DyStruct: Dynamically Structured Diffusion Language Model Decoding via Bayesian Inference

Bian Sun*
Institute for Artificial Intelligence
University of Central Florida
bian.sun@ucf.edu

Kevin Zhai*
Institute for Artificial Intelligence
University of Central Florida
Kevin.Zhai@ucf.edu

Mubarak Shah
Institute for Artificial Intelligence
University of Central Florida
shah@crcv.ucf.edu

Zhenyi Wang†
Institute for Artificial Intelligence
University of Central Florida
zhenyi.wang@ucf.edu

Abstract

Diffusion language models (DLMs) have recently emerged as a promising alternative to autoregressive models, primarily due to their ability to enable parallel decoding. Despite this advantage, most existing DLMs rely on a fixed generation length specified prior to decoding, which restricts their flexibility in real-world applications. While a few recent works attempt to support flexible-length generation, they typically suffer from notable limitations: some require costly retraining to accommodate variable-length outputs, while others depend solely on local confidence signals during decoding. Such local criteria fail to capture the evolving structure of the sequence, often resulting in suboptimal generation quality. In this paper, we propose a training-free, Bayesian structured decoding framework that formulates flexible-length generation as a dynamic structural inference problem. Our approach learns the posterior inference over the dynamic block length, block formations and growth, and block decoding order within a unified Bayesian inference framework to jointly reason about how much to grow, where to grow, and how to organize content structurally. At each window expansion step, the method integrates local uncertainty with structural signals to (i) dynamically expand the sequence via adaptive length growth, (ii) infer block boundaries through Chinese Restaurant Process (CRP)-style partitioning, and (iii) allocate different number of decoding steps for different blocks and determine block decoding order via context-aware scheduling. This yields a unified mechanism that supports dynamic structured generation, including both flexible block expansion and block organization, while maintaining coherence. Extensive experiments across multiple benchmarks demonstrate that our approach significantly improves generation quality and flexibility over existing fixed-length and flexible-length baselines. These results highlight the advantage of Bayesian structured decoding for diffusion language model, providing a principled and efficient solution for structured text generation.

1 Introduction

Most large language models (LLMs) [Brown et al., 2020] rely on autoregressive decoding, where tokens are generated sequentially. This process limits decoding efficiency, especially for long

*Equal contribution.
†Corresponding author

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

sequences, because each new token depends on all previously generated tokens and cannot be produced in parallel. As a result, autoregressive LLMs often suffer from high inference latency and computational cost during deployment. Diffusion language models (DLMs) [Sahoo et al., 2024, Nie et al., 2025b] offer an efficient alternative by enabling parallel decoding. Instead of predicting tokens one by one, diffusion-based approaches iteratively refine multiple token positions simultaneously. This parallel decoding paradigm makes diffusion language models a promising direction for building faster and scalable language generation systems.

However, DLMs typically rely on a fixed, pre-specified generation length. This assumption restricts practical flexibility, as the optimal output length depends on task complexity: complex queries require detailed responses, whereas simpler inputs call for concise outputs. Consequently, fixed-length decoding leads to either truncation or redundancy. More critically, it prevents the model from adapting generation to the evolving semantic context, highlighting the need for mechanisms that dynamically adjust sequence length during generation.

Several recent works attempt to relax this fixed-length assumption in DLMs, but existing approaches exhibit key limitations. FlexMDM [Kim et al., 2026] and DID [Ding et al., 2026] rely on retraining to enable variable-length decoding, which incurs substantial computational cost. DAEDAL [Li et al., 2026] avoids retraining but depends on heuristic, local confidence-based criteria. Crucially, these approaches overlook content organization after sequence expansion. When new tokens are generated, the lack of structural guidance results in fragmented structure.

To overcome these limitations, we formulate variable-length generation as structured decoding via Bayesian inference. We model the joint posterior distribution over the new window expansion size, the partition of the window into contiguous blocks, and the block decoding schedule. We introduce a structured prior over the latent block partition to govern content organization and guide decoding. This prior encourages coherent partition patterns while avoiding rigid assumptions about the number or boundaries of blocks. Specifically, we model block formation through a Chinese Restaurant Process (CRP) [Blei et al., 2010]. The advantages of the CRP are threefold: (1) it removes the need to preset the number of blocks, letting the model determine the quantity adaptively; (2) it does not require predefined partition boundaries, allowing the model to infer splits directly from the data; and (3) it provides a predictive prior over the partition structure, allowing decoding to assess at each step whether the current token should continue the current block or initiate a new one. This framework provides a training-free mechanism for sequence growth, allowing the model to jointly determine how much content to introduce, where to expand, and how newly generated tokens should be organized into contiguous blocks. By unifying local evidence with structural constraints, our approach enables flexible, coherent decoding without modifying the underlying model parameters. The overall algorithm is illustrated in Figure 1.

We evaluate this framework across diverse language generation tasks. The results demonstrate that performing joint structural inference at decoding time actively prevents sequence fragmentation, improving both generation quality and coherence while keeping the model completely frozen.

Our contributions are summarized as follows:

- We introduce a training-free Bayesian framework for dynamic structured decoding in diffusion language models, formulating flexible-length generation as joint inference over the new window size, block partitions, and decoding organization.
- We develop an efficient posterior inference algorithm to estimate dynamic window expansion, block partitioning via the CRP, and blocks decoding scheduling via context-aware prioritization.
- We evaluate the method across multiple datasets, demonstrating improved flexible-length generation quality and coherence without additional training.

2 Related Work

Diffusion Language Models (DLMs). Recent advances establish DLMs by applying denoising diffusion probabilistic models [Ho et al., 2020] through masked discrete formulations, improved training objectives and large-scale pretrained models. These developments demonstrate that diffusion is both a controllable generation framework and a viable foundation-modeling paradigm for language [Hoogeboom et al., 2021, Li et al., 2022, Yu et al., 2022, Savinov et al., 2022, Reid et al., 2023, Gulrajani and Hashimoto, 2023, He et al., 2023, Gong et al., 2023, Lovelace et al., 2023, Gat et al.,

2024, Sahoo et al., 2024, Lou et al., 2024, Liu et al., 2024, Shi et al., 2024, Nie et al., 2025a,b, Liu et al., 2025a, Ye et al., 2025a, Xu et al., 2025, Gong et al., 2025, Deschenaux and Gulcehre, 2025, von Rütte et al., 2025, Liu et al., 2025b, Arriola et al., 2025, Zheng et al., 2025, Sahoo et al., 2025, ZHANG et al., 2025, Kim et al., 2025, Rout et al., 2025, Seo et al., 2025]. Despite this empirical success, DLMs share two practical limitations. First, most existing approaches operate under a fixed-length decoding setting, restricting real-world applicability where generation length must adapt to task complexity. Second, parallel generation in these models introduces distributional drift due to the conditional independence assumption across tokens [Guo and Ermon, 2026]. Recent strategies like Hierarchy-dLLM [Qi et al., 2026] attempt to mitigate this drift via hierarchical decoding. However, it relies on heuristic, position-based rules under a *fixed-length* setting and lacks explicit modeling of decoding structure or content planning. In contrast, we formulate decoding as a Bayesian structural inference problem, jointly inferring new window size, block partitioning, and decoding order within a unified probabilistic framework. This framework enables dynamic length adaptation and coherent content organization, moving beyond local spatial heuristics toward structure-aware generation.

Variable-Length Generation in DLMs. Recent works attempt to relax the fixed-length constraint, but existing approaches exhibit key limitations. DID [Ding et al., 2026] and FlexMDM [Kim et al., 2026] enable dynamic token adjustments during generation, but lack an explicit model of decoding structure and require extensive retraining or alterations to the forward process. DAEDAL [Li et al., 2026] avoids retraining but relies on heuristic, local confidence-based rules to expand the sequence. By contrast, our training-free approach replaces local heuristics with a joint inference framework, allowing the frozen model to determine where to expand, how much to expand, and how new content organizes into contiguous blocks.

3 Problem Formulation

We formulate the structured decoding problem. DLMs generate sequences through iterative refinement. Let x denote the input prompt and let $\mathcal{V}$ denote the vocabulary. We define $t \in \{1, \dots, T_{ext}\}$ as the expansion step index. After step $t-1$, the current response is $y^{(t-1)} = (y_1^{(t-1)}, \dots, y_{n_{t-1}}^{(t-1)})$, where each position is either a vocabulary token or `[MASK]`. To predict the masked values, the model processes the concatenated sequence $[x; y^{(t-1)}]$. To enable flexible-length generation, the decoder appends a new masked window of length L_t at expansion step t: $\tilde{y}^{(t)} = \big[y^{(t-1)}; \underbrace{\texttt{[MASK]}, \dots, \texttt{[MASK]}}_{L_t \text{ new positions}}\big]$.

