arXiv:2611.33896v1 [cs.LG] 25 Nov 2026

Looped Models are Fast Multi-Token Predictors

Ayhan Suleymanzade[1], **Sean McLeish**[1,2], **Sehoon Kim**[3], **Jinwoo Kim**[3], **Antoine Bosselut**[1]

[1]EPFL [2]University of Maryland [3]KAIST

Abstract. Looped language models can adapt computation for predicting a token without a proportional rescaling in unique model parameters by adjusting the number of loops. This property has been used for self-speculative decoding, taking shallow loops as fast drafters and deep loops as verifiers, which is effective since drafting still uses full parameters. In this paper, we show that shallow loops can be also predictive of several future tokens and leverage this to drastically accelerate the inference of looped language models. To that end, we present **Floodgates**, a simple self-speculative decoding method that appends a block of mask tokens as a suffix and learns via low-rank adaptation to draft these tokens in parallel. Across seven models, five benchmarks, a range of batch sizes, and both greedy and sampled decoding, Floodgates is the fastest lossless decoding method in most cases, attaining up to 7.5× acceleration. Through analysis, we show that this speedup is enabled by all-depth training of looped models, with the shallow drafter's ability to leverage full-depth computation by conditioning on deep prefix cache evaluated during verification.

Code: https://github.com/MisakiTaro0414/floodgates

1 Introduction

Autoregressive (AR) transformer language models produce text one token after another. In each layer, the hidden state of the new token is updated by attending to the hidden states of all earlier tokens by the preceding layer. A token can therefore be generated only after all of the tokens before it, leading to a serialized bottleneck in speed (Santilli et al., 2023; Suleymanzade et al., 2026a). Due to local regularities in natural language (Church and Hanks, 1990), many upcoming tokens can be easy to predict, and speculative decoding uses this to accelerate generation (Leviathan et al., 2023; Chen et al., 2023). A fast drafter proposes several tokens, and the model verifies all of them in one forward pass. Yet, the drafter is usually a small separate network (Li et al., 2024a; Chen et al., 2026a) or a subset of the layers (Zhang et al., 2024; Elhoushi et al., 2024; Liu et al., 2024), computed with only a fraction of the parameters of the model. Thus, the drafters are usually prevented from using the full knowledge encoded in all layers of the model, leading to a limited verification rate of their drafts.

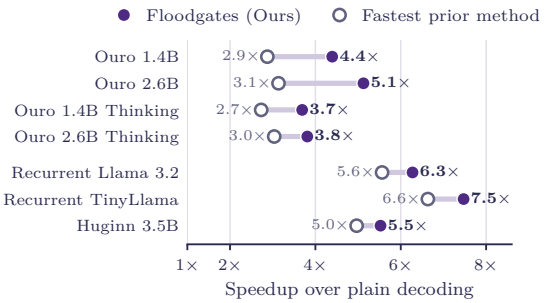

Figure 1 Floodgates is up to 7.5× faster than autoregressive generation and faster than prior methods on every model. Speedup over plain decoding under greedy decoding at batch size one, averaged over the benchmarks of each model. Open circles show the fastest prior lossless method on each model.

Looped language models naturally avoid this limitation. They process each token through a group of layers repeatedly (Dehghani et al., 2019; Geiping et al., 2025; Zhu et al., 2025), so they can reduce their computation for faster generation by shallowing their loop without a proportional decrease in their parameters. This allows for an effective self-speculative decoding scheme that uses the deep loop to verify each token drafted by a shallow loop, in which the drafter now has access to the full knowledge within the model since every loop runs the same block. Yet, existing methods either draft a single token per loop (Cho et al., 2026; Ha and Kim, 2026) or attain good acceptance rate only when the model runs many loops (Li et al., 2026). Moreover, they often rely on *pipelining* for parallelism, starting the next token before the current one has finished its loops. As tokens at different loops are batched together, the methods need to maintain separate, per-depth prefix cache across all batched loops, diminishing their utility when serving many requests at once where memory becomes a bottleneck.

Correspondence to: ayhan.suleymanzade@epfl.ch.

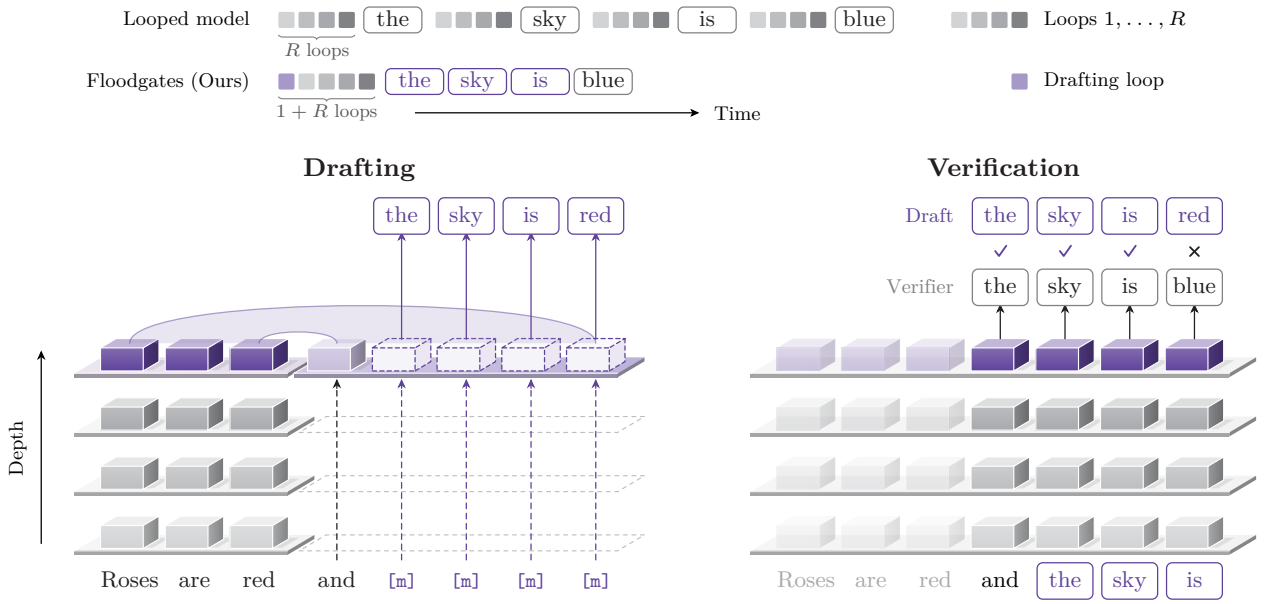

Figure 2 Floodgates drafts a block of tokens with a single loop and verifies it in one decoding step. (Left) A single loop reads the deep prefix cache (purple) computed during the previous verification and drafts several tokens (lilac) in parallel from mask tokens. (Right) The full model verifies the drafts in one decoding step, accepts them up to the first mismatch, and produces the prefix cache for the next draft. (Top) This allows the shallow drafter to leverage full-depth computation for fast speculative decoding.

In this paper, we show that a single loop of a looped language model can be adapted to draft several future tokens in parallel while attending to the final-loop prefix cache computed during verification, drastically accelerating inference. Our first intuition is that looped training can make shallow loops predictive of several future tokens, as it supervises every depth and hence pressures each future token to be predicted as early as possible, encouraging shallow hidden states of past tokens to be already helpful (Section 4.2). The second is that the shallow drafter can make use of full-depth computation through the prefix cache, which holds the keys and values of earlier tokens at the final loop. The model computes this cache during verification, so the drafter can attend to it "for free". We present **Floodgates**, a simple self-speculative decoding method that appends a block of mask tokens as a suffix and learns to read the final-loop prefix cache to draft these tokens in parallel using a single loop iteration. We use low-rank adapters (Hu et al., 2022) that are active only while drafting, so the verifier parameters do not change and thus decoding remains lossless. On models that run many loops, where a single loop costs only a small fraction of a decoding step, Floodgates further refines this draft with additional passes of the pretrained loop (*elastic drafting*) to balance drafting cost and acceptance without extra training. Across seven looped models and five math, coding, and reasoning benchmarks, Floodgates decodes faster than all prior lossless methods under greedy decoding and in most cases under sampling (Figure 1). An illustration is given in Figure 2. In summary, we make the following contributions:

(1) We introduce Floodgates, a simple self-speculative decoding method that adapts a single loop via low-rank adapters to draft several tokens in parallel from mask tokens while reading the deep prefix cache from verification, and refines the draft with elastic passes of the pretrained loop on deep models (Section 3). It overcomes the limitations of existing methods by effective multi-token prediction and its access to the deep prefix cache.

(2) We evaluate Floodgates on seven looped models and five math, coding and reasoning benchmarks. Floodgates is the fastest lossless decoding method on every model and every task under greedy decoding and in most cases under sampling (Section 4.1). It decodes up to 5.1× faster than plain decoding on average with a single loop on 4-loop models and up to 7.5× faster with elastic drafting on 32-loop models, and on some models, it is nearly twice as fast as the best prior method. At the largest batch size that fits in memory, it reaches up to 1.9× the throughput of the strongest baseline method (Section 4.4).

(3) Through controlled comparisons, we show that self-speculative drafting with a single loop iteration gives larger speedups than drafting with a separate smaller network (EAGLE, DFlash) or with the first layers of a non-looped model (Llama 3.2, LayerSkip), supporting the utility of drafting with all of the parameters of the model (Section 4.2). We also show that deep prefix cache accelerates decoding up to 1.3× (Section 4.2), and that elastic drafting maximizes acceleration by optimizing the cost-acceptance tradeoff across verifiers of varying depths (Section 4.3).

2 Background

Let V denote a token vocabulary, and let $\Delta^{|V|-1}$ be the probability simplex over V, whose vectors $\mathbf{p} \in \Delta^{|V|-1}$ we index by $\mathbf{p}(v) \in \mathbb{R}$. An autoregressive language model generates text $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2, \dots)$ with $\mathbf{x}^t \in V$ sequentially, predicting each next token from a conditional distribution $p(\cdot \mid \mathbf{x}^{\le t}) \in \Delta^{|V|-1}$. We write the cross-entropy between a predicted distribution $\mathbf{p} \in \Delta^{|V|-1}$ and a target token $v \in V$ as $\mathsf{CE}(\mathbf{p}, v) \coloneqq -\log \mathbf{p}(v)$.

Looped language models. Looped language models compute a hidden state $\mathbf{z}^t \in \mathbb{R}^d$ for each token by recurrently applying a shared block of layers f_θ (Dehghani et al., 2019; Geiping et al., 2025; Zhu et al., 2025):

$$\mathbf{z}_r^t = f_\theta\big(\mathbf{z}_{r-1}^t;\, \mathbf{c}_r^{<t}\big), \qquad r = 1, \dots, R, \tag{1}$$

where $\mathbf{z}_0^t \in \mathbb{R}^d$ is computed from the token embedding $\mathbf{E}_{\mathbf{x}^t} \in \mathbb{R}^d$, and $\mathbf{c}_r^{<t}$ denotes the context at loop r, that is, the keys and values produced by earlier tokens in loop r. We refer to the context at the final loop, $\mathbf{c}_R^{<t}$, as the *deep prefix cache*. The shared block updates the hidden state at each step, $f_\theta(\cdot\,; \mathbf{c}_r^{<t}) : \mathbb{R}^d \to \mathbb{R}^d$, and an output head $g : \mathbb{R}^d \to \Delta^{|V|-1}$ predicts the next token from the final state, $p(\cdot \mid \mathbf{x}^{\le t}) = g(\mathbf{z}_R^t)$. We call each application of f_θ a loop, the R loops that produce a token a decoding step, and R the depth of the model. In practice, looped models are trained with a next-token loss not only at the final loop R, but also at earlier loops, either by randomizing depth (Geiping et al., 2025; McLeish et al., 2025) or by supervising every loop (Zhu et al., 2025):

$$\mathcal{L}_{\mathsf{loop}}(\theta) \coloneqq \mathbb{E}_{\mathbf{x}}\Big[\sum_t \sum_{r=1}^{R} \lambda_r\, \mathsf{CE}\big(g(\mathbf{z}_r^t),\, \mathbf{x}^{t+1}\big)\Big], \tag{2}$$

where $\lambda_r \ge 0$ weights the loss at loop r. While (2) trains every loop to output the immediate next token, no loop is trained to output the tokens that follow it. An early loop can therefore directly read out the next token, but not several future tokens at once, even though supervising early depths can encourage their hidden states to carry information about future tokens (Pal et al., 2023; Elhoushi et al., 2024), as we demonstrate in Section 4.2.

Speculative decoding. Speculative decoding (Leviathan et al., 2023; Chen et al., 2023) accelerates a target model p using a faster draft model q that proposes upcoming tokens. Given a prefix $\mathbf{x}^{\le t}$, the drafter proposes k draft tokens $\mathbf{d} = (\mathbf{d}^1, \dots, \mathbf{d}^k) \in V^k$, and the target model evaluates all of them in a single decoding step to obtain its own predictions. Under greedy decoding, this is done as:

$$\mathbf{y}^j = \mathsf{argmax}(\mathbf{p}^j), \qquad \mathbf{p}^j \coloneqq p(\cdot \mid \mathbf{x}^{\le t}, \mathbf{d}^{<j}), \qquad j = 1, \dots, k+1, \tag{3}$$

and the drafts are accepted up to the first mismatch $\mathbf{d}^j \neq \mathbf{y}^j$, where the target model substitutes its own prediction $\mathbf{y}^j$. Each step thus produces at least one and at most $k+1$ tokens while keeping the output identical to greedy decoding with p. Under sampling, each draft $\mathbf{d}^j$ is drawn from a draft distribution $\mathbf{q}^j$ and accepted with probability $\min\{1,\, \mathbf{p}^j(\mathbf{d}^j)/\mathbf{q}^j(\mathbf{d}^j)\}$. Upon the first rejection, a replacement token is sampled from the normalization of $\max\{0,\, \mathbf{p}^j - \mathbf{q}^j\}$, which leaves the output distribution of p unchanged (Leviathan et al., 2023). Because verifying $k+1$ tokens in parallel at small batch sizes takes nearly the same time as a single-token decoding step, the speedup grows with the number of accepted tokens τ generated per step and decreases with the relative cost of drafting.

3 Floodgates

In this section, we present Floodgates, a simple self-speculative decoding framework for looped language models. Algorithms 1 and 2 summarize training and decoding.

