

Accelerating Speculative Decoding with Block Diffusion Draft Trees

Liran Ringel¹ & Yaniv Romano^{1,2}

¹Department of Computer Science

²Department of Electrical and Computer Engineering

Technion – Israel Institute of Technology

Haifa, Israel

liranringel@cs.technion.ac.il

Abstract

Speculative decoding accelerates autoregressive language models by using a lightweight drafter to propose multiple future tokens, which the target model then verifies in parallel. DFlash shows that a block diffusion drafter can generate an entire draft block in a single forward pass and achieve state-of-the-art speculative decoding performance, outperforming strong autoregressive drafters such as EAGLE-3. Vanilla DFlash, however, still verifies only a single drafted trajectory per round, potentially limiting its acceptance length. We introduce DDTree (Diffusion Draft Tree), a method that constructs a draft tree directly from the per-position distributions of a block diffusion drafter. Under a fixed node budget, DDTree uses a simple best-first heap algorithm to select the continuations that are most likely to match the target model according to a surrogate defined by the draft model’s output. The resulting tree is verified efficiently in a single target model forward pass using an ancestor-only attention mask. Because DDTree builds on DFlash, a leading draft model for speculative decoding, these gains place DDTree among the leading approaches to speculative decoding.



Project Page



Code




Figure 1: Speedups relative to autoregressive decoding at temperature 0.0 across datasets and target models. DDTree bars use the best tree-node budget for each dataset-model pair.

1 Introduction

Autoregressive language models generate text one token at a time, so the sampling of each new token requires another forward pass through a large model. This sequential dependence makes decoding a major source of inference latency. Speculative decoding addresses this bottleneck by using a lightweight drafter model to propose several future tokens and a large target model to verify them in parallel, while preserving the target model’s output distribution (Leviathan et al., 2023; Chen et al., 2023). As language models continue to grow in size, reducing decoding latency without changing model outputs has made speculative decoding an increasingly important technique for inference.

The effectiveness of speculative decoding depends on the quality of the draft model. To deliver end-to-end speedups, the drafter must be cheap enough that drafting overhead stays small and accurate enough that the target model frequently accepts the drafted tokens. At a high level, the objective is to maximize the expected number of drafted tokens that the target model accepts, i.e., tokens that agree with the target model and can therefore be committed. At the same time, this should be achieved without adding significant overhead that makes the speedup disappear.

Block diffusion (Arriola et al., 2025) is especially attractive in this setting as it can generate an entire draft block in a single forward pass. DFlash (Chen et al., 2026) shows the promise of this approach: it uses a small block diffusion drafter that leverages features derived from the larger target model, achieving state-of-the-art speculative decoding performance. Indeed, DFlash has been shown to outperform strong autoregressive drafters such as EAGLE-3 (Li et al., 2025). These results establish block diffusion as a powerful foundation for speculative decoding. At the same time, these promising results spark a key challenge:

How should we make the best use of the information that the block diffusion draft model produces?

Currently, DFlash verifies only one drafted trajectory per round, even though a single block diffusion forward pass produces a distribution over tokens at each future position. As such, DFlash does not utilize the many plausible continuations that block diffusion produces. Naturally, exploring more continuations could increase the probability that the target model continues along a drafted path. On the other hand, naively using multiple continuations increases the verifier cost and can erase the latency benefit. The challenge, therefore, is to use the draft model’s per-position distributions to choose the continuations that are most worthwhile to verify.

We address this challenge with DDTree (Diffusion Draft Tree). Our method (i) constructs a draft tree directly from the per-position distributions produced by a block diffusion drafter, (ii) selects a compact set of promising continuations under a specified tree-node budget, and (iii) verifies them in a single target model forward pass with tree attention. In this way, DDTree preserves DFlash’s low drafting latency while allowing each round to explore multiple continuations instead of only one. Our experiments show that DDTree consistently improves over vanilla DFlash across models, and Figure 1 summarizes these gains.

Main contributions.

- We introduce DDTree, a speculative decoding method that constructs a draft tree directly from the per-position distributions produced by a single block diffusion forward pass within a fixed node budget.
- We show that the tree we construct provably maximizes the expected number of accepted tokens under the draft model, which we use as a surrogate for the target model’s expected acceptance length. We also show that the optimal tree can be recovered with a simple and efficient best-first heap algorithm. This result builds on OPT-Tree (Wang et al., 2025), but adapts it to the block diffusion setting, where the tree is constructed from the per-position distributions of a single forward pass rather than the multiple forward passes required by autoregressive drafters.
- We implement DDTree on top of the Hugging Face Transformers library and show consistent gains over vanilla DFlash across model sizes and domains.

2 Related Work

Speculative decoding and tree-based verification. Foundational speculative decoding methods verify a drafted continuation against the target model in parallel and preserve the target model’s output distribution (Leviathan et al., 2023; Chen et al., 2023). A subsequent line of work generalizes this setting from a single drafted continuation to a tree of candidate continuations (Spector & Re, 2023; Jeon et al., 2024; Xiong et al., 2025). SpecInfer by Miao et al. (2023) introduces tree attention for efficient target-model verification over token trees, and Medusa (Cai et al., 2024) combines multiple prediction heads with the same style of tree-based verification. The EAGLE family extends this direction using target-model features: EAGLE drafts in feature space (Li et al., 2024a), EAGLE-2 adds dynamic draft-tree construction (Li et al., 2024b), and EAGLE-3 predicts tokens directly from fused multi-layer features (Li et al., 2025). OPT-Tree by Wang et al. (2025) constructs adaptive trees for autoregressive drafters by maximizing an approximate expected acceptance length under a node budget. In that autoregressive setting, tree construction still requires one drafter’s forward pass per tree depth (i.e., per token position), and thus it has higher computational overhead compared to our DDTree approach.

Parallel drafting and block diffusion drafters. A complementary line of work reduces drafting latency by predicting multiple future tokens in a single forward pass. Block Diffusion provides the modeling foundation for prefix-conditioned blockwise denoising with KV-cache-friendly generation (Arriola et al., 2025). PARD (An et al., 2026) studies low-cost parallel drafting by adapting autoregressive models to mimic diffusion-style block prediction. DFlash (Chen et al., 2026) shows that a small block diffusion drafter, conditioned on target-model features, can predict an entire block in one forward pass and then verify a single drafted trajectory losslessly.

Vanilla DFlash, however, explores only one continuation per round. A very recent work, DART (Liu et al., 2026), also constructs draft trees from one-pass parallel logits. Still, it relies on continuity-aware tree pruning with an external N -gram continuity score and a large N -gram trie at runtime (Liu et al., 2026). By contrast, our proposed DDTree keeps the same one-pass DFlash drafter and constructs the tree directly from the per-position probabilities produced by that pass. This avoids auxiliary external scoring and gives an explicit surrogate objective that our best-first construction provably maximizes.

