 2023; Chen et al., 2023); $P_{\text{rej}}(l | v_1)$ denotes the probability that all children

are rejected given that their parent v_1 is accepted; $P_{\text{acc}}(v_1|v_0)$ denotes the probability that the child node v_1 is accepted given that its parent v_0 has been accepted, and p_i denotes the probability that the acceptance length equals i .

To use the dataset for training prediction model, for each data point $x = \{\mathcal{S}, \mathcal{D}\}$, the state sequence $\mathcal{S}$ is input into the prediction model, which produces a complete trajectory denoted as $\tau = (s_1, a_1, \dots, s_T)$. From this trajectory, the number of calls to the draft model, denoted by i , is obtained. Subsequently, the acceptance length ℓ_{acc} is sampled from the distribution $d_i \in \mathcal{D}$. The reward and loss are computed using the reward function defined in Eq. (6) and the loss function defined in Eq. (10), respectively. Finally, the parameters of the prediction model are updated using policy gradient methods. Through the above steps, both the data distribution and the behavior of the predictive model in offline reinforcement learning remain consistent with those in online reinforcement learning, thereby eliminating the impact of extrapolation error on the training process (Fujimoto et al., 2019).




Figure 2: Demonstration of the node in a draft tree for the derivation of the acceptance length distribution. The node v_1 has parent v_0 and children $l_1, l_2, \dots, l_n$, with a depth of i .

5 Experiment

5.1 Setup

Target LLMs. We conduct experiments with state-of-the-art open-source chat and reasoning models, including LLaMA-Instruct 3.1 8B (Grattafiori et al., 2024), Vicuna 13B (Chiang et al., 2023) and DeepSeek-R1-Distill-LLaMA 8B (Guo et al., 2025).

Tasks. Following EAGLE (Li et al., 2024b) and HASS (Zhang et al., 2024), we conduct evaluations on four common tasks, using the same weights for all tasks without fine-tuning on the respective

tasks. For multi-turn conversation, mathematical reasoning, instruction following and code generation, we chose the MT-bench (Zheng et al., 2023), GSM8K (Cobbe et al., 2021), Alpaca (Taori et al., 2023) and MBPP (Austin et al., 2021) datasets, respectively.

Metrics. RADAR does not fine-tune the target LLM’s weights and adopts a strict speculative sampling algorithm, ensuring no loss in performance. Thus, we do not evaluate generation quality. We use the following two metrics to assess acceleration performance:

- **Speedup Ratio:** The actual test speedup ratio relative to vanilla auto-regressive decoding.
- **Average Acceptance Length τ :** The average number of tokens generated per drafting-verification cycle, which corresponds to the number of tokens accepted from the draft.

Comparisons. We use vanilla auto-regressive decoding as the baseline, serving as the benchmark for speedup ratios (1.00x). We compare RADAR with recent lossless speculative decoding methods, including EAGLE-2 (Li et al., 2024a) and EAGLE-3 (Li et al., 2025), as well as recent adaptive drafting methods, including SpecDec++ (Huang et al., 2024), DISCO (Mamou et al., 2024), and LTD (Zhang et al., 2026).

Implementation. Our code is built based on EAGLE-3’s open-source repository. During training, we adopt the REINFORCE algorithm (Williams, 1992) for training prediction model, and set the learning rate to $1e-4$. During evaluation, following EAGLE-3’s setup, k is set to 10, the maximum number of calls to the draft model is 8, temperature is set to 1.0 and batch size is set to 1. More details are provided in Appendix A.3.

5.2 Effectiveness of RADAR

Note that the speedup ratio is hardware-sensitive due to variations in computational capabilities, most of the previous studies (Leviathan et al., 2023; Chen et al., 2023; Miao et al., 2024; Li et al., 2024b) performed inference tests in the same environment. Thereby, in our experiments, all inference tests were run on $2 \times$ NVIDIA RTX3090 GPUs. We additionally report results on LLaMA-Instruct 3.3 70B using $2 \times$ RTX PRO 6000 GPUs in Appendix A.6.

