

Double: Breaking the Acceleration Limit via Double Retrieval Speculative Parallelism

Yuhao Shen [§] Tianyu Liu [◇] Junyi Shen [‡] Jinyang Wu [♣]

Quan Kong [§] Huan Li [§] Cong Wang ^{* §}

[§]Zhejiang University [◇]University of Science and Technology of China

[‡]National University of Singapore [♣]Tsinghua University

{riven, quankong, lihuan.cs, cwang85}@zju.edu.cn

tianyu_liu@mail.ustc.edu.cn j1shen@comp.nus.edu.sg

wu-jy23@mails.tsinghua.edu.cn

Abstract

Parallel Speculative Decoding (PSD) accelerates traditional Speculative Decoding (SD) by overlapping draft generation with verification. However, it remains hampered by two fundamental challenges: (1) a theoretical speedup ceiling dictated by the speed ratio between the draft and target models, and (2) high computational waste and pipeline stall due to mid-sequence token rejections of early errors. To address these limitations, we introduce DOUBLE (Double Retrieval Speculative Parallelism). By bridging the gap between SD and PSD, our framework resolves the Retrieval *Precision-Efficiency Dilemma* through a novel synchronous mechanism. Specifically, we enable the draft model to execute iterative retrieval speculations to break the theoretical speedup limits; to alleviate rejections without rollback, the target model performs authoritative retrieval to generate multi-token guidance. DOUBLE is entirely training-free and lossless. Extensive experiments demonstrate state-of-the-art speedup of **5.3**× on LLaMA3.3-70B and **2.8**× on Qwen3-32B, significantly outperforming the advanced method EAGLE-3 that requires extensive model training. Our code is available at <https://github.com/Sylvan820/Double1>.

1 Introduction

The sequential nature of autoregressive decoding inherently restricts parallelization as each token’s generation is contingent upon its predecessors (Brown et al., 2020a). While techniques like quantization, distillation, and efficient attention (Hinton et al., 2015; Dao et al., 2022; Choi et al., 2018) alleviate computational burdens, inference remains bottlenecked by memory bandwidth.

Speculative Decoding (SD) has emerged as a promising lossless approach to address this

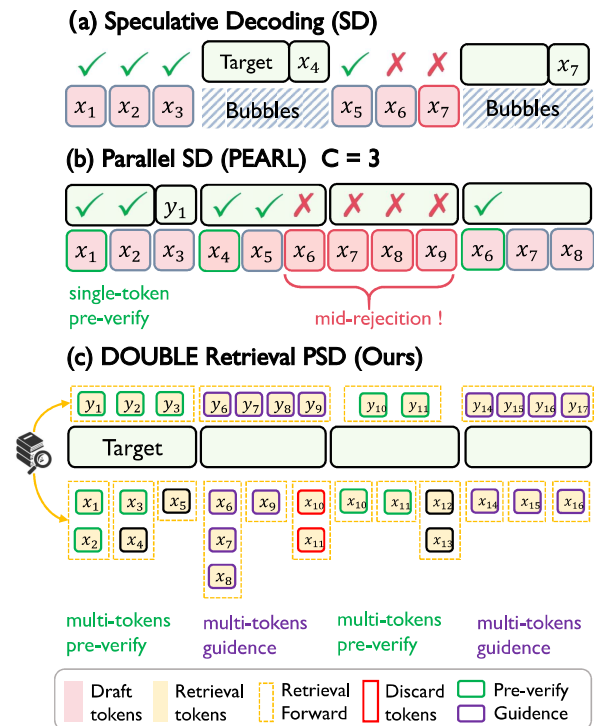


Figure 1: Comparison between SD, PSD and DOUBLE. (a) SD suffers from pipeline bubbles due to sequential dependency. (b) PSD overlaps these processes to reduce latency but struggles with *mid-sequence rejection*, where tokens generated are wasted after an early error (e.g., red boxes x_{6-9}). (c) DOUBLE resolves these issues through a double-retrieval mechanism: draft model leverages retrieval to expand draft length (more tokens x_{1-5}); target model performs retrieval to offer **multi-token pre-verify and guidance** (retrieve y_{1-3} for verifying x_{1-3}) to ensure precision, thereby breaking the speed limit C jointly and mitigating rejection penalties.

memory-bound limitation (Leviathan et al., 2023; Chen et al., 2023). SD utilizes a lightweight *draft model* to propose a sequence of γ tokens, which are then verified in parallel by the larger *target model*. However, the predominant draft-then-verify paradigm in most SD frameworks introduces a significant mutual waiting bottleneck. As illustrated in Fig. 1, this sequential dependency between

*The Corresponding Authors.




Figure 2: Speedup ratios of different methods on HumanEval and CNN/DM. DOUBLE achieves a speedup of $5.3\times$ on LLaMA3.3-70B and $2.8\times$ on Qwen3-32B over than EAGLE3. Full results are available in Table 1.

the draft and verification stages leads to *Bubbles* with underutilized GPU resources. To enhance hardware efficiency, Parallel Speculative Decoding (PSD) frameworks like PEARL (Liu et al., 2024b) and SpecBranch (Shen et al., 2025a) have demonstrated substantial speedup by adopting a draft-when-verify approach, enabling the draft and verification processes to proceed concurrently. Despite these advancements, existing PSD methods face two primary challenges:

- PSD is strictly constrained by a **theoretical upper speedup C** (the draft-to-target latency ratio). Unlike SD which employs tailored distilled models (Li et al., 2025), PSD utilizes off-the-shelf small models (0.6B, 1.3B, 7B). This limits C to a modest $1.5 - 5\times$ relative to large targets (14B, 33B, 70B). Fig 2 shows PEARL achieves a marginal speedup $1.3\times$ on Qwen-0.6B/14B models ($C = 1.6$) due to the autoregressive drafting.
- Existing frameworks struggle with **mid-sequence token rejection**. As shown in Fig. 1, although PEARL employs additional target forwards for pre-verify, it incurs high overhead with low returns for single-token (only y_1 for x_1) predictions, causing the rejection of mid tokens (red boxes x_6-x_9). While SpecBranch attempts to mitigate this via an offline predictor, it introduces additional training costs and yields lossy predictions compared to the target model.

Question. Can we generate longer and high-quality multi-token proposals without paying *autoregressive* draft model compute per token, while keeping high acceptance?

Retrieval is a promising technique (He et al.,

2023): by reusing previously seen continuations, we can propose multiple tokens with minimal neural computation. However, existing retrieval decoding methods face a **Retrieval Precision-Efficiency Dilemma** when combined with speculation. Using retrieval only at the target side is precise but offers limited acceleration (Saxena, 2023; Luo et al., 2024; Liu et al., 2025; Fu et al., 2024), because the expensive target forward pass still remains; using retrieval at the draft side is desirable but often imprecise (Zhao et al., 2024), increasing rejection rates and causing parallel pipeline stalls. As a result, prior approaches typically improve either proposal speed or quality, but NOT both, that remain under the PSD ceiling and rejection rollback.

Based on these insights, we introduce DOUBLE, which inherits the parallel pipeline but comes with a new double retrieval design on both sides of the draft/target models. Specifically, the draft model executes γ iterations of retrieval-based speculation to replace the slow auto-regressive generation, which effectively breaks the speedup limit C . Meanwhile, the target model performs a single-step retrieval to generate multiple tokens, in order to guide the draft process and mitigate mid-sequence rejections. Our contributions are summarized:

- We extend the theoretical speculation process from single-round to multi-round, which unifies both SD and PSD frameworks with a formal proof of the speed limit C . This has laid a theoretical foundation for retrieval.
- We propose a novel Double Retrieval mechanism that resolves the *Precision-Efficiency Dilemma* on parallel architecture. This strategy significantly enlarges the drafting length and leverages the power of target model for multi-token verifi-

cation and forward guidance.

- Across multiple datasets and model pairs, DOUBLE achieves a SOTA speedup of **5.3** \times on LLama3.3-70B and **2.8** \times on Qwen3-32B, surpassing the latest methods such as EAGLE-3 (Li et al., 2025) even with sophisticated training.

2 Background

2.1 Speculative Decoding

Speculative Decoding (SD) operates on a *draft-then-verify* paradigm. A draft model \mathcal{M}_q autoregressively generates γ tentative tokens x_1, \dots, x_γ , which are subsequently evaluated by the target model \mathcal{M}_p in a single parallel forward pass. The acceptance probability for the i -th token is defined as: $\alpha_i = \min\left(1, \frac{p_{i-1}(x_i)}{q_{i-1}(x_i)}\right)$. If a token is rejected at x_i , a correction is sampled from the distribution $\text{norm}(\max(0, p_{i-1} - q_{i-1}))$; otherwise, if all candidates are accepted, an extra token is appended.

Theorem 1 (Single-Round) The expected token generated in a single round is $\mathbb{E}[L_s]$ (Leviathan et al., 2023):

$$\mathbb{E}[L_s] = \frac{1 - \alpha^{\gamma+1}}{1 - \alpha} \quad (1)$$

2.2 Retrieval Speculative Decoding

Retrieval-based SD accelerates the target model by substituting autoregressive drafting with retrieved sequences (He et al., 2023). Specifically, given x , the target model \mathcal{M}_p searches n -gram matched tokens and then retrieves d (retrieval depth) tokens from a datastore \mathcal{D}_R as draft tokens. Then \mathcal{M}_p verifies these tokens to yield s matched tokens and one correction, formalized as the **Retrieval Forward**:

$$\text{RETRIEVAL}(x, n) = \{x_1, x_2, \dots, x_s, x_{s+1}\} \quad (2)$$

Here, s denotes the matched length, ensuring adherence to the target distribution. We define the Average Matched Tokens (AMT) as $\mathbb{E}[s]$, which dictates the effective speedup. Crucially, while existing methods (Luo et al., 2024; Saxena, 2023) vary in \mathcal{D}_R , they share a common oversight: focusing solely on accelerating the target model while neglecting the potential of the draft model.

3 Theoretical Analysis and Motivation

3.1 Speedup Ceiling Theorem

To understand why existing PSD methods plateau, we extend the analysis to a multi-round perspective to unify SD and PSD, formalizing the speedup ceiling that necessitates retrieval-based acceleration.

Definitions Let T_p and T_q be single-token decoding times for target and draft models, with speed ratio $C = T_p/T_q$. Unlike Theorem 1 that does not consider SD and PSD jointly, we unify them by modeling the process as $k - 1$ rounds of continuously accepted tokens followed by a rejection in the k -th round, which approximates the bimodal acceptance distribution (Shen et al., 2025a).

Theorem 2 (Multi-Round) The expected number of tokens generated over k rounds is:

$$\mathbb{E}[L_k] = (k - 1)(\gamma + 1) + \frac{1 - \alpha^{\gamma+1}}{1 - \alpha} \quad (3)$$

Theorem 3 (PSD Speedup Ceiling) The theoretical speedup S_{PSD} strictly dominates S_{SD} but is upper bounded by C :

$$S_{\text{SD}} \leq S_{\text{PSD}} = \frac{(\mathbb{E}[L_k] - k + 1) \cdot C}{k \cdot \gamma + C} \leq \frac{k \cdot \gamma \cdot C}{k \cdot \gamma + C} \leq C \quad (4)$$

Proof. Within one verification step T_p , an autoregressive draft model produces at most $\lfloor T_p/T_q \rfloor \approx C$ tokens. Even with $\alpha = 1$, PSD throughput cannot exceed the draft model’s generation rate. Formal proofs are available in the Appendix B.

Theorem 3 establishes a hard ceiling: standard PSD with off-the-shelf models (e.g., $C \approx 2 - 5$) cannot exceed this factor regardless of acceptance rate. Fig. 3 confirms that both SD and PSD consistently trail this limit. Breaking this barrier requires decoupling draft length from draft latency.

3.2 Why Double Retrieval?

How can we overcome the autoregressive barrier of the draft model to maximize the effective speed ratio C ? While Retrieval-based SD proposes reusing tokens to skip autoregression steps (He et al., 2024), blindly applying existing methods to PSD exposes a fundamental **Retrieval Precision-Efficiency Dilemma** as shown in Fig. 3(b).

1 Target-Side Retrieval (High Precision, Low Efficiency) Methods like Token Recycling (Luo et al., 2024) only leverage the target model to retrieve. Although accurate (high acceptance), they are bottlenecked by the inherent latency of the target model and verification overhead as shown in the bottom-right of Fig. 3, thereby trapping these methods in a “Low Efficiency” regime.

2 Draft-Side Retrieval (High Efficiency, Low Precision) Ouroboros (Zhao et al., 2024) is the only method to shift retrieval to the draft model.

