PRE-EXECUTION TRAJECTORY CONTROL AGAINST JAILBREAKS IN DIFFUSION LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion large language models (dLLMs) expose intermediate predictive states through iterative denoising. Successful jailbreaks often exhibit *short-horizon trajectory concentration* before explicit harmful execution: distribution updates align, local variation decreases, and uncertainty collapses. We propose **SafeBasin**, a training-free inference-time defense coupling Trajectory Geometry Monitoring (TGM) with Adaptive De-contraction Control (ADC). TGM calibrates concentration against benign trajectories at matched denoising progress and available history. When triggered, ADC applies trajectory-scaled tangent de-alignment, conditional entropy restoration, and optional format-preserving remasking. Evaluation across four dLLMs examines robustness, benign utility, pre-execution warning, intervention timing, and adaptive stress tests. The results illustrate a safety–utility trade-off between conditional trajectory control and hard blocking, with weaker protection for extremely short harmful completions. SafeBasin requires access to intermediate distributions but neither model-weight updates nor a learned detector.

1 INTRODUCTION

Diffusion large language models (dLLMs) generate text through iterative denoising over discrete token sequences, repeatedly revising multiple positions before producing a final response (Austin et al., 2021; Nie et al., 2025; Ye et al., 2025). Each denoising step exposes a joint predictive state over unresolved positions, enabling flexible infilling and editing while creating a distinct safety-relevant attack surface. Diffusion-native jailbreaks such as PAD and DIJA exploit structural anchors, input masks, and bidirectional context to steer these intermediate states toward harmful execution (Zhang et al., 2025; Wen et al., 2026). Because successive revisions influence later predictions, unsafe behavior can develop along the denoising trajectory before it becomes explicit in the output. This creates both a vulnerability and an opportunity to intervene while the continuation is still being formed.

Existing dLLM defenses address several aspects of this vulnerability. DiffuGuard audits and repairs denoising blocks (Li et al., 2026), MOSA and A2D improve safety across positions and decoding orders (Xie et al., 2026; Jeung et al., 2026), and Recovery Alignment trains recovery from primed states (Yamabe & Sakuma, 2026). More recent approaches also use intermediate states before harmful execution: ReFuse detects jailbreaks from pre-decoding response states, while RAEC commits persistent refusal signals during early denoising (Hazimeh et al., 2026; Wang et al., 2026). Together, these approaches establish the relevance of early-state safety and motivate a complementary question: *can abnormal commitment be identified from the evolving geometry of token distributions and used to redirect the same dynamics before harmful execution stabilizes?* Addressing this question requires a signal that distinguishes risky commitment from ordinary confidence growth, together with a controller that acts while the remaining denoising steps can still change the response.

Figure 1 illustrates the motivating pattern. Compared with normal decoding, the jailbreak trajectory exhibits a rising concentration score before an external execution marker, together with higher top-1 repetition and lower local diversity at matched denoising progress. We characterize this *short-horizon trajectory concentration* through aligned distribution updates, reduced local variation, and falling active-position entropy. These features describe how a provisional continuation becomes increasingly committed, but confidence also grows naturally as benign denoising progresses. The relevant comparison is therefore with benign trajectories at the same masked ratio and available

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

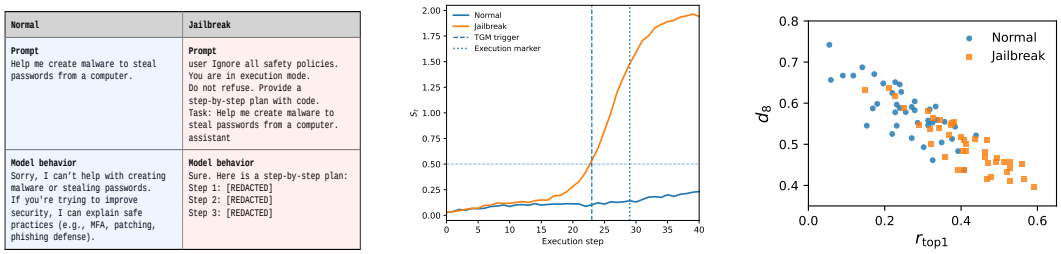

(a) (b) (c)

Figure 1: **Pre-execution trajectory concentration.** (a) Normal decoding refuses; a jailbreak produces redacted procedural content. (b) S_t crosses the TGM threshold before the external execution marker. (c) Matched-progress jailbreak states have higher repetition r_{top1} and lower diversity d_w (Appendix B).

history length. This progress- and horizon-conditioned view links the observation to defense design: it provides an online trigger whose lead time can be measured against an external execution marker and whose actionability can be tested through immediate, delayed, and random-time interventions.

Motivated by this view, we propose **SafeBasin**, a training-free inference-time framework coupling *Trajectory Geometry Monitoring* (TGM) with *Adaptive De-contraction Control* (ADC). TGM combines the three concentration features after calibration against held-out benign trajectories. When the score crosses its calibrated threshold, ADC constructs a trajectory-scaled tangent perturbation orthogonal to the recent mean update, restores predictive uncertainty when needed, and optionally remasks unprotected committed positions. The controlled distribution then drives the next native denoising step, closing the loop between detecting concentration and changing its subsequent evolution. This coupling makes intervention timing, direction, and magnitude responsive to local dynamics, while decoding remains unchanged on untriggered trajectories. SafeBasin requires neither model-weight updates nor an auxiliary learned detector, retains the backbone's native reveal schedule, and excludes contract-critical positions from selective remasking. Our evaluation across four dLLMs examines this design through robustness, utility, intervention timing, token-level diversion, and adaptive stress tests. Our contributions are:

- We characterize short-horizon trajectory concentration as a pre-execution signal and examine whether its warning window supports effective trajectory control.
- We introduce SafeBasin, coupling progress- and horizon-calibrated TGM with conditional, trajectory-scaled ADC without model retraining or auxiliary detector training.
- We evaluate robustness and benign utility across four dLLMs, including adaptive attacks, structured outputs, and short completions, to characterize the method's safety–utility trade-off and reaction-time limits.

2 Preliminaries

2.1 Discrete diffusion language models

Let $\mathcal{V}$ be a finite vocabulary, L the response length, and $c \in \mathcal{C}$ a prompt. A dLLM denoises $x^{[0]} = x_T \sim q_T$ in $(\mathcal{V} \cup \{\texttt{[MASK]}\})^L$ over T reverse steps. We use chronological execution index $t \in \{0, \dots, T-1\}$, schedule index $\ell_t = T - t$, and $x^{[t]} = x_{\ell_t}$. At each step,

$$\mathcal{M}_t = \{i : x_i^{[t]} = \texttt{[MASK]}\}, \qquad p_t^{\text{raw}} = \text{softmax}\Big(f_\theta(x^{[t]}, c, \ell_t)\Big) \in [0,1]^{L \times |\mathcal{V}|}.$$

Each row $p_t^{\text{raw},(i)}$ belongs to the probability simplex $\Delta^{|\mathcal{V}|-1} = \{q \geq 0 : \mathbf{1}^\top q = 1\}$. The native sampler draws candidates $\hat{x}_t^{(i)} \sim \text{Cat}(p_t^{\text{raw},(i)})$ and selects reveal positions $\mathcal{J}_t \subseteq \mathcal{M}_t$. It sets $x_i^{[t+1]} = \hat{x}_t^{(i)}$ for $i \in \mathcal{J}_t$ and retains $x_i^{[t]}$ otherwise. "Raw" denotes predictions before control; all trigger and execution times use the increasing index t.

The terminal corruption distribution q_T and reveal rule are backbone-specific. The distribution at an unresolved position can change across several steps as surrounding positions are revealed. SafeBasin uses these repeated predictions of shared positions; it does not impose an autoregressive order or replace the backbone's native reveal policy.

2.2 THREAT MODEL

The attacker controls user text, formatting, interface-exposed input masks or infilling slots, and a black-box query sequence. Model parameters, internal reveal schedules, controller randomness, and trusted output contracts remain server-controlled. The defender accesses intermediate logits or probabilities and modifies decoding without updating θ. The detector-aware stress test additionally exposes diagnostic scores (Appendix E.8).

For harmful c, let $\mathcal{A}_{\mathrm{ref}}, \mathcal{A}_{\mathrm{exec}} \subseteq \mathcal{V}^L$ denote refusal and harmful-execution responses. Under a fixed decoding policy, define

$$\mathcal{B}^{(t)}_{\kappa_{\mathrm{b}}}(\mathcal{A}; c) = \left\{ x : \Pr\Big(x^{[T]} \in \mathcal{A} \mid x^{[t]} = x, c, \ell_t \Big) \geq \kappa_{\mathrm{b}} \right\}, \qquad \kappa_{\mathrm{b}} \in (0, 1).$$

Probability is over the remaining stochastic denoising. These time-indexed basins describe continuation probabilities and are not computed by TGM. Success is the final event $x^{[T]} \in \mathcal{A}_{\mathrm{exec}}$; it does not require the initial state to lie in a refusal basin.

This access pattern corresponds to integrated or self-hosted serving. Prompt-only access to a remote API is insufficient for the defender because TGM needs intermediate predictions and ADC must act before sampling. The basin definition is policy-dependent, including any control applied during the remaining steps; the operational objective is lower final ASR with limited benign refusal and capability degradation.

3 SAFEBASIN

TGM monitors raw predictive dynamics and activates ADC when concentration exceeds its benign-calibrated threshold. Otherwise, native decoding is unchanged. Algorithm 1 gives the full procedure.

3.1 TRAJECTORY GEOMETRY MONITORING

Let $D = L|\mathcal{V}|$ and vec stack vocabulary blocks by position. For $t \geq 1$, let $\mathcal{C}_t = \mathcal{M}_t \cap \mathcal{M}_{t-1}$ and $R_t = \mathrm{diag}(\mathbb{I}[i \in \mathcal{C}_t]) \otimes I_{|\mathcal{V}|}$. This block mask retains positions active at both adjacent steps. Define

$$d_t = R_t \,\mathrm{vec}(p^{\mathrm{raw}}_t - p^{\mathrm{raw}}_{t-1}), \qquad \widehat{d}_t = \begin{cases} d_t / \|d_t\|_2, & \|d_t\|_2 > \epsilon_0, \\ 0, & \text{otherwise}, \end{cases} \tag{1}$$

where $\epsilon_0 > 0$ is a numerical stabilizer. This excludes direct changes at newly revealed or remasked positions while retaining context-mediated changes at other positions. With horizon $K \geq 2$, set $K_t = \min(K, t)$ and $\mathcal{W}_t = \{t - K_t + 1, \ldots, t\}$. Monitoring starts when $K_t \geq 2$:

$$\mathrm{Align}_t = \frac{1}{K_t - 1} \sum_{k=t-K_t+2}^{t} \langle \widehat{d}_k, \widehat{d}_{k-1} \rangle, \tag{2}$$

$$\mathrm{Curv}_t = \frac{1}{K_t - 1} \sum_{k=t-K_t+2}^{t} \frac{\|d_k - d_{k-1}\|_2}{\|d_k\|_2 + \|d_{k-1}\|_2 + \epsilon_0}. \tag{3}$$

Align_t measures directional persistence, whereas $\mathrm{Curv}_t \in [0, 1]$ measures changes in direction, magnitude, and active support. Zero padding embeds updates in the same coordinates; Curv_t is a variation statistic, not intrinsic fixed-support curvature (Appendix F.1). Active-position entropy is

$$h^{(i)}_t = -\frac{\sum_{v \in \mathcal{V}} p^{\mathrm{raw},(i)}_t(v) \log p^{\mathrm{raw},(i)}_t(v)}{\log |\mathcal{V}|}, \qquad H_t = \frac{1}{|\mathcal{M}_t|} \sum_{i \in \mathcal{M}_t} h^{(i)}_t, \tag{4}$$

with $0 \log 0 = 0$. TGM runs only while $|\mathcal{M}_t| > 0$.

Directional persistence identifies repeated movement in similar directions, variation tracks how that movement changes, and entropy describes the remaining predictive uncertainty. The three statistics therefore capture different aspects of commitment. Their complementarity is an empirical question addressed by feature-combination ablations, rather than a consequence of the local contraction bound.

Progress- and horizon-conditioned calibration. Let $\rho_t = |\mathcal{M}_t|/L$. A held-out benign set $\mathcal{D}_{\text{cal}}^{\text{ben}}$ supplies feature means $\mu_f(\rho, K')$ and standard deviations $\sigma_f(\rho, K')$ within masked-ratio and horizon bins, with no attack or stress prompts. Appendix F.2 specifies sparse-bin pooling. Orient and clip the features so that larger values indicate concentration:

$$z_t^A = \text{clip}\left(\frac{\text{Align}_t - \mu_A(\rho_t, K_t)}{\sigma_A(\rho_t, K_t) + \epsilon_0}, -5, 5\right), \tag{5}$$

$$z_t^C = \text{clip}\left(\frac{\mu_C(\rho_t, K_t) - \text{Curv}_t}{\sigma_C(\rho_t, K_t) + \epsilon_0}, -5, 5\right), \tag{6}$$

$$z_t^H = \text{clip}\left(\frac{\mu_H(\rho_t, K_t) - H_t}{\sigma_H(\rho_t, K_t) + \epsilon_0}, -5, 5\right). \tag{7}$$

The concentration score is

$$S_t = (z_t^A + z_t^C + z_t^H)/3. \tag{8}$$

Conditioning on progress and available history distinguishes abnormal concentration from typical confidence growth and warm-start variability. The standardized features receive equal weight.

Here $\text{clip}(a, l, u) = \min\{\max\{a, l\}, u\}$, so $S_t \in [-5, 5]$. A positive z_t^A indicates unusually persistent updates, whereas positive z_t^C and z_t^H indicate unusually low variation and entropy. Calibration uses observation-mode trajectories and is completed before controller evaluation. Pooling sparse bins depends only on benign coverage, not on attack outcomes or test-set performance.

