

# Speculative Refinement: A Hybrid Autoregressive Diffusion Decoding Strategy and Its Behavior Across Benchmarks

**Aditi Gupta**

IIIT Hyderabad, India

aditi.gu@research.iiit.ac.in

**Neel Mishra**

Microsoft

mishraneel99@gmail.com

**Kushagra Trivedi**

IIIT Hyderabad, India

kushagra.trivedi@students.iiit.ac.in

**Pawan Kumar**

IIIT Hyderabad, India

pawan.kumar@iiit.ac.in

## Abstract

How should we evaluate generation systems that combine autoregressive (AR) and diffusion decoding? We study this question through *Speculative Refinement* (SpecRef), a training-free hybrid method that warm-starts a masked diffusion language model from an AR draft using entropy-guided selective masking. Evaluating SpecRef across six benchmarks (HumanEval, MBPP, GSM8K, BBH, ARC-Challenge, HellaSwag) with three distinct evaluation protocols (execution-based pass@1, exact-match, log-likelihood scoring), we surface several findings relevant beyond our specific system: (1) code benchmarks conflate *structural discovery* with *logical correctness*: providing a syntactic scaffold lifts accuracy from near zero to over 20% without changing the model, indicating that much of the baseline failure is structural; (2) a *refinement tension* phenomenon where multi-stage correction degrades already-correct tokens, exposing benchmark saturation ceilings invisible to single-model evaluation; (3) log-likelihood and generative evaluation produce different model rankings for the same model pair, suggesting they measure different capabilities; (4) standard Python post-processing silently breaks code evaluation for non-AR generators. These observations apply to any multi-stage or non-autoregressive generation pipeline and point toward more diagnostic evaluation practices. Codes are [here](#).

## 1 Introduction

An 8-billion-parameter diffusion language model scores 0% on HumanEval when given 16 denoising steps. Hand it the same problem with a syntactic scaffold from a 2.7B autoregressive drafter, and it scores over 20% without any retraining. The model did not suddenly learn to code; the benchmark was testing whether it could discover Python’s indentation rules from scratch, and that turned out to be the hard part. This is one of several findings that emerged when we tried to evaluate a hybrid

generation system using standard benchmarks and discovered that the benchmarks themselves were not designed for this.

Diffusion language models (dLMs) have moved rapidly from research curiosity to a competitive generation paradigm. LLaDA (Nie et al., 2025) demonstrated that masked diffusion scales to 8B parameters and matches LLaMA-3 on in-context learning. MDLM (Sahoo et al., 2024) showed that simple masked diffusion with modern training recipes closes much of the perplexity gap with AR models. Dream 7B (Ye et al., 2025) further narrowed this gap while retaining diffusion-native capabilities like parallel generation. On the commercial side, Inception Labs’ Mercury (Labs et al., 2025) achieved over 1,000 tokens per second on NVIDIA H100 GPUs while matching the quality of speed-optimized AR models, demonstrating that dLMs are viable at production scale. These models offer structural advantages that AR models lack: bidirectional context at every step, arbitrary position infilling, and the ability to correct multiple tokens in parallel. The question is no longer whether dLMs can work, but how to combine their strengths with the sequential planning ability of AR models.

Several recent works have begun exploring this direction. Christopher et al. (2025) used diffusion as a drafter for speculative decoding. Horvitz et al. (2024) conditioned diffusion paraphrasers on AR output. Warm-starting diffusion from corrupted inputs is well-studied in images (Meng et al., 2022) and was recently formalized for language (Scholz and Turner, 2025). These efforts suggest that hybrid AR-diffusion pipelines are a promising research direction, but they also raise a question that has received less attention: *how should we evaluate them?*

The NLP community has built mature evaluation infrastructure (Gehrmann et al., 2021; Liang et al., 2023): execution-based code benchmarks (Chen et al., 2021), math reasoning suites (Cobbe et al.,

