

Novelty Assessment Report

Paper: Any-Order Flexible Length Masked Diffusion

PDF URL: <https://openreview.net/pdf?id=ttuNnMRI6H>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-01

Abstract

Masked diffusion models (MDMs) have recently emerged as a promising alternative to autoregressive models over discrete domains. MDMs generate sequences in an any-order, parallel fashion, enabling fast inference and strong performance on non-causal tasks. However, a crucial limitation is that they do not support token insertions and are thus limited to fixed-length generations. To this end, we introduce **Flexible Masked Diffusion Models (FlexMDMs)**, a discrete diffusion paradigm that simultaneously can model sequences of flexible length while provably retaining MDMs' flexibility of any-order inference. Grounded in an extension of the stochastic interpolant framework, FlexMDMs generate sequences by inserting mask tokens and unmasking them. Empirically, we show that FlexMDMs match MDMs in perplexity while modeling length statistics with much higher fidelity. On a synthetic maze planning task, they achieve $\approx 60\%$ higher success rate than MDM baselines. Finally, we show pretrained MDMs can easily be retrofitted into FlexMDMs: on 16 H100s, it takes only three days to fine-tune LLaDA-8B into a FlexMDM, achieving superior performance on math (GSM8K, 58% to 67%) and code infilling performance (52% to 65%).

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Variable-Length Sequence Generation with Masked Diffusion Models**

A total of **39 papers** were analyzed and organized into a taxonomy with **17 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Core Masked Diffusion Architectures and Training**
- **Variable-Length Generation and Flexible Decoding**
- **Inference and Sampling Optimization**
- **Application Domains and Specialized Modalities**
- **Surveys and Comparative Studies**

Complete Taxonomy Tree

- Variable-Length Sequence Generation with Masked Diffusion Models Survey Taxonomy
- Core Masked Diffusion Architectures and Training
 - General-Purpose Masked Diffusion Language Models (5 papers)
 - [1] Simple and effective masked diffusion language models (Marianne Arriola, 2024) [View paper](#)
 - [4] Simplified and generalized masked diffusion for discrete data (Arnaud Doucet, 2024) [View paper](#)
 - [19] Tess: Text-to-text self-conditioned simplex diffusion (Mahabadi, 2024) [View paper](#)
 - [30] Non-Autoregressive Text Generation Using Diffusion Models (Krishnamurthy, 2025) [View paper](#)
 - [31] Conditional Discrete Diffusion Language Model (H Koh, 2025) [View paper](#)
 - Hybrid and Unified Generative Frameworks (3 papers)
 - [5] Diffusion forcing: Next-token prediction meets full-sequence diffusion (Boyuan Chen, 2024) [View paper](#)
 - [13] Ladir: Latent diffusion enhances llms for text reasoning (Kang, 2025) [View paper](#)
 - [33] Unifying Autoregressive and Diffusion-Based Sequence Generation (Fathi, 2025) [View paper](#)
- Variable-Length Generation and Flexible Decoding
 - Adaptive Length Prediction and Control ★ (3 papers)
 - [0] Any-Order Flexible Length Masked Diffusion (Anon et al., 2026) [View paper](#)
 - [11] Diffusion llm with native variable generation lengths: Let lead the way (Y Yang, 2025) [View paper](#)
 - [17] Beyond fixed: Training-free variable-length denoising for diffusion large language models (Li Jinsong, 2025) [View paper](#)
 - Sequential and Block-Based Generation (1 papers)
 - [7] Sequential diffusion language models (Liu Yangzhou, 2025) [View paper](#)
 - Insertion and Infilling with Arbitrary Constraints (2 papers)
 - [3] Insertion Language Models: Sequence Generation with Arbitrary-Position Insertions (Patel, 2025) [View paper](#)
 - [24] Discrete Diffusion for Text Infilling (Zhang, 2025) [View paper](#)
- Inference and Sampling Optimization
 - Parallel Sampling and Conditional Independence (1 papers)
 - [9] Parallel Sampling from Masked Diffusion Models via Conditional Independence Testing (Azangulov, 2025) [View paper](#)
 - Self-Speculative and Iterative Refinement (1 papers)
 - [23] Self-Speculative Masked Diffusions (Campbell, 2025) [View paper](#)
 - Masking Schedules and Error Bounds (2 papers)
 - [14] Error Bounds and Optimal Schedules for Masked Diffusions with Factorized Approximations (Lavenant, 2025) [View paper](#)
 - [16] Path Planning for Masked Diffusion Models with Applications to Biological Sequence Generation (FZ Peng, 2025) [View paper](#)