We index positions inside this newly appended window locally by $j = 1, \dots, L_t$, corresponding to global indices $n_{t-1} + j$ in $\tilde{y}^{(t)}$. Appending this window introduces a structural inference problem. At each expansion step t, the decoder needs to infer: (i) the allocated window length L_t, (ii) a partition $\mathcal{P}^{(t)} = \{B_1^{(t)}, \dots, B_{M_t}^{(t)}\}$ dividing the window into M_t contiguous blocks, and (iii) a schedule $\tau^{(t)}$, which is a permutation of $\{1, \dots, M_t\}$ denoting the decoding order. To determine the partition $\mathcal{P}^{(t)}$, the decoder utilizes a CRP prior, governed by a concentration parameter α, to evaluate whether adjacent positions should extend an existing block or initialize a new block. Once the structure $(\mathcal{P}^{(t)}, \tau^{(t)})$ is established, the decoder decodes each block through a series of unmasking iterations, the total count of which is dynamically determined based on block instability. We detail the notation in Appendix Table 5.

4 Method

4.1 Dynamic Structured Decoding as Bayesian Inference

We frame flexible-length diffusion decoding as a dynamic structural inference problem. We formulate this organization as a unified Bayesian inference over latent variables $Z^{(t)} = \{L_t, \mathcal{P}^{(t)}, \tau^{(t)}\}$. We denote $O^{(t)}$ as the set of statistics derived from a temporary diagnostic pass that summarizes uncertainty and structural boundary evidence in the current window. We model the prior over these latent variables as a structured factorization: $p\big(Z^{(t)}\big) = p\big(L_t\big)\; p\big(\mathcal{P}^{(t)} \mid L_t, \alpha\big)\; p\big(\tau^{(t)} \mid \mathcal{P}^{(t)}\big)$. Here, $p(L_t)$ defines a prior over the window growth size L_t. The term $p(\mathcal{P}^{(t)} \mid L_t, \alpha)$ corresponds to a Chinese Restaurant Process (CRP) prior [Blei et al., 2010] over block partitions $\mathcal{P}^{(t)}$ across the local indices $W^{(t)} = \{1, \dots, L_t\}$, where α denotes the concentration parameter. This CRP prior

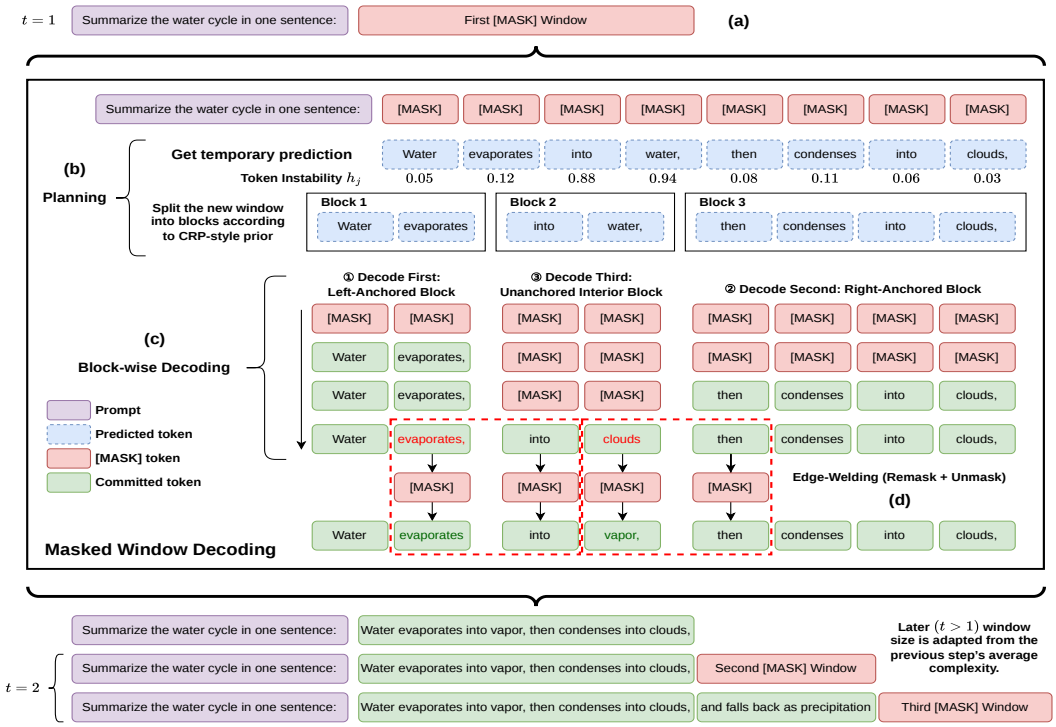


Figure 1: **Overview of DyStruct.** The framework performs flexible-length decoding by iteratively appending masked windows and executing structural inference. **(a) Window Expansion:** The next window size adaptively scales based on the mean instability ($\bar{h}$) of previously decoded tokens. **(b) CRP-Style Partitioning:** A short temporary pass extracts token-level instability scores (h_j), which a CRP-style prior uses to partition the unanchored window into contiguous blocks. **(c) Context-Aware Scheduling:** Partitioned blocks are decoded according to a schedule that prioritizes stable, anchored segments. **(d) Local Boundary Repair (Edge-Welding):** To align predictive distributions at block interfaces, the decoder performs localized remasking (red dashed boxes) to ensure structural consistency.

favors coherent contiguous blocks while allowing flexible splitting when supported by edge evidence. Finally, $p(\tau^{(t)} \mid \mathcal{P}^{(t)})$ defines a preference over the block decoding schedule $\tau^{(t)}$.

Given the prompt x, the previously generated sequence $y^{(t-1)}$, and the diagnostic observations $O^{(t)}$, we perform posterior inference over the latent structure:

$$p\big(Z^{(t)} \mid O^{(t)}, y^{(t-1)}, x\big) \propto p\big(O^{(t)} \mid L_t, \mathcal{P}^{(t)}, \tau^{(t)}, y^{(t-1)}, x\big)\; p\big(L_t\big)\; p\big(\mathcal{P}^{(t)} \mid L_t, \alpha\big)\; p\big(\tau^{(t)} \mid \mathcal{P}^{(t)}\big).$$

The structural progression of this inference is depicted in Figure 1 and summarized in Algorithm 1. A detailed algorithm is provided in Appendix B. To estimate $p\big(Z^{(t)} \mid O^{(t)}, y^{(t-1)}, x\big)$, in the following sections, we will decompose $p\big(Z^{(t)} \mid O^{(t)}, y^{(t-1)}, x\big)$ and present how to estimate the three decomposed posteriors: $p(L_t|O^{(t)}, y^{(t-1)}, x)$, $p\big(\mathcal{P}^{(t)} \mid O^{(t)}, L_t, \alpha\big)$, and $p\big(\tau^{(t)} \mid \mathcal{P}^{(t)}, O^{(t)}\big)$.

4.2 Latent Block Formation and Growth

Posterior over the new window size. At each step t, the decoder determines the expansion length L_t. A stable preceding window permits a larger window expansion, whereas an unstable window restricts expansion to a smaller window. To capture this dynamic scaling, we summarize the preceding window using its finalized mean instability value $\bar{h}^{(t-1)} \in [0, 1]$, where larger values indicate structural instability. We model the posterior distribution over the next window length as a Poisson-distributed random variable:

$$L_t \sim p(L_t \mid O^{(t)}, y^{(t-1)}, x) \approx \text{Poisson}(\mu_t), \quad \text{where} \quad \mu_t = L_{\min} + \big(1 - \bar{h}^{(t-1)}\big)\big(L_{\max} - L_{\min}\big), \tag{1}$$

clipped to $[L_{\min}, L_{\max}]$.

