

A Survey on Diffusion Language Models

Tianyi Li, Mingda Chen, Bowei Guo, and Zhiqiang Shen

Abstract—Diffusion Language Models (DLMs) are rapidly emerging as a powerful and promising alternative to the dominant autoregressive (AR) paradigm. By generating tokens in parallel through an iterative denoising process, DLMs possess inherent advantages in reducing inference latency and capturing bidirectional context, thereby enabling fine-grained control over the generation process. While achieving several-fold speedups, recent advancements have allowed DLMs to show performance comparable to their autoregressive counterparts, making them a compelling choice for various natural language processing tasks. Despite their growing prevalence, DLMs present challenges and opportunities that warrant further exploration, requiring a detailed and systematic understanding of their principles, techniques, and limitations. In this survey, we provide a holistic overview of the current DLM landscape. We trace its evolution and relationship with other paradigms, such as autoregressive and masked language models, and cover both foundational principles and state-of-the-art models. Our work offers an up-to-date, comprehensive taxonomy and an in-depth analysis of current techniques, from pre-training strategies to advanced post-training methods. Another contribution of this survey is a thorough review of DLM inference strategies and optimizations, including improvements in decoding parallelism, caching mechanisms, and generation quality. We also highlight the latest approaches to multimodal extensions of DLMs and delineate their applications across various practical scenarios. Furthermore, our discussion addresses the limitations and challenges of DLMs, including efficiency, long-sequence handling, and infrastructure requirements, while outlining future research directions to sustain progress in this rapidly evolving field. Project GitHub is available at <https://github.com/VILA-Lab/Awesome-DLMs>.

Index Terms—Diffusion Language Model, Large Language Model, Diffusion Model, Diffusion Large Language Model, Language Modeling, Multimodal Language Model



1 INTRODUCTION

RECENT advancements toward artificial general intelligence (AGI) have been largely driven by the emergence of autoregressive large language models (LLMs) [1]–[7] and diffusion models for image and video generation [8]–[12]. These models exhibit remarkable capabilities in both understanding and generation across diverse modalities, achieving levels of performance that were previously unimaginable. The unprecedented scale of these models, reflected in massive parameter counts, vast datasets, substantial efforts in training, and significant computational demands during inference, has pushed AI to new heights, equipping these models with broad general knowledge and a deep understanding of language and the real world.

The rise of the GPT series [1], [13], [14], particularly with the public release of ChatGPT [2], has propelled autoregressive (AR) language models to a dominant position in natural language processing. By training to predict the next token using causal attention and teacher forcing, AR models [4], [15], [16] can effectively scale to large datasets and model sizes. Generating text in a sequential, token-by-token fashion, AR models excel at supporting a wide range of tasks, from simple question answering to complex reasoning and creative writing. However, this sequential nature imposes a major bottleneck on inference speed. The autoregressive generation process, which produces one token at a time, inherently limits parallelism and significantly constrains computational efficiency and throughput.

Diffusion models are another highly promising generative paradigm. They are trained to recover data from progressively noised versions through a denoising process, and generate new samples by reversing this stochastic corruption step by step. Excelling at modeling complex data distributions, diffusion models have achieved state-of-the-art results in image and video synthesis [17]. Academic breakthroughs in diffusion modeling [18]–[21] have established a solid theoretical foundation for training and inference. Concurrently, large-scale practical models like Stable Diffusion [8], [10], [11], Imagen [9], and Sora [12] demonstrate the remarkable scalability and generalization of diffusion paradigm, enabling generation of high-fidelity, art-level images and videos from simple text prompts—often with just a few words. Beyond their strong capacity for modeling complex data distributions, diffusion models offer an inherent advantage in parallelism. Through an iterative denoising process, they can generate multiple tokens or an entire sequence simultaneously, potentially leading to superior inference throughput and better utilization of modern parallel computing hardware. While challenges remain, particularly in modeling discrete data and handling dynamic sequence lengths, Diffusion Language Models (DLMs) have emerged as a compelling alternative to address the trade-off between generation quality and speed.

To adapt diffusion for discrete language data, several key approaches have been proposed. In the early stages, the development of DLMs was primarily driven by diffusion models’ success in continuous domains like image synthesis. Continuous DLMs map tokens into embeddings and perform denoising in continuous space, as in pioneering works Diffusion-LM [22] and SED [23]. Discrete DLMs, on the other hand, define the diffusion process directly in token

- Tianyi Li, Mingda Chen, Bowei Guo, and Zhiqiang Shen are with VILA Lab, Mohamed bin Zayed University of Artificial Intelligence. Mingda Chen is also with Department of Automation, Tsinghua University. E-mail: {Tianyi.Li, Bowei.Guo, Zhiqiang.Shen}@mbzuai.ac.ae, cmd22@mails.tsinghua.edu.cn

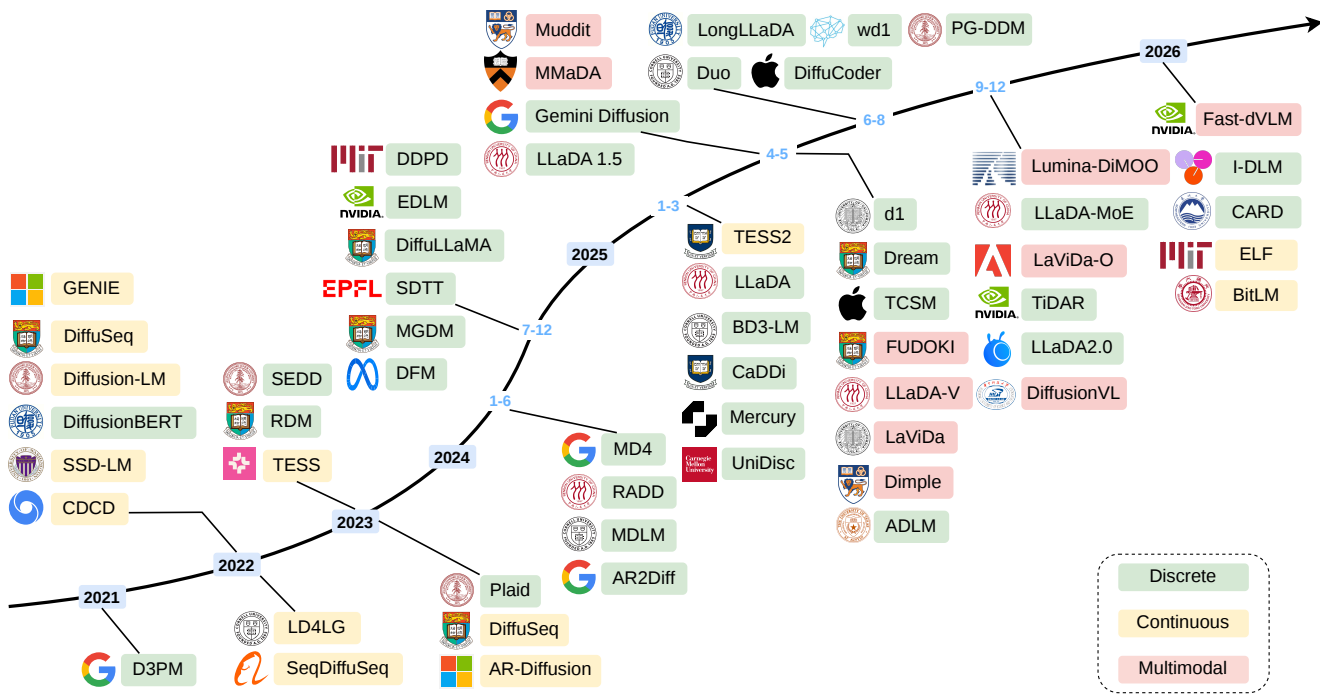


Fig. 1. Timeline of Diffusion Language Models. This figure highlights key milestones in the development of DLMs, categorized into three groups: continuous DLMs, discrete DLMs, and recent multimodal DLMs. We observe that while early research predominantly focused on continuous DLMs, discrete DLMs have gained increasing popularity in more recent years.

space. Early efforts such as D3PM [24] introduced structured transition matrices with absorbing states, allowing token-level corruption and iterative denoising. Subsequent work like DiffusionBERT [25] integrated pre-trained masked language models (e.g., BERT) to enhance denoising quality, and proposed tailored noise schedules (e.g., the spindle schedule) to better align token corruption with token frequency. These early models demonstrated the feasibility of applying iterative denoising to non-autoregressive text generation, offering controllability and parallelism, though their performance still lagged behind strong autoregressive baselines. As core challenges in DLMs are gradually addressed and the paradigm matures, larger-scale DLMs have been developed. By initializing from autoregressive models, 7B-level models like Dream [26] and DiffuLLaMA [27] have shown that DLMs can be effectively adapted from existing models while achieving competitive performance. LLaDA-8B [28] further demonstrates the potential of training DLMs from scratch, achieving performance comparable to similarly sized LLaMA3-8B models. Multimodal DLMs, also known as diffusion multimodal large language models (dMLLMs), have also shown promise in modeling hybrid data such as text and images. Built upon open-source DLMs, models like LLaDA-V [29], Dimple [30], and MMaDA [31] integrate cross-modal reasoning and generation into the diffusion framework. Meanwhile, industry efforts have also shown growing interest in DLMs. The Mercury series [32], Gemini Diffusion [33], and Seed Diffusion [34] report strong performance while achieving inference speeds of thousands of tokens per second. These developments highlight the growing practicality and commercial potential of DLMs. We provide a timeline of DLMs’ development in Fig. 1, ranging

from representative models to recent advancements [35]–[41], followed by a visualization of DLM trends in Fig. 2.

Diffusion language models also present unique challenges and opportunities in both training and inference. Pretraining typically follows strategies similar to those used in autoregressive language models or image diffusion models [26], [30], [31]. To accelerate training and reuse previous training efforts, many DLMs are initialized from pretrained autoregressive model weights [26], [27]. Supervised fine-tuning (SFT) in DLMs also mirrors that of autoregressive models: clean prompt data is provided, and the model learns to generate the target completion. Reinforcement learning (RL) is also adopted for DLMs post-training to improve performance on complex tasks. Variants of GRPO [42] algorithm such as diffu-GRPO [43] and UniGRPO [31] have been proposed to enhance the reasoning capabilities and alignment of DLMs at scale. During inference, various strategies and optimizations have been developed to fully utilize the capabilities of DLMs. Continuous DLMs can leverage ODE/SDE solvers or other few-step generation techniques to accelerate the iterative denoising process [44]. As discrete DLMs face more challenges in parallel generation, specialized parallel decoding strategies [30], [45], [46] have been proposed to enable acceptance of multiple tokens at a single step and overcome the parallelism curse. Unmasking and remasking strategies [28], [47] further improve generation quality by selectively revealing low-confidence tokens, while caching techniques [48], [49] can significantly reduce computation and enhance inference speed for both paradigms.

Compared to autoregressive models, diffusion language models are widely believed to offer several distinct advan-

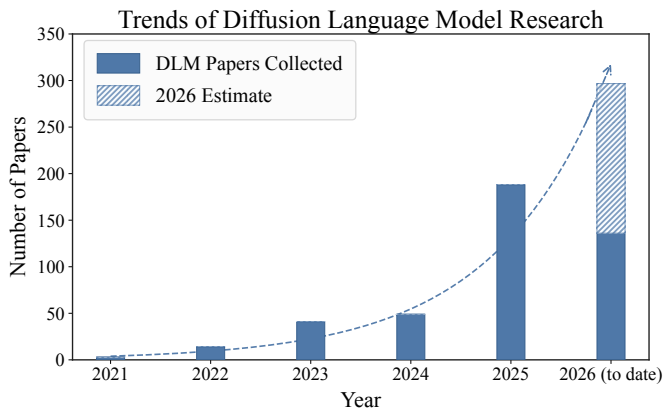


Fig. 2. Trend of diffusion language model papers. Counts are based on unique dated arXiv entries collected in the associated Awesome-DLMs repository. For 2026, the lighter part shows a linear estimate based on papers collected by mid-May. The statistics are for reference only.

tages as follows:

- **Parallel Generation:** DLMs can generate multiple tokens in parallel through an iterative denoising process, significantly improving inference speed and throughput over autoregressive models.
- **Bidirectional Context:** DLMs naturally incorporate bidirectional context, enabling more nuanced language understanding and generation. They also produce richer contextual embeddings, which are beneficial for cross-modal generation tasks. This enables fine-grained control over the generation process as well.
- **Iterative Refinement:** The iterative denoising process allows DLMs to update their perceptions over multiple steps. By accepting high-confidence tokens early and retaining low-confidence regions as masked, Masked DLMs can progressively improve uncertain areas, often resulting in more coherent and higher-quality text generation.
- **Controllability:** DLMs can be conditioned on specific token positions or structures, making them well-suited for tasks like infilling and structured generation. Additionally, guidance techniques (e.g., classifier-free guidance) enable better control over style and semantic relevance.
- **Unified Modeling Across Modalities:** By applying a shared denoising-based modeling framework, DLMs naturally support unified text and vision generation tasks. This makes them particularly promising for multimodal applications that require both generation and understanding within a single model.

Despite the recent rise in popularity of DLMs, there remains a lack of a comprehensive survey that systematically covers the entire DLM ecosystem. We structured our survey as follows: Section 2 provides a comprehensive overview of modern language modeling paradigms, including autoregressive, masked, and diffusion-based approaches. Section 3 delves into the training methodologies for diffusion language models, covering both pre-training and subsequent fine-tuning techniques such as SFT and RL alignment. Sec-

tion 4 details various inference strategies and optimizations, focusing on techniques tailored for continuous and discrete space models. Section 5 explores the extension of diffusion models to multimodal contexts, surveying state-of-the-art models and architectures like LLaDA-V [29], MMaDA [31], and Dimple [30]. Section 6 presents and visualizes performance comparisons of DLMs. Section 7 showcases the diverse applications of DLMs in tasks ranging from text and code generation to computational biology. Section 8 highlights the challenges and limitations of diffusion language models, including issues of efficiency, reasoning, agent capability, and infrastructure, and also outlines promising directions for future research. To provide a consolidated overview, a taxonomy of DLMs is presented in Fig. 3.

2 PARADIGMS OF DIFFUSION LANGUAGE MODELS

Diffusion Language Models have emerged as a powerful non-autoregressive paradigm that balances generative quality with inference parallelism. Inspired by principles from non-equilibrium thermodynamics [171], DLMs learn to reverse a gradual noising process. This iterative refinement approach allows for parallel generation of the entire sequence, offering a potential solution to the inference bottleneck of AR models. DLMs can be broadly categorized based on the space in which the diffusion process operates: either continuous or discrete. Additionally, there are hybrid AR-Diffusion models that combine autoregressive and diffusion in various forms, aiming to leverage the complementary strengths of both paradigms. We present model information from several works in Table 1 and provide a comparison of different paradigms in Fig. 4.

2.1 Preliminaries of Modern Language Modeling

The field of language modeling has evolved through several distinct paradigms, each characterized by unique architectural choices, training objectives, and associated trade-offs. In this subsection, we provide a brief overview of recent transformer-based paradigms at scale, highlighting their core principles, mathematical formulations, and representative models. Earlier approaches are not included, as we focus on modern, large-scale designs here. This review serves to establish the conceptual foundation for understanding the emergence of diffusion language models as a novel and promising alternative that addresses key limitations of prior methods.

2.1.1 Masked Language Models

Masked Language Models (MLMs), popularized by BERT [172], represent a foundational paradigm that scales pretrained language models using transformer-based encoder-only architectures. Conceptually simple yet empirically powerful, MLMs learn bidirectional contextual representations by predicting randomly masked tokens within an input sequence, leveraging both preceding and succeeding context. This approach follows a denoising autoencoding framework, where a subset of input tokens is masked, and the model is trained to reconstruct them:

$$\mathcal{L}_{\text{MLM}} = \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{\mathcal{M} \sim \text{Mask}(x)} \left[- \sum_{i \in \mathcal{M}} \log P_{\theta}(x_i | x_{\setminus \mathcal{M}}) \right] \quad (1)$$

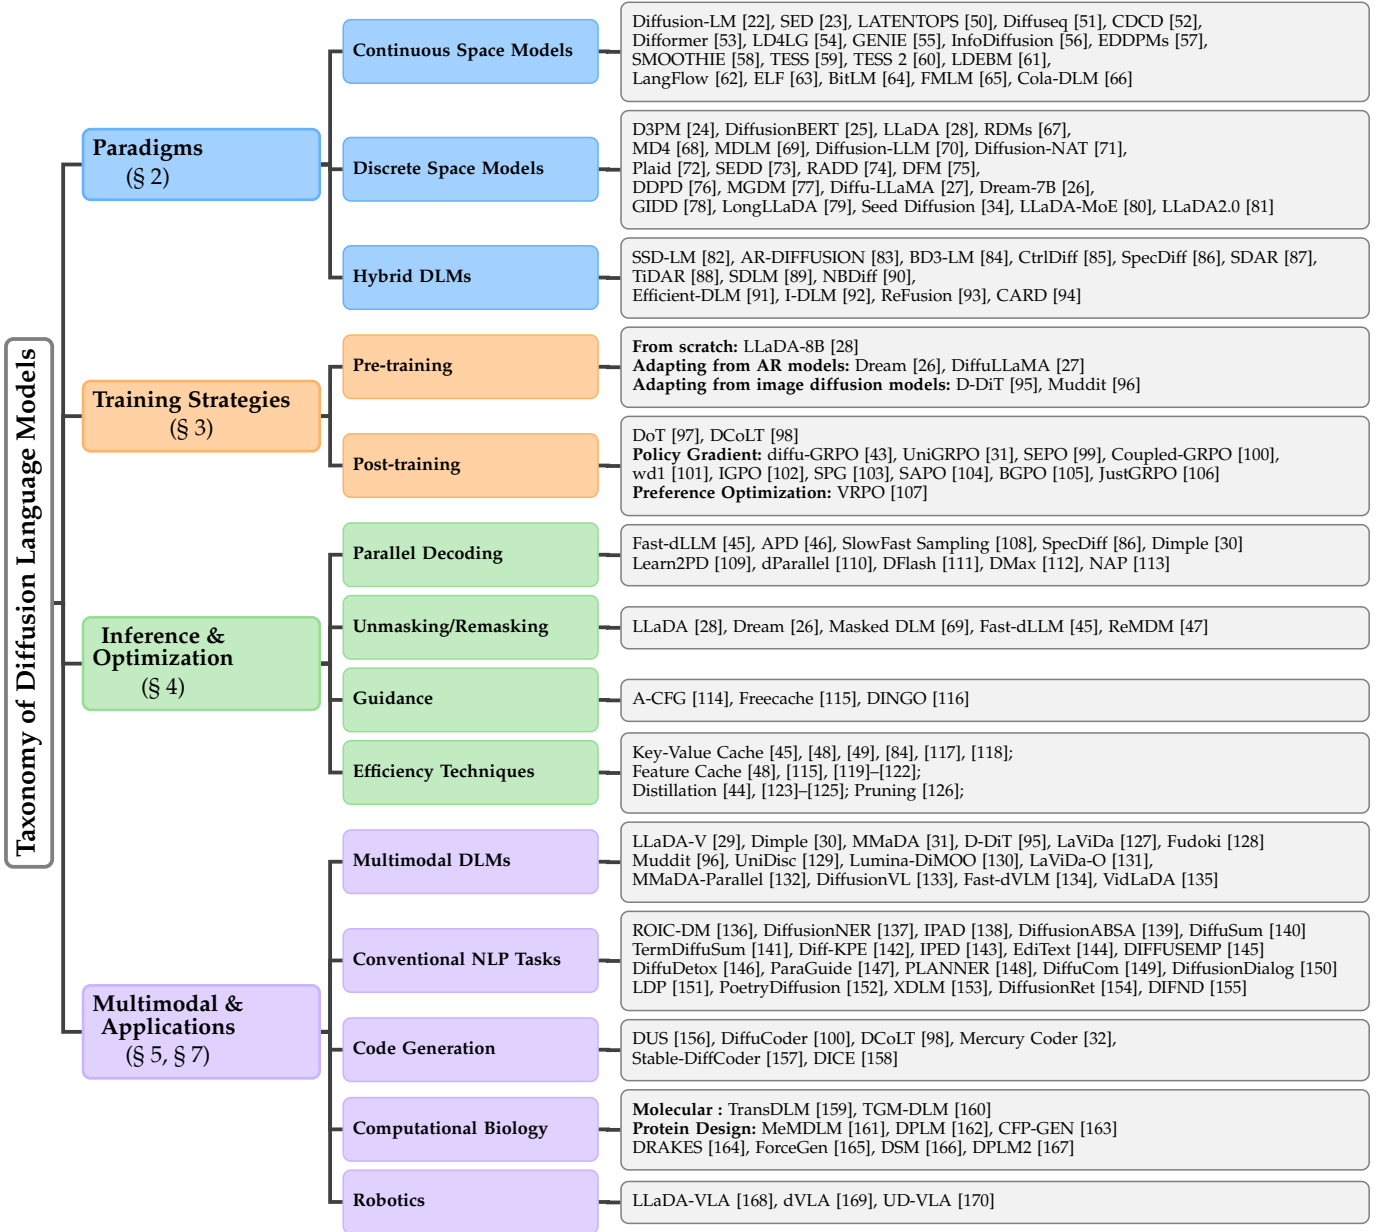


Fig. 3. A taxonomy of Diffusion Language Models, covering foundations, training and inference strategies, and key applications. The section numbers (§) correspond to the sections in this survey.