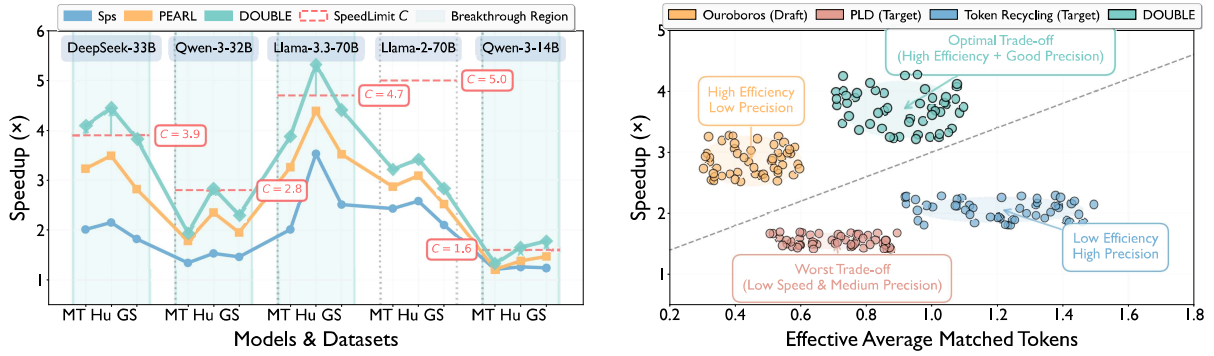


Figure 3: Motivation of DOUBLE. (a) Breaking the theoretical speedup ceiling C across five model pairs on three benchmarks. Green regions indicate where DOUBLE surpasses the speedup limit. (b) Retrieval precision-efficiency trade-off comparison on Deepseek-1.3B&33B, showing DOUBLE achieves optimal balance between effective matched tokens and speedup compared to draft-side (Ouroboros) and target-side (PLD, Token Recycling) methods.

Unfortunately, it faces two fundamental limitations: 1) Small draft models lack semantic understanding and capacity for accurate retrieval. The candidates fail verification frequently, resulting in low *effective* AMT after target verification despite long retrieval lengths. This exacerbates computational waste from **mid-sequence token rejection** (Shen et al., 2025a); 2) Ouroboros performs purely serial execution that fails to eliminate the mutual waiting bottleneck between models.

Motivation The goal is to achieve speed at draft-side retrieval of generating long chains ($> C$) and the precision of target-side retrieval to prevent those chains from being rejected. We propose DOUBLE to fill the precision-efficiency gap that runs retrieval on both draft/target sides simultaneously as detailed in the next section.

4 Methodology

To break the speedup ceiling, we introduce DOUBLE (Fig. 4), which consists of: 1) an iterative retrieval drafter that generates the candidate chains; (2) a parallel retrieval architecture that hides latency, and (3) a target-guided verification that turns potential rejections into forward extensions.

4.1 Parallel Retrieval Decoding

Iterative Retrieval Drafter To transcend the theoretical bound C , the draft model must generate more than C tokens with a single target forward pass T_p . Instead of standard autoregressive sampling of one token $x \sim q(\cdot|\mathbf{x})$, the draft model executes γ iterations of *Retrieval Forward* defined in Sec. 2.2 as shown in (Fig. 4(a)). Given context $\mathbf{x}^{(0)}$, in the j -th step iteration ($1 \leq j \leq C$), the

draft model retrieves phrase R_j from the datastore,

$$\mathcal{R}_j = \text{RETRIEVAL}(\mathbf{x}^{(j-1)}, d) = \{x_1^{(j)}, \dots, x_{s_j}^{(j)}\}, \quad (5)$$

where s_j represents the number of matched tokens in the j -th step, ensuring identical to the original distribution. Then we append R_j to the context and update for the next iteration $\mathbf{x}^{(j)} = \mathbf{x}^{(j-1)} \oplus \mathcal{R}_j$.

Why retrieval helps break this limit? In the time it takes an autoregressive draft model to generate C single tokens, the retrieval drafter generates C phrases. For example (Fig. 4(b)), retrieving “ARR is” ($j = 2$) immediately updates the context to retrieve “a top” ($j = 3$). If the average matched phrase length is AMT, the total draft length becomes $L_{\text{draft}} \approx C \times (1 + \text{AMT}) \gg C$, which breaks the limit and further pushes the effective proposal length beyond the theoretical speed ratio.

Notably, this sequential update forms a chain structure. Unlike tree-based SD methods that require complex tree attention masks for verification, our linear design remains compatible with standard attention mechanisms. This simplicity eliminates overhead, facilitating flexible deployment and high-concurrency scenarios.

Synchronous Parallel Execution We employ a *draft-when-verify* paradigm (Fig. 4(c)) to eliminate pipeline bubbles. While the target model \mathcal{M}_p verifies the candidate sequence from the previous round $k - 1$, the draft model \mathcal{M}_q constructs the candidate chain for the next round k . This synchronous process fully hides the verification latency, as \mathcal{M}_q prepares candidates of length $L_{\text{draft}} \gg C$ in parallel, effectively resolving the mutual waiting bottleneck.

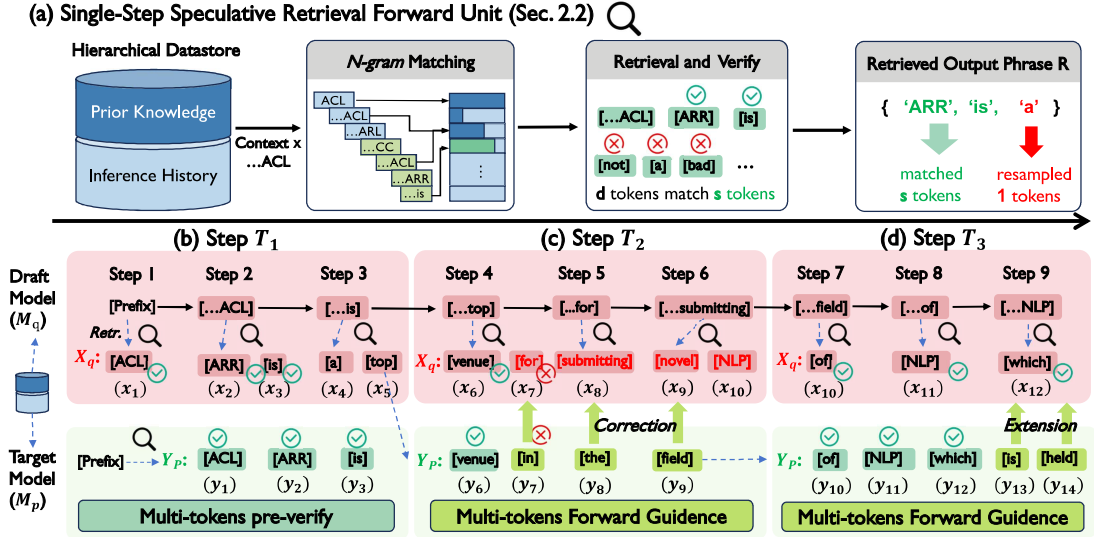


Figure 4: **Workflow of DOUBLE.** (a) **Retrieval Unit:** Utilizes hierarchical dastores to propose d candidates with match length s . (b) **Step T_1 :** At speed ratio $C = 3$, M_q executes iterative retrieval to draft 5 tokens, while M_p provide the multi-tokens pre-verify. (c) **Step T_2 :** Target retrieval rectifies x_{7-9} (“for submitting novel” \rightarrow “in the field”) as a *Correction*. (d) **Step T_3 :** Target retrieval directly extends the sequence beyond the draft as an *Extension*.

4.2 Target-Guided Verification

The previous design only expands the draft length but does not handle the low precision of draft-side retrieval. E.g., in standard PSD, if the draft model makes a mistake at token x_7 , PSD rejects x_7, \dots, x_N and the computational efforts are wasted. To mitigate such rollbacks, we propose a target-guided verification: the target model M_p performs a single-step retrieval forward to verify the next round unverified tokens and fetch high-confidence $Y_p = \text{RETRIEVAL}(x, d) = \{y_1, \dots, y_{s+1}\}$ to serve as multi-token pre-verify and forward guidance as detailed below.

1 Multi-token Pre-verify We employ speculative sampling to validate X_q against the distribution of Y_p . Unlike rigid equality checks, this process dynamically identifies a valid length i based on the acceptance criteria and allows divergent draft paths to be filtered out early before stalling the pipeline.

2 Forward Guidance We illustrate forward guidance via an example first. Consider a scenario where the draft model predicts a *low-quality* token at position 7 ($x_7 = \text{“for”}$). For standard PSD, the target model calculates $p(x_7)$, rejects “for”, and samples a new token, but this causes a pipeline stall. In contrast, in our framework, the target model retrieves a chain with higher quality, e.g., $Y_p = \{y_6 = \text{“venue”}, y_7 = \text{“in”}, y_8 = \text{“the”}, y_9 = \text{“field”}\}$. By matching the draft chain X_q against the target guide Y_p , we detect a mismatch at x_7 . Instead

of simple rejection, we swap the erroneous x_7 with the target’s prediction y_7 , and append the remaining retrieved tokens to the final output, maintaining the original distribution under greedy sampling.

In short, the target sequence Y_p acts as an active guidance to extend the generation. Formally, we construct the final output sequence by concatenating the verified draft prefix with the remainder of the target sequence. This operation unifies *correction* (Fig. 4(c): y_i mismatch x_i) and *extension* (Fig. 4(d): $s + 1 > L_{\text{draft}}$) into a single step:

$$x_{\text{next}} = \{x_1, \dots, x_{i-1}\} \oplus \{y_i, \dots, y_{s+1}\} \quad (6)$$

Here, the suffix $\{y_i, \dots, y_{s+1}\}$ can be considered as a pre-verified *bonus*, as the target model is no longer a passive verifier, but serves as an active guidance for the draft model with fast retrieval.

4.3 Hierarchical Dastore Construction

Existing retrieval methods also face a trade-off in dastore construction challenge. Large-scale external dastores (e.g., REST) incur high memory and training costs (He et al., 2024), while on-the-fly dastores (e.g., PLD) suffer from cold-start problems of context sparsity during early inference (Saxena, 2023). Thus, we propose a hierarchical dastore strategy.

Prior Knowledge Initialization To mitigate the cold-start problem without the latency of external retrieval, we pre-populate the dastore with a compact n -gram index derived from limited inference on generic corpora (ShareGPT (Zheng et al.,

2023b)). Our analysis (Table 3) demonstrates that this lightweight prior suffices to provide robust cross-domain generalization, establishing an immediate retrieval basis with negligible memory cost.

Hybrid Datastore Integration We construct a hybrid datastore that combines static contexts with dynamic inference history to optimize retrieval coverage. Instead of separate indices, such a datastore is shared by both draft and target models to eliminate redundancy and synchronization overhead. It is initialized with the static prior and user prompt to resolve early-stage sparsity. During decoding, tokens verified by the target model are prioritized, but the rejected tokens are also fused. This ensures retrieval with high quality and diversity.

4.4 Theoretical Speedup Analysis

Recall that standard PSD is bounded by $S_{\text{PSD}} \leq C$. DOUBLE breaks this barrier by decoupling the number of draft tokens from the number of forward passes, which effectively scales up the speed ratio to $C' = C(1 + \text{AMT})$ via retrieval-based generation. Incorporating the draft gain $\mathbb{E}[L_d]$ and target guidance $\mathbb{E}[L_{\text{bonus}}]$, the speedup becomes:

$$S_{\text{DOUBLE}} = \frac{(\mathbb{E}[L_d] + \mathbb{E}[L_{\text{bonus}}]) \cdot C(1 + \text{AMT})}{k \cdot \gamma + C(1 + \text{AMT})} \quad (7)$$

This demonstrates that DOUBLE can break the standard limit and elevate the new theoretical upper bound to a new level of $C(1 + \text{AMT})$.

5 Experiments

5.1 Experimental Setting

Tasks and Datasets We evaluate DOUBLE with a diverse suite of LLM configurations, including LLaMA-2 (7B/70B) (Grattafiori et al., 2024), LLaMA-3 (8B/70B) (Dubey et al., 2024), Deepseek-Coder (1.3B/33B) (Guo et al., 2024), and Qwen3 (0.6B/14B/32B) (Yang et al., 2025a). Our evaluation spans five benchmarks covering code, math, summarization, and chat: HumanEval (Chen et al., 2021), GSM8K (Cobbe et al., 2021), CN-N/DM (Nallapati et al., 2016), Alpaca (Taori et al., 2023), and MT-Bench (Zheng et al., 2023a).

