NEURAL INFORMATION PROCESSING SYSTEMS

Machine & Language Learning Lab @ Yonsei

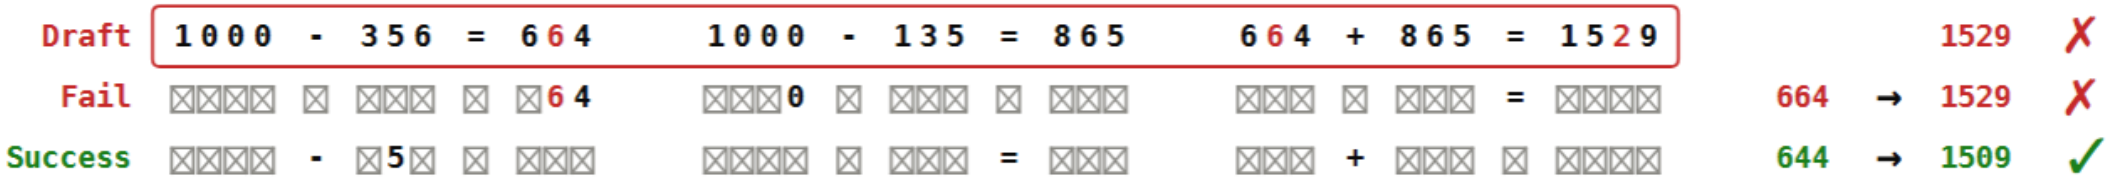
Renoise Consistency: Unlocking Efficient Self-Correction for Diffusion Large Language Models

Jihoon Lee[1], Yeongbin Seo[2,3], Jaehyung Kim[1]

[1]Yonsei University, [2]University of Toronto, [3]Vector Institute

Self-Correction Potential of Diffusion Large Language Models (DLMs)

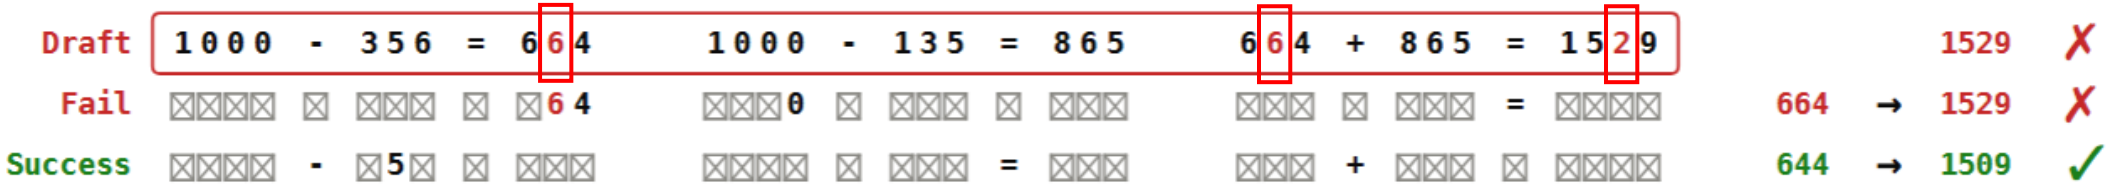
Errors introduced during decoding propagate to the final answer.
DLMs have the potential to self-correct this errors.



Examples of errors in mathematical reasoning traces and successful or failed corrections.

Self-Correction Potential of Diffusion Large Language Models (DLMs)

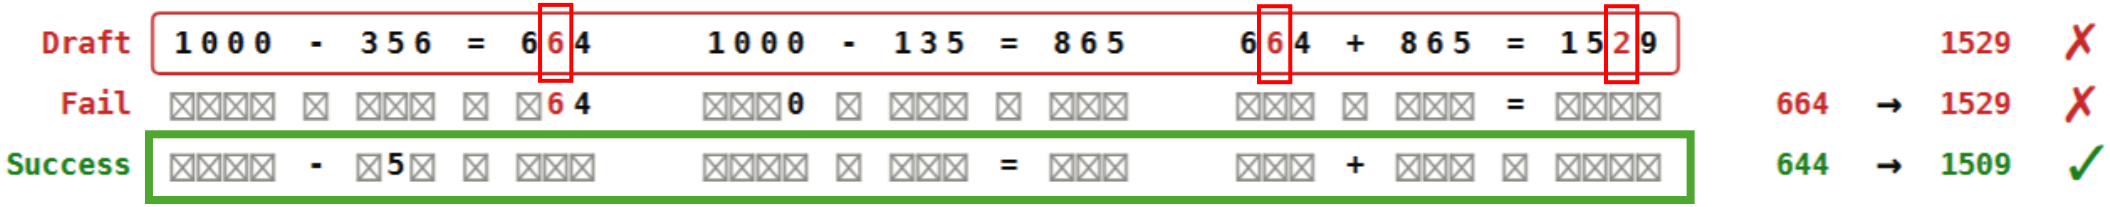
However, detecting specific errors (664, 1529) in reasoning traces is challenging.



Examples of errors in mathematical reasoning traces and successful or failed corrections.

Self-Correction Potential of Diffusion Large Language Models (DLMs)

However, detecting specific errors (664, 1529) in reasoning traces is challenging.



Examples of errors in mathematical reasoning traces and successful or failed corrections.

How can we construct a masked context for self-correction without explicitly targeting errors?

(Errors masked, allowing an opportunity for correction)