Algorithm 1 DyStruct: Dynamic Structured Decoding

Input: Prompt x, global length limit $N_{\max}$, hyperparameters $(\alpha_0, \gamma, r_{\text{weld}})$
Initialize: $y^{(0)} \leftarrow x$, step $t \leftarrow 1$, previous window instability $\bar{h}^{(0)} \leftarrow 0.5$
while `[EOS]` not generated **and** $|y^{(t-1)}| < N_{\max}$ **do**
  Sample expansion length $L_t \sim p(L_t \mid O^{(t)}, y^{(t-1)}, x)$ via Eq. 1
  Append L_t `[MASK]` tokens to $y^{(t-1)}$ to form window indices $W^{(t)}$
  Execute temporary diagnostic pass over $W^{(t)}$ to extract feature signals $\phi^{(t)}$
  Compute instability $h_j^{(t)}$ (Eq. 2) and edge scores $q_g^{(t)}$ for all $j, g \in W^{(t)}$
  Partition $W^{(t)}$ into blocks $\mathcal{P}^{(t)}$ using CRP (Eq. 9)
  Order blocks into schedule $\tau^{(t)}$ via Gibbs distribution (Eq. 10)
  for each block $B \in \tau^{(t)}$ **do**
    Predict and commit tokens over $T(B)$ refinement steps
  end for
  Perform localized edge-welding at shared block boundaries (Eq. 11)
  Calculate finalized mean instability $\bar{h}^{(t)}$ over $W^{(t)}$, update sequence $y^{(t)}$, increment t
end while
Output: Final generated sequence $y^{(\text{final})}$

Statistics for characterizing block boundary changes. Before decoding the new window, the model executes a short sequence of temporary diagnostic steps over indices $W^{(t)}$ to assess positional instability under partial context. At each step, the frozen model predicts all masked positions, commits a fraction of the tokens with the highest confidence, and remasks the remainder. For each position $j \in \{1, \ldots, L_t\}$, this diagnostic pass produces a feature vector $\phi_j^{(t)} \in \mathbb{R}^R$ capturing observable signals including entropy, prediction shifts, hidden state variation, and confidence. This vector is projected to a scalar $u_j^{(t)}$ using an estimated weight w (illustrated in Appendix D), and normalized via a logistic function to obtain a local instability score $h_j^{(t)}$:

$$u_j^{(t)} = w^\top \phi_j^{(t)}, \quad h_j^{(t)} = \sigma\left(u_j^{(t)} - \frac{1}{L_t}\sum_{r=1}^{L_t} u_r^{(t)}\right). \tag{2}$$

A larger $h_j^{(t)}$ indicates higher uncertainty relative to the window. To determine block boundaries, we evaluate the gaps between adjacent tokens. For each gap $g \in \{1, \ldots, L_t - 1\}$, we construct a feature vector $\psi_g^{(t)}$ (illustrated in Appendix C) and project it using an estimated boundary weight vector w_b to compute an edge score:

$$\ell_g^{(t)} = w_b^\top \psi_g^{(t)}, \quad q_g^{(t)} = \sigma\left(\ell_g^{(t)}\right). \tag{3}$$

Larger $q_g^{(t)}$ values indicate stronger probabilistic evidence for placing a boundary at gap g. The window is partitioned using a CRP-inspired prior. We define a local concentration parameter (where α_0 is the empirical base constant):

$$\alpha_g^{(t)} = \alpha_0 \exp\left(\bar{h}^{(t-1)} + \ell_g^{(t)} - \frac{1}{L_t - 1}\sum_{r=1}^{L_t-1} \ell_r^{(t)}\right). \tag{4}$$

Here, $\bar{h}^{(t-1)}$ controls the overall split rate (higher instability results in more blocks), while $\ell_g^{(t)}$ scores the likelihood of splitting at specific gaps.

Prior over block partitions. To model how tokens are grouped into blocks, we place a prior over partitions using the Chinese Restaurant Process (CRP) [Blei et al., 2010]. The key idea is: tokens are assigned sequentially, and each token either joins the current block or starts a new one. Joining the current block is more likely when the block is already large, while starting a new block is controlled by a concentration parameter $\alpha_g^{(t)}$. This naturally balances *block growth* and *block creation*. At expansion step t, the newly allocated window has length L_t and is partitioned into contiguous blocks $\mathcal{P}^{(t)} = \{B_1^{(t)}, \ldots, B_{M_t}^{(t)}\}$. We implement the CRP prior through decisions at each of the $L_t - 1$ gaps between adjacent positions. At each gap g, the decoder decides whether to continue the current block

("stay") or start a new block ("cut"). Let $b_g \in \{0, 1\}$ denote this decision, with $b_g = 1$ indicating a cut. If the current block has length m_g, we define:

$$p(b_g = 0 \mid m_g, \alpha_g^{(t)}) = \frac{m_g}{m_g + \alpha_g^{(t)}}, \quad p(b_g = 1 \mid m_g, \alpha_g^{(t)}) = \frac{\alpha_g^{(t)}}{m_g + \alpha_g^{(t)}}. \tag{5}$$

This formulation has two practical advantages that are important for decoding. First, it *does not require fixing the number of blocks* in advance; the model automatically determines how many blocks are needed. Second, it *does not assume fixed boundaries*; instead, boundaries are inferred dynamically based on local evidence through $\alpha_g^{(t)}$. As a result, the prior encourages coherent block growth (by favoring "stay" for large m_g) while still allowing new blocks when necessary. The prior probability of a partition $\mathcal{P}^{(t)}$ is then given by:

$$p\big(\mathcal{P}^{(t)} \mid L_t, \alpha_g^{(t)}\big) = \prod_{g=1}^{L_t - 1} \left(\frac{\alpha_g^{(t)}}{m_g + \alpha_g^{(t)}}\right)^{b_g} \left(\frac{m_g}{m_g + \alpha_g^{(t)}}\right)^{1 - b_g}. \tag{6}$$

Likelihood of diagnostic observations given a block partition. For each gap g, we compute an edge probability $q_g^{(t)} = \sigma(\ell_g^{(t)})$, which we interpret as: $p(b_g = 1 \mid O^{(t)}) = q_g^{(t)}, \quad p(b_g = 0 \mid O^{(t)}) = 1 - q_g^{(t)}$. The likelihood of a given partition is:

$$p\big(O^{(t)} \mid \mathcal{P}^{(t)}\big) = \prod_{g=1}^{L_t - 1} \big(q_g^{(t)}\big)^{b_g} \big(1 - q_g^{(t)}\big)^{1 - b_g}. \tag{7}$$

Posterior over block partitions. Combining the likelihood and the prior, the posterior over partitions is defined as: $p\big(\mathcal{P}^{(t)} \mid O^{(t)}, L_t, \alpha_g^{(t)}\big) \propto p\big(O^{(t)} \mid \mathcal{P}^{(t)}\big)\, p\big(\mathcal{P}^{(t)} \mid L_t, \alpha_g^{(t)}\big)$.

Taking the logarithm, we obtain the objective function:

$$\begin{aligned} \log p\big(\mathcal{P}^{(t)} \mid O^{(t)}, L_t, \alpha_g^{(t)}\big) &= \sum_{g=1}^{L_t - 1} \Big[b_g \log q_g^{(t)} + (1 - b_g) \log \big(1 - q_g^{(t)}\big)\Big] \\ &+ \sum_{g=1}^{L_t - 1} \Big[b_g \log \frac{\alpha_g^{(t)}}{m_g + \alpha_g^{(t)}} + (1 - b_g) \log \frac{m_g}{m_g + \alpha_g^{(t)}}\Big] + A, \quad \text{where } A \text{ is a constant.} \end{aligned} \tag{8}$$

Maximum a posteriori inference determines the block split positions. The final partition is obtained via maximum a posteriori (MAP) inference:

$$\arg\max_{\mathcal{P}^{(t)}} \Big[\log p\big(\mathcal{P}^{(t)} \mid O^{(t)}, L_t, \alpha_g^{(t)}\big)\Big] \tag{9}$$

This objective explicitly grounds the algorithm: the gap feature vectors provide the likelihood evidence for a split, while the CRP prior enforces contiguous block partitions. Given the resulting split points $g_1 < \cdots < g_{M_t - 1}$, the contiguous blocks $\mathcal{P}^{(t)} = \{B_1^{(t)}, \ldots, B_{M_t}^{(t)}\}$ are defined as:

$$B_1^{(t)} = \{1, \ldots, g_1\}, \quad B_2^{(t)} = \{g_1 + 1, \ldots, g_2\}, \quad \ldots, \quad B_{M_t}^{(t)} = \{g_{M_t - 1} + 1, \ldots, L_t\}.$$

4.3 Blockwise Decoding Schedule and Edge Welding

Posterior over the block decoding schedule. Given a fixed window partition, the decoder assigns each block a refinement budget and a decoding order. For a block $B \subseteq W^{(t)}$, we define the block instability as $H(B) = \frac{1}{|B|} \sum_{j \in B} h_j^{(t)}$. The total refinement steps $T(B)$ allocated to the block is obtained by linearly interpolating between $T_{\min}$ and $T_{\max}$ using $H(B)$.