Baselines and Implementation We benchmark against representative methods across different SD categories: 1) **Standard SD** (Chen et al., 2023); 2) **Target-side Retrieval** (Lookahead (Fu et al., 2024), PLD (Saxena, 2023), Token Recycling (Luo et al., 2024)); 3) **Draft-side Retrieval** (Ouroboros (Zhao

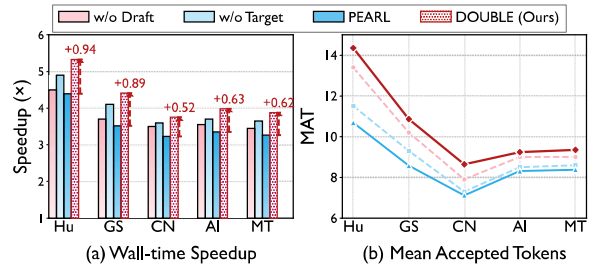


Figure 5: Ablation study on LLaMA-3.3-70B: both retrieval components are indispensable for achieving optimal speedup and accepted length.

et al., 2024)); 4) **Parallel SD** (PEARL (Liu et al., 2024b)); and 5) the SOTA **Training-based** method, EAGLE-3 (Li et al., 2025). Experiments are conducted on 8 NVIDIA A100 (80GB) GPUs under greedy sampling. We report Wall-time Speedup and Mean Accepted Tokens (M), where M denotes the continuously accepted length for parallel frameworks (Liu et al., 2024b). We set the draft length $\gamma = \lceil C \rceil$, with C representing the average speed ratio across datasets. Details are in Appendix D.1.

5.2 Main Results

Table 1 demonstrates that DOUBLE achieves consistent speedups ranging from $1.6\times$ to $5.3\times$ across five benchmarks, highlighting three key advantages: **(I) Breaking the PSD Speed Ceiling.** Existing parallel frameworks (e.g., PEARL) are strictly bounded by the draft-to-target speed ratio C . By decoupling draft length from latency, DOUBLE breaks this theoretical limit. This is evident on the constrained Qwen3-14B pair ($C \approx 1.6$), where DOUBLE achieves a $1.89\times$ speedup on CN/DM, surpassing the inherent ceiling and outperforming PEARL ($1.28\times$) by over 47%. **(II) Tackling the Precision-Efficiency Dilemma.** DOUBLE effectively overcomes the limitations of single-sided retrieval. Compared to Target-side methods (e.g., TR) limited by verification overhead, DOUBLE drastically boosts throughput (e.g., improving Deepseek-33B speedup from $1.96\times$ to $4.05\times$). Conversely, different from Draft-side only methods (e.g., Ouroboros) that are plagued by low precision, our target-guided mechanism ensures high acceptance. This leads to $3.42\times$ on LLaMA2-70B compared to Ouroboros’s $2.08\times$. **(III) Advantages over Training-Based Methods.** Despite *zero* training, DOUBLE matches or exceeds the SOTA training-based method, EAGLE-3. On models (LLaMA-3.3, Qwen3-32B), it attains speedups

Table 1: **Main results on five benchmarks.** Performance comparison between DOUBLE and existing baselines across diverse model configurations. Note that LLaMA-3.1-8B serves as the draft model for LLaMA-3.3-70B. **Bold** numbers denote the best performance and Underline denotes when DOUBLE breaks the speed limit.

Models	Methods	HumanEval		GSM8K		CNN/DM		Alpaca		MT-Bench		Avg.
		M	Speedup	M	Speedup	M	Speedup	M	Speedup	M	Speedup	
Deepseek-33B (C ≈ 3.9)	Lookahead	2.34	1.75×	1.83	1.45×	1.87	1.66×	1.69	1.34×	1.65	1.36×	1.51×
	PLD	1.86	1.68×	1.84	1.62×	2.07	1.87×	1.89	1.74×	1.69	1.47×	1.68×
	TR	2.83	2.14×	2.67	1.98×	2.88	2.17×	2.45	1.83×	2.36	1.72×	1.96×
	Ouroboros	6.38	3.15×	5.95	2.88×	5.97	2.92×	5.30	2.64×	5.12	2.36×	2.79×
	Sps	5.09	2.15×	3.94	1.82×	4.18	2.05×	4.22	2.16×	3.92	2.01×	2.04×
	PEARL	14.51	3.49×	7.83	2.82×	7.76	2.77×	8.71	2.86×	7.92	3.23×	3.03×
	DOUBLE	16.47	4.45×	10.86	3.83×	10.64	4.05×	10.24	3.84×	9.85	4.10×	4.05×
△ (↑, %)	↑ 13.5%	↑ 27.5%	↑ 38.7%	↑ 35.8%	↑ 37.1%	↑ 46.2%	↑ 17.6%	↑ 34.3%	↑ 24.4%	↑ 26.9%	↑ 33.7%	
LLaMA2-70B (C ≈ 5.0)	Lookahead	1.86	1.58×	2.04	1.62×	1.57	1.27×	1.69	1.31×	1.72	1.45×	1.45×
	PLD	1.66	1.53×	1.64	1.52×	1.97	1.77×	1.67	1.54×	1.49	1.45×	1.57×
	TR	2.43	2.02×	2.27	1.98×	2.58	2.21×	1.86	1.73×	1.75	1.67×	1.92×
	Ouroboros	5.08	2.08×	4.95	1.97×	5.21	2.25×	4.90	1.84×	4.82	1.76×	1.98×
	Sps	5.42	2.58×	4.34	2.10×	4.18	2.03×	4.22	2.06×	4.32	2.43×	2.24×
	PEARL	7.24	3.09×	6.83	2.52×	6.76	2.37×	7.71	2.66×	6.98	2.87×	2.70×
	DOUBLE	8.65	3.42×	7.86	2.83×	7.84	3.05×	8.24	2.94×	8.35	3.22×	3.09×
△ (↑, %)	↑ 19.5%	↑ 10.7%	↑ 15.1%	↑ 12.3%	↑ 16.0%	↑ 28.7%	↑ 6.9%	↑ 10.5%	↑ 19.6%	↑ 12.2%	↑ 14.4%	
LLaMA3.3-70B (C ≈ 4.7)	PLD	1.76	1.58×	1.54	1.42×	1.87	1.67×	1.69	1.54×	1.59	1.46×	1.53×
	Sps	5.40	3.53×	4.97	2.51×	4.83	2.46×	4.61	2.23×	4.30	2.01×	2.55×
	EAGLE3	6.51	4.56×	6.02	4.21×	4.82	3.03×	5.89	4.13×	5.23	4.02×	3.99×
	PEARL	10.69	4.39×	8.57	3.52×	7.12	3.23×	8.31	3.35×	8.38	3.26×	3.55×
	DOUBLE	14.35	5.33×	10.86	4.41×	8.64	3.75×	9.24	3.98×	9.35	3.88×	4.27×
	△ (↑, %)	↑ 34.2%	↑ 21.4%	↑ 26.7%	↑ 25.3%	↑ 21.3%	↑ 16.1%	↑ 11.2%	↑ 18.8%	↑ 11.6%	↑ 19.0%	↑ 20.3%
Qwen3-14B (C ≈ 1.6)	PLD	1.52	1.38×	1.44	1.31×	1.67	1.53×	1.53	1.35×	1.47	1.35×	1.38×
	Sps	3.37	1.26×	3.44	1.24×	3.38	1.13×	3.21	1.23×	3.17	1.21×	1.21×
	EAGLE3	2.99	1.91×	3.02	1.95×	2.41	1.58×	2.49	1.65×	2.83	1.89×	1.80×
	PEARL	7.24	1.38×	7.36	1.47×	5.62	1.28×	5.76	1.31×	3.77	1.21×	1.33×
	DOUBLE	12.65	1.65×	10.53	1.78×	9.64	1.89×	8.24	1.54×	5.13	1.32×	1.63×
	△ (↑, %)	↑ 74.7%	↑ 19.6%	↑ 43.1%	↑ 21.1%	↑ 71.5%	↑ 47.7%	↑ 43.1%	↑ 17.6%	↑ 36.1%	↑ 9.1%	↑ 22.6%
Qwen3-32B (C ≈ 2.8)	PLD	1.46	1.16×	1.41	1.12×	1.56	1.47×	1.49	1.31×	1.43	1.29×	1.27×
	Sps	4.27	1.53×	4.04	1.46×	3.38	1.33×	3.71	1.43×	3.47	1.34×	1.42×
	EAGLE3	2.83	2.27×	2.99	2.53×	2.57	1.73×	2.78	1.80×	3.03	2.23×	2.11×
	PEARL	7.78	2.35×	7.53	1.95×	4.12	1.72×	3.86	1.83×	3.68	1.78×	1.92×
	DOUBLE	10.09	2.84×	8.83	2.30×	6.64	2.56×	6.24	2.04×	5.83	1.94×	2.33×
	△ (↑, %)	↑ 29.7%	↑ 20.9%	↑ 17.3%	↑ 17.9%	↑ 61.2%	↑ 48.8%	↑ 61.7%	↑ 11.5%	↑ 58.4%	↑ 9.0%	↑ 21.4%

of 5.33× and 2.84× on HumanEval, respectively. This validates our strategy yields acceleration that can often exceed expensive architectural modifications, offering a promising alternative to EAGLE-3.

5.3 Ablation Study

Component Analysis To evaluate the component-wise contribution, we analyze the impact of Draft-side and Target-side retrieval on LLaMA-3.3. As shown in Fig. 5, the results demonstrate that both components are critical:

- **Impact of Draft Retrieval:** Removing draft retrieval (*w/o Draft*) reverts the system to autoregressive drafting, effectively constraining speedup to the theoretical ceiling C . By enabling iterative retrieval, DOUBLE decouples proposal length from latency, elevating the speedup on HumanEval from 4.50× to **5.33×**.
- **Impact of Target Retrieval:** Excluding target retrieval (*w/o Target*) degrades precision due to the draft model’s limitations, reducing Mean Ac-

cepted Tokens (MAT) from 14.35 to 11.5. Target guidance functions as a correction filter, saving tokens that are meant to be rejected in order to maintain high acceptance rates.

The synergy between them allows us to surpass PEARL (4.39×, 10.69 MAT), effectively resolving the Precision-Efficiency Dilemma by balancing high speed with robust verification precision.

Temperature Sampling Under non-greedy settings, strictly preserving the target distribution necessitates a modified verification protocol. Upon rejecting a token x_i , we sample a correction from the residual distribution $\text{norm}(\max(0, p_{i-1} - q_{i-1}))$. Although distributional consistency mandates discarding subsequent guidance $y_{>i}$, the *Multi-token Pre-verify* mechanism remains active for the prefix, effectively pruning divergent tokens to minimize overhead. Table 2 validates performance under stochastic sampling ($T = 1.0$): on LLaMA3.3-70B, DOUBLE achieves **4.70×** speedup on Hu-

Table 2: Performance comparison between DOUBLE and existing baseline under non-greedy settings.

Methods	HumanEval		GSM8K		CNN/DM		Alpaca		MT-Bench		Avg.	
	M	Speedup	M	Speedup	M	Speedup	M	Speedup	M	Speedup	M	Speedup
LLaMA3.3-70B, $Temperature = 1.0$												
EAGLE3	6.05	4.24×	5.60	3.91 ×	4.49	2.82×	5.48	3.84 ×	4.86	3.74 ×	5.30	3.71×
PEARL	10.06	4.13×	8.06	3.31×	6.69	3.04×	7.82	3.15×	7.89	3.07×	8.10	3.34×
DOUBLE	12.66	4.70 ×	9.58	3.89×	7.62	3.31 ×	8.15	3.51×	8.25	3.42×	9.25	3.77 ×
Qwen3-32B, $Temperature = 1.0$												
EAGLE3	2.63	2.11×	2.78	2.35 ×	2.39	1.61×	2.58	1.67×	2.81	2.07 ×	2.64	1.96×
PEARL	7.31	2.21×	7.07	1.83×	3.86	1.62×	3.63	1.72×	3.45	1.67×	5.06	1.81×
DOUBLE	8.90	2.50 ×	7.80	2.03×	5.86	2.26 ×	5.51	1.80 ×	5.15	1.72×	6.64	2.06 ×

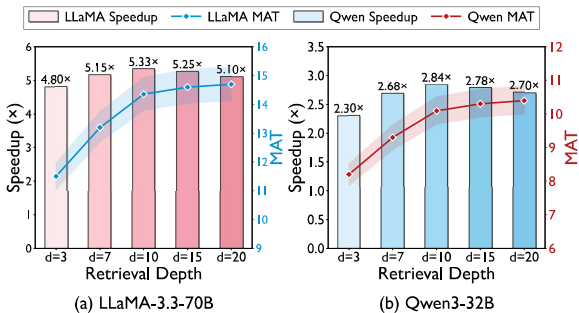


Figure 6: Impact of retrieval depth d . MAT saturates and speedup fluctuates beyond $d = 10$.

manEval and averages **3.77**×, consistently outperforming EAGLE3 (3.71×) and PEARL (3.34×). These results confirm that DOUBLE maintains high acceleration efficiency even in stochastic regimes.

