

000 SEMANTIC-AWARE DIFFUSION LLM INFERENCE 001 002 WITH ADAPTIVE BLOCK SIZE 003

004 **Anonymous authors**

005 Paper under double-blind review

006 007 ABSTRACT 008 009

010 Diffusion-based large language models (dLLMs) are gaining attention for their
011 inherent capacity for parallel decoding, offering a compelling alternative to
012 autoregressive LLMs. Among various decoding strategies, blockwise semi-
013 autoregressive (semi-AR) approaches are widely adopted due to their natural sup-
014 port for KV caching and their favorable accuracy–speed trade-off. However, this
015 paper identifies two fundamental limitations in the conventional semi-AR decod-
016 ing approach that applies a fixed block size: *i*) late decoding overhead, where the
017 unmasking of high-confidence tokens outside the current block is unnecessarilly
018 delayed; and *ii*) premature decoding error, where low-confidence tokens inside
019 the current block are committed too early, leading to incorrect tokens. This paper
020 presents the first systematic investigation challenging the fixed block size assump-
021 tion in semi-AR decoding. Through a statistical analysis of confidence dynamics
022 during the denoising process, we identify a volatility band (VB) region during
023 dLLM decoding, which encodes local semantic structure and can be used to guide
024 adaptive block sizing. Leveraging these insights, we introduce *AdaBlock-dLLM*,
025 a training-free, plug-and-play scheduler that adaptively aligns block boundaries
026 with semantic steps by adjusting block size during runtime. Extensive experi-
027 ments across diverse benchmarks show that *AdaBlock-dLLM* achieves up to 5.3%
028 accuracy improvement under the same throughput budget. Beyond inference-
029 time optimization, we hope our semantics-aware adaptive scheduling approach
030 and confidence-based analysis will inspire future training strategies for dLLMs.
031

032 1 INTRODUCTION

033 Diffusion-based Large Language models (dLLMs) have recently emerged as a promising alternative
034 to autoregressive models, offering parallel decoding, improved controllability, and greater data effi-
035 ciency in low-resource settings (Zhang et al., 2025; Prabhudesai et al., 2025). Open-source dLLMs
036 such as LLaDA (Nie et al., 2025; Zhu et al., 2025) and Dream (Ye et al., 2025) have shown com-
037 parable performance to autoregressive models of similar scale. Notably, in structured generation
038 tasks such as coding, proprietary models including Seed Diffusion (Song et al., 2025b) and Gem-
039 ini Diffusion (Gemini Diffusion, 2025) demonstrate throughput exceeding 1,400 tokens per second.
040 These advances highlight the potential of dLLMs to deliver efficient inference while maintaining
041 competitive algorithmic performance.

042 Recent works have widely adopted a semi-autoregressive (semi-AR) decoding paradigm that com-
043 bines blockwise KV caching (Wu et al., 2025; Chen et al., 2025; Song et al., 2025a) and confidence-
044 based dynamic sampling (Wang et al., 2025d; Wu et al., 2025; Wei et al., 2025) to improve inference
045 efficiency. However, semi-AR decoding enforces block-level causality: the current block must be fi-
046 nalized before decoding the next block. We take the first attempt to identify two **fundamental issues**
047 introduced by conventional semi-AR decoding with a fixed block size: ***i*) Late Decoding Overhead**.
048 As shown in the upper-left of Figure 1, semi-AR decoding delays the unmasking of high-confidence
049 tokens outside the current block. For instance, the second and third blocks are decoded in separate
050 iterations, incurring unnecessary computational overhead to generate a simple complete sentence. ***ii)***
051 **Premature Decoding Error.** As shown in the lower-left of Figure 1, the autoregressiveness across
052 blocks forces early commitment to low-confidence tokens within each block, and this suboptimal
053 sampling often yields incorrect token predictions (Figure 5), particularly in reasoning tasks.

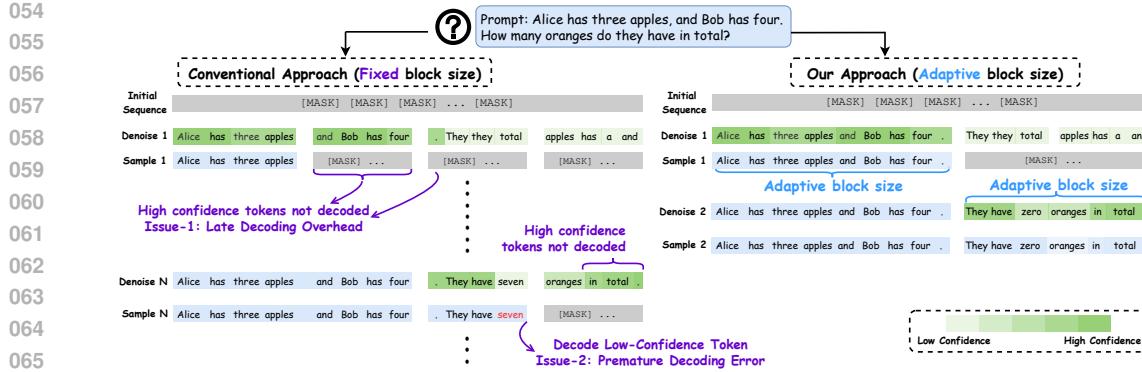


Figure 1: Illustrative examples of two key issues (left) and how they can be overcome with *AdaBlock-dLLM* (right). A real case study is provided in Appendix A.1.

To address these limitations, we first investigate how confidence scores evolve during dLLM denoising and sampling. Our statistical analysis in Section 4.1 reveals a volatility band (VB), a region where confidence fluctuates dynamically. VB regions exhibit semantics structure, which can be exploited to dynamically adapt block size during runtime. Based on these insights and observations, we propose an adaptive block-size decoding method for dLLM inference, termed *AdaBlock-dLLM*, which adopts a semantic-aware approach that adaptively adjusts block boundaries (Figure 1, right). Specifically, *AdaBlock-dLLM* dynamically aligns block size with the length of semantic blocks, as categorized by special semantic tokens (e.g., periods and \n), enabling dLLM to perform efficient decoding while mitigating the limitations of fixed-size approaches. Importantly, *AdaBlock-dLLM* is a training-free and plug-and-play enhancement to the existing semi-AR decoding paradigm. Across comprehensive experiments on various benchmarks, we demonstrate that *AdaBlock-dLLM* improves accuracy by up to 5.3% while achieving throughput comparable to prior dLLM acceleration methods (Figure 2). Gains are especially pronounced under KV caching, where fixed block sizes further compromise semantic consistency. Our results motivate semantics-aware training objectives for block-diffusion models that emphasize context preservation.

In summary, our contributions are threefold:

- We systematically analyze the semi-autoregressive sampling paradigm, and identify the inaccuracy and inefficiency behind fixed block size settings (Section 4.1 and Section 4.2).
- We propose *AdaBlock-dLLM*, a training-free, plug-and-play technique that enhances existing semi-autoregressive decoding paradigm, which dynamically adjusts block sizes based on the confidence of semantic delimiter tokens (Section 4.3).
- We conduct extensive experiments demonstrating that *AdaBlock-dLLM* improves accuracy by up to 5.3% over state-of-the-art methods under the same speed budget (Section 5).

2 RELATED WORKS

2.1 DIFFUSION LANGUAGE MODELS

Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021; Karras et al., 2022) have achieved high-fidelity generation across continuous data domains, including images and video (Peebles & Xie, 2023; Ho et al., 2022). Motivated by this success, a growing line of work adapts diffusion to NLP tasks, giving rise to masked diffusion models (MDMs) that iteratively de-noise an initially masked token sequence into coherent output (Austin et al., 2021a; Hoogeboom et al., 2021; Lou et al., 2024). Recent efforts have scaled MDMs up to 8B parameters (Nie et al.,

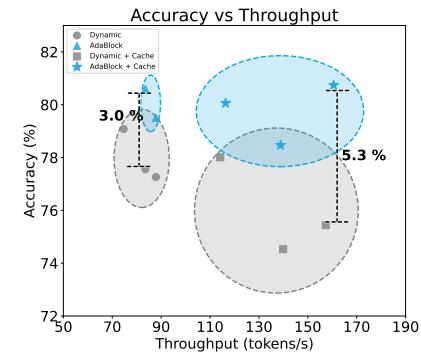


Figure 2: Performance improvement over Fast-dLLM (Wu et al., 2025).

108 2025; Ye et al., 2025), highlighting their robustness and scalability. Current diffusion language
 109 models can be categorized into *i*) models trained from scratch (Nie et al., 2025; Yang et al., 2025); *ii*)
 110 models adapted from autoregressive (AR) models (Ye et al., 2025); and *iii*) block-diffusion
 111 models (Cheng et al., 2025; Wang et al., 2025d), which combine the training efficiency of AR models
 112 with the sampling efficiency of diffusion models.