Algorithm 1 Training

Input: Looped block f_θ, head g, draft length k, weights $w_{1:k}$, learning rate η, initial adapters ϕ and mask $\mathbf{m}$

repeat
  $\mathbf{x} \sim \mathcal{D}$ ▷ self-generated
  $t \sim \mathrm{Unif}\{1, \dots, |\mathbf{x}|\}$
  run (1) on $\mathbf{x}^{<t}$ to obtain $\mathbf{c}_R^{\leq t}$ ▷ prefix cache
  slots $\leftarrow (\mathbf{z}_0^t, \mathbf{m}, \dots, \mathbf{m})$
  $(\hat{\mathbf{z}}^t, \dots, \hat{\mathbf{z}}^{t+k}) \leftarrow f_\theta(\text{slots}; \mathbf{c}_R^{\leq t})$ with ϕ ▷ draft
  $\mathbf{q}^j \leftarrow g(\hat{\mathbf{z}}^{t+j})$ for $j \leq k$
  $\mathcal{L} \leftarrow \sum_{j=1}^{k} w_j \, \mathsf{CE}(\mathbf{q}^j, \mathbf{x}^{t+j})$ ▷ (5)
  $(\phi, \mathbf{m}) \leftarrow (\phi, \mathbf{m}) - \eta \nabla_{\phi,\mathbf{m}} \mathcal{L}$ ▷ θ frozen
until converged
return $(\phi, \mathbf{m})$

Algorithm 2 Decoding step

Input: Prefix $\mathbf{x}^{\leq t}$, deep prefix cache $\mathbf{c}_R^{\leq t}$, model (f_θ, g), drafter $(\phi, \mathbf{m})$, draft length k, elastic draft budget E, window W

slots $\leftarrow (\mathbf{z}_0^t, \mathbf{m}, \dots, \mathbf{m})$
$(\hat{\mathbf{z}}^t, \dots, \hat{\mathbf{z}}^{t+k}) \leftarrow f_\theta(\text{slots}; \mathbf{c}_R^{\leq t})$ with ϕ ▷ draft
$\mathbf{q}^j \leftarrow g(\hat{\mathbf{z}}^{t+j})$; $\mathbf{d}^j \leftarrow \mathsf{argmax}(\mathbf{q}^j)$ for $j \leq k$
$\mathbf{d}^j \leftarrow \mathbf{d}^k$ and $\hat{\mathbf{z}}^{t+j} \leftarrow \mathbf{z}_0^{t+j}$ for $k < j \leq W$ ▷ extend
for $e = 1, \dots, E$ **do**
  $(\hat{\mathbf{z}}^t, \dots, \hat{\mathbf{z}}^{t+W}) \leftarrow f_\theta\big((\hat{\mathbf{z}}^{t+j}, \mathbf{E}_{\mathbf{d}^j})_{j=0}^{W}; \mathbf{c}_R^{\leq t}\big)$ ▷ (6)
  $\mathbf{d}^j \leftarrow \mathsf{argmax}(g(\hat{\mathbf{z}}^{t+j}))$ for $j \leq W$
end for
$\mathbf{p}^j \leftarrow p(\cdot \mid \mathbf{x}^{\leq t}, \mathbf{d}^{<j})$ for $j \leq W+1$ ▷ verify
$\mathbf{y}^j \leftarrow \mathsf{argmax}(\mathbf{p}^j)$ for $j \leq W+1$
$a \leftarrow \max\{\, j : \mathbf{d}^1 = \mathbf{y}^1, \dots, \mathbf{d}^j = \mathbf{y}^j \,\}$ ▷ accept
return $(\mathbf{x}^{\leq t}, \mathbf{d}^1, \dots, \mathbf{d}^a, \mathbf{y}^{a+1})$

Drafting. Let $\mathbf{x}^t$ be the latest generated token and $\mathbf{m} \in \mathbb{R}^d$ be a learnable vector. To draft k tokens in a single loop iteration, we append k *mask tokens* after $\mathbf{x}^t$ by repeating $\mathbf{m}$, forming $k+1$ positions that we call slots, and run the shared block f_θ once over all slots in parallel while attending to the *deep prefix cache* $\mathbf{c}_R^{\leq t}$. We then read out the resulting slot states with the head g to obtain the draft distributions $\mathbf{q}^j$:

$$(\hat{\mathbf{z}}^t, \hat{\mathbf{z}}^{t+1}, \dots, \hat{\mathbf{z}}^{t+k}) = f_\theta\big((\mathbf{z}_0^t, \mathbf{m}, \dots, \mathbf{m}); \mathbf{c}_R^{\leq t}\big), \qquad \mathbf{q}^j = g(\hat{\mathbf{z}}^{t+j}), \quad j = 1, \dots, k, \tag{4}$$

where $\hat{\mathbf{z}}^{t+j} \in \mathbb{R}^d$ is the hidden state of slot j, and $\mathbf{d}^j = \mathsf{argmax}(\mathbf{q}^j)$ under greedy decoding, for instance. While mask tokens have been used for parallel drafting with dense (non-looped) models (Monea et al., 2023; Bhendawade et al., 2024; Samragh et al., 2025; Sahoo et al., 2026), a dense model must pass the mask tokens through all of its layers. In our approach, the slots pass through the shared block only once, using all parameters of the model at the cost of a single loop, and they attend to the deep prefix cache $\mathbf{c}_R^{\leq t}$ instead of the first-loop cache $\mathbf{c}_1^{<t}$. Here, $\mathbf{c}_R^{\leq t}$ provides the model's full-depth representation of all earlier tokens $\mathbf{x}^{<t}$, which is already in the cache from verifying tokens from previous steps. To teach a single loop to predict these tokens without changing the verifier, we use lightweight adapters that are active only while drafting.

Training. Our insight is that all-depth training (2) already encourages the shared block f_θ to build features that are predictive of several future tokens in its early loops (see Section 4.2). During pretraining, however, the first loop is trained to output only the immediate next token and attends only to the first-loop cache $\mathbf{c}_1^{<t}$. Therefore, our approach further learns to read the deep prefix cache $\mathbf{c}_R^{\leq t}$ and decode those future tokens into the mask slots. To this end, we train the slots with a cross-entropy objective applied j tokens ahead:

$$\mathcal{L}(\phi, \mathbf{m}) = \mathbb{E}_{\mathbf{x},t}\Big[\sum_{j=1}^{k} w_j \, \mathsf{CE}\big(g(\hat{\mathbf{z}}^{t+j}), \mathbf{x}^{t+j}\big)\Big], \tag{5}$$

where $\mathbf{x} \sim \mathcal{D}$ is a sequence generated by the target model itself and t is a sampled position in its continuation after which the slots are placed. We train on the model's own generations (Cai et al., 2024) because verification (3) checks drafts against the model predictions. Because a draft at slot j can be accepted only if all earlier drafts $< j$ are also accepted, we set exponentially decaying weights $w_j \propto e^{-(j-1)/4}$ normalized to sum to one to incentivize the drafter to prioritize accuracy on earlier positions. We parameterize the drafter with low-rank adapters ϕ (Hu et al., 2022) that are active only at the slots (Samragh et al., 2025), keeping the base model weights θ and output head g frozen.

Elastic drafting. When a looped model runs many loops, either by design or through test-time scaling, a single loop may be too shallow to draft a long continuation that the full model will accept. We find that the pretrained loop f_θ can be applied out of the box to iteratively refine the initial draft over a window of $W \geq k$ positions. Specifically, each position $j \in \{0, \dots, W\}$ in the window maintains a hidden feature $\hat{\mathbf{z}}^{t+j} \in \mathbb{R}^d$ and a

draft token $\mathbf{d}^j$, with $\mathbf{d}^0 = \mathbf{x}^t$. Before the first pass, the first k positions ($1 \le j \le k$) take their features $\hat{\mathbf{z}}^{t+j}$ and drafted tokens $\mathbf{d}^j$ directly from the single-loop draft of (4), while the remaining $W - k$ extended positions ($k < j \le W$) repeat the last drafted token $\mathbf{d}^j \leftarrow \mathbf{d}^k$ and initialize $\hat{\mathbf{z}}^{t+j}$ to its initial state $\mathbf{z}_0^{t+j}$. We then run $E \ge 0$ elastic passes that jointly update the features and draft tokens across the window:

$$(\hat{\mathbf{z}}^t, \dots, \hat{\mathbf{z}}^{t+W}) \leftarrow f_\theta\big((\hat{\mathbf{z}}^{t+j}, \mathbf{E}_{\mathbf{d}^j})_{j=0}^{W};\, \mathbf{c}_R^{<t}\big), \qquad \mathbf{d}^j \leftarrow \mathsf{argmax}\big(g(\hat{\mathbf{z}}^{t+j})\big), \tag{6}$$

where $\mathbf{E}_{\mathbf{d}^j} \in \mathbb{R}^d$ is the token embedding of $\mathbf{d}^j$, combined with the feature $\hat{\mathbf{z}}^{t+j}$ through the model's input projection so that each pass revises both the features and the draft tokens using the latest tokens in the window. These passes do not use the adapters ϕ and therefore require no extra training. While this update resembles Jacobi decoding (Santilli et al., 2023; Fu et al., 2024) and DAS-Wave (Li et al., 2026), which refine uninitialized or repeated tokens from scratch, elastic drafting starts from the trained multi-token draft of (4) and therefore requires far fewer passes.

Iterating (6) creates a direct tradeoff between drafting cost and acceptance, since running more passes E improves the draft at the cost of one additional loop per pass. Because each pass updates all W draft tokens in parallel, running E passes costs only E/W effective loops per position, which is much cheaper than the R loops of tokenwise generation and yields higher speedup when E and W are chosen adaptively to the verifier depth R. On shallow models with, e.g., $R = 4$, where a single loop already costs a quarter of a decoding step, single-pass drafting with $E = 0$ is fastest, whereas on deep models with, e.g., $R = 32$, where each loop is a small fraction of a step, additional passes $E > 0$ pay off. We treat E as a hyperparameter and select the value that maximizes speedup on held-out problems for each model (Appendix B.1), following standard practice for choosing the draft block size in speculative decoding (Leviathan et al., 2023; Chen et al., 2024).

Concurrent serving. Because an untrained shallow loop only predicts the immediate next token, prior self-speculative methods for looped models (Cho et al., 2026; Ha and Kim, 2026) rely on *depth pipelining*, starting the next token at the first loop before earlier tokens have finished their R loops. While pipelining parallelizes drafting and verification, it batches together tokens at different loop depths, say $B \subseteq [R]$ (e.g., one every four loops), and therefore must store and read the prefix cache $\mathbf{c}_r^{<t}$ of all such loops $r \in B$. When serving many concurrent requests, maintaining and reading B separate caches per request saturates runtime memory bandwidth, causing diminishing return in throughput. By drafting a block of future tokens in parallel from a single loop, Floodgates avoids pipelining altogether. All drafted slots of a sequence attend only to the deep prefix cache $\mathbf{c}_R^{<t}$, reading the keys and values of past tokens once per step for the entire draft block.

At batch size one, decoding is memory-bound, so verifying all k draft tokens in parallel takes nearly the same time as verifying a single token. As the batch size B grows, verification becomes compute-bound and its runtime increases with the total number of verified drafts, so checking all kB candidates across the batch wastes compute on low-confidence tokens at later positions (Sahoo et al., 2026). To keep throughput high across batch sizes, we use single-pass drafting ($E = 0$) and verify only the most promising drafts in the batch. Specifically, let $\alpha_{b,j}$ be the probability $\mathbf{q}^j(\mathbf{d}^j)$ assigned by the drafter to slot j of sequence b. Because slot j is accepted only if all earlier slots $i \le j$ in sequence b are also accepted, we estimate its acceptance probability as $\pi_{b,j} = \prod_{i=1}^{j} \alpha_{b,i}$. Given a total budget of M drafts across the batch, we select the subset $\mathcal{S}_M$ of M slots with the largest $\pi_{b,j}$; verifying $\mathcal{S}_M$ takes time $T(M)$ and yields $B + \sum_{(b,j)\in\mathcal{S}_M} \pi_{b,j}$ tokens in expectation (B guaranteed tokens from verification plus the expected accepted drafts). At each step, we therefore choose the budget $M^\star$ that maximizes expected throughput:

$$M^\star = \underset{M \in \{B, 2B, \dots, kB\}}{\mathsf{argmax}} \frac{B + \sum_{(b,j)\in\mathcal{S}_M} \pi_{b,j}}{T(M)}, \tag{7}$$

where the verification time $T(M)$ is profiled once at initialization for multiples M of B. Since $\pi_{b,j}$ decreases with j, the selected slots form a prefix of the draft block of each sequence. The remaining drafts are discarded without running further loops, and each sequence drafts a new block after its last verified token at the next step.

4 Experiments

Through experiments, we aim to answer the following core questions:

Table 1 Floodgates is the fastest lossless method on every model and task under greedy decoding, and in 16 of 24 settings under sampling. We report wall-clock speedup (×) over plain decoding on an A100 at batch size one, with accepted tokens per step τ in parentheses. Sampling uses temperature 1.0 and top-p 0.7. We bold the fastest lossless method and underline the second fastest. LoopSpec* uses one proposal (Cho et al., 2026), WaveFront (Ha and Kim, 2026) and DAS/DAS-Wave (Li et al., 2026) are training-free looped baselines, JD denotes Jacobi decoding (Santilli et al., 2023), and PLD denotes prompt lookup decoding (Saxena, 2023).