2021), logical deduction tasks (Suzgun et al., 2023), and LLM-as-judge protocols (Zheng et al., 2023). But this infrastructure was built for left-to-right autoregressive (AR) decoding (Vaswani et al., 2017). Post-processing scripts assume tokens arrive in order. Execution sandboxes assume well-formed syntax. Log-likelihood scoring assumes a single forward pass. These assumptions break quietly when applied to dLMs (Nie et al., 2025; Austin et al., 2021a; Lou et al., 2024), speculative decoding pipelines (Leviathan et al., 2023; Chen et al., 2023), and hybrid systems that combine both. Benchmark contamination (Xu et al., 2024) is already a known concern; we show that *protocol mismatch* is equally serious and largely unexamined.

We expose these issues through *Speculative Refinement* (SpecRef), a training-free hybrid method that uses a small AR model (Phi-2, 2.7B) to draft an initial answer, then selectively masks the least-confident tokens based on their Shannon entropy and lets a large diffusion model (LLaDA-8B (Nie et al., 2025)) refine them (Figure 1). The method works with any off-the-shelf AR drafter and requires no training. Evaluating it across six benchmarks with three evaluation protocols surfaced findings that go well beyond our specific system.

**Contributions.** To the best of our knowledge, this is the first systematic study that pairs an AR language model as a drafter with a masked diffusion model (LLaDA) as a refiner, and evaluates the resulting hybrid across multiple benchmarks and evaluation protocols. Our main contributions are evaluation insights rather than the method itself:

- **Structural discovery vs. logical correctness.** Code benchmarks (HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021b)) conflate two capabilities: discovering syntactic structure (indentation, brackets, function signatures) and writing correct logic. Providing an AR scaffold separates these: accuracy jumps from near 0% to 20%+ at the same diffusion step budget (Table 1), indicating that a substantial portion of the standalone diffusion baseline’s failure on code is structural rather than logical.
- **Refinement tension.** On benchmarks where the AR drafter is already accurate (BBH (Suzgun et al., 2023): 82.8%), multi-stage refinement can *decrease* accuracy (to 80.4%) because the refiner overwrites correct tokens.

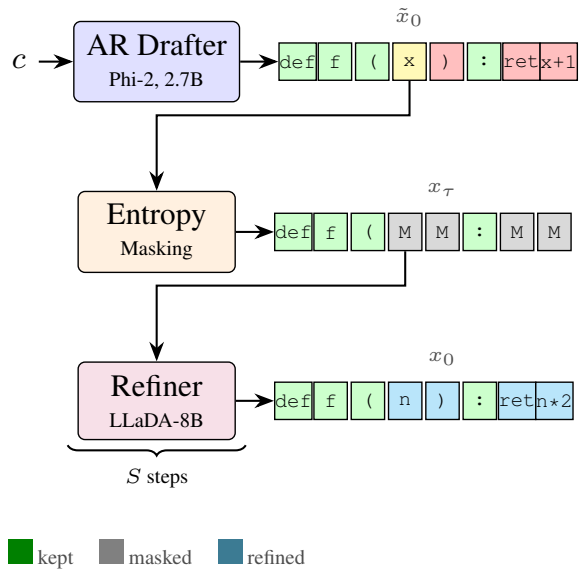


Figure 1: SpecRef pipeline. The AR drafter generates a candidate; per-token entropy identifies uncertain positions (yellow/red), which are replaced with [MASK]. The diffusion refiner denoises from this warm-started state in  $S$  steps.

This ceiling effect is invisible when evaluating either model alone.

- **PPL vs. generative evaluation gap.** On multiple-choice tasks (ARC-Challenge, HellaSwag), log-likelihood scoring produces rankings that differ from those observed on generative benchmarks for the same model pair (Table 2), suggesting that PPL-mode and generative evaluation measure different capabilities in diffusion models.
- **Post-processing pitfalls.** Standard code post-processing (`.strip()`, `textwrap.dedent()`) silently destroys Python indentation for non-AR generators, lowering pass@1 scores. We document this and other practical pitfalls encountered when adapting AR-era evaluation to hybrid systems (Section 3.3).