114 2.2 EFFICIENT INFERENCE FOR DIFFUSION LLMs

116 Recent work has advanced inference for diffusion LLMs (dLLMs) in both speed and accuracy. In
 117 inference acceleration, research has focused on two directions: caching mechanisms and parallel
 118 decoding. Multiple caching mechanisms have been proposed to avoid recomputing key–value
 119 (KV) pairs at each denoising step, including delayed caching (Ma et al., 2025), verification-based
 120 caching (Liu et al., 2025), and block-level caching (Wu et al., 2025). Advances in parallel decoding
 121 include efficient threshold-based dynamic sampling (Wu et al., 2025; Wang et al., 2025c; Wei et al.,
 122 2025; Yu et al., 2025), sampler scheduling (Luxembourg et al., 2025), and guided diffusion (Hu
 123 et al., 2025; Israel et al., 2025). To improve accuracy, test-time strategies such as voting (Wang
 124 et al., 2025b), early committing (Li et al., 2025a) and remasking (Hong et al., 2025; Wang et al.,
 125 2025a; He et al., 2025) have also been explored.

126 2.3 BLOCKWISE SEMI-AUTOREGRESSIVE DECODING

128 Semi-autoregressive (semi-AR) decoding partitions the sequence into blocks. Decoding is autoregressive
 129 at the block level, but non-autoregressive within blocks, allowing tokens inside the block to
 130 be sampled in arbitrary order. This paradigm was first introduced in Block Diffusion (Arriola et al.,
 131 2025), which interpolates between autoregressive and fully diffusion-based decoding by applying
 132 a block-causal attention mask. Various diffusion LLMs, including LLaDA (Nie et al., 2025) and
 133 MMaDA (Yang et al., 2025), adopt semi-AR decoding; however, they predefine a fixed generation
 134 budget. Compared with random-order decoding over the full sequence, a key advantage of semi-AR
 135 decoding is that it naturally supports block-level KV caching (Wu et al., 2025; Chen et al., 2025;
 136 Song et al., 2025a), making it a prevalent scheme in recent dLLM research. To the best of our
 137 knowledge, prior semi-AR decoding uses a **fixed** block size. In contrast, this paper takes the first
 138 attempt to explore **adaptive** block-size decoding with **semantic-aware, training-free** method.

139 3 PRELIMINARIES

142 The decoding process of diffusion LLMs comprises two operations: *denoise* and *sample*. In text
 143 generation, the decoding process starts from a fully masked sequence, and the model iteratively
 144 employs the denoise–sample cycle until no mask token is left. Taking LLaDA (Nie et al., 2025) as
 145 an example, we formalize the decoding process.

146 **Setup.** Let \mathcal{V} denote the vocabulary, which includes a special mask token $[\text{MASK}] \in \mathcal{V}$. Given
 147 a prompt $\mathbf{q} = (q_0, \dots, q_{L_p-1}) \in \mathcal{V}^{L_p}$ and a generation budget L , define the index set $\mathcal{J} \triangleq$
 148 $\{0, 1, \dots, L_p+L-1\}$. Let T be the total number of denoise–sample iterations. At step $t \in$
 149 $\{T, T-1, \dots, 0\}$, the sequence state is $\mathbf{y}^t = (y_i^t)_{i \in \mathcal{J}} \in \mathcal{V}^{L_p+L}$. The initial state of the sequence is
 150 $\mathbf{y}^T = (q_0, \dots, q_{L_p-1}, \underbrace{[\text{MASK}], \dots, [\text{MASK}]}_{L \text{ times}})$.

152 **Denoise.** A mask predictor p_θ predicts a sequence $\hat{\mathbf{y}}^t \in (\mathcal{V} \setminus [\text{MASK}])^{L_p+L}$ using greedy decoding:

$$155 \hat{y}_i^t = \arg \max_{v \in \mathcal{V}} p_\theta(v \mid \mathbf{y}^t, i), \quad i \in \mathcal{J}. \quad (1)$$

157 **Sample.** Define the masked-position set:

$$160 \mathcal{M}_t \triangleq \{i \in \mathcal{J} : y_i^t = [\text{MASK}]\}. \quad (2)$$

161 A sampler selects $S_t \subseteq \mathcal{M}_t$ to unmask. For all $i \in \mathcal{J}$, update

162
163
164
165
166
167

$$y_i^{t-1} = \begin{cases} \hat{y}_i^t, & i \in S_t \text{ (unmask),} \\ [\text{MASK}], & i \in \mathcal{M}_t \setminus S_t \text{ (stay masked),} \\ y_i^t, & i \notin \mathcal{M}_t \text{ (already unmasked; keep).} \end{cases} \quad (3)$$

168 Then, the sequence state becomes

169
170

$$\mathbf{y}^{t-1} = (y_i^{t-1})_{i \in \mathcal{J}} \in \mathcal{V}^{L_p+L}. \quad (4)$$

171
172Repeat for $t = T, T-1, \dots, 1$; terminate when $\mathcal{M}_t = \emptyset$.173
174
175
176
177
178
179
180

Sampling Methods. LLaDA (Nie et al., 2025) proposes a linear-schedule sampler that unmasks a fixed number of tokens L/T at each step, either at random or by confidence. LaViDa (Li et al., 2025b) applies a shifting schedule that unmasks a variable number of [MASK] positions in each step. Fast-dLLM (Wu et al., 2025) and Dimple (Yu et al., 2025) adopt dynamic sampling by introducing a confidence threshold τ . At each step t , the model unmask all positions with confidence $c_i^t \triangleq p_\theta(\hat{y}_i^t \mid \mathbf{y}^t, i)$ that $c_i^t \geq \tau$. Compared to the linear-schedule or the shifting schedule, dynamic sampling adapts to token-level uncertainty and improves the accuracy–throughput balance.

181
182

4 METHODOLOGY

183
184
185
186
187
188
189
190

Semi-AR decoding has been widely adopted in dLLMs, including LLaDA (Nie et al., 2025) and MMaDA (Yang et al., 2025). This decoding paradigm naturally supports a block-level KV cache, addressing a key bottleneck in dLLM inference. Additionally, combining semi-AR decoding with dynamic sampling (Wu et al., 2025) enables multi-token prediction within a single denoise–sample cycle (Wu et al., 2025; Wang et al., 2025c). This paradigm introduces two hyperparameters: **confidence threshold** τ and **block size** B . The confidence threshold maintains a balance between sampling speed and quality, whereas the effect of block size has not been systematically explored.

191
192
193
194

In Section 4.1, we analyze the confidence dynamics that govern dynamic sampling. Using these patterns, we characterize the inaccuracies and inefficiencies caused by the misalignment between a fixed block size and the model’s inductive decoding preferences in Section 4.2. This motivates an adaptive block-size scheduler that minimizes this mismatch.

195
196

4.1 ANALYSIS OF CONFIDENCE DYNAMICS

197
198
199
200
201
202

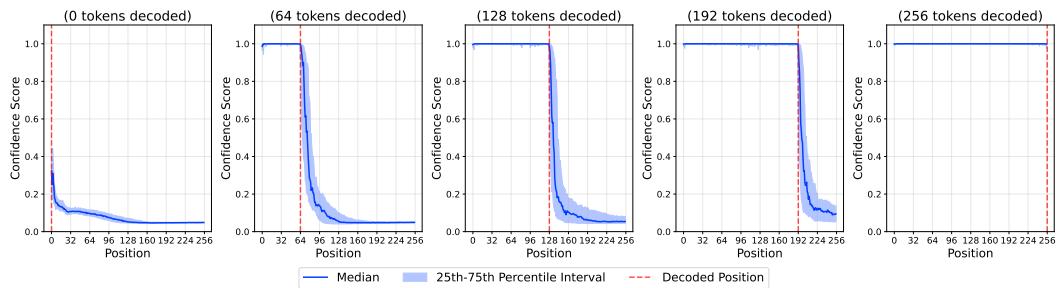
Confidence Scores. Confidence scores are a key metric in dynamic sampling of dLLM. A high confidence score indicates that the model is more certain about the prediction (Wei et al., 2025). The decoding process consists of iterative denoise–sample cycles (Section 3). At each denoise step, the model predicts the tokens in all masked positions via greedy decoding. Afterwards, tokens with confidence $c_i \geq \tau$ are unmasked, where τ is a hyperparameter specifying the confidence threshold.