Effect of Retrieval Depth We analyze the impact of retrieval depth d on efficiency. As shown in Fig. 6, increasing d initially boosts the Mean Accepted Tokens (M), but the performance saturates beyond $d = 10$. This is intuitive because multi-step prediction reduces accuracy. As the retrieval chain extends, distant tokens become less relevant to the current context, leading to a lower acceptance rate. Thus, we select $d = 10$ as the optimal balance.

Prior Retrieval Knowledge Table 3 analyzes the efficacy of our prior knowledge datastore, serialized as a lightweight local .pkl file. We demonstrate that a single datastore derived from generic ShareGPT data achieves robust generalization across diverse benchmarks (HumanEval, GSM8K, MT-Bench), obviating the need for task-specific adaptation. Performance gains saturate at 10 rounds (9.5 MB); extending initialization to 20 rounds yields diminishing returns while linearly increasing storage overhead. Consequently, we establish $K = 10$ as the optimal configuration to

Table 3: Impact of initialization rounds K on storage and speedup (Qwen3-32B).

Rounds (K)	Storage (MB)	Wall-time Speedup		
		HumanEval	GSM8K	MT-Bench
w/o Prior	0.0	2.35×	1.95×	1.78×
5	4.7	2.72×	2.21×	1.88×
10	9.5	2.84 ×	2.30 ×	1.94 ×
15	14.3	2.86×	2.31×	1.94×
20	19.3	2.86×	2.33×	1.96×

balance the overhead.

More Discussions. Due to space constraints, we provide extensive supplementary results and in-depth analyses in the Appendix. This includes scalability under high batch sizes (D.3), comparisons with well-engineered baselines (D.1), disentanglement of performance gains (D.2), and empirical proofs of lossless acceleration (D.4). Furthermore, we detail our hierarchical datastore with rigorous data decontamination protocols (D.5), system latency profiling (D.6), and adaptive deployments across memory-constrained and heterogeneous hardware scenarios (D.7, D.8).

6 Related Work

Speculative Decoding Standard SD accelerates inference by drafting tokens in a lossless manner, yet maximizing the acceptance rate remains a key challenge. Strategies range from training-based auxiliary modules (Medusa (Cai et al., 2024), EAGLE (Li et al., 2024)) to training-free verification trees (SpecInfer (Chen et al., 2023)). Although recent advancements like EAGLE-3 (Li et al., 2025) have optimized small model scaling, these methods adhere to a sequential *draft-then-verify* paradigm, which imposes an unavoidable mutual waiting bottleneck between models.

Parallel Speculative Decoding To address serial inefficiencies, PSD frameworks such as PEARL (Liu et al., 2024b) and SpecBranch (Shen et al., 2025a) introduce pipelined execution, enabling concurrent drafting and verification. While this effectively utilizes idle compute, the overall speedup remains strictly upper-bounded by the inherent speed ratio between draft and target models. Furthermore, these pipelines are highly sensitive to latency penalties once mid-token rejections occur.

Retrieval Speculative Decoding Retrieval-based methods utilize n -gram matching but typically isolate acceleration to either the **Target-Side** (Lookahead (Fu et al., 2024), PLD (Saxena, 2023)) or **Draft-Side** (Ouroboros (Zhao et al., 2024)). This isolation creates a fundamental *Precision-Efficiency Dilemma*. In contrast, we propose **DOUBLE**, which orchestrates synergistic retrieval across both models to resolve this trade-off, effectively breaking the theoretical speedup ceiling of existing parallel frameworks.

7 Conclusion

We propose a parallel speculative decoding framework with a double-retrieval mechanism. **DOUBLE** utilizes iterative drafting to expand candidate lengths beyond theoretical limits and leverages target-side retrieval to repair potential rejections actively, thus compensating for the limitations of each model. Our approach achieves training-free and lossless speedup of $5.3\times$ even against EAGLE-3 and establishes a new SOTA for LLM inference.

Limitations

Our study is comprehensive, but has certain limitations that we plan to address in future research. In this study, we employ a unified datastore and synchronized retrieval depth for both draft and target models. While this configuration simplifies deployment and optimizes memory management, it may not fully exploit the distinct capabilities of each model. We believe these are minor issues and we will explore model-specific optimizations, such as adaptive retrieval depths and decoupled high-efficiency datastores, to further unlock the potential of parallel speculative decoding in future.

Acknowledgements

This work is supported by National Science Foundation of China under grants 62576310, 62394341,

and Zhejiang Provincial National Science Foundation of China under Grant No. LZ25F020007 and Open Research Project of the State Key Laboratory of Industrial Control Technology, China (Grant No. ICT2026A20)..

Ethics Statement

The data and models utilized in this work are derived solely from publicly accessible resources with proper citations, so no sensitive information is involved. As **DOUBLE** is a speculative decoding framework designed to losslessly accelerate LLMs without parameter modification, it inherently inherits the biases and safety risks of the underlying models. It neither introduces new harmful capabilities nor mitigates existing ones. Therefore, standard safety guardrails and alignment techniques remain essential when deploying target models with **DOUBLE**.

References

- Sudhanshu Agrawal, Wonseok Jeon, and Mingu Lee. 2024. Adaedl: Early draft stopping for speculative decoding of large language models via an entropy-based lower bound on token acceptance probability. *arXiv preprint arXiv:2410.18351*.
- Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. 2024. Quarot: Outlier-free 4-bit inference in rotated llms. *arXiv preprint arXiv:2404.00456*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020a. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020b. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- 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, and Greg Brockman et al. 2021. *Evaluating large language models trained on code*.
- Xiaodong Chen, Yuxuan Hu, Jing Zhang, Yanling Wang, Cuiping Li, and Hong Chen. 2024a. *Streamlining redundant layers to compress large language models*. Preprint, arXiv:2403.19135.
- Zhuoming Chen, Avner May, Ruslan Svirschevski, Yu-Hsun Huang, Max Ryabinin, Zhihao Jia, and Beidi Chen. 2024b. *Sequoia: Scalable and robust speculative decoding*. *Advances in Neural Information Processing Systems*, 37:129531–129563.
- Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. 2018. *Pact: Parameterized clipping activation for quantized neural networks*. *arXiv preprint arXiv:1805.06085*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. *Training verifiers to solve math word problems*. *Cornell University - arXiv, Cornell University - arXiv*.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. *Flashattention: Fast and memory-efficient exact attention with io-awareness*. *Advances in neural information processing systems*, 35:16344–16359.
- Cunxiao Du, Jing Jiang, Xu Yuanchen, Jiawei Wu, Sicheng Yu, Yongqi Li, Shenggui Li, Kai Xu, Liqiang Nie, Zhaopeng Tu, and 1 others. 2024. *Glide with a cape: A low-hassle method to accelerate speculative decoding*. *arXiv preprint arXiv:2402.02082*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, and Aiesha Letman et al. 2024. *The llama 3 herd of models*. Preprint, arXiv:2407.21783.
- Elias Frantar and Dan Alistarh. 2023. *Sparsegpt: Massive language models can be accurately pruned in one-shot*. In *International Conference on Machine Learning*, pages 10323–10337. PMLR.
- Yichao Fu, Peter Bailis, Ion Stoica, and Hao Zhang. 2024. *Break the sequential dependency of llm inference using lookahead decoding*. *arXiv preprint arXiv:2402.02057*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, and Anirudh Goyal et al. 2024. *The llama 3 herd of models*. Preprint, arXiv:2407.21783.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. *Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning*. *arXiv preprint arXiv:2501.12948*.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y.K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. 2024. *Deepseek-coder: When the large language model meets programming – the rise of code intelligence*.
- Zhenyu He, Zexuan Zhong, Tianle Cai, Jason Lee, and Di He. 2024. *Rest: Retrieval-based speculative decoding*. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1582–1595.
- Zhenyu He, Zexuan Zhong, Tianle Cai, Jason D Lee, and Di He. 2023. *Rest: Retrieval-based speculative decoding*. *arXiv preprint arXiv:2311.08252*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. *Distilling the knowledge in a neural network*. *arXiv preprint arXiv:1503.02531*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. *Lora: Low-rank adaptation of large language models*. *ICLR*, 1(2):3.
- Haiduo Huang, Fuwei Yang, Zhenhua Liu, Yixing Xu, Jinze Li, Yang Liu, Xuanwu Yin, Dong Li, Pengju Ren, and Emad Barsoum. 2025. *Jakiro: Boosting speculative decoding with decoupled multi-head via moe*. *arXiv preprint arXiv:2502.06282*.
- Kaixuan Huang, Xudong Guo, and Mengdi Wang. 2024. *Specdec++: Boosting speculative decoding via adaptive candidate lengths*. *arXiv preprint arXiv:2405.19715*.
- Yicheng Ji, Jun Zhang, Jinpeng Chen, Cong Wang, Lidan Shou, Gang Chen, and Huan Li. 2026. *See the forest for the trees: Loosely speculative decoding via visual-semantic guidance for efficient inference of video llms*. Preprint, arXiv:2604.05650.
- Quan Kong, Yuhao Shen, Yicheng Ji, Huan Li, and Cong Wang. 2026a. *Parallelvllm: Lossless video-llm acceleration with visual alignment aware parallel speculative decoding*. *arXiv preprint arXiv:2603.19610*.
- Quan Kong, Yanru Xiao, Yuhao Shen, and Cong Wang. 2026b. *Vision-ttt: Efficient and expressive visual representation learning with test-time training*. *arXiv preprint arXiv:2603.00518*.
- 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. 2024. Eagle: Speculative sampling requires rethinking feature uncertainty. *arXiv preprint arXiv:2401.15077*.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2025. Eagle-3: Scaling up inference acceleration of large language models via training-time test. *arXiv preprint arXiv:2503.01840*.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Weiming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:87–100.
- Fangcheng Liu, Yehui Tang, Zhenhua Liu, Yunsheng Ni, Duyu Tang, Kai Han, and Yunhe Wang. 2024a. Kangaroo: Lossless self-speculative decoding for accelerating llms via double early exiting. *Advances in Neural Information Processing Systems*, 37:11946–11965.
- Tianyu Liu, Yun Li, Qitan Lv, Kai Liu, Jianchen Zhu, and Winston Hu. 2024b. Parallel speculative decoding with adaptive draft length. *arXiv preprint arXiv:2408.11850*.
- Tianyu Liu, Qitan Lv, Hao Li, Xing Gao, and Xiao Sun. 2025. Logitspec: Accelerating retrieval-based speculative decoding via next next token speculation. *arXiv preprint arXiv:2507.01449*.
- Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. 2024c. Spinquant-llm quantization with learned rotations. *arXiv preprint arXiv:2405.16406*.
- Kevin Lu and Thinking Machines Lab. 2025. **On-policy distillation**. *Thinking Machines Lab: Connectionism*. <https://thinkingmachines.ai/blog/on-policy-distillation>.
- Xianzhen Luo, Yixuan Wang, Qingfu Zhu, Zhiming Zhang, Xuanyu Zhang, Qing Yang, Dongliang Xu, and Wanxiang Che. 2024. Turning trash into treasure: Accelerating inference of large language models with token recycling. *arXiv preprint arXiv:2408.08696*.
- Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. 2024. **Shortgpt: Layers in large language models are more redundant than you expect**. *Preprint*, arXiv:2403.03853.
- R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors. 1983. *Machine Learning: An Artificial Intelligence Approach, Vol. I*. Tioga, Palo Alto, CA.
- Saurav Muralidharan, Sharath Turuvekere Sreenivas, Raviraj Joshi, Marcin Chochowski, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro, Jan Kautz, and Pavlo Molchanov. 2024. **Compact language models via pruning and knowledge distillation**. *Preprint*, arXiv:2407.14679.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. **Abstractive text summarization using sequence-to-sequence rnns and beyond**. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Apoorv Saxena. 2023. **Prompt lookup decoding**.
- Yuhao Shen, Junyi Shen, Quan Kong, Tianyu Liu, Yao Lu, and Cong Wang. 2025a. Speculative decoding via hybrid drafting and rollback-aware branch parallelism. *arXiv preprint arXiv:2506.01979*.
- Yuhao Shen, Zichen Wang, Tianyi Wang, Chaojie Gu, Zhenyu Wen, Yuanchao Shu, and Cong Wang. 2025b. Hetero 2 pipe: Pipelining multi-dnn inference on heterogeneous mobile processors under co-execution slowdown. In *2025 IEEE 45th International Conference on Distributed Computing Systems (ICDCS)*, pages 483–493. IEEE.
- Sharath Turuvekere Sreenivas, Saurav Muralidharan, Raviraj Joshi, Marcin Chochowski, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro, Jan Kautz, and Pavlo Molchanov. 2024. Llm pruning and distillation in practice: The minitron approach. *arXiv preprint arXiv:2408.11796*.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. 2024. **A simple and effective pruning approach for large language models**. *Preprint*, arXiv:2306.11695.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Nadav Timor, Jonathan Mamou, Daniel Korat, Moshe Berchansky, Oren Pereg, Moshe Wasserblat, Tomer Galanti, Michal Gordon, and David Harel. 2024. Distributed speculative inference (dsi): Speculation parallelism for provably faster lossless language model inference. *arXiv preprint arXiv:2405.14105*.

