

Algorithm 1 Adaptive Classifier-Free Guidance (A-CFG) for one generation step k

```

1: Input: Current sequence  $\mathbf{x}^{(k)}$ , conditioning  $c$ , model  $M_\theta$ , guidance  $w$ , re-masking proportion  $\rho$ .
2: Output: Guided logits  $L_{\text{guided}}^{(k)}$ .
3:  $L_{\text{cond}}^{(k)} \leftarrow M_\theta(\mathbf{x}^{(k)})$  ▷ Compute conditional logits
4:  $\mathcal{C}_{\text{remaskable}}^{(k)} \leftarrow \{j \mid (\mathbf{x}^{(k)})_j \neq [\text{MASK}]\}$  ▷ Identify all non- [MASK] token indices
5:  $\mathcal{CONF}^{(k)} \leftarrow \emptyset$ 
6: for  $j \in \mathcal{C}_{\text{remaskable}}^{(k)}$  do
7:    $c_j^{(k)} \leftarrow \max_v (\text{softmax}(L_{\text{cond}}^{(k)}))_{j,v}$  ▷ Assess confidence for remaskable tokens
8:   Add  $(c_j^{(k)}, j)$  to  $\mathcal{CONF}^{(k)}$ 
9: end for
10:  $\mathcal{S}_{\text{low-conf}}^{(k)} \leftarrow \emptyset$ 
11: if  $|\mathcal{C}_{\text{remaskable}}^{(k)}| > 0$  then
12:    $N_m^{\text{target}} \leftarrow \lceil \rho \cdot |\mathcal{C}_{\text{remaskable}}^{(k)}| \rceil$ 
13:    $N_m^{\text{actual}} \leftarrow \min(N_m^{\text{target}}, |\mathcal{C}_{\text{remaskable}}^{(k)}|)$ 
14:   if  $N_m^{\text{actual}} > 0$  then
15:     Sort  $\mathcal{CONF}^{(k)}$  by confidence values  $c_j^{(k)}$  in ascending order.
16:      $\mathcal{S}_{\text{low-conf}}^{(k)} \leftarrow$  indices  $j$  of the first  $N_m^{\text{actual}}$  elements in sorted  $\mathcal{CONF}^{(k)}$ .
17:   end if
18: end if
19:  $\mathbf{x}_{\text{uncond}}^{(k)} \leftarrow \mathbf{x}^{(k)}$ 
20: for  $j \in \mathcal{S}_{\text{low-conf}}^{(k)}$  do
21:    $(\mathbf{x}_{\text{uncond}}^{(k)})_j \leftarrow [\text{MASK}]$  ▷ Create dynamic unconditional input
22: end for
23:  $L_{\text{uncond}}^{(k)} \leftarrow M_\theta(\mathbf{x}_{\text{uncond}}^{(k)})$  ▷ Compute unconditional logits
24:  $L_{\text{guided}}^{(k)} \leftarrow L_{\text{uncond}}^{(k)} + (w + 1) \cdot (L_{\text{cond}}^{(k)} - L_{\text{uncond}}^{(k)})$  ▷ Apply CFG formula
25: return  $L_{\text{guided}}^{(k)}$ 

```

B Standard Benchmarks and Evaluation

We follow widely-adopted large-language-model (LLM) practices and group tasks into three thematic suites:

General ability. MMLU, ARC-C, Hellaswag, TruthfulQA, WinoGrande and PIQA provide broad coverage of factual recall, commonsense reasoning and causal inference.

Mathematical & scientific reasoning. GSM8K and GPQA probe arithmetic, symbolic manipulation and graduate-level science questions.

Planning ability. Countdown and Sudoku benchmarks assess long-horizon symbolic planning under strict, multi-constraint settings. Following Dream’s evaluation protocol, we employ few-shot prompts and report the exact-solution rate across graded difficulty levels, thereby stressing the model’s capacity to generate coherent step-wise action sequences.

Evaluation mode. Closed-form tasks supply a prompt with a finite set of candidate answers; we compute each candidate’s conditional log-likelihood and select the most likely. Open-ended tasks require free-form generation; we sample responses and score them with task-specific metrics such as exact-match accuracy.

Likelihood estimation. For likelihood-based evaluations we approximate the conditional perplexity bound with Monte-Carlo sampling. A single sample suffices when only one target token is queried

447 (e.g. MMLU). We adopt the same setting as LLaDA, for all other multiple-token tasks we draw 128
448 samples, which we found to stabilise variance without adding prohibitive cost.

449 **Generation hyper-parameters.** Unless otherwise stated, we set the answer length to 256 tokens and
450 run the reverse diffusion process for 256 steps (one token revealed per step).