To determine the decoding order, we measure how well a block is anchored by neighboring decoded tokens. Let $C(B) \in \{0, 0.5, 1\}$ represent the context proximity: $C(B) = 1$ if both sides are anchored,

0.5 if one side is anchored, and 0 if bounded entirely by masks. The schedule $\tau^{(t)}$ follows a Gibbs distribution:

$$p\Big(\tau^{(t)} \mid \mathcal{P}^{(t)}, O^{(t)}, y^{(t-1)}, x\Big) \propto \exp\left(\sum_{B\in\mathcal{P}^{(t)}} \rho(B)\right), \quad \rho(B) = -H(B) + \gamma C(B). \tag{10}$$

This schedule prioritizes anchored blocks ($\gamma C(B)$) with low instability ($H(B)$), where γ is a context weight. This ordering allows the resulting decoded tokens to serve as stable context that constrains the subsequent decoding of regions with higher instability. Within each scheduled block, the model iteratively commits tokens with high confidence while refining the remaining masked positions.

Boundary reconciliation via edge-welding. To ensure distributional consistency across independently scheduled blocks, we apply a local edge-welding step. For neighboring blocks $B_m^{(t)} = [a, b)$ and $B_{m+1}^{(t)} = [b, c)$, we define an interval around the boundary:

$$E_m^{(t)} = \big[\max(a, b - r_{\text{weld}}), \min(c, b + r_{\text{weld}})\big]. \tag{11}$$

where r_{weld} defines the fixed boundary repair radius. Within this interval, tokens with low confidence are remasked and locally refined, while all positions outside the interval remain fixed. This step aligns boundary predictions without modifying the established blocks. After welding is complete, the decoder calculates the updated mean instability $\bar{h}^{(t)}$ to control the subsequent expansion step.

5 Experiments

Baselines and implementation. We evaluate DyStruct using LLaDA-8B-Base [Nie et al., 2025b] and Dream-7B-Base [Ye et al., 2025b]. To ensure a fair computational comparison, we cap the base unmasking iterations and maximum sequence length at 256. For DyStruct, this iteration limit serves as the budget for dynamic window expansion, whereas baseline models are constrained to generate a rigid 256-token window. We also compare against DAEDAL [Li et al., 2026], implementing it as a heuristic plug-in expansion module to represent existing training-free variable-length diffusion methods. DyStruct maintains uniform hyperparameters across all tasks (detailed in Appendix E). Experiments run on a single NVIDIA H100 GPU.

Benchmarks and datasets. To assess generalizability, we benchmark DyStruct across three domains: mathematical reasoning, code synthesis, and symbolic logic. We quantify mathematical reasoning using GSM8K [Cobbe et al., 2021] and MATH [Hendrycks et al., 2021], reporting strict-match accuracy. For code generation, we use MBPP [Austin et al., 2021] and HumanEval [Chen et al., 2021], reporting greedy pass@1 accuracy. Finally, we evaluate multi-step logical reasoning on Big-Bench Hard (BBH) [Suzgun et al., 2023] using exact match accuracy. All evaluations run within the LM-Evaluation-Harness [Gao et al., 2023].

5.1 Main Results

As shown in Table 1, DyStruct improves across all five benchmark tests, among which the score on BBH increases by 4.4 % (44.9 $\rightarrow$ 49.3). DAEDAL shows limited improvement on LLaDA-8B-Base and even degrades performance on the BBH task. Besides, to further verify whether this improvement is statistically reliable at the prompt level, we additionally conduct paired McNemar tests between DyStruct and DAEDAL on the LLaDA-8B-Base backbone, as reported in Appendix H. The paired McNemar tests show that DyStruct achieves statistically significant prompt-level improvements over DAEDAL on BBH and two Math datasets. For Code, DyStruct also exhibits a positive prompt-level trend. Beyond accuracy, this performance gap suggests the instability of DAEDAL when applied to stronger backbone models. Meanwhile, the results on the more advanced Dream-7B-Base backbone further show that DyStruct can be effectively transferred beyond LLaDA-8B-Base, maintaining consistent improvements under a stronger model setting. This comparison demonstrates the limitations of heuristic rules in cross-model applications, while highlighting the adaptability of DyStruct in achieving robust gains through a joint adaptive mechanism of sequence length control and structured partitioning during the reasoning stage.

Dynamic structure and instability-aware scheduling improve inference performance. Beyond the aggregate accuracy gains, the results also suggest that DyStruct improves efficiency by allocating computation according to structural instability rather than applying a uniform decoding

Table 1: **Dynamic structured decoding outperforms baselines.** DAEDAL serves as a heuristic length-adjustment decoding baseline. Across these settings, we keep all shared decoding and evaluation parameters identical wherever applicable, so that performance differences primarily reflect the effect of the method-specific decoding mechanisms. Besides, to ensure a fair computational comparison, all models utilize a base budget of 256 unmasking iterations and a maximum generation limit of 256 tokens. DyStruct dynamically scales this base iteration budget across adaptively sized blocks based on block instability. Numbers in parentheses indicate default few-shot examples. Values in gray are standard errors (SE) from `lm-eval` [Gao et al., 2023].

	Code Ability		Mathematics		General
Model	HumanEval (0)	MBPP (3)	GSM8K (5)	MATH (4)	BBH (3)
LLaDA-8B-Base	32.3 ±3.6	39.8 ±2.1	70.3 ±1.2	30.5 ±0.6	44.9 ±0.5
LLaDA-8B-Base + DAEDAL	33.5 ±3.7	40.2 ±2.1	70.8 ±1.2	31.2 ±0.6	43.7 ±0.5
LLaDA-8B-Base + DyStruct	**34.8** ±3.7	**41.4** ±2.2	**72.1** ±1.2	**31.4** ±0.6	**49.3** ±0.5
Dream-7B-Base	40.2 ±2.7	57.2 ±2.2	74.9 ±1.1	38.2 ±0.6	51.7 ±0.5
Dream-7B-Base + DAEDAL	34.7 ±2.7	54.4 ±2.2	74.3 ±1.2	38.6 ±0.6	44.8 ±0.5
Dream-7B-Base + DyStruct	**47.0** ±3.9	**59.8** ±2.1	**75.1** ±1.9	**38.8** ±0.6	**52.5** ±0.5

Table 2: **Ablation study of structural decoding components.** Both the Bayesian block partitioning and the local boundary repair step are necessary to maintain reasoning accuracy. W/o Block Decoding Schedule replaces the context-aware block scheduling with a fixed left-to-right block order.

	Coding Ability		Mathematics		General
Configuration	HumanEval	MBPP	GSM8K	MATH	BBH
Full DyStruct	**34.8**	**41.4**	**72.1**	**31.4**	**49.3**
w/o Block Decoding Schedule	33.5	41.2	71.8	30.3	49.3
w/o Boundary Repair	32.9	40.8	70.5	31.2	49.1
w/o Block Decoding Schedule & Boundary Repair	32.9	41.0	69.7	31.0	49.1
LLaDA-8B-Base	32.3	40.2	70.3	30.5	44.9

budget. This is particularly reflected on BBH, the most reasoning-intensive benchmark in our evaluation, where DyStruct achieves the largest improvement. Since BBH contains more complex multi-step reasoning and logical composition, this gain is consistent with the role of instability-aware scheduling: high-difficulty regions, such as logical operators, intermediate reasoning steps, and structured expressions, receive more refinement, while easier regions are decoded with fewer redundant updates. Figure 2 shows that this adaptive allocation also leads to faster inference on GSM8K. Because GSM8K mainly consists of structured mathematical reasoning problems, many output regions become stable early once the key arithmetic or reasoning pattern is resolved. DyStruct can therefore avoid repeatedly denoising the entire fixed-length window and instead spend computation only where additional refinement is needed. This explains why DyStruct achieves lower per-question inference time than both the base decoding setup and DAEDAL on LLaDA-8B and Dream-7B. Together, these results show that DyStruct does not merely improve accuracy on difficult reasoning tasks such as BBH, but also reduces unnecessary computation on mathematical tasks by dynamically matching the decoding budget to block instability.

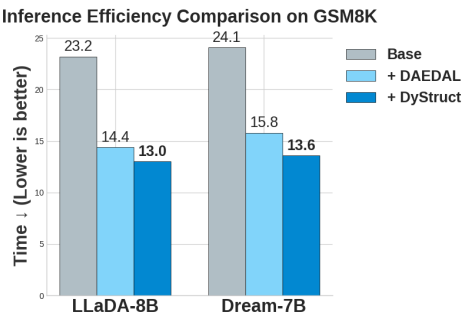


Figure 2: **Inference efficiency comparison.** The proposed DyStruct achieves the lowest inference time across different backbone models on the GSM8K dataset. Time is reported in seconds per iteration (s/it).