We also describe SpecRef itself (Section 2) since it is needed to interpret the evaluation findings, and report that it improves absolute accuracy by up to 25.8% over standalone diffusion on code generation while cutting wall-clock latency by up to 2.43×.

## 2 Background: Hybrid AR-Diffusion Generation

We describe the components of our hybrid system, which serves as the evaluation subject throughout the paper.

### 2.1 AR drafter and uncertainty estimation

Let  $c$  denote a prompt and  $x_0 = (x_{0,1}, \dots, x_{0,L})$  be a sequence of  $L$  tokens from vocabulary  $V$ . An AR drafter with parameters  $\phi$  defines a left-to-right factored distribution:

$$q_\phi(x_0 | c) = \prod_{i=1}^L q_\phi(x_{0,i} | c, x_{0,<i}). \quad (1)$$

At each position  $i$ , we record the Shannon entropy as a measure of uncertainty:

$$\begin{aligned} u_i &\triangleq H(q_\phi(\cdot | c, x_{0,<i})) \\ &= - \sum_v q_\phi(v | \cdot) \log q_\phi(v | \cdot). \end{aligned} \quad (2)$$

A high-entropy position means the AR model is unsure about that token, making it a natural candidate for diffusion refinement. In our experiments we use `microsoft/phi-2` (2.7B parameters), but the method works with any AR model that outputs probability logits.

### 2.2 Masked diffusion language modeling

A masked diffusion LM such as LLaDA (Nie et al., 2025) defines a forward process that progressively masks a clean token sequence toward a fully corrupted state using a token-level masking kernel:

$$q_t(x_{t,i} | x_{0,i}) = \begin{cases} x_{0,i} & \text{with prob. } (1 - \alpha_t) \\ [\text{MASK}] & \text{with prob. } \alpha_t \end{cases}, \quad (3)$$

where  $\alpha_t \in [0, 1]$  is a monotone noise schedule with  $\alpha_0 = 0$  (clean) and  $\alpha_T = 1$  (fully masked). The learned reverse process  $p_\theta(x_{t-1} | x_t, c)$  is a bidirectional BERT-style Transformer (LLaDA-8B in our case) that attends to all positions at once and iteratively denoises from  $t = T$  to  $t = 0$ . Unlike AR models, it can condition on “future” tokens at position  $j > i$ , and it can infill arbitrary subsets of tokens.

### 2.3 SpecRef: warm-start with selective masking

The main bottleneck when sampling from a fully masked prior is *structure discovery*: at  $t = T$ , every

token is `[MASK]`, so the model must recover code skeletons, reasoning chains, and answer formats from scratch. With a small step budget  $S$ , it often fails to produce syntactically valid output.

SpecRef warm-starts the diffusion model from an AR draft instead. It samples a draft  $\tilde{x}_0 \sim q_\phi(\cdot | c)$ , then applies the forward corruption kernel at time  $\tau$ :

$$x_\tau \sim q_\tau(\cdot | \tilde{x}_0), \quad (4)$$

and runs the learned reverse process from  $\tau$  to 0:

$$x_{t-1} \sim p_\theta(\cdot | x_t, c) \quad \text{for } t = \tau, \tau-1, \dots, 1. \quad (5)$$

In practice, we use a subset of time indices (e.g., DDIM-style (Song et al., 2022)) to reduce the number of forward passes to  $S \ll \tau$ .

Rather than masking globally, SpecRef masks *selectively*: only the top  $k\%$  highest-entropy positions are replaced with `[MASK]`. Two heuristics extend the mask set  $\mathcal{I}$ :

1. **Math-aware block expansion** ( $e = 2$ ): positions aligned to numeric or math-operator characters are identified via regex. All tokens within  $\pm 2$  positions are added:  $\mathcal{I} \leftarrow \mathcal{I} \cup \{i | |i-j| \leq 2, j \in \mathcal{M}\}$ , where  $\mathcal{M}$  indexes math-character positions. This ensures arithmetic context is regenerated as a unit.
2. **Tail truncation** ( $M = 150$ ): all tokens past index 150 are unconditionally masked to force the diffusion model to regenerate conclusions, where AR drift is worst.