Here, x denotes the input sequence, \mathcal{M} is the set of masked positions, and $x_{\setminus \mathcal{M}}$ represents the visible (unmasked) context. BERT also introduces a next sentence prediction (NSP) objective to model inter-sentence relationships:

$$\mathcal{L}_{\text{NSP}} = \mathbb{E}_{(A,B,y) \sim \mathcal{D}} [-\log P_{\theta}(y | A, B)] \quad (2)$$

where (A, B) is a pair of text segments, and $y \in \{0, 1\}$ indicates whether B follows A in the original text.

BERT’s effectiveness in language understanding tasks such as sentiment analysis, named entity recognition, and question answering has inspired numerous improved variants. For instance, RoBERTa [173] removes the NSP objective and adopts more aggressive training strategies, while ALBERT [174] introduces parameter sharing and matrix factorization for efficiency. DeBERTa [175] further enhances

contextual encoding with disentangled attention and improved decoding mechanisms for masked token prediction.

Despite their strengths in understanding tasks, MLMs are not inherently designed for generative tasks, generating text requires specialized fine-tuning strategies or decoding schemes, making them unsuitable for open-ended generation without significant architectural modifications.

2.1.2 Autoregressive Language Models

Illustrated by GPT series [1], [2], [13], [14] and Transformer-XL [176], further advanced by subsequent LLMs [3]–[5], [177], autoregressive language models have become the backbone of modern generative AI, characterized by their unidirectional, left-to-right token generation process. Unlike bidirectional models, Autoregressive LMs factorize the joint

probability of a text sequence into a product of conditional probabilities:

$$P(x) = \prod_{i=1}^n P_{\theta}(x_i | x_1, x_2, \dots, x_{i-1}) \quad (3)$$

Given a token sequence $X = (x_1, x_2, \dots, x_n)$, the training objective is to maximize the log-likelihood of the sequence under this factorization:

$$\mathcal{L}_{\text{AR}} = \mathbb{E}_{X \sim \mathcal{D}} \left[- \sum_{i=1}^n \log P_{\theta}(x_i | x_1, \dots, x_{i-1}) \right] \quad (4)$$

This is typically implemented using a decoder-only Transformer architecture with causal attention masking and teacher forcing during training, ensuring that each token prediction is conditioned only on preceding tokens while enabling parallel computation of the loss.

The sequential generation formulation is both a strength and a limitation. On one hand, it aligns with text generation tasks and facilitates straightforward sampling, naturally suits various applications. On the other hand, it imposes a fundamental bottleneck on inference speed, as token generation is inherently sequential and cannot be parallelized. This trade-off between generation quality and latency has become a central challenge in advancing AR models. Beyond the standard next-token prediction (NTP), recent research has explored multi-token prediction (MTP) [16], [178] to accelerate inference by generating multiple tokens per step. These efforts share conceptual similarities with parallel decoding strategies employed in DLMs, while some other works are directly inspired by diffusion process to align LLMs [179].

2.1.3 Other Paradigms

Sequence-to-Sequence Models. Sequence-to-sequence (Seq2Seq) models [180], an early yet powerful paradigm, are built on an encoder-decoder architecture and serve as a versatile framework for conditional text generation tasks such as machine translation and summarization. Modern models like T5 [181] and BART [182] are prominent examples.

In this architecture, the encoder processes the source sequence to produce an intermediate representation, which the decoder then uses to generate the target sequence, typically in an autoregressive manner. While standard Seq2Seq decoders are autoregressive, the framework itself is highly flexible. Many DLMs, such as DiffuSeq [51] and SeqDiffuSeq [183], adapt this architecture by replacing the autoregressive decoder with a non-autoregressive diffusion decoder, leveraging the encoder’s strong conditioning ability to guide the denoising process in generation.

Permutation Language Models. Permutation Language Models (PLM), exemplified by XLNet [184], offer an alternative approach to incorporating bidirectional context within a generative framework. PLMs are trained to predict tokens in a sequence, but in a random, permuted order rather than a fixed left-to-right order. The objective is to maximize the expected log-likelihood over all possible permutations of the factorization order:

$$\mathcal{L}_{\text{PLM}} = \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}_T} \left[- \sum_{t=1}^T \log P_{\theta}(x_{z_t} | \mathbf{x}_{\mathbf{z}_{<t}}) \right] \quad (5)$$

where \mathcal{Z}_T denotes the set of all possible permutations of a sequence of length T , and $\mathbf{z}_t, \mathbf{z}_{<t}$ refer to the t -th and first $t-1$ elements of a given permutation $\mathbf{z} \in \mathcal{Z}_T$. This formulation allows the model to capture bidirectional context for each token, combining the advantages of bidirectional context (like MLMs) with a coherent autoregressive generation process. This contrasts with DLMs, which achieve bidirectionality through a parallel iterative refinement process.

2.2 Continuous Diffusion Language Models

Continuous-space DLMs model language by first mapping discrete tokens into a continuous embedding space. A diffusion process then models the data distribution in this continuous space [22], [23]. Typically, diffusion models define a generative process by learning to reverse a predefined corruption process that gradually transforms data into noise. This process consists of a **forward (noising) process** and a **reverse (denoising) process**. The forward process gradually transforms a data sample \mathbf{x}_0 into noise over T timesteps via a fixed Markov chain:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}) \quad (6)$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \mu_t(\mathbf{x}_{t-1}), \Sigma_t), \quad (7)$$

where μ_t and Σ_t define the noise schedule. In many implementations, such as DDPM [18] and Rectified Flow [21], the marginal distribution at each timestep has a closed-form expression:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + b_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (8)$$

where α_t and b_t are deterministic functions of time t .

The reverse process learns to invert the corruption, starting from noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ and gradually denoising to recover a sample close to \mathbf{x}_0 . This is parameterized by a neural network $f_{\theta}(\mathbf{x}_t, t)$, typically implemented as a Transformer, which predicts a target quantity \mathbf{z} associated with the forward process (e.g., clean data, noise, or velocity). A common training objective takes the form:

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{z}} \left[\|f_{\theta}(\mathbf{x}_t, t) - \mathbf{z}\|^2 \right], \quad (9)$$

where \mathbf{x}_t is sampled via the forward process given \mathbf{x}_0 , and \mathbf{z} is the corresponding regression target derived from \mathbf{x}_0 and t .

After training, generation proceeds by sampling from the learned reverse process, starting from noise $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$. At each timestep $t = T, T-1, \dots, 1$, the model defines a conditional distribution $p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t)$ which aims to approximate the true reverse transition $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$. Sampling iteratively from these learned conditionals produces progressively less noisy latent states until an estimate of the original data \mathbf{x}_0 is recovered. After generating a denoised embedding $\hat{\mathbf{x}}_0$, a **rounding step** maps it back to a discrete token. This is typically done by nearest-neighbor searching in the embedding space, using a decoder head, or thresholding techniques [185].

Diffusion-LM [22] first introduces a diffusion process in the embedding space to create a non-autoregressive language generation model. By using a classifier-guidance mechanism similar to those in image diffusion models, it

achieves highly controllable text generation and infilling. LDEBM [61] presents a novel symbiosis of latent space EBMs and diffusion models in a variational learning framework to address the learning issues of energy-based priors, with a focus on interpretable text modeling. LATENTOPS [50] proposes an efficient framework for composable text operations by working within a compact latent space. It introduces an efficient sampler based on ordinary differential equation (ODE) to generate latent vectors guided by arbitrary plug-in control operators, which are then decoded into the desired text. Later, Diffuseq [51], a classifier-free DLM for sequence-to-sequence tasks is proposed, which corrupts only the target sequence embeddings in the forward process to achieve strong and diverse conditional text generation. The Self-conditioned Embedding Diffusion (SED) [23] framework conducts diffusion directly on a fixed, continuous token embedding space. By incorporating a self-conditioning mechanism, it achieves strong performance in both conditional and unconditional text generation, rivaling standard autoregressive models. CDCD [52] applies continuous diffusion to categorical data by embedding tokens into a continuous space. It proposes score interpolation, which uniquely allows the model to be trained with a cross-entropy loss, and time warping, an adaptive strategy to efficiently schedule noise levels during training. To address optimization challenges in the embedding space, Difformer [53] introduces an anchor loss to prevent embedding collapse and a noise rescaling framework to mitigate model degeneration. LD4LG [54] leverages a pretrained language model as a powerful autoencoder to create a compact latent space, where a continuous diffusion model is then trained for high-quality text generation. GENIE [55] proposes a large-scale pre-training framework for diffusion language models, introducing a novel continuous paragraph denoise objective to effectively learn from large corpora by reconstructing corrupted text paragraphs. InfoDiffusion [56] introduces an information entropy-aware noise schedule to guide the model toward a more human-like “key-info-first” process that prioritizes generating core content. EDDPMs [57] unify generation, reconstruction, and representation by generalizing the diffusion process with a parameterized encoder-decoder, enabling stable, joint training of all components within a single framework. SMOOTHIE [58] proposes a novel diffusion process that progressively smooths token embeddings based on semantic similarity, combining the advantages of continuous latent spaces and discrete token handling.

Continuous diffusion processes can also be formulated in the logit space rather than the embedding space. TESS [59] introduces a fully non-autoregressive framework that diffuses over a k-logit simplex representation of tokens and employs a novel self-conditioning mechanism tailored to this setting. Extending this, TESS 2 [60] scales the approach by adapting pretrained large autoregressive models into general-purpose diffusion language models through a diffusion-specific pretraining recipe and instruction tuning, enabling strong instruction-following capabilities.

Recent work revisits continuous-space DLMs with more explicit choices of representation, transport path, and decoding interface. LangFlow [62] connects embedding-space DLMs with Flow Matching via Bregman divergence and

studies continuous flow-based language modeling with an ODE-based likelihood bound, information-uniform noise scheduling, and self-conditioning. ELF [63] formulates language modeling as continuous-time flow matching in embedding space, keeping the denoising trajectory continuous until a final mapping back to discrete tokens. Cola-DLM [66] instead uses a hierarchical latent formulation: a Text VAE maps text into continuous latents, a block-causal DiT models a global semantic prior, and a conditional decoder realizes text from that latent representation. BitLM [64] replaces vocabulary-level token prediction with bitwise denoising over fixed-length binary token codes while preserving causal attention across blocks.

2.3 Discrete Diffusion Language Models

Discrete space DLMs define the diffusion process directly on the vocabulary of tokens, avoiding the need for a continuous embedding space during the diffusion itself. D3PM [24] first illustrates this by introducing a structured diffusion process over discrete tokens. The **forward process** corrupts a sequence by applying a transition matrix \mathbf{Q}_t at each step. This matrix defines the probability of a token transitioning to any other token in the vocabulary. The probability of a state \mathbf{x}_t given an initial state \mathbf{x}_0 is given by a categorical distribution:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \text{Cat}(\mathbf{x}_t; \mathbf{p} = \mathbf{x}_0 \bar{\mathbf{Q}}_t), \quad \text{where} \quad \bar{\mathbf{Q}}_t = \prod_{i=1}^t \mathbf{Q}_i$$

A common choice for \mathbf{Q}_t is an absorbing state transition, where each token has a probability of either remaining unchanged or transitioning to a special ‘[MASK]’ token. The **reverse process** learns to reverse these transitions, predicting the probability distribution of the original tokens given the corrupted sequence.

Over time, masked DLMs have emerged as a modern and highly effective evolution of discrete diffusion language models, forming the foundation for several recent large-scale efforts [27], [28]. We take LLaDA [28], the most representative model of this kind as an example. Inspired by earlier work on reparameterized and simplified training objectives [67], [68], [74], LLaDA is trained from scratch using a cross-entropy loss that is computed only over masked tokens:

$$\mathcal{L}(\theta) \triangleq -\mathbb{E}_{t, x_0, x_t} \left[\frac{1}{t} \sum_{i=1}^L \mathbf{1}[x_t^i = \text{M}] \log p_\theta(x_0^i | x_t) \right], \quad (10)$$

where x_0 is sampled from the training corpus, t is sampled uniformly from $[0, 1]$, and x_t is obtained by corrupting x_0 through the forward process. The indicator function $\mathbf{1}[\cdot]$ ensures that the loss is applied only to positions that have been masked. During inference, the generation process starts with a fully masked sequence of desired length. In each iterative step, the model takes the current sequence (containing a mix of generated tokens and ‘[MASK]’ tokens) and predicts a complete sequence of tokens. Based on the model’s prediction confidence and noise schedule, a certain number of the highest-confidence predictions are unmasked and fixed, while the remaining positions are re-masked. This refinement process continues iteratively until all ‘[MASK]’ tokens are resolved. This approach elegantly

TABLE 1
Summary of diffusion language models, configurations, and their design choices.

Model	Parameters	Diffusion type	Task	Training data
D3PM [24]	70M	Discrete	Language	65B tokens
Diffusion-LM [22]	100M & 300M	Continuous	Language	-
Diffuseq [51]	91M	Continuous	Language	565K sentence pairs
SSD-LM [82]	400M	Continuous	Language	123B tokens
DiffusionBERT [25]	110M	Discrete	Language	16B tokens
CDCD [52]	1.3B	Continuous	Language	315B tokens
LD4LG [54]	188M	Continuous	Language	5.2M sentence pairs
SeqDiffuSeq [183]	65M & 110M	Continuous	Language	45B tokens
TESS [59]	125M & 355M	Continuous	Language	-
MDLM [69]	110M	Discrete	Language	622B tokens
DFM [75]	1.7B	Discrete	Language & Code	2.5T tokens
TESS-2 [60]	7B	Continuous	Language	360B tokens
LLaDA [28]	1B & 8B	Discrete	Language & Code	2.3T tokens
Mercury [32]	-	Discrete	Code	Trillions of tokens
LLaDA-1.5 [107]	8B	Discrete	Language	2.3T tokens
MMaDA [31]	8B	Discrete	Multimodal Unified	900B image-text tokens
Dream [26]	7B	Discrete	Language & Code	580B tokens
LLaDA-V [29]	8.4B	Discrete	Multimodal	3M image-text samples
LaViDa [127]	8.4B	Discrete	Multimodal	1.6M image-text samples
Dimple [30]	7B	Discrete	Multimodal	0.8B tokens
LongLLaDA [79]	8B	Discrete	Language & Code	2.3T tokens
DiffuCoder [100]	7B	Discrete	Code	130B tokens
LaViDa-O [131]	10.4B	Discrete	Multimodal Unified	> 200M image-text pairs
Lumina-DiMOO [130]	8B	Discrete	Multimodal Unified	> 110M image-text pairs
LLaDA-MoE [80]	7B (1.4B activated)	Discrete	Language	20T tokens
SDAR [87]	1.7B-30B series	Discrete	Language	54B tokens
TiDAR [88]	1.5B & 8B	Discrete	Language	150B tokens
SDLM [89]	3B & 32B	Discrete	Language	2.3B tokens
LLaDA2.0 [81]	16B & 100B	Discrete	Language	-
Stable-DiffCoder [157]	8B	Discrete	Code	1.3T tokens
Efficient-DLM [91]	1.5B, 4B & 8B	Discrete	Language	300B-500B tokens
DiffusionVL [133]	3B & 7B	Multimodal	Vision-language	738K instruction samples
VidLaDA [135]	8B	Multimodal	Video understanding	2.8M samples
LangFlow [62]	130M	Continuous	Language	-
ELF [63]	105M, 342M & 652M	Continuous	Language	45.2B tokens
BitLM [64]	0.6B, 1.7B, 4B & 8B	Continuous	Language	FineWeb-350B tokens

combines the bidirectional context of MLMs with a controllable, parallel generation process. LLaDA-8B, in particular, exhibits strong scalability and instruction-following ability, achieving performance on par with powerful autoregressive models such as LLaMA3-8B. This challenges the long-standing dominance of autoregressive models in large-scale language generation.

DiffusionBERT [25] combines a pre-trained BERT with a discrete diffusion process, leveraging its powerful denoising capabilities to learn the reverse process from a masked state. The model is further enhanced by a novel spindle noise schedule that considers token informativeness, achieving significant improvements in generation quality compared with previous DLMs. A different approach, Reparameterized Discrete diffusion Models (RDMs) [67], establishes an alternative formulation for the reverse process, which simplifies the training objective to a weighted cross-entropy loss. This enables more flexible and adaptive decoding strategies, leading to significant performance gains over previous discrete diffusion models. Similarly, MD4 [68] derives a simple weighted integral of cross-entropy losses as the continuous-time variational objective of masked diffusion models, providing a simple and generalized framework for training DLMs. Another analogous approach is MDLM [69], which introduces a simplified, Rao-Blackwellized objective that takes the form of a weighted average of masked lan-

guage modeling losses. Diffusion-LLM [70] demonstrates the scalability of DLMs by adapting pre-trained masked language models to diffusion paradigm and further task-specific finetuning and instruction finetuning, unlocking their versatility in solving general language tasks. Diffusion-NAT [71] unifies a discrete diffusion model with a PLM by reformulating the denoising process as a non-autoregressive masked token recovery task, allowing BART to act as an effective denoiser. Plaid [72] is the first diffusion language model trained to maximize data likelihood, demonstrating through scaling laws that it can outperform autoregressive models like GPT-2 on standard benchmarks. To improve the training objective, SEDD [73] introduces a score entropy loss to directly learn the ratios of the data distribution, which serves as a discrete extension of score matching. Reparameterized Absorbing Discrete Diffusion (RADD) [74] reveals that the concrete score in absorbing diffusion can be expressed as a time-independent conditional probability of the clean data, multiplied by an analytic, time-dependent scalar. It also formally unifies the training objectives of absorbing discrete diffusion and any-order autoregressive models. Discrete Flow Matching (DFM) [75] introduces a novel generative paradigm for discrete data that is analogous to continuous Flow Matching. The method learns a generating probability velocity to transform samples along a general family of probability paths from a source to a target distribution.