Motivation: How DLMs Are Trained

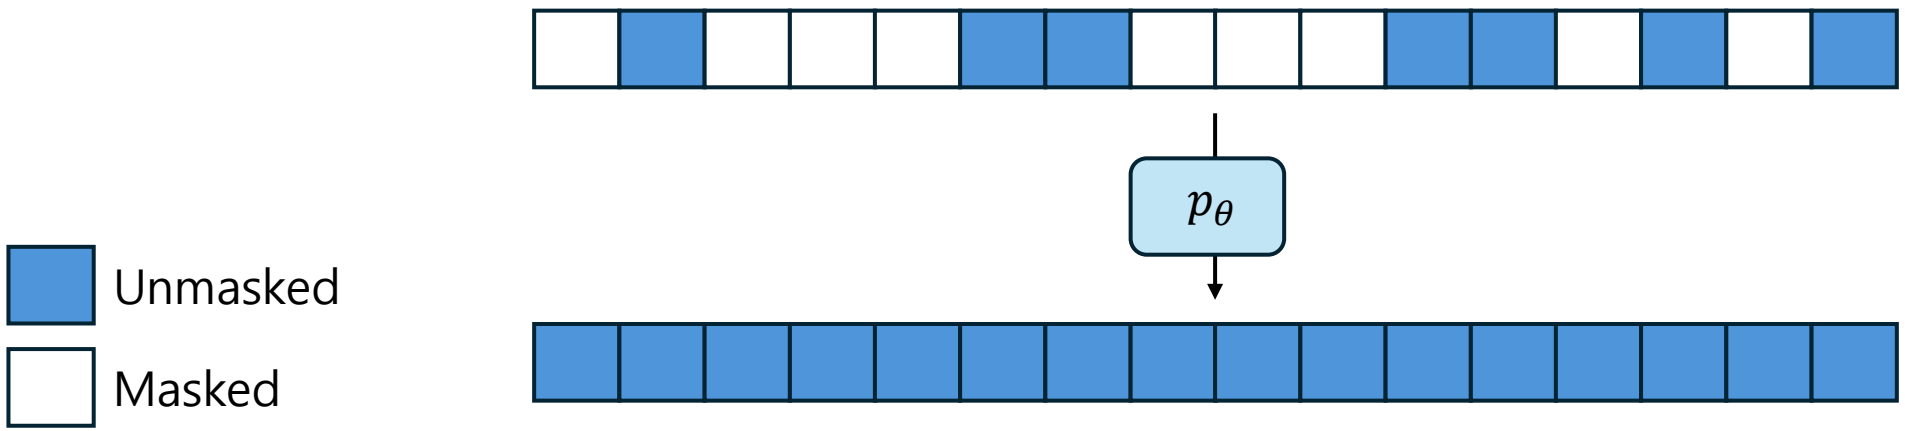
During training, DLMs learned..

1) 1 step generation
2) From scattered token contexts

Train

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

Training Objective of DLMs. Each token target sequence is independently, randomly (t%) masked.



Motivation: How DLMs Are Trained

During training, DLMs learned..

1) 1 step generation
2) From scattered token contexts

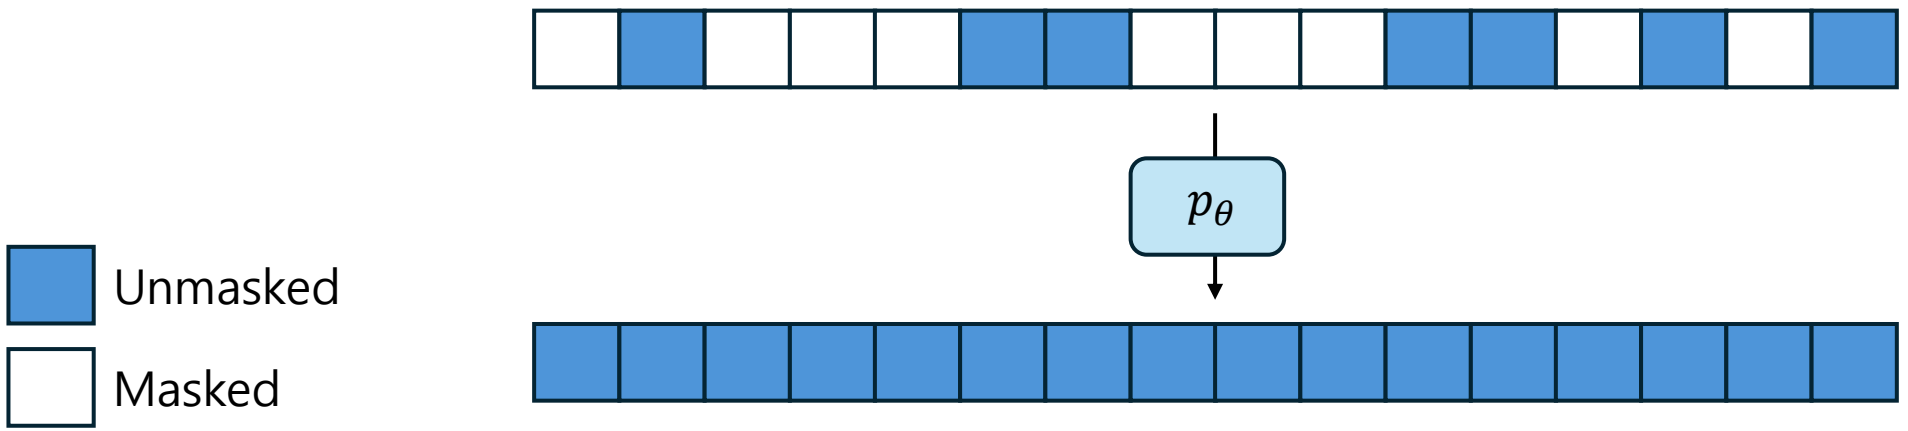
During verification, DLMs can..

1) 1 step reconstruction on a draft
2) From scattered token contexts (contexts for correction)

Train

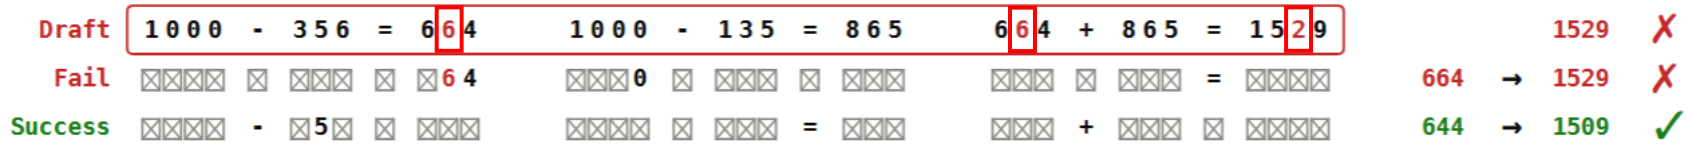
$$\mathcal{L}(\theta) \triangleq - \mathop{\mathbb{E}}_{t,x_0,x_t} \left[\frac{1}{t} \sum_{i=1}^{L} \mathbb{1}\left[x_t^i = M\right] \log p_\theta\left(x_0^i \mid x_t\right) \right]$$