- Guidance and Conditioning Techniques (2 papers)
- [15] Masked diffusion with task-awareness for procedure planning in instructional videos (Fang Fen, 2023) [View paper](#)
- [28] Improved Sampling from Masked Diffusion Models with Position Contrastive Guidance (D Patel, 2025) [View paper](#)
- Application Domains and Specialized Modalities
 - Motion and Human Movement Generation (4 papers)
 - [2] Mmm: Generative masked motion model (Ekkasit Pinyoanuntapong, 2024) [View paper](#)
 - [12] Length-aware motion synthesis via latent diffusion (Sampieri, 2024) [View paper](#)
 - [36] Uniform Text-Motion Generation and Editing via Diffusion Model (R Wang, n.d.) [View paper](#)
 - [39] Smooth and Flexible Camera Movement Synthesis via Temporal Masked Generative Modeling (C Xu, n.d.) [View paper](#)
 - Video and Visual Sequence Generation (3 papers)
 - [10] Avid: Any-length video inpainting with diffusion model (Zhixing Zhang, 2024) [View paper](#)
 - [29] From Narrative to Frames: AI-Assisted Storyboarding with personalized Diffusion Models (Ocleppo, 2025) [View paper](#)
 - [37] ReGenText: Joint Generation and Restoration for Diverse Text Image Super-Resolution (B Wang, n.d.) [View paper](#)
 - Biological Sequence Generation (2 papers)
 - [34] Joint Generation of Protein Sequence and Structure with RoseTTAFold Sequence Space Diffusion (S. Lisanza, 2023) [View paper](#)
 - [35] ProtDiff: Function-Conditioned Masked Diffusion Models for Robust Directed Protein Generation (V Thoutam, n.d.) [View paper](#)
 - Scientific and Multimodal Data (5 papers)
 - [8] A diffusion model for regular time series generation from irregular data with completion and masking (Gal Fadlon, 2025) [View paper](#)
 - [20] Eegmirror: Leveraging eeg data in the wild via montage-agnostic self-supervision for eeg to video decoding (XH Liu, 2025) [View paper](#)
 - [22] Latent diffusion for neural spiking data (Kapoor, 2024) [View paper](#)
 - [27] Diffusion Autoencoders with Perceivers for Long, Irregular and Multimodal Astronomical Sequences (Shen, 2025) [View paper](#)
 - [38] A Comprehensive Evaluation of Watermarking for Time Series Diffusion Models (Timmer, n.d.) [View paper](#)
 - Recommendation and Sequential Decision-Making (3 papers)
 - [21] STCDM: Spatio-Temporal Contrastive Diffusion Model for Check-In Sequence Generation (Letian Gong, 2025) [View paper](#)
 - [25] Diffusion models in offline reinforcement learning (Kyriazopoulos, 2025) [View paper](#)
 - [26] Masked Diffusion for Generative Recommendation (Kulin Shah, 2025) [View paper](#)
 - Autonomous Systems and Robotics (1 papers)
 - [18] WAM-Diff: A Masked Diffusion VLA Framework with MoE and Online Reinforcement Learning for Autonomous Driving (Mingwang Xu, 2025) [View paper](#)
 - Adaptive Tokenization and Representation Learning (1 papers)
 - [32] ALTo: Adaptive-Length Tokenizer for Autoregressive Mask Generation (Wang Lingfeng, 2025) [View paper](#)
- Surveys and Comparative Studies (1 papers)
 - [6] Diffusion-based Large Language Models Survey (Chiung-Yi Tseng, 2025) [View paper](#)