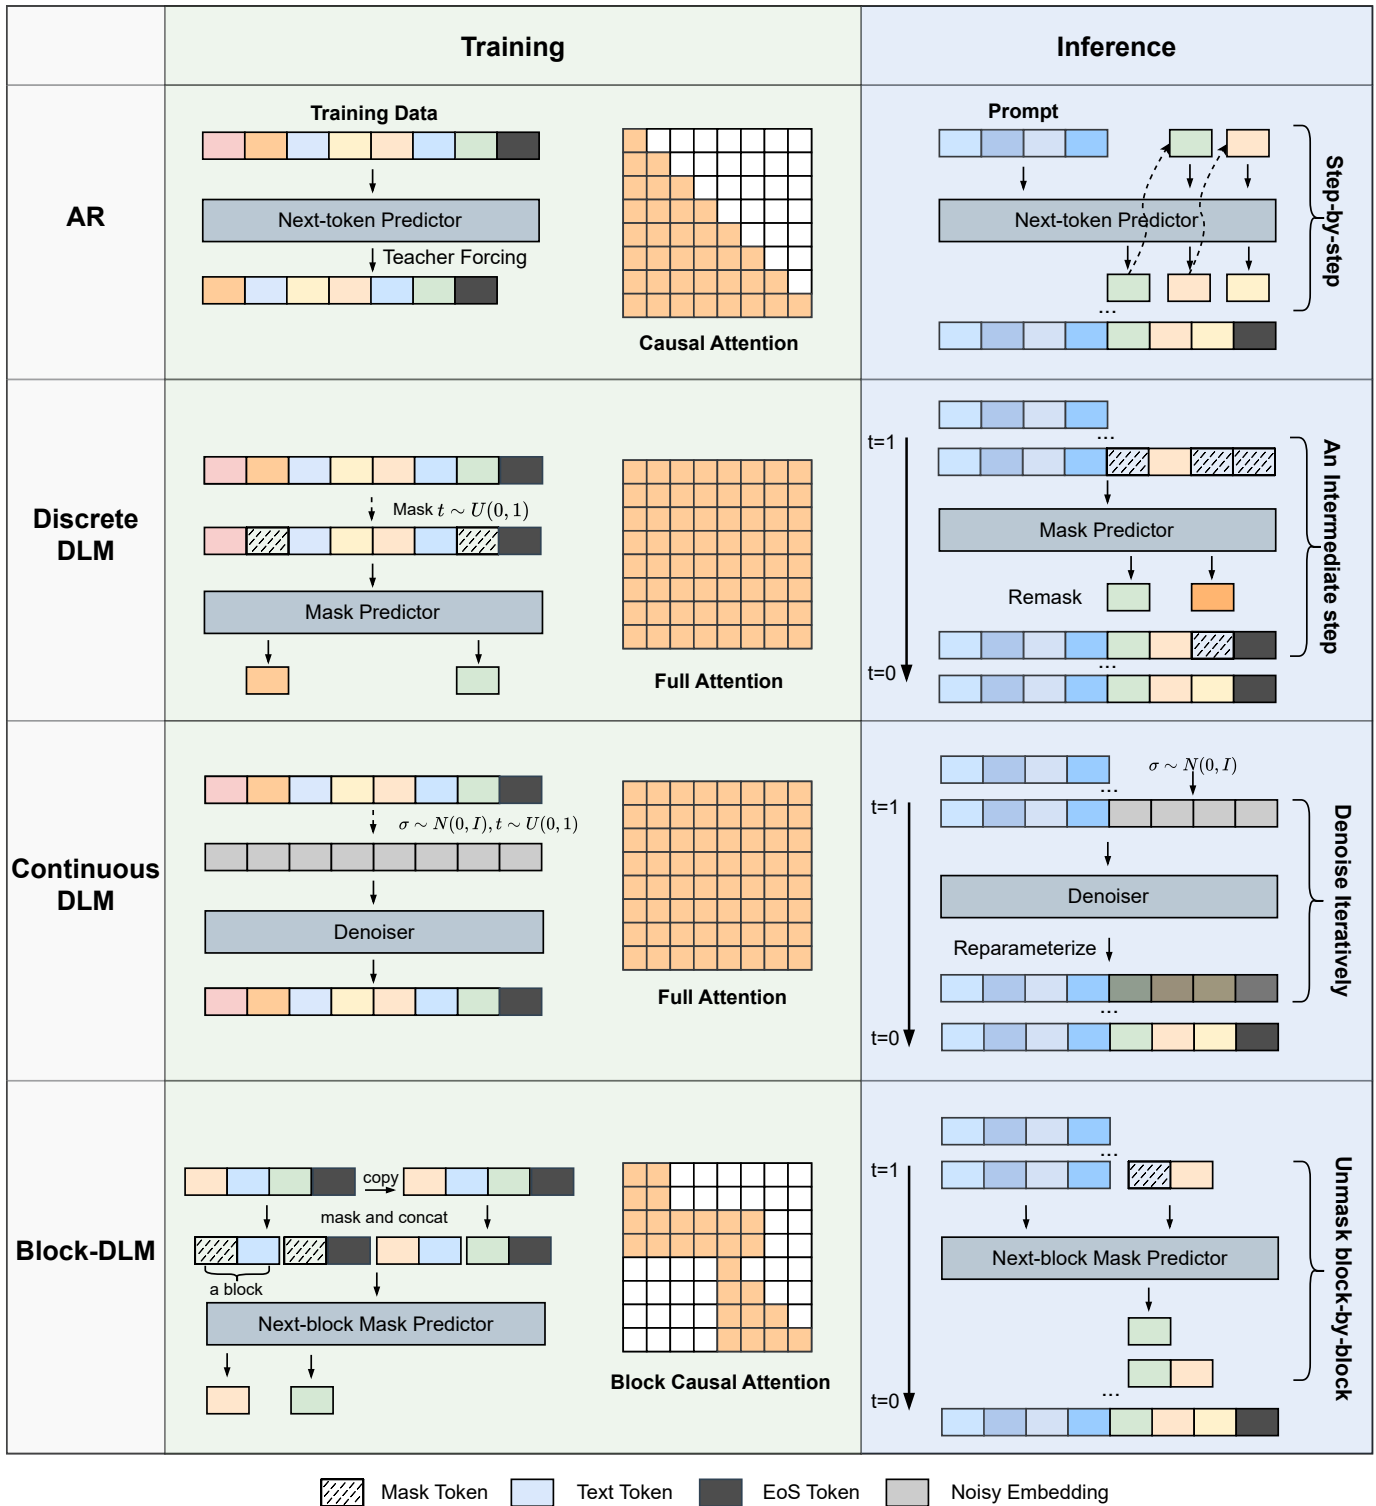


Fig. 4. An overview of training and inference procedures across different paradigms of Diffusion Language Models, with autoregressive (AR) models included for comparison. AR models are trained using teacher forcing and causal attention, whereas both discrete and continuous DLMs employ fully bidirectional attention mechanisms. Block-wise diffusion models, exemplified by BD3-LM [84], integrate autoregressive and diffusion strategies, and are trained using a specially designed block-causal attention mask.

By scaling the model architecture, DFM significantly closes the performance gap with autoregressive models on various benchmarks. DDPD [76] presents a framework that decouples the generation process into two specialized models: a planner and a denoiser. At each step, the planner identifies

the most corrupted token positions needing refinement, after which the denoiser predicts their values. To improve performance on complex reasoning tasks, MGDM [77] is introduced to address the problem of subgoal imbalance. This approach enhances discrete diffusion by prioritizing

more difficult subgoals during the learning process through a token-level reweighting mechanism. To address the challenge of scaling, a continual pre-training approach [27] is proposed to adapt existing autoregressive models, such as LLaMA, into diffusion language models. The resulting models, named DiffuGPT and DiffuLLaMA, are competitive with their AR counterparts while gaining diffusion-native capabilities like flexible infilling. Building on this observation, Dream-7B [26] is initialized from Qwen2.5 7B [186] and further trained with 580B tokens, largely outperforming existing DLMs and matching the performance of top-tier AR models. GIDD [78] is introduced to overcome the limitation that masked diffusion models cannot revise generated tokens. This framework generalizes the noising process by combining masking with uniform noise, which unlocks the model’s ability to self-correct mistakes and improves sample quality. Recently, to address long-context capabilities, LongLLaDA [79] provides the first systematic analysis of DLMs in this domain. It reveals that DLMs can maintain stable perplexity during direct context extrapolation and have better retrieval capabilities. LongLLaDA also introduces a training-free NTK-based RoPE extrapolation method, which significantly improves the extrapolation performance of DLMs, validating that established extrapolation scaling laws remain effective for DLMs. UltraLLaDA [187] extends this line of work by introducing a diffusion-aware NTK RoPE scaling and lightweight long-context post-training, enabling diffusion LLMs to reach 128K context windows and achieving substantially better retrieval and perplexity performance than training-free extrapolation methods. LLaDA-MoE [80] is the first work to integrate a sparse Mixture-of-Experts (MoE) architecture into diffusion language models, training a new MoE-based DLM from scratch on 20T tokens. Despite activating only about 1.4B parameters during inference, it surpasses larger dense diffusion models and achieves performance comparable to Qwen2.5-3B-Instruct across knowledge, coding, and reasoning benchmarks.

2.4 Hybrid AR-Diffusion Language Models

Hybrid AR-Diffusion models aim to strike a balance between the full parallelism of non-autoregressive models and the strong causal dependency modeling of autoregressive models. A prominent strategy for hybrid AR-diffusion modeling adopts a block-wise semi-autoregressive generation process. In this setting, the model generates blocks of tokens autoregressively, while the tokens within each block are generated in parallel using a diffusion-like iterative process. Early efforts such as SSD-LM [82] pioneered hybrid approaches by a block-wise continuous diffusion process on simplex representations, AR-DIFFUSION [83] illustrates a multi-level diffusion process and achieves semi-autoregressive generation by adjusting timestep according to token position. Recent representative model BD3-LM [84] further advances this direction on discrete models, demonstrating strong performance compared to pure AR and diffusion models. CtrlDiff [85] improves this paradigm by introducing dynamic block prediction techniques to enhance block-level efficiency and control. SDAR [87] further strengthens this hybrid paradigm by converting a pretrained autoregressive model into a blockwise diffusion

model through a lightweight adaptation stage. It preserves AR-level performance while enabling efficient parallel intra-block generation, achieving scalable speedups without sacrificing quality.

The generation process in these models usually consists of two nested loops. In the outer loop, blocks of tokens are generated autoregressively, with each block conditioned on previously generated blocks. Within each block, the inner loop performs parallel token-wise generation through a diffusion-style iterative denoising process. In BD3-LM, the training objective is formalized as:

$$\mathcal{L}_{\text{BD}}(\mathbf{x}, \theta) := - \sum_{b=1}^B \mathbb{E}_{t \sim [0,1]} \mathbb{E}_q \frac{1}{t} \log p_{\theta}(\mathbf{x}^b | \mathbf{x}_t^b, \mathbf{x}^{<b}) \quad (11)$$

This hybrid strategy enables the model to capture long-range dependencies across blocks via autoregression, while simultaneously accelerating generation within each block through parallel diffusion. The design also supports flexible output lengths and KV-Cache which is widely used in AR models [84].

Notably, recent masked diffusion language models [28], [31] also adopt similar semi-autoregressive block-based decoding strategies, which can be seen as instances of hybrid AR-diffusion modeling.

Beyond block-based approaches that combine AR and diffusion at sequence level, hybridization can also occur at the architectural level, where some part of the neural network, typically the encoder, diffuses the entire sequence altogether to an intermediate representation, then an autoregressive decoder generates the final sequence [188]. LADIDA [189] is a slightly different approach that diffuses at document level but decodes sentences by an AR decoder. SpecDiff [86] proposes a collaborative speculative decoding framework, where a lightweight diffusion model drafts candidate outputs, which are then validated and finalized by a large AR model. TiDAR [88] proposes a sequence-level hybrid architecture that integrates diffusion-based parallel drafting and autoregressive sampling within a single forward pass through structured causal-bidirectional attention. It effectively unifies the efficiency of diffusion models with the quality of AR decoding, achieving up to $5\times$ throughput improvements while maintaining AR-level performance. SDLM [89] introduces the Next Sequence Prediction (NSP) paradigm, which unifies next-token and next-block prediction to enable adaptive-length generation. By retrofitting pretrained autoregressive models with parallel block training and confidence-based dynamic decoding, SDLM achieves efficient diffusion-style intra-block generation while remaining KV-cache compatible.

3 DLMS: PRE-TRAINING AND POST-TRAINING

3.1 Pre-training and Supervised Fine-tuning

The pretraining process of DLMs largely follows procedures similar to those used in autoregressive language models (for discrete DLMs) or image diffusion models (for continuous DLMs), with relatively fewer design spaces. This section briefly summarizes existing approaches for DLM pretraining, aiming to bridge the methodological gap between DLMs and AR models.

To accelerate training, particularly for large-scale models, it is common practice to initialize DLMs from pretrained AR language models or image diffusion models [27], [190]. DiffuGPT and DiffuLLaMA [27] try to initialize masked DLMs with open-source LLMs from 127M to 7B parameters, found that DLMs can be efficiently adapted from AR models, significantly reducing training time and cost while achieving comparable or even superior performance to their AR counterparts. Building on this insight, Dream-7B is initialized from Qwen 2.5 7B [186], and is reported to outperform both LLaDA-8B and LLaMA3-8B on various benchmarks. Some multimodal DLMs, on the other hand, are initialized from pretrained image diffusion models. D-DiT [95] and Muddit [96] are initialized from pretrained MM-DiT backbones from SD3 [11] and Meissonic [191] respectively. Although these models are not originally designed for text generation, their latent representations contain intrinsic language-aligned knowledge, which can effectively facilitate the training of language modeling while retaining strong visual generation capabilities.

In terms of scaling properties, recent scaling-law analyses [192], [193] reveal that DLMs exhibit distinct compute-data tradeoffs from AR models: they are substantially more data-hungry under compute constraints, yet possess far greater data reuse potential under multi-epoch training, offering a principled foundation for designing optimal DLM training regimes.

Supervised fine-tuning in DLMs generally mirrors that of AR models. For masked DLMs like LLaDA [28], prompt tokens are left unmasked while response tokens are selectively masked, enabling the model to learn conditional response generation in a manner compatible with pre-training. In continuous DLMs, SFT can also be performed by corrupting only the response segment, as demonstrated in TESS2 [60].

Despite the overall similarity with AR training paradigms, DLMs face several unique challenges due to their diffusion-based formulation. A major issue lies in the loss computation efficiency of masked DLMs. In typical masked DLM training, only $\sim 50\%$ of tokens (on average) are involved in the loss computation, if timesteps are sampled uniformly. This reduces data utilization and may lead to suboptimal gradients, particularly if critical answer tokens are excluded from the loss. To address this, LaViDa [127] proposes a complementary masking strategy: each training sample is duplicated with two disjoint masking patterns, ensuring that all tokens are included in the loss computation at least once. Furthermore, due to the training-inference discrepancy, as illustrated in [194], the model performs significantly better during training than at inference time. The authors propose a two-step diffusion process and an improved scheduling technique to mitigate this issue.

3.2 Post-training for Reasoning Capabilities

Exploration of reasoning capabilities is becoming increasingly popular in DLMs as their performance on language tasks improves. Typically, reasoning capabilities are gained through fine-tuning on reasoning datasets. For DLMs, this presents a unique and formidable challenge. Traditional Chain-of-Thought (CoT) methods are based on the sequential nature of AR models to reason step-by-step, but

DLMs generate tokens in parallel. The most successful post-training techniques in the AR domain, particularly those based on reinforcement learning (RL) and policy gradient methods, are built upon the ability to efficiently compute the log-probability of a generated sequence. This is straightforward in AR models due to their factorizable, sequential nature. In DLMs, where generation is an iterative, non-sequential process, the log-likelihood is intractable, creating a significant technical barrier to applying the mature suite of RL algorithms developed for AR models to DLMs. Intuitively, we categorize these works into three main streams, which form the structure of this subsection: (1) Parallelizing the reasoning chain, where CoT in AR models is adapted to DLMs in parallel generation. (2) Adapting policy gradient methods, where variants of popular algorithms like GRPO are introduced to DLMs. (3) Adapting preference optimization methods such as DPO to DLMs.

3.2.1 DoT and DCoLT: Parallelizing the Reasoning Chain

One of the pioneering works to elicit complex reasoning in DLMs is Diffusion-of-Thought (DoT) [97], which adapts the popular Chain-of-Thought paradigm to the diffusion framework. Instead of generating reasoning steps sequentially like autoregressive models, DoT formulates them as intermediate thoughts that are refined in parallel throughout the diffusion denoising process. The approach is implemented by fine-tuning pre-trained DLMs such as Plaid [72] and SEDD [73] on datasets containing problems and their corresponding step-by-step rationales. To enhance the model’s ability to recover from its own mistakes, DoT introduces specialized training techniques like scheduled sampling and coupled sampling, which exposes the model to its own generated errors during training to improve its self-correction capabilities. This post-training methodology enables smaller DLMs to achieve impressive reasoning performance, even outperforming significantly larger autoregressive models on certain mathematical and logical reasoning benchmarks.

A more recent approach, Diffusion Chain of Lateral Thought (DCoLT) [98], introduces a distinct RL-based reasoning framework inspired by the cognitive concept of lateral thinking, which contrasts with the step-by-step vertical thinking of traditional CoT methods. Instead of supervising intermediate steps, DCoLT treats each step of reverse diffusion process as a latent thinking action, but optimizes the entire multi-step denoising trajectory with outcome-based RL to maximize a reward on the final answer. When applied to masked DLMs like LLaDA, DCoLT innovatively introduces an Unmasking Policy Module (UPM), which learns the optimal order for revealing tokens as part of the RL action space. This approach significantly boosts the reasoning capabilities of DLMs, with the DCoLT-reinforced LLaDA model achieving gains of +9.8% on GSM8K and +19.5% on HumanEval.

3.2.2 Adapting Policy Gradient Methods to DLMs

Score Entropy Policy Optimization (SEPO) [99] introduces RLHF to discrete DLMs, proposing a theoretically grounded framework to fine-tune discrete diffusion models using policy gradient methods and non-differentiable rewards. Operating within the score entropy framework, SEPO adapts modern policy gradient methods like PPO and GRPO by

TABLE 2

A brief summary of current post-training methods for DLMs’ reasoning capabilities, focusing on their algorithm type, major goal, key technical innovations, and applicable model types. Notably, most of these methods are based on policy gradient, and are designed for discrete DLMs.

Method	Algorithm Type	Core Goal	Key Technical Innovation	Model Type
DoT [97]	Non-RL Fine-tuning	Enable parallel Chain-of-Thought reasoning	Converts serial CoT into a parallel diffusion process; training-time self-correction	Continuous/Discrete
DCoLT [98]	Outcome-based RL	Enable non-linear latent reasoning	Lateral thought; outcome-based RL; Unmask Policy Module	Continuous/Discrete
SEPO [99]	Policy Gradient Framework(PPO/GRPO)	Finetune discrete DLMs with non-differentiable rewards	Low-variance gradient estimator via score entropy & importance sampling	Discrete
diffu-GRPO [43]	Policy Gradient (GRPO)	Introduce policy gradient method to DLMs	Efficient one-step log-probability estimator for applying GRPO to masked DLMs	Discrete
coupled-GRPO [100]	Policy Gradient (GRPO)	Reduce variance and maintain training efficiency	Coupled-sampling with complementary masks	Discrete
UniGRPO [31]	Policy Gradient (GRPO)	Unified reinforcement learning	Structured noising strategy; diversified reward modeling	Multimodal Discrete
VRPO [107]	Preference Optimization (DPO)	Align with human preferences	Sample budget allocation; antithetic sampling	Discrete
IGPO [102]	Policy Gradient (GRPO)	Use inpaint ability in DLMs to guide exploration	Inpainting-guided sampling; insert partial gt reasoning traces when sampling	Discrete
wd1 [101]	Policy Gradient	Mitigate computational overhead and bias	Weighted likelihood requiring only one approximation	Discrete
SAPO [104]	Policy Gradient (GRPO)	Learn structured coherent reasoning	Process-based reward function	Discrete
SPG [103]	Policy Gradient	Reduce bias in single bound methods	Sandwiched by upper and lower bound; block-wise masking	Discrete
BGPO [105]	Policy Gradient	Reduce memory overhead, scale to larger MC sample size	Boundary-guided lower bound; constant memory via gradient accumulation	Discrete
JustGRPO [106]	Policy Gradient (GRPO)	Expose arbitrary-order limits in reasoning	Applies standard GRPO while avoiding harmful arbitrary-order exploration	Discrete

using importance sampling to derive a stable and low-variance gradient estimate. This allows the model’s policy to be iteratively updated to maximize a reward function, making it a general framework for both conditional and unconditional generation. The objective function of SEPO is defined as follows:

$$l^A(\theta) = \mathbb{E}_{x \sim \pi_{\theta_{old}}} \left[\sum_{\substack{y \in \mathcal{X} \\ y \neq x}} w_{x,y} \log s_{\theta}(x, T - T_0)_y \right] \quad (12)$$

where the model parameters θ are optimized to maximize the expected log-likelihood of the score entropy s_{θ} weighted by $w_{x,y} = \pi_{\theta}(y) f(r_{x,y}^{T-T_0})$. The expectation is taken over samples x from the previous policy $\pi_{\theta_{old}}$. The function f can be selected to recover different policy gradient variants; for example, a clipped function yields PPO, while group-standardized rewards yield GRPO. This formulation enables stable and low-variance gradient estimation, even with non-differentiable rewards, and provides a flexible objective for fine-tuning discrete diffusion models. Numerical experiments across several discrete generative tasks showcase scalability and efficiency of SEPO, demonstrating that policy gradient RL can be soundly applied to discrete diffusion models.

d1 [43] provides a two-stage post-training framework for masked DLMs that combines supervised finetuning (SFT) with a novel policy gradient algorithm, diffu-GRPO. To adapt GRPO for DLMs, which lack a factorized likelihood, it introduces novel methods for both sequence log-probability

and per-token log-probability estimation. d1 uses a simple mean-field decomposition to approximate sequence log-probability by a product of independent per-token probabilities, while per-token log-probability is computed by performing a single forward pass on a fully masked completion, conditioned on a randomly masked prompt during each policy gradient update. Using different random masks for the prompt in each inner gradient update step serves as a form of regularization, improving training efficiency and stability. The full d1 pipeline, leveraging SFT followed by diffu-GRPO, demonstrates significant performance improvements on mathematical and planning reasoning tasks for the LLaDA model.