203
204
205
206
207
208
209
210
211
212
213
214
215

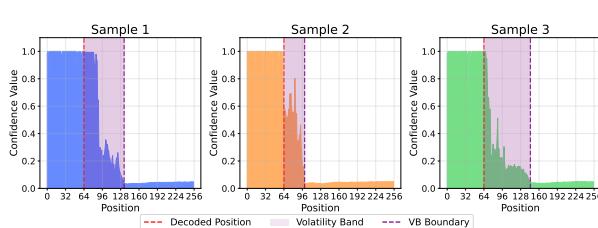
Confidence Dynamics and Global Autoregressiveness. Figure 3 visualizes confidence dynamics at early, middle, and late decoding stages of LLaDA-8B-Base inference on GSM8K Benchmark. Based on these statistical analyses, we summarize the following observations and patterns:

- For all stages, a high-confidence region emerges near the decoded (or prompt) tokens. This indicates that masked positions adjacent to decoded (high-confidence) tokens are more likely to attain high confidence and thus to be decoded. We attribute this to **confidence locality**: dLLMs exhibit higher confidence in regions where local semantic meaning is complete.
- Although the model may generate in arbitrary order, its decoding trace shows a **global autoregressive tendency**. This pattern suggests a chain-like progression, which arises from semantic causality where subsequent predictions depend on prior semantics.
- As the decoding process unfolds, confidence for decoded tokens remains high, and the high-confidence region extends toward adjacent positions. In contrast, positions outside this region maintain consistently low confidence, forming what we term the **low-confidence floor**.

216 **Volatility Band and Local Stochasticity.** Building on these observations, we partition the
 217 position-wise confidence landscape into three regimes: *i*) a high-confidence plateau with consis-
 218 tently large scores; *ii*) a volatility band (VB) characterized by unstable, variable scores; and *iii*) a
 219 low-confidence floor with persistently small scores. Among these regimes, the VB is the key region
 220 where the current decoding steps take place. As illustrated in Figure 4, scores within the VB fluctu-
 221 ate over decoding steps (time) and across positions (space), yet remain clearly separated from the
 222 low-confidence floor. Additionally, the width of the VB varies from case to case. In contrast to the
 223 left-to-right expansion of the high-confidence plateau that is driven by global autoregressive causal-
 224 ity, the decoding order within the VB is locally stochastic: the positional preference diminishes, and
 225 the sampling choice depends more heavily on the immediate semantic context.



236 Figure 3: Confidence scores across sequence positions for LLaDA-8B-Base, evaluated on 100 sam-
 237 ples from the GSM8K benchmark. The high confidence plateau expands as decoding progresses,
 238 while positions beyond the decoded prefix exhibit high variance.



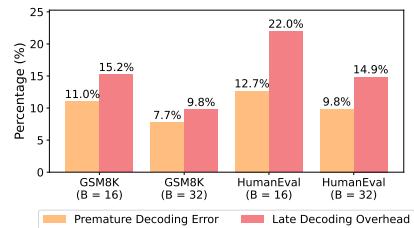
247 Figure 4: Illustration of the high confidence plateau, the
 248 volatility band (VB), and the low confidence floor across
 249 three samples. Within VB, the distribution of confidence
 250 scores and the width of the band vary across samples.

252 4.2 MOTIVATION FOR ADAPTIVE BLOCK SIZE.

253 Despite the local stochasticity exhibited during decoding, prevailing semi-AR decoding strategies
 254 use manually set, fixed block sizes that often fail to capture this stochasticity. This misalignment
 255 leads to **Late Decoding Overhead** and **Premature Decoding Error**, as illustrated in Figure 1,
 256 thereby reducing accuracy and efficiency. This motivates incorporating an adaptive block size.
 257

258 **Late Decoding Overhead.** A fixed block size poses a hard constraint on number of tokens that
 259 can be sampled in a single step. By fixing the block size, the sampler must exclude nearby higher-
 260 confidence positions outside the block, especially when a small block size is used (Wang et al.,
 261 2025d; Cheng et al., 2025). Deferred high-confidence positions undergo additional denoising in
 262 later iterations, incurring unnecessary computation overhead and resulting in degraded throughput.
 263

264 **Premature Decoding Error.** With a fixed block size, the sampler (Algorithm 3) must decode
 265 all masked positions in the current block before advancing. This constraint forces the sampler to
 266 decode low-confidence positions inside the block, instead of the high-confidence positions outside
 267 the block. The result is systematic inaccuracy: decoding premature tokens increases token-level
 268 error rates, propagates mismatches to subsequent steps via block-level autoregressive dependencies,
 269 and biases the hypothesis toward poorly calibrated regions of the confidence landscape. This usually
 contributes to degraded accuracy.



562 Figure 5: Proportion of sampling steps
 563 affected by late decoding overhead and
 564 premature decoding error on GSM8K
 565 and HumanEval for fixed block sizes.

270
271

4.3 SEMANTIC-AWARE ADAPTIVE BLOCK-SIZE SCHEDULER

272
273
274
275
276

Challenge for predicting block size. The VB characterizes the region of active decoding: positions before the VB are typically decoded and exhibit consistently dynamic confidence. As shown in Figure 3, the VB boundary encodes semantic structure due to the confidence locality (Section 4.1). In contrast, tokens in the low-confidence floor are repeatedly assigned placeholders or formatting symbols (e.g., $\langle \text{EOS} \rangle$, spaces).

277
278
279
280
281
282

Although the VB region contains semantic information, it is often too wide. As shown in Figure 7, the token “GB” repeatedly appears with confidence scores mostly between 0.1 and 0.3, lying within the VB. Although such tokens are contextually related, their significance to the current context is weak, providing little guidance for the current sampling step. Consequently, VB regions often provide limited actionable signals for local decisions, requiring a fine-grained segmentation to identify tokens that relate closely to the current decoding step.

283
284
285
286
287
288
289
290
291
292

Align Block Size With Semantic Steps. To obtain a fine-grained segmentation that reflects the context of the current decoding step, we partition the sequence into **semantic steps**, which are contiguous spans whose provisional tokens exhibit local semantic coherence. We then couple the scheduler to the semantic step, setting the block size guided by the current semantic step length. This allows the sampler to finalize higher-confidence positions within the step while deferring lower-confidence positions until the semantic step is ready to close. Across semantic steps, dependencies are enforced by the semi-AR paradigm, since each downstream step conditions on completed predecessors. This alignment curbs premature commitments outside the active step and prevents splitting a step across iterations, thereby reducing both error propagation and computational overhead.

293

Algorithm 1 Semantic-Aware Block Size Determination294
295
296

Inputs: predicted sequence \hat{y} ; confidences \mathbf{c} ; generation budget L ; default block size B_0 ; delimiter set \mathcal{D} ; delimiter threshold $\tau_{\mathcal{D}}$; current position g .
Output: block size B

```

1: function COMPUTEBLOCKLENGTH( $\hat{y}$ ,  $\mathbf{c}$ ,  $L$ ,  $B_0$ ,  $\mathcal{D}$ ,  $\tau_{\mathcal{D}}$ ,  $g$ )
2:    $\triangleright$  Sampling window boundary
3:    $start, remaining \leftarrow g, L - g$ 
4:    $w \leftarrow \min(\max(1, \lfloor 0.25 \cdot g \rfloor), remaining)$ 
5:    $W \leftarrow \{start, \dots, start + w - 1\}$   $\triangleright$  window token indices
6:    $\triangleright$  Find highest-confidence delimiter
7:    $\mathcal{I} \leftarrow \{i \in W \mid \hat{y}_i \in \mathcal{D}\}$ 
8:   if  $\mathcal{I} \neq \emptyset$  then
9:      $pos \leftarrow \arg \max_{i \in \mathcal{I}} c_i$   $\triangleright$  Select position with max delimiter token confidence
10:     $c_{\max} \leftarrow c_{pos}$ 
11:   else
12:     $c_{\max} \leftarrow -\infty$ 
13:   end if
14:    $\triangleright$  Determine block size
15:   if  $c_{\max} \geq \tau_{\mathcal{D}}$  then
16:      $B \leftarrow (pos - start + 1)$   $\triangleright$  inclusive length up to the delimiter
17:   else
18:      $B \leftarrow \min(B_0, remaining)$ 
19:   end if
20:   return  $B$ 
21: end function

```

317

318
319
320
321
322
323

Semantic-Aware Block Size Scheduling With AdaBlock-dLLM. To facilitate the dynamic adjustment of block size B , we insert an additional block-size determination procedure between denoising and sampling at the start of each block (Algorithm 1). Given the current predicted sequence \hat{y} and confidence scores $\{c_i\}$, we collect indices whose predicted tokens fall in the delimiter set \mathcal{D} (line 7). Among these, we choose the delimiter \hat{y}_{\max} with the highest confidence c_{\max} . If $c_{\max} \geq \tau_{\mathcal{D}}$, we set B to the position of \hat{y}_{\max} (lines 15–16), indicating that the model has a reliable preference for the end of the current semantic step. If no delimiters appear in \hat{y} within W_t

(lines 11–12), or if all delimiter predictions are low-confidence ($c_{\max} < \tau_D$, lines 17–18), we fall back to the default block size. Additionally, we apply an index window W that masks distant positions to avoid decoding the $\langle \text{EOS} \rangle$ token in the early stage, a cause for severe performance drop (Nie et al., 2025). This procedure yields step-aware blocks when evidence is strong, while remaining conservative in ambiguous regions.