		GSM8K		MATH-500		HumanEval+		MBPP+		BBH	
Model	Method	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling
Ouro 1.4B	LoopSpec*	2.51 (3.1)	2.10 (2.8)	2.60 (3.2)	2.34 (3.0)	3.01 (3.7)	2.67 (3.5)	2.62 (3.3)	2.23 (3.0)	2.67 (3.3)	2.25 (2.9)
	LoopSpec	2.77 (3.5)	2.40 (3.3)	2.81 (3.5)	2.53 (3.4)	3.10 (3.9)	2.77 (3.8)	2.83 (3.6)	2.51 (3.5)	2.86 (3.6)	2.46 (3.4)
	WaveFront	2.34 (3.1)	1.85 (2.4)	2.47 (3.2)	2.10 (2.7)	2.77 (3.7)	2.31 (3.0)	2.38 (3.2)	2.00 (2.6)	2.52 (3.4)	1.96 (2.5)
	DAS	1.91 (5.8)	1.74 (5.3)	2.07 (6.4)	1.97 (6.0)	2.59 (7.9)	2.40 (7.3)	2.06 (6.3)	1.83 (5.6)	2.16 (6.6)	2.07 (6.3)
	DAS-Wave	0.44 (1.8)	0.54 (1.9)	0.44 (1.8)	0.53 (1.9)	0.43 (1.8)	0.52 (1.8)	0.44 (1.8)	0.50 (1.8)	0.42 (1.7)	0.50 (1.8)
	JD	0.95 (3.3)	0.83 (4.6)	0.94 (5.3)	0.87 (4.9)	0.90 (5.1)	0.87 (4.9)	0.91 (3.2)	0.78 (4.4)	0.93 (5.3)	0.90 (5.0)
	PLD	2.08 (2.0)	1.97 (1.9)	1.96 (2.0)	1.86 (1.9)	3.27 (3.3)	2.91 (2.9)	1.37 (1.4)	1.33 (1.3)	3.38 (3.4)	3.23 (3.2)
	Floodgates	**5.04** (7.5)	**4.16** (6.3)	**4.69** (7.0)	**4.41** (6.7)	**4.83** (7.2)	**4.49** (6.8)	**3.01** (4.6)	**2.75** (4.3)	**4.40** (6.6)	**4.11** (6.2)
Ouro 2.6B	LoopSpec*	2.86 (3.4)	2.68 (3.3)	2.90 (3.4)	2.79 (3.4)	3.28 (3.8)	3.03 (3.6)	2.94 (3.5)	2.79 (3.3)	2.95 (3.5)	2.65 (3.2)
	LoopSpec	3.06 (3.6)	2.89 (3.6)	3.10 (3.7)	2.96 (3.6)	3.33 (3.9)	3.14 (3.8)	3.09 (3.7)	2.95 (3.7)	3.10 (3.7)	2.86 (3.6)
	WaveFront	2.53 (3.3)	2.01 (2.6)	2.50 (3.3)	2.20 (2.9)	2.83 (3.8)	2.32 (3.0)	2.51 (3.3)	2.18 (2.9)	2.61 (3.5)	2.13 (2.8)
	DAS	2.24 (6.7)	2.04 (6.1)	2.35 (7.1)	2.22 (6.7)	2.71 (8.2)	2.47 (7.4)	2.19 (6.6)	1.96 (5.9)	2.36 (7.1)	2.28 (6.8)
	DAS-Wave	0.55 (2.2)	0.66 (2.3)	0.49 (2.0)	0.56 (2.0)	0.51 (2.0)	0.59 (2.1)	0.50 (2.1)	0.58 (2.1)	0.49 (2.0)	0.59 (2.1)
	JD	0.95 (3.2)	0.81 (4.5)	0.94 (5.3)	0.87 (4.9)	0.90 (5.1)	0.85 (4.8)	0.86 (3.0)	0.75 (4.3)	0.91 (5.2)	0.88 (5.0)
	PLD	2.04 (2.0)	1.94 (1.9)	2.79 (2.7)	2.75 (2.7)	3.17 (3.1)	2.55 (2.5)	1.31 (1.3)	1.34 (1.3)	3.40 (3.4)	3.28 (3.3)
	Floodgates	**5.56** (8.2)	**4.50** (6.7)	**6.04** (9.0)	**4.90** (7.3)	**5.57** (8.3)	**4.82** (7.2)	**3.39** (5.2)	**3.23** (4.9)	**5.07** (7.5)	**4.77** (7.1)
Recurrent Llama 3.2	LoopSpec*	3.27 (7.7)	2.98 (7.5)	5.03 (11.3)	3.86 (9.3)	5.14 (11.7)	4.10 (10.0)	3.91 (9.1)	3.79 (9.2)	4.27 (10.1)	3.86 (9.6)
	LoopSpec	4.61 (12.2)	**3.80** (12.3)	6.11 (14.5)	**4.73** (13.5)	6.03 (14.5)	**4.81** (13.9)	5.41 (13.5)	**4.63** (13.4)	5.65 (13.8)	4.77 (13.7)
	DAS	3.82 (7.5)	2.77 (5.4)	4.30 (8.4)	3.45 (6.7)	4.32 (8.4)	3.58 (7.0)	4.01 (7.8)	3.38 (6.6)	4.22 (8.2)	3.83 (7.5)
	DAS-Wave	4.56 (17.9)	2.78 (6.1)	5.98 (23.6)	3.71 (14.7)	5.92 (23.1)	3.95 (15.4)	5.19 (20.3)	3.52 (13.7)	5.84 (23.1)	4.63 (18.2)
	JD	0.76 (4.9)	0.68 (2.6)	0.80 (3.2)	0.71 (2.8)	0.80 (3.1)	0.74 (2.9)	0.78 (3.1)	0.74 (2.9)	0.81 (3.2)	0.78 (3.1)
	PLD	1.83 (1.9)	1.55 (1.6)	3.74 (4.0)	2.66 (2.8)	3.54 (3.7)	3.03 (3.2)	2.23 (2.3)	2.31 (2.4)	3.48 (3.7)	3.24 (3.5)
	Floodgates	**5.04** (20.4)	3.28 (7.6)	**6.97** (28.5)	4.66 (11.0)	**6.91** (28.0)	4.69 (10.9)	**5.84** (23.5)	4.34 (10.1)	**6.62** (27.1)	**5.13** (12.1)
Huginn 3.5B	LoopSpec*	1.35 (2.4)	1.14 (2.1)	2.54 (5.3)	1.44 (2.9)	2.20 (3.9)	1.76 (3.2)	1.87 (3.2)	1.69 (3.0)	1.48 (3.0)	1.20 (2.5)
	LoopSpec	2.20 (5.4)	1.64 (4.7)	3.07 (9.0)	2.16 (7.4)	3.87 (9.0)	2.89 (7.8)	3.21 (7.3)	2.62 (6.7)	2.02 (7.2)	1.43 (6.3)
	WaveFront	2.29 (3.1)	1.62 (2.5)	2.99 (5.3)	1.76 (3.0)	3.53 (5.8)	2.27 (3.6)	2.85 (4.6)	2.03 (3.2)	2.55 (4.4)	1.96 (3.3)
	DAS	2.74 (6.2)	2.20 (5.0)	3.47 (7.9)	2.88 (6.6)	3.59 (8.2)	3.15 (7.2)	3.05 (6.9)	2.57 (5.9)	3.38 (7.7)	3.06 (7.0)
	DAS-Wave	4.23 (17.6)	2.57 (5.7)	5.38 (22.8)	3.57 (15.0)	5.45 (22.7)	4.15 (17.3)	4.57 (19.0)	3.02 (6.7)	5.20 (22.0)	4.22 (17.8)
	JD	0.83 (3.1)	0.74 (2.8)	0.82 (3.1)	0.77 (3.0)	0.80 (3.1)	0.80 (3.0)	0.81 (3.1)	0.74 (2.8)	0.81 (3.1)	0.78 (3.0)
	PLD	2.11 (2.2)	1.88 (1.9)	4.16 (4.3)	3.51 (3.7)	4.77 (4.9)	3.51 (3.6)	2.26 (2.3)	2.08 (2.1)	3.43 (3.6)	3.28 (3.4)
	Floodgates	**4.48** (19.0)	**3.16** (7.3)	**6.21** (26.6)	**4.60** (10.8)	**6.18** (26.2)	**4.70** (10.9)	**5.06** (21.4)	**3.52** (8.1)	**5.70** (24.6)	**4.48** (19.3)
Recurrent TinyLlama	LoopSpec*	3.92 (7.7)	3.47 (7.2)	5.39 (10.2)	4.70 (9.2)	6.25 (11.8)	4.71 (9.5)	4.97 (9.7)	4.38 (8.9)	prompts exceed max context length	
	LoopSpec	5.79 (12.6)	**4.95** (12.4)	6.92 (13.9)	**5.90** (13.3)	7.35 (14.7)	**6.13** (13.9)	6.48 (13.7)	**5.89** (13.6)		
	DAS	4.20 (7.6)	3.06 (5.5)	4.53 (8.2)	3.57 (6.4)	4.66 (8.4)	3.93 (7.1)	4.46 (8.0)	3.75 (6.8)		
	DAS-Wave	6.07 (19.2)	3.40 (6.5)	6.86 (22.1)	3.82 (12.2)	7.10 (22.6)	4.81 (15.2)	6.43 (20.4)	4.13 (13.1)		
	JD	0.87 (3.2)	0.77 (2.9)	0.86 (3.3)	0.77 (2.9)	0.84 (3.1)	0.79 (2.9)	0.85 (3.2)	0.80 (3.0)		
	PLD	1.98 (2.1)	1.77 (1.9)	2.76 (2.9)	2.34 (2.5)	3.86 (4.1)	2.85 (3.0)	2.14 (2.2)	2.29 (2.4)		
	Floodgates	**6.73** (22.2)	4.01 (8.1)	**7.79** (26.0)	4.83 (9.9)	**8.20** (27.0)	5.61 (11.4)	**7.17** (23.5)	5.25 (10.6)		

Q1. Can Floodgates accelerate looped decoding across models and decoding regimes (Section 4.1)?

Q2. How do all-depth supervision and deep prefix conditioning contribute to multi-token drafting (Section 4.2)?

Q3. Does elastic drafting adapt its depth to maximize performance (Section 4.3)?

Q4. Do the decoding gains persist under high-throughput batched serving (Section 4.4)?

The dataset, hyperparameters, and hardware configurations for each model can be found in Appendix A.2.

4.1 Accelerating looped decoding across models and regimes (Q1)

Setup. We evaluate Floodgates on seven looped models from two families. Ouro 1.4B and Ouro 2.6B (Zhu et al., 2025), together with their reasoning variants Ouro 1.4B Thinking and Ouro 2.6B Thinking, loop their entire stack of layers $R = 4$ times. Recurrent Llama 3.2, Recurrent TinyLlama (McLeish et al., 2025), and

Table 2 On the reasoning models, Floodgates is the fastest method on both tasks under greedy decoding and sampling. We report wall-clock speedup (×) over plain decoding on an A100 at batch size one, with accepted tokens per step τ in parentheses. Sampling uses temperature 0.6 and top-p 0.95. We bold the fastest lossless method and underline the second fastest.

		GSM8K		MATH-500	
Model	Method	Greedy	Sampling	Greedy	Sampling
Ouro 1.4B Thinking	LoopSpec	<u>2.72</u> (3.3)	<u>2.40</u> (3.3)	<u>2.74</u> (3.5)	<u>2.47</u> (3.5)
	WaveFront	2.24 (2.9)	2.13 (2.8)	2.31 (3.2)	2.18 (3.0)
	DAS	1.76 (5.5)	1.65 (5.1)	1.97 (6.0)	1.80 (5.5)
	DAS-Wave	0.47 (1.8)	0.47 (1.8)	0.48 (1.9)	0.46 (1.8)
	JD	0.86 (5.3)	0.77 (4.7)	0.85 (5.3)	0.76 (4.8)
	PLD	1.65 (1.7)	1.59 (1.6)	2.29 (2.3)	1.81 (1.9)
	Floodgates	**3.70** (5.6)	**3.46** (5.2)	**3.68** (5.5)	**3.23** (4.9)
Ouro 2.6B Thinking	LoopSpec	<u>3.09</u> (3.5)	<u>2.87</u> (3.6)	<u>2.98</u> (3.6)	<u>2.85</u> (3.6)
	WaveFront	2.42 (3.1)	2.39 (3.1)	2.42 (3.3)	2.36 (3.2)
	DAS	2.05 (6.2)	1.89 (5.8)	2.23 (6.8)	2.02 (6.1)
	DAS-Wave	0.57 (2.2)	0.56 (2.1)	0.51 (2.0)	0.50 (2.0)
	JD	0.86 (5.3)	0.76 (4.7)	0.87 (5.5)	0.77 (4.8)
	PLD	1.62 (1.7)	1.60 (1.6)	2.03 (2.1)	1.84 (1.9)
	Floodgates	**3.90** (5.9)	**3.71** (5.5)	**3.72** (5.5)	**3.41** (5.1)

Huginn 3.5B (Geiping et al., 2025) run their initial layers ("prelude") once, loop their middle layers ("core") $R = 32$ times by default, and run their final layers ("coda") once. Because these three models were trained with randomized loop counts, they can be run at any depth at inference time. We evaluate on five standard benchmarks spanning mathematical reasoning on GSM8K (Cobbe et al., 2021) and MATH-500 (Hendrycks et al., 2021; Lightman et al., 2024), code generation on HumanEval+ and MBPP+ (Chen et al., 2021; Austin et al., 2021; Liu et al., 2023), and general reasoning on BBH (Suzgun et al., 2023), under both greedy decoding and sampling following the protocol of LoopSpec (Cho et al., 2026). Using 100 held-out problems, we set the elastic drafting budget E as a single loop iteration, $E = 0$, on 4-loop Ouro models, and $E = 24$ on the remaining 32-loop models. We leave a comparison of $E = 0$ and 24 across every model and task in Appendix B.1 for reference.

We compare Floodgates with existing training-free self-speculative decoding methods for looped models, namely LoopSpec (Cho et al., 2026), WaveFront (Ha and Kim, 2026), DAS, and DAS-Wave (Li et al., 2026), as well as Jacobi decoding (JD) (Santilli et al., 2023; Fu et al., 2024), and prompt lookup decoding (PLD) (Saxena, 2023). We report wall-clock speedup over plain decoding and the average number of accepted tokens per step τ on one NVIDIA A100-80GB GPU at batch size one. Full experimental details are provided in Appendix A.

Results. Tables 1 and 2 show the results. Under greedy decoding, Floodgates is the fastest lossless method on every model and task, with the largest margin on the shallow, 4-loop Ouro models (Ouro 1.4B and Ouro 2.6B). There, Floodgates drafts with a single loop iteration without elastic passes (Section 3) to accept 4.6–9.0 tokens per step, decoding 3.0–6.0× faster than plain decoding while no prior method exceeds 3.4×. On the deeper, 32-loop models (Recurrent Llama 3.2, Huginn 3.5B, and Recurrent TinyLlama), Floodgates uses elastic drafting with 24 elastic passes (Section 3) to accept between 19 and 29 tokens per step, which reaches 4.5–8.2× speedup and outperforms the strongest training-free baselines, LoopSpec and DAS-Wave, by 6–15% on every task.

Under sampling, stochastic verification at each draft position reduces the number of accepted tokens per step for all block drafters. Nevertheless, in Table 1, Floodgates remains the fastest lossless method across all five benchmarks on Ouro 1.4B (2.8–4.5×), Ouro 2.6B (3.2–4.9×), and Huginn 3.5B (3.2–4.7×), as well as on BBH for Recurrent Llama 3.2 (5.1×), while LoopSpec is faster under sampling on Recurrent TinyLlama and four tasks of Recurrent Llama 3.2. On the long-horizon reasoning models Ouro 1.4B Thinking and Ouro 2.6B Thinking (Table 2), Floodgates achieves the highest speedup on every task under both greedy decoding (3.7–3.9× vs. 2.7–3.1× for LoopSpec) and sampling (3.2–3.7× vs. 2.4–2.9×). In Appendix B, we provide confidence intervals for all speedups, verify that decoding is lossless through task accuracy measurements, and further compare Floodgates against lossy acceleration methods.

4.2 Contribution of all-depth supervision and deep prefix conditioning (Q2)

In Sections 1 and 3, we hypothesized that a shallow loop of a looped model is a strong multi-token drafter because all-depth supervision (2) encourages shallow states to encode future tokens while sharing all model parameters, and because conditioning on the deep prefix cache $\mathbf{c}_R^{<t}$ provides full-depth context representations. We now test both factors through controlled comparisons.

All-depth supervision and full-parameter sharing. We compare a single loop iteration of Floodgates ($E = 0$) against separate draft networks and early layers of non-looped models trained on the same data with the same objective (5), evaluated in a shared PyTorch implementation at batch size one.

First, on Ouro 1.4B (Table 3), a single loop of Floodgates runs 25% of the layers of a decoding step, yet its higher acceptance (7.0–7.6 tokens per step) yields a 5.5× speedup on both GSM8K and MATH-500. This outperforms the 5-layer DFlash block drafter (Chen et al., 2026a), which reaches 4.3× and 4.1× despite having 4.5× more trained parameters than our adapters (278M vs. 61M), as well as 1-layer and 4-layer EAGLE-style separate drafters (Li et al., 2024a) at 2.6–4.1×. Even though a separate drafter runs fewer layers per pass, its small network has limited capacity and must learn representations from scratch, whereas adapting a single loop leverages all pretrained parameters of the target model to achieve a much higher acceptance with far fewer trained parameters.

Table 3 Floodgates gives larger speedups than a separate drafter. Speedup (×) over plain decoding, with accepted tokens per step in parentheses. Params. is the number of trained parameters, and depth is the number of layer evaluations in one drafting pass relative to a full decoding step.

Drafter	Params.	Depth	GSM8K	MATH-500
Separate, 1 layer	56M	1%	2.55 (2.9)	2.85 (2.9)
Separate, 4 layers	210M	4%	3.70 (4.5)	4.09 (4.4)
Separate, DFlash	278M	5%	4.32 (5.1)	4.05 (4.6)
Floodgates	61M	25%	**5.48** (7.6)	**5.54** (7.0)

Second, to separate the effect of early-depth supervision from that of looped parameter sharing, we compare Recurrent Llama 3.2 against the 16-layer non-looped model it was built from, Llama 3.2 1B, and against LayerSkip Llama 3.2 1B (Elhoushi et al., 2024), which finetunes Llama 3.2 1B with next-token losses at early layers (Figure 3). On both non-looped models, we apply our method to the first $L_d \in \{1, \ldots, 16\}$ layers by training LoRA adapters on those L_d layers to draft k mask tokens in parallel, which are then verified by all 16 layers. By comparing Llama 3.2 1B and LayerSkip, we isolate the benefit of supervising early depths. When drafting with only the first layer ($L_d = 1$, which matches the 7% relative depth of one loop of Recurrent Llama 3.2), Llama 3.2 1B reaches only 1.6–1.8× speedup on GSM8K and MATH-500, whereas early-exit supervision in LayerSkip raises the speedup of our drafter to 1.9× on the first layer and from 2.5–3.0× to 3.5–4.0× on the first four layers ($L_d = 4$). Consistent with this gain, training linear probes (Pal et al., 2023) on each layer of both non-looped models to predict the next three tokens shows that early-exit supervision improves accuracy at all three future positions across all five benchmarks (Appendix B.5, Figure 8). Even so, the non-looped drafters peak at 3.5–4.3× with half or all 16 layers, whereas a single loop of Recurrent Llama 3.2 reaches 4.9× and 5.8× while running just 7% of the layers of a decoding step. This is because the first L_d layers of a non-looped model use only $L_d/16$ of the parameters and attend only to their own shallow prefix states, whereas one loop of Recurrent Llama 3.2 uses all parameters of the shared block and attends to the deep prefix cache $\mathbf{c}_R^{<t}$.