MMaDA [31], a unified multimodal diffusion model, presents a three-stage training pipeline. After first-stage pre-training, MMaDA employs a mixed Long chain-of-thought fine-tuning strategy, where reasoning trajectories from diverse tasks are curated into a unified format to align reasoning processes across modalities. This facilitates code start training of the third stage, which introduces UniGRPO, a policy-gradient reinforcement learning algorithm tailored for diffusion language models. UniGRPO overcomes the limitations of baseline methods like d1, by leveraging a structured noising strategy which samples a masking ratio $p_i \in [0, 1]$ uniformly rather than masking all response tokens. This ensures the model is exposed to various stages of multi-step diffusion denoising process, from almost fully masked to nearly unmasked, which is consistent with conventional diffusion training and boosts the utilization of the model’s multi-step denoising capabilities. Additionally, the

sequence-level log-likelihood is approximated by averaging over masked tokens.

DiffuCoder [100] is a 7B-parameter DLM specifically developed and analyzed for code generation. This work introduces an RL algorithm named coupled-GRPO, which is designed to be diffusion-native by leveraging the unique properties of the DLM generation process. The central innovation of coupled-GRPO is its coupled-sampling scheme for log-likelihood estimation. To obtain a more robust and lower-variance estimate, it constructs paired, complementary masks for each completion sequence in a training batch. For a given sequence, two masks are generated such that every token position is masked in exactly one of the two masks. The log-probability estimate is then derived by averaging the losses from these two complementary forward passes. This ensures that every token is evaluated in a partial-masking context during training, providing full token coverage and a more stable gradient signal compared to methods that use a single random mask or a full mask. Coupled-GRPO is shown to substantially improve DiffuCoder’s performance on code generation tasks, while also encouraging more parallel, less autoregressive generation patterns.

Sandwiched Policy Gradient (SPG) [103] leverages both an upper and a lower bound of the true log-likelihood to reduce bias in single-sided approximation policy gradient methods for DLMs. The two bounds of likelihood are estimated via Monte Carlo using a block-wise masking strategy to improve training stability. SPG reports to achieve state-of-the-art performance compared with baseline methods on various reasoning benchmarks when applied to LLaDA.

wd1 [101] introduces a novel policy optimization approach that reformulates the objective as a weighted likelihood, requiring only a single approximation for the current parametrized policy likelihood. This formulation reduces bias and improves both stability and training efficiency, outperforming prior diffusion-based RL methods by up to 16% accuracy on reasoning tasks.

IGPO [102] leverages the unique inpainting ability of masked diffusion models to guide exploration during reinforcement learning. By partially injecting ground-truth reasoning traces, IGPO alleviates the zero-advantage problem in group-based RL.

SAPo [104] proposes a step-aware policy optimization scheme that introduces fine-grained process rewards aligned with the latent reasoning hierarchy, mitigating “unstructured refinement” and yielding more interpretable multi-step reasoning traces.

To improve memory efficiency in ELBO-based RL for DLMs, BGPO [105] introduces a boundary-guided lower bound that allows large Monte Carlo sample sizes without increased memory usage, achieving stronger reasoning performance under the same hardware limits.

A recent work, JustGRPO [106], shows that effective reasoning in dLLMs can be better elicited by intentionally forgoing arbitrary-order generation during RL training and applying standard GRPO. This finding challenges the common view that flexible token orders are inherently advantageous: for reasoning tasks, generation order can become a critical training design choice rather than a free benefit.

3.2.3 Adapting Preference Optimization to DLMs

LLaDA 1.5 [107] proposes a novel framework called Variance-Reduced Preference Optimization (VRPO) to adapt preference optimization methods to discrete DLMs. The work identifies that applying Direct Preference Optimization (DPO) to discrete DLMs is challenging due to the high variance of the Evidence Lower Bound (ELBO) used to approximate log-likelihoods. VRPO addresses this by introducing two key unbiased variance reduction techniques: (1) Optimal allocation of the Monte Carlo sampling budget by sampling more diffusion timesteps rather than multiple masked versions per timestep, i.e. $n_t = n$ and $n_{y_t} = 1$ (2) Antithetic sampling, where the same timesteps and masked data are shared between the ELBO estimates of the current policy π_θ and the reference policy π_{ref} for the same input y_w or y_l . By applying VRPO to LLaDA, the resulting LLaDA 1.5 model shows significant and consistent improvements across mathematics, code, and alignment benchmarks.

4 INFERENCE STRATEGIES

Inference strategies for DLMs serve three key goals: (i) boosting generation quality like unmasking and remasking schedules, (ii) enabling finer content control, and (iii) improving efficiency via techniques such as KV/feature cache and step distillation. A brief overview is presented in Fig. 5.

4.1 Parallel Decoding

Parallel decoding naturally aligns with DLMs, leveraging their inherent mask-predict capability to generate multiple tokens simultaneously rather than sequentially. However, naïve parallelization can degrade coherence, motivating a series of adaptive strategies that balance efficiency and quality. Fast-dLLM [45] employs confidence-aware decoding, selectively unmasking tokens whose predicted probabilities exceed a threshold, and realizes up to 27.6× speedups without compromising quality. Adaptive Parallel Decoding (APD) [46] modulates the degree of parallelism on-the-fly by consulting a lightweight autoregressive auxiliary model, thus trading throughput for fidelity when necessary. Slow-Fast Sampling [108] introduces a two-stage schedule. Firstly, a cautious “slow” phase to locate stable tokens, then an aggressive “fast” phase to finalize them in bulk, achieving up to 34× acceleration when combined with caching. SpecDiff [86] further pushes throughput by using a discrete diffusion model as a fully parallel ‘drafter’ whose output is quickly verified (and corrected if needed) by a larger autoregressive model, yielding up to 7.2× speedups over vanilla AR generation. Dimple [30] employs confident parallel decoding, dynamically adapting the number of tokens revealed per step and cutting generation iterations by 1.5 - 7×. Recent research has further advanced the efficiency of parallel decoding in diffusion language models by introducing learnable and training-level optimization strategies. Learn2PD [109] introduces a learned adaptive parallel decoding policy, where a lightweight filter model predicts whether each token should be unmasked, replacing fixed confidence thresholds. dParallel [110] enhances decoding efficiency through certainty-forcing distillation, training diffusion models to reach high confidence for multiple tokens

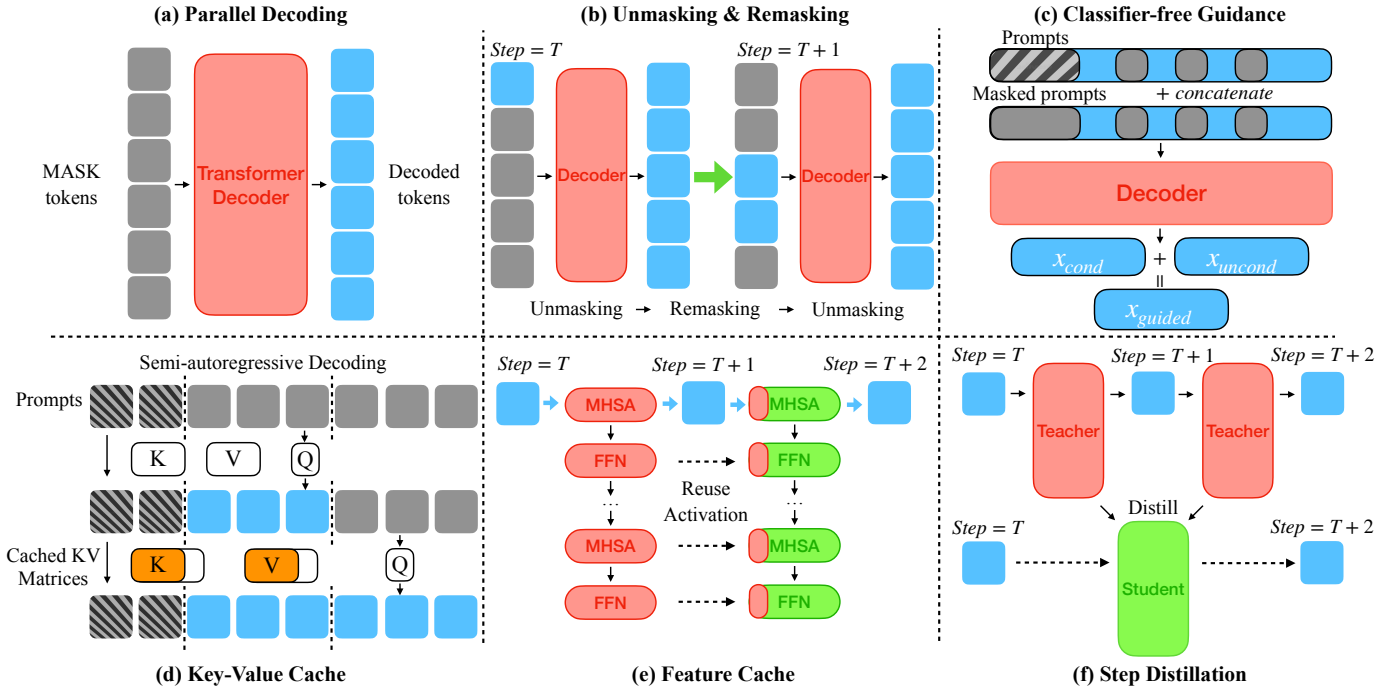


Fig. 5. Inference Techniques of Diffusion Language Models. We illustrate six different strategies here, including: (a) Parallel Decoding; (b) Unmasking & Remasking; (c) Classifier-free Guidance; (d) Key-Value Cache; (e) Feature Cache; and (f) Step Distillation.

in parallel. Collectively, these parallel decoding approaches substantially narrow the latency gap between diffusion and autoregressive models while preserving, and in some cases improving, generation quality.

4.2 Unmasking/Remasking

State-of-the-art open-source discrete DLMs such as LLaDA [28] and Dream [26] adopt a mask-predict paradigm: at each diffusion step they unmask high-confidence tokens and remask uncertain positions, iteratively refining the sequence. The choice of unmasking/remasking policy, i.e., low-confidence sampling, random selection, or adaptive temperature, therefore dominates both generation quality and convergence speed, making it one of the most critical inference levers. Early work Masked DLM [69] formalized two baselines: random remasking and confidence-ranked remasking, showing that prioritizing low-confidence positions yields better quality at no extra cost. Building on this insight, Fast-dLLM [45] introduces confidence-aware parallel decoding: every step it unmask all positions whose predicted probabilities exceed a global threshold, realizing up to 13 \times speedups while maintaining accuracy. Most recently, ReMDM [47] proposes a principled inference-time remasking sampler that can remask already decoded tokens for further refinement; by scaling the remasking budget, it offers a smooth compute-quality trade-off and closes the quality gap with autoregressive models under fixed compute. Collectively, these adaptive unmasking/remasking strategies substantially boost the efficiency and quality of diffusion language models, and they integrate cleanly with orthogonal accelerators that will be discussed later such as caching and step distillation.

4.3 Guidance

Guidance is a pivotal inference technique in diffusion models, steering the generative trajectory toward desired attributes and thereby enhancing output quality. In diffusion models, guidance refers to any technique that modifies the model’s denoising trajectory so that samples conform to a desired condition, such as a text prompt, a class label, or a stylistic attribute. The idea was popularized by classifier guidance [17], where gradients from an external classifier are added to the score estimate to nudge the sample toward a target class. Soon after, classifier-free guidance [114] removed the need for an extra classifier: the model is trained once with and without conditioning, and at inference the two score estimates are combined:

$$s_{\text{guided}} = s_{\text{uncond}} + \lambda(s_{\text{cond}} - s_{\text{uncond}}), \quad (13)$$

where λ is the guidance scale that balances fidelity to the condition against sample diversity. This simple formulation now underpins most text-to-image systems (e.g., Stable Diffusion [8]) and has been adopted by DLMs for prompt-controlled generation. Subsequent work refines CFG along several axes: dropout-augmented CFG smooths the quality-diversity curve; particle-based guidance blends multiple conditions; and p_2 -weighting rescales the noise term to stabilize high- λ sampling. In the text domain, newer schemes extend guidance to structural and semantic constraints. FreeCache [115] couples a lightweight autoregressive verifier with a discrete DLM: the verifier approves (or vetoes) draft tokens before they are committed, simultaneously enforcing coherence and enabling aggressive feature caching. DINGO [116] formulates regular-expression control as a dynamic-programming search over a DFA, guaranteeing constraint satisfaction without altering the model

distribution. In other discrete DLMs, guidance can also be applied at each diffusion step, optionally combined with masking/remasking or caching, to steer content (e.g., topic, sentiment) while preserving efficiency. Overall, guidance has become a cornerstone of diffusion inference, offering a lightweight, tunable handle for aligning model outputs with user intent.

4.4 Efficient Inference

Recent state-of-the-art diffusion language models [26], [32], [107] integrate the canonical Transformer architectures [195] with the step-wise stochastic inference procedures of diffusion processes. Consequently, efforts to accelerate inference in DLMs have converged on two complementary strategies: (1) lowering the per-step computational overhead of the Transformer backbone, e.g., through Key-Value (KV) Cache or Feature Cache. (2) reducing the total number of diffusion sampling steps, e.g., via Step Distillation.

Key-Value Cache. The conventional KV cache leverages the strictly autoregressive decoding pattern of LLMs and is therefore ill-suited to the bidirectional, multi-step generation paradigm of DLMs [49]. Recent work, however, shows that carefully redesigning the decoding schedule can recover much of its benefit. Block Diffusion [84] introduces Block Discrete Denoising Diffusion Language Models (BD3-LMs), which decode text autoregressively across coarse blocks while running diffusion within each block; once a block is finished, its keys and values are frozen and reused, enabling variable-length generation and measurable speedups. Fast-dLLM [45] keeps the blockwise view but adds a training-free, approximate DualCache that exploits the near-identity of KV activations across successive diffusion steps for both prefix and suffix tokens, delivering up to $27\times$ end-to-end throughput gains on LLaDA and Dream with $< 1\%$ accuracy loss. Complementing these block-based schemes, dKV-Cache [49] observes that token representations stabilize only after a position is decoded and therefore deploys a delayed, conditional cache that stores KVs one step later; this design achieves $2\text{-}10\times$ speedups on the same models with negligible quality drop. d^2 Cache [118] introduces a fine-grained dual adaptive caching scheme that adaptively refreshes only rapidly changing KV states while reusing stable ones. Elastic-Cache [117] proposes an attention/depth-aware adaptive refresh mechanism that selectively updates deeper layers and reuses stable shallow-layer caches. It performs attention-based drift detection to trigger cache refreshes only when the most-attended tokens exhibit significant changes, yielding up to $45\times$ speedup with minimal quality loss. Together, these results show that semi-autoregressive scheduling and delayed caching provide practical bridges between diffusion’s bidirectional conditioning and Transformer tricks originally devised for autoregression.

Feature Cache. Feature caching was first introduced by DeepCache [119], which leverages the strong similarity of intermediate U-Net activations across consecutive diffusion steps to avoid redundant computation. Follow-up work Δ -DiT [120], Learning-to-Cache [121], and FasterCache [122] demonstrate that the same principle transfers cleanly to Transformer-based diffusion models, yielding comparable

speedups without retraining. With the rise of diffusion language models, dLLM-Cache [48] extends feature caching to text by distinguishing two redundancies: prompt tokens remain almost static throughout denoising, whereas response tokens evolve only sparsely. It therefore pairs a long-interval prompt cache with an adaptive short-interval response cache refreshed only when a lightweight value-similarity test (“V-verify”) detects substantial change, achieving up to $9\times$ end-to-end speedups on LLaDA-8B and Dream-7B. Most recently, FreeCache [115] caches the KV/feature projections of already “clean” tokens and refreshes only dynamic positions, pushing acceleration further to $34\times$ while preserving fidelity. Collectively, these advances illustrate that feature caching can bring diffusion language models within striking distance of autoregressive LLMs in inference latency without sacrificing output quality.

Step Distillation. Step distillation is a widely adopted acceleration technique for diffusion models, collapsing the typical thousand-step denoising process into only a few and sometimes even a single sampling step, thereby drastically reducing inference time. Unlike the training-free methods discussed earlier, it imposes an offline cost: a compact student network must first be trained to mimic the teacher. Early work such as Progressive Distillation [124], followed by ADD [196] and LADD [197], progressively halves the step count or aligns intermediate distributions to preserve fidelity. Di4C [123] extends the framework to discrete diffusion by explicitly distilling inter-token correlations, enabling four to ten steps students that match teacher quality while providing $\sim 2\times$ speedups. Most recently, DLM-One [44] employs score-based distillation with adversarial regularization to train a continuous diffusion language model that generates an entire sequence in a single forward pass, achieving up to $500\times$ acceleration with near-teacher quality. Collectively, these works establish step distillation as the principal route toward closing the latency gap between diffusion and autoregressive language models.

5 MULTIMODAL AND UNIFIED APPROACHES

This section explores recent developments in extending DLMs to multimodal and unified architectures. Similar to autoregressive LLMs, DLMs can be naturally adapted to handle multimodal inputs and outputs. A straightforward approach is to accept vision inputs through a pretrained vision encoder. Following the success of LLaVA [198] in the AR domain, models such as LLaDA-V [29], LaViDa [127], and Dimple [30] employ vision encoders to extract image features, which are then projected into the same embedding space as text tokens. Beyond simple visual understanding, DLMs offer a promising pathway toward unified multimodal generation and understanding. Thanks to their shared denoising diffusion framework, DLMs naturally support joint modeling of different modalities. Visual inputs can be discretized using VQ-VAE, enabling training on multimodal inputs and outputs in a unified token space. Representative models such as MMaDA [31], Fudoki [128], and Muddit [96] exemplify this direction.

LLaDA and LLaDA’s Derivatives. We begin by introducing the LLaDA [28] family and its derivatives, which are built on the architecture and pretrained weights of the base LLaDA

model. LLaDA-V [29] integrates a vision encoder with an MLP-based projector that maps visual features into the language token embedding space, enabling effective visual instruction tuning. Following LLaVA-NeXT [199], LLaDA-V adopts three-stage tuning strategies. In the first stage, they only train the MLP projector to align visual representations with text embeddings using LLaVA’s training data. In the second stage, the model is further tuned by large-scale visual instruction data [200] using DLM objective. The third stage is to enhance multimodal reasoning capabilities by training on QA pairs with reasoning chains. Although the LLaDA backbone is slightly weaker than LLaMA3-8B [201] on pure text tasks, LLaDA-V achieves strong performance and better scalability across various benchmarks compared with LLaMA3-V trained on the same data. It narrows the performance gap with Qwen2-VL [202] and outperforms both hybrid and pure DLM-based models [95], [203], [204], demonstrating the effectiveness of diffusion architectures in multimodal understanding.