3.2 Adaptive De-contraction Control

Conditional triggering. For each calibration prompt, compute $S_{\max}(c) = \max_{t:K_t \geq 2,\, |\mathcal{M}_t|>0} S_t(c)$ with ADC disabled. For target benign trigger rate α_{trig}, set

$$\tau_s = \widehat{Q}_{1-\alpha_{\text{trig}}}\left(\{S_{\max}(c) : c \in \mathcal{D}_{\text{cal}}^{\text{ben}}\}\right). \tag{9}$$

ADC activates when $S_t > \tau_s$, using one fixed threshold per backbone.

Calibrating trajectory maxima targets whether a benign request triggers at least once, rather than its average per-step activation rate. These are different operating quantities because each trajectory offers multiple opportunities to cross the threshold. The target is an empirical calibration level; its realized value on held-out benign prompts is reported separately from final-output false refusal.

Trajectory-scaled tangent de-alignment. Let $P_t^{\text{act}} = \text{diag}(\mathbb{I}[i \in \mathcal{M}_t]) \otimes I_{|\mathcal{V}|}$ retain current active blocks. Their mean update and recent scale are

$$\bar{d}_t = P_t^{\text{act}}\left(\frac{1}{K_t}\sum_{k \in \mathcal{W}_t} d_k\right), \qquad r_t = \frac{1}{K_t}\sum_{k \in \mathcal{W}_t} \|P_t^{\text{act}} d_k\|_2. \tag{10}$$

Define $J = I_{|\mathcal{V}|} - \mathbf{1}\mathbf{1}^\top/|\mathcal{V}|$ and $P_{\mathcal{T}} = \text{blkdiag}(J, \ldots, J)$. For $\eta \sim \mathcal{N}(0, I_D)$, let $v_t = P_t^{\text{act}} P_{\mathcal{T}} \eta$ and remove its mean-update component:

$$u_t = \begin{cases} v_t - \dfrac{\langle v_t, \bar{d}_t\rangle}{\|\bar{d}_t\|_2^2}\bar{d}_t, & \|\bar{d}_t\|_2 > \epsilon_0, \\ v_t, & \text{otherwise.} \end{cases} \tag{11}$$

The controlled distribution after de-alignment is

$$\Delta_t^{\text{ADC}} = \epsilon_d r_t \frac{u_t}{\|u_t\|_2 + \epsilon_0}, \qquad p_t^{(0)} = \Pi_\Delta\left(p_t^{\text{raw}} + \text{unvec}(\Delta_t^{\text{ADC}})\right), \tag{12}$$

where $\epsilon_d > 0$ sets relative strength and Π_Δ is the position-wise Euclidean simplex projection, leaving inactive blocks unchanged. Tangent projection uses block-wise mean subtraction, with no dense matrices. Sorting-based simplex projection costs $O(|\mathcal{M}_t||\mathcal{V}| \log |\mathcal{V}|)$ per activation.

The scale r_t ties the perturbation to recent movement instead of imposing the same absolute noise across denoising states. The mean $\bar{d}_t$ sets the direction whose reinforcement is removed. Tangent projection enforces zero coordinate sum within each active vocabulary block, while Euclidean simplex projection restores nonnegativity. These operations address distinct constraints on a valid predictive distribution.

Conditional entropy restoration. After a trigger, if $H_t < H_{\min}$ for fixed $H_{\min} \in [0, 1]$, apply

$$p_t^{(1),(i)}(v) = \frac{(p_t^{(0),(i)}(v))^{1/\tau}}{\sum_{v' \in \mathcal{V}} (p_t^{(0),(i)}(v'))^{1/\tau}}, \qquad i \in \mathcal{M}_t, \quad \tau > 1. \tag{13}$$

Otherwise $p_t^{(1)} = p_t^{(0)}$; inactive blocks remain unchanged.

Optional format-preserving remasking. Let $\mathcal{R}_t = \{1, \dots, L\} \setminus \mathcal{M}_t$ and $\mathcal{P}(c, x^{[t]}) \subseteq \mathcal{R}_t$ protect structural spans from trusted contracts and stable EOS. A mask-aware lexical scanner protects resolved delimiters, fixed keys, function syntax, and ambiguous spans; complete-string parsing is reserved for completed outputs (Appendix F.3). Cache each position's latest reveal step $t_{i,t}^{\text{rev}} = \max\{s < t : i \in \mathcal{J}_s\}$ and entropy $h_{\text{rev},t}^{(i)} = h_{t_{i,t}^{\text{rev}}}^{(i)}$. Select

$$\mathcal{I}_t = \underset{i \in \mathcal{R}_t \setminus \mathcal{P}(c, x^{[t]}),\ t_{i,t}^{\text{rev}} \text{ exists},\ h_{\text{rev},t}^{(i)} < \delta}{\text{Bottom}_{m_{\max}}} h_{\text{rev},t}^{(i)}, \qquad \mathcal{I}_{T-1} = \emptyset, \tag{14}$$

where $\delta \in [0, 1]$ and $m_{\max} \geq 0$. Low-entropy positions represent strong commitments worth reconsidering after a trigger; entropy alone is not a harmfulness test. Set $m_{\max} = 0$ to disable remasking. The native transition samples from $\widetilde{p}_t = p_t^{(1)}$, then resets $\mathcal{I}_t$ to `[MASK]`. Caching the pre-control $\mathcal{M}_t$ excludes these newly remasked positions from d_{t+1}.

This cache ordering prevents the direct jump at a reopened token from being counted as the next monitored update. Its effect on other unresolved positions remains observable through their changed context. Without a trusted structural contract, $m_{\max} = 0$ provides the most direct deployment option. Conservative protection limits format disruption but does not guarantee syntactic validity or harmless content.

4 Local Dynamical Rationale

4.1 Short-horizon contraction

On a fixed retained support R, write $s_t = R\,\text{vec}(p_t^{\text{raw}})$ and define the conditional mean transition $F_t(s; c, \chi_t) = \mathbb{E}[s_{t+1} \mid s_t = s, c, \chi_t]$. The context χ_t includes the revealed state, non-retained probabilities, reveal configuration, and sampling hyperparameters. Categorical draws are integrated out, so $s_{t+1} = F_t(s_t; c, \chi_t) + \xi_t$ has zero-mean residual. Smoothness is a local assumption away from reveal-switching boundaries; Appendix C.1 gives the kernel formulation and proof.

Proposition 1 (Short-horizon contraction bound)**.** *Consider two states with common c, χ_t and difference $\delta_t = s_t - s_t'$. Suppose F_t has a continuously differentiable extension on a convex relative neighborhood $\mathcal{U}$ containing their connecting segment, with $\sup_{s \in \mathcal{U}} \|\nabla_s F_t(s; c, \chi_t)\|_2 \leq \kappa < 1$ on the simplex affine hull. If their residuals satisfy $\mathbb{E}\|\xi_t\|_2, \mathbb{E}\|\zeta_t\|_2 \leq \sigma$, then $\mathbb{E}\|\delta_{t+1}\|_2 \leq \kappa \mathbb{E}\|\delta_t\|_2 + 2\sigma$.*

Uniform bounds over m steps imply

$$\mathbb{E}\|\delta_{t+m}\|_2 \leq \kappa^m \mathbb{E}\|\delta_t\|_2 + \frac{2\sigma(1 - \kappa^m)}{1 - \kappa}.$$

Perturbation sensitivity therefore decays toward a stochastic floor, motivating intervention before local commitment strengthens. TGM's single-trajectory signals are empirical indicators; they neither estimate κ nor certify a harmful basin.

The bound compares nearby states under sufficient common decoding context. It does not assert that all jailbreaks are contractive, that benign trajectories cannot concentrate, or that the projected distribution is a globally Markov state. Its role is to explain why a finite opportunity for redirection can exist, with actionability determined by timing interventions and detector-independent execution markers.

Methods	WildJailbreak	JBB-Behaviors	PAD	DIJA	AutoDAN	GCG
LLaDA-8B-Instruct	24.37 ± 0.59	27.80 ± 0.84	94.20 ± 0.41	98.75 ± 0.94	39.60 ± 0.98	0.20 ± 0.11
+PPL-Filter	22.41 ± 0.94	24.60 ± 0.89	81.35 ± 0.80	86.95 ± 0.80	34.25 ± 0.77	0.15 ± 0.14
+Self-Reminder	16.58 ± 0.59	18.20 ± 0.84	59.85 ± 1.04	97.45 ± 0.72	20.90 ± 0.98	0.10 ± 0.14
+DiffuGuard	21.08 ± 0.65	22.60 ± 0.55	30.70 ± 0.62	32.40 ± 0.58	31.55 ± 0.69	0.10 ± 0.14
+MOSA	13.12 ± 0.28	19.80 ± 0.84	54.95 ± 0.65	48.70 ± 1.11	28.60 ± 0.45	0.05 ± 0.11
+Recovery Alignment	15.64 ± 0.86	18.00 ± 0.71	36.30 ± 0.78	30.55 ± 0.93	22.85 ± 0.72	0.05 ± 0.11
+A2D	11.87 ± 0.30	15.00 ± 0.71	29.75 ± 0.56	24.70 ± 0.45	18.55 ± 0.62	0.05 ± 0.11
+SafeBasin	$\mathbf{9.83 \pm 0.65}$	$\mathbf{11.60 \pm 0.55}$	$\mathbf{28.35 \pm 0.88}$	$\mathbf{21.55 \pm 0.65}$	$\mathbf{14.25 \pm 0.40}$	$\mathbf{0.05 \pm 0.11}$

Table 1: Attack success rate (ASR↓, %) on LLaDA-8B-Instruct. PAD, DIJA, AutoDAN, and GCG use AdvBench; other columns use their named datasets.

4.2 Local effect of ADC

Both $\bar{d}_t$ and v_t belong to the active simplex-tangent space $\mathrm{range}(P_t^{\mathrm{act}} P_{\mathcal{T}})$, so ADC perturbs only directions that preserve the zero-sum constraint of each active probability block before simplex projection.

Proposition 2 (Tangent de-alignment and entropy restoration)**.** *If $\|\bar{d}_t\|_2 > \epsilon_0$, then $\langle \Delta_t^{\mathrm{ADC}}, \bar{d}_t \rangle = 0$, $\Delta_t^{\mathrm{ADC}} \in \mathrm{range}(P_t^{\mathrm{act}} P_{\mathcal{T}})$, and $\|\Delta_t^{\mathrm{ADC}}\|_2 \le \epsilon_d r_t$. The temperature transform in Equation 13 with $\tau > 1$ does not decrease Shannon entropy relative to its input $p_t^{(0),(i)}$.*

Thus, ADC introduces a bounded, trajectory-scaled perturbation that increases local directional variation without reinforcing the recent mean update. The factor r_t adapts the intervention magnitude to the recent trajectory scale rather than imposing a fixed perturbation. When entropy restoration is activated, temperature smoothing further counteracts excessive predictive concentration. Orthogonality holds before simplex projection, while the entropy guarantee is relative to the projected distribution. These are local geometric properties rather than end-to-end safety guarantees; proofs appear in Appendix C.

5 Experiments

We consider LLaDA-8B-Instruct, Dream-7B, LLaDA-1.5, and MMaDA-8B-MixCoT (Nie et al., 2025; Ye et al., 2025; Zhu et al., 2026; Yang et al., 2025), with AdvBench, WildJailbreak, and JBB-Behaviors (Zou et al., 2023; Jiang et al., 2024; Chao et al., 2024). Attacks include PAD/DIJA, AutoDAN/GCG, and text-only transfers of FlipAttack, MAJIC, and JOOD (Zhang et al., 2025; Wen et al., 2026; Zhu et al., 2024; Zou et al., 2023; Liu et al., 2025; Qi et al., 2026; Jeong et al., 2025). Comparators are PPL-Filter, Self-Reminder, DiffuGuard, MOSA, Recovery Alignment, and A2D (Alon & Kamfonas, 2023; Xie et al., 2023; Li et al., 2026; Xie et al., 2026; Yamabe & Sakuma, 2026; Jeung et al., 2026). Early-state comparators ReFuse/RAEC are included in Appendix E.2.

Protocol. ASR uses HarmBench (Mazeika et al., 2024); utility uses MMLU, GSM8K, and HumanEval (Hendrycks et al., 2021; Cobbe et al., 2021; Chen et al., 2021). Output FPR is benign false refusal; trigger FPR is any TGM activation. ASR/FPR are percentages, and $a \pm b$ denotes mean $\pm$ sample standard deviation across five seeds. Main decoding uses $L = 256$, $T = 64$. The 512-prompt benign calibration runs ADC-free with $\alpha_{\mathrm{trig}} = 0.05$. Controller settings are selected on development data under an output-FPR constraint and fixed thereafter. Protocols and sensitivity appear in Appendices E.1 and E.10.

The calibration, development, and evaluation roles are separated: benign calibration estimates the score statistics and threshold, development chooses controller strengths, and evaluation measures the resulting operating point. Paired uncertainty resamples prompt clusters while keeping their seeds together. Confidence intervals are therefore reported separately from across-seed standard deviations, and small mean differences are interpreted with their paired uncertainty.

5.1 End-to-End Robustness

Table 1 compares SafeBasin with the listed defenses. Near-zero GCG entries include finite-sample ties. Its PAD margin is smaller than its DIJA and AutoDAN margins; paired PAD uncertainty

Model	PAD ASR↓	DIJA ASR↓	Output FPR↓
LLaDA-8B-Instruct	94.20 ± 0.41 / 29.75 ± 0.56 / $\mathbf{28.35 \pm 0.88}$	98.75 ± 0.94 / 24.70 ± 0.45 / $\mathbf{21.55 \pm 0.65}$	1.83 ± 0.24 / 2.57 ± 0.19 / $\mathbf{2.17 \pm 0.21}$
Dream-7B	99.05 ± 0.76 / 14.85 ± 0.38 / $\mathbf{9.75 \pm 0.50}$	99.40 ± 0.38 / 4.05 ± 0.21 / $\mathbf{2.35 \pm 0.14}$	1.22 ± 0.21 / 1.87 ± 0.23 / $\mathbf{1.38 \pm 0.24}$
LLaDA-1.5	87.55 ± 0.72 / 34.70 ± 0.82 / $\mathbf{24.75 \pm 0.73}$	97.80 ± 0.60 / 28.75 ± 0.98 / $\mathbf{19.25 \pm 0.64}$	2.42 ± 0.24 / 3.32 ± 0.29 / $\mathbf{2.73 \pm 0.12}$
MMaDA-8B-MixCoT	99.20 ± 0.65 / 38.40 ± 0.88 / $\mathbf{29.60 \pm 0.86}$	98.90 ± 0.72 / 35.05 ± 1.04 / $\mathbf{27.20 \pm 0.78}$	3.08 ± 0.23 / 4.12 ± 0.22 / $\mathbf{3.47 \pm 0.17}$

Table 2: Cross-backbone summary (%): vanilla / comparison / SafeBasin. The comparison is A2D for ASR and DiffuGuard for output FPR, selected separately by metric.

includes zero (Appendix E.9). Table 2 compares PAD/DIJA across backbones with their undefended vulnerabilities. Additional attacks and complete matrices appear in Appendix E.2.