Token scores for the top- $k$  selection are:

$$m_i = \text{clip}\left(\frac{u_i}{\log |V|}, 0, 1\right), \quad (6)$$

where  $u_i$  is mapped to refiner-token positions via character-level offset alignment (needed because the drafter and refiner use different tokenizers). We use  $k = 60\%$  throughout. Algorithm 1 gives the full procedure.

### 2.4 Compute cost and the handoff point

The total cost is  $L C_{\text{AR}} + S C_{\text{D}}$ , where  $C_{\text{AR}}$  is the amortized cost of one AR token (negligible under vLLM serving) and  $C_{\text{D}}$  is the cost of one diffusion forward pass over all  $L$  positions. A standalone diffusion sampler with  $N$  steps costs  $N C_{\text{D}}$ ; SpecRef targets  $S \ll N$ .

---

**Algorithm 1** Statistical SpecRef (Top- $k$  Entropy Masking)

---

**Require:** Prompt  $c$ ; AR drafter  $q_\phi$ ; refiner  $p_\theta$ ; mask percentile  $k$ ; math-expansion window  $e$ ; tail truncation  $M$ .

**Ensure:** Refined sample:  $x_0$

- 1: Sample draft:  $\tilde{x}_0 \sim q_\phi(\cdot | c)$
- 2: Extract logit entropies:  
 $u_i = H(q_\phi(\cdot | c, \tilde{x}_{0, < i}))$
- 3: Compute the set of top  $k\%$  tokens by entropy:  $\mathcal{I}_{\text{top}}$
- 4: **Math Block Expansion:** Let  $\mathcal{M}$  be indices of tokens aligned to numeric/math characters. Set  $\mathcal{I} = \mathcal{I}_{\text{top}} \cup \{i \mid |i - j| \leq e, j \in \mathcal{M}\}$
- 5: **Tail Truncation:** Set  $\mathcal{I} = \mathcal{I} \cup \{i \mid i > M\}$
- 6: Form intermediate mask state  $x_\tau$ :

$$x_{\tau, i} = \begin{cases} [\text{MASK}] & \text{if } i \in \mathcal{I} \\ \tilde{x}_{0, i} & \text{otherwise} \end{cases}$$

- 7: **for** refinement steps  $S$  **do**
  - 8:   Sample proposal:  $x'_{t-1} \sim p_\theta(\cdot | x_t, c)$
  - 9: **end for**
  - 10: **return**  $x_0$
- 

The mask ratio  $k\%$  (equivalently the warm-start time  $\tau$ ) controls a tradeoff:  $k \rightarrow 0$  preserves the entire draft;  $k \rightarrow 100\%$  reduces to standard diffusion from scratch. We use  $k = 60\%$ : enough tokens survive to carry global structure, while enough are masked for correction. On GSM8K (RTX 4090), this yields 6.79s per query vs. 16.49s for standalone LLaDA at  $S = 64$ , a  $2.43\times$  speedup.

### 3 Evaluation Setup

#### 3.1 Models and infrastructure

All experiments ran on the  $8\times$  NVIDIA V100, 16GB each. The refiner (LLaDA-8B-Instruct) is split across 2 GPUs via pipeline parallelism. The drafter (Phi-2, 2.7B) runs under vLLM with 4-bit quantization; its latency is under 1% of total generation time.

#### 3.2 Evaluation protocols

We deliberately chose benchmarks that require three different evaluation protocols, since the protocol itself turned out to be a significant variable.

**Protocol A: Execution-based (pass@1).** HumanEval (Chen et al., 2021) and MBPP (Austin

et al., 2021b). Generated Python code is executed against ground-truth unit tests in a sandbox with a 10-second SIGALRM timeout (infinite loops are common in early diffusion steps). We report pass@1.