3 Background: Block Diffusion Drafting

Speculative decoding methods proceed in rounds. At the beginning of each round, we already have one new token produced by the target model; we denote this ‘bonus’ token by b . This token is guaranteed to exist, even if no drafted tokens were accepted in the previous round. Importantly, although b has been selected by the target model, the target model has not yet been run on b , i.e., no forward pass has been performed with b appended to the context. As a result, the drafter can condition on the identity of b , but any target-model features it uses, as in DFlash, are only available for the preceding context.

With b fixed, the drafter’s task is to predict the next tokens after b . A block diffusion drafter does this in parallel, producing a short block of future tokens rather than generating them autoregressively one token at a time. Given the context and b , the drafter takes a masked block of the form $[b, m, \dots, m]$ and predicts tokens for the next L masked positions in one forward pass. This produces logits

$$\ell_i \in \mathbb{R}^{|\mathcal{V}|}, \quad i = 1, \dots, L,$$

or, equivalently, per-position token distributions

$$q_i(v) = \text{softmax}(\ell_i)_v, \quad v \in \mathcal{V}.$$

Crucially, each q_i is a *marginal* distribution for position i , not a path-conditioned distribution. Because the block is predicted in parallel, the prediction at position i conditions on the context before the drafted block (and, in DFlash, on target-model features), but not on the specific token choices at positions $1, \dots, i - 1$ inside that same block. Thus, a single

DFlash pass provides one marginal distribution per future position, rather than conditional probabilities for every partial continuation.

This distinction is central to DDTree. Let c denote the current context, let b denote the bonus token available at the start of the round, and let $y_{1:L}$ denote a candidate continuation after b . Under the target model, the continuation distribution factorizes autoregressively as

$$p(y_{1:L} \mid c, b) = \prod_{i=1}^L p(y_i \mid c, b, y_{1:i-1}). \quad (1)$$

A one-pass block diffusion drafter does not expose these continuation-conditioned factors. Instead, it provides only per-position marginals $\{q_i\}_{i=1}^L$, where each $q_i(\cdot \mid c, b)$ predicts the token at position i from the shared conditioning context (c, b) , without conditioning on the realized tokens at earlier positions within the same block. Accordingly, the natural distribution associated with a one-pass block diffusion drafter is the factorized distribution

$$Q(y_{1:L} \mid c, b) := \prod_{i=1}^L q_i(y_i \mid c, b). \quad (2)$$

Thus, the target model provides a path-conditioned autoregressive distribution, whereas the drafter provides only a factorized distribution over the next L positions.

Our proposed DDTree builds on the factorized distribution (2): we use it to define the surrogate objective for selecting a draft tree from a single block diffusion drafter forward pass. As in standard speculative decoding with a block diffusion drafter, DDTree uses one lightweight drafter pass followed by target-model verification. The key difference is how we use the resulting one-pass marginals $\{q_i\}_{i=1}^L$. Rather than collapsing the marginals into a single continuation, we use them to construct a compact draft tree for verification.

4 The Proposed Method: Diffusion Draft Tree

4.1 Overview

DDTree builds a draft tree under a node budget B , in which a node at depth i represents a candidate token for the i -th future position. Ideally, in each decoding round, we would choose a draft tree with B nodes that maximizes the expected number of speculative tokens accepted by the target model. A single block diffusion drafter forward pass outputs, in parallel, one drafter-predicted marginal token distribution for each future position in the next block. It does not provide the target-conditioned continuation probabilities in (1) needed to optimize the expected target-model acceptance length directly. We therefore can only optimize a surrogate objective, which is the expected acceptance length under the drafter’s factorized approximation (2). In Section 4.2, we formalize this surrogate and derive the corresponding tree-construction objective.

At each decoding round, let c denote the full current context, including the prompt and all previously generated tokens. Recall that the bonus token b is already obtained by the target model: in the first round it comes from the prefill pass, and, in later rounds it is the previous round’s bonus token. With this in place, conditioned on c and b , our proposed DDTree method implements speculative decoding for the next block by applying four steps:

1. Run the block diffusion drafter once to obtain per-position distributions for the next L draft positions after b .
2. Build a draft tree with B nodes from those distributions. The node budget B controls how many candidate tokens the verifier (target model) evaluates.
3. Compile the tree into input tensors for the target model and run one target-model forward pass with tree attention.
4. Walk the tree as illustrated in Figure 2: at each step, use the target model’s decoding rule to choose the next token and check whether it matches a child in the tree; accept the




Figure 2: Illustration of one DDTree decoding round. The bonus token a from the previous round is carried into the current round and serves as the tree root. One block diffusion drafter pass produces per-position marginals for the next three draft positions, from which DDTree selects a draft tree. During verification, the target model follows its own decoding rule, and the walk continues whenever the selected token appears as a child in the tree. In this example, two speculative nodes are matched (green), the matched path is added to the output, and the first unmatched target token a' becomes the bonus token for the next round.

matched drafted path, and carry the first unmatched target token (the next bonus token) to the next round.

4.2 Surrogate objective for draft-tree selection

Consider a decoding round with current context c and bonus token b . A single block diffusion drafter forward pass produces L per-position distributions $q_i(\cdot)$, $i = 1, \dots, L$, for the next L positions after b . For $1 \leq d \leq L$, let

$$u = (u_1, \dots, u_d) \in \mathcal{V}^d \quad (3)$$

denote a candidate continuation prefix after b . We identify each tree node with such a prefix, so a node at depth d represents the path $u_1, \dots, u_d$ from the root b .¹

A draft tree T is a set of such prefixes. The tree is valid if it is prefix-closed: whenever $(u_1, \dots, u_d) \in T$ with $d \geq 2$, the parent prefix $(u_1, \dots, u_{d-1})$ also lies in T . In other words, if a node is included in the tree, then so are all of its ancestors back to the root b . The root b itself does not count toward the node budget. Let

$$N = \sum_{d=1}^L |\mathcal{V}|^d$$

be the total number of nonempty prefixes of length at most L . In the regime relevant for DDTree, this number is enormous, and we always choose node budgets with $B \ll N$. Accordingly, throughout the rest of this section we treat $B < N$ as a standing assumption.

For a candidate continuation $y_{1:L}$, we define its acceptance length under T as the longest prefix of $y_{1:L}$ that appears in the tree:

$$\alpha_T(y_{1:L}) = \max\{d : y_{1:d} \in T\}, \quad (4)$$

with $\alpha_T(y_{1:L}) = 0$ when no depth-1 node matches. With this in place, the ideal tree-construction objective is

$$\max_{T: |T| \leq B, T \text{ valid}} \mathbb{E}_{Y_{1:L} \sim p(\cdot | c, b)} [\alpha_T(Y_{1:L})], \quad (5)$$

namely, the expected acceptance length under the target model p . This objective, however, is infeasible to optimize as it depends on the target-conditioned continuation probabilities (1), which are unavailable when the tree is constructed.

¹Throughout the paper, we use "node" to refer either to the full prefix u or, when the context is unambiguous, to its final token u_d .

We therefore use a surrogate based on the factorized distribution induced by the block diffusion drafter $Q(\cdot | c, b)$, constructed solely from the per-position marginals $\{q_i\}_{i=1}^L$ in (2). This surrogate does not require full path-conditioned probabilities, which are unavailable from a single block diffusion forward pass.