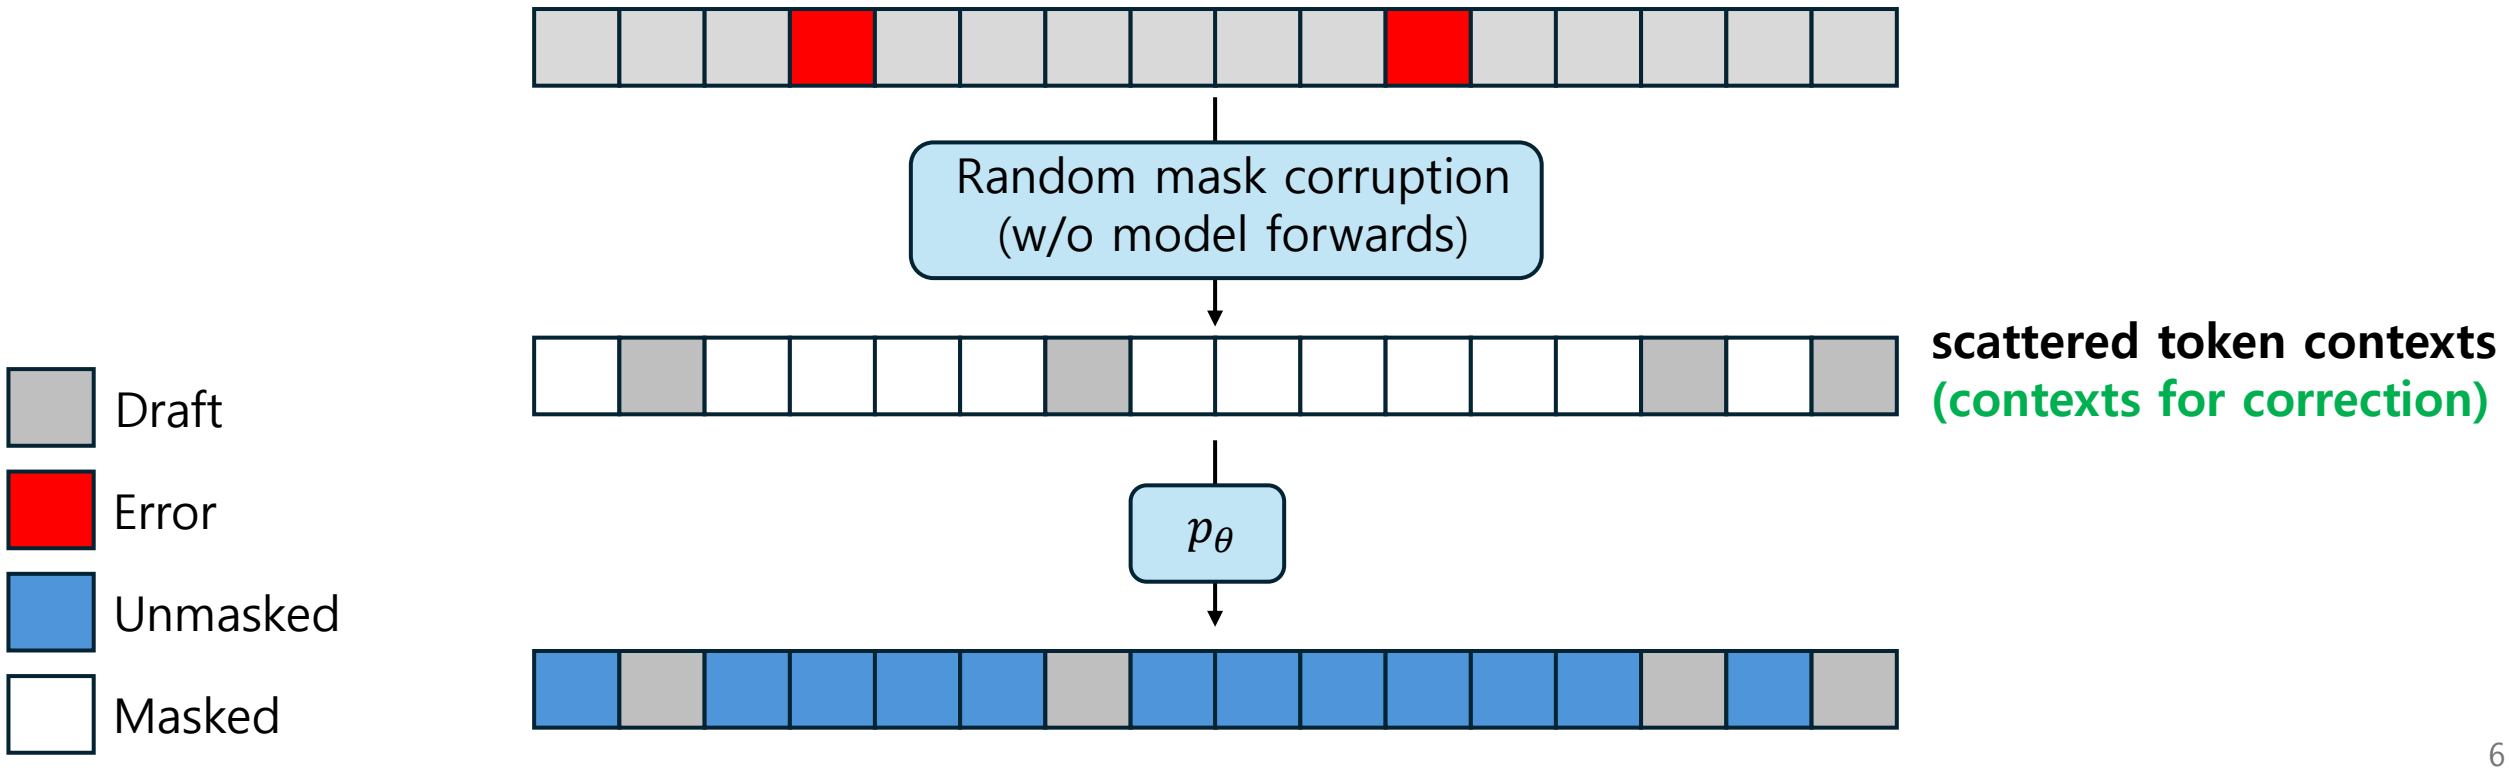
Training Objective of DLMs. Each token target sequence is independently, randomly (t%) masked.