Narrative

Core task: variable-length sequence generation with masked diffusion models. The field has coalesced around several complementary directions. At the foundation, works such as Simple Masked Diffusion[1] and Simplified Masked Diffusion[4] establish core architectures and training procedures for discrete masked diffusion, while Diffusion Forcing[5] explores unified frameworks that bridge autoregressive and diffusion paradigms. A second major branch addresses variable-length generation and flexible decoding, where methods like Insertion Language Models[3] and Variable Length Diffusion LLM[11] tackle adaptive length prediction and non-autoregressive ordering. Inference optimization forms another active area, with Parallel Masked Diffusion Sampling[9] and Self-Speculative Masked Diffusions[23] accelerating sampling through speculative or parallel strategies. Finally, diverse application domains—from motion synthesis (Generative Masked Motion[2], Length-Aware Motion Synthesis[12]) to specialized modalities like protein design (RoseTTAFold Sequence Diffusion[34]) and recommendation systems (Masked Diffusion Recommendation[26])—demonstrate the breadth of masked diffusion’s reach.

Within the variable-length generation branch, a central challenge is balancing flexibility in output length with coherent structure and efficient sampling. Works like Variable Length Diffusion LLM[11] and Variable-Length Denoising[17] explore explicit length conditioning and adaptive stopping criteria, while Insertion Language Models[3] investigates order-agnostic generation that can insert tokens at arbitrary positions. Any-Order Flexible Masked Diffusion[0] sits naturally in this cluster, emphasizing flexible decoding orders that allow dynamic control over sequence construction without rigid left-to-right constraints. Compared to Variable Length Diffusion LLM[11], which focuses on length prediction mechanisms, and Variable-Length Denoising[17], which addresses denoising schedules for varying lengths, the original work appears to prioritize the ordering flexibility itself—enabling any-order masking patterns that adapt to task-specific needs. This positioning reflects ongoing debates about whether length control, decoding order, or sampling efficiency should take precedence in making masked diffusion practical for real-world variable-length tasks.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Diffusion llm with native variable generation lengths: Let lead the way

Authors: Y Yang, C Wang, S Wang, Z Wen, B Qi, et al. (6 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

Compared with the traditional block diffusion model, it face challenges in handling variable-length sequences and efficient masking schedule such that EOS tokens are always

Relationship Analysis

Both papers belong to the Adaptive Length Prediction and Control category, addressing variable-length sequence generation in masked diffusion models through dynamic length determination mechanisms. They overlap in their core objective of enabling flexible-length generation without fixed maximum lengths, both introducing methods for handling end-of-sequence tokens and length adaptation. The key difference is that the original paper (FlexMDM) uses a joint interpolant framework with insertion and unmasking operations based on stochastic interpolants, while the candidate paper appears to focus on EOS token handling and masking schedules as the primary mechanism for variable-length control.

2. Beyond fixed: Training-free variable-length denoising for diffusion large language models

Authors: Li Jinsong, Dong, Xiaoyi, Jinsong Li, Zang, et al. (16 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

Diffusion Large Language Models (DLLMs) are emerging as a powerful alternative to the dominant Autoregressive Large Language Models, offering efficient parallel generation and capable global context modeling. However, the practical application of DLLMs is hindered by a critical architectural constraint: the need for a statically predefined generation length. This static length allocation leads to a problematic trade-off: insufficient lengths cripple performance on complex tasks, while excessive ...

Relationship Analysis

Both papers address adaptive length prediction and control in masked diffusion models for variable-length sequence generation. The original paper (Any-Order Flexible Length Masked Diffusion) introduces FlexMDM, a training-based approach that learns insertion expectations and unmasking posteriors through a joint interpolant framework, requiring model retraining or fine-tuning. The candidate paper (DAEDAL) takes a fundamentally different approach by proposing a training-free inference strategy that dynamically expands generation length using internal model signals, without requiring any model modification or retraining.