- Chengbing Wang, Yang Zhang, Wenjie Wang, Xiaoyan Zhao, Fuli Feng, Xiangnan He, and Tat-Seng Chua. 2025. Think-while-generating: On-the-fly reasoning for personalized long-form generation. *arXiv preprint arXiv:2512.06690*.
- Chengbing Wang, Wuqiang Zheng, Yang Zhang, Fengbin Zhu, Junyi Cheng, Yi Xie, Wenjie Wang, and Fuli Feng. 2026. Perm: Psychology-grounded empathetic reward modeling for large language models. *arXiv preprint arXiv:2601.10532*.
- Tongxi Wang. 2026. Fbs: Modeling native parallel reading inside a transformer. *arXiv preprint arXiv:2601.21708*.
- Jinyang Wu, Changpeng Yang, Yuhao Shen, Fangzhi Xu, Bolin Ni, Chonghua Liao, Yuchen Liu, Hongzhen Wang, Shuai Nie, Shuai Zhang, and 1 others. 2026a. Ssl: Sweet spot learning for differentiated guidance in agentic optimization. *arXiv preprint arXiv:2601.22491*.
- Jinyang Wu, Shuo Yang, Changpeng Yang, Yuhao Shen, Shuai Zhang, Zhengqi Wen, and Jianhua Tao. 2026b. Spark: Strategic policy-aware exploration via dynamic branching for long-horizon agentic learning. *arXiv preprint arXiv:2601.20209*.
- Jinyang Wu, Guocheng Zhai, Ruihan Jin, Jiahao Yuan, Yuhao Shen, Shuai Zhang, Zhengqi Wen, and Jianhua Tao. 2026c. Atlas: Orchestrating heterogeneous models and tools for multi-domain complex reasoning. *arXiv preprint arXiv:2601.03872*.
- Heming Xia, Yongqi Li, Jun Zhang, Cunxiao Du, and Wenjie Li. 2024a. Swift: On-the-fly self-speculative decoding for llm inference acceleration. *arXiv preprint arXiv:2410.06916*.
- Heming Xia, Zhe Yang, Qingxiu Dong, Peiyi Wang, Yongqi Li, Tao Ge, Tianyu Liu, Wenjie Li, and Zhifang Sui. 2024b. Unlocking efficiency in large language model inference: A comprehensive survey of speculative decoding. *arXiv preprint arXiv:2401.07851*.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pages 38087–38099. PMLR.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Penghui Yang, Cunxiao Du, Fengzhuo Zhang, Haonan Wang, Tianyu Pang, Chao Du, and Bo An. 2025b. Longspec: Long-context lossless speculative decoding with efficient drafting and verification. *arXiv preprint arXiv:2502.17421*.
- Bowen Zeng, Feiyang Ren, Jun Zhang, Xiaoling Gu, Ke Chen, Lidan Shou, and Huan Li. 2026. Hybridkv: Hybrid kv cache compression for efficient multimodal large language model inference. *Preprint, arXiv:2604.05887*.
- Jun Zhang, Yicheng Ji, Feiyang Ren, Yihang Li, Bowen Zeng, Zonghao Chen, Ke Chen, Lidan Shou, Gang Chen, and Huan Li. 2026a. Efficient inference for large vision-language models: Bottlenecks, techniques, and prospects. *Preprint, arXiv:2604.05546*.
- Jun Zhang, Jue Wang, Huan Li, Lidan Shou, Ke Chen, Gang Chen, and Sharad Mehrotra. 2024a. Draft&verify: Lossless large language model acceleration via self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11263–11282.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026b. Logical phase transitions: Understanding collapse in llm logical reasoning. *arXiv preprint arXiv:2601.02902*.
- Yingtao Zhang, Haoli Bai, Haokun Lin, Jialin Zhao, Lu Hou, and Carlo Vittorio Cannistraci. 2024b. Plug-and-play: An efficient post-training pruning method for large language models. In *The Twelfth International Conference on Learning Representations*.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Zikai Song. 2025. From ambiguity to verdict: A semiotic-grounded multi-perspective agent for llm logical reasoning. *arXiv preprint arXiv:2509.24765*.
- Weilin Zhao, Yuxiang Huang, Xu Han, Wang Xu, Chaojun Xiao, Xinrong Zhang, Yewei Fang, Kaihuo Zhang, Zhiyuan Liu, and Maosong Sun. 2024. Ouroboros: Generating longer drafts phrase by phrase for faster speculative decoding. *arXiv preprint arXiv:2402.13720*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023a. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric.P Xing, Hao Zhang, JosephE. Gonzalez, and Ion Stoica. 2023b. Judging llm-as-a-judge with mt-bench and chatbot arena.
- Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. 2024. Distserve: Disaggregating prefill and decoding for goodput-optimized large language model serving.

Contents

A	Procedures of DOUBLE	13
B	Detailed Theoretical Analysis	13
B.1	Preliminaries and Notations	13
B.2	Proof of Theorem 2 (Multi-Round)	13
B.3	Proof of Theorem 3 (Speedup Ceiling)	14
C	Evaluation Details	14
C.1	Experimental Setting	14
D	More Experimental Results and Discussions	15
D.1	Comparisons with Well-Engineered Baselines	15
D.2	Disentanglement of Performance Gains	16
D.3	Scalability to High Batch Sizes	16
D.4	Proof of Lossless Acceleration	16
D.5	Hierarchical Datastore	17
D.6	Time Consumption	19
D.7	Precision-Efficiency Dilemma	20
D.8	Analysis of Memory-Constrained Scenarios	20
E	Detailed Related Works	21

A Procedures of DOUBLE

The detailed procedures of DOUBLE are described in Algorithm 1 and Algorithm 2.

B Detailed Theoretical Analysis

In this section, we provide the formal derivation and proofs for the theorems presented in Section 3.1.

B.1 Preliminaries and Notations

To be consistent with the main text, let $\alpha \in [0, 1]$ denote the token acceptance rate and γ the draft length. Let T_p and T_q represent the single-token inference latency for the target and draft models, respectively. The speed ratio is defined as $C = T_p/T_q$.

Following standard modeling (Chen et al., 2023; Liu et al., 2024b), we observe the accepted token distribution is bimodal (Shen et al., 2025a) and model the decoding process as a sequence of $k - 1$ fully accepted rounds followed by a single rejected round in the k -th iteration. Here, k represents the

Algorithm 1 Double Retrieval Speculative Parallelism (DOUBLE) - Part I.

Input: Draft model M_q , Target model M_p , Datastore \mathcal{D} , Prefix \mathbf{x} , Gamma γ .

- 1: \triangleright Initialization
- 2: $\text{mode} \leftarrow \text{“pre-verify”}$, $\text{prev_tokens} \leftarrow \gamma$
- 3: **while** not EndOfSequence **do**
- 4: $L \leftarrow |\mathbf{x}|$
- 5: \triangleright Parallel Generation
- 6: $X_q, n_q \leftarrow \text{RETRIEVAL}(M_q, \mathbf{x}, \mathcal{D})$
- 7: $X_p, n_p \leftarrow \text{RETRIEVAL}(M_p, \mathbf{x}, \mathcal{D})$
- 8: Gather probabilities:
- 9: $\text{Prob}_q[L - \text{prev_tokens} - 1 : L]$,
- 10: $\text{Prob}_p[L - \text{prev_tokens} - 1 : L]$.
- 11: \triangleright Pre-verify Mode
- 12: **if** $\text{mode} = \text{“pre-verify”}$ **then**
- 13: Find first reject position n in $X_q[-n_q :]$ using $\text{Prob}_p, \text{Prob}_q$
- 14: **if** $n = n_p$ (all accepted) **then**
- 15: $\checkmark \mathbf{x} \leftarrow \mathbf{x} + X_q[-n_q :]$, $\text{mode} \leftarrow \text{“post-verify”}$
- 16: $\text{prev_tokens} \leftarrow n_q$
- 17: **else**
- 18: $\times \mathbf{x} \leftarrow \mathbf{x} + X_p[: n]$, **rollback** M_q
- 19: $\text{prev_tokens} \leftarrow n_p$
- 20: **end if**
- 21: **end if**
- 22: Continued in Algorithm 2
- 23: **end while**

expected number of rounds until a rejection occurs:

$$k = \mathbb{E}[\text{rounds}] = \frac{1}{1 - \alpha^\gamma} \quad (8)$$

B.2 Proof of Theorem 2 (Multi-Round)

Theorem 2. *The expected number of tokens generated over k rounds is,*

$$\mathbb{E}[L_k] = (k - 1)(\gamma + 1) + \frac{1 - \alpha^{\gamma+1}}{1 - \alpha}. \quad (9)$$

Proof . In standard SD, verification is sequential. The total expected length $\mathbb{E}[L_k]$ consists of $k - 1$ full rounds (each contributing $\gamma + 1$ tokens: γ accepted draft tokens + 1 verification token) and one final truncated round. The expected length of the final truncated round $\mathbb{E}[L_{\text{last}}]$ follows a geometric sum based on the acceptance rate α :

$$\begin{aligned} \mathbb{E}[L_k] &= (k - 1)(\gamma + 1) + \mathbb{E}[L_{\text{last}}] \\ &= (k - 1)(\gamma + 1) + \sum_{i=0}^{\gamma} \alpha^i \\ &= (k - 1)(\gamma + 1) + \frac{1 - \alpha^{\gamma+1}}{1 - \alpha} \end{aligned} \quad (10)$$

Assuming verification time dominates (Under compute-bound, $T_{AR} \approx T_p$), the speedup S_{SD} is:

$$\begin{aligned} S_{SD} &= \frac{\mathbb{E}[L_k] \cdot T_{AR}}{k(\gamma T_q + T_p)} \\ &= \frac{\mathbb{E}[L_k] \cdot C}{k(\gamma + C)} \end{aligned} \quad (11)$$

B.3 Proof of Theorem 3 (Speedup Ceiling)

Theorem 3. *The theoretical speedup S_{PSD} strictly dominates S_{SD} but is upper bounded by C : $S_{SD} \leq S_{PSD} \leq C$.*

1. Formulation of PSD Speedup Parallel Speculative Decoding (PSD) pipelines the draft and verification phases. The expected accepted length $\mathbb{E}[L_p]$ in PSD differs from SD because the verification token of the previous round is not "free" for drafting. Based on the relationship derived in Liu et al. (2024b):

$$\mathbb{E}[L_p] = \mathbb{E}[L_k] - (k - 1) \quad (12)$$

Assume optimal pipelining where the cost per round is dominated by $\max(\gamma T_q, T_p)$ (typically $\gamma T_q \approx T_p$, i.e., $\gamma \approx C$), the speedup S_{PSD} is:

$$\begin{aligned} S_{PSD} &= \frac{\mathbb{E}[L_p] \cdot C}{k\gamma + C} \\ &= \frac{(\mathbb{E}[L_k] - k + 1) \cdot C}{k\gamma + C} \end{aligned} \quad (13)$$

2. Proof of $S_{PSD} \geq S_{SD}$ Let $N_{SD} = \mathbb{E}[L_k] \cdot C$ and $D_{SD} = k(\gamma + C)$ be the numerator and denominator of S_{SD} . Define a shift term $\Delta = (k - 1)C$. Since $k \geq 1, C > 0$, we have $\Delta \geq 0$. Expressing S_{PSD} in terms of S_{SD} components:

$$N_{PSD} = (\mathbb{E}[L_k] - k + 1)C = N_{SD} - \Delta \quad (14)$$

$$D_{PSD} = k\gamma + C = D_{SD} - \Delta \quad (15)$$

We analyze the function $f(x) = \frac{N_{SD} - x}{D_{SD} - x}$. Its derivative is $f'(x) = \frac{N_{SD} - D_{SD}}{(D_{SD} - x)^2}$. Given a valid speedup $S_{SD} \geq 1 \implies N_{SD} \geq D_{SD}$, thus $f'(x) \geq 0$. Since $\Delta \geq 0$, it follows that $f(\Delta) \geq f(0)$, implying:

$$S_{PSD} = \frac{N_{SD} - \Delta}{D_{SD} - \Delta} \geq \frac{N_{SD}}{D_{SD}} = S_{SD} \quad (16)$$

