Towards Symbolic Planning via Diffusion

Arman Mohammadi, Markus Fritzsche, Jendrik Seipp

June 28, 2026

Linköping University, Sweden

Table of contents

Autoregressive Generation

Autoregressive Generation [1]

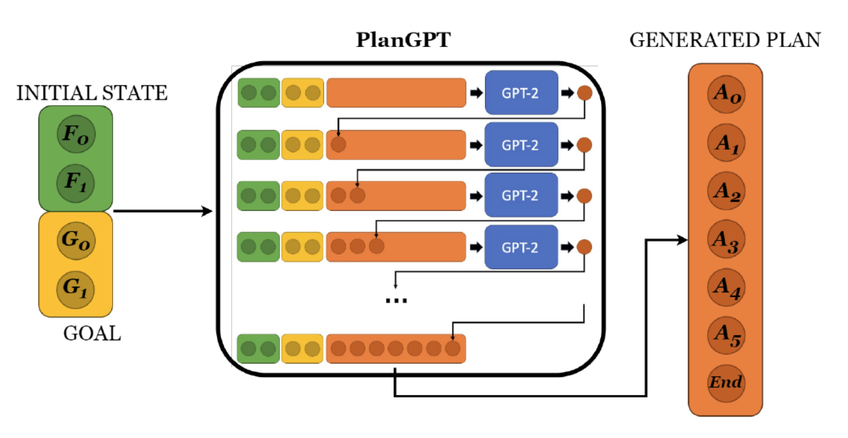

[1] Rossetti, N., *et al.* (2024). Learning General Policies for Planning through GPT Models. ICAPS 2024, pp. 500–508.

Diffusion Generation

Overview

The underlying principle of diffusion models is to progressively perturb the observed data with a forward diffusion process, and then recover the original data through a reverse process.

Normally, the reverse process is parameterized by a neural network.

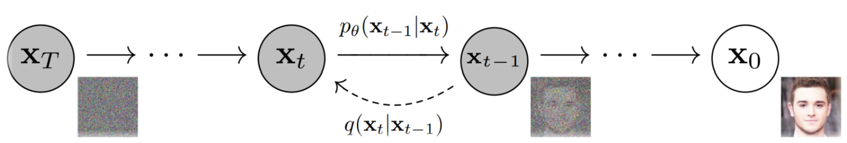

Reverse Process $\Longrightarrow$

$\Longleftarrow$ **Forward process**

Masked Diffusion Modeling

Unlike continuous diffusion models, masked diffusion modeling operates in discrete token space.

This is suited for text and other symbolic data.

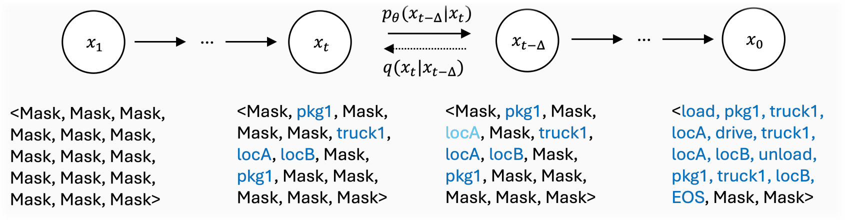

Reverse Process $\Longrightarrow$

$\Longleftarrow$ **Forward process**

Training Objective

Our model $M_\theta(x_t, c)$ is a decoder-only transformer that predicts complete action sequences conditioned on a planning problem representation.

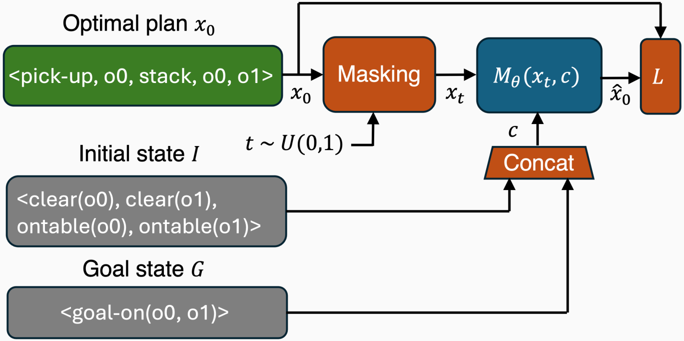

Plan Generation

To generate samples from a learned model, we follow a series of (reverse) state transitions $x_1 \to x_{1-\Delta} \to \cdots \to x_0$. At each state x_t, we apply M_θ to predict the clean plan $\hat{x}_0$.