Structural components drive stable decoding performance. Table 2 isolates the contribution of each structural mechanism. Removing block decoding leads to lower mathematical reasoning accuracy on both benchmarks, as the procedure loses the benefit of dynamically prioritizing structurally stable blocks during decoding. Furthermore,

Table 3: **Effect of concentration parameter.** The default setting is shaded in gray. *Toks* denotes the average number of newly generated tokens per sample, and *Blks* denotes the rounded average number of finalized blocks per sample. We vary the concentration parameter α_0 and the results remain stable across different α_0.

	HumanEval			BBH		
α_0	Pass@1↑	Toks↓	Blks	Acc.↑	Toks↓	Blks
1.2	35.4	217	7	49.4	246	7
1.5	34.8	219	7	49.3	246	7
1.8	33.5	220	7	49.5	245	6
2.0	34.1	221	6	49.3	245	6

Table 4: **Effect of initial sequence length.** We vary the initial sequence length to evaluate the stability of our method. Performance remains relatively stable across different initial sequence length. The default setting is shaded in gray.

	HumanEval			BBH		
Len.	Pass@1↑	Toks↓	Blks	Acc.↑	Toks↓	Blks
32	33.5	220	8	48.8	239	7
48	34.8	219	7	49.3	246	7
64	33.5	216	6	49.0	246	6
128	35.4	225	5	48.9	248	5
256	36.6	256	7	46.3	256	3

removing the local boundary repair step degrades code generation performance, reducing HumanEval accuracy by 1.9 points. This performance degradation demonstrates that independently scheduled blocks require localized edge-welding to align their predictive distributions; without it, adjacent blocks remain structurally incompatible. When both the block decoding schedule and boundary repair are removed, the degradation becomes more apparent, with GSM8K further decreasing to 69.7% and MATH to 31.0%. Notably, even after removing these components, the remaining CRP-based block partitioning mechanism still outperforms the LLaDA-8B-Base baseline on most benchmarks and maintains strong performance on BBH. This indicates that the core structural partitioning itself already provides a robust benefit, while the block decoding schedule and boundary repair further improve decoding stability and task performance.

5.2 Ablation, Sensitivity and Qualitative Analysis

Adaptive window expansion mitigates fragmentation without sacrificing accuracy. As demonstrated in Table 4, DyStruct remains robust across different initial window lengths. Although HumanEval reaches its highest Pass@1 under the full 256-token setting, the selected length of 48 provides a stronger overall balance, maintaining competitive HumanEval performance while achieving the best BBH accuracy. Table 3 further supports the mechanism behind this behavior. Across different α_0 and initial length settings, DyStruct maintains stable performance while using a compact effective token budget and a small number of blocks in most configurations, rather than relying on uniformly expanding the full 256-token window. This indicates that the model can adapt the active decoding region according to the local structure of the output. By partitioning the window into contiguous blocks and scheduling computation based on local instability h_j, DyStruct decouples generation length from the global computational budget. Instead of assigning the same number of unmasking iterations to the entire sequence, the total iterations for each block are interpolated from a base iteration count and the mean instability score $H(B)$. As a result, high-instability segments receive near-maximum iteration counts, whereas stable linguistic transitions and low-complexity regions receive fewer unmasking iterations. This explains why DyStruct can improve difficult reasoning tasks such as BBH while also reducing redundant computation on simpler regions.

Hyperparameters balance granularity and context. The effect of the CRP concentration prior α_0 has been analyzed above in Table 3, where DyStruct shows stable performance under variations of the partition granularity. Beyond α_0, DyStruct relies on several architectural hyperparameters to govern block resolution. The welding radius r_{weld} defines the local interval used for boundary repair. We provide additional ablations of these architectural hyperparameters, including the welding radius r_{weld}, in Appendix G. These results show that DyStruct remains stable under variations of their values; only overly small or overly large settings lead to slight performance changes. Thus, we adopt fixed values that provide stable partitioning across all datasets rather than tuning them for task-specific optimality, demonstrating the robustness of the structural formulation without heavy hyperparameter sweeping.

Qualitative Results To illustrate how DyStruct treats flexible-length generation as a problem of dynamic structure, we provide examples in Appendix F. Cases demonstrate the execution of CRP-inspired block partitioning, context-aware scheduling, and localized edge-welding. While baseline models frequently commit to context-brittle tokens early in the sequence and propagate structural

failure forward, DyStruct utilizes local boundary repair to identify and revise conflicting boundaries as the context evolves, recovering structurally valid outputs on complex mathematical prompts.

Efficiency Analysis As shown in Figure 2 , DyStruct reduces inference latency on GSM8K from 23.2s to 13.0s for LLaDA-8B-Base and from 24.1s to 13.6s for Dream-7B-Base, outperforming both standard decoding and DAEDAL. Although DyStruct introduces slightly more model forward passes, its adaptive window expansion starts with 48 tokens and progressively extends the sequence, avoiding repeated computation over a fixed 256-token window. As shown in Table 3, the average realized generation lengths are approximately 219 tokens on HumanEval and 246 on BBH. Thus, the additional forward passes process shorter sequences on average and therefore are cheaper in terms of FLOPs.

6 Conclusion

This paper presents a principled Bayesian framework for flexible-length diffusion language models (DLMs). We formulate flexible-length generation as a joint posterior inference problem over dynamic window expansion, latent block structure, and decoding organization. Extensive experiments across multiple benchmarks show that our approach consistently outperforms both fixed-length and existing flexible-length DLMs decoding methods on a variety of datasets.

Limitations Our method operates purely at inference time without modifying model parameters. While this enables broad applicability, integrating structural inference into training may further enhance performance, which we leave for future work.

Impact Statements This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

7 Acknowledgments

We acknowledge the University of Central Florida Advanced Research Computing Center (ARCC) for providing computational resources, as well as the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program for computational resources.

References

Marianne Arriola, Subham Sekhar Sahoo, Aaron Gokaslan, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Justin T Chiu, and Volodymyr Kuleshov. Block diffusion: Interpolating between autoregressive and diffusion language models. In *The Thirteenth International Conference on Learning Representations*, 2025.

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, 2021. URL `https://arxiv.org/abs/2108.07732`.

David M Blei, Thomas L Griffiths, and Michael I Jordan. The nested chinese restaurant process and bayesian nonparametric inference of topic hierarchies. *Journal of the ACM (JACM)*, 57(2):1–30, 2010.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,

Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, et al. Evaluating large language models trained on code, 2021. URL https://arxiv.org/abs/2107.03374.

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, 2021. URL https://arxiv.org/abs/2110.14168.

Justin Deschenaux and Caglar Gulcehre. Beyond autoregression: Fast LLMs via self-distillation through time. In *The Thirteenth International Conference on Learning Representations*, 2025.

Fangyu Ding, Ding Ding, Sijin Chen, Kaibo Wang, Peng Xu, Zijin Feng, Haoli Bai, Kai Han, Youliang Yan, Binhang Yuan, and Jiacheng Sun. Beyond masks: Efficient, flexible diffusion language models via deletion-insertion processes. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=VbvXjs5f72.

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. A framework for few-shot language model evaluation, 12 2023. URL https://zenodo.org/records/10256836.

Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385, 2024.

Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and Lingpeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. In *The Eleventh International Conference on Learning Representations*, 2023.

Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong. Scaling diffusion language models via adaptation from autoregressive models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Ishaan Gulrajani and Tatsunori Hashimoto. Likelihood-based diffusion language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Gabe Guo and Stefano Ermon. Self-speculative decoding accelerates lossless inference in any-order and any-subset autoregressive models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=hZnibTOke7.

Zhengfu He, Tianxiang Sun, Qiong Tang, Kuanning Wang, Xuan-Jing Huang, and Xipeng Qiu. Diffusionbert: Improving generative masked language models with diffusion models. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 4521–4534, 2023.

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 *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in neural information processing systems*, 34:12454–12465, 2021.

Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, and Sitan Chen. Train for the worst, plan for the best: Understanding token ordering in masked diffusions. In *Forty-second International Conference on Machine Learning*, 2025.

Jaeyeon Kim, Lee Cheuk Kit, Carles Domingo-Enrich, Yilun Du, Sham M. Kakade, Timothy Ngotiaoco, Sitan Chen, and Michael Samuel Albergo. Any-order flexible length masked diffusion. In *The Fourteenth International Conference on Learning Representations*, 2026.