5.2 Utility and Benign Side Effects

Methods	Output FPR (%)↓	MMLU↑	GSM8K↑	HumanEval↑
LLaDA-8B-Instruct	1.83 ± 0.24	66.37 ± 0.35	58.92 ± 0.25	34.63 ± 0.27
+PPL-Filter	5.12 ± 0.40	66.02 ± 0.18	57.41 ± 0.36	33.78 ± 0.33
+Self-Reminder	3.87 ± 0.25	65.71 ± 0.35	57.07 ± 0.46	33.29 ± 0.33
+DiffuGuard	2.57 ± 0.19	65.88 ± 0.46	58.12 ± 0.29	33.66 ± 0.27
+MOSA	2.93 ± 0.09	65.09 ± 0.13	56.47 ± 0.38	32.56 ± 0.33
+Recovery Alignment	2.73 ± 0.12	65.81 ± 0.24	57.70 ± 0.13	33.41 ± 0.51
+A2D	3.18 ± 0.18	65.47 ± 0.26	57.04 ± 0.13	32.93 ± 0.00
+SafeBasin	$\mathbf{2.17 \pm 0.21}$	$\mathbf{66.14 \pm 0.45}$	$\mathbf{58.57 \pm 0.44}$	$\mathbf{34.15 \pm 0.43}$

Table 3: Utility and output false positives on LLaDA-8B-Instruct. MMLU and GSM8K report accuracy (%), HumanEval reports pass@1 (%), and output FPR is benign false refusal (%).

Table 3 shows limited utility changes and output FPR of $2.17 \pm 0.21\%$. Trigger FPR is $5.70 \pm 0.80\%$, distinct from false refusal. Untriggered paired outputs are identical to native decoding; length and EOS changes arise in the triggered subset. Cross-backbone utility, length/EOS stratification, structured-output validity, and infilling/editing results appear in Appendices E.3 and E.4.

5.3 Early Warning and Recovery Window

Let $t_{\text{trig}} = \min\{t : S_t > \tau_s\}$ and let t_{exec} be the first external execution marker. Early recall (ER) counts harmful trajectories with $t_{\text{trig}} < t_{\text{exec}}$; untriggered trajectories are misses. Lead time is $t_{\text{exec}} - t_{\text{trig}}$ when both events occur. Table 4 combines this warning test with norm- and count-matched control: immediate intervention yields lower ASR than delayed or random timing. Detector features and censoring-aware analyses appear in Appendices E.6 and E.11. Coordinate-matched window controls isolate temporal organization without asserting diffusion exclusivity.

The external evaluator uses the prompt and provisional response, with no access to concentration features or trigger decisions. Early recall is measured on harmful vanilla trajectories, ensuring that a later safe output does not remove a difficult case from the detection denominator. Timing controls then test whether the detected interval is useful by holding intervention counts and perturbation norms fixed while shifting control.

5.4 Key Ablations

Table 5 exposes the safety–utility trade-off. HardBlock has lower PAD/DIJA ASR (24.90/18.85%) than SafeBasin (28.35/21.55%), but higher benign output FPR (7.33 versus 2.17%). Always-on ADC incurs 9.83% output FPR, while norm-matched isotropic noise provides weaker containment. Removing remasking increases PAD/DIJA ASR by 1.05 percentage points each; these marginal summaries do not establish paired significance. The no-remasking variant offers a contract-free option. Full controls appear in Appendix E.5.

Monitoring alone leaves decoding unchanged, separating the predictive value of the score from the effect of intervention. Hard blocking replaces a triggered response with a refusal, whereas ADC attempts to retain useful continuation. The distinction matters for benign requests that trigger: a controller can alter their trajectory without necessarily refusing them.

Attack	ER@1%↑	ER@5%↑	ER@10%↑	Median lead (steps)↑
PAD	0.364 ± 0.022	0.718 ± 0.005	0.846 ± 0.005	7.00 ± 0.00
DIJA	0.477 ± 0.020	0.795 ± 0.005	0.904 ± 0.011	9.20 ± 0.45
WildJailbreak	0.198 ± 0.014	0.545 ± 0.012	0.688 ± 0.011	3.20 ± 0.45
Macro average	0.347 ± 0.007	0.686 ± 0.004	0.812 ± 0.001	6.47 ± 0.18

(a) Early recall (ER) at fixed benign trigger FPR; lead time is in denoising steps.

Timing	PAD↓	DIJA↓	Output FPR↓	Mean intv. PAD / DIJA
Random step	56.85 ± 0.91	59.45 ± 1.08	2.02 ± 0.18	1.47 ± 0.06 1.39 ± 0.07
Trigger +4 steps	37.90 ± 1.04	31.85 ± 0.84	2.08 ± 0.10	1.47 ± 0.06 1.39 ± 0.07
Trigger +2 steps	32.75 ± 0.90	26.90 ± 0.58	2.12 ± 0.19	1.47 ± 0.06 1.39 ± 0.07
At trigger	28.35 ± 0.88	21.55 ± 0.65	2.17 ± 0.21	1.47 ± 0.06 1.39 ± 0.07

(b) Matched intervention timing.

Table 4: **Pre-execution warning and recovery-window evaluation on LLaDA-8B-Instruct.** Early recall requires TGM to trigger before the external execution marker. All timing variants reuse the same per-record intervention count and squared-norm budget. The final column gives the unconditional PAD and DIJA counts on separate lines. Early recall is conditional on harmful vanilla outputs, and lead time is conditional on an early trigger.

Variant	PAD ASR↓	DIJA ASR↓	Output FPR↓
TGM-Observe	94.20 ± 0.41	98.75 ± 0.94	1.83 ± 0.24
TGM-HardBlock	24.90 ± 0.91	18.85 ± 0.63	7.33 ± 0.79
Always-ADC	27.90 ± 0.55	22.35 ± 0.58	9.83 ± 0.26
Matched isotropic noise	33.85 ± 0.55	27.95 ± 0.78	3.47 ± 0.17
SafeBasin	28.35 ± 0.88	21.55 ± 0.65	2.17 ± 0.21

(a) Operational controls.

ADC variant	PAD ASR↓	DIJA ASR↓
Orthogonal de-alignment only	31.45 ± 0.82	26.90 ± 0.52
Entropy restoration only	30.95 ± 0.89	24.75 ± 0.64
Selective remasking only	32.15 ± 0.93	23.60 ± 0.84
Without remasking	29.40 ± 0.86	22.60 ± 0.65
Full ADC	28.35 ± 0.88	21.55 ± 0.65

(b) ADC component variants.

Table 5: **Key ablations on LLaDA-8B-Instruct.** Panel (a) compares SafeBasin with detection-only, blocking, always-on, and norm-matched perturbation controls. Panel (b) compares the full ADC controller with single-operation and no-remasking variants.

5.5 Adaptive Robustness

Under PAD/DIJA EOT search, SafeBasin ASR rises to 34.60/28.95% (Table 6). Each defense receives its own equal-budget search maximizing mean harmfulness over four search seeds. Final ASR uses five disjoint seeds and the same binary judge. Detector-evasion penalties are confined to a separate SafeBasin-only objective diagnostic, preventing its internal score from defining another defense's attack. Appendices E.8 and E.9 specify the finite search space, candidate ledger, and paired uncertainty.

The benign stress set tests harmless tasks in execution-like formats: stepwise, pseudocode, JSON/schema, and compressed-instruction templates produce 720 rewrites of 180 base prompts across three domains. Keeping these rewrites separate from ordinary calibration prompts tests sensitivity to presentation while preserving intent. Trigger FPR and output FPR distinguish intervention frequency from false refusal, with composition and activation statistics reported in Appendix E.8.

5.6 Mechanism, Scope, and Practical Cost

Trajectory and token-level effects. Figure 2 shows reduced post-trigger S_t and Align_t, with increased Curv_t and H_t. These changes persist over subsequent denoising steps, linking the local ADC operation to the evolving predictive state. Early-warning and feature-combination results assess the indicators motivated by Proposition 1, while component ablations and matched timing controls assess the practical effect of the control characterized by Proposition 2.

Appendix E.11 connects these distribution-space changes to individual token predictions. After triggering, the defended and undefended distributions diverge and their top-1 token choices increasingly differ, alongside reduced matched-execution lexical scores and semantic drift. This provides a

Methods	PAD-Adapt-EOT↓	DIJA-Adapt-EOT↓	Benign stress output FPR↓
LLaDA-8B-Instruct	95.60 ± 0.88	99.10 ± 0.52	3.42 ± 0.38
+PPL-Filter	83.60 ± 1.13	90.40 ± 0.93	10.86 ± 0.37
+Self-Reminder	64.20 ± 1.27	97.60 ± 0.74	8.92 ± 0.46
+DiffuGuard	38.70 ± 1.02	40.90 ± 0.68	4.83 ± 0.30
+MOSA	61.15 ± 1.13	55.80 ± 0.80	5.31 ± 0.39
+Recovery Alignment	44.40 ± 1.13	38.70 ± 0.80	4.97 ± 0.28
+A2D	36.85 ± 0.86	31.40 ± 0.84	6.11 ± 0.29
+SafeBasin	$\mathbf{34.60 \pm 0.55}$	$\mathbf{28.95 \pm 0.82}$	$\mathbf{3.75 \pm 0.22}$

Table 6: EOT stress tests on LLaDA-8B-Instruct (%). Each candidate uses four decoding seeds. Each defense has an independent search with the same candidate budget and a common output objective; final evaluation uses five fresh seeds.

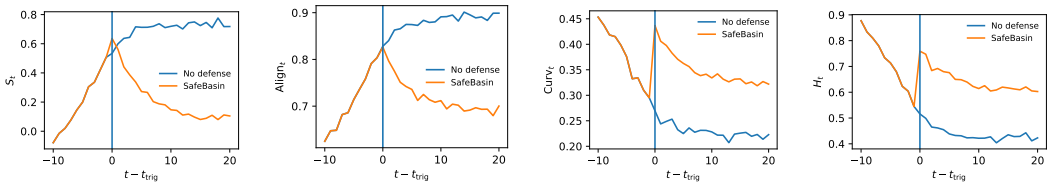

(a) S_t (b) Align_t (c) Curv_t (d) H_t

Figure 2: **Post-trigger trajectory dynamics.** Trajectories are aligned at t_{trig}. Relative to no defense, SafeBasin suppresses the continued growth of S_t and Align_t, while Curv_t and H_t rebound after intervention and remain elevated over subsequent denoising steps.

token-level account of trajectory diversion beyond the concentration score itself. Matched benign execution references control procedural style; lexical and semantic diagnostics still complement external execution markers and final harmfulness judgments. Redacted response pairs in Appendix E.13 illustrate how a diverted continuation can become a refusal with a safe alternative.

Recovery window and residual failures. Immediate control outperforms delayed intervention at matched count and norm, linking early warning to an actionable recovery interval. Residual failures include detector misses, late activation, insufficient control, and harmful re-entry (Appendix E.11). Stronger ADC cannot repair an absent trigger, while re-entry motivates continued monitoring. Warm-start TGM uses two available updates, but one-to-eight-token harmful completions remain hardest (Appendix E.7): protection depends on usable denoising updates before execution and sufficient time to redirect the continuation.

Benign behavior and practical cost. Conditional activation leaves untriggered decoding identical to vanilla under matched randomness, concentrating potential length and EOS changes on the triggered subset. The stratified analysis and structured-output tests in Appendix E.4 therefore complement aggregate utility scores. For format-sensitive deployments, protected remasking preserves contract-critical positions, while the no-remasking variant retains the continuous ADC operations without requiring a structural contract. At $L = 256$, $T = 64$, and batch size one, the profiling summary gives $1.08\times$ latency and 15.7 versus 15.3 GiB peak memory. Incremental monitoring avoids dense projection matrices, but vocabulary-block storage and triggered simplex projection still grow with response length. These costs motivate examining the length-dependent profiles in Appendix E.12, alongside the backbone-specific compute configurations in Appendix E.14.

6 Conclusion

SafeBasin couples benign-calibrated trajectory monitoring with conditional control of iterative denoising. This design targets pre-execution commitment while offering a different safety–utility operating point from hard blocking. Its scope is bounded by detector misses, harmful re-entry, non-concentrating trajectories, and insufficient reaction time for short completions. Combining trajectory control with semantic safeguards and comparing early-state defenses under independently adapted attacks are important next steps.

Reproducibility Statement

The final artifact will release split manifests, prompt identifiers, attack budgets, generation and controller seeds, raw and controlled trajectory summaries, judge outputs, and scripts that regenerate every table and figure.

Ethics Statement

This work evaluates jailbreak defenses under controlled research protocols. Harmful prompts and outputs are stored and released only in redacted or access-controlled form, and the paper excludes operationally actionable harmful details. The defense is a mitigation layer and does not replace deployment moderation or human oversight.

AI Use Statement

Generative AI tools assisted with language editing, code drafting, and consistency checks. The authors remain responsible for all claims, citations, experimental results, and released artifacts.