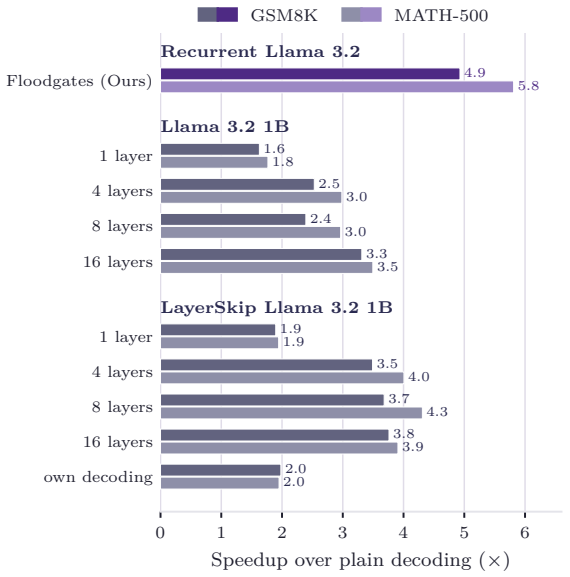

Figure 3 Floodgates gives larger speedups than any number of layers of a non-looped model. Speedup over plain decoding on GSM8K and MATH-500. On Llama 3.2 1B and LayerSkip Llama 3.2 1B, we train our drafter on the first 1 to 16 layers, and "own decoding" is the native self-speculative decoding of LayerSkip.

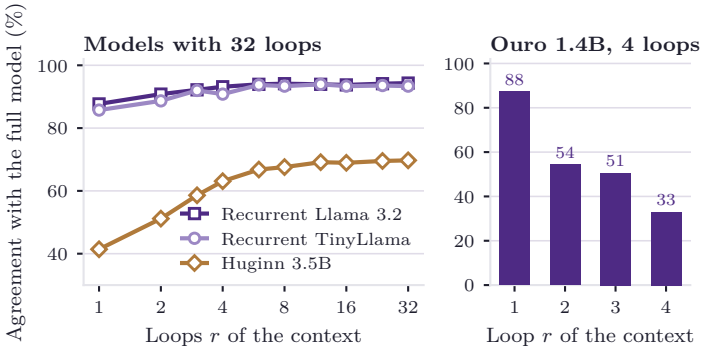

Figure 4 Without training, a deeper prefix cache helps one loop on the 32-loop models but not on Ouro 1.4B. Next-token agreement on GSM8K between the full model and one untrained loop that attends to the prefix cache after r loops.

Benchmark	Final loop $\mathbf{c}_R^{<t}$	First loop $\mathbf{c}_1^{<t}$
Recurrent Llama 3.2		
GSM8K	**4.93** (5.9)	4.35 (5.2)
MATH-500	**6.47** (7.8)	5.89 (7.1)
HumanEval+	**5.37** (6.5)	4.61 (5.5)
MBPP+	**4.53** (5.4)	4.01 (4.8)
BBH	**4.47** (5.4)	3.39 (4.1)
Ouro 1.4B		
GSM8K	**5.04** (7.5)	4.42 (6.6)
MATH-500	**4.69** (7.0)	4.05 (6.1)
HumanEval+	**4.83** (7.2)	4.33 (6.5)
MBPP+	**3.01** (4.6)	2.70 (4.1)
BBH	**4.40** (6.6)	3.62 (5.4)

Table 4 After training, the final-loop cache gives larger speedups on both models. Speedup ($\times$) over plain decoding with a single pass ($E = 0$), with τ in parentheses. The two drafters of each model differ only in the cache they read.

Conditioning on the deep prefix cache. During pretraining, each loop attends to the context from the same loop depth, so the first loop only sees the shallowest states of earlier tokens ($\mathbf{c}_1^{<t}$) and never sees the final-loop cache ($\mathbf{c}_R^{<t}$). To see how cache depth affects the first loop before any training, we run one untrained loop over the current token with the prefix cache after r loops and measure how often its predicted next token agrees with the full model (Figure 4). On the 32-loop models, moving from the first-loop cache ($r = 1$) to the final-loop cache ($r = 32$) raises agreement from 86–88% to 93–94% on Recurrent Llama 3.2 and Recurrent TinyLlama, and from 41% to 70% on Huginn 3.5B. On Ouro 1.4B, however, agreement drops from 88% at the first-loop cache ($r = 1$) to 33% at the final-loop cache ($r = 4$), because Ouro loops its entire layer stack and its first loop cannot directly read the final-loop cache without adaptation. To test whether the deep prefix cache is still better once the drafter is trained, we train drafters under identical settings conditioned on either $\mathbf{c}_R^{<t}$ or $\mathbf{c}_1^{<t}$ (Table 4). Once adapted, Ouro 1.4B reverses the untrained trend and accepts 4.6–7.5 tokens per step from $\mathbf{c}_R^{<t}$ compared to 4.1–6.6 from $\mathbf{c}_1^{<t}$, decoding 11–22% faster across all five benchmarks, while Recurrent Llama 3.2 decodes 10–32% faster with $\mathbf{c}_R^{<t}$ than with $\mathbf{c}_1^{<t}$. Together, these results confirm that both the deep prefix cache and adapter training are essential for fast drafting.

4.3 Elastic drafting across verifier depths (Q3)

A key advantage of looped models is that their compute can be scaled at test time by varying the number of loops R without retraining (Geiping et al., 2025; Suleymanzade et al., 2026b). In Section 3, we showed that because each elastic pass of Floodgates costs a single loop iteration ($1/R$ of a decoding step) and requires no extra parameters, a single trained drafter can adjust its elastic drafting budget E to the depth of the verifier. We now test how the optimal budget E changes with verifier depth and how Floodgates compares with prior methods across test-time depths.

Elastic drafting budget. Figure 5 shows how speedup varies with the number of elastic passes E as the verifier depth of Recurrent Llama 3.2 scales from 4 to 64 loops, using a single drafter trained at 32 loops. On shallower verifiers running up to 16 loops, where each loop is a large fraction of a decoding step, single-pass drafting ($E = 0$) is fastest and additional passes reduce speedup, whereas for deeper verifiers running 32 and 64 loops, the speedup rises up to 16–24 elastic passes before declining. At every depth and budget $E \leq 24$, Floodgates outperforms DAS-Wave, which runs the same loops starting from copies of the last token. The gap is largest when only a few elastic passes are run (63–89% higher speedup at 4 passes, 37–49% at 8 passes, and 16–21% at 16 passes), because Floodgates starts its elastic passes from a trained multi-token draft (Appendix B.6 shows per-position acceptance rates across E).

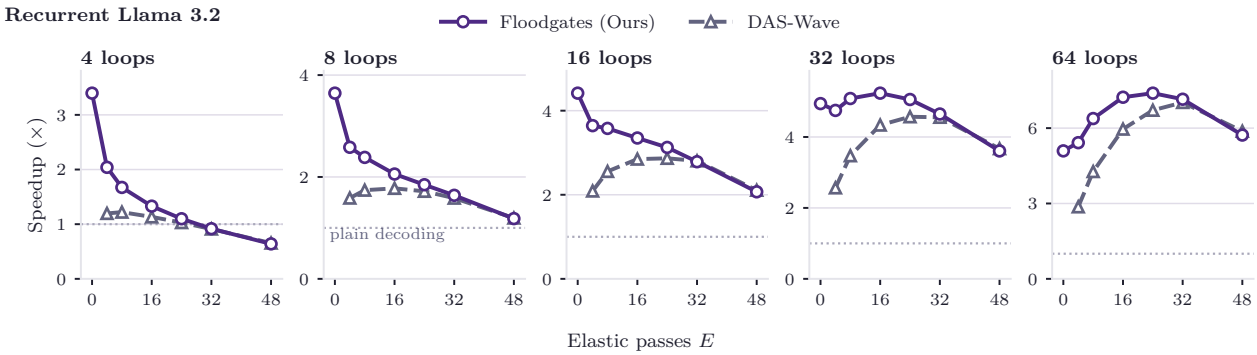

Figure 5 Elastic passes increase the speedup only on deep verifiers. Speedup over plain decoding on GSM8K against the elastic drafting budget E. Recurrent Llama 3.2 runs at 4 to 64 loops with a single drafter trained at 32 loops. Up to 16 loops, a single pass ($E = 0$) is fastest, whereas at 32 and 64 loops the speedup peaks at 16 to 24 passes. DAS-Wave runs the same loops starting from copies of the last token.

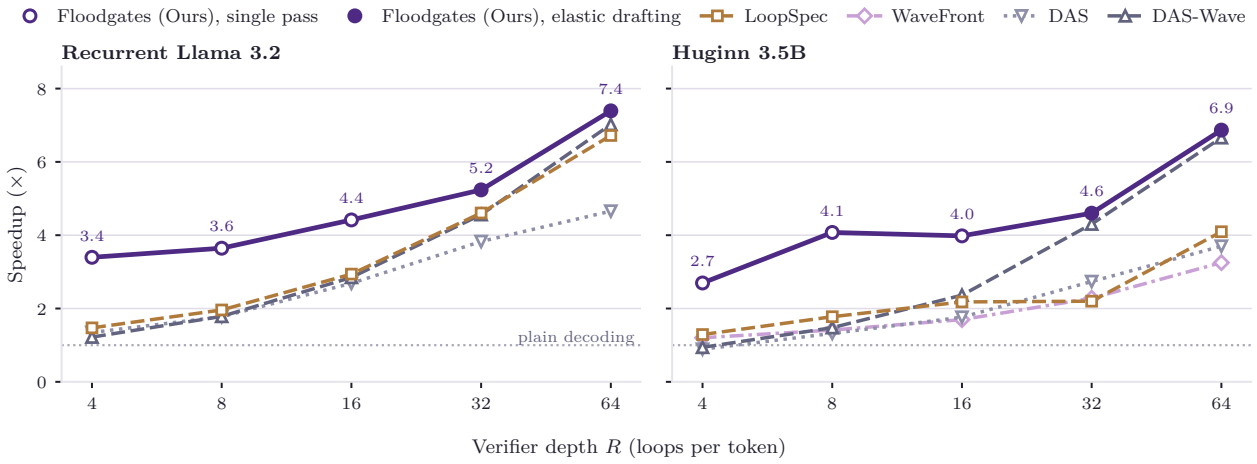

Figure 6 Floodgates is the fastest method at every verifier depth. Speedup over plain decoding on GSM8K when Recurrent Llama 3.2 and Huginn 3.5B run at 4 to 64 loops. Each method uses its best setting, and Floodgates uses a single drafter per model trained at 32 loops. At each depth, we select the elastic drafting budget E on held-out problems. Open circles mark the depths where a single pass ($E = 0$) is selected. Filled circles mark the depths where elastic passes ($E > 0$) are preferred (the deeper depths).

Effect of model depth. Choosing the elastic drafting budget E by optimizing on held-out problems at each depth, we compare Floodgates against all baselines as the test-time depth of Recurrent Llama 3.2 and Huginn 3.5B varies from 4 to 64 loops (Figure 6). Floodgates is the fastest method at every depth on both models. From 4 to 16 loops, where single-pass drafting ($E = 0$) is selected, Floodgates already achieves 3.4–4.4× speedup on Recurrent Llama 3.2 and 2.7–4.1× on Huginn 3.5B. At 32 and 64 loops, adding elastic passes raises the speedup to 5.2× and 7.4× on Recurrent Llama 3.2 and to 4.6× and 6.9× on Huginn 3.5B. In contrast, training-free methods use untrained loops that predict only one new token per step, making them slow on shallow models (1.3–1.5× for LoopSpec, 0.9–1.4× for DAS, and 0.9–1.2× for DAS-Wave at 4 loops) and competitive only at 32 loops and beyond.

4.4 High-throughput batched serving (Q4)

In Section 3, we argued that Floodgates is well suited to concurrent serving because it avoids depth pipelining and reads only the final-loop prefix cache $\mathbf{c}_R^{<t}$ once per draft block, while dynamically adapting the number of verified drafts $M^\star$ (7) as the batch size grows. To test whether these properties translate to high-throughput

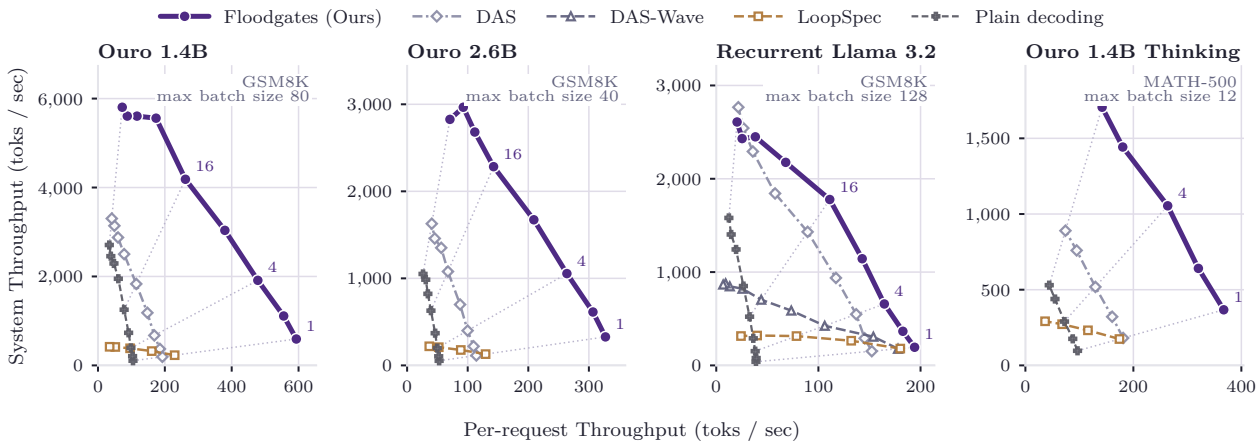

Figure 7 Floodgates has the highest serving throughput at every batch size on the Ouro models, and up to batch size 64 on Recurrent Llama 3.2. System throughput (tokens/s over all requests) against per-request throughput on an A100 under continuous batching. Each point is one batch size, which grows from 1 at the bottom right to the largest that fits in memory at the top left.

serving under hardware saturation, we evaluate continuous batching across varying batch sizes.

Setup. Following Sahoo et al. (2026), we evaluate continuous batching on an A100-80GB GPU across batch sizes up to the device memory limit, decoding all 1,319 GSM8K test problems with Ouro 1.4B, Ouro 2.6B, and Recurrent Llama 3.2, and all 500 MATH-500 problems with Ouro 1.4B Thinking under greedy decoding. At each batch size, we report total system throughput (tokens/s across all requests) and per-request throughput (system throughput divided by batch size). Since LoopSpec's official implementation decodes one request at a time, we run as many instances of its server on the GPU as fit in memory.

Results. Figure 7 shows the results. On Ouro 1.4B and Ouro 2.6B, Floodgates achieves the highest system throughput at every batch size. At the largest batch that fits in memory, it generates 2.1× and 2.7× as many tokens per second as plain decoding (5,805 and 2,827 tokens/s), outperforming the strongest baseline, DAS, which reaches 1.2× and 1.5×, while replicated LoopSpec servers exhaust memory early and peak at 0.15–0.21×. On Recurrent Llama 3.2, Floodgates leads up to batch size 64, while DAS slightly leads at the two largest batch sizes (1.75× vs. 1.65× at batch size 128). Across these models, our adaptive verification budget $M^\star$ (7) matches fixed-block verification (which verifies all k drafts for every request) at small batch sizes and increases throughput at the largest batch size by 19% on Ouro 1.4B and 54% on Recurrent Llama 3.2 (Appendix B.7). On Ouro 1.4B Thinking, where long reasoning chains limit the maximum batch size to 12, Floodgates reaches 3.2× the throughput of plain decoding and 1.9× that of DAS.