Jinsong Li, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Jiaqi Wang, and Dahua Lin. Beyond fixed: Training-free variable-length denoising for diffusion large language models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Ic2A2gCseC.

Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *Advances in neural information processing systems*, 35: 4328–4343, 2022.

Anji Liu, Oliver Broadrick, Mathias Niepert, and Guy Van den Broeck. Discrete copula diffusion. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Guangyi Liu, Yu Wang, Zeyu Feng, Qiyu Wu, Liping Tang, Yuan Gao, Zhen Li, Shuguang Cui, Julian McAuley, Zichao Yang, Eric P. Xing, and Zhiting Hu. Unified generation, reconstruction, and representation: Generalized diffusion with adaptive latent encoding-decoding. In *Forty-first International Conference on Machine Learning*, 2024.

Sulin Liu, Juno Nam, Andrew Campbell, Hannes Stark, Yilun Xu, Tommi Jaakkola, and Rafael Gomez-Bombarelli. Think while you generate: Discrete diffusion with planned denoising. In *The Thirteenth International Conference on Learning Representations*, 2025b.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024.

Justin Lovelace, Varsha Kishore, Chao Wan, Eliot Shekhtman, and Kilian Q Weinberger. Latent diffusion for language generation. *Advances in Neural Information Processing Systems*, 36: 56998–57025, 2023.

Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, and Chongxuan Li. Scaling up masked diffusion models on text. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, JUN ZHOU, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b.

Xiaojing Qi, Lun Du, Xinyuan Zhang, Lanning Wei, Tao Jin, and Da Zheng. Hierarchy decoding: A training-free parallel decoding strategy for diffusion large language models. In *The Fourteenth International Conference on Learning Representations*, 2026.

Machel Reid, Vincent Josua Hellendoorn, and Graham Neubig. DiffusER: Diffusion via edit-based reconstruction. In *The Eleventh International Conference on Learning Representations*, 2023.

Litu Rout, Constantine Caramanis, and Sanjay Shakkottai. Anchored diffusion language model. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

Subham S Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.

Subham Sekhar Sahoo, Justin Deschenaux, Aaron Gokaslan, Guanghan Wang, Justin T Chiu, and Volodymyr Kuleshov. The diffusion duality. In *Forty-second International Conference on Machine Learning*, 2025.

Nikolay Savinov, Junyoung Chung, Mikolaj Binkowski, Erich Elsen, and Aaron van den Oord. Step-unrolled denoising autoencoders for text generation. In *International Conference on Learning Representations*, 2022.

Yeongbin Seo, Dongha Lee, Jaehyung Kim, and Jinyoung Yeo. Fast and fluent diffusion language models via convolutional decoding and rejective fine-tuning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized masked diffusion for discrete data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada, July 2023. Association for Computational Linguistics.

Dimitri von Rütte, Janis Fluri, Yuhui Ding, Antonio Orvieto, Bernhard Schölkopf, and Thomas Hofmann. Generalized interpolating discrete diffusion. In *Forty-second International Conference on Machine Learning*, 2025.

Minkai Xu, Tomas Geffner, Karsten Kreis, Weili Nie, Yilun Xu, Jure Leskovec, Stefano Ermon, and Arash Vahdat. Energy-based diffusion language models for text generation. In *The Thirteenth International Conference on Learning Representations*, 2025.

Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong. Beyond autoregression: Discrete diffusion for complex reasoning and planning. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. Dream 7b: Diffusion large language models, 2025b. URL `https://arxiv.org/abs/2508.15487`.

Peiyu Yu, Sirui Xie, Xiaojian Ma, Baoxiong Jia, Bo Pang, Ruiqi Gao, Yixin Zhu, Song-Chun Zhu, and Ying Nian Wu. Latent diffusion energy-based model for interpretable text modelling. In *International Conference on Machine Learning*, pages 25702–25720. PMLR, 2022.

Ruixiang ZHANG, Shuangfei Zhai, Yizhe Zhang, James Thornton, Zijing Ou, Joshua M. Susskind, and Navdeep Jaitly. Target concrete score matching: A holistic framework for discrete diffusion. In *Forty-second International Conference on Machine Learning*, 2025.

Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Masked diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical sampling. In *The Thirteenth International Conference on Learning Representations*, 2025.

A Notation

Table 5 provides a concise mathematical reference for the notation used in our method.

Table 5: Mathematical notation for DyStruct.

Symbol	Description
Decoding and Sequence Parameters	
t	Window expansion step index.
L_t	Window expansion length at step t.
$L_{\min}, L_{\max}$	Bounds for the window expansion length.
$N_{\max}$	Global sequence token limit.
$W^{(t)}$	Local indices in the newly appended window.
$Z^{(t)}$	Set of latent structural variables $\{L_t, \mathcal{P}^{(t)}, \tau^{(t)}\}$.
Diagnostic Temporary Pass and Uncertainty Metrics	
$O^{(t)}$	Observations from the diagnostic temporary pass.
p_g^{temp}	Initial distribution from the diagnostic pass.
$\mathcal{H}_i$	Token predictive entropy at position i.
$\phi_j^{(t)}$	Local feature vector at token position j.
w	Learned weights for position features $\phi_j^{(t)}$.
$h_j^{(t)}$	Position instability score at position j.
$\bar{h}^{(t)}$	Mean instability of the previously decoded window.
Partitioning and CRP Prior	
g	Gap index between token positions g and $g+1$.
$\psi_g^{(t)}$	Gap feature vector for the boundary at gap g.
w_b	Learned weights for gap features $\psi_g^{(t)}$.
$\ell_g^{(t)}, q_g^{(t)}$	Logit edge score and boundary split probability.
$\alpha_g^{(t)}$	Local CRP concentration parameter at gap g.
α_0	Base CRP concentration parameter.
m_g	Length of the block immediately preceding gap g.
$\mathcal{P}^{(t)}$	Partition of $W^{(t)}$ into contiguous token blocks.
$B_m^{(t)}$	The m-th block in partition $\mathcal{P}^{(t)}$.
Scheduling and Edge-Welding	
$\tau^{(t)}$	Context-aware decoding schedule for partition $\mathcal{P}^{(t)}$.
$H(B)$	Mean instability score of block B.
$C(B)$	Context-adjacency metric for block B.
γ	Weight for context-adjacency priority.
$T(B)$	Interpolated refinement steps allocated to block B.
$T_{\min}, T_{\max}$	Bounds for block refinement steps.
r_{weld}	Spatial radius for localized edge-welding.
$E_m^{(t)}$	Localized edge-welding interval between adjacent blocks.

B Algorithm

Algorithm 2 details the execution DyStruct. The process transforms flexible-length generation into an iterative cycle of adaptive expansion, structural partitioning, and context-aware resolution.

Algorithm 2 DyStruct: Dynamic Structured Decoding

Input: Prompt x, global length limit $N_{\max}$, hyperparameters $(\alpha_0, \gamma, r_{\text{weld}})$
Initialize: $y^{(0)} \leftarrow x$, step $t \leftarrow 1$, previous window instability $\bar{h}^{(0)} \leftarrow 0.5$
while `[EOS]` not generated **and** $|y^{(t-1)}| < N_{\max}$ **do**
  # 1. Adaptive Window Expansion
  Sample expansion length $L_t \sim p(L_t \mid O^{(t)}, y^{(t-1)}, x)$ via $\text{Poisson}(\mu_t)$ scaled by $(1 - \bar{h}^{(t-1)})$
  Append L_t `[MASK]` tokens to $y^{(t-1)}$ to form local window indices $W^{(t)}$
  # 2. Diagnostic Pass & Partitioning
  Execute temporary diagnostic pass over $W^{(t)}$ to extract feature signals $\phi^{(t)}$
  Compute position instability $h_j^{(t)}$ and gap split probabilities $q_g^{(t)}$ for all $j, g \in W^{(t)}$
  Evaluate local CRP parameters $\alpha_g^{(t)}$ against continuation and cutting scores
  Partition $W^{(t)}$ into contiguous blocks $\mathcal{P}^{(t)} = \{B_1^{(t)}, \ldots, B_{M_t}^{(t)}\}$
  # 3. Instability-Aware Scheduling
  Compute mean instability $H(B)$ and context-adjacency $C(B)$ for all $B \in \mathcal{P}^{(t)}$
  Order blocks into schedule $\tau^{(t)}$ by sorting $-H(B) + \gamma C(B)$
  # 4. Block-wise Resolution & Edge-Welding
  for each block $B \in \tau^{(t)}$ **do**
    Predict and commit tokens over interpolated $T(B)$ refinement steps
  end for
  for each shared boundary between adjacent blocks in $\mathcal{P}^{(t)}$ **do**
    Remask lowest-confidence positions strictly inside radius r_{weld} and refine
  end for
  # 5. State Update
  Calculate finalized mean instability $\bar{h}^{(t)}$ over the decoded window $W^{(t)}$
  Update sequence $y^{(t)}$ and increment $t \leftarrow t + 1$
end while
Output: Final generated sequence $y^{(\text{final})}$

Algorithmic Description. At each window expansion step, the decoder determines the length of the new masked window by evaluating the stability of the previously generated segment; high instability restricts expansion to prevent the propagation of structural errors. Before unmasking begins, a temporary diagnostic pass extracts local feature signals $\phi^{(t)}$. These diagnostic features inform a Bayesian partitioning step where a CRP prior groups tokens into contiguous blocks based on local instability. To ensure stable conditioning, the decoder resolves these partitioned blocks in a context-aware order, prioritizing segments that exhibit low-instability or exist adjacent to established context. Finally, an edge-welding step performs localized remasking at the interfaces of adjacent blocks to reconcile the predictive distributions and ensure sequence coherence.