**Protocol B: Exact match with parsing.** GSM8K (Cobbe et al., 2021) and MATH (supplementary, evaluated on a 5,000-sample subset). Models are prompted to produce answers in `\boxed{\dots}` format. We extract the final answer with regex and compute exact-match accuracy.

**Protocol C: Log-likelihood scoring (PPL-mode).** ARC-Challenge and HellaSwag. Each answer choice is formatted as `context + choice_text`. The conditional log-likelihood is estimated via 8-round Monte Carlo masking (randomly masking a subset of continuation tokens per round, scoring under the model, and averaging), following the LLaDA evaluation protocol (Nie et al., 2025). The highest-scoring choice is selected. This protocol is diffusion-step-independent: no iterative denoising is performed, giving a  $4\times$  speedup over generative evaluation.

#### 3.3 Evaluation pitfalls encountered

Adapting standard AR-era evaluation to a hybrid system revealed several pitfalls that would affect any non-AR or multi-stage generator:

1. **Post-processing breaks indentation.** Standard code post-processing (`.strip()`, `textwrap.dedent()`) destroys Python’s semantic indentation. For diffusion outputs, which may contain irregular whitespace patterns unlike AR outputs, this lowered pass@1 scores significantly. We replaced these with stop-sequence truncation (`[\nndef, \nclass, \nif __name__]`).
2. **Infinite loops require sandboxing.** At low step counts ( $S \leq 16$ ), diffusion models frequently produce syntactically valid but non-terminating code. Without the SIGALRM timeout, evaluation hangs indefinitely, a problem AR models rarely cause.
3. **Proxy metrics mislead for PPL-mode benchmarks.** During development, we tried string-matching the first 40 characters of the target as a proxy for PPL-mode benchmarks.

The proxy and true metric disagreed substantially, producing misleading intermediate results.

4. **Tokenizer mismatch.** The drafter and refiner use different tokenizers. Entropy scores must be mapped between tokenizations via character-level offset alignment; naive token-index transfer produces incorrect masks and unreliable evaluation.

## 4 Results

We organize results around evaluation findings rather than per-dataset performance. All numerical results are in Tables 1–3; we do not repeat numbers already in the tables except when interpreting them.

### 4.1 What code benchmarks actually test: structural discovery

The most striking pattern in Table 1 is on code benchmarks. Standalone LLaDA produces 0% valid code on HumanEval at  $S \leq 16$  and only 8.70% at  $S = 64$ . But when warm-started from an AR draft, SpecRef reaches 20.12% at just 8 steps, without changing the diffusion model at all.

This gap tells us something about what HumanEval and MBPP actually measure. This suggests the diffusion model at 8B parameters is not entirely incapable of code logic; rather, it cannot discover the right syntactic structure (indentation, brackets, function signatures) from a fully masked starting point in a few steps. Once that structure is provided, accuracy jumps immediately. These benchmarks appear to conflate two distinct capabilities: *structural discovery* and *logical correctness*. For diffusion models at low step counts, the structural component accounts for a large share of the failure.

A near-zero score on HumanEval does not mean a diffusion model “cannot code.” It may mean the model cannot discover Python’s formatting conventions quickly enough. Separating structural and logical evaluation (e.g., providing scaffolds and measuring only logic fill-in) would give more diagnostic information.

On MBPP the pattern is even clearer: SpecRef’s absolute gain expands to 25.80% at  $S = 64$  (33.00% vs. 7.20%). On GSM8K, SpecRef at  $S = 64$  reaches 54.28%, outperforming both the standalone drafter (47.38%) and refiner (44.81%), suggesting complementary strengths on math reasoning.