3. Proof of Upper Bound C We examine the upper bound of S_{PSD} from Eq. (13). In the ideal scenario (perfect acceptance, $\alpha = 1$), the maximum length generated per round is limited to γ . Thus, $\mathbb{E}[L_p] \leq k\gamma$. Substituting this:

$$\begin{aligned} S_{PSD} &\leq \frac{k\gamma \cdot C}{k\gamma + C} \\ &= C \cdot \left(\frac{k\gamma}{k\gamma + C} \right) \end{aligned} \quad (17)$$

Algorithm 2 Double Retrieval Speculative Parallelism (DOUBLE) - Part II.

```

1: ▷ Post-verify Mode
2: if mode = "post-verify" then
3:   Find first reject  $n$  in  $X_q[0 : prev\_tokens - 1]$  using  $\text{Prob}_p, \text{Prob}_q$ 
4:   if  $n = prev\_tokens - 1$  then
5:     ▷ Previous verified, check new tokens
6:     if  $n_p \geq n_q$  then
7:       ✓  $\mathbf{x} \leftarrow \mathbf{x} + X_p, \text{mode} \leftarrow$  "pre-verify"
8:     else
9:       Find reject  $n$  in  $X_q[prev\_tokens - 1 : prev\_tokens + n_p - 1]$ 
10:      if all verified then
11:        ✓  $\mathbf{x} \leftarrow \mathbf{x} + X_q[-n_q : ], prev\_tokens \leftarrow n_q$ 
12:      else
13:        ✗  $\mathbf{x} \leftarrow \mathbf{x} + X_p, \text{mode} \leftarrow$  "pre-verify"
14:      end if
15:    end if
16:  else
17:    ✗ Reject at  $n$ : sample  $t \sim \max(\text{Prob}_p[n] - \text{Prob}_q[n], 0)$ 
18:     $\mathbf{x} \leftarrow \mathbf{x}[: L - prev\_tokens + n + 1] + t$ 
19:    mode  $\leftarrow$  "pre-verify", rollback both models
20:  end if
21: end if
22:  $\mathcal{D} \leftarrow \text{Update}(\mathcal{D}, \mathbf{x})$ 

```

Since $k\gamma > 0$ and $C > 0$, the term $\frac{k\gamma}{k\gamma + C}$ is strictly less than 1. Therefore, $S_{PSD} < C$. Combining paragraphs 2 and 3, we conclude:

$$S_{SD} \leq S_{PSD} \leq C. \quad (18)$$

C Evaluation Details

For reproducibility, we discuss the experimental setup (Section 5) in detail and the source code of this project will be made available at a later time.

C.1 Experimental Setting

Tasks and Datasets We evaluate DOUBLE across a diverse suite of LLM configurations, ranging from lightweight to large-scale architectures: LLaMA-2 (7B/70B) (Grattafiori et al., 2024), LLaMA-3 (8B/70B) (Dubey et al., 2024), Deepseek-Coder (1.3B/33B) (Guo et al., 2024), and Qwen3 (0.6B/14B/32B) (Yang et al., 2025a). To

ensure comprehensive coverage, our benchmarks span code generation, mathematical reasoning, summarization, and general instruction following: HumanEval (Chen et al., 2021), GSM8K (Cobbe et al., 2021), CNN/DM (Nallapati et al., 2016), Alpaca (Taori et al., 2023), and MT-Bench (Zheng et al., 2023a) following the set of EAGLE-3.

Baselines To validate the effectiveness of DOUBLE, we compare it against a comprehensive set of representative methods categorized by their acceleration mechanism:

- **Standard SD** (Chen et al., 2023): The canonical *draft-then-verify* framework. We utilize the same draft models as DOUBLE but execute them sequentially without retrieval augmentation.
- **Target-side Retrieval**: We include methods that accelerate inference solely using the target model. **Lookahead** (Fu et al., 2024) employs Jacobi iteration for multi-branch generation without a draft model. **PLD** (Saxena, 2023) matches the current prefix against the prompt to reuse computations. **Token Recycling** (Luo et al., 2024) leverages past key-value pairs to predict future tokens.
- **Draft-side Retrieval**: We compare against **Ouroboros** (Zhao et al., 2024), which constructs an n -gram index from the draft model’s generation history to propose candidate suffixes, aiming to extend the draft length.
- **Parallel SD**: We use **PEARL** (Liu et al., 2024b) as the primary parallel baseline. PEARL pipelines the drafting and verification phases to hide draft latency but relies on standard autoregressive drafting without retrieval guidance.
- **Training-based Methods**: We include **EAGLE-3** (Li et al., 2025), the current state-of-the-art method that utilizes a lightweight training layer to incorporate feature-level history for more accurate drafting.

All baselines are reproduced from their official codebases using the optimal configurations. Experiments are conducted under identical conditions to avoid implementation bias.

Implementation and Metrics All experiments are conducted on 8 NVIDIA A100 (80GB) GPUs. Model parallelism is applied where necessary: models under 33B parameters run on a single device, and 70B models are distributed across TWO devices. We use greedy sampling with a batch size of 1 for the main experiments. The draft length is

dynamically set to $\gamma = \lceil C \rceil$, where C represents the average speed ratio T_p/T_q across datasets in off-the-shelf framework. For **EAGLE-3**, it utilizes a lightweight training layer to reduce the draft model cost and the C is different. We set the default depth $d = 10$, prior round $K = 10$, and n-gram $n = 3$.

We report three key metrics: **Wall-Time Speedup**, and **Mean Accepted Length** (M). In the context of parallel frameworks (PEARL and DOUBLE), M denotes the *continuously accepted length* (Liu et al., 2024b) achieved through multiple rounds of pipelined generation, rather than the single-round acceptance of vanilla SD. Additionally, we analyze the Average Matched Tokens (AMT) defined as the expectation of s , which directly determines the effective speedup, serving as a proxy for the precision-efficiency trade-off.

Models	Layers	dim	FFN dim	Vocabulary size
Deepseek 1.3B	24	2048	5504	32256
Deepseek 33B	62	7168	19200	32256
LLaMA-2 7B	32	4096	11008	32000
LLaMA-2 70B	80	8192	28672	32000
LLaMA-3.1 8B	32	4096	14336	128256
LLaMA-3.3 70B	80	8192	28672	128256
Qwen-3 0.6B	28	1024	3072	151936
Qwen-3 14B	40	5120	17408	151936
Qwen-3 32B	64	5120	25600	151936

Table 4: Model configurations.

D More Experimental Results and Discussions

D.1 Comparisons with Well-Engineered Baselines

To further validate the superiority of DOUBLE, we expand our evaluations to include a broader set of well-engineered speculative decoding approaches from the same period, including Medusa (Cai et al., 2024), SpecInfer (Chen et al., 2023), SEQUOIA (Chen et al., 2024b), REST (He et al., 2024), and Swift (Xia et al., 2024a).

We conducted comparative experiments on the LLaMA-2-70B target model (with LLaMA-2-7B as the drafter). As demonstrated in Table 5, DOUBLE significantly outperforms these well-engineered baselines across all evaluated benchmarks. For instance, it achieves a $3.42\times$ speedup compared to Medusa’s $2.51\times$ on HumanEval, and $3.05\times$ versus SEQUOIA’s $2.25\times$ on CNN/DM. This substantial advantage proves DOUBLE’s superior ability to decouple draft length from autoregressive latency, effectively breaking the speedup ceiling that constrains prior tree-based and retrieval-based meth-

ods.

Table 5: **Comparisons with well-engineered approaches on LLaMA2-70B.** All baselines were run under their official optimal settings. DOUBLE consistently achieves the highest acceleration.

Methods	Medusa	SpecInfer	SEQUOIA	REST	Swift	DOUBLE
HumanEval	2.51×	1.92×	2.37×	1.78×	1.56×	3.42×
GSM8k	2.23×	1.73×	2.03×	1.46×	1.43×	2.83×
CNN/DM	1.73×	1.69×	2.25×	1.47×	1.45×	3.05×

D.2 Disentanglement of Performance Gains

To clarify the exact sources of our massive end-to-end acceleration, we disentangle the performance gains by mapping them to our ablation studies (Section 5):

- **Draft-side Retrieval (Breaking Length Barrier):** As shown in the “*w/o Draft-Retrieval*” ablation, without draft retrieval, the system reverts to standard autoregressive drafting. This strictly caps the speedup to the theoretical limit (dropping from 4.41× to 2.23× on LLaMA-3.3-70B). Draft retrieval is the core driver for assembling exceptionally long candidate spans.
- **Target-side Guidance (Correcting Errors):** Shown in the “*w/o Target-Retrieval*” ablation, without authoritative target guidance, mid-sequence rejections cause severe pipeline stalls. Target-side retrieval turns potential rollbacks into valid extensions, preserving the pipeline efficiency.
- **Algorithmic Synergy vs. System Overhead:** Our profiling confirms that the Python-based transformers library actually incurs pipeline overheads (e.g., KV-cache updates). The fact that DOUBLE achieves state-of-the-art wall-clock speedups *despite* these engineering overheads proves that the acceleration is entirely driven by the profound algorithmic synergy of dual retrieval, rather than low-level engineering tricks.

D.3 Scalability to High Batch Sizes

Our primary evaluation focuses on latency-critical scenarios ($bs = 1$), which represent the standard use case for real-time interaction. However, in high-throughput settings with larger batch sizes, the GPU’s compute-bound nature may saturate, and the overhead of maintaining diverse dynamic tree structures for each request can become non-trivial.

To assess scalability, we leverage NANO-PEARL (Liu et al., 2024b), a novel Parallel SD framework built upon Nano-vLLM. It features a

high-concurrency engine optimized with custom kernels for industrial acceleration. We benchmark NANO-PEARL against EAGLE-3 on Qwen3-32B across batch sizes $bs \in \{1, \dots, 16\}$. As detailed in Table 6, EAGLE-3 suffers significant degradation at $bs = 16$ (dropping to 1.13×). This decline is primarily attributed to the memory access overhead and compute-bound caused by the complex, non-contiguous tree attention masks.

In contrast, NANO-PEARL unlocks high-throughput potential through chained drafting and parallel decoding, sustaining a robust speedup of 1.73× even at $bs = 16$. Similarly, DOUBLE is inherently suited for this architecture; its linear retrieval structure eliminates the overhead of dynamic masking. By integrating with NANO-PEARL and dynamically tuning retrieval depth, DOUBLE can further enhance speedups in high-throughput environments. We are currently integrating DOUBLE into NANO-PEARL and plan to release it as a feature branch to the community, contributing to robust industrial deployment. Future work will also explore migrating this parallel SD paradigm to comprehensive inference engines such as vLLM and SGLang.

Table 6: High-Batch performance comparison on Qwen3-32B. While EAGLE-3 degrades significantly due to mask overheads, NANO-PEARL maintains robust acceleration under high concurrency ($bs = 16$).

	Batch Size (HumanEval)				
	1	2	4	8	16
EAGLE-3	1.98×	1.83×	1.67×	1.34×	1.13×
NANO-PEARL	2.64×	2.33×	2.17×	1.94×	1.73×

D.4 Proof of Lossless Acceleration

In this section, we provide both theoretical and empirical verification to demonstrate the lossless nature of DOUBLE. **We show that our Target-Guided Verification mechanism strictly preserves the target model’s output distribution $p(x)$ under both greedy ($T = 0$) and stochastic ($T > 0$) decoding regimes.**

D.4.1 Theoretical Guarantee

The fundamental premise of Speculative Decoding (SD) and Retrieval-based SD is universally acknowledged as **lossless**: the final output distribution must rigorously match the target distribution $p(x)$, independent of the draft source. Our framework strictly adheres to this standard verification protocol (Leviathan et al., 2023; Chen et al., 2023).

We analyze the validity of our mechanism under two distinct settings:

Greedy Decoding ($T = 0$) In the greedy setting, the target distribution degenerates into a Dirac delta function: $p(x) = \delta(x - \arg \max p(\cdot))$. We provide the formal proof for the *Correction* and *Extension* phases:

- **Correction Mechanism** (y_i): Standard rejection sampling requires drawing a correction from the residual distribution when a draft token x_i diverges from the target. In greedy decoding, the target distribution $p(x)$ is a Dirac delta function centered at the top-1 prediction y_i . This forces the residual sampling to deterministically select y_i . Therefore, replacing the rejected x_i with y_i is mathematically equivalent to the formal rejection sampling process.
- **Validity of Extension** ($y_{i+1} \dots y_{s+1}$): We explicitly address the concern regarding whether appending tokens from Y_p strictly preserves the target distribution. It is crucial to clarify that Y_p is **not** a heuristic sequence retrieved directly from the datastore. As defined in Sec. 2.2 (Eq. 2), Y_p represents the **verified output** of the target model \mathcal{M}_p computed during the parallel verification step.