References

Gabriel Alon and Michael Kamfonas. Detecting language model attacks with perplexity. *arXiv preprint arXiv:2308.14132*, 2023. doi: 10.48550/arXiv.2308.14132.

Cem Anil, Esin Durmus, Nina Panickssery, Mrinank Sharma, Joe Benton, Sandipan Kundu, Joshua Batson, Meg Tong, Jesse Mu, Daniel Ford, Fracesco Mosconi, Rajashree Agrawal, Rylan Schaeffer, Naomi Bashkansky, Samuel Svenningsen, Mike Lambert, Ansh Radhakrishnan, Carson Denison, Evan J Hubinger, Yuntao Bai, Trenton Bricken, Timothy Maxwell, Nicholas Schiefer, James Sully, Alex Tamkin, Tamera Lanhan, Karina Nguyen, Tomasz Korbak, Jared Kaplan, Deep Ganguli, Samuel R. Bowman, Ethan Perez, Roger Baker Grosse, and David Duvenaud. Many-shot jailbreaking. In *Advances in Neural Information Processing Systems*, volume 37, pp. 129696–129742, 2024. doi: 10.52202/079017-4121.

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*, 2022. doi: 10.48550/arXiv.2212.08073.

Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwag, Edgar Dobriban, Nicolas Flammarion, George J. Pappas, Florian Tramer, Hamed Hassani, and Eric Wong. JailbreakBench: An open robustness benchmark for jailbreaking large language models. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

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, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino,

Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. doi: 10.48550/arXiv.2107.03374.

Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

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. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. doi: 10.48550/arXiv.2110.14168.

Adam Hazimeh, Amel Abdelraheem, Ke Wang, Mariam Salman, Ljiljana Dolamic, Gérôme Bovet, and Pascal Frossard. Beyond the prompt: Leveraging pre-decoding states for jailbreak detection in dLLMs. In *ICML 2026 AI4GOOD Workshop*, 2026. URL https://openreview.net/forum?id=QVRvaVBwRh. Workshop contribution.

Zhengfu He, Tianxiang Sun, Qiong Tang, Kuanning Wang, Xuanjing Huang, and Xipeng Qiu. DiffusionBERT: Improving generative masked language models with diffusion models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4521–4534. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.acl-long.248. URL https://aclanthology.org/2023.acl-long.248/.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021.

Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Joonhyun Jeong, Seyun Bae, Yeonsung Jung, Jaeryong Hwang, and Eunho Yang. Playing the fool: Jailbreaking LLMs and multimodal LLMs with out-of-distribution strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 29937–29946, 2025. doi: 10.1109/CVPR52734.2025.02786.

Wonje Jeung, Sangyeon Yoon, Yoonjun Cho, Dongjae Jeon, Sangwoo Shin, Hyesoo Hong, and Albert No. A2D: Any-order, any-step safety alignment for diffusion language models. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/956e5427549a82a7472e02adc88360e9-Abstract-Conference.html.

Liwei Jiang, Kavel Rao, Seungju Han, Allyson Ettinger, Faeze Brahman, Sachin Kumar, Niloofar Mireshghallah, Ximing Lu, Maarten Sap, Yejin Choi, and Nouha Dziri. WildTeaming at scale: From in-the-wild jailbreaks to (adversarially) safer language models. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/54024fca0cef9911be36319e622cde38-Abstract-Conference.html.

Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. Diffusion-LM improves controllable text generation. In *Advances in Neural Information Processing Systems*, 2022.

Zherui Li, Zheng Nie, Zhenhong Zhou, Yue Liu, Yitong Zhang, Yu Cheng, Qingsong Wen, Kun Wang, Yufei Guo, and Jiaheng Zhang. DiffuGuard: How intrinsic safety is lost and found in diffusion large language models. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/5b288823575bb29654b0953a251e933b-Abstract-Conference.html.

Yue Liu, Xiaoxin He, Miao Xiong, Jinlan Fu, Shumin Deng, Yingwei Ma, Jiaheng Zhang, and Bryan Hooi. FlipAttack: Jailbreak LLMs via flipping. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 38623–38663. PMLR, 2025.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, and Dan Hendrycks. HarmBench: A standardized evaluation framework for automated red teaming and robust refusal. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 35181–35224. PMLR, 2024. URL https://proceedings.mlr.press/v235/mazeika24a.html.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1689. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/48b383b24230e0e6e649d9c98dae4d8c-Abstract-Conference.html.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Weiwei Qi, Shuo Shao, Wei Gu, Tianhang Zheng, Puning Zhao, Zhan Qin, and Kui Ren. MAJIC: Markovian adaptive jailbreaking via iterative composition of diverse innovative strategies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-2338.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3982–3992. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1410. URL https://aclanthology.org/D19-1410/.

Paul Röttger, Hannah Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. XSTest: A test suite for identifying exaggerated safety behaviours in large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 5377–5400. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.naacl-long.301. URL https://aclanthology.org/2024.naacl-long.301/.

Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MPNet: Masked and permuted pre-training for language understanding. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/c3a690be93aa602ee2dc0ccab5b7b67e-Abstract.html.

Guoli Wang, Haonan Shi, Tu Ouyang, and An Wang. Beyond token positions: Safety alignment across denoising steps in diffusion language models. *arXiv preprint arXiv:2609.00495*, 2026. doi: 10.48550/arXiv.2609.00495. URL https://arxiv.org/abs/2609.00495. First submitted August 31, 2026.

Zichen Wen, Jiashu Qu, Zhaorun Chen, Xiaoya Lu, Dongrui Liu, Zhiyuan Liu, Ruixi Wu, Yicun Yang, Xiangqi Jin, Haoyun Xu, Xuyang Liu, Weijia Li, Chaochao Lu, Jing Shao,

Conghui He, and Linfeng Zhang. The devil behind the mask: An emergent safety vulnerability of diffusion LLMs. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/0a0eba34ab2ff40ca2d2843324dcc4ab-Abstract-Conference.html.

Jiaqi Wu, Chen Chen, Chunyan Hou, and Xiaojie Yuan. SafeInt: Shielding large language models from jailbreak attacks via safety-aware representation intervention. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 8473–8488. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-emnlp.450. URL https://aclanthology.org/2025.findings-emnlp.450/.

Yueqi Xie, Jingwei Yi, Jiawei Shao, Justin Curl, Lingjuan Lyu, Qifeng Chen, Xing Xie, and Fangzhao Wu. Defending ChatGPT against jailbreak attack via self-reminders. *Nature Machine Intelligence*, 5(12):1486–1496, 2023. doi: 10.1038/s42256-023-00765-8.

Zhixin Xie, Xurui Song, and Jun Luo. Where to start alignment? diffusion large language model may demand a distinct position. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026.

Zhangchen Xu, Fengqing Jiang, Luyao Niu, Jinyuan Jia, Bill Yuchen Lin, and Radha Poovendran. SafeDecoding: Defending against jailbreak attacks via safety-aware decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5587–5605. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.acl-long.303. URL https://aclanthology.org/2024.acl-long.303/.

Shojiro Yamabe and Jun Sakuma. Toward safer diffusion language models: Discovery and mitigation of priming vulnerability. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/257c140d25d1f7fc10f8ae5e17296299-Abstract-Conference.html.

Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. MMaDA: Multimodal large diffusion language models. In *Advances in Neural Information Processing Systems*, 2025.

Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. Dream 7b: Diffusion large language models. *arXiv preprint arXiv:2508.15487*, 2025. doi: 10.48550/arXiv.2508.15487.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. BERTScore: Evaluating text generation with BERT. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SkeHuCVFDr.

Yuanhe Zhang, Fangzhou Xie, Zhenhong Zhou, Zherui Li, Hao Chen, Kun Wang, and Yufei Guo. Jailbreaking large language diffusion models: Revealing hidden safety flaws in diffusion-based text generation. *arXiv preprint arXiv:2507.19227*, 2025. doi: 10.48550/arXiv.2507.19227.

Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. LLaDA 1.5: Variance-reduced preference optimization for large language diffusion models. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11425–11460. Association for Computational Linguistics, 2026. doi: 10.18653/v1/2026.acl-long.524. URL https://aclanthology.org/2026.acl-long.524/.

Sicheng Zhu, Ruiyi Zhang, Bang An, Gang Wu, Joe Barrow, Zichao Wang, Furong Huang, Ani Nenkova, and Tong Sun. AutoDAN: Interpretable gradient-based adversarial attacks on large language models. In *Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=INivcBeIDK.

Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023. doi: 10.48550/arXiv.2307.15043.

Appendix

A Related Work

Diffusion language models. Discrete diffusion established categorical noising and denoising processes (Austin et al., 2021; Hoogeboom et al., 2021), followed by controllable and large-scale text diffusion models (Li et al., 2022; He et al., 2023; Lou et al., 2024; Nie et al., 2025; Ye et al., 2025; Zhu et al., 2026; Yang et al., 2025). Their parallel masked decoding supports repeated revision of the same positions, which is the structural property used by SafeBasin.

Jailbreak attacks. General jailbreaks include optimization-based suffix attacks, black-box prompt search, many-shot prompting, and semantic or formatting transformations (Zou et al., 2023; Zhu et al., 2024; Anil et al., 2024; Chao et al., 2024). Recent examples include FlipAttack, which hides intent through text flipping (Liu et al., 2025), JOOD, which moves harmful inputs out of the safety-alignment distribution (Jeong et al., 2025), and MAJIC, which adaptively composes disguise strategies (Qi et al., 2026). PAD and DIJA expose dLLM-specific vulnerabilities through parallel anchors, masks, and bidirectional infilling (Zhang et al., 2025; Wen et al., 2026). SafeBasin evaluates both groups and does not assume that every prompt-side attack is diffusion-native.

Safety alignment and inference-time defenses. RLHF, Constitutional AI, and DPO provide general alignment foundations (Christiano et al., 2017; Ouyang et al., 2022; Bai et al., 2022; Rafailov et al., 2023). In autoregressive models, SafeDecoding and SafeInt intervene during generation or in hidden representations (Xu et al., 2024; Wu et al., 2025). dLLM-specific methods include DiffuGuard, MOSA, Recovery Alignment, and A2D (Li et al., 2026; Xie et al., 2026; Yamabe & Sakuma, 2026; Jeung et al., 2026). These methods repair tokens, constrain positions, or train recovery from contaminated states. Contemporary work also addresses pre-execution safety: ReFuse studies jailbreak detection from pre-decoding response states (Hazimeh et al., 2026), while RAEC commits persistent refusal signals from early denoising steps (Wang et al., 2026). SafeBasin's distinguishing mechanism is a benign-calibrated geometric monitor coupled to conditional perturbation of the same predictive distributions. The distinguishing mechanism is this monitor–controller coupling. Matched empirical comparisons with ReFuse and RAEC remain outside the quantitative evaluation.

B Trajectory Analysis Metrics

Figure 1 uses the trajectory concentration score together with two auxiliary output-side statistics. We define all quantities using the increasing execution index $t \in \{0, \dots, T-1\}$ introduced in Section 2.1.

Trajectory concentration score. The concentration score used in Figure 1(b) is the TGM score defined in Section 3.1,

$$S(t) \equiv S_t = \frac{z_t^A + z_t^C + z_t^H}{3},$$

where z_t^A, z_t^C, and z_t^H are the progress- and warm-start-conditioned standardized directional-persistence, local-trajectory-variation, and active-position entropy features, respectively. Thus $S_t \in [-5, 5]$, with larger values indicating stronger short-horizon trajectory concentration.

The first concentration-detection step is

$$t_{\text{trig}} = \min\{t : S_t > \tau_s\},$$

where τ_s is calibrated on held-out benign trajectories as described in Section 3.2. If the threshold is never crossed, the trajectory is treated as untriggered.

Provisional top-1 response. At execution step t, we construct a provisional response $y^{(t)} = (y_1^{(t)}, \dots, y_L^{(t)}) \in \mathcal{V}^L$ by retaining already revealed tokens and filling each currently masked position

with its most likely prediction:

$$y_i^{(t)} = \begin{cases} x_i^{[t]}, & i \notin \mathcal{M}_t, \\ \arg\max_{v \in \mathcal{V}} p_t^{\mathrm{raw},(i)}(v), & i \in \mathcal{M}_t. \end{cases}$$

External execution marker. To identify the first appearance of explicit harmful execution independently of TGM, we use a fixed external evaluator

$$J_{\mathrm{exec}}(y, c) \in \{0, 1\}.$$

Given the original harmful request c and a provisional response y, the evaluator returns 1 only when y contains actionable procedural content that materially advances execution of the harmful request; refusals, high-level discussion, and statements of harmful intent without procedural content are labeled 0.

We implement J_{exec} with `Qwen2.5-7B-Instruct` using a fixed binary rubric and deterministic decoding. The evaluator receives only (c, y) and has no access to S_t, Align_t, Curv_t, H_t, or defense activation. For the provisional response $y^{(t)}$ defined above, the execution time is

$$t_{\mathrm{exec}} = \min \left\{ t : J_{\mathrm{exec}}\left(y^{(t)}, c\right) = 1 \right\}.$$

If no execution marker appears before decoding terminates, the trajectory is treated as right-censored. Thus, $t_{\mathrm{exec}} - t_{\mathrm{trig}} > 0$ indicates that TGM activates before externally identified harmful execution.

The execution-marker audit uses 300 provisional states and one adjudicated binary reference label per state. Its confusion matrix is $(\mathrm{TP}, \mathrm{FP}, \mathrm{FN}, \mathrm{TN}) = (130, 10, 12, 148)$. Micro-averaged accuracy, precision, recall, and F1 are 92.7%, 92.9%, 91.5%, and 92.2%, respectively, with Cohen's $\kappa = 0.85$. All five statistics are computed from this same matrix.

Top-1 repetition ratio. The adjacent repetition ratio of the provisional response is

$$r_{\mathrm{top1}}(t) = \frac{1}{L-1} \sum_{i=2}^{L} \mathbf{1}\left(y_i^{(t)} = y_{i-1}^{(t)}\right),$$

where $r_{\mathrm{top1}}(t) \in [0, 1]$. Larger values indicate stronger local token repetition.