Method: How DLMs Are Trained

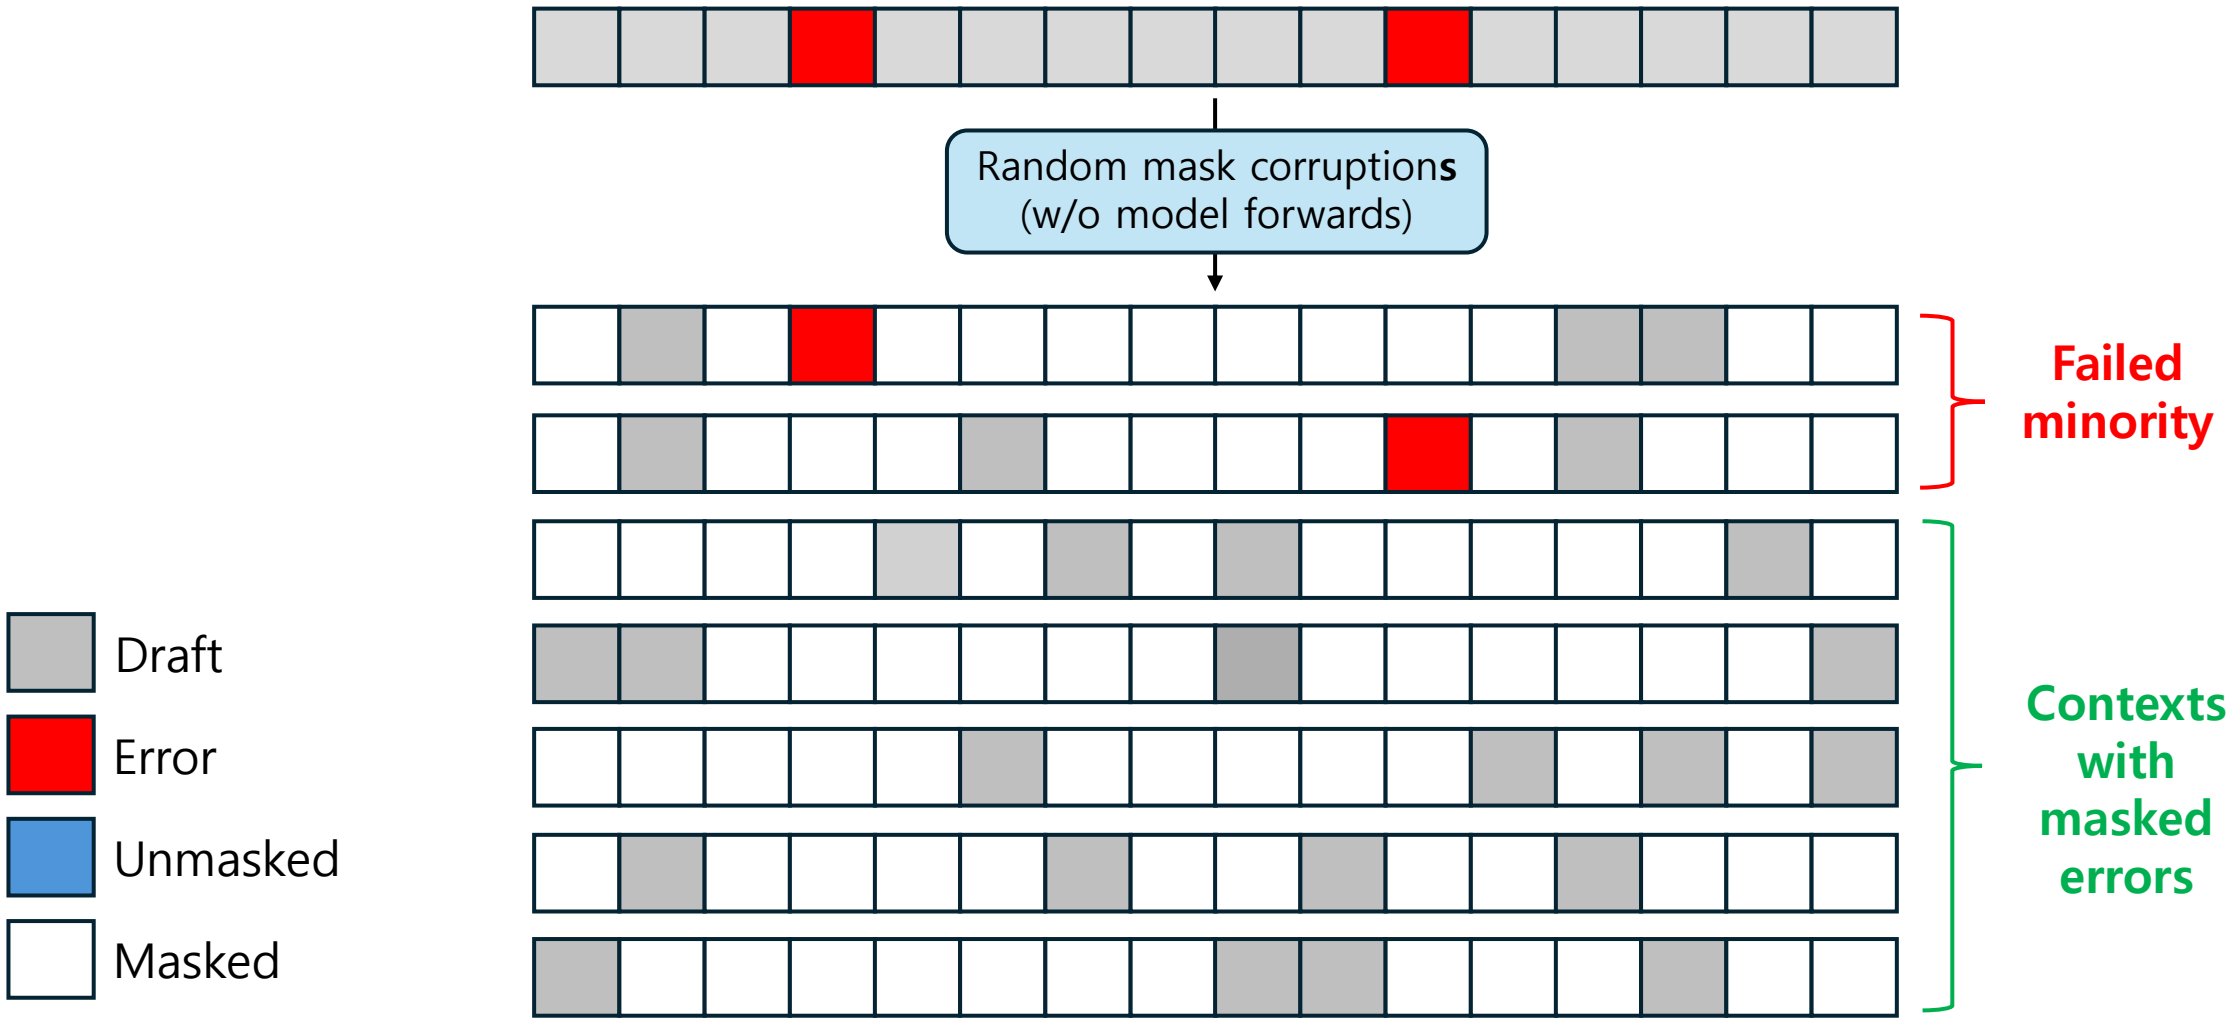
Errors are typically localized and scattered.
By randomly masking most (e.g. 90%) of the generated draft output, we can mask these errors with high probability.