As shown in Table 1, we present different methods’ speedup ratios and acceptance lengths across

Model	Method	MT-bench		GSM8K		Alpaca		MBPP	
		Speedup	τ	Speedup	τ	Speedup	τ	Speedup	τ
L3 8B	Eagle-2	2.56x	3.37	3.43x	3.88	2.89x	4.08	3.29x	4.55
	Eagle-3	3.08x	4.55	4.68x	5.50	3.86x	5.62	4.21x	6.12
	Eagle-3+DISCO	3.10x	4.07	4.51x	5.08	3.68x	5.32	3.71x	5.74
	Eagle-3+Spec++	3.12x	4.40	4.44x	5.39	3.63x	5.36	3.98x	5.81
	Eagle-3+LTD	2.98x	5.12	4.61x	6.04	3.72x	6.47	4.28x	6.89
	Eagle-3+RADAR	3.41x	4.48	4.82x	5.32	4.04x	5.51	4.44x	6.00
V 13B	Eagle-2	2.89x	4.38	3.18x	4.46	2.83x	4.09	3.56x	4.93
	Eagle-3	3.74x	5.69	4.24x	5.94	3.50x	5.51	4.55x	6.44
	Eagle-3+DISCO	3.68x	5.67	3.64x	5.86	3.47x	5.63	4.25x	6.62
	Eagle-3+Spec++	3.60x	5.54	3.64x	5.81	3.37x	5.61	4.38x	6.45
	Eagle-3+LTD	3.78x	6.62	3.82x	6.60	3.31x	6.12	4.56x	7.98
	Eagle-3+RADAR	4.05x	5.67	4.36x	5.87	3.84x	5.48	4.75x	6.42
DSL 8B	Eagle-3	3.42x	4.88	4.39x	6.39	3.08x	4.49	3.71x	5.35
	Eagle-3+LTD	3.56x	4.16	4.15x	4.61	3.39x	4.03	3.97x	4.40
	Eagle-3+RADAR	3.86x	4.85	4.71x	6.33	3.17x	4.44	3.99x	5.31

Table 1: Speedup ratios and average acceptance lengths τ of different methods setting temperature=1. L3 represents LLaMA-Instruct 3.1, V represents Vicuna and DSL represents DeepSeek-R1-Distill-LLaMA. Spec++ stands for SpecDec++. The bold text in the speedup ratio column indicates the highest value for each dataset and model.

four datasets. RADAR provides a speedup of approximately 3.17x–4.82x compared to vanilla autoregressive generation, with a 3%–12.9% improvement over EAGLE-3. Besides, RADAR maintains a high average acceptance length, only about 1.2% lower than that of EAGLE-3. On most of the tested tasks and target LLMs, RADAR achieves the highest speedup ratio.

Generally, given the same model and dataset, a higher acceptance length implies a greater speedup ratio, however, RADAR deviates from this pattern by reducing redundant draft calls. As shown in Table 2, we present the average number of calls to the draft model per draft-verification cycle across four different datasets with three LLMs as target model, respectively. Under the configuration where the maximum number of calls is set to 8, RADAR reduces the average number of calls by 9.3%–34.3% compared to EAGLE-3’s fixed 8 calls, with an average reduction of 18.7%. RADAR effectively decreases the frequency of draft model invocations and reduces the time overhead during the draft phase.

5.3 Case Study of Acceptance Length and Draft-Model Calls

Figure 3 illustrates the distribution of acceptance length and the number of calls to the draft model

Model	MT-bench	GSM8K	Alpaca	MBPP
L3 8B	5.25	6.19	6.20	6.60
V 13B	6.88	7.26	6.83	7.26
DSL 8B	6.10	7.20	5.85	6.47

Table 2: Average number of calls to draft model of RADAR under temperature=1.0. L3 represents LLaMA-Instruct 3.1, V represents Vicuna and DSL represents DeepSeek-R1-Distill-LLaMA. Following the original papers’ setup, EAGLE-3 has an preset number of calls of 8 across all tested models and datasets.

on MT-bench, using LLaMA-Instruct 3.1 8B as the target model. As shown in Figure 3(a), while maintaining a high acceptance length, RADAR significantly reduces the probability of an acceptance length of 0 compared with EAGLE-3. It also increases the frequency of acceptance lengths from 1 to 7, while the frequency at length 8 decreases slightly. This is because the dynamic drafting mechanism not only reduces draft-model calls but also transforms the draft tree from a tall and thin structure into a shorter and wider one, thereby increasing the acceptance probability at lower depths.

Figure 3(b) shows how RADAR dynamically determines the number of draft-model calls. The prediction model stops draft-tree generation early when further drafting is deemed inefficient. This




Figure 3: Distribution of acceptance length and draft-model calls on MT-bench with LLaMA-Instruct 3.1 8B as the target model under temperature = 1.0.

adaptive stopping reduces redundant computation while maintaining the acceptance length, thereby accelerating inference.