Contributions Analysis

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Flexible Masked Diffusion Models (FlexMDMs)

Description: The authors propose FlexMDMs, a new discrete diffusion framework that extends masked diffusion models to handle variable-length sequences through token insertion and unmasking operations, while preserving the any-order generation capability of standard MDMs. This addresses a key limitation of existing MDMs which are restricted to fixed-length generations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Diffuseq: Sequence to sequence text generation with diffusion models

URL: [View paper](#)

Brief Assessment

DiffuSeq[42] focuses on sequence-to-sequence text generation using diffusion models for conditional generation tasks, not on variable-length sequence generation with any-order inference capabilities or token insertion operations that FlexMDMs address.

2. Non-autoregressive diffusion-based temporal point processes for continuous-time long-term event prediction

URL: [View paper](#)

Brief Assessment

Temporal Point Processes[47] focuses on continuous-time event prediction using diffusion models for temporal point processes, not discrete sequence generation with variable-length token insertion and any-order unmasking capabilities.

3. Sequential diffusion language models

URL: [View paper](#)

Brief Assessment

Sequential Diffusion Language[7] focuses on block-based diffusion with adaptive subsequence decoding for efficiency and KV-cache compatibility, not on extending masked diffusion models to handle variable-length sequences through token insertion and unmasking operations while preserving any-order generation capability.

4. Block diffusion: Interpolating between autoregressive and diffusion language models

URL: [View paper](#)

Brief Assessment

Block Diffusion[40] focuses on interpolating between autoregressive and diffusion models with block-based generation, not on extending masked diffusion models to handle variable-length sequences through token insertion and unmasking operations while preserving any-order generation capability.

5. Vector quantized diffusion model with codeunet for text-to-sign pose sequences generation

URL: [View paper](#)

Brief Assessment

Text-to-Sign Diffusion[48] focuses on sign language pose sequence generation using vector quantization and diffusion in a domain-specific application, not on general variable-length discrete sequence generation with any-order inference capabilities.

6. dkv-cache: The cache for diffusion language models

URL: [View paper](#)

Brief Assessment

dkv-cache[45] focuses on accelerating inference in diffusion language models through key-value caching mechanisms, not on enabling variable-length sequence generation with token insertion capabilities.

7. DiffSound: Discrete diffusion model for text-to-sound generation

URL: [View paper](#)

Brief Assessment

DiffSound[41] applies discrete diffusion to audio token generation in a fixed mel-spectrogram grid, not to variable-length sequence modeling with token insertion capabilities as in FlexMDMs.

8. Discrete diffusion models for language generation

URL: [View paper](#)

Brief Assessment

Discrete Diffusion Language[43] focuses on D3PM models for fixed-length language generation and compares them with autoregressive models. It does not address variable-length sequence generation or token insertion mechanisms that are central to FlexMDMs.

9. DiffListener: Discrete Diffusion Model for Listener Generation

URL: [View paper](#)

Brief Assessment

DiffListener[46] focuses on listener head generation using discrete diffusion for a specific multimodal task (generating nonverbal listener responses), not on general variable-length sequence generation frameworks with any-order inference capabilities.

10. DiffRhythm: Blazingly fast and embarrassingly simple end-to-end full-length song generation with latent diffusion

URL: [View paper](#)

Brief Assessment

DiffRhythm[44] focuses on latent diffusion for music generation with vocal and accompaniment synthesis. It does not address discrete diffusion models for variable-length sequence generation or any-order inference capabilities.

Contribution 2: Joint interpolant framework for variable-length modeling

Description: The authors develop a theoretical foundation based on extending the stochastic interpolant framework to discrete spaces. They introduce a joint interpolant that augments the process with an auxiliary variable tracking token positions, enabling closed-form rate matrix characterization for variable-length sequence generation.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Quantitative properties of sovereign default models: solution methods matter

URL: [View paper](#)

Brief Assessment

Sovereign Default Models[60] focuses on numerical methods for sovereign default models using interpolation over asset positions in economic modeling. This is fundamentally different from the original paper's stochastic interpolant framework for discrete sequence generation in machine learning.