Accordingly, we choose the draft tree by maximizing the expected acceptance length under the factorized distribution:

$$\max_{T: |T| \leq B, T \text{ valid}} \mathbb{E}_{Y_{1:L} \sim Q(\cdot | c, b)} [\alpha_T(Y_{1:L})]. \quad (6)$$

To analyze this surrogate objective, we express it in terms of prefix probabilities under Q . For a prefix $u = (u_1, \dots, u_d)$, the probability that a sampled continuation under Q begins with u is

$$q(u | c, b) = \prod_{i=1}^{|u|} q_i(u_i | c, b). \quad (7)$$

The above is the key quantity for tree construction, because the acceptance length is determined by the longest prefix of the sampled continuation that matches a path in the tree. The next proposition shows that the surrogate objective decomposes as an additive sum of prefix masses over the nodes in T .

Proposition 1. *For any valid draft tree T ,*

$$\mathbb{E}_{Y_{1:L} \sim Q(\cdot | c, b)} [\alpha_T(Y_{1:L})] = \sum_{u \in T} q(u | c, b). \quad (8)$$

All proofs are in Appendix A. Proposition 1 shows that the surrogate objective is an additive sum of prefix probabilities $\sum_{u \in T} q(u | c, b)$. Therefore, selecting a draft tree reduces to choosing a valid set of at most B prefixes that maximizes this sum. Since the objective is additive over nodes and all terms are nonnegative, the optimal solution is to include the highest-probability prefixes up to the budget B , subject to the prefix-closure constraint. The next proposition shows that these top- B prefixes automatically form a valid tree and are therefore optimal.

For the next proposition and the results that follow, we assume $q_i(v | c, b) \in (0, 1)$ for all depths i and tokens v . This is the standard case when the draft model outputs probabilities via a softmax over finite logits. If needed, the construction can be modified slightly to handle edge cases without changing the main results.

Proposition 2. *Let $u^{(1)}, u^{(2)}, \dots$ be all nonempty prefixes of length at most L , ordered so that*

$$q(u^{(1)} | c, b) \geq q(u^{(2)} | c, b) \geq \dots,$$

with ties broken arbitrarily. Define

$$T_B = \{u^{(1)}, \dots, u^{(B)}\}.$$

Then, T_B is a valid draft tree with $|T_B| \leq B$. Moreover, T_B maximizes $\mathbb{E}_{Y_{1:L} \sim Q(\cdot | c, b)} [\alpha_T(Y_{1:L})]$ among all valid draft trees with $|T| \leq B$.

The above result is analogous to OPT-Tree’s expected acceptance length objective (Wang et al., 2025), with the crucial difference that in our case all required probabilities are obtained from a single block diffusion drafter forward pass rather than from multiple autoregressive passes.

Remark 1. The result stated in Proposition 2 is exact for the surrogate objective induced by the factorized draft distribution $Q(\cdot | c, b)$, not for the true target-model distribution. The task here is not to recover the unavailable target model path-conditioned probabilities, but to make the best use of the information available from one block diffusion drafter forward pass.

The correspondence between the surrogate expectation in (8) and target-model acceptance is examined empirically in Appendix C.3, which evaluates each fixed DDTree tree under draft and target probabilities and compares both expectations with realized acceptance.

4.3 Efficient and optimal tree construction

Proposition 2 reveals that an optimal draft tree is obtained by taking the B highest-probability prefixes. The remaining challenge is therefore algorithmic: how can we recover these prefixes efficiently, without enumerating the exponentially many possible prefixes up to depth L ? In what follows, we show that this can be done with a simple best-first search procedure.

At each depth i , let $v_i^{(1)}, v_i^{(2)}, \dots$ be the tokens ordered so that

$$q_i(v_i^{(1)} \mid c, b) \geq q_i(v_i^{(2)} \mid c, b) \geq \dots,$$

and let $q_i^{(k)} = q_i(v_i^{(k)} \mid c, b)$. We index prefixes by token ranks rather than by vocabulary ids. Specifically, a rank tuple

$$\rho = (\rho_1, \dots, \rho_d), \quad 1 \leq d \leq L, \quad 1 \leq \rho_i \leq |\mathcal{V}|,$$

denotes the depth- d prefix $(v_1^{(\rho_1)}, \dots, v_d^{(\rho_d)})$. For example, $\rho = (1, 3, 2)$ denotes the prefix obtained by taking the most probable token at position 1, the third-most probable token at position 2, and the second-most probable token at position 3. The probability of a prefix ρ is

$$q(\rho) = \prod_{i=1}^d q_i^{(\rho_i)},$$

and its log-probability is

$$\sigma(\rho) = \log q(\rho) = \sum_{i=1}^d \log q_i^{(\rho_i)}.$$

We will use $\sigma(\rho)$ to rank prefixes in a max-heap, which we use as a priority queue that returns the prefix with the largest value first. Taking logarithms converts products into sums, improving numerical stability while preserving the ordering of prefixes.

Let $K = \min(B, |\mathcal{V}|)$, and let

$$\mathcal{S}_K = \{(\rho_1, \dots, \rho_d) : 1 \leq d \leq L, 1 \leq \rho_i \leq K \text{ for } i = 1, \dots, d\}.$$

Thus, $\mathcal{S}_K$ contains exactly the prefixes that use only the top- K tokens at every depth.

Lemma 1. *There exists an optimal valid draft tree that maximizes the surrogate expected acceptance length in (6), such that every node in the tree lies in $\mathcal{S}_K$. Equivalently, there exists an optimal valid draft tree in which every node uses only the top- K tokens at each depth.*

Lemma 1 reduces the search space to the top- K tokens at each depth. Algorithm 1 enumerates the retained prefixes $\mathcal{S}_K$ in descending order of $\sigma(\rho)$ using a max-heap. It starts from the tuple (1) , i.e., the length-1 tuple containing only rank 1. When a tuple $\rho = (\rho_1, \dots, \rho_d)$ is popped, the algorithm generates at most two new tuples: its next sibling, which changes only the last rank to $\rho_d + 1$, and its first child, which appends rank 1 at the next depth. Intuitively, the sibling explores an alternative token at the current position, while the child extends the current prefix with the best available token at the next position. After B pops, the algorithm returns the top- B prefixes. The next proposition establishes the optimality of Algorithm 1.