5.4 Ablation Study

In this section, we aim to validate the effectiveness of RADAR through a series of analyses, focusing on the robustness of the reward design and the choices of policy network architectures and input features. We also provide supplementary studies on LLaMA-Instruct 3.3 70B (Grattafiori et al., 2024) in Appendix A.6 for scalability, on HASS (Zhang et al., 2024) in Appendix A.7 for generality, and across decoding temperatures in Appendix A.5.

Robustness of α . We evaluate the impact of α on MT-bench across three target models, as shown in Figure 4. The throughput curves for all evaluated models exhibit a consistent bell-shaped trend within the range $[0.01, 0.08]$, indicating that overly small α values fail to sufficiently penalize redundant draft calls, while overly large values lead to premature stopping. The consistently strong performance around $\alpha \approx 0.05$ suggests that our reward design is robust across different model architectures. Furthermore, owing to the high efficiency

of our offline training pipeline, the entire training process for each model finishes within 30 minutes. This efficiency makes it practical to tune α via a simple grid search.




Figure 4: Throughput (tokens/sec) across different α values on MT-bench across three LLMs.

Cross-Device Robustness. We evaluate the sensitivity to the device-specific latency parameter in the reward on MT-bench using LLaMA-Instruct 3.1 8B. We train three RADAR variants, denoted as RADAR-3090, RADAR-4090, and RADAR-5090, where the latency parameter is profiled on the corresponding device. As shown in Table 4, RADAR maintains high performance even with mismatched parameters. This is because acceptance length is hardware-independent, and the relative magnitude of T_o , T_d , and T_p remains consistent across devices. While matched parameters yield the best results, device-specific tuning is trivially affordable since the training cost for the prediction model requires only about 60 TFLOPs and 30 minutes.

Policy Networks and Input Features. We evaluate different network architectures and input features on LLaMA-Instruct 3.1 8B, with results detailed in Table 3. Comparing network architectures, the LSTM-based policy consistently outperforms the MLP-based policy across all datasets, which validates that the draft phase is intrinsically a sequential decision-making process. By capturing temporal dependencies across generation steps, the LSTM architecture achieves more accurate termination decisions than the memoryless MLP. Regarding input features, using confidence scores yields a higher speedup than using the draft model’s hidden states, despite the latter containing richer contextual information. This efficiency stems from a strong positive correlation between confidence scores and token acceptance rates (Li et al., 2024a),

Policy	Input	MT-bench		GSM8K		Alpaca		MBPP	
		Speedup	τ	Speedup	τ	Speedup	τ	Speedup	τ
MLP	scores	3.01x	4.55	4.63x	5.26	3.59x	5.30	4.15x	5.87
LSTM	scores	3.41x	4.48	4.82x	5.32	4.04x	5.51	4.44x	6.00
LSTM	hidden states	3.10x	4.60	4.74x	5.49	3.89x	5.67	4.29x	6.05

Table 3: Ablation study on different policy networks and input features evaluated on LLaMA-Instruct 3.1 8B. "Scores" indicates the confidence scores of the top- k candidate tokens generated by the draft model, and "hidden states" stands for the hidden states from the draft model.

Method	RTX3090	RTX4090	RTX5090
Eagle-3	48.12	91.15	121.41
RADAR-3090	53.87	99.80	123.02
RADAR-4090	47.51	100.40	124.32
RADAR-5090	51.91	99.10	131.82

Table 4: Throughput (tokens/sec) with LLaMA-Instruct 3.1 8B as the target model across different GPUs on MT-bench. RADAR-3090, RADAR-4090 and RADAR-5090 represent models trained with latency parameters from NVIDIA RTX3090, NVIDIA RTX4090 and NVIDIA RTX5090, respectively.

which we further confirm by calculating the Pearson correlation of $r = 0.4719$ ($p = 7.45e^{-137}$) between the top-1 score and acceptance length. In contrast, leveraging hidden states demands sufficient semantic understanding, inherently necessitating more training data, extended training time, and a larger prediction model.

6 Conclusion

We propose RADAR, a speculative sampling method with RL-based dynamic draft trees. We formulate the draft tree generation as a MDP and employ a prediction model to make real-time decisions of calls to the draft model. To handle the lack of labeled data, which is challenging to obtain due to the randomness and sparsity, we estimate the acceptance length distribution to construct a novel dataset for offline reinforcement learning in training prediction model. Experimental results on three LLMs and four tasks demonstrate that RADAR achieves a speedup of 3.17x–4.82x over vanilla auto-regressive decoding and reduces the calls to the draft model by 18.7% averaged over all tested datasets, outperforming previous methods.