### 4.2 Refinement tension: when correction hurts

A second finding concerns what happens when the AR drafter is already accurate. On BBH Boolean Expressions (Suzgun et al., 2023), the drafter scores 82.80%. SpecRef improves over the standalone diffusion baseline at every step count but plateaus at 80.40% for  $S \geq 32$ , below the drafter itself. The same pattern appears on MATH, where the baseline edges ahead of SpecRef at  $S = 64$  (22.32% vs. 21.38%).

We call this *refinement tension*: the entropy-based mask sometimes selects tokens that the drafter already had correct, and the refiner replaces them with worse alternatives. The drafter’s entropy is a noisy signal of actual error; on saturated tasks, high entropy does not always mean the token is wrong.

This phenomenon is invisible when evaluating either model in isolation. It only appears in the pipeline, and it reveals a benchmark saturation ceiling: once a component model is near-ceiling, aggregate pipeline scores can *decrease*. For evaluation of any multi-stage system (not just SpecRef), this argues for reporting component-level scores alongside pipeline scores. It also suggests that adaptive, per-instance mask rate selection (lowering  $k$  when the drafter is already confident) could mitigate the effect, though we leave this to future work.

### 4.3 The log-likelihood vs. generative evaluation disconnect

Table 2 shows PPL-mode evaluation on ARC-Challenge and HellaSwag. Two observations are worth noting.

First, PPL-mode is diffusion-step-independent. The same score emerges regardless of step count, because log-likelihood is computed in a single masked forward pass, not through iterative denoising. This makes PPL-mode evaluation 4× cheaper on our cluster, but it also means PPL-mode and generative evaluation measure fundamentally different things. Whether a model can *score* a correct answer highly is not the same as whether it can *generate* that answer.

Second, the AR drafter outperforms LLaDA on HellaSwag in PPL mode (54.94% vs. 45.62%), yet on generative benchmarks (Table 1) the relationship between the two models is more complex and step-dependent. While these are different tasks, the pattern suggests that PPL-mode rankings do

Table 1: Step-wise scaling on generative benchmarks. SpecRef sets a quality floor at low step counts ( $S = 8, 16$ ) where standalone diffusion fails. All results use `microsoft/phi-2` as drafter and `GSAI-ML/LLaDA-8B-Instruct` as refiner.

Dataset	Phi-2 (Draft)	LLaDA Baseline				SpecRef (Ours)			
		S=8	S=16	S=32	S=64	S=8	S=16	S=32	S=64
GSM8K	47.38%	8.72%	12.43%	22.14%	44.81%	39.50%	45.26%	50.95%	<b>54.28%</b>
HumanEval	44.51%	0.00%	0.00%	1.74%	8.70%	20.12%	20.73%	20.73%	<b>25.61%</b>
MBPP	43.00%	0.20%	0.60%	2.00%	7.20%	20.80%	22.80%	26.00%	<b>33.00%</b>
MATH*	5.22%	10.76%	12.72%	16.12%	<b>22.32%</b>	14.16%	17.70%	20.44%	21.38%
BBH (Bool)	82.80%	54.80%	53.60%	58.80%	66.00%	71.20%	76.80%	80.40%	<b>80.40%</b>

\* Note: The MATH evaluation was performed on a 5,000-sample representative subset of the final test suite.

Table 2: Likelihood (PPL) performance. SpecRef in PPL mode matches the baseline because both use log-likelihood scoring on the full sequence.

Dataset	Phi-2 (Draft)	Baseline (S=64)	SpecRef
ARC-Chall.	58.10%	51.45%	51.50%
HellaSwag	54.94%	45.62%	45.62%

not straightforwardly transfer to generative settings. For diffusion models specifically, PPL-mode may overstate or understate generation ability depending on the task.

#### 4.4 Latency-quality tradeoffs

Table 3: Latency (seconds per query) on GSM8K, NVIDIA RTX 4090. SpecRef times include both the AR drafting pass (vLLM) and the diffusion steps. Because fewer diffusion steps are needed to match baseline quality, total wall-clock time drops.