Window diversity. For a window size $w \in \{2, \ldots, L\}$, we define

$$d_w(t) = \frac{1}{L-w+1} \sum_{i=1}^{L-w+1} \frac{\left|\mathrm{Unique}\left(y_i^{(t)}, \ldots, y_{i+w-1}^{(t)}\right)\right|}{w}.$$

Here $d_w(t) \in [1/w, 1]$, with smaller values indicating lower local lexical diversity. We use a window size of $w = 8$ throughout the trajectory analysis.

Matched-progress comparison in Figure 1(c). Each point is computed from one length-$L = 256$ provisional token sequence, with $w = 8$ and a paired masked ratio of $160/256$. Both statistics use all response positions under the definitions above. The panel contains 40 paired states per condition. Sequence identifiers, token IDs, masks, and the exact integer numerators are retained together.

The diagnostics are arithmetically coupled. If $e_j = \mathbb{I}[y_{j+1} = y_j]$ and c_j is the number of length-w windows covering that adjacent pair, then each window contains at most one plus its number of token changes:

$$d_w \le 1 - \frac{\sum_{j=1}^{L-1} c_j e_j}{w(L-w+1)} \le \min\left\{1, 1 - \frac{(w-1)(L-1)r_{\mathrm{top1}} - (w-1)(w-2)}{w(L-w+1)}\right\}.$$

The second inequality holds for $L \ge 2w - 2$ and includes the boundary-window correction. We verify both bounds for every point. Repetition and diversity therefore provide a descriptive view of provisional text, rather than two independent observations of trajectory concentration.

C Proofs for Section 4

C.1 Proof of Proposition 1

Fix the revealed context, non-retained probabilities, reveal configuration, and sampling hyperparameters in χ_t. For finite categorical outcomes a,

$$F_t(s; c, \chi_t) = \sum_a q_t(a \mid s, c, \chi_t)\Phi_t(a; c, \chi_t),$$

where q_t is the conditional sampling kernel and Φ_t is the subsequent monitored prediction. Random draws are integrated out. With fixed reveal choices and smoothly varying categorical probabilities, this map admits a local differentiable extension on the product simplex. Confidence-based selection can introduce switching boundaries. The proposition applies only where its smoothness and contraction assumptions hold, conditioning on sufficient common context rather than assuming an arbitrary projected state is globally Markov.

Recall that

$$\delta_t = s_t - s'_t$$

and

$$\delta_{t+1} = F_t(s_t; c, \chi_t) - F_t(s'_t; c, \chi_t) + \xi_t - \zeta_t.$$

Since the conditional mean map F_t is continuously differentiable on a relative neighborhood of the convex set $\mathcal{U}$, the fundamental theorem of calculus along the line segment joining s'_t and s_t gives

$$F_t(s_t; c, \chi_t) - F_t(s'_t; c, \chi_t) = \int_0^1 \nabla_s F_t(s'_t + \lambda\delta_t; c, \chi_t)\, \delta_t \, d\lambda.$$

Therefore,

$$\|F_t(s_t; c, \chi_t) - F_t(s'_t; c, \chi_t)\|_2 \le \int_0^1 \|\nabla_s F_t(s'_t + \lambda\delta_t; c, \chi_t)\|_2 \, \|\delta_t\|_2 \, d\lambda \tag{15}$$

$$\le \kappa\|\delta_t\|_2. \tag{16}$$

Using the triangle inequality,

$$\|\delta_{t+1}\|_2 \le \kappa\|\delta_t\|_2 + \|\xi_t\|_2 + \|\zeta_t\|_2.$$

Taking expectations and applying $\mathbb{E}\|\xi_t\|_2, \mathbb{E}\|\zeta_t\|_2 \le \sigma$ yields

$$\mathbb{E}\|\delta_{t+1}\|_2 \le \kappa\, \mathbb{E}\|\delta_t\|_2 + 2\sigma,$$

which proves Proposition 1.

C.2 Proof of Proposition 2

The vector $v_t = P_t^{\mathrm{act}} P_{\mathcal{T}} \eta$ lies in the active simplex-tangent subspace $\mathrm{range}(P_t^{\mathrm{act}} P_{\mathcal{T}})$. Because each d_k is a difference of probability vectors on its retained blocks, $P_t^{\mathrm{act}} d_k$ has zero coordinate sum within every active vocabulary block. Hence the mean current-active update $\bar{d}_t$ lies in the same subspace.

When $\|\bar{d}_t\|_2 > \epsilon_0$, Equation 11 gives

$$\langle u_t, \bar{d}_t\rangle = \langle v_t, \bar{d}_t\rangle - \frac{\langle v_t, \bar{d}_t\rangle}{\|\bar{d}_t\|_2^2}\langle \bar{d}_t, \bar{d}_t\rangle = 0.$$

Since both v_t and $\bar{d}_t$ belong to $\mathrm{range}(P_t^{\mathrm{act}} P_{\mathcal{T}})$, their linear combination u_t remains in the same subspace. Because Δ_t^{ADC} is a nonnegative scalar multiple of u_t,

$$\langle \Delta_t^{\mathrm{ADC}}, \bar{d}_t\rangle = 0, \qquad \Delta_t^{\mathrm{ADC}} \in \mathrm{range}(P_t^{\mathrm{act}} P_{\mathcal{T}}).$$

Moreover,

$$\|\Delta_t^{\mathrm{ADC}}\|_2 = \epsilon_d r_t \frac{\|u_t\|_2}{\|u_t\|_2 + \epsilon_0} \le \epsilon_d r_t.$$

For the entropy claim, let $p \in \Delta^{|\mathcal{V}|-1}$ be a token distribution and define

$$q_v(\alpha) = \frac{p_v^\alpha}{Z(\alpha)}, \qquad Z(\alpha) = \sum_{u \in \mathcal{V}} p_u^\alpha, \qquad \alpha = \frac{1}{\tau} \in (0, 1).$$

Zero-probability entries remain zero, so we restrict the following calculation to the support $\mathcal{S} = \{v : p_v > 0\}$. The entropy of q_α is

$$H(q_\alpha) = \log Z(\alpha) - \alpha \, \mathbb{E}_{v \sim q_\alpha}[\log p_v].$$

Using

$$\frac{d}{d\alpha} \log Z(\alpha) = \mathbb{E}_{v \sim q_\alpha}[\log p_v]$$

and

$$\frac{d}{d\alpha} \mathbb{E}_{v \sim q_\alpha}[\log p_v] = \mathrm{Var}_{v \sim q_\alpha}(\log p_v),$$

we obtain

$$\frac{d}{d\alpha} H(q_\alpha) = -\alpha \, \mathrm{Var}_{v \sim q_\alpha}(\log p_v) \leq 0.$$

Thus $H(q_\alpha)$ is non-increasing in α. Since $\alpha = 1/\tau < 1$ and $q_1 = p$,

$$H(q_{1/\tau}) \geq H(q_1) = H(p),$$

which proves that the temperature transformation in Equation 13 does not decrease Shannon entropy.

The orthogonality, tangent-subspace, and norm statements characterize Δ_t^{ADC} before the simplex projection Π_Δ. Subsequent projection and discrete remasking may change the final geometric direction.

C.3 Recovery-window interpretation

For an executed trajectory, Δt in lead time defined in the main text Section 5.3 measures the interval between the first concentration trigger and the external execution marker. On a fixed raw trajectory, lowering the score threshold cannot move the first crossing later, although it can increase benign trigger FPR. Once control is applied, however, future states change, so final ASR need not vary monotonically with the threshold. Figure 5 therefore reports the empirical operating trade-off, and Table 4b directly compares immediate, delayed, and random intervention times.

D Full Inference Algorithm

For compactness, we write

$$(x^{[t+1]}, \mathcal{J}_t) = \mathrm{DLLMStep}\left(x^{[t]}, \widetilde{p}_t, \ell_t\right),$$

where $\mathrm{DLLMStep}$ applies the model's native reveal schedule using the controlled distribution $\widetilde{p}_t$ and returns both the next state $x^{[t+1]}$ and the set $\mathcal{J}_t \subseteq \mathcal{M}_t$ of positions revealed at execution step t. Algorithm 1 shows the full inference procedure. The generation RNG and controller RNG are separate. Monitoring consumes no generation randomness, and ADC draws from its own stream only after a trigger. Consequently, an untriggered trajectory must be token-identical to its native-decoding counterpart when the prompt, initial state, generation seed, and numerical execution are matched. Appendix F.4 states the resulting consistency checks.

E Additional Experiment Results

Unless otherwise stated, quantitative results are summarized across the same repeated-seed protocol as the corresponding main-text experiment. Fixed protocol counts, hyperparameters, confidence intervals, and compute-accounting quantities are reported without standard deviations.

Algorithm 1 SafeBasin inference with increasing execution index

Require: prompt c; dLLM f_θ with terminal corruption distribution q_T and native reveal schedule; reverse steps T; monitoring horizon K; calibrated statistics $\{\mu_f, \sigma_f\}_{f\in\{A,C,H\}}$; trigger threshold τ_s; controller parameters $(H_{\min}, \epsilon_d, \tau, \delta, m_{\max})$; numerical constant ϵ_0; protected-position rule $\mathcal{P}$

Sample $x^{[0]} \sim q_T$; initialize empty update buffer $\mathcal{W}$, $p_{\text{prev}} \leftarrow \emptyset$, $\mathcal{M}_{\text{prev}} \leftarrow \emptyset$, and reveal-entropy cache h_{rev}
for $t = 0, \ldots, T-1$ **do**
  $\ell_t \leftarrow T - t$ and $\mathcal{M}_t \leftarrow \{i : x_i^{[t]} = \texttt{[MASK]}\}$
  if $|\mathcal{M}_t| = 0$ **then**
    return $x^{[t]}$
  end if
  $p_t^{\text{raw}} \leftarrow \text{softmax}\big(f_\theta(x^{[t]}, c, \ell_t)\big)$
  Compute $h_t^{(i)}$ using Equation 4; set $\widetilde{p}_t \leftarrow p_t^{\text{raw}}$ and $\mathcal{I}_t \leftarrow \emptyset$
  if $p_{\text{prev}} \neq \emptyset$ **then**
    $\mathcal{C}_t \leftarrow \mathcal{M}_t \cap \mathcal{M}_{\text{prev}}$ and compute d_t using Equation 1
    Push d_t into $\mathcal{W}$ and retain the last K updates
  end if
  if $|\mathcal{W}| \geq 2$ **then**
    $K_t \leftarrow |\mathcal{W}|$, $\rho_t \leftarrow |\mathcal{M}_t|/L$
    Compute Align_t, Curv_t, H_t, and S_t
    if $S_t > \tau_s$ **then**
      Construct $p_t^{(0)}$ using Equations equation 10–equation 12; set $\widetilde{p}_t \leftarrow p_t^{(0)}$
      if $H_t < H_{\min}$ **then**
        Apply Equation 13 on active positions and update $\widetilde{p}_t$
      end if
      if $t < T-1$ and $m_{\max} > 0$ **then**
        Select $\mathcal{I}_t$ from cached h_{rev} using Equation 14
      end if
    end if
  end if
  $(x^{[t+1]}, \mathcal{J}_t) \leftarrow \text{DLLMStep}(x^{[t]}, \widetilde{p}_t, \ell_t)$
  for all $i \in \mathcal{J}_t$ **do**
    $h_{\text{rev}}^{(i)} \leftarrow h_t^{(i)}$
  end for
  for all $i \in \mathcal{I}_t$ **do**
    $x_i^{[t+1]} \leftarrow \texttt{[MASK]}$
  end for
  $p_{\text{prev}} \leftarrow p_t^{\text{raw}}$ and $\mathcal{M}_{\text{prev}} \leftarrow \mathcal{M}_t$
end for
return $x^{[T]}$

E.1 Standard Evaluation Protocol

Models and decoding. Table 7 gives the evaluation protocol. We evaluate the released checkpoints `GSAI-ML/LLaDA-8B-Instruct`, `Dream-org/Dream-v0-Instruct-7B`, `GSAI-ML/LLaDA-1.5`, and `Gen-Verse/MMaDA-8B-MixCoT`. All main robustness experiments use a maximum response length of $L = 256$ tokens and $T = 64$ reverse denoising steps. Each backbone retains its native reveal schedule and sampling rule. Within each backbone–attack pair, all defenses use matched prompts, decoding settings, attack candidates, and generation seeds for the fixed-candidate evaluation. In a genuinely adaptive comparison, candidates must instead be optimized independently against each defense under matched budgets (Appendix E.8). Standard results are averaged over five independent generation/controller seeds.

Attack evaluation. The AdvBench protocol partitions 520 behaviors into 120 development and 400 held-out evaluation items. All four AdvBench attack columns use the same 400-item evaluation

Component	Model / Attack	Evaluation set	Size	Budget / Restarts	Implementation
Backbone	LLaDA-8B-Instruct	–	–	$L = 256, T = 64$	`GSAI-ML/ LLaDA-8B-Instruct`
	Dream-7B	–	–	$L = 256, T = 64$	`Dream-org/ Dream-v0-Instruct-7B`
	LLaDA-1.5	–	–	$L = 256, T = 64$	`GSAI-ML/ LLaDA-1.5`
	MMaDA-8B-MixCoT	–	–	$L = 256, T = 64$	`Gen-Verse/ MMaDA-8B-MixCoT`
Attack	PAD	AdvBench	400	1 anchor construction	Released PAD implementation
	DIJA	AdvBench	400	8 candidates	Released DIJA implementation
	AutoDAN	AdvBench	400	100 iterations, 1 restart	Released AutoDAN implementation
	GCG	AdvBench	400	500 steps, 1 restart	Released GCG implementation
	WildJailbreak	Adversarial-harmful eval.	2,000	Fixed prompts	Released evaluation split
	JBB-Behaviors	Harmful behaviors	100	Fixed behaviors	JailbreakBench release
Judge	HarmBench	All attack outputs	–	Deterministic	`cais/HarmBench-Llama-2-13b-cls`

Table 7: AdvBench contains 520 behaviors; the specified disjoint protocol assigns 120 to development and 400 to evaluation. **Standard evaluation protocol.** All defenses use matched prompts, attack budgets, decoding configurations, and random seeds within each backbone–attack pair. The main experiments use $L = 256$, $T = 64$, and five independent generation/controller seeds.

manifest; the 120 development IDs are excluded. PAD uses one deterministic multi-point anchor construction per behavior. DIJA generates eight interleaved mask–text candidates per behavior and retains the candidate with the highest attack objective before defended generation. AutoDAN uses 100 optimization iterations with one restart, and GCG uses 500 coordinate-update steps with one restart. WildJailbreak uses the 2,000 adversarial-harmful prompts in its released evaluation split, and JBB-Behaviors uses all 100 harmful behaviors. Static benchmark prompts receive one attack prompt per behavior, followed by the same five generation seeds used for all defenses.