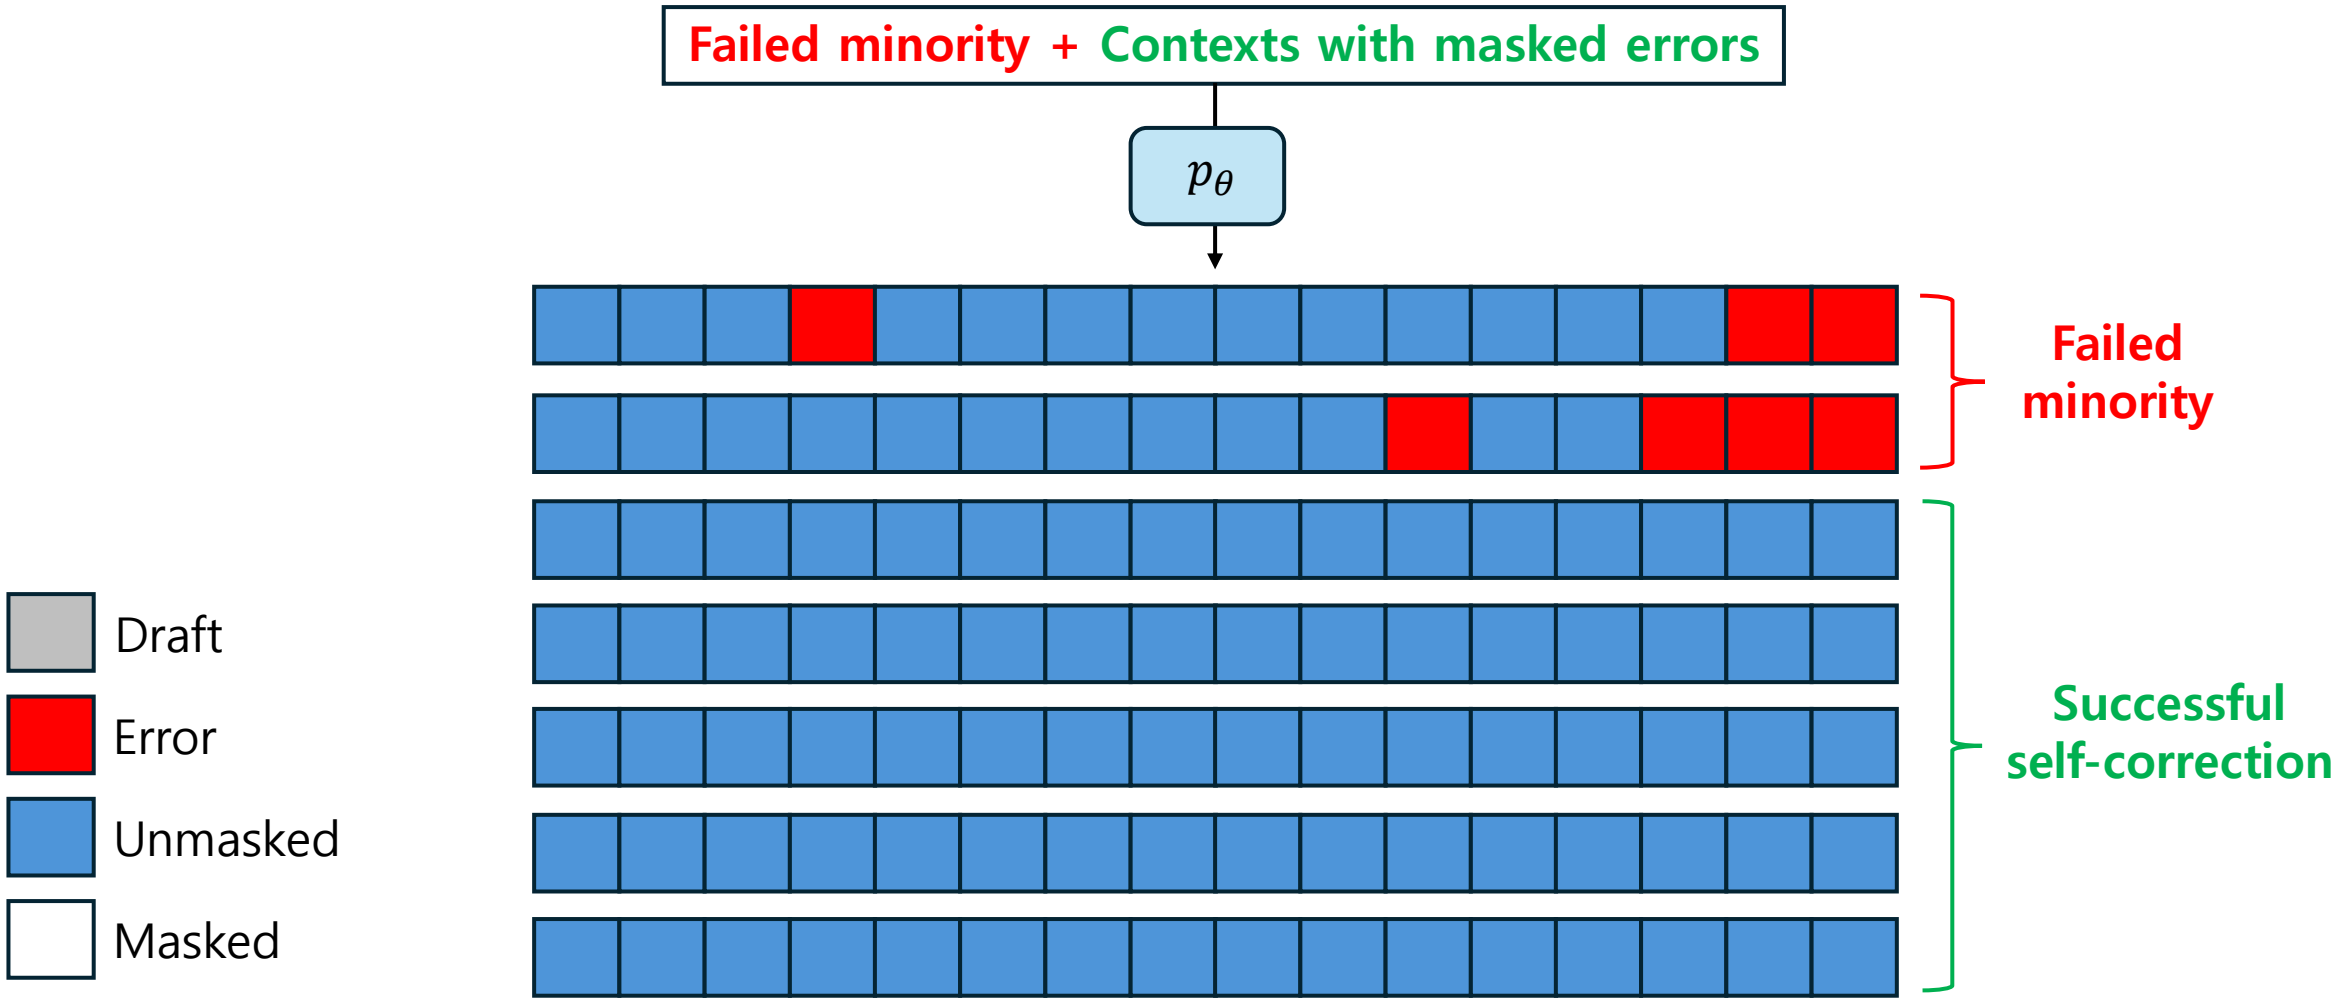
Method: Renoise Consistency (ReCon)

Of course, some masking patterns may fail to cover all errors…
However, diverse masking patterns will cover all errors in most cases.