LaViDa [127] introduces a family of VLM based on LLaDA and Dream-7B [26]. Also utilizing a pretrained vision encoder, LaViDa uses a two-stage training strategy to train the projector and finetune the model respectively. LaViDa makes notable contributions to address training and inference challenges of multimodal DLMs. Typically, in masked DLMs, only about 50% of the tokens are masked for loss computation on average, which reduces efficiency and may omit critical answer tokens during VLM training, thereby causing gradient misalignment. LaViDa introduces complementary masking for effective training: For each sample, two masked versions with disjoint corrupted spans are generated, ensuring all tokens are eventually used in training and improving sample efficiency and gradient flow. During inference, LaViDa employs Prefix KV-Cache to cache the keys and values of visual and prompt tokens, significantly reducing latency and achieving a maximum speedup of $3.9\times$ with a marginal performance drop. Additionally, timestep shifting is used to unmask tokens earlier, further boosting generation quality. Empirical results show that LaViDa achieves competitive or superior performance to AR-based VLMs, while enjoying significant inference speedup.

Lavida-O [131] further extends LaViDa into a full-spectrum unified multimodal model capable of both high-quality image generation and fine-grained understanding tasks. It introduces a novel Elastic Mixture-of-Transformers (Elastic-MoT) architecture that decouples the model into a lightweight generation branch and a more powerful understanding branch, enabling scalable training and inference. Lavida-O uniquely supports localized object-level understanding, instruction-based image editing, high-resolution text-to-image synthesis (1024px), and interleaved reasoning and planning within a single unified diffusion framework.

Building upon LLaDA, MMaDA [31] further generalizes the architecture to support both multimodal understanding and generation. Unlike prior models, MMaDA eliminates the need for an explicit vision encoder by tokenizing images into discrete codes using VQ-VAE, and modeling all modalities jointly with a modality-agnostic diffusion transformer. This design allows seamless integration across text and image modalities without modality-specific components.

MMaDA also implements a mixed long CoT fine-tuning strategy that aligns CoT reasoning format across modalities. Moreover, UniGRPO, a unified policy-gradient based RL algorithm, is tailored specially for diffusion language models, making it possible to reason across modalities. Not only surpass similar-sized models like LLaMA3 for textual reasoning and Show-o [203] for multimodal understanding, MMaDA even excels professional image generation models like SDXL [10] in image generation.

MMaDA-Parallel [132] replaces the sequential reasoning-then-generation pipeline in MMaDA with a fully parallel multimodal diffusion framework, enabling text and images to interact bidirectionally at every denoising step. By jointly generating reasoning traces and visual outputs and further optimizing cross-modal consistency via a trajectory-level Parallel RL (ParaRL) algorithm, MMaDA-Parallel substantially improves semantic alignment and thinking-aware image synthesis performance.

Dimple. Dimple [30] introduces a large multimodal DLM, combining a vision encoder with a discrete DLM backbone. The authors identify that a pure discrete diffusion training approach suffers from significant instability, poor performance, and severe length bias. To overcome these challenges, Dimple proposes a novel two-phase training paradigm called Autoregressive-then-Diffusion. In the first phase, the model undergoes standard autoregressive training to effectively align the vision and language modalities. In the second phase, it switches to diffusion-based training to restore its parallel decoding capabilities. This hybrid strategy ensures stable and efficient training while achieving performance comparable to or even better than contemporary autoregressive models like LLaVA-NEXT.

For inference, Dimple introduces several techniques to improve efficiency and controllability. Confident Decoding dynamically adjusts the number of tokens generated in each step based on a confidence threshold, which reduces the total number of generation iterations. The model also successfully re-implements the prefilling technique, common in autoregressive models, to cache prompt tokens and achieve a speedup of up to $7\times$ with minimal performance loss. Furthermore, Dimple explores the use of Structure Priors, allowing for precise, fine-grained control over the response format and length, a feature that is difficult to achieve in autoregressive models.

D-DiT. Dual Diffusion Transformer (D-DiT) [95] is a large-scale fully end-to-end unified multimodal diffusion model that supports both text-to-image (T2I) and image-to-text (I2T) tasks. It directly addresses the challenges previous diffusion models faced in visual understanding tasks, which have been largely dominated by autoregressive models. The architecture is inspired by the Multimodal Diffusion Transformer (MM-DiT), featuring a dual-branch transformer that processes image and text tokens, with attention mechanisms allowing interaction between modalities in every layer. The model uses a frozen VAE for image processing and a frozen T5 encoder for text, and the major backbone MM-DiT is initialized from pretrained SD3 [11] weight.

One core innovation of D-DiT is its joint training objective, which combines continuous latent-space diffusion for images and discrete masked-token diffusion for text by jointly optimizing the sum of both modalities’ losses. Unlike

prior multimodal diffusion models that required an autoregressive component to decode text latents, D-DiT is fully diffusion-based and demonstrates competitive performance against other unified models.

UniDisc. Unified Multimodal Discrete Diffusion (UniDisc) [129] is proposed as a unified generative model for the joint text and image modeling, building upon discrete diffusion as an alternative to dominant AR approaches. Different from previously discussed D-DiT, UniDisc employs an entire masked diffusion process jointly on text and image tokens with full attention, learning to map a sequence of masked tokens back to a clean sequence from a shared vocabulary. Training is performed using a unified discrete diffusion objective from scratch, where tokens from both modalities are randomly masked and the model is supervised with a re-weighted cross-entropy loss.

A key advantage of UniDisc is its superior performance in conditional generation tasks, which is largely attributed to the effective use of classifier-free guidance. One of the most notable capabilities of UniDisc is its ability to perform joint image and text inpainting in a zero-shot manner, a feature not possible with previous AR or unified generative models. The author performs scaling analysis by scaling up the model up to 1.4B, demonstrating UniDisc outperforms AR models in terms of both performance and inference-time compute, with enhanced controllability and editability. However, UniDisc is found to be less training-efficient than a comparable AR model in terms of achieving the same validation loss.

Fudoki. Fudoki [128] is introduced as the first general-purpose unified multimodal model built entirely on the discrete flow matching framework, challenging the dominance of autoregressive (AR) and masking-based diffusion models. Instead of relying on a simple masking corruption process, Fudoki leverages a more general metric-induced probability path with kinetic optimal velocities, which allows for a more semantically meaningful corruption process and enables the model to continuously self-correct its predictions during iterative refinement. This self-correction capability is a key distinction from masked DLMs, where unmasked tokens are typically fixed and cannot be revised.

To reduce the high cost of training from scratch, Fudoki is initialized from a pre-trained AR-based MLLM, Janus-1.5B [205], and is then adapted to the discrete flow matching paradigm in a two-stage process. Its architecture is based on Janus-1.5B but uses a full attention mask to better capture global context and removes time embedding layers, as the model can implicitly infer the timestep from the corrupted input. Fudoki achieves performance comparable to state-of-the-art AR models in both visual understanding and image generation tasks, demonstrating a flexible trade-off between inference speed and quality. The model shows significant performance gains when test-time inference scaling techniques are applied, suggesting the potential of this architecture to be further explored for next-generation unified models.

Muddit. Muddit [96] is a pure unified discrete diffusion transformer that integrates a strong text-to-image backbone with a lightweight text decoder, enabling flexible and high-quality multimodal generation under a truly unified architecture. Initialized from pretrained MM-DiT from Meis-

sonic [191], the model is trained using a unified discrete diffusion objective, where text and image tokens are stochastically masked according to a cosine schedule and the model learns to predict the original tokens via a re-weighted cross-entropy loss. By a combination of the strength from a semantically rich visual prior and parallel discrete diffusion, Muddit achieves competitive or superior performance compared to significantly larger AR models across generation and understanding benchmarks. It also demonstrates several times speedup over AR baseline, highlighting the efficiency and scalability of a discrete diffusion approach when properly initialized.

Lumina-DiMOO. Lumina-DiMOO [130] is a state-of-the-art open-source unified multimodal diffusion model that achieves fast and high-quality multi-modal generation and understanding through a fully discrete diffusion framework. Built on LLaDA, it expands the vocabulary to include 8,192 visual tokens from aMUSEd-VQ [206] and employs a unified training objective over mixed text-image sequences. Lumina-DiMOO supports a wide range of tasks, including text-to-image generation, image editing, subject-driven and controllable generation, and advanced image understanding. It introduces innovations such as Max Logit-based Cache (ML-Cache) for sampling acceleration, parallel and block-wise sampling for efficient decoding, and an end-of-line special token to support arbitrary image resolutions. The training of Lumina-DiMOO is performed in four stages, culminating with Self-GRPO, a self-improving reinforcement learning algorithm that enhances generation and understanding alignment. Lumina-DiMOO ranks first among open-source models on the UniGenBench [207] leaderboard, offering $32\times$ speedup over AR baselines while delivering superior generation quality.

6 PERFORMANCE STUDY

In this section, we briefly compare the performance of various DLMs with AR models. We present visualizations based on several widely used benchmarks for evaluating DLMs, including PIQA [208] and HellaSwag [209] for general language understanding, HumanEval [210] for code generation, and GenEval [211], MME [212], MMMU [213] and GQA [214] for multimodal generation and comprehension. We also include GSM8K [215], a popular benchmark in DLM literature for assessing mathematical reasoning capabilities. The corresponding performance visualizations are shown in Fig. 6.

The DLMs surveyed range in size from under 1B to 8B parameters. For comparison, we also report the performance of representative AR models of similar scale. Performance data are primarily taken from original publications. If results were not available in the source papers, we consulted subsequent works that reported comparable evaluations.

Our findings suggest that DLMs generally perform competitively with AR models of comparable size. On general language understanding benchmarks such as PIQA and HellaSwag, models like LLaDA achieve performance that is slightly below or on par with AR models such as LLaMA2 [4] and Qwen2.5 [216]. However, DLMs exhibit stronger performance in math and science-related benchmarks, including GSM8K, GPQA [217], and MATH [218],

where models such as LLaDA and Dream consistently outperform similarly sized AR counterparts. In multimodal tasks, models like MMaDA [31] and LLaDA-V [29] often surpass AR-based multimodal models, highlighting the potential of DLMS in unified and cross-modal reasoning. On code generation tasks, DLMS also demonstrate competitive capabilities. Notably, DiffuCoder [100] achieves competitive HumanEval performance among open-source models, illustrating the potential of DLMS in structured, logic-heavy domains. Furthermore, closed-source DLMS such as Gemini Diffusion [33] and Mercury [32] achieve state-of-the-art results among all DLMS, rivaling top-tier AR models like GPT-4o.

Given the relatively limited training data and computational resources used to train most current DLMS, these results suggest that DLMS hold strong potential as viable alternatives to AR models in many real-world applications.

Recent scaling studies further show that DLMS tend to outperform AR models in data-constrained, multi-epoch regimes, likely because their any-order denoising objective enables more effective reuse of limited data [193].

7 APPLICATIONS ON DOWNSTREAM TASKS

7.1 Conventional NLP Tasks

Before the emergence of large-scale DLMS for general-purpose language generation, DLMS have already been applied to various conventional NLP tasks, such as text classification [136], named entity/scene recognition [137], [138], sentiment analysis [139], document summarization [140], [141], style transfer [147], [219], constrained generation [148]–[152], and machine translation [153], [220], etc.

ROIC-DM [136] is the first work to adapt diffusion models for robust text classification and inference. It applies the diffusion process directly to the class labels and conditions the denoising process on the input text, which can be further enhanced by incorporating traditional language models as advisors. DiffusionNER [137] formulates Named Entity Recognition as a boundary-denoising task. It applies a diffusion process to the start and end boundaries of entities, generating entity spans from random noise through an iterative refinement process. For scene text recognition, IPAD [138] introduces a parallel, iterative network that frames the task as conditional text generation, employing discrete diffusion and an easy-first decoding method to effectively balance recognition accuracy and inference speed. For aspect-based sentiment analysis, DiffusionABSA [139] employs a diffusion model to progressively extract the aspects step-by-step. DiffuSum [140] proposes a novel paradigm for extractive summarization by using a diffusion model to directly generate desired summary sentence representations. The final summary is then formed by extracting document sentences that best match these generated representations. For legal document summarization, TermDiffuSum [141] proposes a term-guided diffusion model that prioritizes sentences with legal terminology via a multifactor fusion noise weighting schedule. For keyphrase extraction, Diff-KPE [142] enhances phrase representations by guiding a text diffusion process with a Variational Information Bottleneck to generate and inject keyphrase information. IPED [143] treats relational

triple extraction as an implicit block diffusion task. EditText [144] introduces a controllable coarse-to-fine text editing framework by integrating an SEdit-based technique with a novel self-conditioning method for precise editing control. To generate more specific empathetic responses, DIFFUSEMP [145] utilizes a conditional diffusion model guided by multi-grained control signals (e.g., intent and semantic frames) that are integrated via a special masking strategy. DiffuDetox [146] utilizes a mixed diffusion approach for text detoxification, combining a conditional model to reduce toxicity with an unconditional model to ensure the fluency of the output text. A finetuned DiffuSeq model is shown to achieve state-of-the-art performance on fine-grained text style transfer tasks [219], while ParaGuide [147] introduces a more flexible plug-and-play framework that guides a paraphrase-conditioned diffusion model with off-the-shelf classifiers and style embedders at inference time. To generate fluent and diverse paragraphs while avoiding repetition, PLANNER [148] combines a latent diffusion planning module to generate semantic paragraph embeddings with an autoregressive decoding module to render the final text. DiffuCom [149] presents an efficient diffusion model for comment generation that uses context-aware attention mechanism and self-conditioning technology. DiffusionDialog [150] tackles the one-to-many problem in dialogue generation by performing a diffusion process with continuous latent variables, improving response diversity and inference speed. For paraphrase generation, LDP [151] models diffusion in a pretrained model’s latent space, avoiding the typical rounding step to achieve greater efficiency. For the highly constrained task of poetry generation, PoetryDiffusion [152] uniquely separates the task by using the diffusion model to generate semantics while a novel, independently trained metrical controller enforces structural rules like format and rhyme. In machine translation, XDLM [153] pioneers a cross-lingual pre-training objective for diffusion models, enabling them to effectively learn the mapping between languages in the pretraining stage. DiffusionRet [154] proposes a two-stage generative retrieval method that first utilizes a diffusion model to generate a pseudo-document from a query, which then serves as input for an n-gram-based model to retrieve the final document. DIFND [155] employs a diffusion model to generate debunking evidence and a multi-agent MLLM system for chain-of-debunk reasoning to improve accuracy and interpretability for multimodal fake news detection.

7.2 Code Generation

Although DLMS are rarely explicitly designed for code generation, the global planning and iterative refinement capabilities of them are particularly well-suited for the non-sequential nature of code generation. Foundational models like DiffuCoder [100], a 7B open-source model, have been developed specifically for this domain. DiffuCoder’s analysis reveals unique decoding behaviors, such as generation order becoming more flexible at higher temperatures. It also proposes coupled-GRPO, a novel sampling scheme that constructs complementary mask noise for completions used in training, which significantly improves the model’s performance on code generation tasks. Building on the

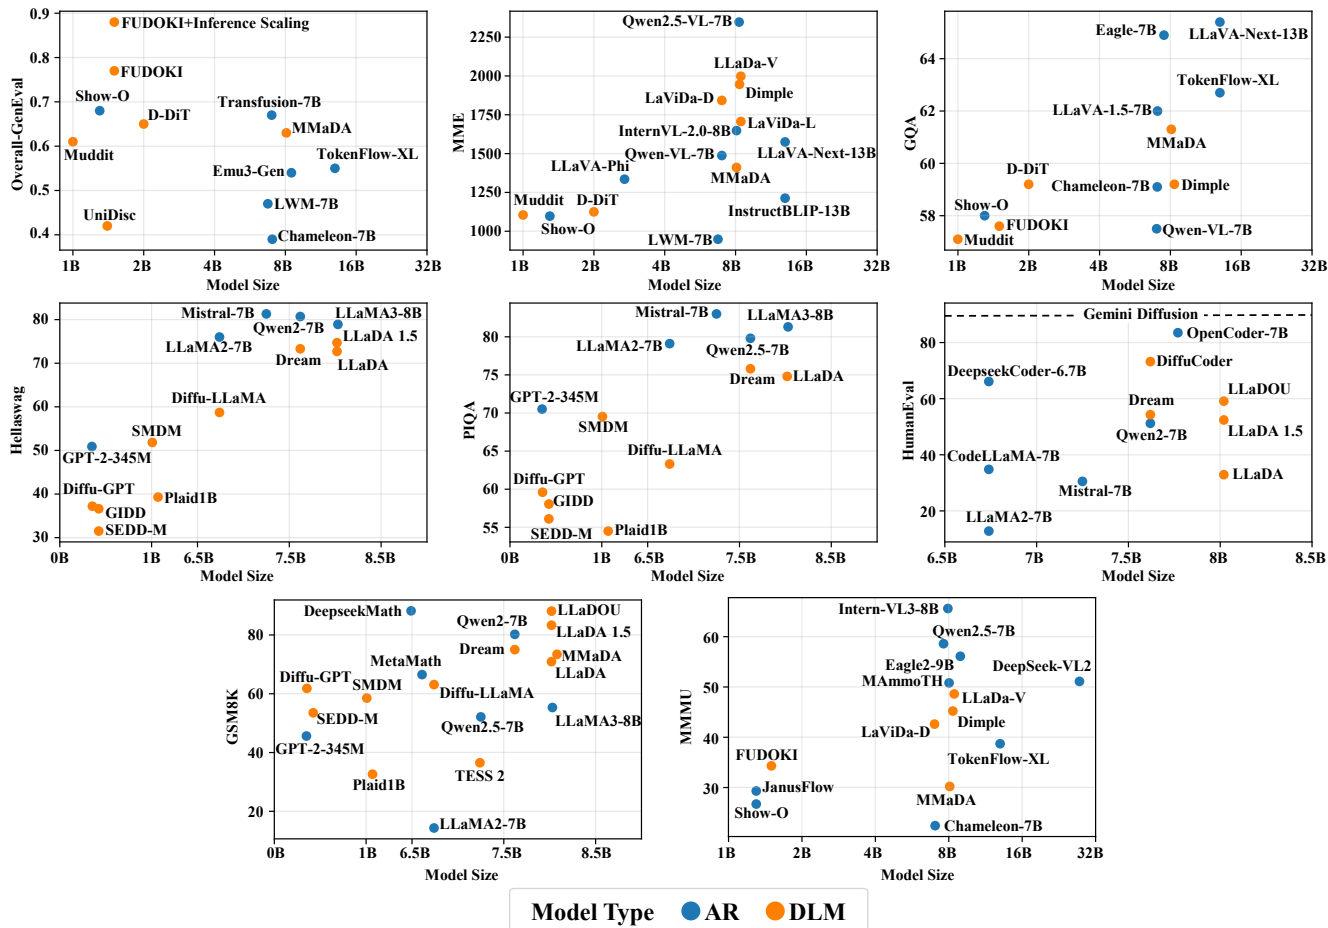


Fig. 6. Performance comparison on eight benchmarks: Overall-GenEval, MME, CQA, Hellaswag, PIQA, HumanEval, GSM8K, and MMMU. The horizontal axis in each subplot represents the model size, measured in the number of parameters. The vertical axis indicates the score under the corresponding benchmark, with higher scores reflecting better performance. Model types are distinguished by color: blue represents AR language models, while orange represents DLMs.

reasoning aspect, DCoLT [98] treats the entire reverse diffusion process as a form of non-linear, “lateral” thinking. With outcome-based RL and unmasking policy module, it achieves strong results on complex coding tasks. Dilated Unmasking Scheduler (DUS) [156] offers an inference-only, planner-free method that unmarks tokens in a non-adjacent pattern to minimize an upper bound on joint entropy gain at each denoising step, achieving promising results on code generation while improving speed-quality trade-off. Demonstrating the real-world potential of DLMs’ speed, Mercury Coder [32] is a commercial-scale diffusion model that achieves state-of-the-art throughput, outperforming speed-optimized autoregressive models by up to $10\times$ while maintaining comparable quality on major code benchmarks. Stable-DiffCoder [157] builds a block diffusion code model on the Seed-Coder training pipeline with block diffusion continual pretraining. DICE [158] targets CUDA kernel generation with the CuKe dataset and a bi-phase curated RL framework.