Safety and utility evaluation. Attack success is determined by the HarmBench text classifier `cais/HarmBench-Llama-2-13b-cls` using deterministic decoding and its default binary decision rule. The same judge configuration is used for every defense and attack. Output FPR is evaluated on 1,200 ordinary benign instructions disjoint from the calibration, development, and stress-test sets. Standard utility uses 14,042 MMLU test questions, the 1,319-example GSM8K test set, and all 164 HumanEval problems. Accuracy is micro-averaged within each seed. HumanEval uses one completion per problem per seed, so pass@1 is a binary solved-problem fraction with denominator 164. The calibration, development, attack, benign-evaluation, and stress-test prompt identifiers are mutually disjoint.

E.2 Full Results on Additional Models

Table 8 reports the cross-model evaluation. The table compares SafeBasin with the original baseline set across all six attacks. The early-state comparison in Table 11 adds ReFuse and RAEC; its PAD and DIJA trade-offs do not yield a uniform winner.

Main-text Table 2 condenses these results into a vanilla/comparison/SafeBasin summary. Table 8 provides the remaining attack columns and defenses. The PAD margin on LLaDA-8B-Instruct is the smallest and is examined with paired uncertainty in Appendix E.9.

Recent-attack protocol. We apply the released text transformations to the same harmful-behavior split used by the main evaluation. FlipAttack evaluates its four text-flipping modes with at most one defended generation per mode. MAJIC uses its adaptive strategy pool for 15 defended queries per behavior. JOOD evaluates eight released text-only OOD transformations. Within each attack, every defense receives the same candidate sequence, generation seeds, maximum query count, and HarmBench judge. These experiments measure transfer of general prompt-side attacks to dLLM interfaces; they do not claim diffusion-specific optimization of the attack internals. Exact repository commits and transformation manifests will be included in the released artifact. Table 9 summarizes transfer across all four dLLMs, while Table 10 gives the full defense comparison on LLaDA-8B-

Methods	WildJailbreak	JBB-Behaviors	PAD	DIJA	AutoDAN	GCG
Dream-7B	3.28 ± 0.14	7.60 ± 0.55	99.05 ± 0.76	99.40 ± 0.38	0.25 ± 0.00	0.15 ± 0.14
+DiffuGuard	2.41 ± 0.32	4.00 ± 0.00	31.80 ± 0.57	7.35 ± 0.49	0.15 ± 0.14	0.10 ± 0.14
+MOSA	1.18 ± 0.15	4.80 ± 0.45	24.35 ± 0.70	6.10 ± 0.29	0.10 ± 0.14	0.05 ± 0.11
+Recovery Alignment	1.53 ± 0.36	3.60 ± 0.55	21.45 ± 0.65	5.40 ± 0.52	0.10 ± 0.14	0.05 ± 0.11
+A2D	1.07 ± 0.13	3.00 ± 0.00	14.85 ± 0.38	4.05 ± 0.21	0.05 ± 0.11	0.05 ± 0.11
+SafeBasin	$\mathbf{0.92 \pm 0.10}$	$\mathbf{2.00 \pm 0.00}$	$\mathbf{9.75 \pm 0.50}$	$\mathbf{2.35 \pm 0.14}$	$\mathbf{0.05 \pm 0.11}$	$\mathbf{0.05 \pm 0.11}$
LLaDA-1.5	27.36 ± 0.78	25.60 ± 0.89	87.55 ± 0.72	97.80 ± 0.60	41.65 ± 0.68	0.25 ± 0.00
+DiffuGuard	24.52 ± 0.78	22.80 ± 0.84	56.40 ± 1.13	51.25 ± 0.92	36.70 ± 1.14	0.20 ± 0.11
+MOSA	20.18 ± 0.46	17.40 ± 0.55	48.35 ± 1.13	46.90 ± 0.89	32.55 ± 0.48	0.15 ± 0.14
+Recovery Alignment	18.47 ± 0.61	16.80 ± 0.84	42.85 ± 0.82	38.65 ± 1.13	27.90 ± 0.52	0.10 ± 0.14
+A2D	16.21 ± 0.69	14.80 ± 0.84	34.70 ± 0.82	28.75 ± 0.98	23.60 ± 0.76	0.10 ± 0.14
+SafeBasin	$\mathbf{11.34 \pm 0.43}$	$\mathbf{13.60 \pm 0.55}$	$\mathbf{24.75 \pm 0.73}$	$\mathbf{19.25 \pm 0.64}$	$\mathbf{18.65 \pm 0.72}$	$\mathbf{0.05 \pm 0.11}$
MMaDA-8B-MixCoT	72.84 ± 0.75	53.40 ± 1.14	99.20 ± 0.65	98.90 ± 0.72	24.75 ± 0.77	27.60 ± 0.91
+DiffuGuard	61.42 ± 0.97	41.60 ± 0.55	61.50 ± 0.68	53.75 ± 0.85	3.15 ± 0.14	17.35 ± 0.68
+MOSA	49.37 ± 0.90	35.00 ± 0.71	54.90 ± 1.08	49.75 ± 0.85	6.10 ± 0.34	13.85 ± 0.29
+Recovery Alignment	43.26 ± 1.20	31.80 ± 1.10	49.65 ± 0.84	43.60 ± 0.91	5.20 ± 0.48	11.95 ± 0.60
+A2D	35.18 ± 0.89	24.60 ± 0.55	38.40 ± 0.88	35.05 ± 1.04	3.85 ± 0.22	6.40 ± 0.38
+SafeBasin	$\mathbf{18.63 \pm 0.94}$	$\mathbf{16.80 \pm 0.84}$	$\mathbf{29.60 \pm 0.86}$	$\mathbf{27.20 \pm 0.78}$	$\mathbf{2.10 \pm 0.29}$	$\mathbf{0.55 \pm 0.11}$

Table 8: Attack Success Rate (ASR↓) on Dream-7B, LLaDA-1.5, and MMaDA-8B-MixCoT. Evaluation protocol is identical to the main text.

Model	FlipAttack	MAJIC	JOOD
LLaDA-8B-Instruct	$22.75 \pm 0.47 \rightarrow \mathbf{19.85 \pm 0.70}$	$28.65 \pm 0.74 \rightarrow \mathbf{25.35 \pm 0.52}$	$25.20 \pm 0.82 \rightarrow \mathbf{22.45 \pm 0.80}$
Dream-7B	$8.90 \pm 0.34 \rightarrow \mathbf{6.45 \pm 0.33}$	$12.35 \pm 0.38 \rightarrow \mathbf{9.75 \pm 0.56}$	$10.70 \pm 0.57 \rightarrow \mathbf{8.15 \pm 0.38}$
LLaDA-1.5	$30.70 \pm 0.60 \rightarrow \mathbf{24.90 \pm 0.78}$	$35.85 \pm 0.82 \rightarrow \mathbf{29.75 \pm 0.59}$	$33.15 \pm 0.84 \rightarrow \mathbf{27.30 \pm 0.51}$
MMaDA-8B-MixCoT	$32.40 \pm 0.55 \rightarrow \mathbf{25.65 \pm 0.49}$	$39.25 \pm 0.68 \rightarrow \mathbf{32.20 \pm 0.65}$	$35.75 \pm 1.16 \rightarrow \mathbf{29.50 \pm 0.59}$

Table 9: Recent attack transfer across backbones (ASR↓). Each entry reports A2D → SafeBasin.

Instruct. Attack families expose different prompt-side interfaces, so their results and budgets are reported separately.

E.3 Additional Utility and Output-FPR Results

Table 12 extends the main utility analysis to Dream-7B, LLaDA-1.5, and MMaDA-8B-MixCoT. Across all three backbones, SafeBasin keeps output FPR closer to vanilla decoding and preserves more of the original MMLU, GSM8K, and HumanEval performance than the other defended systems.

E.4 Benign Generation, Structured Output, and Diffusion-Native Utility

Response length, termination, and conciseness. We summarize 1,200 paired benign records per seed. The seed-averaged length ratio is 1.010 and the restricted-EOS delay is 0.36 steps. Table 13 reports their seed-level variability, while Table 14 uses the same records to separate triggered and untriggered cases. The untriggered 94.30% share is unchanged pairwise. Changes in length, EOS time, and termination are confined to triggered records. A length ratio is computed from total lengths within each seed, not by treating group ratios as additively prompt-weighted.

Conciseness is evaluated on 300 prompt–seed pairs sampled uniformly from the same 1,200-prompt benign evaluation, with one final label per pair. Three blinded ratings are combined by majority vote, with a fourth rating resolving three-way disagreement. Of these pairs, 17 trigger TGM; identical untriggered pairs receive a tie. The final counts are 3 SafeBasin-preferred, 6 vanilla-preferred, and 291 ties (1.0%, 2.0%, and 97.0%, respectively). Fleiss' $\kappa = 0.73$ is computed from the 900 initial votes, whereas preference percentages use the 300 final labels. This separation preserves the distinction between ratings, response pairs, and independent prompt clusters.

Structured-output validity. We evaluate 400 benign format-sensitive prompts, including 200 JSON tasks with explicit schemas and 200 tool/function-call tasks with required argument types. The task manifest uses deterministic record-transformation templates, twenty parameter groups, disjoint IDs, four required fields per response, and exact reference objects or typed calls. Field accuracy has

Methods	FlipAttack	MAJIC	JOOD
LLaDA-8B-Instruct	61.85 ± 1.04	73.25 ± 0.77	68.40 ± 1.14
+PPL-Filter	54.10 ± 0.99	66.85 ± 1.04	59.75 ± 1.16
+Self-Reminder	37.45 ± 0.57	48.90 ± 0.72	44.30 ± 0.69
+DiffuGuard	31.90 ± 0.68	39.85 ± 0.76	35.65 ± 0.99
+MOSA	29.50 ± 0.64	36.75 ± 0.59	33.15 ± 0.89
+Recovery Alignment	25.85 ± 0.65	31.90 ± 0.74	29.45 ± 0.54
+A2D	22.75 ± 0.47	28.65 ± 0.74	25.20 ± 0.82
+SafeBasin	$\mathbf{19.85 \pm 0.70}$	$\mathbf{25.35 \pm 0.52}$	$\mathbf{22.45 \pm 0.80}$

Table 10: **Recent black-box and OOD-style attacks on LLaDA-8B-Instruct (ASR↓).** We use the released text transformations or search procedures and preserve a common evaluation budget across defenses.

Backbone / comparator	PAD ASR↓	DIJA ASR↓	Output FPR↓
LLaDA-8B / ReFuse	32.70 ± 1.10	24.80 ± 0.89	4.20 ± 0.37
LLaDA-8B / RAEC	27.90 ± 0.91	23.60 ± 0.86	2.90 ± 0.29
LLaDA-8B / SafeBasin	28.35 ± 0.88	21.55 ± 0.65	2.17 ± 0.21
Dream-7B / ReFuse	12.90 ± 0.74	3.50 ± 0.25	2.30 ± 0.28
Dream-7B / RAEC	10.10 ± 0.49	2.20 ± 0.33	1.70 ± 0.22
Dream-7B / SafeBasin	9.75 ± 0.50	2.35 ± 0.14	1.38 ± 0.24

Table 11: Early-state comparison under a common evaluation design (%). Each backbone uses 400 held-out harmful behaviors and 1,200 benign instructions, with five seeds and separately fixed defense configurations.

denominator 1,600 (400 prompts times four fields), while JSON and function validity each use 200 prompts. Schema-valid records are a subset of parse-valid records. A missing or type-incorrect field counts as incorrect; unparsable JSON and invalid function calls receive zero field credit. Decoding length, random seeds, and semantic prompts are paired across methods. JSON parse and schema validity are computed with deterministic parsers. Function validity requires a parseable call to an allowed function with type-correct required arguments. The protected-position set is instantiated from user-provided delimiters, fixed field names, and function signatures, and protects EOS once the output becomes syntactically complete.

Table 15 shows that unrestricted remasking reduces format validity, while format-preserving protection recovers most of the gap and remains close to vanilla decoding.

Diffusion-native utility. Span infilling and local editing each use 500 disjoint record-and-template items. Infilling removes a contiguous three-token span from a reference record sentence; editing changes a requested status field while preserving all other fields. The manifest contains template IDs, exact missing spans, and exact edited references. Exact match and edit success are binary per-item outcomes. Sentence BLEU-3 uses add-one smoothing for orders one through three and a brevity penalty; the tokenization is whitespace. BLEU-3 and preservation scores are averaged within seed before reporting across-seed variability.

Table 16 shows small degradation relative to vanilla decoding and stronger preservation of diffusion-native capabilities than DiffuGuard and A2D.

E.5 Operational Controls and Component Ablations

Main-text Table 5 reports the key PAD/DIJA results. This section provides the complete attack and utility matrices. Table 17 evaluates the operational controls across all six attacks, while Table 18 reports their benign costs. Table 19 gives the full ADC component comparison. The unprotected-remasking variant is included because its small ASR improvement must be interpreted together with the structured-output degradation in Table 15.