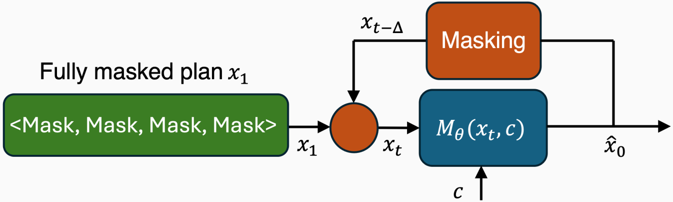

Visualization of Unmasking Process

Step	Token sequence	#MASK	#PAD
1	<BOS>, <MASK>, <MASK>, ..., <MASK>, <PAD>, <MASK>, ...	209	11
10	<BOS>, <MASK>, <MASK>, ..., <MASK>, <PAD>, <MASK>, ...	110	110
15	<BOS>, <MASK>, <MASK>, ..., <MASK>, <PAD>, <MASK>, ...	55	165
17	<BOS>, unstack, <MASK>, <MASK>, put-down, o3, <MASK>, <MASK>, <MASK>, ..., <PAD>, <MASK>, ...	33	184
18	<BOS>, unstack, o3, o1, put-down, o3, <MASK>, o1, <MASK>, <MASK>, o1, <MASK>, ..., <PAD>, <MASK>, ...	22	191
19	<BOS>, unstack, o3, o1, put-down, o3, unstack, o1, o0, <MASK>, o1, <MASK>, o0, <MASK>, stack, o0, <MASK>, <MASK>, o1, <MASK>, o1, <MASK>, pick-up, <MASK>, stack, <MASK>, <MASK>, <EOS>, ...	11	192
20	<BOS>, unstack, o3, o1, put-down, o3, unstack, o1, o0, put-down, o1, unstack, o0, o2, stack, o0, o3, pick-up, o1, stack, o1, o0, pick-up, o2, stack, o2, o1, <EOS>, <PAD>, ...	0	193

Results

Domain	Dataset	PlanGPT [2]		SymT [2]		Diffusion	
		Coverage ($\mu \pm \sigma$)	Quality ($\mu \pm \sigma$)	Coverage ($\mu \pm \sigma$)	Quality ($\mu \pm \sigma$)	Coverage ($\mu \pm \sigma$)	Quality ($\mu \pm \sigma$)
Blocks	validation	$.00 \pm .00$	$.00 \pm .00$	$1.00 \pm .00$	$1.00 \pm .00$	$.11 \pm .19$	$.11 \pm .19$
	interpolation	$.56 \pm .16$	$.56 \pm .16$	$1.00 \pm .00$	$1.00 \pm .00$	$.56 \pm .51$	$.56 \pm .51$
Gripper	validation	$.00 \pm .00$	$.00 \pm .00$	$.17 \pm .24$	$.17 \pm .24$	$.00 \pm .00$	$.00 \pm .00$
	interpolation	$.00 \pm .00$	$.00 \pm .00$	$.67 \pm .00$	$.67 \pm .00$	$.00 \pm .00$	$.00 \pm .00$
Visitall	validation	$.00 \pm .00$	$.00 \pm .00$	$.33 \pm .09$	$.32 \pm .10$	$.00 \pm .00$	$.00 \pm .00$
	interpolation	$.05 \pm .04$	$.05 \pm .04$	$.87 \pm .01$	$.87 \pm .02$	$.11 \pm .00$	$.11 \pm .00$
Logistics	validation	$.00 \pm .00$	$.00 \pm .00$	$.00 \pm .00$	$.00 \pm .00$	$.00 \pm .00$	$.00 \pm .00$
	interpolation	$.07 \pm .05$	$.07 \pm .05$	$.22 \pm .31$	$.22 \pm .31$	$.00 \pm .00$	$.00 \pm .00$

[2] Fritzsche, M., Gestrin, E., & Seipp, J. (2026). Symmetry-Aware Transformer Training for Automated Planning. AAAI 2026.

Inference Step Effect

Coverage and inference time as a function of the number of diffusion inference steps for the Blocksworld domain.