2. Atmospheric water vapor maps generations from stochastic interpolation of GNSS Zenith Tropospheric Delays

URL: [View paper](#)

Brief Assessment

GNSS Tropospheric Delays[62] focuses on atmospheric water vapor mapping using stochastic interpolation of GNSS tropospheric delays through kriging methods. This is entirely unrelated to discrete diffusion models or variable-length sequence generation in machine learning contexts.

3. Elimination of systematic error in digital image correlation caused by intensity interpolation by introducing position randomness to subset points

URL: [View paper](#)

Brief Assessment

Digital Image Correlation[59] focuses on eliminating systematic errors in image correlation through intensity interpolation and position randomness in subset points. This is a computer vision technique unrelated to stochastic interpolants for discrete sequence generation or variable-length modeling in generative models.

4. Uncertainty Analysis of the Dynamic Response of a Randomly Parametrized Corrugated Skin

URL: [View paper](#)

Brief Assessment

Corrugated Skin Uncertainty[63] focuses on uncertainty quantification for structural dynamics using sparse grid collocation methods, not on stochastic interpolants for discrete sequence generation or variable-length modeling in generative models.

5. Stochastic signal processing in adaptive measurement systems with rough space-time statistics: Method of invertible spectral analysis

URL: [View paper](#)

Brief Assessment

Invertible Spectral Analysis[61] focuses on signal processing in adaptive measurement systems with space-time statistics. The candidate discusses stochastic interpolation in discrete processing contexts but does not address variable-length sequence generation, discrete diffusion models, or the theoretical framework for modeling distributions over sequences of different lengths that is central to the original paper's contribution.

Contribution 3: Efficient retrofitting of pretrained MDMs into FlexMDMs

Description: The authors demonstrate that existing pretrained masked diffusion models can be efficiently converted into FlexMDMs through minimal architectural modifications and fine-tuning. This enables scaling to 8B parameters with only three days of training on 16 H100 GPUs, achieving substantial performance improvements on downstream tasks.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Plan for Speed--Dilated Scheduling for Masked Diffusion Language Models

URL: [View paper](#)

Brief Assessment

Dilated Scheduling[57] focuses on inference-time scheduling strategies for existing MDMs without architectural modifications or fine-tuning. The original paper's contribution involves architectural changes (time-embedding layers, scalar heads, LoRA adapters) and fine-tuning pretrained models, which is fundamentally different from an inference-only method.

2. Unified auto-encoding with masked diffusion

URL: [View paper](#)

Brief Assessment

Unified Auto-Encoding[56] focuses on combining patch-based masking with noise-based corruption for image generation and representation learning, not on retrofitting pretrained masked diffusion models for variable-length text generation through architectural modifications and fine-tuning.

3. Fine-tuning diffusion models with limited data

URL: [View paper](#)

Brief Assessment

Fine-tuning Limited Data[50] focuses on fine-tuning pretrained diffusion models for image generation tasks with limited data, using attention block tuning and adapter modules. This is fundamentally different from retrofitting masked diffusion models for variable-length text generation as described in the original paper.

4. Steering masked discrete diffusion models via discrete denoising posterior prediction

URL: [View paper](#)

Brief Assessment

Discrete Denoising Posterior[53] focuses on steering pretrained masked diffusion models via reward-based fine-tuning for controlled generation, not on architectural retrofitting for variable-length modeling capabilities.

5. Maskdiffusion: Exploiting pre-trained diffusion models for semantic segmentation

URL: [View paper](#)

Brief Assessment

MaskDiffusion[49] focuses on semantic segmentation using pretrained Stable Diffusion models without additional training, not on retrofitting masked diffusion models for text generation tasks.