Algorithm 1 Best-first draft-tree construction from one block diffusion drafter pass

Require: Top-K tokens $\{v_i^{(k)}\}_{i=1,k=1}^{L,K}$ and their probabilities $\{q_i^{(k)}\}_{i=1,k=1}^{L,K}$; node budget B

- 1: Initialize max-heap $H \leftarrow \{((1), \sigma((1)))\}$
- 2: Initialize draft tree $T \leftarrow \emptyset$
- 3: **while** $|T| < B$ and $H \neq \emptyset$ **do**
- 4: Pop the rank tuple $\rho = (\rho_1, \dots, \rho_d)$ with largest score $\sigma(\rho)$
- 5: Add prefix $(v_1^{(\rho_1)}, \dots, v_d^{(\rho_d)})$ to T
- 6: **if** $\rho_d + 1 \leq K$ **then**
- 7: Push next sibling $(\rho_1, \dots, \rho_{d-1}, \rho_d + 1)$ with score $\sigma(\rho) - \log q_d^{(\rho_d)} + \log q_d^{(\rho_d+1)}$
- 8: **end if**
- 9: **if** $d < L$ **then**
- 10: Push first child $(\rho_1, \dots, \rho_d, 1)$ with score $\sigma(\rho) + \log q_{d+1}^{(1)}$
- 11: **end if**
- 12: **end while**
- 13: **return** draft tree T

Proposition 3. *Algorithm 1 returns an optimal valid draft tree for the surrogate objective in (6) under node budget B .*

Remark 2. The heap stage costs $O(B \log B)$: the algorithm performs at most B pops and at most $2B$ pushes, and the heap size is $O(B)$ throughout.

4.4 Efficient verification and cache update

Armed with an algorithm that efficiently constructs the optimal draft tree under the node budget, we now describe how the selected draft tree is compiled for verification, how the verifier walk is performed, and how the KV cache is updated afterward.

To verify the selected draft tree in one target-model forward pass, we flatten it into a sequence of token ids rooted at the bonus token b . We assign position ids by tree depth so that the verifier applies the correct positional encoding. We then use tree attention (Miao et al., 2023), under which each drafted node attends to the past context through the KV cache and, within the drafted tree, only to the root, its ancestors, and itself. Together, these inputs let the verifier score all selected tree-branches in a single forward pass.

Verification then follows the target model’s own decoding rule, whether greedy or temperature-based sampling. Figure 2(b) illustrates the process. Starting from the bonus token b , we check whether the token selected by the target model at the current node matches one of that node’s children in the draft tree. If it does, that child is accepted and the walk continues. Since the verifier has already produced outputs for all drafted nodes in the same forward pass, this continuation requires no additional target-model call. If no child matches, the walk stops. The accepted path is appended to the output sequence, the first unmatched target-model token becomes the bonus token for the next round, and the KV cache is compacted to retain only the accepted path.

5 Experiments

We evaluate whether DDTree improves end-to-end speculative decoding over vanilla DFlash. Our experiments focus on decoding speed and acceptance behavior across target model size, domain, and decoding temperature.

5.1 Experimental setup

We evaluate four target models: Qwen3-4B, Qwen3-8B, Qwen3-Coder-30B-A3B-Instruct (Yang et al., 2025), and Gemma 4 31B (Gemma Team et al., 2026), each paired with its corresponding DFlash checkpoint available at <https://huggingface.co/collections/z-lab/dflash>. The Gemma 4 results are reported in Appendix C.1. Our benchmark suite

Dataset	Qwen3-4B				Qwen3-8B				Qwen3-Coder-30B			
	DFlash		+DDTree		DFlash		+DDTree		DFlash		+DDTree	
	Speed	τ	Speed	τ	Speed	τ	Speed	τ	Speed	τ	Speed	τ
<i>Temperature = 0.0</i>												
AIME 2024	5.56×	7.54	7.27×	10.37	5.38×	7.46	7.35×	10.42	3.95×	5.16	5.93×	7.87
AIME 2025	5.33×	7.37	7.23×	10.23	5.32×	7.39	6.99×	9.86	3.98×	5.13	5.88×	7.63
Alpaca	2.03×	3.11	3.32×	5.35	2.07×	3.12	3.36×	5.09	1.53×	2.22	2.46×	3.55
GSM8K	4.77×	6.51	6.58×	9.33	4.78×	6.57	6.75×	9.54	4.00×	5.18	5.83×	7.55
HumanEval	4.81×	6.62	6.81×	9.44	4.84×	6.61	6.90×	9.67	6.09×	8.02	8.22×	10.72
LiveCodeBench	4.97×	7.02	6.78×	9.79	5.02×	7.22	7.10×	10.28	4.72×	6.32	6.54×	8.63
MATH-500	5.54×	7.72	7.50×	10.71	5.56×	7.79	7.52×	10.73	4.29×	5.58	6.21×	8.10
MBPP	4.38×	6.10	6.42×	9.16	4.33×	5.99	6.39×	9.07	5.61×	7.19	7.68×	9.94
MT-Bench	2.64×	4.38	4.18×	6.70	2.56×	4.28	4.10×	6.58	2.04×	3.52	3.27×	5.36
SWE-bench Lite	2.70×	3.66	4.25×	5.99	2.65×	3.60	4.23×	5.91	2.77×	3.61	4.38×	5.71
<i>Temperature = 1.0</i>												
AIME 2024	3.50×	5.01	5.31×	7.82	3.46×	4.98	5.36×	7.79	3.12×	4.23	4.94×	7.32
AIME 2025	3.38×	4.79	5.08×	7.22	3.36×	4.79	5.25×	7.71	3.17×	4.26	4.98×	6.62
Alpaca	1.97×	2.96	3.23×	4.97	1.95×	2.94	3.20×	4.89	1.51×	2.14	2.43×	3.51
GSM8K	4.35×	5.99	6.17×	8.87	4.33×	5.93	6.27×	8.95	3.94×	5.08	5.66×	7.38
HumanEval	4.34×	5.97	6.43×	9.18	4.00×	5.48	6.09×	8.59	5.64×	7.60	7.88×	10.42
LiveCodeBench	4.53×	6.53	6.44×	9.44	4.52×	6.61	6.46×	9.53	3.77×	5.57	5.47×	7.79
MATH-500	4.65×	6.60	6.60×	9.61	4.56×	6.46	6.59×	9.54	4.01×	5.32	5.91×	7.76
MBPP	4.03×	5.56	6.09×	8.72	3.83×	5.30	5.87×	8.34	5.47×	7.03	7.55×	9.76
MT-Bench	2.46×	4.05	3.91×	6.27	2.30×	3.77	3.75×	6.02	1.96×	3.43	3.10×	5.14
SWE-bench Lite	2.29×	3.07	3.71×	5.20	2.10×	2.82	3.47×	4.86	2.42×	3.16	3.80×	4.96

Table 1: Speedup over autoregressive decoding and mean acceptance length (τ). Results are reported separately for temperature 0.0 and temperature 1.0. For each dataset, model, and temperature, the DDTree entry uses the best node budget from $\{16, 32, 64, 128, 256, 512, 1024\}$.

spans reasoning tasks such as MATH-500 (Lightman et al., 2023), GSM8K (Cobbe et al., 2021), AIME 2024, and AIME 2025; code tasks such as HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), LiveCodeBench (Jain et al., 2025), and SWE-bench Lite (Jimenez et al., 2024); and general instruction and dialogue tasks such as MT-Bench (Zheng et al., 2023) and Alpaca (Taori et al., 2023). We run the benchmark on 8 H200 GPUs at temperatures 0.0 and 1.0 and report speedup relative to autoregressive decoding, mean acceptance length τ including the bonus token, and, for the case study below, the acceptance length histogram. Additional benchmark details, including per-dataset sample counts, decoding hyperparameters, warmup, and backend selection, are provided in Appendix B.