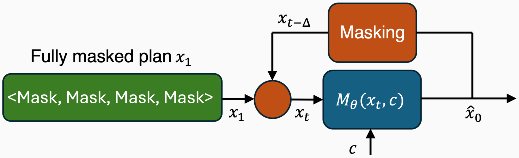

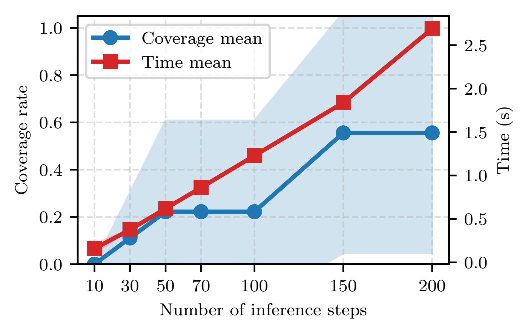

Future Research

Symmetry-Aware Training [2]

We can incorporate symmetry-aware constraints during training so the model explicitly accounts for these invariances.

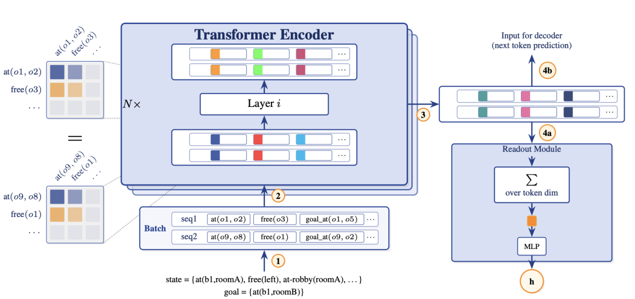

[2] Fritzsche, M., Gestrin, E., & Seipp, J. (2026). Symmetry-Aware Transformer Training for Automated Planning. AAAI 2026.

Curriculum-Based Masking

A curriculum masking schedule: gradually shifting the masking ratio distribution toward higher masking.

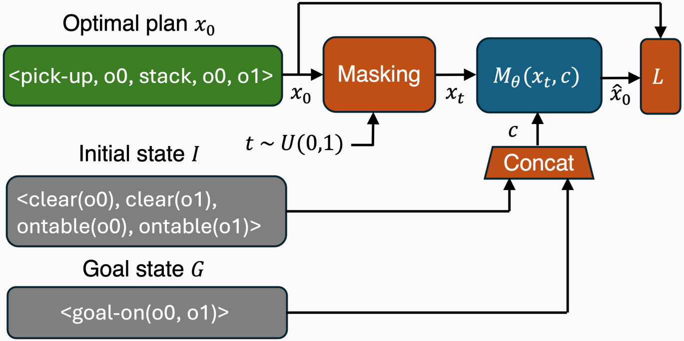

Diffusion in Embedding Space

Rather than masking discrete tokens, one can apply Gaussian noise directly in a semantic embedding space, closer to the original vision diffusion formulation.

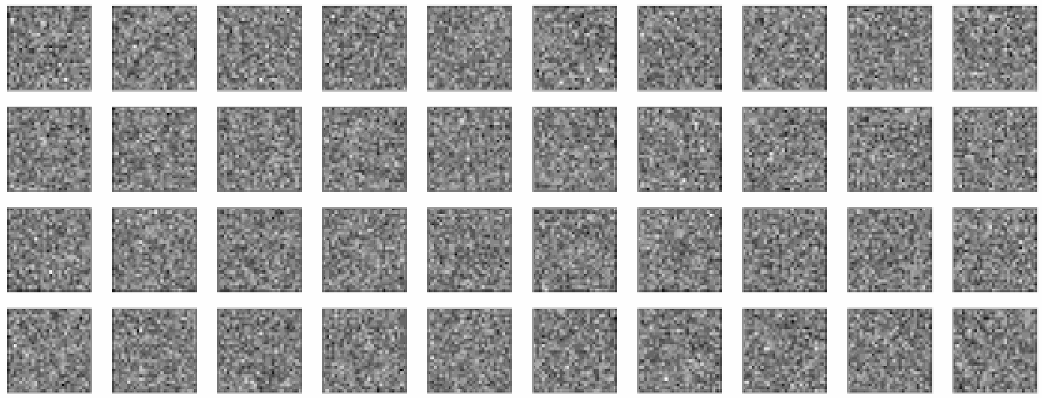

N. Rossetti, M. Tummolo, A. E. Gerevini, L. Putelli, I. Serina, M. Chiari, and M. Olivato, “Learning general policies for planning through GPT models,” in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 34, 2024, pp. 500–508.

M. Fritzsche, E. Gestrin, and J. Seipp, “Symmetry-aware transformer training for automated planning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026, pp. 36 236–36 244.