Furthermore, because Floodgates does not batch tokens at different loop depths as pipelined methods do, it can be combined directly with cross-loop cache sharing (Prairie et al., 2026), where all R verification loops reuse only the final-loop key-value cache. In Appendix B.8 (Table 14 and figure 10), we show that serving Huginn 3.5B with a single shared cache across loops preserves accuracy within 2 points while scaling the maximum batch size on GSM8K from 16 to 320, where Floodgates achieves the highest throughput at every batch size across all five benchmarks (1.08–1.33× over the strongest baseline at the largest batch size).

5 Related work

Speculative decoding and drafter capacity. Speculative decoding accelerates generation without changing the target distribution by using a fast draft model to propose candidate tokens and verifying them in a single forward pass of the target model (Leviathan et al., 2023; Chen et al., 2023). To achieve a large speedup, the drafter must be both much cheaper to run than the target model and accurate enough that most of its proposals

are accepted (Leviathan et al., 2023; Zhou et al., 2024). In dense models, where every layer is applied once per token, these two requirements conflict. Fast drafters are therefore either small separate networks trained on target features (Li et al., 2024a,b, 2025; Chen et al., 2026a; Hui et al., 2026; Cheng et al., 2026; Huang et al., 2026; Chen et al., 2026b), or self-speculative methods that skip or exit early from layers of the target model (Zhang et al., 2024; Elhoushi et al., 2024; Liu et al., 2024; Xia et al., 2025; Lin et al., 2026). In both cases, the drafter uses only a fraction of the parameters of the target model, which limits its verification rate. On looped models, Cho et al. (2026) even found a separate drafter slower than training-free self-speculation. Floodgates avoids this tradeoff by using a shallow loop of the model's own shared block as the drafter, reducing computation while keeping all parameters of the model.

Multi-token prediction and mask-token drafting. To propose multiple tokens without running the drafter sequentially, a model can be trained to predict several future tokens in parallel. One line of work adds auxiliary output heads that predict future tokens from the final hidden state (Cai et al., 2024; Ankner et al., 2024; Gloeckle et al., 2024; DeepSeek-AI, 2024), or pretrains a looped model so that later loops predict later tokens (Shomali et al., 2026). A second line of work, which is closest to our drafting pass, appends placeholder or mask tokens as a suffix and adapts the model with a small number of trained parameters to fill them in at once (Monea et al., 2023; Bhendawade et al., 2024; Samragh et al., 2025; Sahoo et al., 2026). On a dense model, however, the mask tokens pass through all layers, so a separate drafting pass costs about as much as verification, which can cancel the speedup (Javat and Kazakov, 2026). In a looped model, a drafting pass can be as cheap as a single loop iteration, so Floodgates drafts from the accepted tokens at a small fraction of a decoding step.

Looped language models and their decoding. Looped language models apply a shared block of layers repeatedly to every token (Dehghani et al., 2019; Geiping et al., 2025; Zhu et al., 2025; McLeish et al., 2025; Bae et al., 2025), and the number of loops can be chosen at inference time (Yu et al., 2026; Schwethelm et al., 2026). Prior methods for fast decoding of looped models use the loop without training. LoopSpec and WaveFront take the prediction after an early loop as the draft for the next token (Cho et al., 2026; Ha and Kim, 2026), building on pipelined speculative execution (Hooper et al., 2023). Because the model has only been trained to predict the immediate next token, this yields at most one draft token per loop, and because later tokens start their loops before earlier tokens finish, it requires storing a separate key-value cache for every loop depth. DAS and DAS-Wave instead let shallow passes attend to the final-loop prefix cache (Li et al., 2026), but because the loop is not trained to look ahead, DAS drafts one token at a time and DAS-Wave refines a window of drafts over many shallow passes, which pays off only on models with many loops. Similarly, the parallel sampler of Geiping et al. (2026) decodes multiple tokens in parallel without a verification step, which changes the outputs of the model and thus is not lossless. Floodgates shares with DAS-Wave the deep prefix cache and the use of the model's own loops for elastic passes, but starts from a trained multi-token draft, so that a single loop iteration already yields several tokens and elastic passes are used only when a single loop is cheap.

6 Conclusion

We presented Floodgates, a simple self-speculative decoding method that adapts a shallow loop of a looped language model to draft several upcoming tokens in parallel over suffix mask tokens. Across seven looped models and five benchmarks, Floodgates outperforms prior lossless decoding methods at batch size one, and its speedups carry over to high-throughput batched serving. Our analysis shows that a shallow loop makes an effective drafter because all-depth training already encourages early states to carry information about future tokens, and because conditioning on the deep prefix cache lets a single loop build on full-depth context computation. These findings suggest that looped models are naturally suited for fast multi-token generation once their shallow loops are trained to look ahead.

References

Zachary Ankner, Rishab Parthasarathy, Aniruddha Nrusimha, Christopher Rinard, Jonathan Ragan-Kelley, and William Brandon. Hydra: Sequentially-dependent draft heads for Medusa decoding. In *First Conference on Language Modeling (COLM)*, 2024. (page 12)

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021. (pages 7, 18)

Sangmin Bae, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, and Tal Schuster. Relaxed recursive transformers: Effective parameter sharing with layer-wise LoRA. In *International Conference on Learning Representations (ICLR)*, 2025. (page 12)

Nikhil Bhendawade, Irina Belousova, Qichen Fu, Henry Mason, Mohammad Rastegari, and Mahyar Najibi. Speculative streaming: Fast LLM inference without auxiliary models. *arXiv preprint arXiv:2402.11131*, 2024. (pages 4, 12)

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 *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pages 5209–5235, 2024. (pages 4, 12)

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. (pages 1, 3, 11)

Jian Chen, Yesheng Liang, and Zhijian Liu. DFlash: Block diffusion for flash speculative decoding. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, volume 306 of *Proceedings of Machine Learning Research*, pages 15343–15355, 2026a. (pages 1, 8, 12, 20)

Kecheng Chen, Yuyang He, Cheng Gong, Hui Liu, Guoping Long, Jiajun Li, Shi Wu, Suiyun Zhang, Haoliang Li, Ziru Liu, and Rui Liu. D-Loop: Looped diffusion drafting for speculative decoding. *arXiv preprint arXiv:2610.06011*, 2026b. (page 12)

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. (pages 7, 18)

Zhuoming Chen, Avner May, Ruslan Svirschevski, Yuhsun Huang, Max Ryabinin, Zhihao Jia, and Beidi Chen. Sequoia: Scalable, robust, and hardware-aware speculative decoding. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. (page 5)

Xin Cheng, Xingkai Yu, Chenze Shao, Jiashi Li, Yunfan Xiong, Yi Qian, Jiaqi Zhu, Shirong Ma, Xiaokang Zhang, Jiasheng Ye, Qinyu Chen, Chengqi Deng, Jiping Yu, Damai Dai, Zhengyan Zhang, Yixuan Wei, Yixuan Tan, Wenkai Yang, Runxin Xu, Yu Wu, Zhean Xu, Xuanyu Wang, Muyang Chen, Rui Tian, Xiao Bi, Zhewen Hao, Shaoyuan Chen, Huanqi Cao, Wentao Zhang, Anyi Xu, Huishuai Zhang, Dongyan Zhao, and Wenfeng Liang. DSpark: Confidence-scheduled speculative decoding with semi-autoregressive generation. *arXiv preprint arXiv:2607.05147*, 2026. (pages 12, 28)

SangLyul Cho, Langqing Cui, Sehoon Kim, Dongsu Han, and Insu Han. LoopSpec: Pipelined self-speculative decoding for looped transformers. *arXiv preprint arXiv:2609.17184*, 2026. (pages 1, 5, 6, 7, 12, 18, 20)

Kenneth Ward Church and Patrick Hanks. Word association norms, mutual information, and lexicography. *Computational Linguistics*, 16(1):22–29, 1990. (page 1)

Karl Cobbe, Vineet Kosaraju, Mohammad 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. (pages 7, 18)

DeepSeek-AI. DeepSeek-V3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. (page 12)

Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. In *International Conference on Learning Representations (ICLR)*, 2019. (pages 1, 3, 12)

Mostafa Elhoushi, Akshat Shrivastava, Diana Liskovich, Basil Hosmer, Bram Wasti, Liangzhen Lai, Anas Mahmoud, Bilge Acun, Saurabh Agarwal, Ahmed Roman, Ahmed A. Aly, Beidi Chen, and Carole-Jean Wu.

LayerSkip: Enabling early exit inference and self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024. (pages 1, 3, 8, 12, 20, 22)

Yichao Fu, Peter Bailis, Ion Stoica, and Hao Zhang. Break the sequential dependency of LLM inference using lookahead decoding. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024. (pages 5, 7, 20)

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness. Zenodo, version v0.4.3, July 2024. URL https://zenodo.org/records/12608602. (page 18)

Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. (pages 1, 3, 7, 9, 12, 18)

Jonas Geiping, Xinyu Yang, and Guinan Su. Efficient parallel samplers for recurrent-depth models and their connection to diffusion language models. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, volume 306 of *Proceedings of Machine Learning Research*, pages 34442–34462, 2026. (pages 12, 22, 23)

Fabian Gloeckle, Badr Youbi Idrissi, Baptiste Rozière, David Lopez-Paz, and Gabriel Synnaeve. Better & faster large language models via multi-token prediction. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pages 15706–15734, 2024. (page 12)

Hyeongju Ha and Jae-Joon Kim. WaveFront decoding: Parallelized self-speculative decoding for looped language models. *arXiv preprint arXiv:2609.23033*, 2026. (pages 1, 5, 6, 7, 12)

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2021. (pages 7, 19)

Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Hasan Genc, Kurt Keutzer, Amir Gholami, and Yakun Sophia Shao. SPEED: Speculative pipelined execution for efficient decoding. *arXiv preprint arXiv:2310.12072*, 2023. (page 12)

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*, 2022. (pages 2, 4, 19)

Jianuo Huang, Yaojie Zhang, Qituan Zhang, Hao Lin, Hanlin Xu, and Linfeng Zhang. Domino: Decoupling causal modeling from autoregressive drafting in speculative decoding. *arXiv preprint arXiv:2605.29707*, 2026. (page 12)

Mude Hui, Xin Huang, Jaime Campos Salas, Yue Sun, Nathan Pemberton, Xiang Song, Ashish Khetan, and George Karypis. P-EAGLE: Parallel-drafting EAGLE with scalable training. *arXiv preprint arXiv:2602.01469*, 2026. (page 12)

Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. CodeSearchNet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436*, 2019. (page 19)

Abdurrahman Javat and Allan Kazakov. Accepted prefixes are not all you need: A negative result on PEFT-based block-diffusion drafting. *arXiv preprint arXiv:2607.12422*, 2026. (page 12)

Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023. (pages 1, 3, 5, 11, 12)

Guanghao Li, Zihan Su, Hao Yu, Jinyang Jiang, Tao Ren, Zehao Li, Feng Lu, Ming Tang, and Chun Yuan. Shallow queries, mature values: Depth-asynchronous self-speculation for looped transformers. *arXiv preprint arXiv:2609.34538*, 2026. (pages 1, 5, 6, 7, 12, 20)

Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE: Speculative sampling requires rethinking feature uncertainty. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pages 28935–28948, 2024a. (pages 1, 8, 12, 20)

Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE-2: Faster inference of language models with dynamic draft trees. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7421–7432, 2024b. doi: 10.18653/v1/2024.emnlp-main.422. (page 12)

Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. EAGLE-3: Scaling up inference acceleration of large language models via training-time test. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. (page 12)

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri 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 (ICLR)*, 2024. (pages 7, 18)

Hankun Lin, Patrick Pynadath, and Ruqi Zhang. SEED: Self-speculative decoding via implicit encoder–decoder. *arXiv preprint arXiv:2609.36590*, 2026. (page 12)

Fangcheng Liu, Yehui Tang, Zhenhua Liu, Yunsheng Ni, Duyu Tang, Kai Han, and Yunhe Wang. Kangaroo: Lossless self-speculative decoding for accelerating LLMs via double early exiting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. (pages 1, 12)

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. (pages 7, 18)

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. (page 19)

Sean McLeish, Ang Li, John Kirchenbauer, Dayal Singh Kalra, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Jonas Geiping, Tom Goldstein, and Micah Goldblum. Teaching pretrained language models to think deeper with retrofitted recurrence. *arXiv preprint arXiv:2511.07384*, 2025. (pages 3, 6, 12, 18)

Giovanni Monea, Armand Joulin, and Edouard Grave. PaSS: Parallel speculative sampling. In *3rd Workshop on Efficient Natural Language and Speech Processing (ENLSP), NeurIPS 2023*, 2023. (pages 4, 12)

Koyena Pal, Jiuding Sun, Andrew Yuan, Byron Wallace, and David Bau. Future lens: Anticipating subsequent tokens from a single hidden state. In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 548–560. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.conll-1.37. (pages 3, 8, 23)

Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The FineWeb datasets: Decanting the web for the finest text data at scale. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2024. (page 19)

Hayden Prairie, Jacob M. Ryan, Taylor Berg-Kirkpatrick, and Daniel Y. Fu. BeNice! closing looping's inference-memory gap via KV-sharing. *arXiv preprint arXiv:2610.31633*, 2026. (pages 11, 26)

Subham Sekhar Sahoo, Lingjie Chen, Khiem Pham, Jonathan Geuter, Chaitanya Dwivedi, Varad Pimpalkhute, Yash Akhauri, Alexander Moreno, Mikhail Yurochkin, Zhenting Wang, Mostafa Elhoushi, Nolan Dey, Shane Bergsma, Joel Hestness, John Thickstun, Eric Xing, and Zhengzhong Liu. Unlocking lossless speedups in LLMs via discrete diffusion. *arXiv preprint arXiv:2609.04010*, 2026. (pages 4, 5, 11, 12)

Mohammad Samragh, Arnav Kundu, David Harrison, Kumari Nishu, Devang Naik, Minsik Cho, and Mehrdad Farajtabar. Your LLM knows the future: Uncovering its multi-token prediction potential. *arXiv preprint arXiv:2507.11851*, 2025. (pages 4, 12)

Andrea Santilli, Silvio Severino, Emilian Postolache, Valentino Maiorca, Michele Mancusi, Riccardo Marin, and Emanuele Rodolà. Accelerating transformer inference for translation via parallel decoding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023. (pages 1, 5, 6, 7, 20)

Apoorv Saxena. Prompt lookup decoding. https://github.com/apoorvumang/prompt-lookup-decoding, 2023. (pages 6, 7, 20)

Kristian Schwethelm, Daniel Rueckert, and Georgios Kaissis. Depth-adaptive inference of looped language models via continuous depth batching. *arXiv preprint arXiv:2608.09444*, 2026. (page 12)

Behzad Shomali, Markus Frey, David Berghaus, Joachim Koehler, and Mehdi Ali. LoopMTP: A looped transformer guided by latent multi-token prediction. *arXiv preprint arXiv:2608.03624*, 2026. (page 12)

Ayhan Suleymanzade, Halil Alperen Gozeten, Michael Bronstein, İsmail İlkan Ceylan, and Jinwoo Kim. MUX: Continuous reasoning via multiplexed tokens. In *Advances in Neural Information Processing Systems*, 2026a. (page 1)

Ayhan Suleymanzade, Chanhyuk Lee, Floor Eijkelboom, Nicholas M. Boffi, İsmail İlkan Ceylan, and Jinwoo Kim. Thinking with looped flows. *arXiv preprint arXiv:2609.11801*, 2026b. (page 9)

Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging BIG-Bench tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL 2023*, 2023. (pages 7, 18)

Heming Xia, Yongqi Li, Jun Zhang, Cunxiao Du, and Wenjie Li. SWIFT: On-the-fly self-speculative decoding for LLM inference acceleration. In *International Conference on Learning Representations (ICLR)*, 2025. (page 12)