5 EXPERIMENT

5.1 EXPERIMENTAL SETUP

Implementation Details. We evaluate *AdaBlock-dLLM* on representative diffusion LLMs (dLLMs): LLaDA-8B-Instruct (Nie et al., 2025), LLaDA-1.5 (Zhu et al., 2025), and Dream-v0-Base-7B (Ye et al., 2025). Unless otherwise noted, all experiments run on NVIDIA H100 GPUs.

Hyperparameters Settings. We use the generation budget $L = 512$ for all benchmarks. For dynamic sampling, we use a confidence threshold $\tau = 0.9$. We sweep default block sizes $B \in \{16, 32, 64\}$. *AdaBlock-dLLM* introduces two hyperparameters: the delimiter set \mathcal{D} and the delimiter confidence threshold τ_D . We set $\mathcal{D} = \{\backslash n\}$, since newline tokens commonly mark the end of reasoning steps in test-time search (Snell et al., 2024) style prompting. We use $\tau_D = 0.3$ for LLaDA-series models and $\tau_D = 0.5$ in Dream-series models. This is tuned on a small subset of the GSM8K benchmark, and the selection is discussed in Section 5.3.

Benchmarks and Metrics. We evaluate *AdaBlock-dLLM* on standard LLM benchmarks. For math reasoning, we use GSM8K (5-shot) (Cobbe et al., 2021) and MATH (4-shot) (Hendrycks et al., 2021). For code generation, we use HumanEval (0-shot) (Chen et al., 2021) and MBPP (3-shot) (Austin et al., 2021b). Generation quality is measured by accuracy: pass@1 for code generation and answer accuracy for math reasoning. All accuracy values are reported as percentages. We evaluate accuracy across five sampling methods in total, including three baselines based on Fast-dLLM (Wu et al., 2025): **Vanilla** (top-1 confidence sampling), **Dynamic** (dynamic sampling), and **+Cache** (dynamic sampling with DualCache); and two variants enhanced with *AdaBlock-dLLM*: **+Ada** (dynamic sampling with *AdaBlock-dLLM*) and **+Ada+Cache** (dynamic sampling with DualCache and *AdaBlock-dLLM*).

5.2 MAIN RESULTS

Generation Quality Across Models. Table 1 reports accuracy to quantify generation quality. *AdaBlock-dLLM* achieves accuracy gains across most model and dataset pairs. Notably, on GSM8K with LLaDA-Instruct, accuracy is improved by 3.0% without caching and 5.3% with caching.

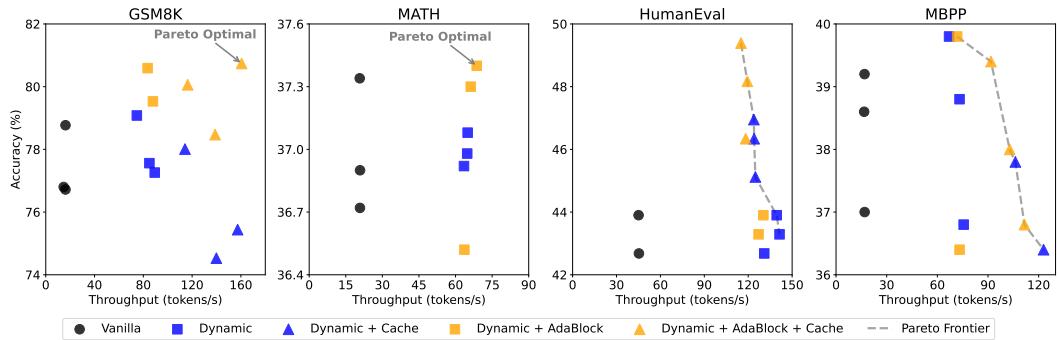


Figure 6: Accuracy and throughput of different sampling methods evaluated on LLaDA-Instruct. Integrating *AdaBlock-dLLM* into Fast-dLLM (Wu et al., 2025) yields accuracy gains across all benchmarks while maintaining little throughput overhead. Notably, Fast-dLLM with *AdaBlock-dLLM* is Pareto-optimal on the GSM8K and MATH datasets.

378 Table 1: Accuracy (%) across sampling methods, evaluated on LLaDA-1.5, LLaDA-Instruct, and
 379 Dream-Base under default block sizes $B_0 \in \{16, 32, 64\}$. Differences are shown in gray. Compar-
 380 isons are reported relative to **Dynamic** and **+Cache** (Wu et al., 2025).

Method	LLaDA-Instruct			LLaDA-1.5			Dream-Base		
	$B_0 = 16$	$B_0 = 32$	$B_0 = 64$	$B_0 = 16$	$B_0 = 32$	$B_0 = 64$	$B_0 = 16$	$B_0 = 32$	$B_0 = 64$
GSM8K									
Vanilla	78.8	76.7	76.8	82.3	82.3	80.4	76.3	76.4	75.1
Dynamic	79.1	77.6	77.3	82.6	82.2	80.7	75.5	75.5	75.6
+Ada	80.6 _{+1.5}	80.6 _{+3.0}	79.5 _{+2.2}	83.0 _{+0.4}	82.4 _{+0.2}	80.3 _{-0.4}	75.7 _{+0.2}	75.7 _{+0.2}	75.9 _{+0.3}
+Cache	78.0	74.5	75.4	80.7	80.2	80.0	75.6	74.5	74.6
+Ada+Cache	80.0 _{+2.0}	78.5 _{+4.0}	80.7 _{+5.3}	81.3 _{+0.6}	81.7 _{+1.5}	79.7 _{-0.3}	76.5 _{+0.9}	75.1 _{+0.6}	74.6 _{+0.0}
HumanEval									
Vanilla	43.9	43.9	42.7	39.0	36.6	38.4	53.7	52.4	54.3
Dynamic	42.7	43.9	43.3	36.6	37.8	36.6	53.0	51.2	52.4
+Ada	43.3 _{+0.6}	43.3 _{-0.6}	43.9 _{+0.6}	37.8 _{+1.2}	38.4 _{+0.6}	38.4 _{+1.8}	53.7 _{+0.7}	51.2 _{+0.0}	53.7 _{+1.3}
+Cache	45.1	46.3	47.0	33.5	36.0	34.1	50.0	53.0	56.1
+Ada+Cache	49.4 _{+4.3}	46.3 _{+0.0}	48.2 _{+1.2}	36.0 _{+2.5}	39.0 _{+3.0}	36.0 _{+1.9}	52.4 _{+2.4}	53.0 _{+0.0}	57.3 _{+1.2}
MATH									
Vanilla	36.7	36.9	37.3	36.3	37.0	34.4	39.8	40.2	40.1
Dynamic	37.0	36.9	37.1	36.3	36.7	34.4	39.7	39.9	39.9
+Ada	36.5 _{-0.5}	37.3 _{+0.4}	37.4 _{+0.3}	36.8 _{+0.5}	36.7 _{+0.0}	34.1 _{-0.3}	39.6 _{-0.1}	39.9 _{+0.0}	39.9 _{+0.0}
+Cache	35.4	35.8	36.0	34.9	33.2	32.1	38.0	38.5	38.8
+Ada+Cache	35.8 _{+0.4}	35.3 _{-0.5}	35.6 _{-0.4}	35.2 _{+0.3}	33.9 _{+0.7}	32.4 _{+0.3}	37.8 _{-0.2}	38.4 _{-0.1}	38.4 _{-0.4}
MBPP									
Vanilla	39.2	38.6	37.0	38.2	37.0	23.2	12.4	12.4	12.8
Dynamic	39.8	38.8	36.8	38.2	37.0	24.6	12.6	12.4	12.2
+Ada	40.2 _{+0.4}	39.8 _{+1.0}	36.4 _{-0.4}	39.4 _{+1.2}	37.6 _{+0.6}	29.8 _{+5.2}	12.8 _{+0.2}	14.2 _{+1.8}	12.4 _{+0.2}
+Cache	35.6	37.8	36.4	38.0	34.8	19.8	12.8	11.6	9.6
+Ada+Cache	39.4 _{+3.8}	38.0 _{+0.2}	36.8 _{+0.4}	36.6 _{-1.4}	36.4 _{+1.6}	36.6 _{+6.8}	12.8 _{+0.0}	11.6 _{+0.0}	12.4 _{+2.8}

407 **Pronounced Accuracy Improvement With Cache.** We observe that accuracy gains are partic-
 408 ularly pronounced when KV caching is used. Unlike autoregressive decoding, where caching is
 409 effectively lossless, block-level KV caching in dLLMs is approximated because key and value ten-
 410 sors vary across time steps, and the decoding order within each block is non-sequential. Notably,
 411 accuracy degrades markedly at large block sizes.