C Method Details

We form a gap feature vector:

$$\psi_g^{(t)} = \left[h_g^{(t)}, h_{g+1}^{(t)}, \left| h_g^{(t)} - h_{g+1}^{(t)} \right|, \text{JSD}(p_g^{\text{temp}} \parallel p_{g+1}^{\text{temp}}) \right], \tag{12}$$

where p_g^{temp} is the initial diagnosis distribution at position g, and $\text{JSD}(\cdot \parallel \cdot)$ represents the Jensen-Shannon divergence. This gap feature vector quantifies whether the positions bounding the gap operate as a single block or require structural separation.

D Calibration of Instability Coefficients

To determine the instability coefficients (w), we construct a calibration dataset $\mathcal{D} = \{(\phi_i, d_i)\}_{i=1}^{N}$ by extracting N token-level observations from validation trajectories with accessible ground-truth sequences. We define the feature vector $\phi_i \in \mathbb{R}^7$ such that each vector contains the diagnostic metrics recorded for a specific token:

$$\phi_i = [\mathcal{H}_i, R_i, \Omega_i, \mathrm{JSD}_i, \Delta s_i, F_i, G_i]^\top \tag{13}$$

where the features correspond to predictive entropy ($\mathcal{H}$), remasking frequency (R), logit oscillation (Ω), Jensen-Shannon divergence (JSD), hidden state jump (Δs), confidence (F), and probability margin (G). We define the binary diagnostic targets $d_i \in \{0, 1\}$ using the ground-truth sequences: $d_i = 1$ if the predicted token mismatches the ground truth or triggers a remasking event during decoding. Conversely, $d_i = 0$ if the token correctly matches the ground truth and remains committed.

The components of ϕ_i are defined as follows:

$$\mathcal{H}_i = -\sum_{v \in \mathcal{V}} p_{i,v} \log p_{i,v}, \tag{14a}$$

$$R_i = \frac{1}{K} \sum_{k=1}^{K} \mathbf{1}\left[\kappa_i^{(k)} = 1 \wedge \eta_i^{(k)} = 0\right], \tag{14b}$$

$$\Omega_i = \frac{1}{\max(K-1, 1)} \sum_{k=2}^{K} \mathbf{1}\left[\hat{y}_i^{(k)} \neq \hat{y}_i^{(k-1)}\right], \tag{14c}$$

$$\mathrm{JSD}_i = \frac{1}{\max(K-1, 1)} \sum_{k=2}^{K} \mathrm{JSD}\left(p_i^{(k)} \parallel p_i^{(k-1)}\right), \tag{14d}$$

$$\Delta s_i = \frac{1}{d} \sum_{r=1}^{d} \left| s_{i,r} - s_{i-1,r} \right|, \tag{14e}$$

$$F_i = p_{i,\hat{y}_i}, \tag{14f}$$

$$G_i = \log p_{i,(1)} - \log p_{i,(2)}. \tag{14g}$$

Here $z_i \in \mathbb{R}^{|\mathcal{V}|}$ denotes the model logits at position i, $s_i \in \mathbb{R}^d$ denotes the corresponding final-layer hidden state, $p_i = \mathrm{softmax}(z_i)$, $\hat{y}_i = \arg\max_v z_{i,v}$, and $p_{i,(1)}$ and $p_{i,(2)}$ denote the largest and second-largest probabilities at position i, respectively. The boolean variable $\kappa_i^{(k)}$ indicates if the position is masked at refinement step k, and $\eta_i^{(k)}$ indicates if the token is accepted. The Jensen–Shannon divergence is defined as:

$$\mathrm{JSD}(p \parallel q) = \frac{1}{2}\mathrm{KL}(p \parallel m) + \frac{1}{2}\mathrm{KL}(q \parallel m), \qquad m = \frac{1}{2}(p + q). \tag{15}$$

We estimate the optimal coefficient vector w^* by minimizing an L_2-regularized binary cross-entropy loss over the calibration dataset $\mathcal{D}$:

$$\mathcal{J}(w) = -\frac{1}{N} \sum_{i=1}^{N} \left[d_i \log \sigma(w^\top \phi_i) + (1 - d_i) \log(1 - \sigma(w^\top \phi_i)) \right] + \lambda_{\mathrm{reg}} \|w\|_2^2 \tag{16}$$

where $\sigma(\cdot)$ is the sigmoid function and λ_{reg} is the regularization penalty.

During inference, the runtime algorithm evaluates the unconstrained logit u_i:

$$u_i = (w^*)^\top \phi_i \tag{17}$$

The linear projection u_i preserves the relative ranking of token instability established during calibration. The generation pipeline then normalizes the projection u_i using the window-centered logistic function defined in Equation 2 to compute the final positional instability score h_i.

E Implementation Details and Hyperparameters

To ensure the generalizability of DyStruct, DyStruct maintains a uniform set of hyperparameters across all benchmarks, domains, and model scales. We do not tune these parameters for task-specific optimality.

Window Expansion. The dynamic expansion loop is bounded by a minimum burst length of $L_{\min} = 8$ and a maximum burst length of $L_{\max} = 48$.

Partitioning and Scheduling. The Bayesian partitioning utilizes an empirical base CRP concentration prior of $\alpha_0 = 1.5$. During schedule evaluation, the context-adjacency priority weight is set to $\gamma = 2.0$.

Blockwise Decoding and Welding. Based on the block instability score $H(B) \in [0, 1]$, the total refinement steps $T(B)$ are interpolated between $T_{\min} = 6$ and $T_{\max} = 18$. Finally, the localized edge-welding step operates within a fixed spatial repair radius of $r_{\text{weld}} = 4$ tokens and executes for 4 refinement steps.

F Qualitative Examples

We provide decoding traces for code synthesis, logical reasoning, and linguistic disambiguation to demonstrate the mechanical execution of CRP-inspired partitioning, context-aware scheduling, and localized edge-welding.

Localized edge-welding resolves boundary fragmentation through targeted re-evaluation. Baseline diffusion models frequently produce distributional shifts at boundaries during flexible-length generation. Figure 3 illustrates this shift in a **HumanEval** programming task. When Block 2 is resolved independently, the partial context `if abs` provides an insufficient conditional signal, resulting in incoherent variables (`x - y`). DyStruct utilizes localized edge-welding to target the predictive entropy ($\mathcal{H}$) spike at this specific interface. By forcing a targeted re-evaluation strictly within the boundary overlap, the framework reconciles the adjacent distributions to recover context-grounded syntax `(numbers[i] - numbers[j])`.

Dataset:	`HumanEval`
Prompt:	`Given an array numbers, determine if any two elements have an absolute difference less than a threshold.`
Context:	`def check_diff(numbers, threshold):` `  for i in range(len(numbers)):`
Pre-Welding: *(Independent Blocks)*	`for j in range(i + 1, len(numbers)):` `  if abs(x - y) < threshold:`
Post-Welding: *(Final Output)*	`for j in range(i + 1, len(numbers)):` `  if abs(numbers[i] - numbers[j]) < threshold:`

Figure 3: **DyStruct Resolves Boundary Fragmentation via Edge-Welding.** Independent block decoding produces structurally incompatible boundaries. In this trace, Block 2 produces incoherent variables (`x - y`) due to weak boundary conditioning. DyStruct targets the entropy spike at this interface, replacing the incoherent segment with context-grounded syntax. (Red: incoherent variables; Green: localized repair.)