7.3 Biological and Scientific Applications

TransDLM [159] performs molecular optimization guided by a textual description of target properties to avoid the

error propagation. Another text-guided approach, TGM-DLM [160], focuses on molecular generation by collectively and iteratively updating token embeddings of SMILES strings. Without relying on additional data resources, TGM-DLM surpasses MolT5-Base in generation performance. DRAKES [164] introduces an RL-based fine-tuning method for discrete diffusion models that backpropagate rewards using the Gumbel-Softmax trick for DNA and protein design. For protein modeling, ForceGen [165] enables de novo protein design by using a protein language diffusion model to generate sequences that meet complex, nonlinear mechanical property-design objectives. MeMDLM [161] introduces a masked diffusion language model for de novo membrane protein design by fine-tuning the ESM-2 protein language model to generate novel and realistic transmembrane sequences. Inspired by LLaDA, DSM [166] introduces a enabling both high-quality representation learning and effective generative protein design. DPLM [162] offers a versatile protein language model that exhibits strong generative and predictive capabilities for protein sequences, and demonstrates superior performance in representation learning. DPLM2 [167] further extends the model into a multimodal protein foundation model that can simultaneously process both sequences and structures. By converting 3D

structural coordinates into discrete tokens, DPLM-2 learns the joint distribution of these two modalities. This enables the simultaneous co-generation of compatible protein sequences and their 3D structures, in addition to supporting conditional tasks such as protein folding and inverse folding. CFP-GEN [163] is a novel diffusion language model designed for Combinatorial Functional Protein Generation. It facilitates de novo protein design by integrating multimodal constraints, including functional, sequence, and structural information. CFP-GEN supports high-throughput generation of novel proteins with functionality comparable to that of natural proteins and achieves a high success rate in the design of multifunctional proteins.

7.4 Robotics

Recently, DLM-based vision-language-action (VLA) models have demonstrated strong potential in unifying perception, reasoning, and control. Built upon LLaDA, LLaDA-VLA [169] incorporates localized special-token classification alongside hierarchical action-structured decoding, leading to significant improvements over autoregressive VLA baselines in both simulation and real-world evaluations. dVLA [169] leverages pretrained MMaDA as a diffusion backbone to jointly generate visual subgoal images, textual chain-of-thought, and discretized actions, further introducing prefix-attention masking and KV caching for efficient long-horizon manipulation. Unified Diffusion VLA (UD-VLA) [170] proposes a Joint Discrete Denoising Diffusion Process that synchronously denoises future image and action tokens within a shared token space, achieving state-of-the-art performance on benchmarks with substantially improved inference speed.

8 CHALLENGES AND FUTURE DIRECTIONS

While diffusion language models have shown considerable promise across a wide range of tasks, several key challenges still remain and limit their practical deployment and broader application. In this section, we outline and discuss critical areas that require further research and innovation.

8.1 Major Challenges

1) Parallelism–Performance Trade-off. Diffusion language models are designed to generate multiple tokens in parallel. However, this parallelism often comes at the expense of generation quality and consistency. In discrete DLMs, unmasking multiple tokens simultaneously in a single step increases the denoising burden, which can lead to error accumulation. A central issue is the interdependence between tokens, known as the **Parallel Decoding Curse** [45]. When predicting multiple tokens at once, the model produces a distribution for each position and samples from them independently, failing to account for dependencies among positions. Consider a simple example where the training data consists only of two sequences: “ABABAB” and “BABABA”. Statistically, “A” and “B” appear with equal frequency at each position in the training data, leading DLMs to assign them similar probabilities during prediction. In autoregressive models, once the first “A” is generated, the model is likely to predict “B” next, preserving consistency. In

contrast, a DLM generating tokens in parallel may independently sample “A” for both the first and second positions, producing a sequence like “AAABBA”, which deviates from valid training patterns. Empirical studies show that this issue significantly affects DLM performance, particularly when the number of denoising steps is reduced [27]. This phenomenon is illustrated in Fig. 7. Future work may focus on mitigating this trade-off. Potential directions include introducing structured constraints, modeling inter-token dependencies more explicitly, or refining sampling strategies to improve coherence during parallel generation.

2) Infrastructure. While the training, fine-tuning, and inference of AR models have been significantly simplified and accelerated by open-source, highly optimized libraries and frameworks (e.g., Hugging Face Transformers [221]), DLMs still lag behind in this regard. Currently, major machine learning ecosystems offer little to no native support for DLMs, posing practical challenges for researchers and developers. Furthermore, during inference, DLMs lack mature, open-source deployment infrastructure akin to vLLM [222], making efficient serving of DLMs difficult.

3) Long Sequence and Dynamic-Length Generation. DLMs are typically trained to denoise fixed-length sequences under a diffusion-based objective, which makes it challenging to generalize to longer or dynamically sized sequences at inference time. Most existing DLMs are limited to a maximum context length of 4,096 tokens, and widely used extrapolation techniques in AR models for longer sequences remain underexplored in the DLM setting. This limitation hinders the applicability of DLMs in tasks requiring long-context understanding or complex reasoning. In addition, DLMs generally require the generation length to be predetermined during inference, making them ill-suited for dynamic-length generation. Although DLMs can predict an [EOS] token and omit displaying tokens generated afterward, the entire sequence is still fully updated throughout the denoising process, regardless of whether the generation has logically ended, which leads to unnecessary computational overhead. Recent work has begun to address this limitation through both training-based [223], [224] and training-free [225], [226] approaches. In addition, masked DLMs utilize full bidirectional attention at every denoising step, which incurs a computational cost of $\mathcal{O}(N^2)$ per step, where N is the sequence length. Assuming a fixed number of tokens are unmasked at each step, the total number of denoising steps scales linearly with N , leading to an overall inference complexity of $\mathcal{O}(N^3)$. Without architectural optimizations such as KV-Cache, this cubic time complexity severely limits the scalability of DLMs for long-sequence generation in real-world applications.

4) Scalability. Scalability remains an underexplored challenge for diffusion language models, particularly in comparison to autoregressive models. Although DLMs have shown promising results on certain metrics and benchmarks, they have yet to be scaled to the same extent as AR counterparts. Earlier public DLMs were mostly concentrated around the 7B–8B scale, while recent models such as LLaDA2.0 [81] have begun to extend public DLMs to larger model sizes. Nevertheless, compared with leading AR models that have been scaled to hundreds of billions or even trillions of parameters, such as Llama-3.1-405B [201], DeepSeek-V3-671B-

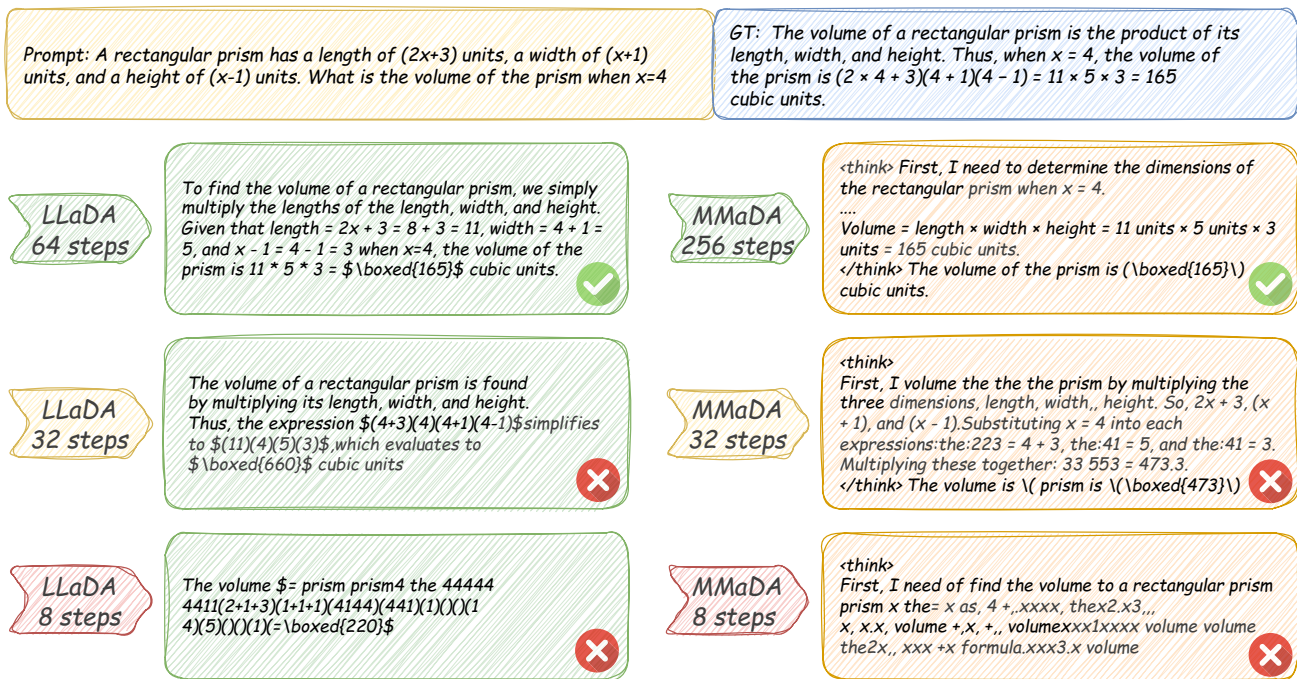


Fig. 7. Generation results of LLaDA [28] and MMaDA [31] under different denoising step settings. Note that the generation length is set to 128 tokens and 256 tokens for LLaDA and MMaDA respectively. Both models generate a correct and coherent response only when 1 or 2 tokens are unmasked at each step. With fewer steps and more parallelism, the responses are either incorrect or lack fluency and consistency. This illustrates the trade-off between parallelism and output quality in DLMs. We omit part of the thinking process of MMaDA with 256 steps for simplicity.

A37B MoE [16], Qwen3-235B-A22B MoE [227], and Kimi-K2-1T-A32B MoE [228], large-scale DLMs remain relatively less explored. Closed-source DLMs, such as Mercury and Gemini Diffusion, also fall short of state-of-the-art AR models across a wide range of benchmarks. Furthermore, many existing DLMs are trained either from previously pretrained AR models or built upon baseline DLMs (e.g., LLaDA) using limited datasets, which further constrains their scalability and performance. Therefore, the ability to further scale up DLMs still needs to be validated or explored.

8.2 Future Directions

Despite the challenges discussed above, DLMs present many promising directions for future exploration. Below, we briefly outline several under-explored directions and opportunities that could significantly advance the field:

- **Training Efficiency:** Current DLMs generally exhibit lower training efficiency compared to AR models, due to factors such as limited token usage during loss computation. Future research could explore hybrid DLM architectures or improved training schemes that match or exceed AR models in efficiency.
- **Quantization and Binarization (Low-bit DLMs):** While extensively studied in AR models, low-bit quantization and binarization remain largely unexplored in DLMs. Adapting these techniques to the diffusion paradigm could yield faster inference and reduced memory consumption, benefiting deployment in real-world systems.
- **Pruning and Distillation:** Model compression techniques such as pruning and knowledge distillation

have been successful in reducing model size and inference cost for AR models. Applying these techniques to DLMs could enhance their deployability, especially in resource-constrained or latency-critical environments.

- **Multimodal Unified Reasoning:** Although recent multimodal DLMs demonstrate impressive capabilities in cross-modal understanding and generation, most models are still limited to reasoning within a single modality at a time. Future efforts can focus on building unified DLMs capable of performing complex reasoning across multiple modalities in a truly integrated manner.
- **DLM-based Agents:** The potential of DLMs in powering intelligent agents remains largely underexplored. Leveraging their bidirectional context modeling, parallel decoding, and iterative refinement capabilities, DLM-based agents could offer greater flexibility and adaptability in dynamic environments, making them a promising alternative to traditional AR-based agent approaches.

9 CONCLUSION

In this survey, we present an in-depth overview of the entire landscape of diffusion language models. We outline the fundamental principles, taxonomy, and modeling paradigms of DLMs, and compare them with mainstream autoregressive models, highlighting their unique characteristics and advantages. We further explore the design space of training and inference, covering various training strategies and inference techniques for both quality and efficiency. Moreover, we

highlight recent advances in multimodal diffusion language models, demonstrating their capabilities in handling diverse data modalities. Finally, we discuss the limitations and challenges in this field, and outline promising directions for future research. We hope this survey serves as a comprehensive reference for researchers interested in diffusion-based language modeling, offering valuable insights about the current state of the field and its future prospects. We also encourage further exploration and innovation in this exciting area of research, as diffusion language models continue to evolve and push the boundaries of language understanding and generation.

REFERENCES

- [1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [2] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [3] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, "Palm: Scaling language modeling with pathways," *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.
- [4] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [5] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, "Qwen technical report," *arXiv preprint arXiv:2309.16609*, 2023.
- [6] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong *et al.*, "A survey of large language models," *arXiv preprint arXiv:2303.18223*, vol. 1, no. 2, 2023.
- [7] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi *et al.*, "Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning," *arXiv preprint arXiv:2501.12948*, 2025.
- [8] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 684–10 695.
- [9] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," *Advances in neural information processing systems*, vol. 35, pp. 36 479–36 494, 2022.
- [10] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach, "Sdxl: Improving latent diffusion models for high-resolution image synthesis," in *The Twelfth International Conference on Learning Representations*.
- [11] P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel *et al.*, "Scaling rectified flow transformers for high-resolution image synthesis," in *Forty-first international conference on machine learning*, 2024.
- [12] T. Brooks, B. Peebles, C. Holmes, W. DePue, Y. Guo, L. Jing, D. Schnurr, J. Taylor, T. Luhman, E. Luhman *et al.*, "Video generation models as world simulators," *OpenAI Blog*, vol. 1, p. 8, 2024.
- [13] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, "Improving language understanding by generative pre-training," 2018.
- [14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [15] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican *et al.*, "Gemini: a family of highly capable multimodal models," *arXiv preprint arXiv:2312.11805*, 2023.
- [16] A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan *et al.*, "Deepseek-v3 technical report," *arXiv preprint arXiv:2412.19437*, 2024.
- [17] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [18] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [19] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, 2020.
- [20] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *International Conference on Learning Representations*.
- [21] X. Liu, C. Gong *et al.*, "Flow straight and fast: Learning to generate and transfer data with rectified flow," in *The Eleventh International Conference on Learning Representations*.
- [22] X. Li, J. Thickstun, I. Gulrajani, P. S. Liang, and T. B. Hashimoto, "Diffusion-lm improves controllable text generation," *Advances in neural information processing systems*, vol. 35, pp. 4328–4343, 2022.
- [23] R. Strudel, C. Tallec, F. Alché, Y. Du, Y. Ganin, A. Mensch, W. Grathwohl, N. Savinov, S. Dieleman, L. Sifre *et al.*, "Self-conditioned embedding diffusion for text generation," *arXiv preprint arXiv:2211.04236*, 2022.
- [24] J. Austin, D. D. Johnson, J. Ho, D. Tarlow, and R. Van Den Berg, "Structured denoising diffusion models in discrete state-spaces," *Advances in neural information processing systems*, vol. 34, pp. 17 981–17 993, 2021.
- [25] Z. He, T. Sun, Q. Tang, K. Wang, X. Huang, and X. Qiu, "Diffusionbert: Improving generative masked language models with diffusion models," in *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- [26] J. Ye, Z. Xie, L. Zheng, J. Gao, Z. Wu, X. Jiang, Z. Li, and L. Kong, "Dream 7b," 2025. [Online]. Available: <https://hkunlp.github.io/blog/2025/dream>
- [27] S. Gong, S. Agarwal, Y. Zhang, J. Ye, L. Zheng, M. Li, C. An, P. Zhao, W. Bi, J. Han *et al.*, "Scaling diffusion language models via adaptation from autoregressive models," in *The Thirteenth International Conference on Learning Representations*.
- [28] S. Nie, F. Zhu, Z. You, X. Zhang, J. Ou, J. Hu, J. Zhou, Y. Lin, J.-R. Wen, and C. Li, "Large language diffusion models," *arXiv preprint arXiv:2502.09992*, 2025.
- [29] Z. You, S. Nie, X. Zhang, J. Hu, J. Zhou, Z. Lu, J.-R. Wen, and C. Li, "Llada-v: Large language diffusion models with visual instruction tuning," *arXiv preprint arXiv:2505.16933*, 2025.
- [30] R. Yu, X. Ma, and X. Wang, "Dimple: Discrete diffusion multimodal large language model with parallel decoding," *arXiv preprint arXiv:2505.16990*, 2025.
- [31] L. Yang, Y. Tian, B. Li, X. Zhang, K. Shen, Y. Tong, and M. Wang, "Mmada: Multimodal large diffusion language models," *arXiv preprint arXiv:2505.15809*, 2025.
- [32] I. Labs, S. Khanna, S. Kharbanda, S. Li, H. Varma, E. Wang, S. Birnbaum, Z. Luo, Y. Miraoui, A. Palrecha *et al.*, "Mercury: Ultra-fast language models based on diffusion," *arXiv preprint arXiv:2506.17298*, 2025.
- [33] DeepMind, "Gemini diffusion," <https://deepmind.google/technologies/gemini>, 2024, accessed: 2025-07-09.
- [34] Y. Song, Z. Zhang, C. Luo, P. Gao, F. Xia, H. Luo, Z. Li, Y. Yang, H. Yu, X. Qu *et al.*, "Seed diffusion: A large-scale diffusion language model with high-speed inference," *arXiv preprint arXiv:2508.02193*, 2025.
- [35] M. Xu, T. Geffner, K. Kreis, W. Nie, Y. Xu, J. Leskovec, S. Ermon, and A. Vahdat, "Energy-based diffusion language models for text generation," *arXiv preprint arXiv:2410.21357*, 2024.
- [36] J. Deschenaux and C. Gulcehre, "Beyond autoregression: Fast llms via self-distillation through time," in *The Thirteenth International Conference on Learning Representations*.
- [37] K. Han, K. Kenealy, A. Barua, N. Fiedel, and N. Constant, "Transfer learning for text diffusion models," *arXiv preprint arXiv:2401.17181*, 2024.
- [38] S. S. Sahoo, J. Deschenaux, A. Gokaslan, G. Wang, J. Chiu, and V. Kuleshov, "The diffusion duality," *arXiv preprint arXiv:2506.10892*, 2025.
- [39] Y. Zhang, S. He, D. Levine, L. Zhao, D. Zhang, S. A. Rizvi, E. Zappala, R. Ying, and D. van Dijk, "Non-markovian dis-