412 Table 3 reports accuracy gains with both PrefixCache and DualCache (Wu et al., 2025). Improve-
 413 ments are twofold: *i*) For large default block sizes B_0 , the resulting average block size \bar{B} is smaller,
 414 reducing KV-cache approximation error. *ii*) By enhancing semantic locality within each block, inter-
 415 block dependencies are reduced, making decoding less sensitive to stale cache entries. The second
 416 effect is more impactful: on GSM8K, \bar{B} for $B_0 = 64$ is 33.98, yet accuracy still exceeds that of
 417 $B_0 = 32$ without *AdaBlock-dLLM* by 1.90% (no cache) and 6.20% (DualCache). Given that block-
 418 level KV caching is a core advantage of semi-AR decoding, these results indicate that the method
 419 integrates seamlessly with existing techniques that improve inference efficiency.

420 **Throughput Overhead.** Table 2 reports the accuracy, throughput (measured in tokens per second,
 421 TPS), and the average number of function evaluations (NFE). The product of throughput and NFE
 422 remains stable across methods and block sizes, indicating an approximately inverse relationship
 423 between these metrics. This observation suggests that throughput can be improved primarily by
 424 reducing NFE, for example, by increasing the parallelism capacity of the denoiser (Wang et al.,
 425 2025c) or by improving the sampling efficiency of the sampler.

426 We observe that the Vanilla decoding, which fixes NFE to the generation budget, maintains almost
 427 identical throughput across block sizes. In contrast, dynamic sampling with either fixed or adap-
 428 tive block sizes benefits from increasing B_0 from 4 to 64, but exhibits a drop in throughput when
 429 the block size is increased further. We attribute this initial positive correlation between block size
 430 and throughput to the mitigation of late-stage decoding overhead. The subsequent degradation in
 431 throughput is mainly due to the noisy output of the denoiser: when a large block size B is used, the
 432 semantic coherence within each block weakens, resulting in fewer tokens being decoded per step

432 and thereby requiring more denoise–sample iterations. Given the approximate inverse relationship
 433 between NFE and throughput, this increase in NFE leads to reduced throughput. The noisy output
 434 is also reflected in the accuracy performance, with the large block size $B_0 = 128$ experiencing a
 435 significant performance degradation.

436 With *AdaBlock-dLLM*, throughput improves for small default block sizes ($B_0 \in \{4, 8\}$) and slightly
 437 decreases for larger defaults ($B_0 \in \{16, 32, 64, 128\}$). *AdaBlock-dLLM* attempts to align the block
 438 size with semantic steps, leading to $B > B_0$ for small defaults and $B < B_0$ for larger defaults.
 439 For small B_0 , *AdaBlock-dLLM* effectively reduces late decoding overhead. As B_0 increases, applying
 440 $B < B_0$ strengthens the local semantic context and improves sampling quality, at the cost of
 441 a modest throughput decrease. Nevertheless, both **Dynamic** and **Dynamic+Ada** achieve substan-
 442 tial speedups over **Vanilla** (top-1) decoding. Figure 6 presents the trade-off between accuracy and
 443 throughput, demonstrating the improvements induced by our method.

444
 445 **Table 2:** Performance comparison across default block sizes B_0 under a generation budget of
 446 $L = 512$. The product of throughput and NFE is nearly identical across methods and block sizes,
 447 indicating an approximately inverse relationship between these quantities when no block-level KV
 448 caching is applied. *AdaBlock-dLLM* yields throughput gains for $B_0 \in \{4, 8\}$. Boldface indicates
 449 superior performance; this convention applies to all tables unless noted otherwise.

	Acc. (%)	TPS	Avg. NFE	TPS×NFE (10^3)	Acc. (%)	TPS	Avg. NFE	TPS×NFE (10^3)	Acc. (%)	TPS	Avg. NFE	TPS×NFE (10^3)
Method	$B_0 = 4$				$B_0 = 8$				$B_0 = 16$			
Vanilla	80.9	16.1	512.0	8.2	80.5	16.1	512.0	8.2	78.8	16.1	512.0	8.2
Dynamic	81.2	43.0	189.4	8.2	80.6	60.0	135.5	8.1	79.1	74.7	109.2	8.2
+Ada	81.6	51.3	159.8	8.2	81.8	63.9	128.3	8.2	80.6	73.9	102.4	8.2
Method	$B_0 = 32$				$B_0 = 64$				$B_0 = 128$			
Vanilla	76.8	16.1	512.0	8.2	76.8	16.1	512.0	8.2	71.0	16.0	512.0	8.2
Dynamic	77.6	85.0	94.9	8.1	77.3	89.4	91.6	8.2	70.7	81.2	101.2	8.2
+Ada	80.6	83.5	98.5	8.2	79.5	87.9	93.4	8.2	72.8	80.6	101.6	8.2

461
 462 **Table 3:** Accuracy (%) on GSM8K for LLaDA-
 463 Instruct across caching methods with $L = 512$.

464 **Table 4:** Accuracy (%) on GSM8K for LLaDA-
 465 Instruct under different generation budgets.

Method	$B_0 = 16$	$B_0 = 32$	$B_0 = 64$	Method	$L = 256$	$L = 512$	$L = 1024$
+PrefixCache	78.2	76.9	75.0	Dynamic	78.1	77.6	77.4
+Ada+PrefixCache	81.4	79.8	77.6	+Ada	78.5	80.6	79.3
+DualCache	78.0	74.5	75.4	+Cache	77.4	74.5	75.8
+Ada+DualCache	80.0	78.5	80.7	+Ada+Cache	79.2	78.5	78.1

470
 471 **Table 5:** Accuracy (%) on GSM8K for LLaDA and Dream
 472 across delimiter thresholds $\tau_D \in \{0.3, 0.5, 0.7\}$ with $B_0 =$
 473 32. A smaller τ_D provides sufficient semantic guidance for
 474 dLLMs trained from scratch (e.g., LLaDA).

Model	$\tau_D = 0.3$	$\tau_D = 0.5$	$\tau_D = 0.7$
LLaDA-Instruct	80.59	79.08	77.94
Dream-Base	75.66	75.74	75.74

475 **Table 6:** Accuracy (%) on IFEval for LLaDA-1.5 across sam-
 476 pling methods. *AdaBlock-dLLM* also improves performance.

Method	$B_0 = 16$	$B_0 = 32$	$B_0 = 64$
Vanilla	69.1	66.7	61.3
Dynamic	69.0	66.7	61.2
+Ada	68.4 _{-0.6}	67.5 _{+0.8}	64.4 _{+3.2}
+Cache	67.5	64.6	59.4
+Ada+Cache	68.9 _{+1.4}	66.4 _{+1.8}	62.7 _{+3.3}

477 **Table 7: Accuracy (%) on GSM8K** across eight different delimiter sets with $B_0 = 32$. Results show
 478 that using the newline token (\n) as the delimiter accounts for most of the accuracy gains, while addition-
 479 ally including the comma and period further improves performance.

Delimiter Set	Acc. (%)
None (+Cache)	74.5
{[\n]}	78.5
{[,]}	75.1
{[.]}	74.5
{[,], [.]}	75.1
{[\n], [,]}	78.5
{[\n], [.]}	78.3
{[\n], [,], [.]}	78.7

486 **Discussion on Difference Between Models.** The gains from *AdaBlock-dLLM* vary across mod-
 487 els. In particular, *AdaBlock-dLLM* yields larger improvements on the LLaDA family of models. By
 488 setting block size according to local semantics, *AdaBlock-dLLM* groups tokens that constitute a se-
 489 mantic step into the same block and focuses on refining the local context. This effect is strongest for
 490 dLLMs trained from scratch, which exhibit greater local stochasticity. In contrast, dLLMs adapted
 491 from AR models display global autoregressive order and a high degree of local autoregressiveness
 492 (Gong et al., 2025). In such a scenario, the improvements from *AdaBlock-dLLM* are correspond-
 493 ingly smaller. Additionally, the fundamental generation quality is limited by the model’s denoising
 494 quality; our work focuses on mitigating the performance loss due to semi-AR decoding.

495 **5.3 ABLATION STUDIES**

496 **Performance Across Different Generation Budgets.** We evaluate *AdaBlock-dLLM* under three
 497 generation budgets: $L \in \{256, 512, 1024\}$, as shown in Table 4. Across all generation budgets
 498 L and decoding settings, *AdaBlock-dLLM* improves generation quality. These results motivate a
 499 semantics-aware block-size scheduling design for dLLMs.