5.2 Main results

Figure 1 and Table 1 summarize the main benchmark results. Figure 1 summarizes the temperature 0.0 results across datasets and target models, using the best DDTree node budget for each dataset-model pair. Table 1 reports the exact numbers at both temperatures.

DDTree improves every entry in Table 1, covering all $10 \times 3 \times 2 = 60$ dataset-model-temperature settings. The gains are consistent across all three target models, across reasoning, code, and general instruction tasks, and at both temperatures.

5.3 Budget-quality tradeoff

Figure 3 shows a case study on MATH-500 with Qwen3-8B at temperature 0.0. As our DDTree node budget grows, acceptance length increases steadily, and the end-to-end speedup improves until it peaks around budgets of 256 to 512. Pushing the budget to 1024 increases acceptance length further, but the tradeoff is no longer favorable as the additional overhead of verifying more drafted tokens outweighs the gain from the longer

accepted prefix. Vanilla DFlash uses a single block of size 16, so the figure also shows that DDTree provides better speedup under the same conceptual budget. This highlights the importance of a front-heavy tree that does not waste budget on low-probability trajectories. Note that the optimal budget can shift across hardware platforms and implementations (Appendix C.4).




Figure 3: Budget tradeoff on MATH-500 with Qwen3-8B at temperature 0.0. Acceptance length increases steadily with the DDTree node budget, while speedup peaks at an intermediate budget once verifier cost becomes dominant.

5.4 Acceptance length distribution

Figure 4 shows the histogram of acceptance lengths on MATH-500 with Qwen3-8B at temperature 0.0. The plotted DDTree distribution uses the best speedup budget, $B = 512$. Compared with vanilla DFlash, DDTree shifts substantial probability mass toward longer accepted prefixes. With DDTree, it becomes much rarer to observe acceptance lengths below 4, while full-block acceptance at length 16 becomes substantially more common. This shift explains the end-to-end speedup improvement: DDTree makes long accepted prefixes substantially more common, so the verifier needs fewer rounds per generated token.




Figure 4: Acceptance length distribution on MATH-500 with Qwen3-8B at temperature 0.0. The DDTree histogram uses the best speedup budget, $B = 512$, and shifts mass toward longer accepted prefixes, especially full block acceptances.

Acknowledgments

L. R. and Y. R. were supported by the European Union (ERC, SafetyBounds, 101163414). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This research was also partially supported by the Israel Science Foundation (ISF grant 729/21). Y. R. acknowledges additional support from the Career Advancement Fellowship at the Technion. The contribution of the first author is part of a PhD thesis research conducted at the Technion.

References

- Zihao An, Huajun Bai, Ziqiong Liu, Dong Li, and Emad Barsoum. PARD: Accelerating LLM inference with low-cost PARallel draft model adaptation. In *International Conference on Learning Representations*, 2026.
- Marianne Arriola, Subham Sekhar Sahoo, Aaron Gokaslan, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Justin T Chiu, and Volodymyr Kuleshov. Block diffusion: Interpolating between autoregressive and diffusion language models. In *International Conference on Learning Representations*, 2025.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Tianle Cai, Yuhong Li, Zhengyang Geng, Hongwu Peng, Jason D. Lee, Deming Chen, and Tri Dao. Medusa: Simple LLM inference acceleration framework with multiple decoding heads. In *International Conference on Machine Learning*, 2024.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*, 2023.
- Jian Chen, Yesheng Liang, and Zhijian Liu. DFlash: Block diffusion for flash speculative decoding. *arXiv preprint arXiv:2602.06036*, 2026.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé, Jared Kaplan, Harrison Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mo Bavarian, Clemens Winter, Phil Tillet, Felipe Petroski Such, David W. Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William H. Guss, Alex Nichol, Igor Babuschkin, Suchir Balaji, Shantanu Jain, Andrew Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew M. Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Karl Cobbe, Vineet Kosaraju, Mo Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations*, 2024.
- Gemma Team, Sherif El Abd, Vaibhav Aggarwal, Robin Algayres, Alek Andreev, Olivier Bachem, Ian Ballantyne, Cormac Brick, Victor Cărbune, Michelle Casbon, et al. Gemma 4 technical report. *arXiv preprint arXiv:2607.02770*, 2026.

- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. In *International Conference on Learning Representations*, 2025.
- Wonseok Jeon, Mukul Gagrani, Raghavv Goel, Junyoung Park, Mingu Lee, and Christopher Lott. Recursive speculative decoding: Accelerating LLM inference via sampling without replacement. In *Workshop on Large Language Model (LLM) Agents, ICLR*, 2024.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations*, 2024.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*. PMLR, 2023.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE: Speculative sampling requires rethinking feature uncertainty. In *International Conference on Machine Learning*, 2024a.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. Eagle-2: Faster inference of language models with dynamic draft trees. In *Conference on empirical methods in natural language processing*, pp. 7421–7432, 2024b.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE-3: Scaling up inference acceleration of large language models via training-time test. In *Conference on Neural Information Processing Systems*, 2025.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *International conference on learning representations*, 2023.
- Fuliang Liu, Xue Li, Ketai Zhao, Yinxi Gao, Ziyan Zhou, Zhonghui Zhang, Zhibin Wang, Wanchun Dou, Sheng Zhong, and Chen Tian. Dart: Diffusion-inspired speculative decoding for fast llm inference. *arXiv preprint arXiv:2601.19278*, 2026.
- Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Xinhao Cheng, Zeyu Wang, Zhengxin Zhang, Rae Ying Yee Wong, Alan Zhu, Lijie Yang, Xiaoxiang Shi, Chunan Shi, Zhuoming Chen, Daiyaan Arfeen, Reyna Abhyankar, and Zhihao Jia. Specinfer: Accelerating large language model serving with tree-based speculative inference and verification. *ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2023.
- Benjamin Frederick Spector and Christopher Re. Accelerating LLM inference with staged speculative decoding. In *Workshop on Efficient Systems for Foundation Models, ICML*, 2023.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Jikai Wang, Yi Su, Juntao Li, Qingrong Xia, Zi Ye, Xinyu Duan, Zhefeng Wang, and Min Zhang. Opt-tree: Speculative decoding with adaptive draft tree structure. *Transactions of the Association for Computational Linguistics*, 13:188–199, 2025.
- Yunfan Xiong, Ruoyu Zhang, Yanzeng Li, and Lei Zou. Dyspec: Faster speculative decoding with dynamic token tree structure. *World Wide Web*, 28(3):36, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxin Yang, Jingren Zhou, Jingren Zhou, Junyan Lin, Kai Dang, Keqin Bao, Ke-Pei Yang, Le Yu, Li-Chun Deng, Mei Li, Min Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shi-Qiang Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng

Ren, Yang Fan, Yang Su, Yi-Chao Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

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

A Mathematical proofs

Proof of Proposition 1. We expand

$$\alpha_T(Y_{1:L}) = \sum_{d=1}^L \mathbf{1}\{\alpha_T(Y_{1:L}) \geq d\}. \quad (9)$$

For a fixed depth d , the event $\{\alpha_T(Y_{1:L}) \geq d\}$ holds if and only if the sampled depth- d prefix $Y_{1:d}$ is one of the depth- d nodes in T . These depth- d events are disjoint because $Y_{1:d}$ can equal only one sequence. Therefore

$$\Pr[\alpha_T(Y_{1:L}) \geq d] = \sum_{u \in T: |u|=d} \Pr[Y_{1:d} = u] = \sum_{u \in T: |u|=d} q(u \mid c, b). \quad (10)$$

Summing over d gives the claim. $\square$

Proof of Proposition 2. Let v be any strict descendant of u . Then

$$q(v \mid c, b) = q(u \mid c, b) \prod_{i=|u|+1}^{|v|} q_i(v_i \mid c, b) < q(u \mid c, b). \quad (11)$$

Therefore, every strict ancestor has a strictly larger probability than any of its descendants. It follows that, in any nonincreasing ordering of prefixes, every ancestor must appear before its descendants, which implies that T_B is prefix-closed and hence valid.