Limitations

While RADAR advances speculative decoding through RL-based dynamic draft trees, it has sev-

eral limitations that offer avenues for future work. Currently, our method controls only the depth of the draft tree. Extending RADAR to finer-grained topological control, such as per-branch pruning or width adaptation, within an offline RL framework remains an open challenge. Additionally, our prediction model consumes only the top- k confidence scores through a lightweight LSTM. Scaling up the prediction model and incorporating additional input features, such as positional context, may further improve the accuracy of termination decisions, and are left for future work.

Acknowledgments

We are greatly appreciative of all the anonymous reviewers for their valuable feedback and insightful suggestions.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. 2024. On-policy distillation of language models: Learning from self-generated mistakes. In *The twelfth international conference on learning representations*.
- Zachary Ankner, Rishab Parthasarathy, Aniruddha Nrusimha, Christopher Rinard, Jonathan Ragan-Kelley, and William Brandon. 2024. Hydra: Sequentially-dependent draft heads for medusa decoding. *arXiv preprint arXiv:2402.05109*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.

- Tianle Cai, Yuhong Li, Zhengyang Geng, Hongwu Peng, Jason D Lee, Deming Chen, and Tri Dao. 2024. Medusa: Simple llm inference acceleration framework with multiple decoding heads. *arXiv preprint arXiv:2401.10774*.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. 2023. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Wei-Lin Chiang, Zhuohan Li, Ziqing Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Scott Fujimoto, David Meger, and Doina Precup. 2019. Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, pages 2052–2062. PMLR.
- Xiangxiang Gao, Weisheng Xie, Yiwei Xiang, and Feng Ji. 2025. Falcon: Faster and parallel inference of large language models through enhanced semi-autoregressive drafting and custom-designed decoding tree. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23933–23941.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Kaixuan Huang, Xudong Guo, and Mengdi Wang. 2024. Specdec++: Boosting speculative decoding via adaptive candidate lengths. *arXiv preprint arXiv:2405.19715*.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2024a. EAGLE-2: Faster inference of language models with dynamic draft trees. In *Empirical Methods in Natural Language Processing*.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2024b. EAGLE: Speculative sampling requires rethinking feature uncertainty. In *International Conference on Machine Learning*.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2025. EAGLE-3: Scaling up inference acceleration of large language models via training-time test. In *Annual Conference on Neural Information Processing Systems*.
- Fangcheng Liu, Yehui Tang, Zhenhua Liu, Yunsheng Ni, Kai Han, and Yunhe Wang. 2024. Kangaroo: Lossless self-speculative decoding via double early exiting. *arXiv preprint arXiv:2404.18911*.
- Tianyu Liu, Yun Li, Qitan Lv, Kai Liu, Jianchen Zhu, Winston Hu, and Xiao Sun. 2025. PEARL: Parallel speculative decoding with adaptive draft length. In *The Thirteenth International Conference on Learning Representations*.
- Jonathan Mamou, Oren Pereg, Daniel Korat, Moshe Berchansky, Nadav Timor, Moshe Wasserblat, and Roy Schwartz. 2024. Accelerating speculative decoding using dynamic speculation length. *arXiv preprint arXiv:2405.04304*.
- Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Xinhao Cheng, Zeyu Wang, Zhengxin Zhang, Rae Ying Yee Wong, Alan Zhu, Lijie Yang, Xiaoxiang Shi, et al. 2024. Specinfer: Accelerating large language model serving with tree-based speculative inference and verification. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, pages 932–949.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.

Jiebin Zhang, Zhenghan Yu, Liang Wang, Nan Yang, Eugene J Yu, Zheng Li, Yifan Song, Dawei Zhu, Xingxing Zhang, Furu Wei, et al. 2026. Learning to draft: Adaptive speculative decoding with reinforcement learning. *arXiv preprint arXiv:2603.01639*.

Lefan Zhang, Xiaodan Wang, Yanhua Huang, and Ruiwen Xu. 2024. Learning harmonized representations for speculative sampling. *arXiv preprint arXiv:2408.15766*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.

Wangchunshu Zhou, Yuchen Eleanor Jiang, Ryan Cotterell, and Mrinmaya Sachan. 2023. Efficient prompting via dynamic in-context learning. *arXiv preprint arXiv:2305.11170*.