Method: Renoise Consistency (ReCon)

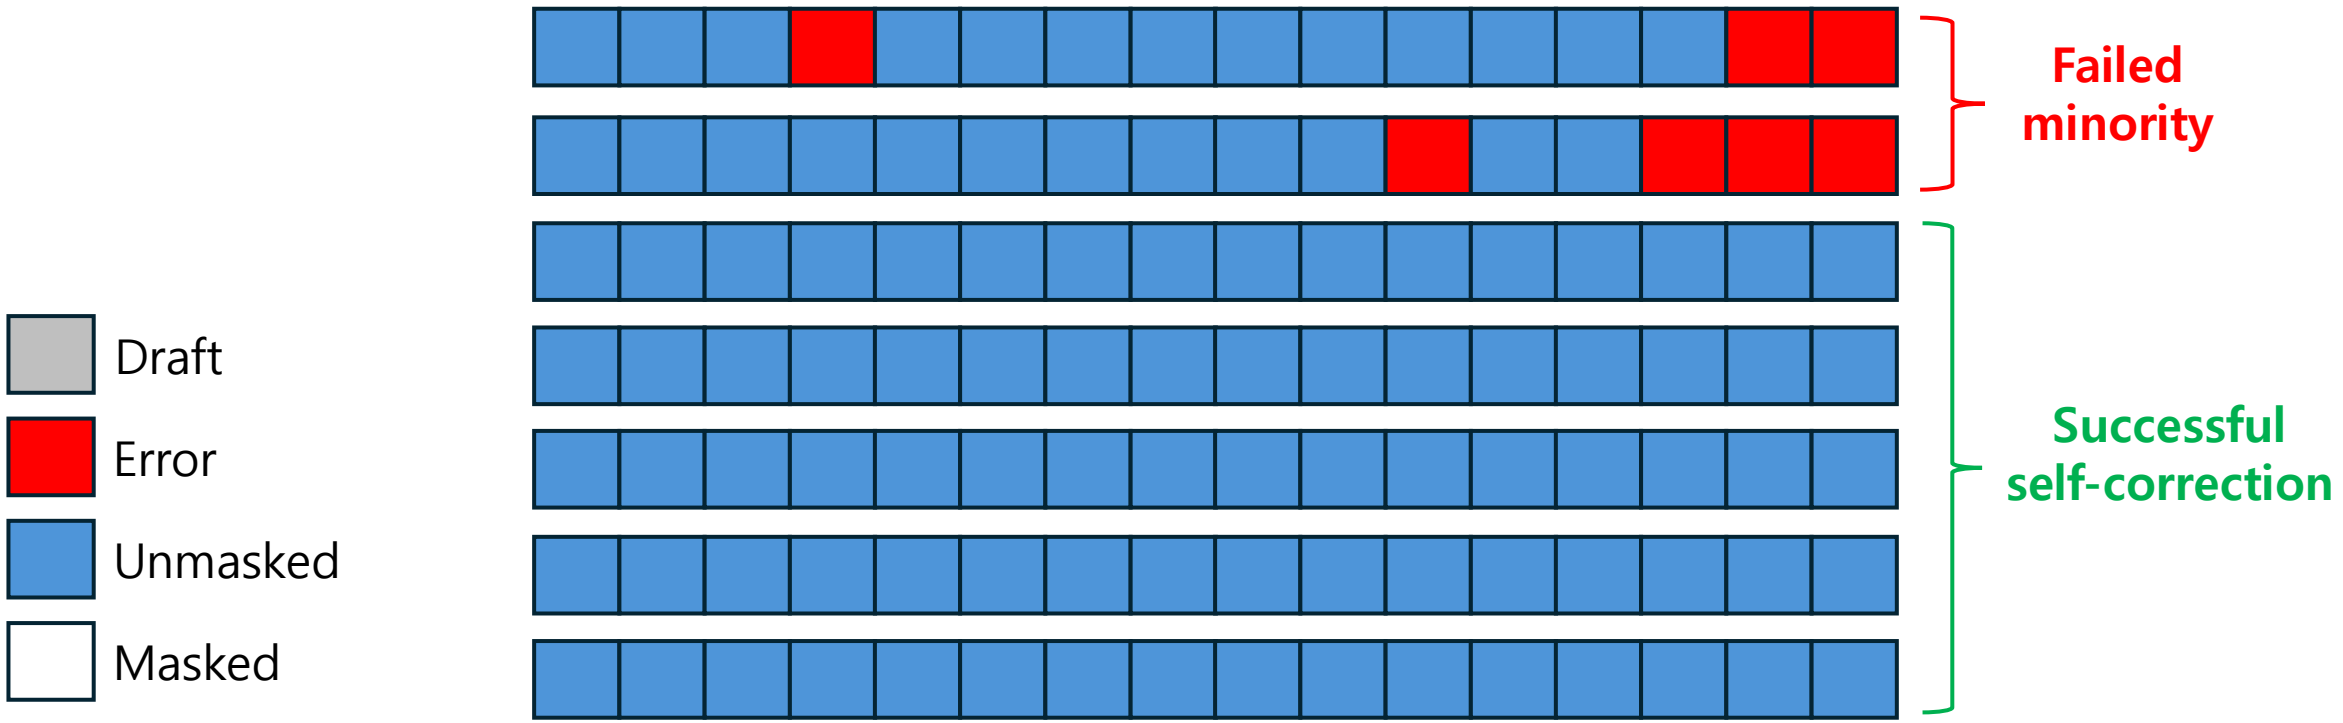
Ultimately, most masking patterns get a chance to correct errors, and their scattered contexts enable high-quality self-correction with just an efficient few forward passes.



Method: Renoise Consistency (ReCon)

Now, what if we select one of the successfully self-corrected candidates?
Since there is only one correct answer,
most successful candidates predict the same answer.