500 **Effect on Delimiter Threshold τ_D .** We evaluate three delimiter thresholds for each model family,
 501 as shown in Table 5. We observe that $\tau_D = 0.3$ yields the best performance in most cases for
 502 LLaDA, whereas a higher $\tau_D = 0.5$ is optimal for Dream. We attribute this difference to the distinct
 503 confidence distributions of the two models. LLaDA is trained purely from scratch and exhibits lower
 504 variance within the volatility band, whereas Dream is adapted from autoregressive models and shows
 505 substantially higher variance (Figure 8). Consequently, different thresholds are required to track the
 506 boundary of a semantic step. Additionally, overly high thresholds (e.g., $\tau_D = 0.9$) often cause the
 507 scheduler to revert to its default behavior, reducing its effectiveness.

508 **Selection of Delimiter Set \mathcal{D} .** We apply more delimiter sets \mathcal{D} to include additional tokens
 509 (comma and period), which often mark the termination of local semantic context. Table 7 shows
 510 that although accuracy improvements vary, the inclusion of the comma and period achieves higher
 511 accuracy than the dynamic sampling baseline. These results highlight the importance of aligning
 512 block size with semantic steps in semi-AR decoding. We provide further analysis in A.4.

513 **Performance on Non-Reasoning Benchmarks.** We further evaluate the performance of
 514 *AdaBlock-dLLM* on IFEval (Zhou et al., 2023), a benchmark that examines the instruction-following
 515 capability of LLMs. As shown in Table 6, with the delimiter set $\mathcal{D} = \{[\backslash n], [,], [.]\}$,
 516 *AdaBlock-dLLM* yields accuracy improvements, especially when integrated with block-level KV
 517 caching. These results suggest that aligning block sizes with semantic steps effectively enhances
 518 sampling quality, leading to improved performance on tasks other than math reasoning or coding.

519 **Limitations.** *AdaBlock-dLLM* effectively enhances the existing semi-AR decoding paradigm by
 520 aligning block sizes with semantic steps. However, the proposed block scheduler may be less effec-
 521 tive when the generation budget is small (e.g., multiple-choice questions), where semi-AR decoding
 522 is not particularly beneficial. Additionally, the decoding process of dLLMs includes both denoising
 523 and sampling. *AdaBlock-dLLM* primarily improves sampling quality; when the denoising outputs
 524 are poor, the benefit of *AdaBlock-dLLM* diminishes.

525 **6 CONCLUSION**

526 This work proposes *AdaBlock-dLLM*, a training-free, plug-and-play scheduler that enhances the
 527 existing semi-autoregressive decoding paradigm. We identify two fundamental limitations of con-
 528 ventional semi-AR decoding (late decoding overhead and premature decoding errors) that motivate
 529 adaptive block-size scheduling. Building on an analysis of confidence dynamics, *AdaBlock-dLLM*
 530 adaptively adjusts the block size at runtime, aligning block sizes with semantic steps. Extensive ex-
 531 periments across benchmarks demonstrate improvements in generation quality of up to 5.3% under
 532 a comparable speed budget. We hope that our semantics-aware adaptive approach and statistical
 533 analysis will inspire future training and inference strategies for dLLMs.

540 ETHICS STATEMENT
541

542 This work adheres to the ICLR Code of Ethics. In this study, no research involving human subjects
543 or animals was conducted. All datasets used, including GSM8K, MATH, HumanEval, and MBPP,
544 were obtained in compliance with relevant usage guidelines, ensuring no violations of privacy. We
545 took care to avoid the bias and discriminatory outcomes throughout our research process. No person-
546 ally identifiable information was used, and no experiments were conducted that could raise privacy
547 or security concerns. We are committed to maintaining transparency and integrity throughout the
548 research process.

549
550 REPRODUCIBILITY STATEMENT
551

552 We ensure that all reported results are reproducible. The code repository is provided in the supple-
553 mentary material to facilitate replication and verification. The experimental setup, including model
554 configurations and hardware details, is described in Section 5.1. We also provide a full description
555 of our methodology, including detailed pseudocode in Algorithm 1, to support understanding and
556 reproduction. Additionally, the open-source models used in this work (e.g., LLaDA and Dream) are
557 publicly available, enabling consistent and reproducible evaluation.

558
559 REFERENCES
560

561 Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Sub-
562 ham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between autore-
563 gressive and diffusion language models. In *The Thirteenth International Conference on Learning
564 Representations*, 2025. URL <https://arxiv.org/abs/2503.09573>.

565 Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Struc-
566 tured denoising diffusion models in discrete state-spaces. In *NeurIPS*, 2021a. URL <https://arxiv.org/abs/2107.03006>.

567 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
568 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language
569 models. *arXiv preprint arXiv:2108.07732*, 2021b.

570 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
571 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
572 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

573 Xinhua Chen, Sitao Huang, Cong Guo, Chiyue Wei, Yintao He, Jianyi Zhang, Hai "Hellen" Li,
574 and Yiran Chen. Dpad: Efficient diffusion language models with suffix dropout, 2025. URL
575 <https://arxiv.org/abs/2508.14148>.

576 Shuang Cheng, Yihan Bian, Dawei Liu, Yuhua Jiang, Yihao Liu, Linfeng Zhang, Wenghai
577 Wang, Qipeng Guo, Kai Chen, Biqing Qi*, and Bowen Zhou. Sdar: A synergistic dif-
578 fusion-autoregression paradigm for scalable sequence generation, 2025. URL <https://github.com/JetAstra/SDAR>.

579 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
580 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
581 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

582 Gemini Diffusion. Gemini diffusion: A state-of-the-art, experimental text diffusion model. <https://deepmind.google/models/gemini-diffusion/>, 2025. Accessed: 2025-09-22.

583 Shansan Gong, Ruixiang Zhang, Huangjie Zheng, Jiatao Gu, Navdeep Jaitly, Lingpeng Kong, and
584 Yizhe Zhang. Diffucoder: Understanding and improving masked diffusion models for code gen-
585 eration. *arXiv preprint arXiv:2506.20639*, 2025.

586 Haoyu He, Katrin Renz, Yong Cao, and Andreas Geiger. Mdpo: Overcoming the training-inference
587 divide of masked diffusion language models. *arXiv preprint arXiv:2508.13148*, 2025.

594 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 595 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*,
 596 2021.

597 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *34th Con-
 598 ference on Neural Information Processi*, 2020. URL <https://arxiv.org/abs/2006.11239>.

600 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
 601 Fleet. Video diffusion models. *Advances in neural information processing systems*, 35:8633–
 602 8646, 2022.

603 Feng Hong, Geng Yu, Yushi Ye, Haicheng Huang, Huangjie Zheng, Ya Zhang, Yanfeng Wang, and
 604 Jiangchao Yao. Wide-in, narrow-out: Revokable decoding for efficient and effective dllms. *arXiv
 605 preprint arXiv:2507.18578*, 2025.

606 Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
 607 and multinomial diffusion: Learning categorical distributions. *arXiv preprint arXiv:2102.05379*,
 608 2021. URL <https://arxiv.org/abs/2102.05379>.

609 Zhanqiu Hu, Jian Meng, Yash Akhauri, Mohamed S Abdelfattah, Jae-sun Seo, Zhiru Zhang, and
 610 Udit Gupta. Accelerating diffusion language model inference via efficient kv caching and guided
 611 diffusion. *arXiv preprint arXiv:2505.21467*, 2025.

612 Daniel Israel, Guy Van den Broeck, and Aditya Grover. Accelerating diffusion llms via adaptive
 613 parallel decoding. *arXiv preprint arXiv:2506.00413*, 2025.

614 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
 615 based generative models. *Advances in neural information processing systems*, 35:26565–26577,
 616 2022.

617 Pengxiang Li, Yefan Zhou, Dilxat Muhtar, Lu Yin, Shilin Yan, Li Shen, Yi Liang, Soroush Vosoughi,
 618 and Shiwei Liu. Diffusion language models know the answer before decoding. *arXiv preprint
 619 arXiv:2508.19982*, 2025a.

620 Shufan Li, Konstantinos Kallidromitis, Hritik Bansal, Akash Gokul, Yusuke Kato, Kazuki Kozuka,
 621 Jason Kuen, Zhe Lin, Kai-Wei Chang, and Aditya Grover. Lavida: A large diffusion language
 622 model for multimodal understanding. *arXiv preprint arXiv:2505.16839*, 2025b.

623 Zhiyuan Liu, Yicun Yang, Yaojie Zhang, Junjie Chen, Chang Zou, Qingyuan Wei, Shaobo Wang,
 624 and Linfeng Zhang. dllm-cache: Accelerating diffusion large language models with adaptive
 625 caching. *arXiv preprint arXiv:2506.06295*, 2025.

626 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
 627 of the data distribution, 2024. URL <https://arxiv.org/abs/2310.16834>.

628 Omer Luxembourg, Haim Permuter, and Eliya Nachmani. Plan for speed–dilated scheduling for
 629 masked diffusion language models. *arXiv preprint arXiv:2506.19037*, 2025.

630 Xinyin Ma, Runpeng Yu, Gongfan Fang, and Xinchao Wang. dkv-cache: The cache for diffusion
 631 language models. *arXiv preprint arXiv:2505.15781*, 2025.