For optimality, Proposition 1 shows that the surrogate objective for a tree T is

$$\sum_{u \in T} q(u \mid c, b).$$

This objective is additive over nodes, and every term is nonnegative. Therefore, among all sets of at most B prefixes, the maximum is attained by taking the top- B probabilities $q(u \mid c, b)$, namely the prefixes in T_B . Since T_B is valid, it is optimal among all valid draft trees with $|T| \leq B$. If ties occur between unrelated prefixes, other optimal trees may also exist. $\square$

Proof of Lemma 1. Order all nonempty prefixes by nonincreasing $q(u \mid c, b)$, breaking ties by placing prefixes in $\mathcal{S}_K$ before prefixes outside $\mathcal{S}_K$. By Proposition 2, the first B prefixes in this order form an optimal valid draft tree. It therefore suffices to show that all of these first B prefixes lie in $\mathcal{S}_K$.

If $B \geq |\mathcal{V}|$, then $K = \min(B, |\mathcal{V}|) = |\mathcal{V}|$, so every prefix already lies in $\mathcal{S}_K$. Thus only the case $B < |\mathcal{V}|$ remains, in which case $K = B$.

Assume by contradiction that some prefix $u = (u_1, \dots, u_d) \notin \mathcal{S}_K$ appears among the first B prefixes. Let

$$I = \{j \in \{1, \dots, d\} : \text{the rank of } u_j \text{ at depth } j \text{ exceeds } B\},$$

pick any $i \in I$, and let $\tilde{u}$ be obtained from u by replacing u_j with $v_j^{(1)}$ for every $j \in I$. For each $k \in \{1, \dots, B\}$, let $u^{(k)}$ be obtained from $\tilde{u}$ by replacing its i -th coordinate with $v_i^{(k)}$.

Then $u^{(1)}, \dots, u^{(B)}$ are B distinct prefixes in $\mathcal{S}_K$, and for every k ,

$$q(u^{(k)} \mid c, b) \geq q(u \mid c, b),$$

since each replaced coordinate is changed to a token of rank at most B . By the tie-breaking rule, each $u^{(k)}$ appears before u in the ordering. Hence at least B prefixes appear before u , contradicting that u is among the first B .

Therefore no prefix outside $\mathcal{S}_K$ can appear among the first B prefixes. Hence all first B prefixes lie in $\mathcal{S}_K$, and so some optimal valid draft tree is contained in $\mathcal{S}_K$. $\square$

Proof of Proposition 3. Each $\rho \in \mathcal{S}_K$ corresponds to the prefix $(v_1^{(\rho_1)}, \dots, v_d^{(\rho_d)})$, and its score is

$$\log q(\rho) = \sum_{i=1}^d \log q_i^{(\rho_i)}.$$

By Lemma 1, some optimal valid draft tree is contained in $\mathcal{S}_K$. It is therefore enough to show that Algorithm 1 returns the B highest-scoring elements of $\mathcal{S}_K$, in nonincreasing order of $q(\rho)$.

For every $\rho = (\rho_1, \dots, \rho_d) \in \mathcal{S}_K$ other than (1), define its predecessor by

$$\text{pred}(\rho) = \begin{cases} (\rho_1, \dots, \rho_{d-1}, \rho_d - 1), & \rho_d > 1, \\ (\rho_1, \dots, \rho_{d-1}), & \rho_d = 1, d > 1. \end{cases}$$

This predecessor is unique. If $\rho_d > 1$, then popping $\text{pred}(\rho)$ generates ρ as the next sibling. If $\rho_d = 1$ and $d > 1$, then popping $\text{pred}(\rho)$ generates ρ as the first child.

We next show that the heap pops tuples in nonincreasing order of $q(\rho)$. The first pop is (1), since it is the unique initial heap element. Now consider any later iteration. At that point, the heap contains exactly the unpopped tuples whose predecessor has already been popped. Fix any unpopped tuple ρ , and let $\bar{\rho}$ be the first tuple on the predecessor chain from ρ back toward (1) whose predecessor has already been popped. Then $\bar{\rho}$ is in the heap. Moreover,

$$q(\text{pred}(\rho)) \geq q(\rho)$$

for every noninitial tuple ρ : in the sibling case this replaces $q_d^{(\rho_d)}$ by the larger quantity $q_d^{(\rho_d-1)}$, and in the child case it removes the factor $q_d^{(1)} < 1$. Applying this inequality repeatedly along the chain from ρ to $\bar{\rho}$ gives $q(\bar{\rho}) \geq q(\rho)$. Therefore the maximum-probability unpopped tuple is always present in the heap, and each pop removes a highest-probability remaining tuple. Repeating this argument proves that the pop sequence is nonincreasing in $q(\rho)$.

After B pops, the returned set T therefore consists of the top- B prefixes. It remains to show that T is valid. Let $\rho = (\rho_1, \dots, \rho_d) \in T$ with $d \geq 2$. Repeatedly applying pred to ρ decreases the last coordinate until reaching $(\rho_1, \dots, \rho_{d-1}, 1)$, and one more predecessor step yields the parent tuple $(\rho_1, \dots, \rho_{d-1})$. Since each tuple is generated only after its predecessor is popped, every tuple on this chain was popped earlier and therefore belongs to T . Hence the parent of every returned nonroot prefix is also returned, so T is prefix-closed.

After B pops, the returned set T is exactly the top- B elements of $\mathcal{S}_K$. By Proposition 2, T is a valid draft tree and is optimal among trees contained in $\mathcal{S}_K$. By Lemma 1, some global optimum is contained in $\mathcal{S}_K$. Hence T is optimal for (6).

Finally, Proposition 1 implies that, within $\mathcal{S}_K$, the surrogate objective is maximized by taking the highest-probability prefixes. Since Algorithm 1 returns exactly those prefixes, and some global optimum lies in $\mathcal{S}_K$, the returned tree is optimal for (6) under node budget B . $\square$

B Benchmark details

All runs use block size 16, DDTree node budgets $\{16, 32, 64, 128, 256, 512, 1024\}$, temperatures 0.0 and 1.0, a maximum of 2048 new tokens, and bfloat16 inference. We run the benchmark on 8 H200 GPUs and shard the evaluation set across workers. Before timing the benchmark loop, we run a short warmup prompt through the autoregressive baseline, vanilla DFlash, and each DDTree budget so that one time setup costs are excluded from the reported measurements.