- crete diffusion with causal language models," *arXiv preprint arXiv:2502.09767*, 2025.
- [40] M. Dang, J. Han, M. Xu, K. Xu, A. Srivastava, and S. Ermon, "Inference-time scaling of diffusion language models with particle gibbs sampling," *arXiv preprint arXiv:2507.08390*, 2025.
- [41] L. Rout, C. Caramanis, and S. Shakkottai, "Anchored diffusion language model," *arXiv preprint arXiv:2505.18456*, 2025.
- [42] Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu *et al.*, "Deepseekmath: Pushing the limits of mathematical reasoning in open language models," *arXiv preprint arXiv:2402.03300*, 2024.
- [43] S. Zhao, D. Gupta, Q. Zheng, and A. Grover, "d1: Scaling reasoning in diffusion large language models via reinforcement learning," *arXiv preprint arXiv:2504.12216*, 2025.
- [44] T. Chen, S. Zhang, and M. Zhou, "Dlm-one: Diffusion language models for one-step sequence generation," *arXiv e-prints*, pp. arXiv-2506, 2025.
- [45] C. Wu, H. Zhang, S. Xue, Z. Liu, S. Diao, L. Zhu, P. Luo, S. Han, and E. Xie, "Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache and parallel decoding," *arXiv preprint arXiv:2505.22618*, 2025.
- [46] D. Israel, G. V. d. Broeck, and A. Grover, "Accelerating diffusion llms via adaptive parallel decoding," *arXiv preprint arXiv:2506.00413*, 2025.
- [47] G. Wang, Y. Schiff, S. S. Sahoo, and V. Kuleshov, "Remasking discrete diffusion models with inference-time scaling," in *ICLR 2025 Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*.
- [48] Z. Liu, Y. Yang, Y. Zhang, J. Chen, C. Zou, Q. Wei, S. Wang, and L. Zhang, "dllm-cache: Accelerating diffusion large language models with adaptive caching," *arXiv preprint arXiv:2506.06295*, 2025.
- [49] X. Ma, R. Yu, G. Fang, and X. Wang, "dkv-cache: The cache for diffusion language models," *arXiv preprint arXiv:2505.15781*, 2025.
- [50] G. Liu, Z. Feng, Y. Gao, Z. Yang, X. Liang, J. Bao, X. He, S. Cui, Z. Li, and Z. Hu, "Composable text controls in latent space with odes," *arXiv preprint arXiv:2208.00638*, 2022.
- [51] S. Gong, M. Li, J. Feng, Z. Wu, and L. Kong, "Diffuseq: Sequence to sequence text generation with diffusion models," in *The Eleventh International Conference on Learning Representations*.
- [52] S. Dieleman, L. Sartran, A. Roshannai, N. Savinov, Y. Ganin, P. H. Richemond, A. Doucet, R. Strudel, C. Dyer, C. Durkan *et al.*, "Continuous diffusion for categorical data," *arXiv preprint arXiv:2211.15089*, 2022.
- [53] Z. Gao, J. Guo, X. Tan, Y. Zhu, F. Zhang, J. Bian, and L. Xu, "Empowering diffusion models on the embedding space for text generation," *arXiv preprint arXiv:2212.09412*, 2022.
- [54] J. Lovelace, V. Kishore, C. Wan, E. Shekhtman, and K. Q. Weinberger, "Latent diffusion for language generation," *Advances in Neural Information Processing Systems*, vol. 36, pp. 56 998–57 025, 2023.
- [55] Z. Lin, Y. Gong, Y. Shen, T. Wu, Z. Fan, C. Lin, N. Duan, and W. Chen, "Text generation with diffusion language models: A pre-training approach with continuous paragraph denoise," in *International Conference on Machine Learning*. PMLR, 2023, pp. 21 051–21 064.
- [56] R. Wang, J. Li, and P. Li, "Infodiffusion: Information entropy aware diffusion process for non-autoregressive text generation," in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 13 757–13 770.
- [57] G. Liu, Y. Wang, Z. Feng, Q. Wu, L. Tang, Y. Gao, Z. Li, S. Cui, J. McAuley, Z. Yang *et al.*, "Unified generation, reconstruction, and representation: Generalized diffusion with adaptive latent encoding-decoding," in *International Conference on Machine Learning*. PMLR, 2024, pp. 31 964–31 993.
- [58] A. Shabalin, V. Meshchaninov, and D. Vetrov, "Smoothie: Smoothing diffusion on token embeddings for text generation," *arXiv preprint arXiv:2505.18853*, 2025.
- [59] R. K. Mahabadi, H. Ivison, J. Tae, J. Henderson, I. Beltagy, M. E. Peters, and A. Cohan, "Tess: Text-to-text self-conditioned simplex diffusion," in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 2347–2361.
- [60] J. Tae, H. Ivison, S. Kumar, and A. Cohan, "Tess 2: A large-scale generalist diffusion language model," *arXiv preprint arXiv:2502.13917*, 2025.
- [61] P. Yu, S. Xie, X. Ma, B. Jia, B. Pang, R. Gao, Y. Zhu, S.-C. Zhu, and Y. N. Wu, "Latent diffusion energy-based model for interpretable text modelling," in *International Conference on Machine Learning*. PMLR, 2022, pp. 25 702–25 720.
- [62] Y. Chen, C. Liang, H. Sui, R. Guo, C. Cheng, J. You, and G. Liu, "Langflow: Continuous diffusion rivals discrete in language modeling," *arXiv preprint arXiv:2604.11748*, 2026.
- [63] K. Hu, L. Qiu, Y. Lu, H. Zhao, T. Li, Y. Kim, J. Andreas, and K. He, "Elf: Embedded language flows," *arXiv preprint arXiv:2605.10938*, 2026.
- [64] S. Zhuang, Y. Ai, J. Han, X. Li, H. Huang, X. Yue, X. Hu, K. Xu, Y. Wang, and H. Chen, "Bitlm: Unlocking multi-token language generation with bitwise continuous diffusion," *arXiv preprint arXiv:2605.11577*, 2026.
- [65] C. Lee, J. Yoo, M. Agarwal, S. Shah, J. Huang, A. Raghunathan, S. Hong, N. M. Boffi, and J. Kim, "Flow map language models: One-step language modeling via continuous denoising," *arXiv preprint arXiv:2602.16813*, 2026.
- [66] H. Guo, Q. Zhao, Y. Zhao, S. Nie, R. Zhu, Q. Guo, F. Wang, T. Yang, H. Zhao, G. Wei *et al.*, "Continuous latent diffusion language model," *arXiv preprint arXiv:2605.06548*, 2026.
- [67] L. Zheng, J. Yuan, L. Yu, and L. Kong, "A reparameterized discrete diffusion model for text generation," in *First Conference on Language Modeling*.
- [68] J. Shi, K. Han, Z. Wang, A. Doucet, and M. Titsias, "Simplified and generalized masked diffusion for discrete data," *Advances in neural information processing systems*, vol. 37, pp. 103 131–103 167, 2024.
- [69] S. Sahoo, M. Arriola, Y. Schiff, A. Gokaslan, E. Marroquin, J. Chiu, A. Rush, and V. Kuleshov, "Simple and effective masked diffusion language models," *Advances in Neural Information Processing Systems*, vol. 37, pp. 130 136–130 184, 2024.
- [70] J. Ye, Z. Zheng, Y. Bao, L. Qian, and Q. Gu, "Diffusion language models can perform many tasks with scaling and instruction-finetuning," *arXiv preprint arXiv:2308.12219*, 2023.
- [71] K. Zhou, Y. Li, W. X. Zhao, and J.-R. Wen, "Diffusion-nat: Self-prompting discrete diffusion for non-autoregressive text generation," in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 1438–1451.
- [72] I. Gulrajani and T. B. Hashimoto, "Likelihood-based diffusion language models," *Advances in Neural Information Processing Systems*, vol. 36, pp. 16 693–16 715, 2023.
- [73] A. Lou, C. Meng, and S. Ermon, "Discrete diffusion modeling by estimating the ratios of the data distribution," in *International Conference on Machine Learning*. PMLR, 2024, pp. 32 819–32 848.
- [74] J. Ou, S. Nie, K. Xue, F. Zhu, J. Sun, Z. Li, and C. Li, "Your absorbing discrete diffusion secretly models the conditional distributions of clean data," in *The Thirteenth International Conference on Learning Representations*, 2024.
- [75] I. Gat, T. Remez, N. Shaul, F. Kreuk, R. T. Chen, G. Synnaeve, Y. Adi, and Y. Lipman, "Discrete flow matching," *Advances in Neural Information Processing Systems*, vol. 37, pp. 133 345–133 385, 2024.
- [76] S. Liu, J. Nam, A. Campbell, H. Stark, Y. Xu, T. Jaakkola, and R. Gomez-Bombarelli, "Think while you generate: Discrete diffusion with planned denoising," in *The Thirteenth International Conference on Learning Representations*, 2024.
- [77] J. Ye, J. Gao, S. Gong, L. Zheng, X. Jiang, Z. Li, and L. Kong, "Beyond autoregression: Discrete diffusion for complex reasoning and planning," *arXiv preprint arXiv:2410.14157*, 2024.
- [78] D. von Rütte, J. Fluri, Y. Ding, A. Orvieto, B. Schölkopf, and T. Hofmann, "Generalized interpolating discrete diffusion," in *Forty-second International Conference on Machine Learning*, 2025.
- [79] X. Liu, Z. Liu, Z. Huang, Q. Guo, Z. He, and X. Qiu, "Longllada: Unlocking long context capabilities in diffusion llms," *arXiv preprint arXiv:2506.14429*, 2025.
- [80] F. Zhu, Z. You, Y. Xing, Z. Huang, L. Liu, Y. Zhuang, G. Lu, K. Wang, X. Wang, L. Wei *et al.*, "Llada-moe: A sparse moe diffusion language model," *arXiv preprint arXiv:2509.24389*, 2025.
- [81] T. Bie, M. Cao, K. Chen, L. Du, M. Gong, Z. Gong, Y. Gu, J. Hu, Z. Huang, Z. Lan *et al.*, "Llada2. 0: Scaling up diffusion language models to 100b," *arXiv preprint arXiv:2512.15745*, 2025.
- [82] X. Han, S. Kumar, and Y. Tsvetkov, "Ssd-llm: Semi-autoregressive simplex-based diffusion language model for text generation and modular control," in *Proceedings of the 61st Annual Meeting of the*

- Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 11 575–11 596.
- [83] T. Wu, Z. Fan, X. Liu, H.-T. Zheng, Y. Gong, J. Jiao, J. Li, J. Guo, N. Duan, W. Chen *et al.*, “Ar-diffusion: Auto-regressive diffusion model for text generation,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 39 957–39 974, 2023.
- [84] M. Arriola, A. Gokaslan, J. T. Chiu, Z. Yang, Z. Qi, J. Han, S. S. Sahoo, and V. Kuleshov, “Block diffusion: Interpolating between autoregressive and diffusion language models,” in *The Thirteenth International Conference on Learning Representations*.
- [85] C. Huang and H. Tang, “CtrlDiff: Boosting large diffusion language models with dynamic block prediction and controllable generation,” *arXiv preprint arXiv:2505.14455*, 2025.
- [86] J. K. Christopher, B. R. Bartoldson, T. Ben-Nun, M. Cardei, B. Kailkhura, and F. Fioretto, “Speculative diffusion decoding: Accelerating language generation through diffusion,” in *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2025, pp. 12 042–12 059.
- [87] S. Cheng, Y. Bian, D. Liu, L. Zhang, Q. Yao, Z. Tian, W. Wang, Q. Guo, K. Chen, B. Qi *et al.*, “Sdar: A synergistic diffusion-autoregression paradigm for scalable sequence generation,” *arXiv preprint arXiv:2510.06303*, 2025.
- [88] J. Liu, X. Dong, Z. Ye, R. Mehta, Y. Fu, V. Singh, J. Kautz, C. Zhang, and P. Molchanov, “Tidar: Think in diffusion, talk in autoregression,” *arXiv preprint arXiv:2511.08923*, 2025.
- [89] Y. Liu, Y. Cao, H. Li, G. Luo, Z. Chen, W. Wang, X. Liang, B. Qi, L. Wu, C. Tian *et al.*, “Sequential diffusion language models,” *arXiv preprint arXiv:2509.24007*, 2025.
- [90] Y. Tian, Y. Liang, S. Zhang, Y. Shu, G. Yang, W. He, S. Fang, T. Guo, K. Han, C. Xu *et al.*, “From next-token to next-block: A principled adaptation path for diffusion llms,” *arXiv preprint arXiv:2512.06776*, 2025.
- [91] Y. Fu, L. Whalen, Z. Ye, X. Dong, S. Diao, J. Liu, C. Wu, H. Zhang, E. Xie, S. Han *et al.*, “Efficient-dlm: From autoregressive to diffusion language models, and beyond in speed,” *arXiv preprint arXiv:2512.14067*, 2025.
- [92] Y. Yu, Y. Jian, J. Wang, Z. Zhou, D. Zhuang, X. Fang, S. Yanamandra, X. Wu, Q. Wu, S. L. Song *et al.*, “Introspective diffusion language models,” *arXiv preprint arXiv:2604.11035*, 2026.
- [93] J.-N. Li, J. Guan, W. Wu, and C. Li, “Refusion: A diffusion large language model with parallel autoregressive decoding,” *arXiv preprint arXiv:2512.13586*, 2025.
- [94] J. Ruan, B. Li, Y. Yin, P. Huang, X. Chen, J. Wang, X. Cai, T. Xiao, and J. Zhu, “Causal autoregressive diffusion language model,” *arXiv preprint arXiv:2601.22031*, 2026.
- [95] Z. Li, H. Li, Y. Shi, A. B. Farimani, Y. Kluger, L. Yang, and P. Wang, “Dual diffusion for unified image generation and understanding,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 2779–2790.
- [96] Q. Shi, J. Bai, Z. Zhao, W. Chai, K. Yu, J. Wu, S. Song, Y. Tong, X. Li, X. Li *et al.*, “Muddit: Liberating generation beyond text-to-image with a unified discrete diffusion model,” *arXiv preprint arXiv:2505.23606*, 2025.
- [97] J. Ye, S. Gong, L. Chen, L. Zheng, J. Gao, H. Shi, C. Wu, X. Jiang, Z. Li, W. Bi *et al.*, “Diffusion of thought: Chain-of-thought reasoning in diffusion language models,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 105 345–105 374, 2024.
- [98] Z. Huang, Z. Chen, Z. Wang, T. Li, and G.-J. Qi, “Reinforcing the diffusion chain of lateral thought with diffusion language models,” *arXiv preprint arXiv:2505.10446*, 2025.
- [99] O. Zekri and N. Boullé, “Fine-tuning discrete diffusion models with policy gradient methods,” *arXiv preprint arXiv:2502.01384*, 2025.
- [100] S. Gong, R. Zhang, H. Zheng, J. Gu, N. Jaitly, L. Kong, and Y. Zhang, “Diffucoder: Understanding and improving masked diffusion models for code generation,” *arXiv preprint arXiv:2506.20639*, 2025.
- [101] X. Tang, R. Dolga, S. Yoon, and I. Bogunovic, “wd1: Weighted policy optimization for reasoning in diffusion language models,” *arXiv preprint arXiv:2507.08838*, 2025.
- [102] S. Zhao, M. Liu, J. Huang, M. Liu, C. Wang, B. Liu, Y. Tian, G. Pang, S. Bell, A. Grover *et al.*, “Inpainting-guided policy optimization for diffusion large language models,” *arXiv preprint arXiv:2509.10396*, 2025.
- [103] C. Wang, P. Rashidinejad, D. Su, S. Jiang, S. Wang, S. Zhao, C. Zhou, S. Z. Shen, F. Chen, T. Jaakkola *et al.*, “Spg: Sandwiched policy gradient for masked diffusion language models,” *arXiv preprint arXiv:2510.09541*, 2025.
- [104] S. Xie, L. Kong, X. Song, X. Dong, G. Chen, E. P. Xing, and K. Zhang, “Step-aware policy optimization for reasoning in diffusion large language models,” *arXiv preprint arXiv:2510.01544*, 2025.
- [105] N. Lin, J. Zhang, L. Hou, and J. Li, “Boundary-guided policy optimization for memory-efficient rl of diffusion large language models,” *arXiv preprint arXiv:2510.11683*, 2025.
- [106] Z. Ni, S. Wang, Y. Yue, T. Yu, W. Zhao, Y. Hua, T. Chen, J. Song, C. Yu, B. Zheng *et al.*, “The flexibility trap: Why arbitrary order limits reasoning potential in diffusion language models,” *arXiv preprint arXiv:2601.15165*, 2026.
- [107] F. Zhu, R. Wang, S. Nie, X. Zhang, C. Wu, J. Hu, J. Zhou, J. Chen, Y. Lin, J.-R. Wen *et al.*, “Llada 1.5: Variance-reduced preference optimization for large language diffusion models,” *arXiv preprint arXiv:2505.19223*, 2025.
- [108] Q. Wei, Y. Zhang, Z. Liu, D. Liu, and L. Zhang, “Accelerating diffusion large language models with slowfast: The three golden principles,” *arXiv preprint arXiv:2506.10848*, 2025.
- [109] W. Bao, Z. Chen, D. Xu, and Y. Shang, “Learning to parallel: Accelerating diffusion large language models via learnable parallel decoding,” *arXiv preprint arXiv:2509.25188*, 2025.
- [110] Z. Chen, G. Fang, X. Ma, R. Yu, and X. Wang, “dparallel: Learnable parallel decoding for dllms,” *arXiv preprint arXiv:2509.26488*, 2025.
- [111] J. Chen, Y. Liang, and Z. Liu, “Dflash: Block diffusion for flash speculative decoding,” *arXiv preprint arXiv:2602.06036*, 2026.
- [112] Z. Chen, G. Fang, X. Ma, R. Yu, and X. Wang, “Dmax: Aggressive parallel decoding for dllms,” *arXiv preprint arXiv:2604.08302*, 2026.
- [113] P. Li, D. Muhtar, T. Chen, L. Yin, and S. Liu, “Why diffusion language models struggle with truly parallel (non-autoregressive) decoding?” *arXiv preprint arXiv:2602.23225*, 2026.
- [114] P. Li, S. Yan, J. Tsai, R. Zhang, R. An, Z. Guo, and X. Gao, “Adaptive classifier-free guidance via dynamic low-confidence masking,” *arXiv preprint arXiv:2505.20199*, 2025.
- [115] Z. Hu, J. Meng, Y. Akhauri, M. S. Abdelfattah, J.-s. Seo, Z. Zhang, and U. Gupta, “Accelerating diffusion language model inference via efficient kv caching and guided diffusion,” *arXiv preprint arXiv:2505.21467*, 2025.
- [116] T. Suresh, D. Banerjee, S. Ugare, S. Misailovic, and G. Singh, “Dingo: Constrained inference for diffusion llms,” in *ICML 2025 Workshop on Reliable and Responsible Foundation Models*.
- [117] Q. Nguyen-Tri, M. Ranjan, and Z. Shen, “Attention is all you need for kv cache in diffusion llms,” *arXiv preprint arXiv:2510.14973*, 2025.
- [118] Y. Jiang, Y. Cai, X. Luo, J. Fu, J. Wang, C. Liu, and X. Yang, “d² cache: Accelerating diffusion-based llms via dual adaptive caching,” *arXiv preprint arXiv:2509.23094*, 2025.
- [119] X. Ma, G. Fang, and X. Wang, “Deepcache: Accelerating diffusion models for free,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 15 762–15 772.
- [120] P. Chen, M. Shen, P. Ye, J. Cao, C. Tu, C.-S. Bouganis, Y. Zhao, and T. Chen, “ Δ -dit: A training-free acceleration method tailored for diffusion transformers,” *arXiv preprint arXiv:2406.01125*, 2024.
- [121] X. Ma, G. Fang, M. Bi Mi, and X. Wang, “Learning-to-cache: Accelerating diffusion transformer via layer caching,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 133 282–133 304, 2024.
- [122] Z. Lv, C. Si, J. Song, Z. Yang, Y. Qiao, Z. Liu, and K.-Y. K. Wong, “Fastercache: Training-free video diffusion model acceleration with high quality,” in *The Thirteenth International Conference on Learning Representations*.
- [123] S. Hayakawa, Y. Takida, M. Imaizumi, H. Wakaki, and Y. Mitsufuji, “Distillation of discrete diffusion through dimensional correlations,” in *Forty-second International Conference on Machine Learning*.
- [124] T. Salimans and J. Ho, “Progressive distillation for fast sampling of diffusion models,” in *International Conference on Learning Representations*.
- [125] Y.-Y. Qian, J. Su, L. Hu, P. Zhang, Z. Deng, P. Zhao, and H. Zhang, “d3llm: Ultra-fast diffusion llm using pseudo-trajectory distillation,” *arXiv preprint arXiv:2601.07568*, 2026.
- [126] A. Myrzakhan, T. Li, B. Guo, S. Tang, and Z. Shen, “Sink-aware pruning for diffusion language models,” *arXiv preprint arXiv:2602.17664*, 2026.