Mingqian Yu, Wenpeng Zhang, Shaobo Cui, and Peilin Zhao. T-LoopFormer: Token-level elastic-depth looped transformers for latent reasoning with dynamic routing. *arXiv preprint arXiv:2609.15160*, 2026. (page 12)

Jun Zhang, Jue Wang, Huan Li, Lidan Shou, Ke Chen, Gang Chen, and Sharad Mehrotra. Draft & verify: Lossless large language model acceleration via self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024. (pages 1, 12)

Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. SGLang: Efficient execution of structured language model programs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. (page 19)

Yongchao Zhou, Kaifeng Lyu, Ankit Singh Rawat, Aditya Krishna Menon, Afshin Rostamizadeh, Sanjiv Kumar, Jean-François Kagy, and Rishabh Agarwal. DistillSpec: Improving speculative decoding via knowledge distillation. In *International Conference on Learning Representations*, 2024. (page 12)

Rui-Jie Zhu, Zixuan Wang, Kai Hua, Tianyu Zhang, Ziniu Li, Haoran Que, Boyi Wei, Zixin Wen, Fan Yin, He Xing, Lu Li, Jiajun Shi, Kaijing Ma, Shanda Li, Taylor Kergan, Andrew Smith, Xingwei Qu, Mude Hui, Bohong Wu, Qiyang Min, Hongzhi Huang, Xun Zhou, Wei Ye, Jiaheng Liu, Jian Yang, Yunfeng Shi, Chenghua Lin, Enduo Zhao, Tianle Cai, Ge Zhang, Wenhao Huang, Yoshua Bengio, and Jason Eshraghian. Scaling latent reasoning via looped language models. *arXiv preprint arXiv:2510.25741*, 2025. (pages 1, 3, 6, 12, 18)

Appendix

A Experimental details **18**

A.1 Models and tasks . . . 18

A.2 Training details . . . 19

A.3 Implementation details . . . 19

A.4 Baseline settings . . . 19

B Additional results **20**

B.1 Single-pass vs. elastic drafting . . . 20

B.2 Confidence intervals . . . 20

B.3 Accuracy . . . 21

B.4 Comparison with the parallel sampler . . . 22

B.5 Multi-token prediction with early layer states . . . 22

B.6 Acceptance rate by draft position . . . 23

B.7 System throughput across batch sizes . . . 26

B.8 Serving with a single shared cache across loops . . . 26

C Design choices **28**

A Experimental details

We describe the evaluated models and benchmarks in Appendix A.1, drafter training in Appendix A.2, our inference engine in Appendix A.3, and baseline configurations in Appendix A.4.

A.1 Models and tasks

Models. Table 5 summarizes the seven looped models evaluated in our experiments. Ouro 1.4B and Ouro 2.6B (Zhu et al., 2025), along with their reasoning variants Ouro 1.4B Thinking and Ouro 2.6B Thinking, loop their entire layer stack $R = 4$ times. Recurrent Llama 3.2, Recurrent TinyLlama (McLeish et al., 2025), and Huginn 3.5B (Geiping et al., 2025) run an initial prelude once, loop a middle core block, and run a final coda once. Because these three models were trained with loop counts sampled randomly, they can be evaluated at arbitrary loop depths at inference time. Recurrent Llama 3.2 and Recurrent TinyLlama were adapted from pretrained Llama 3.2 1B and TinyLlama 1.1B checkpoints (McLeish et al., 2025), while Huginn 3.5B was pretrained from scratch (Geiping et al., 2025). All models are evaluated using their public checkpoints in bfloat16 precision.

Table 5 The Ouro models loop all of their layers four times, and the other three models loop a middle block 32 times. Block denotes the looped layers, R the default number of loops, and layers per token the total number of layer evaluations in one decoding step.

Model	Block	R	Layers per token	Width
Ouro 1.4B, Ouro 1.4B Thinking	entire stack, 24 layers	4	96	2048
Ouro 2.6B, Ouro 2.6B Thinking	entire stack, 48 layers	4	192	2048
Recurrent Llama 3.2	prelude 4, core 6, coda 4	32	200	2048
Recurrent TinyLlama	prelude 4, core 8, coda 4	32	264	2048
Huginn 3.5B	prelude 2, core 4, coda 2	32	132	5280

Tasks. Table 6 lists the evaluation benchmarks, comprising GSM8K (Cobbe et al., 2021), MATH-500 (Lightman et al., 2024), HumanEval+ and MBPP+ (Chen et al., 2021; Austin et al., 2021; Liu et al., 2023), and BBH (Suzgun et al., 2023). Following LoopSpec (Cho et al., 2026), we use the prompts, few-shot formats, generation limits, and stop strings implemented in `lm-evaluation-harness` (Gao et al., 2024). Prompts on GSM8K and BBH use 3-shot chain-of-thought examples, and prompts on MATH-500 use 4-shot examples. On HumanEval+ and MBPP+, each prompt provides a Python function signature and docstring, and the model completes the function body. The reasoning models use their native chat template with thinking enabled, capped at 4,096 new tokens on GSM8K and 8,192 on MATH-500. A generation is scored as correct on GSM8K and BBH if the extracted final answer matches the reference, on MATH-500 if it is mathematically equivalent to the reference, and on HumanEval+ and MBPP+ if the completed function passes all test cases. Recurrent TinyLlama has a 2,048-token context limit, which excludes BBH prompts and caps its MATH-500 generation length at 1,024 tokens.

Table 6 All benchmarks follow the evaluation protocol of LoopSpec. We list the number of test problems, few-shot examples, maximum new tokens, and stop strings of each benchmark (Cho et al., 2026).

Task	Problems	Shots	Max. new tokens	Stop strings
GSM8K	first 100 of 1,319	3	512	`Q:`, end of text
MATH-500	first 100 of 500	4	2,048	`Problem:`
HumanEval+	all 164	0	1,024	end of function
MBPP+	all 378	0	1,024	end of function
BBH	20 per task, 540	3	1,024	`Q`, blank line, end of text

A.2 Training details

Data. We train each drafter on greedy generations from its target model. For the five base models, we generate continuations on 7,473 GSM8K training problems, 7,500 MATH training problems (Hendrycks et al., 2021), 3,000 Python function signatures and docstrings from CodeSearchNet (Husain et al., 2019), and 256-token prefixes from 4,000 FineWeb-Edu documents (Penedo et al., 2024). GSM8K and MATH prompts follow the 3-shot and 4-shot evaluation formats, and CodeSearchNet prompts follow the HumanEval+/MBPP+ format. For the reasoning models, we generate continuations on the 7,473 GSM8K training problems and 2,032 MATH training problems using their chat template with thinking enabled and a limit of 4,096 new tokens.

Adapters and loss. We apply rank-64 LoRA adapters (Hu et al., 2022) to all attention and feed-forward projections. On Recurrent Llama 3.2, Recurrent TinyLlama, and Huginn 3.5B, adapters are placed on the prelude, core, and coda layers, as well as the input projection that combines the token embedding with the recurrent state. Adapters and the mask embedding $\mathbf{m}$ are initialized to zero. In the training loss (5), we set slot weights $w_j \propto e^{-(j-1)/4}$ normalized to sum to one. On the Ouro models, each mask slot predicts the token at its own position, producing k drafts from k mask tokens. On the other three models, each position predicts the next token, so the slot of the last verified token $\mathbf{x}^t$ also produces a draft, and k mask tokens yield $k+1$ drafts.

Optimization. We train with AdamW (Loshchilov and Hutter, 2019) using decay rates $(\beta_1, \beta_2) = (0.9, 0.95)$ for the first and second moments, zero weight decay, a peak learning rate of 4×10^{-4} with 100 warmup steps and linear decay to 4×10^{-5}, gradient clipping at 1.0, and a batch size of 8 sequences. For each training sequence, we sample 64 positions uniformly from the generated continuation and append a block of k mask slots after each position. A single forward pass of the frozen target model computes the deep prefix cache $\mathbf{c}_R^{<t}$ for all sampled positions in a sequence. Attention masks prevent slots placed at different sampled positions from attending to one another. The frozen model runs in bfloat16 and the adapters in float32. Table 7 lists the hyperparameters and training time for each model, and Appendix C ablates the main design choices.

Table 7 Each drafter trains 31M to 121M parameters in 4 to 15 hours on one GPU. Drafts per pass is the number of tokens drafted by one loop, and E is the elastic drafting budget selected on held-out problems (Appendix B.1).

	Ouro 1.4B	Ouro 2.6B	1.4B Thinking	2.6B Thinking	Rec. Llama	Rec. TinyLlama	Huginn
Training sequences	21,973	21,973	9,505	9,505	21,973	21,973	21,973
Mask tokens k	12	12	8	8	12	12	12
Drafts per pass	12	12	8	8	13	13	13
Adapter rank	64	64	64	64	64	64	64
Trained parameters	60.6M	121.1M	60.6M	121.1M	34.3M	30.8M	50.2M
Training steps	32K	32K	16K	16K	32K	32K	32K
Elastic passes E	0	0	0	0	24	24	24
Training time	6.3 h	9.2 h	3.8 h	7.3 h	8.4 h	8.9 h	14.7 h
GPU	H200	B300	B300	B300	H100	H100	B300

A.3 Implementation details

We implement our decoding engine in PyTorch with static key-value buffers for each loop depth, CUDA graphs, and Triton kernels. The training-free baselines in Table 1 (except LoopSpec and WaveFront, which use the authors' code) run in this engine and share the same verification step. Plain decoding in our engine is 32% faster on Ouro 1.4B than in SGLang, ensuring that our speedups are measured against a strong baseline. All inference experiments are run on a single NVIDIA A100-80GB GPU in bfloat16.

A.4 Baseline settings

LoopSpec and WaveFront. We use the official implementations from the authors. LoopSpec serves one request at a time in SGLang (Zheng et al., 2024) and is benchmarked against SGLang's plain decoding. Following the authors' settings, we use proposal depths $d_1 = 1$ for one proposal (LoopSpec*) and $(d_1, d_2) = (1, 2)$ for the full method on the Ouro models, and $d_1 = 2$ and $(2, 8)$ on the 32-loop models. WaveFront is evaluated in its lossless setting on the models it supports, drafting a token after every loop on the Ouro models and after every T_d loops on Huginn 3.5B, where we report the better of $T_d = 4$ and $T_d = 8$.

DAS and DAS-Wave. Since neither method has public code, we reimplemented both following Li et al. (2026) in our engine with the same verification step as Floodgates. DAS drafts eight tokens sequentially, where each draft position runs r loops conditioned on the deep prefix cache ($r = 2$ on Recurrent Llama 3.2 and Recurrent TinyLlama, and $r = 4$ on Huginn 3.5B). On the Ouro models, where an untrained first loop cannot read the final-loop cache (Figure 4), DAS runs one loop per draft conditioned on the first-loop cache, which gives its best results. DAS-Wave updates a window of up to γ positions in parallel over J refinement rounds. Each round runs k loops over the window conditioned on the deep prefix cache, where each position takes as input the token predicted at the preceding position in the previous round and the window grows by one position per round. On the 32-loop models, we use the paper's settings ($\gamma = 32$ and $J = 24$, with $k = 1$ on Recurrent Llama 3.2 and Recurrent TinyLlama and $k = 2$ on Huginn 3.5B). On the 4-loop Ouro models, where each round costs 1/4 of a decoding step, we use $J = 12$ rounds over $\gamma = 16$ positions on the base models and $J = 8$ rounds over $\gamma = 32$ positions on the reasoning models, which performed best among the settings we tested.

Jacobi decoding and prompt lookup decoding. Jacobi decoding iteratively refines a window of candidate tokens over multiple passes of the full model (Santilli et al., 2023; Fu et al., 2024). We use a window of 32 positions with two model passes conditioned on the deep prefix cache on the 32-loop models, and report the best of one, two, or four passes on the Ouro models. Prompt lookup decoding searches the prefix for the most recent match of the last n tokens and proposes the γ tokens that followed it (Saxena, 2023). We report the better of $(n, \gamma) = (2, 8)$ and $(3, 12)$.

Trained baselines. The separate drafters in Section 4.2 are trained for Ouro 1.4B on the same four-domain data and for the same 32K steps as Floodgates (Appendix A.2). The 1-layer and 4-layer separate drafters follow EAGLE (Li et al., 2024a), where each position combines its token embedding with the final hidden state of Ouro 1.4B, passes through the drafter layers, and is read out by the frozen head of Ouro 1.4B over $k = 12$ mask tokens using the loss in (5). The DFlash drafter (Chen et al., 2026a), following LoopSpec (Cho et al., 2026), uses five layers at the width of Ouro 1.4B that attend to hidden states from five layers of Ouro 1.4B and drafts a block of 16 tokens from mask tokens, trained with the optimizer settings and loss weights of DFlash. On Llama 3.2 1B and LayerSkip Llama 3.2 1B (Elhoushi et al., 2024), we replace the loop with the first L_d layers of the dense model and train LoRA adapters on those layers over $k = 12$ mask tokens using each model's own generations (Appendix A.2). For LayerSkip's native self-speculative decoding, we draft one token at a time from the exit after layer L_d, verify with the remaining layers, and report the speedup at its best L_d.

B Additional results

B.1 Single-pass vs. elastic drafting

The elastic drafting budget E is the only inference hyperparameter of Floodgates. For each model, we select E on 100 held-out problems (test problems 101–150 of GSM8K and MATH-500, excluded from Tables 1 and 2). We evaluate $E \in \{0, 1, 2\}$ on the Ouro models and $E \in \{0, 8, 16, 24, 32\}$ on the 32-loop models, choosing the E with the highest mean speedup over the two tasks. This selects $E = 0$ on the four Ouro models and $E = 24$ on the three 32-loop models. Table 8 compares both settings across all models and tasks. On the Ouro models ($R = 4$), where each elastic pass costs 1/4 of a decoding step, a single pass ($E = 0$) is faster on every task under both greedy decoding and sampling. On the 32-loop models ($R = 32$), where each pass costs 1/32 of a decoding step, elastic drafting ($E = 24$) is faster on 12 of the 14 tasks under greedy decoding (except GSM8K and MATH-500 on Huginn 3.5B, where $E = 0$ is faster by 0.4% and 8%) and on 12 of the 14 tasks under sampling (except GSM8K on Recurrent Llama 3.2 and Huginn 3.5B).

B.2 Confidence intervals

Table 9 reports 95% bootstrap confidence intervals for the greedy speedups in Table 1. For each task, we resample the test problems 2,000 times with replacement and compute the speedup of each method on the same resamples. The intervals remain within 5% of the speedup for the training-free baselines and within 8% for Floodgates.

Table 8 A single pass is fastest on the 4-loop Ouro models, and elastic drafting is faster on most tasks of the 32-loop models. Speedup (×) over plain decoding at batch size one, with tokens per step τ in parentheses, for a single pass ($E = 0$) and for elastic drafting with E passes. Shaded rows indicate the setting selected on held-out problems and reported in Tables 1 and 2. Reasoning models are sampled at temperature 0.6 and top-p 0.95.