For DDTree, the target model uses standard PyTorch scaled dot product attention, because FlashAttention-2 (Dao, 2024) does not support the required tree attention pattern. The DFlash drafter itself still uses FlashAttention-2. For fairness, for the autoregressive baseline and vanilla DFlash, we evaluate the target model with both standard PyTorch scaled dot product attention and FlashAttention-2, and report the faster result, which can only improve these baselines relative to DDTree.

Table 2 lists the number of evaluated examples for each dataset. We follow the original DFlash benchmark setup for these sample counts.

Dataset	Examples
AIME 2024	30
AIME 2025	30
Alpaca	128
GSM8K	128
HumanEval	164
LiveCodeBench	128
MATH-500	128
MBPP	128
MT-Bench	80
SWE-bench Lite	128

Table 2: Number of evaluated examples per dataset in the benchmark suite.

C Additional experimental results

C.1 Gemma 4 31B results

Table 3 reports the complete Gemma 4 31B evaluation. It uses the same DFlash setup, benchmark suite, and budget selection protocol as the Qwen3 evaluation in Table 1.

Dataset	DFlash		DFlash+DDTree	
	Speedup	τ	Speedup	τ
<i>Temperature = 0.0</i>				
MATH-500	6.62×	7.44	8.54×	9.62
GSM8K	5.94×	6.65	7.88×	9.06
AIME 2024	5.84×	6.58	7.64×	8.62
AIME 2025	5.81×	6.59	7.66×	8.68
HumanEval	7.16×	8.27	9.01×	10.24
MBPP	5.04×	5.59	6.90×	7.69
LiveCodeBench	5.50×	6.36	7.35×	8.45
SWE-bench Lite	3.10×	3.43	4.53×	5.07
MT-Bench	2.91×	3.94	4.15×	5.42
Alpaca	2.61×	3.34	3.78×	4.93
<i>Temperature = 1.0</i>				
MATH-500	6.12×	6.91	8.00×	9.05
GSM8K	5.75×	6.43	7.51×	8.45
AIME 2024	5.27×	5.99	7.02×	7.99
AIME 2025	5.34×	6.04	6.89×	7.88
HumanEval	6.82×	7.97	8.71×	10.01
MBPP	4.84×	5.36	6.64×	7.77
LiveCodeBench	5.06×	5.90	6.70×	7.80
SWE-bench Lite	3.00×	3.32	4.33×	4.85
MT-Bench	2.79×	3.82	3.97×	5.25
Alpaca	2.60×	3.31	3.67×	4.58

Table 3: Gemma 4 31B results. Speedup is relative to autoregressive decoding and τ is mean acceptance length. Each DDTree row uses the best speedup budget.

DDTree improves both speedup and τ in all 20 Gemma dataset-temperature settings. Averaged over the ten datasets, speedup rises from $5.05\times$ to $6.74\times$ at temperature 0.0 and from $4.76\times$ to $6.34\times$ at temperature 1.0. These results show that the gains are not specific to a single model family.

C.2 Controlled comparison with DART

DART (Liu et al., 2026) is a closely related speculative decoding method. Its lightweight drafter predicts logits for multiple future positions in one pass from target-model hidden states, and an expand/prune procedure converts those logits into a draft tree. DART ranks candidates using a scoring function called COMBINE, which uses draft-logit and tree-level weights together with an external N -gram continuity score. DDTree likewise starts from one-pass parallel logits, but derives a best-first tree-construction algorithm by maximizing a stated surrogate expected-acceptance objective.

To isolate the effect of tree construction, all methods in Table 4 use the same Qwen3-4B target model and the same DFlash drafter; DART’s tree search therefore receives the DFlash per-position logits rather than outputs from its own drafter. The DFlash columns report the vanilla single trajectory baseline. We evaluate the same ten datasets at temperatures 0.0 and 1.0. Reported speedup is relative to autoregressive decoding, and τ denotes mean acceptance length.

We compare DDTree with node budget $B = 59$ against DART with maximum tree size $\theta = 59$, using the hyperparameters from DART’s Appendix D. We evaluate both DART’s default N -gram setting, $w_{\text{ng}} = 0.5$, and the ablation $w_{\text{ng}} = 0$. Importantly, DART with $w_{\text{ng}} = 0$ is still not DDTree: it retains the weights of the COMBINE scoring function, $w_{\text{logit}} = 0.9^{\text{level}}$ and $w_{\text{level}} = (\text{level} + 1)^{-0.7}$, as well as its expand/prune procedure. We reuse DART’s official C++ tree-search and N -gram implementation. DDTree has no tree-construction hyperparameter other than B . We restrict the comparison to 59 nodes because DART does not specify how its per-position candidate counts k_i and beam width w should be adjusted for larger budgets, which tend to favor DDTree. Thus, this choice is favorable to DART.

Dataset	DFlash		DDTree		DART $w_{\text{ng}} = 0$		DART $w_{\text{ng}} = 0.5$	
	Speedup	τ	Speedup	τ	Speedup	τ	Speedup	τ
<i>Temperature = 0.0</i>								
AIME 2024	$5.61 \times$	7.56	$6.75 \times$	9.07	$6.50 \times$	9.39	$5.01 \times$	7.40
AIME 2025	$5.57 \times$	7.55	$6.70 \times$	9.07	$6.37 \times$	9.31	$4.74 \times$	7.14
Alpaca	$2.08 \times$	3.12	$2.96 \times$	4.36	$2.48 \times$	3.99	$2.67 \times$	4.31
GSM8K	$4.88 \times$	6.57	$6.24 \times$	8.35	$5.48 \times$	7.97	$4.72 \times$	6.99
HumanEval	$4.87 \times$	6.58	$6.43 \times$	8.70	$5.64 \times$	8.25	$5.08 \times$	7.70
LiveCodeBench	$5.09 \times$	7.03	$6.51 \times$	8.95	$5.57 \times$	8.36	$5.33 \times$	8.12
MATH-500	$5.75 \times$	7.82	$7.12 \times$	9.66	$6.55 \times$	9.54	$5.35 \times$	8.02
MBPP	$4.50 \times$	6.11	$5.97 \times$	8.15	$5.18 \times$	7.66	$4.78 \times$	7.34
MT-Bench	$2.54 \times$	4.42	$3.60 \times$	5.88	$3.01 \times$	5.52	$3.13 \times$	5.38
SWE-bench Lite	$2.80 \times$	3.65	$3.78 \times$	5.15	$3.33 \times$	4.92	$3.20 \times$	4.87
<i>Temperature = 1.0</i>								
AIME 2024	$3.58 \times$	4.97	$5.00 \times$	6.93	$4.04 \times$	6.46	$3.65 \times$	5.98
AIME 2025	$3.54 \times$	4.97	$4.65 \times$	6.44	$4.10 \times$	6.51	$3.71 \times$	5.99
Alpaca	$2.01 \times$	3.00	$2.85 \times$	4.16	$2.33 \times$	3.81	$2.46 \times$	4.16
GSM8K	$4.45 \times$	6.03	$5.81 \times$	7.86	$4.77 \times$	7.34	$4.10 \times$	6.66
HumanEval	$4.45 \times$	5.95	$6.03 \times$	8.10	$4.94 \times$	7.59	$4.49 \times$	7.31
LiveCodeBench	$4.74 \times$	6.62	$5.96 \times$	8.22	$4.96 \times$	7.91	$4.72 \times$	7.78
MATH-500	$4.74 \times$	6.59	$6.26 \times$	8.60	$5.22 \times$	8.23	$4.49 \times$	7.43
MBPP	$4.17 \times$	5.64	$5.55 \times$	7.56	$4.55 \times$	7.02	$4.26 \times$	7.05
MT-Bench	$2.42 \times$	4.03	$3.40 \times$	5.47	$2.72 \times$	4.99	$2.80 \times$	5.03
SWE-bench Lite	$2.32 \times$	3.04	$3.36 \times$	4.43	$2.72 \times$	4.10	$2.77 \times$	4.37
Overall average	$4.01 \times$	5.56	$5.25 \times$	7.25	$4.52 \times$	6.94	$4.07 \times$	6.45