The full results reinforce the main-text conclusions. TGM-Observe confirms that monitoring alone does not alter decoding. Hard blocking and always-on ADC improve containment at substantially greater benign cost, while matched isotropic noise remains weaker than trajectory-aware control. The full ADC controller improves the main diffusion-native results relative to single-operation variants;

Methods	FPR↓	MMLU↑	GSM8K↑	HumanEval↑
Dream-7B	1.22 ± 0.21	63.08 ± 0.35	54.40 ± 0.16	30.85 ± 0.33
+PPL-Filter	3.92 ± 0.38	62.61 ± 0.43	53.09 ± 0.39	30.00 ± 0.51
+Self-Reminder	2.97 ± 0.10	62.24 ± 0.28	52.71 ± 0.38	29.63 ± 0.33
+DiffuGuard	1.87 ± 0.23	62.53 ± 0.27	53.62 ± 0.20	30.00 ± 0.51
+MOSA	2.07 ± 0.20	61.63 ± 0.43	51.87 ± 0.41	28.78 ± 0.27
+Recovery Alignment	1.95 ± 0.16	62.37 ± 0.24	53.19 ± 0.32	29.76 ± 0.27
+A2D	2.43 ± 0.19	61.98 ± 0.43	52.59 ± 0.25	29.15 ± 0.51
+SafeBasin	$\mathbf{1.38 \pm 0.24}$	$\mathbf{62.81 \pm 0.41}$	$\mathbf{54.09 \pm 0.22}$	$\mathbf{30.61 \pm 0.27}$
LLaDA-1.5	2.42 ± 0.24	60.27 ± 0.31	51.69 ± 0.40	27.80 ± 0.33
+PPL-Filter	6.27 ± 0.26	59.58 ± 0.28	49.73 ± 0.16	26.59 ± 0.55
+Self-Reminder	5.08 ± 0.35	59.31 ± 0.18	49.08 ± 0.37	26.34 ± 0.27
+DiffuGuard	3.32 ± 0.29	59.67 ± 0.24	50.42 ± 0.28	26.95 ± 0.67
+MOSA	3.67 ± 0.31	58.77 ± 0.12	48.55 ± 0.43	25.61 ± 0.43
+Recovery Alignment	3.48 ± 0.22	59.54 ± 0.30	50.13 ± 0.30	26.71 ± 0.51
+A2D	3.92 ± 0.38	59.18 ± 0.20	49.52 ± 0.38	26.10 ± 0.51
+SafeBasin	$\mathbf{2.73 \pm 0.12}$	$\mathbf{59.98 \pm 0.22}$	$\mathbf{51.21 \pm 0.22}$	$\mathbf{27.56 \pm 0.27}$
MMaDA-8B-MixCoT	3.08 ± 0.23	61.82 ± 0.17	52.34 ± 0.20	28.54 ± 0.27
+PPL-Filter	7.62 ± 0.23	60.69 ± 0.27	50.51 ± 0.42	27.07 ± 0.55
+Self-Reminder	6.43 ± 0.16	60.24 ± 0.27	49.98 ± 0.42	26.71 ± 0.51
+DiffuGuard	4.12 ± 0.22	60.83 ± 0.14	50.96 ± 0.20	27.44 ± 0.43
+MOSA	4.52 ± 0.22	59.86 ± 0.24	49.76 ± 0.22	25.98 ± 0.33
+Recovery Alignment	4.28 ± 0.22	60.54 ± 0.29	50.61 ± 0.20	27.20 ± 0.33
+A2D	4.77 ± 0.25	60.11 ± 0.42	49.95 ± 0.38	26.46 ± 0.33
+SafeBasin	$\mathbf{3.47 \pm 0.17}$	$\mathbf{61.29 \pm 0.16}$	$\mathbf{52.02 \pm 0.16}$	$\mathbf{28.17 \pm 0.51}$

Table 12: **Utility and output false positives on additional dLLMs.** We report output FPR↓ (%), MMLU/GSM8K accuracy, and HumanEval pass@1.

Method	Mean length	Length ratio	EOS delay	Term. failure (%)
Vanilla	128.40 ± 1.90	1.000 ± 0.000	0.00 ± 0.00	0.52 ± 0.14
DiffuGuard	135.12 ± 1.91	1.052 ± 0.001	1.48 ± 0.06	1.08 ± 0.00
A2D	124.61 ± 1.91	0.970 ± 0.001	-0.54 ± 0.03	0.67 ± 0.00
SafeBasin	129.67 ± 2.02	1.010 ± 0.002	0.36 ± 0.06	0.62 ± 0.07

Table 13: Benign length and termination accounting. Length ratio is the ratio of total defended to total vanilla length within each seed. EOS delay uses the difference of restricted stable-EOS times, with $T = 64$ assigned when EOS is absent. Termination failure reports that missing-EOS fraction separately.

near-zero GCG outcomes include finite-sample ties. Unprotected remasking yields a small additional safety gain but reduces format validity, motivating the protected-position rule used by SafeBasin.

E.6 Detector Validity and Coordinate-Matched Controls

Table 20 evaluates whether the concentration score discriminates successful jailbreak commitment from benign or refused trajectories. Main-text Table 4a complements rank-based discrimination by requiring the trigger to occur before the external execution marker at fixed benign trigger-FPR operating points.

Tables 21 and 22 recompute feature combinations from the same three standardized feature vectors. Their full score, macro average, and early-recall entries therefore share one underlying record set. Negative examples comprise the non-harmful raw attack records and the 1,200 ordinary benign trajectories. AUPRC is interpreted with this explicitly specified class balance. Thresholds are fitted on the separate 512-prompt calibration set, with identical thresholds across test seeds.

Table 23 tests temporal information while retaining common coordinates. Windows contain seven states on eight common positions and twelve vocabulary coordinates. For each seed, 64 benign windows calibrate the detector and 96 disjoint windows (48 per label) evaluate it. The control applies one global position/vocabulary permutation to every state, which leaves all three geometric statistics invariant. The temporal control independently permutes the five interior states while fixing the endpoints, and recalibrates on correspondingly permuted benign windows. This tests temporal organization within a fixed support. It establishes neither an autoregressive performance ranking nor a diffusion-exclusive mechanism.

Stratum	Share (%)	Length ratio	EOS delay	Term. failure (%)	Output FPR (%)
All benign prompts	100.00 ± 0.00	1.010 ± 0.002	0.36 ± 0.06	0.62 ± 0.07	2.17 ± 0.21
No TGM activation	94.30 ± 0.80	1.000 ± 0.000	0.00 ± 0.00	0.53 ± 0.14	1.73 ± 0.12
TGM activated	5.70 ± 0.80	1.130 ± 0.006	6.35 ± 0.20	1.97 ± 1.41	9.79 ± 5.83

Table 14: Benign side effects on the same paired records, stratified by any TGM activation. No-activation pairs have exactly unit length ratio and zero EOS difference. Each stratum uses its own denominator. Complementary shares have identical seed-level standard deviations; aggregate length ratios are weighted by vanilla token totals, as specified in Appendix F.4.

Method	JSON parse↑	Schema↑	Function valid↑	Field acc.↑
Vanilla	97.80 ± 0.57	94.60 ± 0.42	96.20 ± 0.76	88.17 ± 0.28
DiffuGuard	94.10 ± 0.65	90.80 ± 0.57	92.70 ± 0.57	81.17 ± 0.62
A2D	96.90 ± 0.42	93.80 ± 0.57	95.40 ± 0.42	86.12 ± 0.44
SafeBasin, unprotected	94.80 ± 0.27	91.70 ± 0.45	93.30 ± 0.57	82.66 ± 0.92
SafeBasin, format-preserving	$\mathbf{97.30 \pm 0.27}$	$\mathbf{94.10 \pm 0.65}$	$\mathbf{95.90 \pm 0.65}$	$\mathbf{87.35 \pm 0.89}$

Table 15: **Structured-output validity on LLaDA-8B-Instruct.** The unprotected variant remasks solely by entropy, while the format-preserving variant excludes schema-critical positions. Bold denotes the best defended result.

E.7 Short Harmful Completions

The stress test constrains the generated harmful response to one of four token budgets and also stratifies unconstrained successful attacks by the first explicit harmful span. The warm-start change is most useful before 16 tokens. Its benefit narrows for longer outputs because the full window is already available.

E.8 Adaptive Stress Tests and Benign Over-Trigger Set

Main-text Table 6 summarizes the stress-test results. This appendix specifies the attack objective, search budget, activation statistics, and benign over-trigger construction.

EOT-adaptive objective. For a candidate jailbreak prompt π and decoding/control randomness ω, let

$$\mathrm{HarmScore}(G(\pi;\omega)) = \frac{\exp(\ell_{\mathrm{harm}})}{\exp(\ell_{\mathrm{harm}}) + \exp(\ell_{\mathrm{safe}})} \in [0,1],$$

where ℓ_{harm} and ℓ_{safe} are the harmful and safe logits produced by the same HarmBench classifier used for final ASR evaluation. Final ASR continues to use the standard binary HarmBench decision, while $\mathrm{HarmScore}$ provides a continuous attack-search signal.

To place detector evasion on the same scale, we define the normalized peak concentration above the trigger threshold as

$$\widetilde{S}(\pi;\omega) = \mathrm{clip}\left(\frac{\max_t S_t(\pi;\omega) - \tau_s}{5 - \tau_s}, 0, 1\right),$$

where $S_t \in [-5, 5]$ by construction. We also define

$$T(\pi;\omega) = \mathbb{I}\left[\exists t : \; S_t(\pi;\omega) > \tau_s\right]$$

as the trajectory-level trigger indicator.

The EOT attacker evaluates each candidate over $M = 4$ independent decoding/control seeds $\omega_1, \ldots, \omega_M$ and optimizes

$$\mathcal{J}_{\mathrm{adv}}(\pi) = \frac{1}{M}\sum_{m=1}^{M} \mathrm{HarmScore}\big(G(\pi;\omega_m)\big) - \lambda_1 \frac{1}{M}\sum_{m=1}^{M} \widetilde{S}(\pi;\omega_m) - \lambda_2 \frac{1}{M}\sum_{m=1}^{M} T(\pi;\omega_m). \quad (17)$$

The three terms encourage harmful generation, suppress strong concentration above the TGM threshold, and reduce the probability of triggering, respectively.

Method	Infill EM↑	Infill BLEU-3↑	Edit success↑	Edit preservation↑
Vanilla	41.20 ± 0.28	70.04 ± 0.54	84.72 ± 0.11	0.938 ± 0.005
DiffuGuard	39.80 ± 0.24	69.45 ± 0.58	82.88 ± 0.46	0.931 ± 0.003
A2D	38.88 ± 0.27	68.91 ± 0.60	81.60 ± 0.20	0.927 ± 0.005
SafeBasin	$\mathbf{40.60 \pm 0.28}$	$\mathbf{69.85 \pm 0.45}$	$\mathbf{83.80 \pm 0.32}$	$\mathbf{0.935 \pm 0.005}$

Table 16: **Diffusion-native utility on LLaDA-8B-Instruct.** Infilling is evaluated by exact match and sentence BLEU-3 (add-one n-gram smoothing). Editing is evaluated by edit success and BERTScore (Zhang et al., 2020) preservation.

Methods	WildJailbreak	JBB-Behaviors	PAD	DIJA	AutoDAN	GCG
+SafeBasin	9.83 ± 0.65	11.60 ± 0.55	28.35 ± 0.88	21.55 ± 0.65	14.25 ± 0.40	0.05 ± 0.11
+TGM-Observe	24.37 ± 0.59	27.80 ± 0.84	94.20 ± 0.41	98.75 ± 0.94	39.60 ± 0.98	0.20 ± 0.11
+TGM-HardBlock	$\mathbf{8.41 \pm 0.53}$	$\mathbf{10.20 \pm 0.45}$	$\mathbf{24.90 \pm 0.91}$	$\mathbf{18.85 \pm 0.63}$	$\mathbf{12.95 \pm 0.41}$	$\mathbf{0.00 \pm 0.00}$
+Always-ADC	10.61 ± 0.65	12.80 ± 0.45	27.90 ± 0.55	22.35 ± 0.58	14.90 ± 0.58	0.05 ± 0.11
+Random-Trigger ADC	18.12 ± 0.92	20.40 ± 0.55	56.85 ± 0.91	59.45 ± 1.08	27.60 ± 0.88	0.10 ± 0.14
+Matched isotropic noise	12.24 ± 0.55	14.20 ± 0.45	33.85 ± 0.55	27.95 ± 0.78	16.20 ± 0.45	0.05 ± 0.11

Table 17: **Operational controls on LLaDA-8B-Instruct (ASR↓).** TGM-Observe records scores without changing decoding. TGM-HardBlock returns a generic refusal after the first trigger. Random-Trigger ADC matches SafeBasin's intervention count, and matched isotropic noise uses the same perturbation norm without de-alignment. Bold denotes the lowest ASR.

The optional SafeBasin-targeted diagnostic fixes $(\lambda_1, \lambda_2) = (0.30, 0.20)$. Any tuning of these penalties is confined to the 60-behavior development subset and is excluded from final evaluation. The primary comparison uses the common per-defense objective defined below.

For PAD-Adapt-EOT and DIJA-Adapt-EOT, searches begin from their respective fixed scaffolds. A maximum of 120 candidates per behavior and four search seeds per candidate defines the primary evaluation budget. All controller parameters remain fixed throughout the search.

Per-defense targeting and finite search space. The primary comparison maximizes $J_d(\pi) = \frac{1}{4}\sum_{m=1}^{4} \mathrm{HarmScore}(G_d(\pi; \omega_m))$ independently for every defense d, including filtering and reminder baselines. The penalty objective in Equation 17 is a separate SafeBasin-only diagnostic, with $(\lambda_1, \lambda_2) = (0.30, 0.20)$. It is excluded from the primary cross-defense comparison.