Supplementary: Generative Modeling Summary

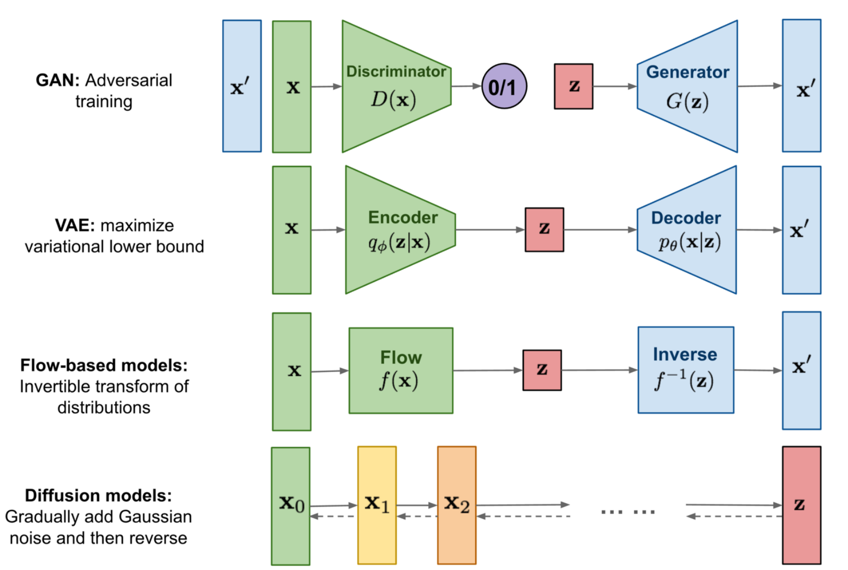

Supplementary: Training Loss

$$L = \min_{\theta} \ \mathbb{E}_{x_0, t, x_t} [\frac{1}{t} \sum_{i \in m} \ \ell_{\mathrm{CE}}(M_\theta(x_t, c)_i, \ x_{0,i})]$$

Visualization of Unmasking Process for Blocks

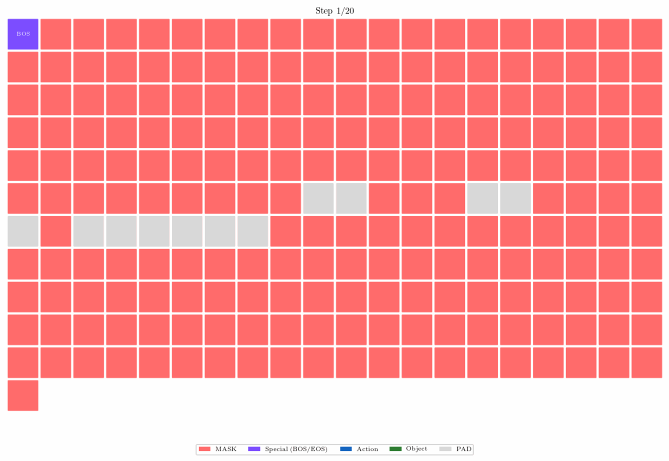

Visualization of Unmasking Process for Gripper with 8 Balls

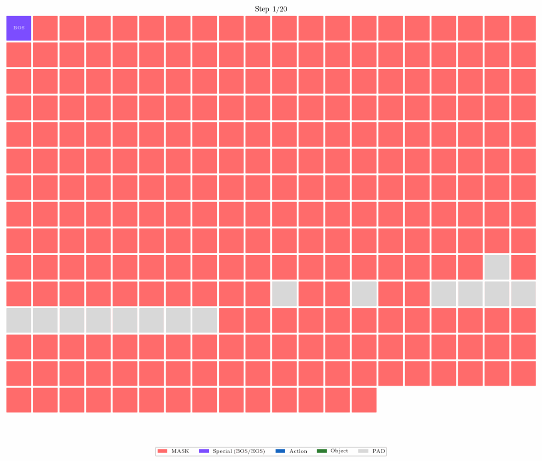

Benchmark Problem Sizes

Domain	Parameter	#Training	Training Sizes	Validation Sizes	Interpolation Sizes
Blocksworld	#blocks	9	4, 6, 7	8	5
Gripper	#balls	4	2, 4, 6, 8	9, 10	3, 5, 7
Logistics	#goals	12	1, 3, 5	6	2, 4
Visitall	#cells	207	1, 3, 4, 6, 10, 11, 12, 14, 16	18, 20	2, 5, 8, 9, 15