632 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai
 633 Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint
 634 arXiv:2502.09992*, 2025.

635 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of
 636 the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

637 Mihir Prabhudesai, Mengning Wu, Amir Zadeh, Katerina Fragkiadaki, and Deepak Pathak. Diffu-
 638 sion beats autoregressive in data-constrained settings. *arXiv preprint arXiv:2507.15857*, 2025.

639 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 640 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

648 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep un-
 649 supervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. URL <https://arxiv.org/abs/1503.03585>.
 650

651 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 652 Poole. Score-based generative modeling through stochastic differential equations. In *International
 653 Conference on Learning Representations*, 2021.

654

655 Yuerong Song, Xiaoran Liu, Ruixiao Li, Zhigeng Liu, Zengfeng Huang, Qipeng Guo, Ziwei He,
 656 and Xipeng Qiu. Sparse-dllm: Accelerating diffusion llms with dynamic cache eviction. *arXiv
 657 preprint arXiv:2508.02558*, 2025a.

658

659 Yuxuan Song, Zheng Zhang, Cheng Luo, Pengyang Gao, Fan Xia, Hao Luo, Zheng Li, Yuehang
 660 Yang, Hongli Yu, Xingwei Qu, et al. Seed diffusion: A large-scale diffusion language model with
 661 high-speed inference. *arXiv preprint arXiv:2508.02193*, 2025b.

662

663 Guanghan Wang, Yair Schiff, Subham Sekhar Sahoo, and Volodymyr Kuleshov. Remasking discrete
 664 diffusion models with inference-time scaling. In *ICLR 2025 Workshop on Deep Generative Model
 665 in Machine Learning: Theory, Principle and Efficacy*, 2025a.

666

667 Wen Wang, Bozhen Fang, Chenchen Jing, Yongliang Shen, Yangyi Shen, Qiuyu Wang, Hao Ouyang,
 668 Hao Chen, and Chunhua Shen. Time is a feature: Exploiting temporal dynamics in diffusion
 669 language models. *arXiv preprint arXiv:2508.09138*, 2025b.

670

671 Xu Wang, Chenkai Xu, Yijie Jin, Jiachun Jin, Hao Zhang, and Zhijie Deng. Diffusion llms can do
 672 faster-than-ar inference via discrete diffusion forcing. *arXiv preprint arXiv:2508.09192*, 2025c.

673

674 Yinjie Wang, Ling Yang, Bowen Li, Ye Tian, Ke Shen, and Mengdi Wang. Revolutionizing re-
 675 enforcement learning framework for diffusion large language models, 2025d. URL <https://arxiv.org/abs/2509.06949>.

676

677 Qingyan Wei, Yaojie Zhang, Zhiyuan Liu, Dongrui Liu, and Linfeng Zhang. Accelerating dif-
 678 fusion large language models with slowfast: The three golden principles. *arXiv preprint
 679 arXiv:2506.10848*, 2025.

680

681 Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song
 682 Han, and Enze Xie. Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache
 683 and parallel decoding. *arXiv preprint arXiv:2505.22618*, 2025.

684

685 Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada:
 686 Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025.

687

688 Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng
 689 Kong. Dream 7b: Diffusion large language models. *arXiv preprint arXiv:2508.15487*, 2025.

690

691 Runpeng Yu, Xinyin Ma, and Xinchao Wang. Dimple: Discrete diffusion multimodal large language
 692 model with parallel decoding, 2025. URL <https://arxiv.org/abs/2505.16990>.

693

694 Lingzhe Zhang, Liancheng Fang, Chiming Duan, Minghua He, Leyi Pan, Pei Xiao, Shiyu Huang,
 695 Yunpeng Zhai, Xuming Hu, Philip S Yu, et al. A survey on parallel text generation: From parallel
 696 decoding to diffusion language models. *arXiv preprint arXiv:2508.08712*, 2025.

697

698 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
 699 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint
 700 arXiv:2311.07911*, 2023.

701

Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Hu, Jun Zhou, Jianfei
 Chen, Yankai Lin, Ji-Rong Wen, et al. Llada 1.5: Variance-reduced preference optimization for
 702 large language diffusion models. *arXiv preprint arXiv:2505.19223*, 2025.

702 **A APPENDIX**
703704 **A.1 CASE STUDY FOR THE TWO FUNDAMENTAL ISSUES IN SEMI-AUTOREGRESSIVE
705 DECODING**
706

Prompt:	
Carla is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40% of the way through the download, Windows forces a restart to install updates, which takes 20 minutes. Then Carla has to restart the download from the beginning. How long does it take to download the file?	
Decoded Tokens First, let's determine how long it would take Carla to download the entire 200 GB file without any interruptions. Carla's download speed is 2 GB per minute. To download 200 GB, she would need: $\frac{200 \text{ GB}}{2 \text{ GB/minute}} = 100 \text{ minutes}$ (56 tokens decoded)	Decoded Tokens First, let's determine how long it would take Carla to download the entire 200 GB file without any interruptions. Carla's download speed is 2 GB per minute. To download 200 GB, she would need: $\frac{200 \text{ GB}}{2 \text{ GB/minute}} = 100 \text{ minutes}$ However, Windows forces a restart 40% of the way through the download. This means she has to restart the download from the beginning after the restart. (120 tokens decoded)
Current Block (Block No. 8) Denoised Sequence Block Start 2(1.000) 0(1.000) 0(1.000) \0.988) text(0.980) {(0.973) GB(0.977) })(0.984) Block End (8 tokens) More tokens can be decoded without an additional denoising step 2(0.988) \0.980) text(0.949) {(0.949) GB(0.965) /min(0.848) ute(0.703) })(0.969) =(0.957) <space>(0.949) 1(0.938) 0(0.941) 0(0.945) \0.949) text(0.934) {(0.945) minutes(0.930))(0.871) \0.816)](0.840) \n(0.770) \n(0.922) Optimal Block End (30 tokens) \n(0.240) (0.492) we(0.235) need(0.152) to(0.185) the(0.140) the(0.209) the(0.169) download(0.104) the(0.087) download(0.104) (0.102) download(0.089) <space>(0.085) <space>(0.093) 0(0.097) 0(0.104) 0(0.116) 0(0.112) 0(0.111) 0(0.108) 0(0.105) 0(0.111) 0(0.096) 0(0.104) 0(0.103) 0(0.111) 0(0.108) 0(0.112) 0(0.123) 0(0.114) 0(0.109) 0(0.108) ... Late Decoding Overhead	Current Block (Block No. 16) Denoised Sequence Block Start The(1.000) download(0.887) time(0.992) for(0.992) the(1.000) first(0.977) <space>(0.973) 4(0.559) Block End (8 tokens) Token "4" with low confidence is decoded, sub-optimal sampling 0(0.934) %0(0.648) of(0.555) the(0.531) file(0.203) is(0.277) :(0.359) \n(0.342) Optimal Block End (15 tokens) "GB" tokens are in the Volatility Band, but have weak significance to current context \n(0.277) \0.184) {(0.158) {(0.146) 0(0.124) 0(0.177) GB(0.162) GB(0.156) GB(0.205) GB(0.248) GB(0.266) GB(0.250) GB(0.217) GB(0.240) GB(0.299) GB(0.355) GB(0.250) GB(0.165) GB(0.136) GB(0.142) GB(0.126) GB(0.185) GB(0.157) \0.180) \0.131) \0.143) minutes(0.139) \0.162) \0.214) \0.150) \n(0.198) \n(0.198) \n(0.143) \n(0.091) \n(0.075) \n(0.066) \n(0.062) <space>(0.075) minutes(0.079) minutes(0.087) ... Premature Decoding Error

731 Figure 7: A case study of the two fundamental issues (Late Decoding Overhead and Premature
732 Decoding Error). The configuration uses dynamic sampling with a generation budget of $L = 512$
733 and a block size of $B = 8$.
734735 **A.2 SEMI-AR DECODING ALGORITHM WITH ADAPTIVE BLOCK SIZE**
736737 **Algorithm 2** Auxiliary: Denoiser
738

739 **Inputs:** mask predictor p_θ ; vocabulary \mathcal{V} ; index set \mathcal{J} ; current sequence \mathbf{y} .
740 **Outputs:** predicted sequence $\hat{\mathbf{y}}$; confidences \mathbf{c}

741 1: **function** DENOISER(p_θ , \mathcal{V} , \mathcal{J} , \mathbf{y})
742 2: **for** $i \in \mathcal{J}$ **do**
743 3: $\hat{y}_i \leftarrow \arg \max_{v \in \mathcal{V}} p_\theta(v | \mathbf{y}, i)$
744 4: $c_i \leftarrow \max_{v \in \mathcal{V}} p_\theta(v | \mathbf{y}, i)$
745 5: **end for**
746 6: **return** $\hat{\mathbf{y}}$, \mathbf{c}
747 7: **end function**