Specifically, \mathcal{M}_p employs a causal mask to process the retrieved candidates in parallel. This means \mathcal{M}_p calculates the conditional probability $p(y_k \mid \text{prefix}, y_1, \dots, y_{k-1})$ for all positions k simultaneously within the same forward pass. Consequently, if y_i is selected as the correction, the subsequent token y_{i+1} has already been computed by \mathcal{M}_p conditioned specifically on the path including y_i . Thus, the sequence $\{y_{i+1}, \dots, y_{s+1}\}$ constitutes the ground-truth greedy path of the target model. Utilizing these pre-calculated tokens is not an approximation but a direct application of the target model’s own generation, ensuring strict losslessness without redundant re-computation.

According to the standard rejection sampling criterion (Leviathan et al., 2023), if a draft token x_i diverges from the target (i.e., $x_i \neq \arg \max p(\cdot)$), a correction must be sampled from the residual distribution,

$$x'_i \sim \text{norm}(\max(0, p(x) - q(x))) \quad (19)$$

Under greedy sampling, since $p(x)$ has a probability mass of 1 ($q(x) = 0$) at the optimal token y_i ,

the sampled correction x'_i collapses deterministically to y_i . Consequently, replacing the rejected draft token x_i with the target’s retrieved token y_i is mathematically equivalent to performing rejection sampling. Furthermore, because the corrected prefix aligns with the target model’s deterministic path, the subsequent tokens in the retrieved chain (y_{i+1}, \dots) remain valid. This allows DOUBLE to utilize **Forward Guidance** (Extension) losslessly in greedy scenarios.

Stochastic Sampling ($T > 0$) Under non-greedy settings, the output distribution is stochastic. To strictly preserve $p(x)$, we adhere to the verified speculative sampling protocol. When a draft token x_i is rejected, we sample a correction x'_i from the residual distribution $p(x) - q(x)$ rather than deterministically selecting the top-1 token. Crucially, unlike the greedy case, the sampled correction x'_i may differ from the pre-retrieved target token y_i . To maintain distributional consistency, we strictly **discard the subsequent Target Guidance** (y_{i+1}, \dots) upon rejection, as these tokens were generated conditionally on a specific path that is no longer valid. However, the **Multi-token Pre-verify** mechanism remains active for the prefix preceding the rejection point. This ensures that the system benefits from early filtering of divergent tokens while strictly adhering to the target distribution.

D.4.2 Empirical Verification

To empirically validate that DOUBLE introduces no deviation from the target model’s output, we conducted exact-match accuracy tests on the GSM8K dataset across three temperature settings ($T \in \{0, 0.5, 1.0\}$).

As presented in Table 7 and 8, DOUBLE achieves accuracy scores identical to the Vanilla autoregressive baseline (within numerical hardware precision limits for BFLOAT16) across all temperatures. For instance, on LLaMA-3.3-70B across GSM8K, the accuracy remains invariant at 0.93 ($T = 0$), 0.92 ($T = 0.5$), and 0.89 ($T = 1.0$) for DOUBLE. These results confirm that DOUBLE delivers substantial speedups $3.67\times$ without compromising generation quality or altering the theoretical output distribution.

D.5 Hierarchical Datastore

To optimize the trade-off between retrieval coverage and memory efficiency, while effectively leveraging verification signals, we implement a hierar-

Table 7: **Verification of Lossless Acceleration.** We compare the Exact Match (EM) accuracy and Speedup of DOUBLE against the Vanilla baseline on GSM8K. The results confirm that DOUBLE maintains output consistency with the target model across varying temperatures. (The minor drops stem from the inherent uncertainty of LLMs and the precision limits of hardware numerics.)

Models	Methods	Temperature=0		Temperature=0.5		Temperature=1.0	
		Acc.	Speedup	Acc.	Speedup	Acc.	Speedup
Qwen-32B	Vanilla	0.94	1.00×	0.92	1.00×	0.89	1.00×
	Sps	0.94	1.00×	0.91	1.00×	0.88	1.00×
	DOUBLE	0.94	2.30×	0.91	2.12×	0.89	2.03×
LLaMA-3.3-70B	Vanilla	0.93	1.00×	0.92	1.00×	0.90	1.00×
	Sps	0.94	1.00×	0.91	1.00×	0.88	1.00×
	DOUBLE	0.93	4.41×	0.92	4.23×	0.89	3.89×

Table 8: **More results of Lossless Acceleration.** We compare the Exact Match (EM) accuracy and Speedup of DOUBLE against the Vanilla baseline on Deepseek1.3&33B across three benchmarks.

Benchmarks	Methods	Temperature=0		Temperature=0.5		Temperature=1.0	
		Acc.	Speedup	Acc.	Speedup	Acc.	Speedup
HumanEval	Vanilla	0.65	1.00×	0.62	1.00×	0.59	1.00×
	Sps	0.65	2.15×	0.63	1.95×	0.58	1.81×
	DOUBLE	0.65	4.45×	0.63	4.22×	0.59	4.13×
GSM8K	Vanilla	0.41	1.00×	0.39	1.00×	0.35	1.00×
	Sps	0.41	1.82×	0.39	1.70×	0.35	1.64×
	DOUBLE	0.41	3.83×	0.39	3.62×	0.35	3.43×
CNN/DM	Vanilla	0.37	1.00×	0.32	1.00×	0.30	1.00×
	Sps	0.37	2.05×	0.32	1.90×	0.31	1.78×
	DOUBLE	0.37	4.05×	0.32	3.93×	0.31	3.85×

chical multi-layered datastore. The detailed procedure is formalized in Algorithm 3.

Data Decontamination and Fairness We explicitly address potential concerns regarding data sourcing and benchmark contamination (test-set leakage). We justify our datastore setup based on two rigorous decontamination protocols:

1) **Static Prior Knowledge:** Our static prior datastore is initialized via inference on a small subset of ShareGPT. We specifically chose ShareGPT because it is strictly disjoint from our specialized evaluation benchmarks (e.g., HumanEval, GSM8K), structurally eliminating the risk of test-set memorization. Furthermore, this guarantees absolute fairness: ShareGPT is the exact corpus used to train the draft models for EAGLE-3 (Li et al., 2025) (our strongest baseline). DOUBLE achieves higher speedups not by leveraging an unfair data advantage, but by utilizing the exact same data source more efficiently via non-parametric retrieval.

2) **Ephemeral Dynamic Cache:** Beyond the static prior, DOUBLE leverages an ephemeral, session-scoped dynamic cache. To strictly prevent

Algorithm 3 Retrieval Cache Mechanism

Input: Input tokens \mathbf{x} , Max n-gram size N , Prediction length K

```

1: ▷ Cache Structure
2:  $\mathcal{C}_{prefix}$  ▷ Persistent prefix cache (warmup data)
3:  $\mathcal{C}_{dynamic}$  ▷ Dynamic cache (current task tokens)
4:  $\mathcal{C}_{rejected}$  ▷ Rejected tokens cache
5: ▷ N-gram Retrieval (Priority Search)
6: for  $n = \min(N, |\mathbf{x}|)$  down to 1 do
7:    $ngram \leftarrow \mathbf{x}[-n : ]$  ▷ Extract suffix
8:   ▷ Search in prefix cache first
9:   if  $ngram \in \mathcal{C}_{prefix}$  then
10:     ✓ Return next  $K$  tokens from  $\mathcal{C}_{prefix}$ 
11:   end if
12:   ▷ Then search in dynamic cache
13:   if  $ngram \in \mathcal{C}_{dynamic}$  then
14:     ✓ Return next  $K$  tokens from  $\mathcal{C}_{dynamic}$ 
15:   end if
16:   ▷ Finally search in rejected cache
17:   if  $ngram \in \mathcal{C}_{rejected}$  then
18:     ✓ Return next  $K$  tokens from  $\mathcal{C}_{rejected}$ 
19:   end if
20: end for
21: ▷ Fallback: search in input sequence itself
22: return Original PLD search in  $\mathbf{x}$ 
23: ▷ Cache Update
24:  $\mathcal{C}_{dynamic} \leftarrow \mathcal{C}_{dynamic} \cup \{\text{accepted tokens}\}$ 
25:  $\mathcal{C}_{rejected} \leftarrow \mathcal{C}_{rejected} \cup \{\text{rejected tokens}\}$ 

```

contamination, the dynamic cache is completely flushed and reset at the beginning of every new inference request. This strict isolation guarantees that the context or generated tokens from one test instance can never pollute the static prior datastore or leak into subsequent evaluations.

Clarification on Memory Overhead Concerns regarding retrieval-based methods often stem from the presumption of unbounded datastore growth. We emphasize that DOUBLE employs a strictly **Lightweight & Ephemeral** storage protocol designed with minimal resource footprint:

- **Static Prior (Initialization):** As detailed in the methodology, the prior knowledge is serialized as a compact local file (approx. 9.5 MB). It is loaded once into CPU RAM as a read-only trie structure, incurring negligible system overhead.
- **Ephemeral Dynamic Datastore:** distinct from RAG frameworks that maintain persistent, high-




Figure 7: Profiling results showing that DOUBLE incurs minimal overhead: retrieval (1.9%) and communication (2.5%) remain negligible compared to model forward (86.2%) in single **Retrieval Forward**.

dimensional vector databases, our dynamic datastore records only integer Token IDs rather than hidden states. Crucially, the dynamic layers ($\mathcal{C}_{\text{dynamic}}$ and $\mathcal{C}_{\text{rejected}}$) are session-scoped: they are initialized at the start of a generation request and immediately *flushed* upon completion. Even during long-context generation (e.g., 8k tokens), this integer-based mapping consumes only a few megabytes of CPU RAM, effectively circumventing VRAM scarcity issues on consumer-grade hardware.

Table 9: **Ablation Study on Rejected Token Cache.** Performance comparison on Qwen3-32B (MT-Bench). Incorporating the Rejected Cache significantly improves the retrieval hit rate, leading to superior speedups.

Configuration	Writing	Roleplay	Reasoning	Math	Coding	Extraction	Avg.
DOUBLE (w/o Rejected Cache)	2.15×	2.10×	2.38×	2.05×	2.18×	1.95×	2.14×
DOUBLE (Full)	2.37×	2.23×	2.55×	2.21×	2.33×	2.03×	2.30×

Analysis of Rejected Tokens A legitimate concern is that caching rejected tokens might pollute the retrieval pool with suboptimal candidates. However, within the context of speculative decoding, a rejection primarily signifies a divergence from the target model’s current output rather than a semantic error. These tokens often represent valid synonymic variations or near misses. By caching these rejected yet high-probability tokens, DOUBLE preserves valid alternative generation paths. This allows the model to recall and reuse them when similar contexts recur, effectively transforming ‘waste’ into valuable candidates and boosting the retrieval hit rate.

We empirically validate this design in Table 9. The inclusion of the Rejected Token

Cache ($\mathcal{C}_{\text{rejected}}$) yields consistent performance gains across all benchmarks (Avg. $2.30\times$ vs. $2.14\times$). This confirms that utilizing the rejected history provides valuable alternative guidance that augments the retrieval process.

D.6 Time Consumption

To assess the runtime overhead of managing dynamic retrieval datastores, we profiled the latency breakdown of a single **Retrieval Forward** step on the Qwen3 pair (0.6B Draft / 14B Target). As detailed in Table 10 and Figure 7, we provide a granular quantification of the system overhead:

Table 10: **Latency Profiling of DOUBLE in a single Retrieval Forward.** The evaluation is conducted on the Qwen3 pair (0.6B Draft / 14B Target). The CPU-based retrieval overhead is strictly negligible compared to the GPU-bound model forward pass.

Component	Model Forward	Verify	Comm.	Preprocess	Update	Retrieval (CPU)
Draft (0.6B)	86.2%	5.8%	2.5%	2.1%	1.5%	1.9%
Target (14B)	93.9%	2.4%	1.0%	1.3%	0.6%	0.8%

Negligible Retrieval Cost. The inference process is overwhelmingly dominated by the GPU-bound **Model Forward** phase (86.2% for the draft model and 93.9% for the target model). Crucially, the CPU-based **Retrieval** mechanism introduces a strictly negligible overhead (only 1.9% and 0.8%, respectively). This high efficiency stems from our architectural design choice. Unlike dense vector databases that require heavy sliding window operations (which are naturally unfriendly to CPUs), our dynamic datastore utilizes a lightweight, local sparse n -gram index. Querying and updating this structure involves simple $\mathcal{O}(1)$ hash map operations, which are extremely fast on modern CPUs compared to the massive matrix multiplications executed by the LLM on the GPU, ensuring that the retrieval-augmented drafting process never becomes a bottleneck.