- [127] S. Li, K. Kallidromitis, H. Bansal, A. Gokul, Y. Kato, K. Kozuka, J. Kuen, Z. Lin, K.-W. Chang, and A. Grover, "Lavida: A large diffusion language model for multimodal understanding," *arXiv preprint arXiv:2505.16839*, 2025.
- [128] J. Wang, Y. Lai, A. Li, S. Zhang, J. Sun, N. Kang, C. Wu, Z. Li, and P. Luo, "Fudoki: Discrete flow-based unified understanding and generation via kinetic-optimal velocities," *arXiv preprint arXiv:2505.20147*, 2025.
- [129] A. Swerdlow, M. Prabhudesai, S. Gandhi, D. Pathak, and K. Fragkiadaki, "Unified multimodal discrete diffusion," *arXiv preprint arXiv:2503.20853*, 2025.
- [130] Y. Xin, Q. Qin, S. Luo, K. Zhu, J. Yan, Y. Tai, J. Lei, Y. Cao, K. Wang, Y. Wang *et al.*, "Lumina-dimoo: An omni diffusion large language model for multi-modal generation and understanding," *arXiv preprint arXiv:2510.06308*, 2025.
- [131] S. Li, J. Gu, K. Liu, Z. Lin, Z. Wei, A. Grover, and J. Kuen, "Lavida-o: Elastic large masked diffusion models for unified multimodal understanding and generation," *arXiv preprint arXiv:2509.19244*, 2025.
- [132] Y. Tian, L. Yang, J. Yang, A. Wang, Y. Tian, J. Zheng, H. Wang, Z. Teng, Z. Wang, Y. Wang *et al.*, "Mmada-parallel: Multimodal large diffusion language models for thinking-aware editing and generation," *arXiv preprint arXiv:2511.09611*, 2025.
- [133] L. Zeng, J. Yao, B. Liao, H. Tao, W. Liu, and X. Wang, "Diffusionvl: Translating any autoregressive models into diffusion vision language models," *arXiv preprint arXiv:2512.15713*, 2025.
- [134] C. Wu, S. Lan, Y. Fu, S. Gao, J. Wang, J. Yu, J. M. Alvarez, P. Molchanov, P. Luo, S. Han *et al.*, "Fast-dvllm: Efficient block-diffusion vlm via direct conversion from autoregressive vlm," *arXiv preprint arXiv:2604.06832*, 2026.
- [135] Z. He, T. Chen, K. Wang, Z. Qin, Y. Shao, C. Gan, S. Li, Z. Wu, and W. Lin, "Vidlada: Bidirectional diffusion large language models for efficient video understanding," *arXiv preprint arXiv:2601.17868*, 2026.
- [136] S. Yuan, W. Yuan, H. Yin, and T. He, "Roic-dm: Robust text inference and classification via diffusion model," *arXiv preprint arXiv:2401.03514*, 2024.
- [137] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang, "Diffusionner: Boundary diffusion for named entity recognition," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 3875–3890.
- [138] X. Yang, Z. Qiao, and Y. Zhou, "Ipad: Iterative, parallel, and diffusion-based network for scene text recognition," *International Journal of Computer Vision*, pp. 1–21, 2025.
- [139] S. Liu, J. Zhou, Q. Zhu, Q. Chen, Q. Bai, J. Xiao, and L. He, "Let's rectify step by step: Improving aspect-based sentiment analysis with diffusion models," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024, pp. 10324–10335.
- [140] H. Zhang, X. Liu, and J. Zhang, "Diffusum: Generation enhanced extractive summarization with diffusion," in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, pp. 13089–13100.
- [141] X. Dong, W. Li, Y. Le, Z. Jiang, J. Zhong, and Z. Wang, "Termdiffusum: a term-guided diffusion model for extractive summarization of legal documents," in *Proceedings of the 31st international conference on computational linguistics*, 2025, pp. 3222–3235.
- [142] Y. Luo, Q. Zhou, and F. Zhou, "Enhancing phrase representation by information bottleneck guided text diffusion process for keyphrase extraction," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024, pp. 6036–6047.
- [143] J. Zhao, C. Xu, and B. Jiang, "Iped: An implicit perspective for relational triple extraction based on diffusion model," in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024, pp. 2080–2092.
- [144] C. H. Lee, H. Kim, J. Yeom, and S. Yoon, "Editext: Controllable coarse-to-fine text editing with diffusion language models," *arXiv preprint arXiv:2502.19765*, 2025.
- [145] G. Bi, L. Shen, Y. Cao, M. Chen, Y. Xie, Z. Lin, and X. He, "Diffusemp: A diffusion model-based framework with multi-grained control for empathetic response generation," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 2812–2831.
- [146] G. Floto, M. M. A. Pour, P. Farinneya, Z. Tang, A. Pesaraghader, M. Bharadwaj, and S. Sanner, "Diffudetox: A mixed diffusion model for text detoxification," in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, pp. 7566–7574.
- [147] Z. Horvitz, A. Patel, C. Callison-Burch, Z. Yu, and K. McKeown, "Paraguide: Guided diffusion paraphrasers for plug-and-play textual style transfer," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 38, no. 16, 2024, pp. 18216–18224.
- [148] Y. Zhang, J. Gu, Z. Wu, S. Zhai, J. Susskind, and N. Jaitly, "Planner: Generating diversified paragraph via latent language diffusion model," *Advances in Neural Information Processing Systems*, vol. 36, pp. 80178–80190, 2023.
- [149] J. Liu, P. Cheng, J. Dai, and J. Liu, "Diffucom: A novel diffusion model for comment generation," *Knowledge-Based Systems*, vol. 281, p. 111069, 2023.
- [150] J. Xiang, Z. Liu, H. Liu, Y. Bai, J. Cheng, and W. Chen, "Diffusiondialog: A diffusion model for diverse dialog generation with latent space," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024, pp. 4912–4921.
- [151] W. Zou, Z. Zhuang, X. Geng, S. Huang, J. Liu, and J. Chen, "Improved paraphrase generation via controllable latent diffusion," *arXiv preprint arXiv:2404.08938*, 2024.
- [152] Z. Hu, C. Liu, Y. Feng, A. T. Luu, and B. Hooi, "Poetrydiffusion: Towards joint semantic and metrical manipulation in poetry generation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 18279–18288.
- [153] L. Chen, A. Feng, B. Yang, and Z. Li, "Xdlm: Cross-lingual diffusion language model for machine translation," *arXiv preprint arXiv:2307.13560*, 2023.
- [154] S. Qiao, X. Liu, and S.-H. Na, "Diffusionret: Diffusion-enhanced generative retriever using constrained decoding," in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 9515–9529.
- [155] K. Yan, M. Liu, Y. Liu, R. Fu, Z. Wen, J. Tao, and X. Liu, "Debunk and infer: Multimodal fake news detection via diffusion-generated evidence and llm reasoning," *arXiv preprint arXiv:2506.21557*, 2025.
- [156] O. Luxembourg, H. Permuter, and E. Nachmani, "Plan for speed-dilated scheduling for masked diffusion language models," *arXiv preprint arXiv:2506.19037*, 2025.
- [157] C. Fan, W. Heng, B. Li, S. Liu, Y. Song, J. Su, X. Qu, K. Shen, and W. Wei, "Stable-diffcoder: Pushing the frontier of code diffusion large language model," *arXiv preprint arXiv:2601.15892*, 2026.
- [158] H. Bai, L. Kong, X. Chen, J. Wang, Z. Tao, and H. Wang, "Dice: Diffusion large language models excel at generating cuda kernels," *arXiv preprint arXiv:2602.11715*, 2026.
- [159] Y. Xiong, K. Li, J. Chen, H. Zhang, D. Lin, Y. Che, and W. Hu, "Text-guided multi-property molecular optimization with a diffusion language model," *arXiv preprint arXiv:2410.13597*, 2024.
- [160] H. Gong, Q. Liu, S. Wu, and L. Wang, "Text-guided molecule generation with diffusion language model," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 1, 2024, pp. 109–117.
- [161] S. Goel, V. Thoutam, E. M. Marroquin, A. Gokaslan, A. Firouzbakht, S. Vincoff, V. Kuleshov, H. T. Kratochvil, and P. Chatterjee, "Memdlm: De novo membrane protein design with masked discrete diffusion protein language models," in *NeurIPS 2024 Workshop on AI for New Drug Modalities*.
- [162] X. Wang, Z. Zheng, D. Xue, S. Huang, Q. Gu *et al.*, "Diffusion language models are versatile protein learners," in *Forty-first International Conference on Machine Learning*.
- [163] J. Yin, C. Zha, W. He, C. Xu, and X. Gao, "Cfp-gen: Combinatorial functional protein generation via diffusion language models," in *Forty-second International Conference on Machine Learning*.
- [164] C. Wang, M. Uehara, Y. He, A. Wang, A. Lal, T. Jaakkola, S. Levine, A. Regev, T. Biancalani *et al.*, "Fine-tuning discrete diffusion models via reward optimization with applications to dna and protein design," in *The Thirteenth International Conference on Learning Representations*.
- [165] B. Ni, D. L. Kaplan, and M. J. Buehler, "Forcegen: End-to-end de novo protein generation based on nonlinear mechanical unfolding responses using a language diffusion model," *Science Advances*, vol. 10, no. 6, p. ead14000, 2024.
- [166] L. Hallee, N. Rafailidis, D. B. Bichara, and J. P. Gleghorn, "Diffusion sequence models for enhanced protein representation and generation," *arXiv preprint arXiv:2506.08293*, 2025.

- [167] X. Wang, Z. Zheng, F. Ye, D. Xue, S. Huang, and Q. Gu, "Dplm-2: A multimodal diffusion protein language model," *arXiv preprint arXiv:2410.13782*, 2024.
- [168] Y. Wen, H. Li, K. Gu, Y. Zhao, T. Wang, and X. Sun, "Llada-*vla*: Vision language diffusion action models," *arXiv preprint arXiv:2509.06932*, 2025.
- [169] J. Wen, M. Zhu, J. Liu, Z. Liu, Y. Yang, L. Zhang, S. Zhang, Y. Zhu, and Y. Xu, "dvla: Diffusion vision-language-action model with multimodal chain-of-thought," *arXiv preprint arXiv:2509.25681*, 2025.
- [170] J. Chen, W. Song, P. Ding, Z. Zhou, H. Zhao, F. Tang, D. Wang, and H. Li, "Unified diffusion *vla*: Vision-language-action model via joint discrete denoising diffusion process," *arXiv preprint arXiv:2511.01718*, 2025.
- [171] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International conference on machine learning*. pmlr, 2015, pp. 2256–2265.
- [172] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 2019, pp. 4171–4186.
- [173] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [174] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "Albert: A lite bert for self-supervised learning of language representations," in *International Conference on Learning Representations*.
- [175] P. He, X. Liu, J. Gao, and W. Chen, "Deberta: Decoding-enhanced bert with disentangled attention," in *International Conference on Learning Representations*.
- [176] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," *arXiv preprint arXiv:1901.02860*, 2019.
- [177] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin *et al.*, "Opt: Open pre-trained transformer language models," *arXiv preprint arXiv:2205.01068*, 2022.
- [178] F. Gloeckle, B. Y. Idrissi, B. Roziere, D. Lopez-Paz, and G. Synnaeve, "Better & faster large language models via multi-token prediction," in *Forty-first International Conference on Machine Learning*.
- [179] R. Chen, W. Chai, Z. Yang, X. Zhang, J. T. Zhou, T. Quek, S. Poria, and Z. Liu, "Diffpo: Diffusion-styled preference optimization for efficient inference-time alignment of large language models," *arXiv preprint arXiv:2503.04240*, 2025.
- [180] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," *Advances in neural information processing systems*, vol. 27, 2014.
- [181] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of machine learning research*, vol. 21, no. 140, pp. 1–67, 2020.
- [182] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7871–7880.
- [183] H. Yuan, Z. Yuan, C. Tan, F. Huang, and S. Huang, "Seqdiffuseq: Text diffusion with encoder-decoder transformers," *arXiv preprint arXiv:2212.10325*, 2022.
- [184] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," *Advances in neural information processing systems*, vol. 32, 2019.
- [185] T. Chen, R. Zhang, and G. Hinton, "Analog bits: Generating discrete data using diffusion models with self-conditioning," *arXiv preprint arXiv:2208.04202*, 2022.
- [186] Q. Team, "Qwen2.5: A party of foundation models," September 2024. [Online]. Available: <https://qwenlm.github.io/blog/qwen2.5/>
- [187] G. He, S. Nie, F. Zhu, Y. Zhao, T. Bai, R. Yan, J. Fu, C. Li, and B. Yuan, "Ultrallada: Scaling the context length to 128k for diffusion large language models," *arXiv preprint arXiv:2510.10481*, 2025.
- [188] X. Zhu, G. Karadzhov, C. Whitehouse, and A. Vlachos, "Segment-level diffusion: A framework for controllable long-form generation with diffusion language models," *arXiv preprint arXiv:2412.11333*, 2024.
- [189] Y. Zihuiwen, Y. Elle Michelle, and B. Phil, "Latent diffusion for document generation with sequential decoding," in *NeurIPS 2023 Workshop on Diffusion Models*, 2023. [Online]. Available: <https://neurips.cc/virtual/2023/74876>
- [190] E. Cetin, T. Zhao, and Y. Tang, "Large language models to diffusion finetuning," *arXiv preprint arXiv:2501.15781*, 2025.
- [191] J. Bai, T. Ye, W. Chow, E. Song, Q.-G. Chen, X. Li, Z. Dong, L. Zhu, and S. Yan, "Meissonic: Revitalizing masked generative transformers for efficient high-resolution text-to-image synthesis," in *The Thirteenth International Conference on Learning Representations*, 2024.
- [192] J. Ni, Q. Liu, C. Du, L. Dou, H. Yan, Z. Wang, T. Pang, and M. Q. Shieh, "Training optimal large diffusion language models," *arXiv preprint arXiv:2510.03280*, 2025.
- [193] J. Ni, Q. Liu, L. Dou, C. Du, Z. Wang, H. Yan, T. Pang, and M. Q. Shieh, "Diffusion language models are super data learners," *arXiv preprint arXiv:2511.03276*, 2025.
- [194] M. Asada and M. Miwa, "Addressing the training-inference discrepancy in discrete diffusion for text generation," in *Proceedings of the 31st International Conference on Computational Linguistics*, 2025, pp. 7156–7164.
- [195] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [196] A. Sauer, D. Lorenz, A. Blattmann, and R. Rombach, "Adversarial diffusion distillation," in *European Conference on Computer Vision*. Springer, 2024, pp. 87–103.
- [197] A. Sauer, F. Boesel, T. Dockhorn, A. Blattmann, P. Esser, and R. Rombach, "Fast high-resolution image synthesis with latent adversarial diffusion distillation," in *SIGGRAPH Asia 2024 Conference Papers*, 2024, pp. 1–11.
- [198] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *Advances in neural information processing systems*, vol. 36, pp. 34 892–34 916, 2023.
- [199] F. Li, R. Zhang, H. Zhang, Y. Zhang, B. Li, W. Li, Z. Ma, and C. Li, "Llava-interleave: Tackling multi-image, video, and 3d in large multimodal models," in *The Thirteenth International Conference on Learning Representations*.
- [200] J. Guo, T. Zheng, Y. Bai, B. Li, Y. Wang, K. Zhu, Y. Li, G. Neubig, W. Chen, and X. Yue, "Mammoth-vl: Eliciting multimodal reasoning with instruction tuning at scale," *arXiv preprint arXiv:2412.05237*, 2024.
- [201] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.
- [202] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge *et al.*, "Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution," *arXiv preprint arXiv:2409.12191*, 2024.
- [203] J. Xie, W. Mao, Z. Bai, D. J. Zhang, W. Wang, K. Q. Lin, Y. Gu, Z. Chen, Z. Yang, and M. Z. Shou, "Show-o: One single transformer to unify multimodal understanding and generation," in *The Thirteenth International Conference on Learning Representations*.
- [204] S. Kou, J. Jin, Z. Liu, C. Liu, Y. Ma, J. Jia, Q. Chen, P. Jiang, and Z. Deng, "Orthus: Autoregressive interleaved image-text generation with modality-specific heads," *arXiv preprint arXiv:2412.00127*, 2024.
- [205] C. Wu, X. Chen, Z. Wu, Y. Ma, X. Liu, Z. Pan, W. Liu, Z. Xie, X. Yu, C. Ruan *et al.*, "Janus: Decoupling visual encoding for unified multimodal understanding and generation," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 12 966–12 977.
- [206] S. Patil, W. Berman, R. Rombach, and P. von Platen, "amused: An open muse reproduction," *arXiv preprint arXiv:2401.01808*, 2024.
- [207] Y. Wang, Z. Li, Y. Zang, Y. Zhou, J. Bu, C. Wang, Q. Lu, C. Jin, and J. Wang, "Pref-grpo: Pairwise preference reward-based grpo for stable text-to-image reinforcement learning," *arXiv preprint arXiv:2508.20751*, 2025.

- [208] Y. Bisk, R. Zellers, J. Gao, Y. Choi *et al.*, "Piqa: Reasoning about physical commonsense in natural language," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 05, 2020, pp. 7432–7439.
- [209] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi, "Hellaswag: Can a machine really finish your sentence?" in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4791–4800.
- [210] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [211] D. Ghosh, H. Hajishirzi, and L. Schmidt, "Geneval: An object-focused framework for evaluating text-to-image alignment," *Advances in Neural Information Processing Systems*, vol. 36, pp. 52 132–52 152, 2023.
- [212] C. Fu, P. Chen, Y. Shen, Y. Qin, M. Zhang, X. Lin, J. Yang, X. Zheng, K. Li, X. Sun *et al.*, "Mme: A comprehensive evaluation benchmark for multimodal large language models," *arXiv preprint arXiv:2306.13394*, 2023.
- [213] X. Yue, Y. Ni, K. Zhang, T. Zheng, R. Liu, G. Zhang, S. Stevens, D. Jiang, W. Ren, Y. Sun *et al.*, "Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 9556–9567.
- [214] D. A. Hudson and C. D. Manning, "Gqa: A new dataset for real-world visual reasoning and compositional question answering," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 6700–6709.
- [215] K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano *et al.*, "Training verifiers to solve math word problems," *arXiv preprint arXiv:2110.14168*, 2021.
- [216] Q. Team, "Qwen2 technical report," *arXiv preprint arXiv:2407.10671*, 2024.
- [217] D. Rein, B. L. Hou, A. C. Stickland, J. Petty, R. Y. Pang, J. Dirani, J. Michael, and S. R. Bowman, "Gpqa: A graduate-level google-proof q&a benchmark," in *First Conference on Language Modeling*, 2024.
- [218] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt, "Measuring mathematical problem solving with the math dataset," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- [219] Y. Lyu, T. Luo, J. Shi, T. C. Hollon, and H. Lee, "Fine-grained text style transfer with diffusion-based language models," *arXiv preprint arXiv:2305.19512*, 2023.
- [220] Y. Demirag, D. Liu, and J. Niehues, "Benchmarking diffusion models for machine translation," in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, 2024, pp. 313–324.
- [221] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, 2020, pp. 38–45.
- [222] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. E. Gonzalez, H. Zhang, and I. Stoica, "Efficient memory management for large language model serving with pagedattention," in *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- [223] Z. Wu, L. Zheng, Z. Xie, J. Ye, J. Gao, Y. Feng, Z. Li, V. W., G. Zhou, and L. Kong, "Dreamon: Diffusion language models for code infilling beyond fixed-size canvas," 2025. [Online]. Available: <https://hkunlp.github.io/blog/2025/dreamon>
- [224] Y. Yang, C. Wang, S. Wang, Z. Wen, B. Qi, H. Xu, and L. Zhang, "Diffusion llm with native variable generation lengths: Let [eos] lead the way," *arXiv preprint arXiv:2510.24605*, 2025.
- [225] J. Li, X. Dong, Y. Zang, Y. Cao, J. Wang, and D. Lin, "Beyond fixed: Training-free variable-length denoising for diffusion large language models," *arXiv preprint arXiv:2508.00819*, 2025.
- [226] X. Chen, S. Huang, C. Guo, C. Wei, Y. He, J. Zhang, H. Li, Y. Chen *et al.*, "Dpad: Efficient diffusion language models with suffix dropout," *arXiv preprint arXiv:2508.14148*, 2025.
- [227] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv *et al.*, "Qwen3 technical report," *arXiv preprint arXiv:2505.09388*, 2025.
- [228] K. Team, Y. Bai, Y. Bao, G. Chen, J. Chen, N. Chen, R. Chen, Y. Chen, Y. Chen, Y. Chen *et al.*, "Kimi k2: Open agentic intelligence," *arXiv preprint arXiv:2507.20534*, 2025.