756

Algorithm 3 Auxiliary: Threshold-Based Dynamic Sampling

757

Inputs: current sequence \mathbf{y} ; predicted sequence $\hat{\mathbf{y}}$; confidences \mathbf{c} ; unmasking threshold τ ; index set \mathcal{J} .
Output: sampled index set \mathcal{S}

```

1: function THRESHOLD-SAMPLE( $\mathbf{y}$ ,  $\hat{\mathbf{y}}$ ,  $\mathbf{c}$ ,  $\tau$ ,  $\mathcal{J}$ )
2:    $\mathcal{S} \leftarrow \emptyset$ 
3:    $\mathcal{J}_{\text{mask}} \leftarrow \{ i \in \mathcal{J} \mid y_i = [\text{MASK}] \}$             $\triangleright$  Select all masked positions in the block
4:   Select index  $i_{\text{top}}$  with highest confidence  $c_i$ 
5:    $\mathcal{S} \cup i_{\text{top}}$ 
6:   for  $i \in \mathcal{J}$  do
7:     if  $i \in \mathcal{J}_{\text{mask}} \wedge c_i \geq \tau$  then
8:        $\mathcal{S} \leftarrow \mathcal{S} \cup \{i\}$ 
9:     end if
10:   end for
11:   return  $\mathcal{S}$ 
12: end function

```

772

Algorithm 4 Adaptive Semi-AR Decoding with Semantic-Aware Block Size

773

Inputs: mask predictor p_θ ; vocabulary \mathcal{V} ; initial sequence $\mathbf{y}^{(T)}$ with index set \mathcal{J} ; generation budget L ; default block size B_0 ; delimiter set \mathcal{D} ; delimiter threshold $\tau_{\mathcal{D}}$; unmasking threshold τ .

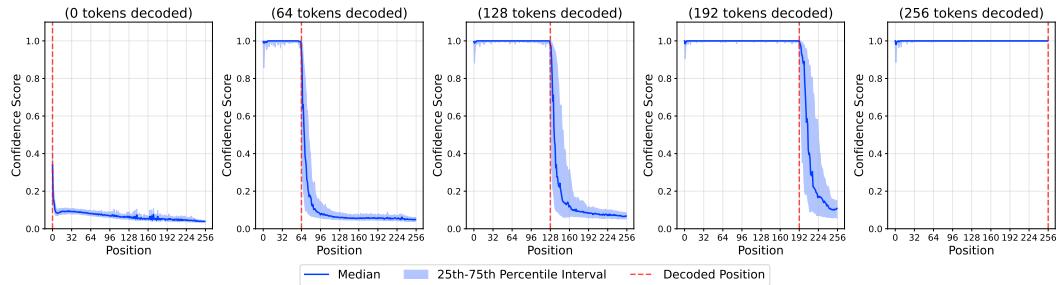
Output: decoded sequence \mathbf{y} .

774

```

1: generated length  $g \leftarrow 0$ , timestep  $t \leftarrow T$ 
2: while  $g < L \wedge t \geq 1$  do
3:
4:    $\triangleright$  First denoising to obtain predicted sequence and confidence at step  $t$ 
5:    $(\hat{\mathbf{y}}^{(t)}, \mathbf{c}^{(t)}) \leftarrow \text{DENOISER}(p_\theta, \mathcal{V}, \mathcal{J}, \mathbf{y}^{(t)})$ 
6:
7:    $\triangleright$  Compute block size
8:    $B \leftarrow \text{COMPUTEBLOCKLENGTH}(\hat{\mathbf{y}}^{(t)}, \mathbf{c}^{(t)}, L, B_0, \mathcal{D}, \tau_{\mathcal{D}}, g)$ 
9:    $\mathcal{J}_{\text{blk}} \leftarrow \{ g, g+1, \dots, g+B-1 \}$ 
10:
11:   $\triangleright$  First sample
12:   $\mathcal{S} \leftarrow \text{THRESHOLD-SAMPLE}(\mathbf{y}^{(t)}, \hat{\mathbf{y}}^{(t)}, \mathbf{c}^{(t)}, \tau, \mathcal{J}_{\text{blk}})$ 
13:  for  $i \in \mathcal{S}$  do
14:     $y_i^{(t-1)} \leftarrow \hat{y}_i^{(t)}$             $\triangleright$  sample tokens with high confidence
15:  end for
16:
17:   $y_j^{(t-1)} \leftarrow y_j^{(t)} \quad \forall j \notin \mathcal{S}, \quad t \leftarrow t - 1$             $\triangleright$  Copy other tokens
18:   $\mathcal{J}_{\text{mask}} \leftarrow \{ i \in \mathcal{J}_{\text{blk}} \mid y_i^{(t)} = [\text{MASK}] \}$ 
19:
20:   $\triangleright$  In-block denoise-sample cycles
21:  while  $\mathcal{J}_{\text{mask}} \neq \emptyset \wedge t \geq 1$  do
22:     $(\hat{\mathbf{y}}^{(t)}, \mathbf{c}^{(t)}) \leftarrow \text{DENOISER}(p_\theta, \mathcal{V}, \mathcal{J}_{\text{mask}}, \mathbf{y}^{(t)})$ 
23:     $\mathcal{S} \leftarrow \text{THRESHOLD-SAMPLE}(\mathbf{y}^{(t)}, \hat{\mathbf{y}}^{(t)}, \mathbf{c}^{(t)}, \tau, \mathcal{J}_{\text{mask}})$ 
24:    for  $i \in \mathcal{S}$  do
25:       $y_i^{(t-1)} \leftarrow \hat{y}_i^{(t)}$ 
26:    end for
27:     $y_j^{(t-1)} \leftarrow y_j^{(t)} \quad \forall j \notin \mathcal{S}, \quad t \leftarrow t - 1$ 
28:     $\mathcal{J}_{\text{mask}} \leftarrow \{ i \in \mathcal{J}_{\text{blk}} \mid y_i^{(t)} = [\text{MASK}] \}$ 
29:  end while
30:
31:   $g \leftarrow g + B$ 
32: end while
33:
34: return  $\mathbf{y}^t$ 

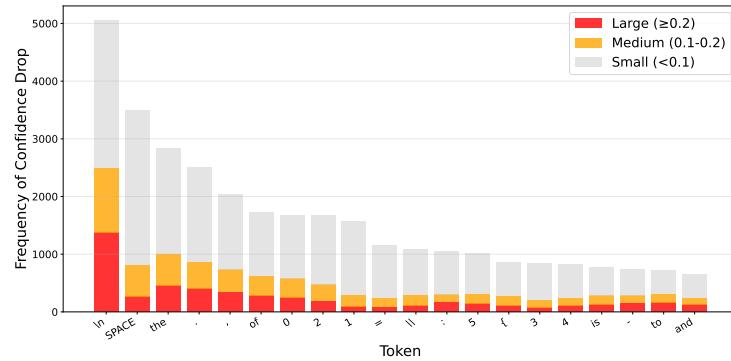
```

810
811 A.3 CONFIDENCE DYNAMICS FOR DREAM-BASE
812

822
823 Figure 8: Confidence scores across sequence positions for Dream-v0-Base-7B, evaluated on 100
824 samples from the GSM8K benchmark. Adapted models such as Dream exhibit a similar degree of
825 global autoregressiveness to LLaDA, but with higher variance in the volatility band. This increased
826 variance motivates the use of a higher delimiter threshold τ_D to provide stronger semantic guidance.
827

828 A.4 FURTHER ANALYSIS OF THE DELIMITER SET
829

830 The choice of `\n` as a delimiter token is also supported by the statistics of confidence drops between
831 consecutive tokens. Large confidence drops indicate sharp semantic boundaries within the volatility
832 band and can thus be used to segment semantic steps. Evaluating 100 samples from the GSM8K
833 dataset, we find that `\n` frequently leads to large confidence drops, indicating semantic boundaries.
834



835
836
837
838
839
840
841
842
843
844
845
846 Figure 9: The frequency of confidence drops between consecutive tokens.
847
848

849 A.5 THE USE OF LARGE LANGUAGE MODELS
850

851 LLMs were used for polishing the manuscript. Specifically, we used an LLM to assist in correcting
852 grammar errors to improve readability. It is important to note that the LLM was not involved in the
853 ideation, research methodology, or experimental design. All intellectual contributions and scientific
854 ideas developed in this work originated from the authors. The contributions of the LLM were solely
855 focused on improving the linguistic quality of the paper, with no involvement in the scientific content
856 or data analysis. The authors take full responsibility for the content of the manuscript, including any
857 text polished by the LLM. We have ensured that the LLM-polished text adheres to ethical guidelines
858 and does not contribute to plagiarism or scientific misconduct.
859
860
861
862
863