→ Majority voting works!

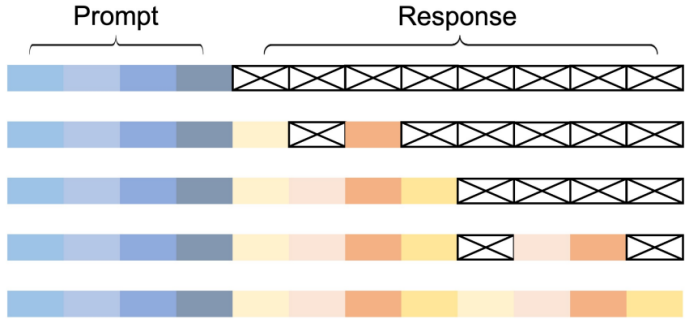


Method: Renoise Consistency (ReCon)

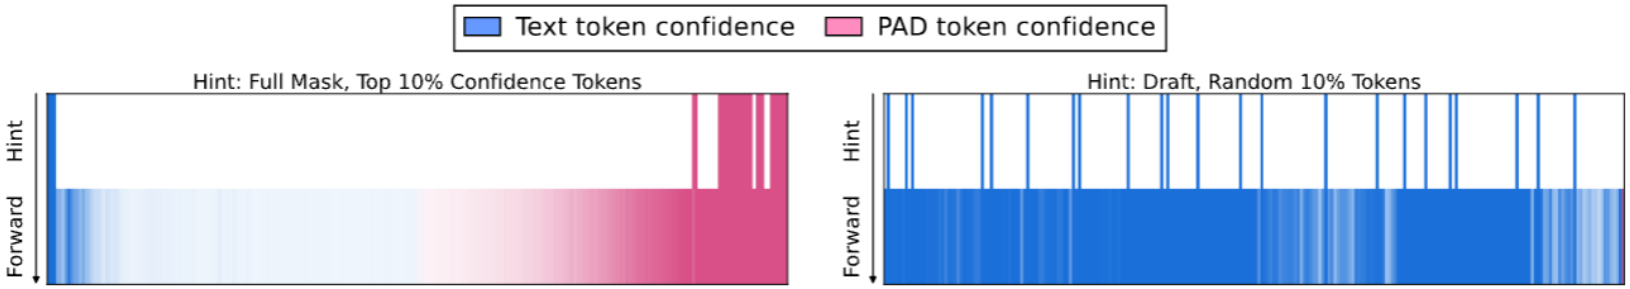
Train

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

Inference



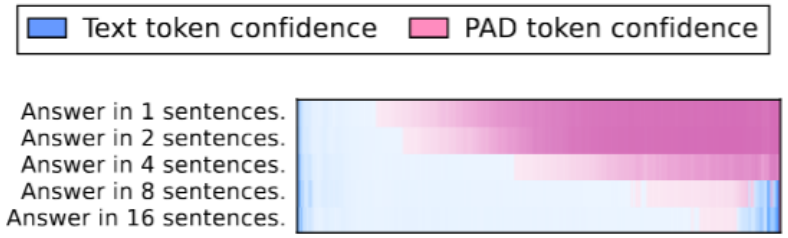
(c) Block Diffusion LLaDA.

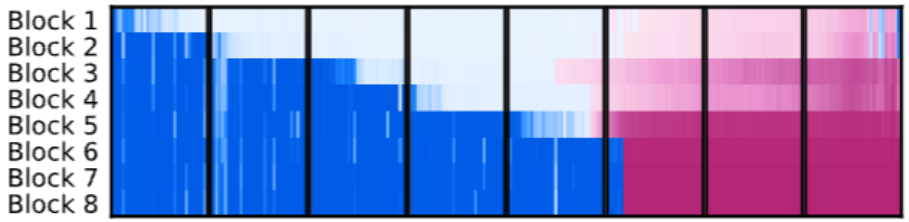


Takeaway. **ReCon's Scattered Anchor Context** resolves the **training-inference mismatch.**

[1] Nie et al., "Large Language Diffusion Models," *NeurIPS 2025*.

Method: Adaptive Block Drafting (ABD)





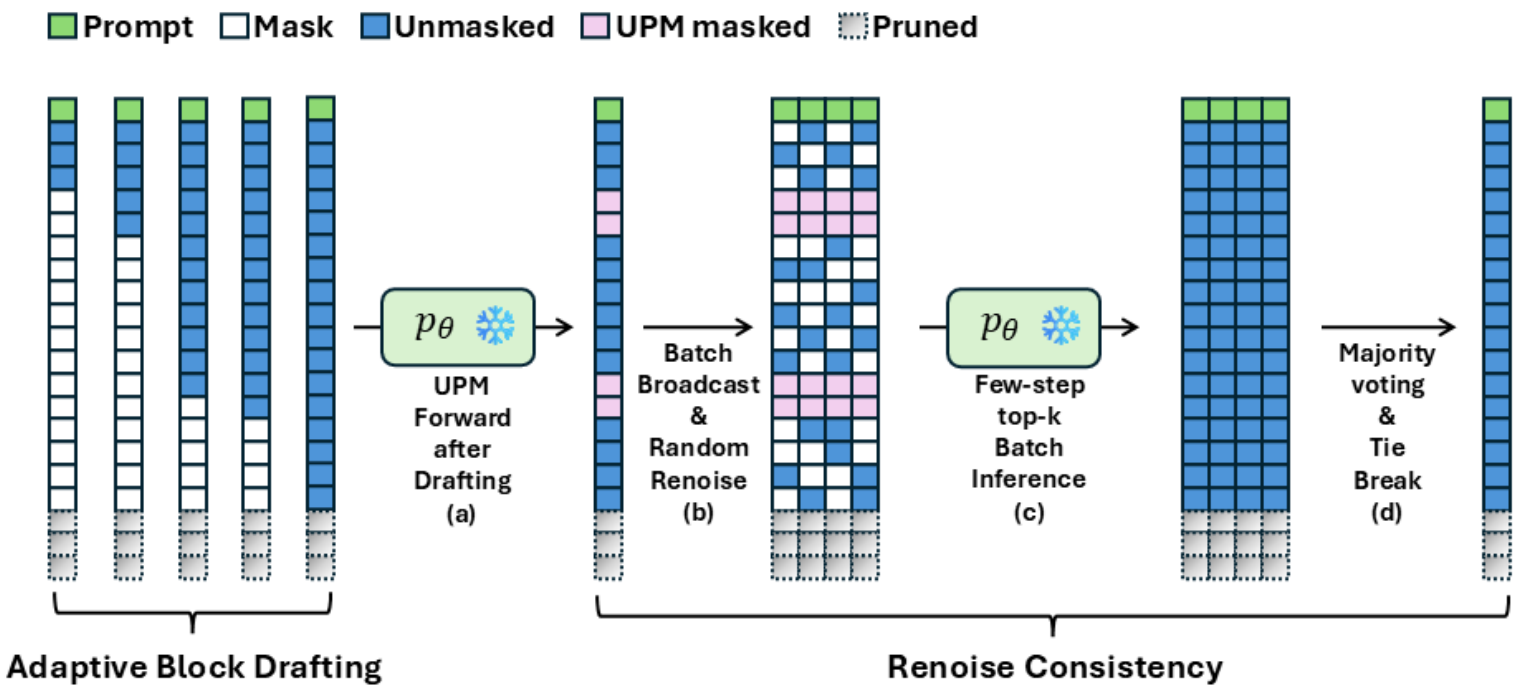
Another Observation. Bidirectional Inertia reveals the optimal **output length** or **block length.**