Table 4: Controlled tree-builder comparison on Qwen3-4B. The overall average is computed across both temperatures.

DDTree gives the highest wall-clock speedup in every setting in Table 4 and also has the best average τ . DART with $w_{\text{ng}} = 0$ is the stronger DART variant on average, but remains

consistently slower than DDTree. Adding the N -gram score reduces both average speedup and average τ .

We do not interpret these results as evidence that N -gram information is inherently harmful. With a strong DFlash drafter, its probabilities appear to be the more useful signal in this setting. N -grams may remain useful in copy-heavy settings such as code editing.

C.3 Multi-sample baseline and surrogate fidelity

We use two complementary analyses to assess the practical value of DDTree’s surrogate objective. The multi-sample comparison tests whether DDTree benefits merely from branching or specifically from selecting the tree that maximizes the surrogate; the fidelity analysis then tests whether the draft-surrogate expectation tracks target-model acceptance.

Multi-sample baseline. The most direct tree extension of DFlash is to sample several continuations from the per-position distributions produced by one DFlash forward pass and merge them into a prefix tree. The multi-sample baseline draws B/L such continuations, using the same DFlash output, the same $B = 128$ maximum node budget, and the same target-model verification as DDTree. Thus, the methods differ in how they select tree nodes: the baseline samples complete paths, whereas DDTree selects the surrogate-optimal prefix set.

Dataset	Multi-sample τ	DDTree τ	$\Delta\tau$	Speedup
MATH-500	7.61	8.95	1.34	1.20×
Alpaca	3.50	4.47	0.96	1.31×

Table 5: DDTree versus a multi-sample tree at temperature 1.0 and $B = 128$. The final column is DDTree’s wall-clock speedup relative to the multi-sample baseline.

DDTree outperforms the multi-sample baseline on both datasets. In these settings, the improvement therefore does not come merely from branching DFlash into several sampled continuations; selecting the tree according to the surrogate objective produces a more effective tree under the same budget.

Surrogate fidelity. DDTree’s optimality guarantee applies to the factorized draft distribution Q , rather than the target distribution p . We therefore test whether the draft-surrogate expectation tracks target-model acceptance. In each decoding round, DDTree first constructs a tree from the draft distributions. Holding that tree fixed, we compute its expected acceptance using both the draft prefix masses and, post hoc, the target-model path probabilities. We average both expectations over decoding rounds. Table 6 compares these expectations with realized mean acceptance τ . Because reported acceptance includes the bonus/root token, both expectations include an added 1.

Dataset	B	Draft $1 + \mathbb{E}_q[\alpha_T]$	Target $1 + \mathbb{E}_p[\alpha_T]$	τ
Alpaca	32	3.86	3.71	3.72
Alpaca	64	4.14	3.96	3.98
Alpaca	128	4.42	4.20	4.19
Alpaca	256	4.65	4.40	4.41
MATH-500	32	7.68	7.54	7.53
MATH-500	64	8.13	7.93	7.91
MATH-500	128	8.66	8.34	8.35
MATH-500	256	9.25	8.82	8.82

Table 6: Draft-surrogate expectation, target-model expectation, and realized mean acceptance at temperature 1.0.

As expected, $1 + \mathbb{E}_p[\alpha_T]$ agrees with realized τ across budgets, with an absolute difference of at most 0.02 tokens. The draft quantity $1 + \mathbb{E}_q[\alpha_T]$ follows the same trend with mild

optimism, indicating that, in these settings, the one-pass draft surrogate closely tracks the acceptance behavior of the trees DDTree constructs.

C.4 Budget and hardware sweeps

Budget sweep. The node budget B trades longer accepted prefixes against the cost of verifying a larger tree. Table 7 reports performance at every evaluated budget for Gemma 4 31B at temperature 0.0, averaged over all ten datasets.

Method	Avg. speedup	Avg. τ
DFlash	$5.05\times$	5.82
DDTree $B = 16$	$5.56\times$	6.41
DDTree $B = 32$	$5.92\times$	6.82
DDTree $B = 64$	$6.29\times$	7.16
DDTree $B = 128$	$6.57\times$	7.48
DDTree $B = 256$	$6.74\times$	7.73
DDTree $B = 512$	$6.61\times$	8.00
DDTree $B = 1024$	$4.10\times$	8.24

Table 7: Gemma 4 31B budget sweep at temperature 0.0, averaged over all ten datasets.

Average acceptance length increases throughout the sweep, while average speedup peaks at $B = 256$ and then declines. This pattern reflects the budget-quality tradeoff: larger trees provide more opportunities for acceptance, but their additional verification cost can eventually outweigh the benefit.

Hardware sweep. To test whether the speedup-optimal budget transfers across hardware, we repeat the budget sweep for Qwen3-4B and Qwen3-8B on NVIDIA H200 and A40 GPUs at both temperatures. For each model, temperature, and platform, Table 8 reports the budget that maximizes speedup averaged over the ten datasets.

Model	Temp.	Best on H200	Best on A40
Qwen3-4B	0	$B = 512, 6.03\times$	$B = 64, 5.19\times$
Qwen3-4B	1	$B = 512, 5.29\times$	$B = 64, 4.54\times$
Qwen3-8B	0	$B = 512, 6.07\times$	$B = 32, 4.63\times$
Qwen3-8B	1	$B = 512, 5.23\times$	$B = 32, 3.97\times$

Table 8: Speedup-optimal DDTree budget averaged over ten datasets on NVIDIA H200 and A40 GPUs.

In this sweep, H200 selects $B = 512$ for both models and temperatures, whereas A40 selects $B = 64$ for Qwen3-4B and $B = 32$ for Qwen3-8B. Thus, the speedup-optimal budget can depend on both the hardware platform and target model and should be tuned for the deployment setting.