6. DiffFit: Unlocking transferability of large diffusion models via simple parameter-efficient fine-tuning

URL: [View paper](#)

Brief Assessment

DiffFit[52] focuses on parameter-efficient fine-tuning of diffusion models for image generation tasks (adapting to new visual domains/resolutions), not on retrofitting masked diffusion models for discrete sequence generation or converting MDMs to FlexMDMs for variable-length text/code tasks.

7. Beyond masked and unmasked: Discrete diffusion models via partial masking

URL: [View paper](#)

Brief Assessment

Partial Masking Diffusion[51] focuses on introducing intermediate token states between masked and unmasked to reduce redundant computation, not on retrofitting pretrained models with minimal architectural changes for scaling.

8. Cross-view Masked Diffusion Transformers for Person Image Synthesis

URL: [View paper](#)

Brief Assessment

Cross-view Person Synthesis[58] focuses on pose-guided human image generation using masked diffusion transformers on latent patches, not on retrofitting pretrained masked diffusion models for variable-length sequence generation or general language modeling tasks.

9. Not All Parameters Matter: Masking Diffusion Models for Enhancing Generation Ability

URL: [View paper](#)

Brief Assessment

Parameter Masking Diffusion[54] focuses on masking U-Net parameters in diffusion models for image generation, not on retrofitting masked diffusion models for discrete sequence generation. The technical domains and objectives are fundamentally different.

10. More control for free! image synthesis with semantic diffusion guidance

URL: [View paper](#)

Brief Assessment

Semantic Diffusion Guidance[55] focuses on guiding pretrained unconditional diffusion models for image synthesis using CLIP-based semantic guidance, not on retrofitting masked diffusion models for text generation or converting them to handle variable-length sequences.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Any-Order Flexible Length Masked Diffusion [View paper](#)
- [1] Simple and effective masked diffusion language models [View paper](#)
- [2] Mmm: Generative masked motion model [View paper](#)
- [3] Insertion Language Models: Sequence Generation with Arbitrary-Position Insertions [View paper](#)
- [4] Simplified and generalized masked diffusion for discrete data [View paper](#)
- [5] Diffusion forcing: Next-token prediction meets full-sequence diffusion [View paper](#)
- [6] Diffusion-based Large Language Models Survey [View paper](#)
- [7] Sequential diffusion language models [View paper](#)
- [8] A diffusion model for regular time series generation from irregular data with completion and masking [View paper](#)
- [9] Parallel Sampling from Masked Diffusion Models via Conditional Independence Testing [View paper](#)
- [10] Avid: Any-length video inpainting with diffusion model [View paper](#)
- [11] Diffusion llm with native variable generation lengths: Let lead the way [View paper](#)
- [12] Length-aware motion synthesis via latent diffusion [View paper](#)
- [13] Ladir: Latent diffusion enhances llms for text reasoning [View paper](#)
- [14] Error Bounds and Optimal Schedules for Masked Diffusions with Factorized Approximations [View paper](#)
- [15] Masked diffusion with task-awareness for procedure planning in instructional videos [View paper](#)
- [16] Path Planning for Masked Diffusion Models with Applications to Biological Sequence Generation [View paper](#)
- [17] Beyond fixed: Training-free variable-length denoising for diffusion large language models [View paper](#)
- [18] WAM-Diff: A Masked Diffusion VLA Framework with MoE and Online Reinforcement Learning for Autonomous Driving [View paper](#)
- [19] Tess: Text-to-text self-conditioned simplex diffusion [View paper](#)
- [20] Eegmirror: Leveraging eeg data in the wild via montage-agnostic self-supervision for eeg to video decoding [View paper](#)
- [21] STCDM: Spatio-Temporal Contrastive Diffusion Model for Check-In Sequence Generation [View paper](#)
- [22] Latent diffusion for neural spiking data [View paper](#)