Model	E	GSM8K	MATH-500	HumanEval+	MBPP+	BBH
Greedy decoding						
Ouro 1.4B	0	5.04 (7.5)	4.69 (7.0)	4.83 (7.2)	3.01 (4.6)	4.40 (6.6)
	1	3.75 (7.0)	2.89 (5.4)	3.37 (6.3)	2.27 (4.4)	3.29 (6.2)
Ouro 2.6B	0	5.56 (8.2)	6.04 (9.0)	5.57 (8.3)	3.39 (5.2)	5.07 (7.5)
	1	3.43 (6.3)	2.92 (5.4)	3.14 (5.8)	2.28 (4.4)	3.15 (5.9)
Ouro 1.4B Thinking	0	3.70 (5.6)	3.68 (5.5)	–	–	–
	1	2.56 (4.9)	2.39 (4.5)	–	–	–
Ouro 2.6B Thinking	0	3.90 (5.9)	3.72 (5.5)	–	–	–
	1	2.19 (4.1)	2.09 (3.9)	–	–	–
Recurrent Llama 3.2	0	4.93 (5.9)	6.47 (7.8)	5.37 (6.5)	4.53 (5.4)	4.47 (5.4)
	24	5.04 (20.4)	6.97 (28.5)	6.91 (28.0)	5.84 (23.5)	6.62 (27.1)
Huginn 3.5B	0	4.50 (5.1)	6.69 (7.5)	4.77 (5.4)	3.72 (4.2)	3.66 (4.1)
	24	4.48 (19.0)	6.21 (26.6)	6.18 (26.2)	5.06 (21.4)	5.70 (24.6)
Recurrent TinyLlama	0	5.22 (6.2)	6.07 (7.2)	5.27 (6.2)	4.72 (5.5)	–
	24	6.73 (22.2)	7.79 (26.0)	8.20 (27.0)	7.17 (23.5)	–
Sampling at temperature 1 *and top-p* 0.7						
Ouro 1.4B	0	4.16 (6.3)	4.41 (6.7)	4.49 (6.8)	2.75 (4.3)	4.11 (6.2)
	1	3.16 (6.0)	2.63 (5.0)	3.17 (6.0)	2.09 (4.1)	3.02 (5.7)
Ouro 2.6B	0	4.50 (6.7)	4.90 (7.3)	4.82 (7.2)	3.23 (4.9)	4.77 (7.1)
	1	3.05 (5.7)	2.57 (4.8)	2.79 (5.2)	2.15 (4.1)	3.01 (5.6)
Ouro 1.4B Thinking	0	3.46 (5.2)	3.23 (4.9)	–	–	–
	1	2.31 (4.5)	2.14 (4.2)	–	–	–
Ouro 2.6B Thinking	0	3.71 (5.5)	3.41 (5.1)	–	–	–
	1	2.06 (3.9)	1.97 (3.8)	–	–	–
Recurrent Llama 3.2	0	3.65 (4.4)	4.47 (5.5)	4.35 (5.3)	4.23 (5.1)	4.00 (4.9)
	24	3.28 (7.6)	4.66 (11.0)	4.69 (10.9)	4.34 (10.1)	5.13 (12.1)
Huginn 3.5B	0	3.24 (3.7)	4.53 (5.1)	3.91 (4.4)	2.93 (3.3)	3.34 (3.8)
	24	3.16 (7.3)	4.60 (10.8)	4.70 (10.9)	3.52 (8.1)	4.48 (19.3)
Recurrent TinyLlama	0	3.73 (4.4)	4.54 (5.4)	4.16 (5.0)	4.00 (4.8)	–
	24	4.01 (8.1)	4.83 (9.9)	5.61 (11.4)	5.25 (10.6)	–

Table 9 The greedy speedups of Table 1 vary by at most 8% when the test problems are resampled. Speedup (×) over plain decoding with 95% bootstrap intervals over problems, for Floodgates and the prior method with the highest mean speedup on each model.

Model	Method	GSM8K	MATH-500	HumanEval+	MBPP+	BBH
Ouro 1.4B	LoopSpec	2.77 [2.73, 2.80]	2.81 [2.78, 2.84]	3.10 [3.08, 3.11]	2.83 [2.79, 2.88]	2.86 [2.84, 2.88]
	Floodgates	**5.04** [4.84, 5.24]	**4.69** [4.49, 4.91]	**4.83** [4.69, 4.99]	**3.01** [2.85, 3.25]	**4.40** [4.28, 4.52]
Ouro 2.6B	LoopSpec	3.06 [3.03, 3.09]	3.09 [3.02, 3.15]	3.33 [3.31, 3.34]	3.09 [3.07, 3.11]	3.10 [3.08, 3.11]
	Floodgates	**5.56** [5.36, 5.77]	**6.04** [5.59, 6.46]	**5.57** [5.43, 5.70]	**3.39** [3.30, 3.49]	**5.07** [4.93, 5.20]
Rec. Llama 3.2	LoopSpec	4.61 [4.45, 4.77]	6.11 [5.88, 6.28]	6.03 [5.90, 6.14]	5.41 [5.21, 5.60]	5.65 [5.57, 5.73]
	Floodgates	**5.04** [4.75, 5.35]	**6.97** [6.68, 7.16]	**6.91** [6.75, 7.06]	**5.84** [5.54, 6.07]	**6.62** [6.54, 6.70]
Huginn 3.5B	DAS-Wave	4.23 [4.03, 4.43]	5.38 [5.12, 5.55]	5.45 [5.36, 5.53]	4.57 [4.43, 4.72]	5.20 [5.14, 5.26]
	Floodgates	**4.48** [4.20, 4.76]	**6.21** [5.83, 6.47]	**6.18** [6.06, 6.30]	**5.06** [4.88, 5.25]	**5.70** [5.62, 5.78]
Rec. TinyLlama	LoopSpec	5.79 [5.58, 6.01]	6.92 [6.65, 7.16]	7.35 [7.16, 7.50]	6.48 [6.20, 6.74]	–
	Floodgates	**6.73** [6.44, 7.01]	**7.79** [7.52, 8.04]	**8.20** [8.01, 8.37]	**7.17** [6.84, 7.47]	–

B.3 Accuracy

Table 10 reports the task accuracy of plain decoding and Floodgates across all models and benchmarks in Tables 1 and 2. On each task, the two methods differ on at most seven test problems due to bfloat16 rounding when

verifying several tokens in parallel instead of one at a time. In our batched serving experiments (Section 4.4), where we use batch-invariant kernels whose floating-point reductions are independent of batch size, Floodgates produces bit-identical outputs to plain decoding across all problems on every model.

Table 10 Floodgates matches the accuracy of plain decoding within 3 points on every task. Accuracy (%) under greedy decoding. For each task, the left column is plain decoding and the shaded column is Floodgates. The two differ on at most seven problems per task, because bfloat16 rounding differs when several tokens are verified in parallel.

Model	GSM8K		MATH-500		HumanEval+		MBPP+		BBH	
Ouro 1.4B	75.0	74.0	76.0	73.0	64.6	64.0	65.1	65.1	70.9	71.1
Ouro 2.6B	85.0	85.0	57.0	59.0	74.4	74.4	69.3	69.0	78.9	80.2
Recurrent Llama 3.2	47.0	45.0	44.0	41.0	11.6	11.0	24.3	23.8	34.3	34.1
Huginn 3.5B	29.0	29.0	25.0	26.0	14.6	14.6	26.2	26.5	31.9	32.0
Recurrent TinyLlama	38.0	40.0	41.0	42.0	14.0	13.4	23.3	23.0	–	–
Ouro 1.4B Thinking	87.0	85.0	88.0	85.0	–	–	–	–	–	–
Ouro 2.6B Thinking	88.0	87.0	89.0	91.0	–	–	–	–	–	–

B.4 Comparison with the parallel sampler

The parallel sampler of Geiping et al. (2026) initializes a window of tokens and updates their hidden states in parallel across multiple recurrent steps, emitting each token once its hidden state changes by less than a threshold. Because the released code does not support the retrofitted models, we reimplemented it in PyTorch with four loops per step and a convergence threshold of 0.03, and measure its speedup relative to plain decoding in the same implementation. Because the parallel sampler does not verify its predictions, it is not lossless (Table 11). Under greedy decoding, its accuracy stays within 3 points of plain decoding on Recurrent Llama 3.2 and Recurrent TinyLlama, but drops on all five tasks on Huginn 3.5B, including by 4 points on GSM8K (23.0% vs. 27.0%) and 9 points on MATH-500 (16.0% vs. 25.0%), where a third of the tokens reach 32 loops before converging. Under sampling, accuracy drops further, by 13 points on GSM8K and 12 points on MBPP+ for Huginn 3.5B, and by 6–8 points on MATH-500 and MBPP+ for Recurrent Llama 3.2. Table 12 compares its speedups with those of Floodgates. Under greedy decoding, the parallel sampler is faster than Floodgates on GSM8K for Recurrent Llama 3.2 and Huginn 3.5B and on MBPP+ for Huginn 3.5B, and slower in the other eleven settings, but it does not preserve the outputs of the model.

Table 11 The parallel sampler does not preserve accuracy, which drops by up to 13 points on Huginn 3.5B. Accuracy (%). For each task, the left column is plain decoding and the shaded column is the parallel sampler (Geiping et al., 2026).

Model	GSM8K		MATH-500		HumanEval+		MBPP+		BBH	
Greedy decoding										
Recurrent Llama 3.2	46.0	47.0	41.0	39.0	11.0	8.5	23.0	23.0	37.0	38.5
Huginn 3.5B	27.0	23.0	25.0	16.0	14.6	12.2	27.0	25.4	31.9	31.1
Recurrent TinyLlama	37.0	40.0	40.0	39.0	14.0	14.0	24.1	24.1	–	–
Sampling at temperature 1 *and top-p* 0.7										
Recurrent Llama 3.2	52.0	50.0	39.0	31.0	13.4	11.6	16.9	11.1	38.5	36.3
Huginn 3.5B	19.0	6.0	18.0	10.0	7.3	6.1	22.0	10.3	30.4	31.1
Recurrent TinyLlama	39.0	34.0	34.0	33.0	7.9	7.9	16.9	16.1	–	–

B.5 Multi-token prediction with early layer states

In Sections 1 and 4.2, we argue that supervising next-token prediction at early depths encourages shallow states to carry information about several upcoming tokens. To test this on dense models, we compare Llama 3.2 1B against LayerSkip Llama 3.2 1B (Elhoushi et al., 2024), which fine-tunes Llama 3.2 1B with next-token losses at shallow layers, and measure how well the hidden state after each layer predicts the next three tokens.

Table 12 Under greedy decoding, Floodgates is faster than the parallel sampler in 11 of 14 settings, without changing the outputs. Wall-clock speedup (×) over plain decoding on an A100 at batch size one, with accepted tokens per step τ in parentheses. Sampling uses temperature 1.0 and top-p 0.7. † marks the parallel sampler (Geiping et al., 2026), which is not lossless. It is faster under sampling, where its accuracy drops the most (Table 11).

		GSM8K		MATH-500		HumanEval+		MBPP+		BBH	
Model	Method	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling	Greedy	Sampling
Recurrent Llama 3.2	Parallel sampler†	5.38	5.03	5.68	5.41	5.67	5.56	5.54	5.52	5.44	5.39
	Floodgates	5.04 (20.4)	3.28 (7.6)	6.97 (28.5)	4.66 (11.0)	6.91 (28.0)	4.69 (10.9)	5.84 (23.5)	4.34 (10.1)	6.62 (27.1)	5.13 (12.1)
Huginn 3.5B	Parallel sampler†	5.37	5.38	5.79	5.79	5.78	5.73	5.24	5.38	5.34	5.13
	Floodgates	4.48 (19.0)	3.16 (7.3)	6.21 (26.6)	4.60 (10.8)	6.18 (26.2)	4.70 (10.9)	5.06 (21.4)	3.52 (8.1)	5.70 (24.6)	4.48 (19.3)
Recurrent TinyLlama	Parallel sampler†	5.72	5.32	6.01	5.87	6.03	5.06	5.86	5.54	prompts exceed	
	Floodgates	6.73 (22.2)	4.01 (8.1)	7.79 (26.0)	4.83 (9.9)	8.20 (27.0)	5.61 (11.4)	7.17 (23.5)	5.25 (10.6)		

Linear readers. Following Future Lens (Pal et al., 2023), we train a separate linear reader for each layer and each lookahead offset $j \in \{1, 2, 3\}$. The reader applies a learned linear map to the hidden state of token $\mathbf{x}^t$ after a given layer, followed by the frozen final normalization and output head of the model, to predict $\mathbf{x}^{t+j}$. Because the output head is frozen, the reader only extracts information already present in the hidden state. We initialize each reader to the identity matrix, which recovers the model's next-token prediction at the final layer, and train it with cross-entropy for four epochs, selecting the checkpoint with the highest accuracy on a held-out 10% validation split.

Data. Because the two models share a tokenizer, we train and evaluate both on the same text. The training set consists of 6,300 sequences from the benchmark training splits (reference solutions of GSM8K and MATH problems, MBPP programs outside MBPP+, and BBH problems outside our evaluation split). The test set consists of reference solutions for 500 GSM8K test problems and MATH-500, canonical programs for HumanEval+ and MBPP+, and the evaluation problems of BBH.

Results. Figure 8 shows the reader accuracy across layers on each benchmark. On all five benchmarks, LayerSkip Llama 3.2 1B predicts the immediate next token ($j = 1$) more accurately than Llama 3.2 1B from every intermediate layer by up to 22 points, while Llama 3.2 1B is equally or more accurate at the final layer. This gap carries over to two and three tokens ahead ($j = 2$ and $j = 3$), though it narrows with distance. The only exception is $j = 3$ on HumanEval+, where both models remain within one point across all layers. These results support the hypothesis in Sections 1 and 4.2 that next-token supervision at early depths causes intermediate states to encode information about future tokens.

B.6 Acceptance rate by draft position

To examine how elastic drafting increases accepted draft length, Figure 9 plots the acceptance rate at each draft position j on Recurrent Llama 3.2 ($R = 32$). Single-pass drafting ($E = 0$) is accurate on early slots, with 88–94% acceptance at the first position, but its acceptance rate drops after position 12. Elastic drafting with $E = 24$ passes extends high acceptance across the window, keeping the acceptance rate above 50% up to positions 23–32 depending on the task. DAS-Wave runs the same loops initialized from copies of the last verified token; it matches Floodgates over the first 16 positions, but falls behind on later positions and drops sharply after position 26 on MATH-500, HumanEval+, and BBH. Because Floodgates initializes the full window from its trained parallel draft, it retains 44–86% acceptance at position 24, compared to 20–67% for DAS-Wave. The last panel of Figure 9 varies the elastic drafting budget E on GSM8K, showing that additional passes extend the 50% acceptance horizon from position 7 at $E = 4$ to position 23 at $E = 24$.

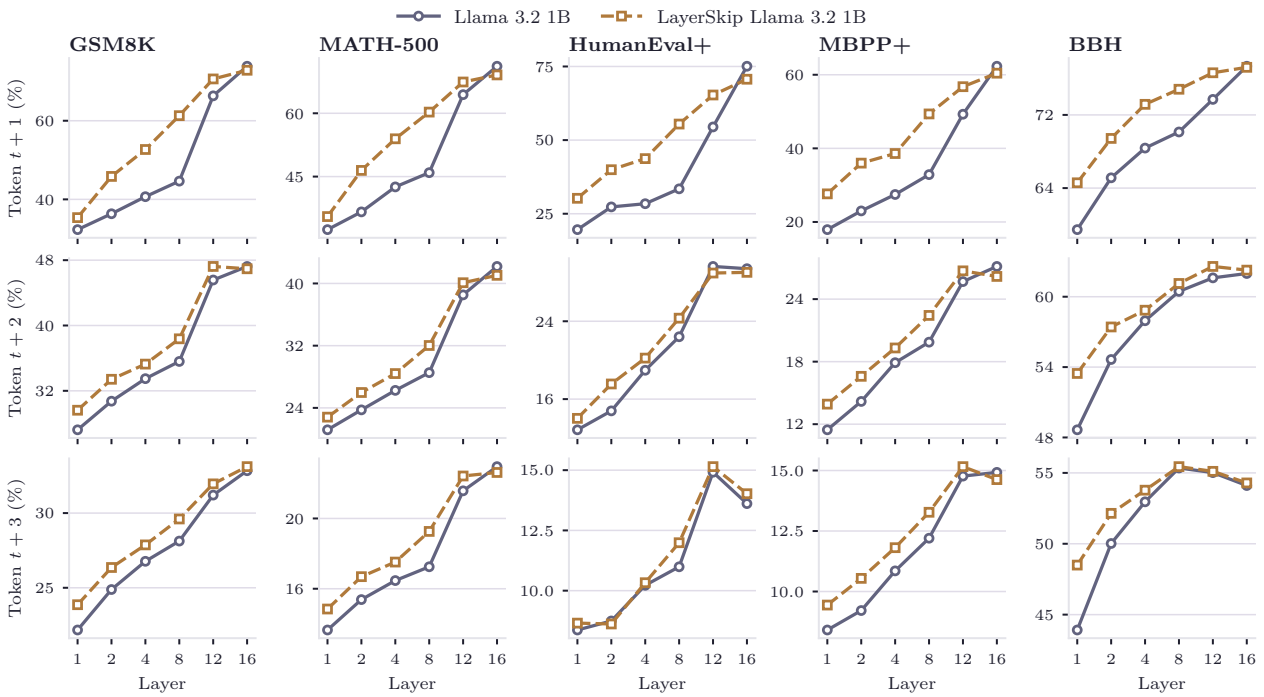

Figure 8 Early-exit supervision makes the early layers more predictive of the next tokens. Accuracy of linear readers that predict the token one, two, or three positions ahead from the hidden state after each layer of Llama 3.2 1B and LayerSkip Llama 3.2 1B.