Steps ( $S$ )	LLaDA Baseline (s)	SpecRef (s)	Speedup Factor
16	4.12	5.18	0.80×
32	8.24	5.95	1.38×
64	16.49	6.79	<b>2.43×</b>

Table 3 shows that SpecRef  $S = 64$  is both more accurate and  $2.43\times$  faster than standalone LLaDA at the same step count. At  $S = 16$  SpecRef, it is slower (the drafting pass dominates when diffusion steps are few) but far more accurate.

For evaluation methodology, the lesson is that **reporting only accuracy is insufficient** for comparing generation paradigms with different cost structures. Step count, wall-clock time, and FLOPs all affect the comparison. Table 1 shows accuracy swings of over 30 percentage points across step counts for the same model; a single-step-count

comparison could be arbitrarily misleading.

## 5 Discussion: Lessons for Evaluating Hybrid Generation

**1. Report evaluation protocol details exhaustively.** Post-processing choices alone can change code evaluation scores by large margins (Section 3.3). Sandbox timeouts, stop-sequence selection, and tokenizer alignment details are not minor implementation notes; they are evaluation-critical parameters. We recommend treating them with the same care as hyperparameters.

**2. Benchmark scores conflate multiple capabilities.** Code benchmarks test structural discovery and logical correctness simultaneously. For AR models, this distinction rarely matters (structure comes for free from left-to-right generation), but for diffusion and hybrid models, it is the dominant factor. Providing syntactic scaffolds as a controlled experiment could separate these components and give more diagnostic evaluations.

**3. Multi-stage systems need component-level reporting.** Aggregate pipeline scores can mask component-level regression (the refinement tension effect). We recommend reporting drafter-only, refiner-only, and pipeline scores together, along with an analysis of where the pipeline helps and where it hurts.

**4. PPL-mode and generative evaluation are not interchangeable.** They measure different aspects of model capability and can rank the same models differently. Papers that use only one mode should be explicit about this limitation.

**5. Step count and compute budget must be reported.** Diffusion model accuracy varies by 30+ pp across step counts (Table 1). A score without

a step count is incomplete; accuracy-vs-compute curves should replace single numbers.

## 6 Related Work

GEM (Gehrmann et al., 2021) and HELM (Liang et al., 2023) established shared evaluation infrastructure; LLM-as-judge methods (Zheng et al., 2023) added scalable evaluation beyond reference matching; and benchmark contamination (Xu et al., 2024) has motivated living benchmarks. Deschenaux and Gulcehre (2024) noted evaluation challenges for discrete diffusion but did not systematically document them, our work fills this gap. Continuous-space diffusion for text was introduced by Li et al. (2022); discrete extensions (Austin et al., 2021a; Ho et al., 2020; Song et al., 2021) led to SEDD (Lou et al., 2024) and MDLM (Sahoo et al., 2024). Speculative decoding (Leviathan et al., 2023; Chen et al., 2023) accelerates AR generation via draft-then-verify; SpecRef inverts this pattern, adapting warm-start diffusion (Meng et al., 2022; Scholz and Turner, 2025) to discrete masked diffusion with a confidence-aware corruption policy.

## 7 Conclusion

We used SpecRef as a lens to examine how standard benchmarks behave when generation is no longer purely autoregressive, surfacing four findings: code benchmarks conflate structural discovery with logical correctness; multi-stage pipelines can degrade accuracy through refinement tension; likelihood and generative evaluation are not interchangeable; and standard post-processing silently breaks evaluation for non-AR generators. We hope these observations contribute to more diagnostic evaluation practices as generation architectures continue to diversify.

## Limitations

Our evaluation uses a single drafter-refiner pair (Phi-2 2.7B + LLaDA 8B); quantitative magnitudes are specific to this combination, though the qualitative findings likely generalize. Phi-2’s 2,048-token context window prevented evaluation on long-prompt benchmarks (e.g., GPQA). The fixed mask rate  $k = 60\%$  was not tuned per task; adaptive masking may reduce refinement tension.