CRP-inspired partitioning isolates high-instability regions to constrain error propagation. Beyond coordinating boundaries, DyStruct adapts the internal structure of the generated region based on local instability. Figure 4 demonstrates the framework partitioning a window during a **Big-Bench Hard (BBH)** evaluation. The CRP prior evaluates the instability signals ($h_j^{(t)}$) from the temporary diagnostic pass and partitions the sequence into two distinct blocks. The context-aware schedule identifies the initial segment as well-anchored and low-instability ($H(B) = 0.27$), prioritizing its

resolution. Prioritizing stable segments first establishes a reliable conditional anchor before the decoder refines the high-instability deductive segment in Block 2 ($H(B) = 0.52$).

Dataset:	`BBH Logical Deduction`
Prompt:	`Evaluate the logical validity: Let A = True. B = (not A). Is A and B True?`
Context:	`Let's evaluate B: B = ( True ) = True. Plugging`
Block 1 (Low $\mathcal{H}$):	`in A and B,`
Schedule Priority 1	*Length: 5 tokens* | *Mean Instability: 0.27* | *Steps: 41*
Block 2 (High $\mathcal{H}$):	`we get True and False, which evaluates to False.`
Schedule Priority 2	*Length: 12 tokens* | *Mean Instability: 0.52* | *Steps: 215*

Figure 4: **DyStruct Isolates Logical Transitions via Partitioning.** The framework splits the unanchored window to isolate segments with high instability scores. The decoding schedule prioritizes the low-instability segment (Block 1) to provide stable conditioning before refining the logical evaluation in Block 2. (Blue: low-instability segment; Red: high-instability deduction.)

Context-aware scheduling utilizes stable anchors to provide bidirectional conditioning. In tasks containing causal dependencies, the decoding order determines the stability of the generated context. Figure 5 details a three-block trace from a **BBH Disambiguation** task. By identifying both the initial setup and the final conclusive answer as stable segments ($H = 0.21$ and $H = 0.18$), the scheduler resolves these terminal blocks first. This terminal-block-first scheduling strategy provides bidirectional context to ground the high-instability inferential segment in Block 2, ensuring the pronoun resolution is consistent with both the preceding setup and the required output format.

Dataset:	`BBH Disambiguation`
Prompt:	`Identify the antecedent: The chief thanked the housekeeper and gave her some tips.`
Context:	`Options: (A) The chief (B) The housekeeper (C) Ambiguous`
Block 1 (Low $\mathcal{H}$):	`The chief thanked the housekeeper`
Schedule Priority 1	*Length: 5 tokens* | *Mean Instability: 0.21* | *Steps: 40*
Block 2 (High $\mathcal{H}$):	`and gave her some tips.`
Schedule Priority 3	*Length: 5 tokens* | *Mean Instability: 0.59* | *Steps: 216*
Block 3 (Low $\mathcal{H}$):	`So the answer is (B).`
Schedule Priority 2	*Length: 5 tokens* | *Mean Instability: 0.18* | *Steps: 35*

Figure 5: **DyStruct Multi-Block Scheduling via Stable Anchors.** The scheduler prioritizes both terminal anchor blocks (1 and 3) to establish a constrained context for the high-instability inferential resolution in Block 2. (Blue: stable anchors; Red: high-instability inference.)

G Ablative Studies

More hyperparameter sensitivity analysis. Appendix Table 6 studies the sensitivity of the welding radius r_{weld} used in Eq. (11). For two neighboring blocks $B_m^{(t)} = [a, b)$ and $B_{m+1}^{(t)} = [b, c)$, edge-welding only operates within the local interval $E_m^{(t)} = [\max(a, b - r_{\text{weld}}), \min(c, b + r_{\text{weld}})]$. Therefore, r_{weld} controls the amount of neighboring context used to reconcile predictions around the block boundary, but it does not change the block partitioning objective in Eq. (9).

This explains the stability observed in the table. When r_{weld} is within a moderate range, the repair interval covers enough cross-boundary context to correct local inconsistencies while preserving the original block structure. On HumanEval, the results remain unchanged from $r_{\text{weld}} = 8$ to $r_{\text{weld}} = 12$, and on MBPP the nearby settings remain close to the default $r_{\text{weld}} = 10$. Since the token counts and block numbers are almost unchanged, the welding radius mainly affects local boundary coherence rather than generation length or segmentation granularity.

Overly small values of r_{weld} may make $E_m^{(t)}$ too narrow to include sufficient cross-boundary dependency, leaving adjacent blocks weakly aligned. Conversely, overly large values expand $E_m^{(t)}$ too far into neighboring blocks, which may remask positions that are already stable under the block score $\rho(B) = -H(B)+C(B)$. This can weaken the locality induced by the Bayesian partition and slightly degrade performance, as observed when $r_{\text{weld}} = 16$ on HumanEval. We therefore use $r_{\text{weld}} = 10$ as a fixed default, which provides stable boundary reconciliation without requiring task-specific tuning.

Table 6: **Sensitivity analysis of the welding radius r_{weld}.** We evaluate the effect of different welding radius on HumanEval and MBPP. The setting used in our main experiments is highlighted in blue. P@1 denotes Pass@1, Toks denotes the average number of newly generated tokens per sample, and Blks denotes the average number of blocks per sample.

r_{weld}	HumanEval			MBPP		
	P@1↑	**Toks↓**	**Blks**	**P@1↑**	**Toks↓**	**Blks**
8	34.8	219	7	41.0	256	7
10	**34.8**	**219**	**7**	**41.4**	**256**	**7**
12	34.8	219	7	41.2	256	7
16	34.1	219	7	41.0	256	7

H Statistical Significance

We compare LLaDA-8B-Base + DyStruct against LLaDA-8B-Base + DAEDAL using McNemar's test, since both methods are evaluated on the same set of prompts. We choose DAEDAL as the comparison baseline because it improves upon LLaDA-8B-Base. This paired test focuses only on discordant examples: prompts solved by DAEDAL but not DyStruct, and prompts solved by DyStruct but not DAEDAL. A significant one-sided McNemar test indicates that DyStruct wins on significantly more prompts than it loses against DAEDAL.

Table 7: **Paired McNemar test comparing DyStruct vs. Daedal.** BBH is pooled over all BBH subtasks; Math is pooled over GSM8K and MATH; Code is pooled over HumanEval and MBPP. We report McNemar's chi-square statistic with continuity correction, its asymptotic p_{CC}-value, and the exact two-sided binomial p_{exact} computed on discordant pairs only. Significant p-values are highlighted in light green.

Benchmark	n	**Acc$_{\text{Daedal}}$**	**Acc$_{\text{DyStruct}}$**	χ^2	p_{CC}	p_{exact}
BBH	6511	43.7%	49.3%	116.9	3.07×10^{-27}	1.09×10^{-27}
Math	6319	44.2%	45.3%	4.6	0.03	0.03
Code	664	38.5%	39.9%	0.9	0.35	0.35

Overall, the paired McNemar test shows that DyStruct consistently improves over DAEDAL across all three evaluation groups. The improvement is especially strong on BBH, where DyStruct achieves a substantially higher accuracy than DAEDAL and the difference is highly significant under both the continuity-corrected McNemar test and the exact binomial test. On the pooled Math benchmarks, DyStruct also obtains a statistically significant improvement, indicating that the gains are not limited to reasoning-heavy BBH tasks but also extend to mathematical problem solving.

For Code, DyStruct achieves higher accuracy than DAEDAL. This group contains only 664 instances, compared with 6,511 for BBH and 6,319 for Math, leading to fewer discordant pairs and lower statistical power in the paired test. The result therefore indicates a positive trend, while statistical significance would require a larger code evaluation set.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a proper justification is given (e.g., error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes].

   Justification: The claims in the abstract and introduction accurately reflect the paper's contributions.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes].

Justification: We have discussed limitations after conclusion.

Guidelines:

- The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A].

   Justification: the paper does not include theoretical results.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide sufficient implementation details in the experiment section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: Codes will be released upon acceptance.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We provide sufficient details for test details since our method is training-free.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: We provide error bars in the experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes].

   Justification: We provide computation resources details in the experimental section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.

- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes].

   Justification: The research conforms, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes].

    Justification: We discuss both potential positive societal impacts and negative societal impacts of the work after conclusion.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A].

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes].

    Justification: The creators or original owners of assets are credited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A].

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A].

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A].

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A].

    Justification: The core method development in this research does not involve LLMs.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.