System Overhead. Auxiliary operations (such as inter-process Communication and KV-cache Verification) constitute a small but noticeable portion of the latency. This is largely attributed to the architectural constraints of the Python-based transformers library. We anticipate that migrating to high-performance inference engines (e.g., SGLang) will further minimize these non-model system overheads, pushing the actual wall-clock speedup even closer to the theoretical ceiling.




Figure 8: **Precision-Efficiency Analysis.** The scatter plot compares DOUBLE against single-sided ablations. DOUBLE (Teal) resolves the dilemma by achieving both high Precision (AMT) and Efficiency (Speedup). Note that architectural overheads in transformers (e.g., KV rollback) create a slight gap between AMT gains and wall-time speedup, which we aim to address in future high-performance implementations.

D.7 Precision-Efficiency Dilemma

To illustrate how DOUBLE navigates the trade-off between retrieval quality and inference speed, we visualize the relationship between Average Matched Tokens (AMT), a proxy for Precision and Wall-time Speedup (Efficiency) in Figure 8. The scatter plot, derived from LLaMA-3.3-70B inference traces, delineates three distinct performance regimes:

- **Low Efficiency Regime (w/o Draft-Retrieval):** Marked in orange, this variant depends on autoregressive drafting. While it maintains respectable precision, its efficiency is strictly constrained by the theoretical speed ceiling C , clustering in the lower-speed region.
- **Low Precision Regime (w/o Target-Retrieval):** Marked in blue, this variant utilizes draft-side retrieval but lacks authoritative target guidance. Constrained by the limited capacity of the lightweight draft model, it suffers from frequent verification failures, resulting in low AMT and suboptimal speedups.
- **Optimal Regime (DOUBLE):** Marked in teal, our method occupies the top-right quadrant, representing the Pareto-optimal frontier. By synergizing the high throughput of draft retrieval with the precise guidance of the target model, DOUBLE simultaneously maximizes AMT (> 0.9) and Speedup ($> 4.0\times$), effectively resolving the

dilemma.

Engineering Gap Analysis Notably, we observe a divergence between the high AMT values and the realized wall-time speedups across all retrieval-based methods. This discrepancy stems from the architectural limitations of the standard transformers library. Operations such as inter-process communication and, more critically, the costly KV-cache rollback mechanisms upon rejection, incur non-trivial latency overheads. These engineering bottlenecks partially offset the throughput gains derived from high retrieval precision. To bridge this gap, future work will focus on migrating DOUBLE to high-performance inference engines like vLLM or SGLang. By leveraging custom CUDA kernels and optimized memory management, we aim to eliminate these overheads and fully translate algorithmic precision into industrial-grade acceleration.

D.8 Analysis of Memory-Constrained Scenarios

This section evaluates the adaptability of DOUBLE across varying resource environments, distinguishing it from traditional sequential methods. Unlike “draft-then-verify” paradigms that often underutilize hardware, DOUBLE is designed to exploit parallel computation. We categorize deployment scenarios based on memory availability:

Resource-Abundant Scenarios DOUBLE is optimal for multi-GPU cloud environments or cloud-edge collaborative setups where computational resources are ample. In these settings, the draft and target models are deployed on independent processors (e.g., separate GPUs or heterogeneous CPU/GPU clusters). This physical isolation eliminates memory contention, preventing the draft model from slowing down the target model’s primary inference loop. Our main experiments (Section 5.2) reflect this ideal configuration. Future integrations with Tensor Parallelism (TP) (Zhong et al., 2024) could further amplify acceleration in these robust environments.

Resource-Constrained Scenarios (Shared GPUs) In scenarios where the draft and target models must share GPUs (e.g., a 33B model spanning two 40GB A100s, with a 1.3B draft model co-located), resource contention becomes a bottleneck. To mitigate this, we adopt a modified Pipeline Parallelism (PP) strategy inspired by PEARL (Liu et al., 2024b).

By interleaving computation, the draft model generates tokens on the idle GPU while the target model executes on the active GPU. For instance, while the target model computes on GPU 0, the draft model operates in parallel on GPU 1; they efficiently swap roles as the target model’s pipeline progresses. Although this introduces minor communication overhead, it effectively circumvents memory contention. As shown in Table 11, experiments with Deepseek 1.3B & 33B on Spec-Bench demonstrate that DOUBLE with PP retains **89.76%** of its peak performance, significantly outperforming vanilla SD (SpS) even under constrained settings.

Table 11: Performance retention of DOUBLE under memory-constrained scenarios (Deepseek 1.3B & 33B on MT-Bench). Using Pipeline Parallelism (PP), DOUBLE maintains $\sim 90\%$ of its ideal acceleration.

Methods	Writing	Roleplay	Reasoning	Math	Coding	Extraction	Avg.
SpS (Vanilla SD)	1.89×	2.08×	2.14×	1.97×	2.23×	1.75×	2.01×
DOUBLE (Ideal)	3.86×	4.25×	4.38×	4.01×	4.56×	3.54×	4.10×
DOUBLE (w/ PP)	3.48×	3.74×	3.94×	3.69×	3.96×	3.19×	3.67×
Retention	90.2%	88.0%	90.0%	92.0%	86.8%	90.1%	89.5%

Single-GPU Scenarios (Extreme Constraint)

In single-GPU environments, DOUBLE gracefully degrades to a serialized workflow, executing retrieval-augmented drafting followed by target verification. Table 12 compares this mode against PEARL (which reverts to Vanilla SD) on Deepseek 1.3B & 33B. Even without parallel execution, DOUBLE’s retrieval mechanism ensures superior drafting precision, outperforming the baseline by a clear margin (1.81× vs. 1.64×). Furthermore, advanced frameworks like Nano-vLLM offer a pathway to restore flexibility: by leveraging Gloo for intra-device process orchestration and utilizing concurrent CUDA streams, DOUBLE can achieve logical parallelism for verification even on a single device, maximizing hardware utilization.

Table 12: Performance in Single-GPU scenarios (Deepseek 1.3B & 33B on MT-Bench). DOUBLE outperforms PEARL (Vanilla SD) even without parallel execution capabilities.

Methods	Writing	Roleplay	Reasoning	Math	Coding	Extraction	Avg.
PEARL (SpS)	1.89×	2.08×	2.14×	1.97×	2.23×	1.75×	2.01×
DOUBLE	2.37×	2.23×	2.55×	2.21×	2.33×	2.03×	2.29×

Heterogeneous Deployment and Future Optimizations Beyond homogeneous GPU setups, DOUBLE is inherently well-suited for heterogeneous hardware environments (e.g., executing the

draft model on consumer-grade GPUs or robust CPUs, while the target model resides on data-center GPUs). Because our architecture allows the draft model to "run ahead" asynchronously in the parallel pipeline, it naturally accommodates and overlaps the communication latency between heterogeneous devices. For future system-level optimizations, we plan to implement asynchronous CPU-GPU execution overlap, which allows the CPU to manage data-store updates concurrently with the target model’s verification, and explore GPU-resident sparse data-stores to completely eliminate CPU contention and PCIe communication bottlenecks.

E Detailed Related Works

While Large Language Models (LLMs) have achieved remarkable success in various benchmarks (Team et al., 2023; Brown et al., 2020b; Yang et al., 2024; Guo et al., 2025; Michalski et al., 1983), their deployment is severely constrained by their auto-regressive token-by-token generation.

Efficient LLM Architectures. Substantial research has been directed toward accelerating inference through structural model optimizations. Key approaches include: 1) **Model Distillation** (Sreenivas et al., 2024; Muralidharan et al., 2024; Lu and Lab, 2025; Sanh et al., 2019; Hinton et al., 2015), which compresses knowledge from a large teacher model into a compact student model to boost speed while maintaining capability; 2) **Quantization** (Frantar and Alistarh, 2023; Xiao et al., 2023; Lin et al., 2024; Liu et al., 2024c; Ashkboos et al., 2024), which reduces parameter precision to lower storage footprint and minimizes data transfer latency from HBM to on-chip memory; and 3) **Pruning** (Frantar and Alistarh, 2023; Dao et al., 2022; Men et al., 2024; Chen et al., 2024a; Hu et al., 2022; Sun et al., 2024; Zhang et al., 2024b), which eliminates redundant parameters. Specifically, structured pruning is often coupled with distillation to train efficient lightweight models to reduce memory access overhead. Despite these advances in reducing computational complexity, LLM inference remains inherently *memory-bound* due to the sequential nature of auto-regressive decoding. Furthermore, these structural modifications often necessitate extensive retraining or specialized hardware support, and few strategies effectively resolve the memory bandwidth bottleneck without compromising performance.

Speculative Decoding While Speculative Decoding (SD) (Agrawal et al., 2024; Chen et al., 2024b; Huang et al., 2024; Liu et al., 2024a; Xia et al., 2024b,a; Yang et al., 2025b; Huang et al., 2025), has demonstrated significant acceleration to alleviate *memory-bound* with lossless generalization, maximizing the draft token acceptance rate remains a critical challenge. Existing literature primarily addresses this through training-based or training-free alignment strategies. For instance, Medusa (Cai et al., 2024) introduces auxiliary decoding heads, whereas EAGLE (Li et al., 2024) and Glide (Du et al., 2024) leverage target model features to enhance drafting precision. Similarly, SpecInfer (Chen et al., 2023) employs tree-based attention to verify multiple candidates. On the training-free front, methods like Draft&Verify (Zhang et al., 2024a) utilize a subset of target model layers as a proxy for the draft model; however, the limited number of skipped layers constrains the overall speedup. More recently, EAGLE3 (Li et al., 2025) has redefined the scaling laws for small model training via test-time training and feature extrapolation, gaining widespread industrial adoption. Nevertheless, a fundamental limitation persists across these approaches: their adherence to a sequential *draft-then-verify* paradigm imposes an inherent mutual waiting bottleneck.

Parallel Speculative Decoding To mitigate the hardware underutilization in serialized execution, recent works have shifted toward parallel paradigms. PEARL (Liu et al., 2024b) introduces a pipelined framework that enables the draft model to generate subsequent tokens concurrently with the target model’s verification of the initial draft. Building on this, SpecBranch (Shen et al., 2025a) incorporates fine-grained dynamic length adjustments and robust multi-branch fallback strategies to further optimize parallel efficiency. In the distributed landscape, DSI (Timor et al., 2024) proposes a parallel framework that orchestrates the temporal overlap between target and drafter instances. However, the efficacy of these approaches remains constrained by the theoretical speedup limits of specific model pairings and the severe penalties associated with rollback, which inevitably disrupt the pipeline.

Retrieval Speculative Decoding Retrieval-based methods, which generate drafts via n -gram matching, have emerged as a distinct branch of SD. These methods generally bifurcate into two categories: 1) Target-Side Self-SD, where

approaches like Lookahead Decoding (Fu et al., 2024), PLD (Saxena, 2023), REST (He et al., 2024), and LogitSpec (Liu et al., 2025) focus on accelerating the target model intrinsically. 2) Draft-Side Acceleration, where Ouroboros (Zhao et al., 2024) stands as the sole work employing lookahead logic to expedite draft model generation. However, existing retrieval mechanisms operate in isolation on either the target or draft model, creating an inherent *Retrieval Precision-Efficiency Dilemma*. To address these challenges, we propose **DOUBLE**, which resolves this dilemma by orchestrating a synergistic dual-retrieval strategy. By leveraging the complementary strengths of both models simultaneously, DOUBLE breaks the theoretical speedup ceiling imposed by existing parallel frameworks.

More Discussions about EAGLE We acknowledge that EAGLE-3’s dedicated draft model is a significant but orthogonal design choice for PSD. DOUBLE currently utilizes a standard draft model; however, this opens a promising avenue for future work to integrate on-policy online distillation or EAGLE’s draft design, developing specialized draft models to further enhance the parallel framework.

Discussions on Future Work. Although speculative decoding has achieved remarkable success in accelerating pure-text generation, its potential across diverse downstream applications remains to be fully explored. Looking forward, we envision that Parallel Speculative Decoding (PSD) paradigms, exemplified by DOUBLE, can be seamlessly adapted into several promising avenues. **1) multimodal LLM acceleration**, where PSD can be tailored to efficiently serve large vision-language and video models through visual-semantic guidance, KV cache compression, and novel parallel decoding architectures (Zhang et al., 2026a; Kong et al., 2026a,b; Ji et al., 2026; Zeng et al., 2026). **2) complex agentic reasoning**, where accelerated on-the-fly logical deduction, dynamic policy exploration, and multi-perspective verification are critical for agent workflows (Wang et al., 2025; Zhang et al., 2026b, 2025; Wu et al., 2026c,a,b; Wang et al., 2026). **3) broader efficiency optimizations**, including synergies with native parallel reading mechanisms inside Transformers and pipelined multi-DNN execution on heterogeneous hardware, to push the limits of modern inference systems (Wang, 2026; Shen et al., 2025b).