## Acknowledgement

This work was carried out at IIIT, Hyderabad. We thank the institute for logistics and compute resources. For GPU rentals, we used vast.ai, we acknowledge the seamless access. This work was funded by Qualcomm faculty award at IIIT, Hyderabad for funding the first author. We thank Qualcomm for such unrestricted grant to support research.

## Ethics Statement

SpecRef is a decoding strategy; it does not change training data or model capabilities. It does, however, reduce the cost of generating text, which can amplify both beneficial and harmful uses. Risks include:

- **Scaling misuse:** faster generation can facilitate spam, disinformation, or automated harassment.
- **Bias propagation:** refinement can preserve biases from either model, and selective masking may “lock in” problematic phrasing from the draft.
- **Over-trust in refinement:** users may assume “refined” outputs are more factual than they are.

Mitigations include maintaining standard safety filters, auditing outputs for bias/toxicity, and using evaluation prompts that stress factuality and sensitive domains.

## References

- Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. 2021a. Structured denoising diffusion models in discrete state-spaces. In *Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS ’21*, Red Hook, NY, USA. Curran Associates Inc.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021b. [Program synthesis with large language models](#). *Preprint*, arXiv:2108.07732.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. 2023. [Accelerating large language model decoding with speculative sampling](#). 2302.01318.

- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Jacob K Christopher, Brian R. Bartoldson, Tal Ben-Nun, Michael Cardei, Bhavya Kailkhura, and Ferdinando Fioretto. 2025. [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)*, pages 12042–12059, Albuquerque, New Mexico. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). 2110.14168.
- Justin Deschenaux and Caglar Gulcehre. 2024. [Promises, outlooks and challenges of diffusion language modeling](#). ArXiv: 2406.11473.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, and 37 others. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). In *Proceedings of the First Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 96–120, Online. Association for Computational Linguistics.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. [Denoising diffusion probabilistic models](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Zachary Horvitz, Ajay Patel, Chris Callison-Burch, Zhou Yu, and Kathleen McKeown. 2024. [Paraguide: guided diffusion paraphrasers for plug-and-play textual style transfer](#). In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence, AAAI'24/IAAI'24/EAAI'24*. AAAI Press.
- Inception Labs, Samar Khanna, Siddhant Kharbanda, Shufan Li, Harshit Varma, Eric Wang, Sawyer Birnbaum, Ziyang Luo, Yanis Miraoui, Akash Palrecha, Stefano Ermon, Aditya Grover, and Volodymyr Kuleshov. 2025. [Mercury: Ultra-fast language models based on diffusion](#). ArXiv: 2506.17298.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. [Fast inference from transformers via speculative decoding](#). In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. 2022. [Diffusion-lm improves controllable text generation](#). In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, and 31 others. 2023. [Holistic evaluation of language models](#). 2211.09110.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. 2024. [Discrete diffusion modeling by estimating the ratios of the data distribution](#). ArXiv: 2310.16834.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2022. [Sdedit: Guided image synthesis and editing with stochastic differential equations](#). ArXiv: 2108.01073.
- Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025. [Large language diffusion models](#). *Preprint*, arXiv:2502.09992.
- Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T Chiu, Alexander Rush, and Volodymyr Kuleshov. 2024. [Simple and effective masked diffusion language models](#). In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24*, Red Hook, NY, USA. Curran Associates Inc.
- Jonas Scholz and Richard E. Turner. 2025. [Warm starts accelerate conditional diffusion](#). ArXiv: 2507.09212.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2022. [Denoising diffusion implicit models](#). ArXiv: 2010.02502.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. [Score-based generative modeling through stochastic differential equations](#). ArXiv: 2011.13456.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. [Challenging BIG-bench](#)

tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.

Ruijie Xu, Zengzhi Wang, Run-Ze Fan, and Pengfei Liu. 2024. [Benchmarking benchmark leakage in large language models](#). ArXiv: 2404.18824.

Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. 2025. [Dream 7b: Diffusion large language models](#). ArXiv: 2508.15487.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.