Yongchao Zhou, Kaifeng Lyu, Ankit Singh Rawat, Aditya Menon, Afshin Rostamizadeh, Sanjiv Kumar, Jean-François Kagy, and Rishabh Agarwal. 2024. [Distillspec: Improving speculative decoding via knowledge distillation](#). In *International Conference on Learning Representations (ICLR)*.

A RADAR

A.1 Inference Pipeline Walkthrough

Here, we detailed the inference pipeline of our approach by the example shown in Figure 1. First, consider "What" as a query into the LLM, the token "can" is generated. Along with some features from the LLM, the token "can" is then fed into the draft model for a forward pass, generating two tokens: "I" and "you", and their confidence scores are returned simultaneously. After that, the prediction model produces a control signal *continue/stop* and a hidden vector h with the confidence scores as inputs. When receiving a *continue* signal, the draft model performs another forward pass with "I" and "you" as inputs, and produces four tokens: "say," "do", "see" and "do", and the top- k confidence scores, corresponding to the tokens "say" and "see" in Figure 1 (a). The above steps repeat iteratively. When receiving a *stop* signal, the draft model halts further forward, and the final draft tree is generated, as shown in Figure 1 (b).

A.2 Breakdown of the Latency

Here we provide a breakdown of the latency in Table 5. All data are collected on $2 \times$ NVIDIA RTX3090 GPUs.

Model	T_o	T_d	T_p
L3 8B	0.068s	0.0027s	0.0004s
V 13B	0.073s	0.0025s	0.0004s
DSL 8B	0.0668s	0.0028s	0.0004s

Table 5: Breakdown of the latency for T_o , T_d , and T_p . L3 represents LLaMA-Instruct 3.1, V represents Vicuna and DSL represents DeepSeek-R1-Distill-LLaMA.

A.3 Implementation Details

Our prediction model is an LSTM-based network, which consists of a single-layer LSTM followed by a two-layer MLP head. The LSTM outputs are concatenated with the original inputs and projected through the two-layer MLP to produce 2-dimensional logits, which are then transformed into a binary action distribution via Softmax. Detailed parameters of LSTM are used throughout: input_size=10, hidden_size=128, num_layers=1, dropout=0.1. Detailed parameters of MLP are as follows: hidden size=128, ReLU activation, dropout=0.1.

For offline dataset construction, we sample 1,000 prompts from ShareGPT and run EAGLE-3 to collect confidence scores and estimate the acceptance-length distributions. During training, we train the prediction model using REINFORCE with a learning rate of $1e-4$, a batch size of 256, for 100 epochs.

A.4 Interpretation of alpha as a Lagrange Multiplier

The latency terms T_o , T_d , and T_p are already included in the terminal throughput reward in Eq. (6); therefore, α is not an additional latency estimate. Instead, the intermediate penalty supplies a dense learning signal for an otherwise sparse delayed reward, eases credit assignment, and regularizes excessive continuation.

More formally, let R_T denote the terminal throughput reward and let N_c denote the number of *continue* actions in a trajectory. Consider the constrained objective in Eq. (14), which maximizes expected terminal throughput subject to an expected draft-computation budget B :

$$\max_{\pi} \mathbb{E}_{\pi}[R_T] \quad \text{s.t.} \quad \mathbb{E}_{\pi}[N_c] \leq B. \quad (14)$$

Its Lagrangian is given by Eq. (15):

$$\mathcal{J}(\pi, \alpha) = \mathbb{E}_{\pi}[R_T] - \alpha(\mathbb{E}_{\pi}[N_c] - B), \quad \alpha \geq 0. \quad (15)$$

For a fixed B , optimizing the policy is equivalent, up to the constant αB , to maximizing $\mathbb{E}_{\pi}[R_T] -$

Aspect	RADAR	LTD
RL paradigm	Offline reinforcement learning	Online reinforcement learning
Policy model	LSTM-based policy	MLP-based policies
Number of policy networks	1	2
Control variables	Draft tree depth via adaptive stopping	Draft tree depth and number of selected draft nodes
Training data	Offline dataset generated from 1,000 prompts sampled from ShareGPT	Real-time rollouts from 80 HumanEval samples (Chen et al., 2021) during training
LLM interactions during policy optimization	0	100,000 for size policy and 1,000,000 for depth policy
Training time	Within 30 minutes	20–30 GPU hours

Table 6: Comparison between RADAR and LTD in methodology and training cost.

$\alpha \mathbb{E}_\pi[N_c]$, which is the objective induced by our per-step penalty. Thus, α acts as a Lagrange multiplier that controls the trade-off between terminal throughput and draft computation. Tuning it requires no additional interaction with the target LLM because all candidates reuse the same offline confidence sequences and acceptance-length distributions.