- [23] Self-Speculative Masked Diffusions [View paper](#)
- [24] Discrete Diffusion for Text Infilling [View paper](#)
- [25] Diffusion models in offline reinforcement learning [View paper](#)
- [26] Masked Diffusion for Generative Recommendation [View paper](#)
- [27] Diffusion Autoencoders with Perceivers for Long, Irregular and Multimodal Astronomical Sequences [View paper](#)
- [28] Improved Sampling from Masked Diffusion Models with Position Contrastive Guidance [View paper](#)
- [29] From Narrative to Frames: AI-Assisted Storyboarding with personalized Diffusion Models [View paper](#)
- [30] Non-Autoregressive Text Generation Using Diffusion Models [View paper](#)
- [31] Conditional Discrete Diffusion Language Model [View paper](#)
- [32] ALTo: Adaptive-Length Tokenizer for Autoregressive Mask Generation [View paper](#)
- [33] Unifying Autoregressive and Diffusion-Based Sequence Generation [View paper](#)
- [34] Joint Generation of Protein Sequence and Structure with RoseTTAFold Sequence Space Diffusion [View paper](#)
- [35] ProtDiff: Function-Conditioned Masked Diffusion Models for Robust Directed Protein Generation [View paper](#)
- [36] Uniform Text-Motion Generation and Editing via Diffusion Model [View paper](#)
- [37] ReGenText: Joint Generation and Restoration for Diverse Text Image Super-Resolution [View paper](#)
- [38] A Comprehensive Evaluation of Watermarking for Time Series Diffusion Models [View paper](#)
- [39] Smooth and Flexible Camera Movement Synthesis via Temporal Masked Generative Modeling [View paper](#)
- [40] Block diffusion: Interpolating between autoregressive and diffusion language models [View paper](#)
- [41] Diffsound: Discrete diffusion model for text-to-sound generation [View paper](#)
- [42] Diffuseq: Sequence to sequence text generation with diffusion models [View paper](#)
- [43] Discrete diffusion models for language generation [View paper](#)
- [44] DiffRhythm: Blazingly fast and embarrassingly simple end-to-end full-length song generation with latent diffusion [View paper](#)
- [45] dkv-cache: The cache for diffusion language models [View paper](#)
- [46] DiffListener: Discrete Diffusion Model for Listener Generation [View paper](#)
- [47] Non-autoregressive diffusion-based temporal point processes for continuous-time long-term event prediction [View paper](#)
- [48] Vector quantized diffusion model with codeunet for text-to-sign pose sequences generation [View paper](#)
- [49] Maskdiffusion: Exploiting pre-trained diffusion models for semantic segmentation [View paper](#)
- [50] Fine-tuning diffusion models with limited data [View paper](#)
- [51] Beyond masked and unmasked: Discrete diffusion models via partial masking [View paper](#)
- [52] DiffFit: Unlocking transferability of large diffusion models via simple parameter-efficient fine-tuning [View paper](#)
- [53] Steering masked discrete diffusion models via discrete denoising posterior prediction [View paper](#)
- [54] Not All Parameters Matter: Masking Diffusion Models for Enhancing Generation Ability [View paper](#)
- [55] More control for free! image synthesis with semantic diffusion guidance [View paper](#)
- [56] Unified auto-encoding with masked diffusion [View paper](#)
- [57] Plan for Speed--Dilated Scheduling for Masked Diffusion Language Models [View paper](#)
- [58] Cross-view Masked Diffusion Transformers for Person Image Synthesis [View paper](#)
- [59] Elimination of systematic error in digital image correlation caused by intensity interpolation by introducing position randomness to subset points [View paper](#)
- [60] Quantitative properties of sovereign default models: solution methods matter [View paper](#)
- [61] Stochastic signal processing in adaptive measurement systems with rough space-time statistics: Method of invertible spectral analysis [View paper](#)
- [62] Atmospheric water vapor maps generations from stochastic interpolation of GNSS Zenith Tropospheric Delays [View paper](#)
- [63] Uncertainty Analysis of the Dynamic Response of a Randomly Parametrized Corrugated Skin [View paper](#)