Left: The **PAD** token confidence distribution **(output length)**

Right: The **text** token confidence distribution **(block length)**

Adaptive **B**lock **D**rafting **(ABD),** efficient drafting.

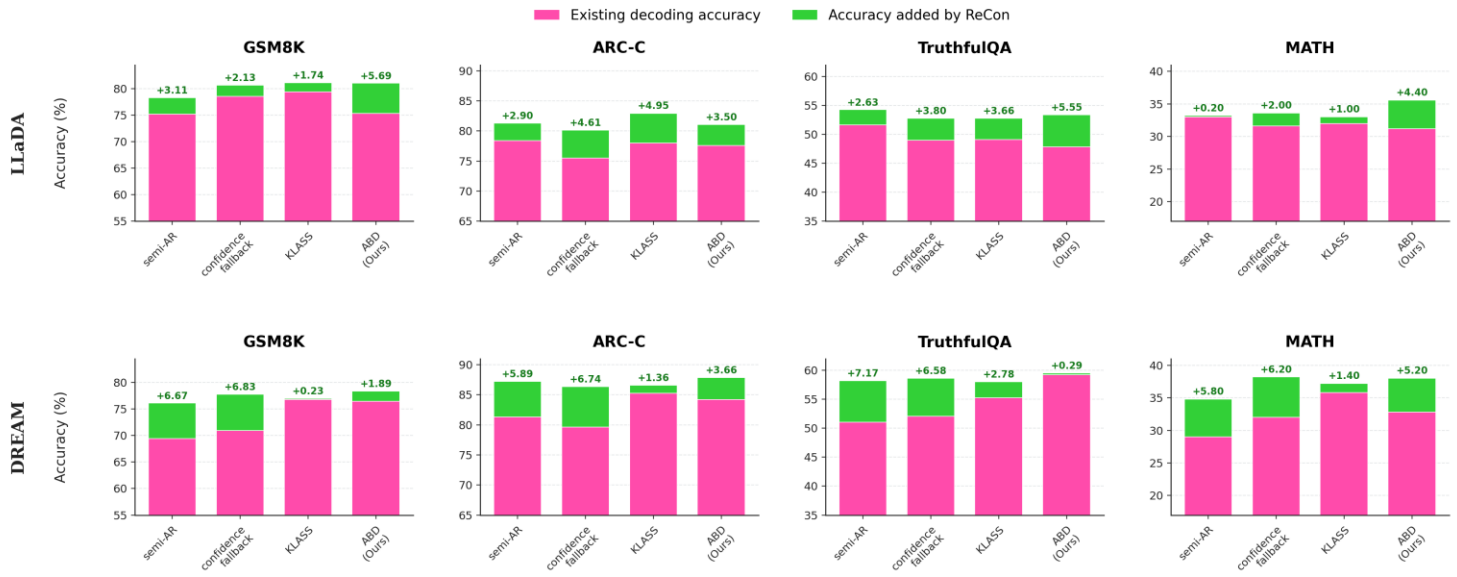
Method: Full Pipeline (ABD + ReCon)

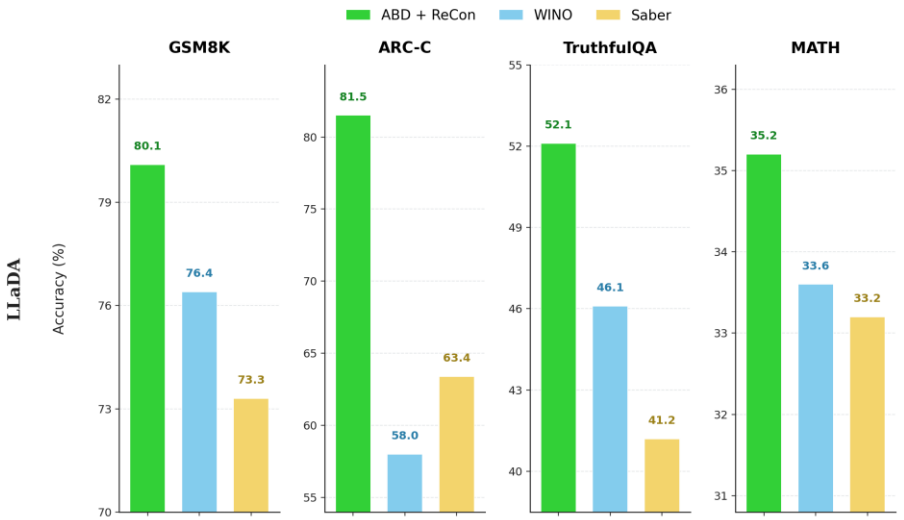


Diffusion Large Language Models (DLMs) can self-correct.
We propose a two-step training-free inference method.

1# Adaptive **B**lock **D**rafting **(ABD),** efficient drafting.
2 # Renoise **Con**sistency **(ReCon),** efficient self-correction.

Experiments: Full Pipeline (ABD + ReCon)





ReCon improves drafts.
Semi-AR: up to **7.17%p** | Conf. fallback: up to **6.83%p**
KLASS: up to **4.95%p** | ABD: up to **5.69%p**

ABD + ReCon outperforms lightweight revocable decodings.
WINO: up to **23.5%p**
Saber: up to **18.1%p**

[3] Hong et al., "Wide-In, Narrow-Out: Revokable Decoding for Efficient and Effective DLLMs," *ICLR 2026*.
[4] Dong et al., "Saber: An Efficient Sampling with Adaptive Acceleration and Backtracking Enhanced Remasking for Diffusion Language Model," *ACL 2026*.

Thank you!

{ jihoonlee98, jaehyungk }@yonsei.ac.kr

Project Page

Project Page

ML^3 Lab

Lab Website