A.5 Robustness across Decoding Temperatures

We further evaluate RADAR with LLaMA-Instruct 3.1 8B under different decoding temperatures. Table 7 reports a controlled temperature sweep. RADAR consistently outperforms EAGLE-3 across all evaluated settings, showing that the learned stopping policy remains effective under both greedy decoding and stochastic sampling.

Dataset	Temperature	EAGLE-3	RADAR
GSM8K	0.0	4.13x	4.18x
HumanEval	0.0	4.40x	4.48x
MBPP	0.0	4.71x	4.72x
MT-bench	0.0	4.01x	4.12x
MT-bench	0.3	3.85x	3.87x
MT-bench	0.5	3.74x	3.95x
MT-bench	0.7	3.44x	3.51x
MT-bench	1.0	3.08x	3.41x

Table 7: Speedup ratios of EAGLE-3 and RADAR under different decoding temperatures with LLaMA-Instruct 3.1 8B.

A.6 Results on LLaMA-Instruct 3.3 70B

To evaluate scalability on larger models, we further conduct experiments with LLaMA-Instruct 3.3 70B (Grattafiori et al., 2024) as the target LLM. These experiments are performed on $2 \times$ RTX PRO 6000 GPUs, and the speedup results are reported in Table 8. Since the hardware differs from that used

in the main experiments, these results serve as a supplementary evaluation.

Method	MT-bench	GSM8K	Alpaca	MBPP
Eagle-3	4.42x	4.83x	4.93x	5.35x
RADAR	4.45x	4.95x	5.07x	5.55x

Table 8: Speedup ratios of Eagle-3 and RADAR across different datasets setting temperature=1 with LLaMA-Instruct 3.3 70B as target LLM on $2 \times$ RTX PRO 6000 GPUs.

A.7 RADAR Effectiveness on HASS

To further validate the generality of our approach across different speculative decoding frameworks, we apply RADAR to HASS (Zhang et al., 2024), another tree-based speculative decoding method.

Method	MT-bench	GSM8K	Alpaca	MBPP
HASS	2.49x	2.70x	2.61x	2.87x
RADAR	2.56x	2.80x	2.73x	3.03x

Table 9: Speedup ratios of HASS and RADAR on LLaMA-Instruct 3.1 8B across four benchmarks on $2 \times$ RTX3090 GPUs setting temperature=1.

As shown in Table 9, RADAR consistently improves the speedup of HASS across all four benchmarks on LLaMA-Instruct 3.1 8B (Grattafiori et al., 2024).

B Comparison

B.1 Comparison with LTD

We provide a more detailed comparison between RADAR and the concurrent work LTD (Zhang et al., 2026), as shown in Table 6. While both methods apply reinforcement learning to adaptive speculative decoding, they differ substantially in

policy architecture, training paradigm, data construction, and training cost. In particular, RADAR is built on an offline RL framework, whereas LTD relies on online RL. Consequently, RADAR avoids repeated interaction with the target LLM during policy optimization, since the offline dataset only needs to be constructed once and can thereafter be reused across different devices.

B.2 Baseline Settings

For EAGLE-2 (Li et al., 2024a) and EAGLE-3 (Li et al., 2025), the code and weights are obtained from their official repository². For LTD (Zhang et al., 2026), the training and evaluation scripts are from its official repository³. For SpecDec++ (Huang et al., 2024) and Disco (Mamou et al., 2024), we train them from scratch using an MLP-based prediction head for 100 epochs. Both models are trained with the AdamW optimizer (lr=1e-4).

C Licenses of Artifacts

The artifacts used in this work are licensed as follows: LLaMA-Instruct 3.1 8B/3.3 70B (LLAMA 3.1 COMMUNITY LICENSE/LLAMA 3.3 COMMUNITY LICENSE), Vicuna 13B v1.3 (Non-commercial license), DeepSeek-R1-Distill-LLaMA (MIT License), EAGLE-2/3 and HASS (Apache 2.0), MT-bench (Apache 2.0), GSM8K (MIT License), Alpaca (CC BY-NC 4.0), MBPP (CC BY 4.0), and ShareGPT (Apache 2.0). All artifacts are used strictly for non-commercial academic research, consistent with their respective licenses. RADAR is also intended solely for academic research on speculative decoding and LLM inference acceleration.

²<https://github.com/SafeAILab/EAGLE>

³<https://github.com/zhzihao/Learning-to-Draft>