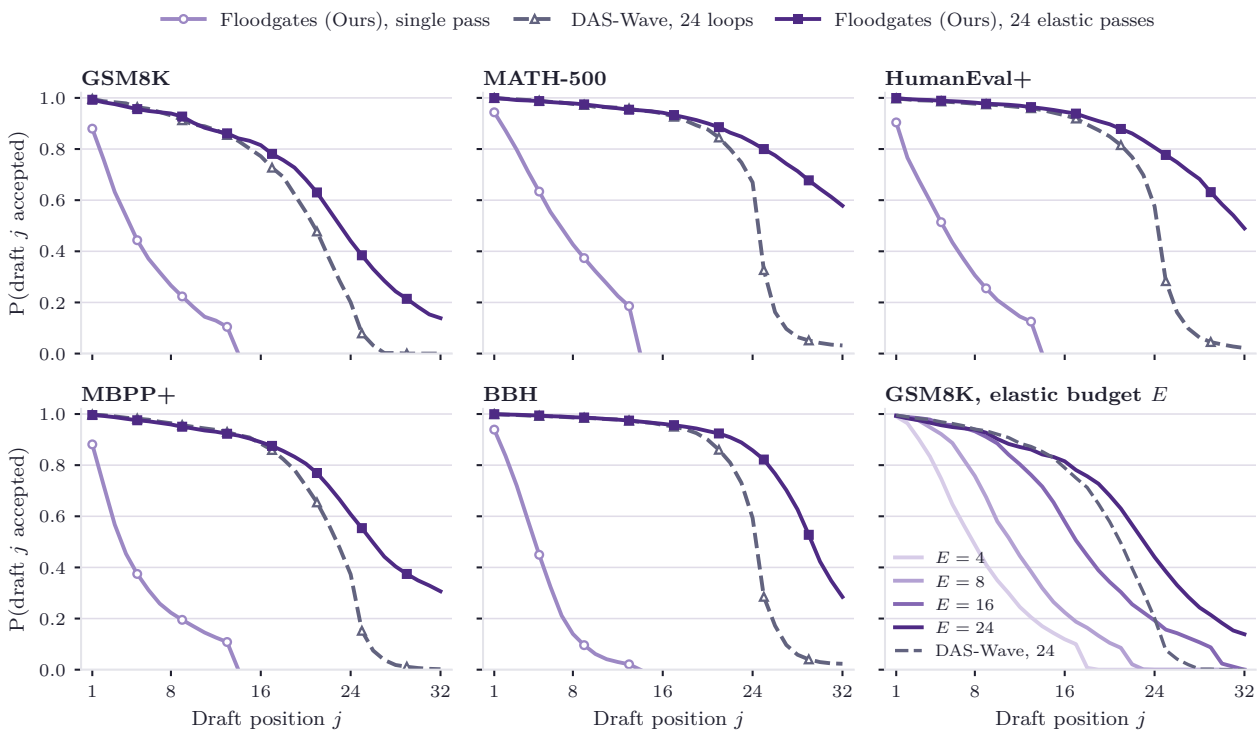

Figure 9 With 24 elastic passes, the acceptance rate stays above 50% up to draft positions 23 to 32. Acceptance rate by draft position on Recurrent Llama 3.2 ($R = 32$), measured as the fraction of verification steps in which the j-th draft is accepted. DAS-Wave runs the same 24 loops starting from copies of the last verified token. The last panel varies the elastic drafting budget E on GSM8K.

Table 13 Floodgates reaches the highest peak throughput on three of the four models, and the adaptive budget adds up to 54% at the largest batch size. System throughput (tokens/s) across batch sizes. The last column reports the speedup of each method's peak throughput over that of plain decoding. "Fixed block" verifies all k drafts of every request instead of the budget $M^\star$ in (7). LoopSpec ($*$) runs as many server instances as fit in memory.

Batch size	1	2	4	8	16	32	Largest	Best
Ouro 1.4B, GSM8K, largest batch size 80								
Plain decoding	105	207	394	733	1,252	1,948	2,708	1.00×
LoopSpec	229	322	384	414	–	–	419*	0.15×
DAS	193	371	674	1,183	1,832	2,502	3,302	1.22×
Prompt lookup decoding	206	387	662	1,045	1,460	1,775	2,079	0.77×
Floodgates, fixed block	588	1,099	1,865	2,980	3,787	4,463	4,859	1.81×
Floodgates	**593**	**1,111**	**1,913**	**3,037**	**4,185**	**5,559**	**5,805**	2.14×
Ouro 2.6B, GSM8K, largest batch size 40								
Plain decoding	53	104	199	370	629	985	1,051	1.00×
LoopSpec	130	177	209	–	–	–	220*	0.21×
DAS	114	219	399	699	1,079	1,458	1,628	1.55×
Prompt lookup decoding	102	192	331	523	725	886	950	0.90×
Floodgates, fixed block	323	603	1,023	1,636	2,073	2,456	2,218	2.34×
Floodgates	**327**	**613**	**1,055**	**1,673**	**2,283**	**2,966**	**2,827**	2.82×
Recurrent Llama 3.2, GSM8K, largest batch size 128								
Plain decoding	39	78	152	290	522	850	1,580	1.00×
LoopSpec	180	264	314	320	–	–	315*	0.20×
DAS	152	291	548	937	1,431	1,843	**2,767**	1.75×
DAS-Wave	178	308	425	587	704	822	866	0.56×
Prompt lookup decoding	73	138	256	445	671	824	1,033	0.65×
Floodgates, fixed block	190	354	631	1,009	1,271	1,426	1,697	1.07×
Floodgates	**194**	**365**	**658**	**1,143**	**1,778**	**2,176**	2,609	1.65×
Ouro 1.4B Thinking, MATH-500, largest batch size 12								
Plain decoding	97	175	290	439	–	–	530	1.00×
LoopSpec	174	232	272	292	–	–	292*	0.55×
DAS	181	321	519	760	–	–	890	1.68×
Prompt lookup decoding	170	302	471	642	–	–	764	1.44×
Floodgates, fixed block	**369**	636	999	1,406	–	–	1,598	3.02×
Floodgates	367	**641**	**1,054**	**1,443**	–	–	**1,705**	3.22×

B.7 System throughput across batch sizes

Table 13 reports the per-batch-size system throughput values underlying Figure 7, along with the maximum batch size that fits in memory and the peak throughput of each method relative to plain decoding. We also evaluate Floodgates with a fixed verification block, which verifies all k drafts for every sequence at all batch sizes instead of using the adaptive budget $M^\star$ in (7). Fixed-block verification matches $M^\star$ at small batch sizes, but falls behind as the batch size grows, reducing throughput at the largest batch size by 19% on Ouro 1.4B and 54% on Recurrent Llama 3.2.

B.8 Serving with a single shared cache across loops

A looped model stores the keys and values of every loop, which limits how many requests fit in memory. Prairie et al. (2026) show that some looped models can instead keep only the final-loop key-value cache and reuse it at every loop without further training. Unlike the deep prefix cache in the main text, which only the drafting pass reads, this setting reuses the final-loop cache across all R verification loops as well. Because LoopSpec and WaveFront start the next token before the current token finishes its loops, they require the cache of every loop and cannot use a single cache. We therefore compare Floodgates with DAS, DAS-Wave, and plain decoding.

Accuracy. Table 14 compares per-loop caches with a single shared cache at batch size one on the three 32-loop models. Using a single cache changes the accuracy of plain decoding by at most 2 points on Huginn 3.5B and Recurrent TinyLlama, and lowers it by 3–5 points on GSM8K and MATH-500 for Recurrent Llama 3.2. With the single cache, Floodgates matches plain decoding within 1 point on almost every task.

Table 14 A single shared cache changes the accuracy of plain decoding by at most 2 points on Huginn 3.5B and Recurrent TinyLlama. Accuracy (%) at batch size one, with a separate cache per loop or with the final-loop cache reused by all loops. Both cache settings are evaluated on the same problems, with the number of problems in parentheses.

Model	Decoding	GSM8K (100)	MATH-500 (100)	HumanEval+ (164)	MBPP+ (378)	BBH (540)
Recurrent Llama 3.2	Plain decoding, per-loop caches	47.0	44.0	11.6	24.3	34.3
	Plain decoding, single cache	44.0	39.0	9.8	24.6	35.2
	Floodgates, single cache	44.0	42.0	10.4	24.3	35.0
Recurrent TinyLlama	Plain decoding, per-loop caches	38.0	41.0	14.0	23.3	–
	Plain decoding, single cache	38.0	42.0	13.4	24.3	–
	Floodgates, single cache	39.0	42.0	14.0	23.8	–
Huginn 3.5B	Plain decoding, per-loop caches	29.0	25.0	14.6	26.2	31.9
	Plain decoding, single cache	27.0	23.0	13.4	26.7	32.6
	Floodgates, single cache	28.0	23.0	14.0	27.2	32.2

Throughput. Since a single cache preserves the accuracy of Huginn 3.5B, we serve it with continuous batching as in Section 4.4 on all problems of the five tasks (Figure 10). With per-loop caches, Floodgates runs out of memory on GSM8K at batch size 24 and peaks at 380 tokens/s at batch size 16. With a single cache, it scales to batch size 320 on GSM8K (reaching 481 tokens/s), 192 on MBPP+, 96 on MATH-500 and BBH, and 64 on HumanEval+. Floodgates achieves the highest system throughput at every batch size on all five tasks. At the largest batch size, it outperforms the strongest baseline by 1.29–1.33× on HumanEval+, GSM8K, and MATH-500, 1.15× on MBPP+, and 1.08× on BBH. Together, these results show that Floodgates combines naturally with cache sharing to serve much larger batches while maintaining the highest throughput.

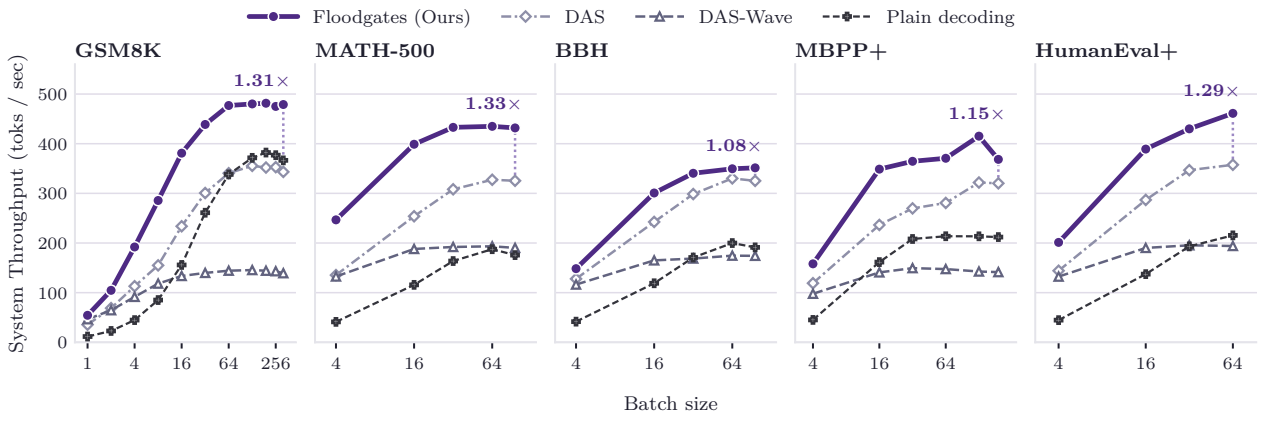

Figure 10 With a single shared cache, Floodgates has the highest throughput at every batch size on all five tasks. System throughput against batch size on an A100 for Huginn 3.5B, served with one key-value cache for all loops. The number in each panel is the throughput of Floodgates relative to the strongest baseline at the largest batch size.

C Design choices

We ablate the design choices of Floodgates on Ouro 1.4B in Table 15. Our default configuration samples 64 training positions per sequence and uses $k = 12$ mask tokens with rank-64 LoRA adapters. Each row of Table 15 modifies a single choice from this default, with all drafters trained for 8K steps on the four-domain dataset of Appendix A.2.

Markov head. In (4), each slot state $\hat{\mathbf{z}}^{t+j}$ is read out independently by the head g, so the draft token $\mathbf{d}^j$ does not depend directly on the preceding draft token $\mathbf{d}^{j-1}$. Following DSpark (Cheng et al., 2026), we also evaluate a Markov head that adds a logit bias conditioned on the previous draft token

$$\mathbf{q}^j(v) \propto g(\hat{\mathbf{z}}^{t+j})(v)\, \exp\big((\mathbf{a}_{\mathbf{d}^{j-1}}^{\top}\mathbf{B})_v\big), \qquad \mathbf{d}^j = \mathsf{argmax}(\mathbf{q}^j), \qquad \mathbf{d}^0 = \mathbf{x}^t, \tag{8}$$

where $\mathbf{a}_u \in \mathbb{R}^{128}$ is a learned token embedding and $\mathbf{B} \in \mathbb{R}^{128\times|V|}$ projects it to a vocabulary bias. We initialize $\mathbf{B}$ to zero so that training starts from the standard drafter. During training, $\mathbf{d}^{j-1}$ is replaced by the ground-truth token $\mathbf{x}^{t+j-1}$. During drafting, tokens are selected from left to right after the parallel pass of f_θ, adding one small matrix-vector product per slot.

Results. Table 15 shows the ablation results. Sampling fewer positions per sequence (16), using fewer mask tokens ($k = 8$), halving the adapter rank (32), and replacing the learned mask embedding with random noise all reduce both accepted tokens per step and decoding throughput. Increasing the number of mask tokens to $k = 16$ raises accepted tokens per step on GSM8K from 7.69 to 8.26, and adding the Markov head improves acceptance on both tasks (from 7.69 to 8.00 on GSM8K and from 6.46 to 6.98 on MATH-500), yielding the highest throughput (482 tokens/s).

Table 15 The default drafter is within 2% of the fastest configuration. Accepted tokens per step on GSM8K and MATH-500, and decoding throughput (tokens/s) on GSM8K, for drafters of Ouro 1.4B trained for 8K steps. Each row changes one choice of the default configuration in the first row.

Configuration	GSM8K	MATH-500	tok/s
Floodgates, 64 positions, $k = 12$, rank 64	7.69	6.46	475
16 sampled positions per sequence	7.30	5.90	449
$k = 8$	6.58	5.84	410
$k = 16$	**8.26**	6.63	477
rank 32	7.50	6.23	465
noise input instead of mask tokens	7.54	6.21	451
Markov head	8.00	**6.98**	**482**