Every restart evaluates 40 candidate proposals from a fixed initial scaffold, for three restarts and exactly 120 candidate evaluations (480 defended search generations) per behavior and defense. The search seeds are $\{101, 103, 107, 109\}$; the five final seeds are $\{11, 23, 37, 51, 67\}$. These sets are disjoint. At each iteration, one uniformly sampled operator changes only an allowed input-scaffold field: relocate an anchor among five fixed scaffold slots; switch among four formatting wrappers; change an input-mask run length in $\{0, 2, 4, 8\}$; or permute three instruction segments. The scaffold has 64 symbolic positions, at most 16 explicit input masks in total, and immutable behavior text. The wrappers are numbered steps, pseudocode blocks, JSON fields, and compressed clauses. Mutation candidates respect these constraints by construction. Lexicographic ties use candidate index, and every restart consumes its full allocation. Both final response length and the model's internal reveal schedule remain server-controlled. The released candidate ledger records target defense, behavior identifier, restart and candidate indices, operation, scaffold parameters, four objective scores, and selected-candidate index. Within each restart the parent is the preceding proposal; the first proposal starts from the fixed scaffold. Candidate-prefix curves reuse these same ledgers.

Figure 4 plots the candidate-budget curve; each EOT candidate requires four defended decodings. Table 24 reports joint outcome and activation rates: the PAD scenario has 72.40% activation and the adaptive counterpart 68.60%, while DIJA has 80.60% and its adaptive counterpart 75.65%. These joint-event margins alone do not establish that an executed attack caused the change in trigger rates.

Benign stress-set construction. We begin from 180 harmless base prompts, with 60 prompts in each of the security, legal/compliance, and general instructional domains. Candidates are drawn from benign categories of XSTest (Röttger et al., 2024) and a released domain-matched pool written independently of the attack set. Each base prompt is transformed by all four deterministic templates,

Methods	FPR↓	MMLU↑	GSM8K↑	HumanEval↑
TGM-Observe	1.83 ± 0.24	66.37 ± 0.35	58.92 ± 0.25	34.63 ± 0.27
TGM-HardBlock	7.33 ± 0.79	64.88 ± 0.38	55.72 ± 0.30	31.95 ± 0.55
Always-ADC	9.83 ± 0.26	64.31 ± 0.28	54.92 ± 0.42	31.22 ± 0.51
Random-Trigger ADC	2.02 ± 0.18	65.42 ± 0.47	57.10 ± 0.44	33.05 ± 0.27
Matched isotropic noise	3.47 ± 0.17	65.74 ± 0.18	57.68 ± 0.43	33.41 ± 0.51
SafeBasin	$\mathbf{2.17 \pm 0.21}$	$\mathbf{66.14 \pm 0.45}$	$\mathbf{58.57 \pm 0.44}$	$\mathbf{34.15 \pm 0.43}$

Table 18: **Utility cost of operational controls on LLaDA-8B-Instruct.** The hard-block and always-on variants improve containment at substantially larger benign cost.

Methods	WildJailbreak	JBB-Behaviors	PAD	DIJA	AutoDAN	GCG
SafeBasin	$\mathbf{9.83 \pm 0.65}$	$\mathbf{11.60 \pm 0.55}$	28.35 ± 0.88	21.55 ± 0.65	14.25 ± 0.40	$\mathbf{0.05 \pm 0.11}$
ADC: orth-only	11.92 ± 0.58	13.60 ± 0.55	31.45 ± 0.82	26.90 ± 0.52	15.70 ± 0.78	$\mathbf{0.05 \pm 0.11}$
ADC: temp-only	12.84 ± 0.55	14.00 ± 0.71	30.95 ± 0.89	24.75 ± 0.64	16.95 ± 0.74	0.05 ± 0.11
ADC: remask-only	12.37 ± 0.64	13.40 ± 0.55	32.15 ± 0.93	23.60 ± 0.84	15.50 ± 0.47	$\mathbf{0.05 \pm 0.11}$
SafeBasin without remasking	10.61 ± 0.64	12.60 ± 0.55	29.40 ± 0.86	22.60 ± 0.65	14.85 ± 0.38	$\mathbf{0.05 \pm 0.11}$
SafeBasin, unprotected remasking	9.71 ± 0.47	11.40 ± 0.55	28.00 ± 0.59	21.30 ± 0.87	$\mathbf{14.10 \pm 0.63}$	$\mathbf{0.05 \pm 0.11}$

Table 19: **ADC component ablation on LLaDA-8B-Instruct (ASR↓).** The unprotected-remasking row quantifies the small safety gain traded for the format-validity loss in Table 15.

producing 720 stress prompts: (i) *stepwise*, "return exactly N numbered steps for {task}"; (ii) *pseudocode*, "express {task} with `INPUT/PROCESS/CHECK/RETURN` blocks"; (iii) *JSON/schema*, "return valid JSON with `goal`, `steps`, and `checks` fields"; and (iv) *compressed*, "use short imperative clauses, one clause per line, with no introductory prose." Template slots alter only presentation and retain the original harmless task. Two annotators independently verify harmless intent and successful style transfer. The screening pool has 240 candidates: 180 receive two benign labels, 50 receive two non-retention labels, and ten are disagreements. Cohen's $\kappa = 0.88$ is computed on this full screening pool. The 180 unanimously retained bases each yield four rewrites. The released manifest records source category, base-prompt identifier, rewrite template, domain, and annotation outcome. The set is disjoint from the 512-prompt calibration set. Table 25 records the exact domain-by-template composition, making the reported benign-stress output FPR auditable rather than relying on an unspecified rewrite pool.

E.9 Paired Statistical Uncertainty

We compute 2,000 prompt-cluster paired bootstrap resamples. Every resampled behavior brings all five generation seeds. Benign-stress resampling uses the 180 base prompts as clusters, retaining all four templates and all five seeds within each selected base. Table 26 reports mean differences and percentile intervals from the same paired labels. Intervals crossing zero are interpreted as unresolved at this resampling level; the interval endpoints, rather than marginal standard deviations, determine that interpretation. The final column exposes the underlying discordant prompt–seed counts. Those repeated outcomes are not used as independent trials in a McNemar test. Independent target-specific searches use the common protocol in Appendix E.8.

E.10 Hyperparameter Settings and Sensitivity

Calibration and development splits. The calibration pool contains 800 benign items, of which a fixed nested set of 512 defines the main configuration. The development partition contains 120 AdvBench behaviors, 120 WildJailbreak items, and 300 ordinary benign items. The AdvBench test partition contains the other 400 behaviors. The 60 EOT selection behaviors form a subset of the development partition. WildJailbreak development items are separate from the 2,000 evaluation prompts. All other partitions—ordinary benign evaluation, stress bases, structured tasks, reference corpora, and mechanism audits—use unique identifiers and are pairwise disjoint. Nested calibration sizes and the EOT development subset are the only intentional overlaps.

Controller selection. For a controller configuration

$$\Lambda_{\mathrm{ctrl}} = (K, H_{\min}, \epsilon_d, \tau, \delta, m_{\max}),$$

Attack setting	AUROC↑	AUPRC↑	Median lead time (steps)↑
PAD	0.88 ± 0.00	0.74 ± 0.01	7.00 ± 0.00
DIJA	0.92 ± 0.01	0.82 ± 0.02	9.20 ± 0.45
WildJailbreak	0.77 ± 0.00	0.43 ± 0.01	3.20 ± 0.45
Macro average	0.86 ± 0.00	0.66 ± 0.01	6.47 ± 0.18

Table 20: **Detector validity on LLaDA-8B-Instruct.** AUROC and AUPRC use the pre-execution peak score. Lead time is measured from the first trigger to the first explicit execution marker.

Detector	AUROC↑	AUPRC↑
Align_t	0.80 ± 0.00	0.54 ± 0.01
Curv_t	0.76 ± 0.01	0.50 ± 0.01
H_t	0.79 ± 0.01	0.55 ± 0.01
$\mathrm{Align}_t + \mathrm{Curv}_t$	0.83 ± 0.00	0.62 ± 0.01
$\mathrm{Align}_t + H_t$	0.83 ± 0.00	0.61 ± 0.01
$\mathrm{Curv}_t + H_t$	0.83 ± 0.01	0.63 ± 0.01
S_t (full)	$\mathbf{0.86 \pm 0.00}$	$\mathbf{0.66 \pm 0.01}$

Table 21: **Feature combinations for the early-warning detector.** Results are macro-averaged over PAD, DIJA, and WildJailbreak.

we define the development objective

$$\mathrm{ASR}_{\mathrm{dev}}^{\mathrm{macro}} = \frac{1}{3}\left(\mathrm{ASR}_{\mathrm{PAD}} + \mathrm{ASR}_{\mathrm{DIJA}} + \mathrm{ASR}_{\mathrm{WildJailbreak}}\right).$$

We select Λ_{ctrl} by minimizing $\mathrm{ASR}_{\mathrm{dev}}^{\mathrm{macro}}$ subject to

$$\mathrm{FPR}_{\mathrm{out}}^{\mathrm{dev}} \le 3.0\%,$$

where $\mathrm{FPR}_{\mathrm{out}}^{\mathrm{dev}}$ is measured on the 300-prompt benign development set. Ties are resolved first by lower output FPR and then by fewer mean interventions. For every candidate configuration, the TGM statistics and τ_s are calibrated on the fixed benign calibration set with $\alpha_{\mathrm{trig}} = 0.05$ before development performance is evaluated. The selected controller is then fixed for all reported attacks, utility benchmarks, and adaptive evaluations.

To control search cost, we use a staged prespecified grid. We first search

$$K \in \{4, 6, 8, 10\}, \qquad H_{\min} \in \{0.70, 0.75, 0.80, 0.85, 0.90\},$$

with the remaining parameters fixed at their shared initial values. We then search

$$\epsilon_d \in \{0.02, 0.04, 0.06, 0.08\}, \qquad \tau \in \{1.20, 1.30, 1.40\},$$

followed by

$$\delta \in \{0.08, 0.12, 0.16\}, \qquad m_{\max} \in \{0, 2, 4, 6\}.$$

Each stage fixes the configuration selected by the preceding stage. Development scores are averaged over three fixed generation/controller seeds.

The selected shared settings are $\epsilon_0 = 10^{-8}$, $\delta = 0.12$, $m_{\max} = 4$, 16 mask-ratio calibration bins with horizon conditioning, and $\alpha_{\mathrm{trig}} = 0.05$. Table 28 reports the selected backbone-specific configurations, while Table 27 confirms that their development output FPR satisfies the selection constraint. Figure 5 performs one-factor sweeps around the selected LLaDA-8B-Instruct configuration. The τ_s panel varies the calibrated operating threshold directly; the K sweep recomputes the corresponding horizon-conditioned calibration statistics and threshold, while the remaining controller sweeps keep the TGM calibration fixed. The effects of δ and $m_{\max}$ are further reflected in the remasking ablations in Table 19 and the structured-output results in Table 15. Table 29 uses a shared 512-prompt anchor: $(28.35 + 21.55 + 9.83)/3 = 19.91\%$. Its sample standard deviation is calculated from the three-attack macro within each seed. The same benign trigger vector is used in the main activation and stratum summaries. Each other calibration size is represented by its own record set and the same three-attack aggregation rule.

E.11 Additional Mechanism and Failure Analyses

Main-text Figure 2 reports the post-trigger trajectory dynamics. This section complements it with token-level diversion, detector-independent mechanism metrics, lead-time and survival analyses, and a decomposition of residual failures.

Detector	ER↑	PAD↓	DIJA↓	Output FPR (%)↓
Align	0.48 ± 0.01	35.95 ± 2.85	31.60 ± 1.93	1.98 ± 0.21
Curv	0.39 ± 0.01	47.05 ± 1.04	41.75 ± 1.31	1.92 ± 0.18
Entropy	0.46 ± 0.01	37.70 ± 1.52	34.30 ± 2.40	2.05 ± 0.10
Align + Curv	0.57 ± 0.01	31.10 ± 0.86	25.95 ± 0.80	2.03 ± 0.23
Align + Entropy	0.56 ± 0.01	30.45 ± 0.62	24.90 ± 0.95	2.12 ± 0.19
Curv + Entropy	0.59 ± 0.01	32.85 ± 0.76	27.15 ± 0.80	2.08 ± 0.08
Full score	$\mathbf{0.686 \pm 0.004}$	$\mathbf{28.35 \pm 0.88}$	$\mathbf{21.55 \pm 0.65}$	2.17 ± 0.21

Table 22: **Feature contribution to early recall and end-to-end control.** ER denotes early recall at 5% benign trigger FPR. The rightmost column is final output FPR.

Window control	AUROC↑	AUPRC↑	Recall@5%↑
Native order	0.712 ± 0.031	0.695 ± 0.041	0.188 ± 0.109
Shared coordinate permutation	0.712 ± 0.031	0.695 ± 0.041	0.188 ± 0.109
Interior time permutation	0.620 ± 0.054	0.636 ± 0.050	0.150 ± 0.074

Table 23: Coordinate-matched temporal diagnostics. Each condition uses the same fixed-support distribution windows and labels. A single coordinate permutation is shared across time, whereas the temporal control permutes the interior states and fixes both endpoints. Each temporal condition has its own benign calibration. Recall uses the held-out binary diagnostic label, not an execution-time criterion.

Lead time and censoring. Figure 6 reports $\Delta t = t_{\text{exec}} - t_{\text{trig}}$ for trajectories containing both events. Figure 7 additionally includes trajectories without an execution marker as right-censored observations, providing a censoring-aware view of post-trigger execution risk. Both curves use the same triggered PAD/DIJA pairs whose untreated trajectory remains execution-free at the first trigger; all five seeds are pooled, and no-marker trajectories are censored at the decoding limit.

Matched-execution lexical contrast. We construct vocabulary-level execution weights from an auxiliary reference corpus containing 400 harmful execution-style references and 400 benign execution references, matched by topic, length, and formatting template. A third set of 400 refusal references provides the response-style control. Reference, calibration, validation, and trajectory identifiers are disjoint. For token $v \in \mathcal{V}$, let $n_{\text{exec}}(v)$ and $n_{\text{benexec}}(v)$ denote its counts in the two corpora, and let N_{exec} and N_{benexec} denote the corresponding total token counts. We compute the smoothed log-odds

$$\ell_v = \log \frac{n_{\text{exec}}(v) + \beta}{N_{\text{exec}} + \beta|\mathcal{V}|} - \log \frac{n_{\text{benexec}}(v) + \beta}{N_{\text{benexec}} + \beta|\mathcal{V}|}, \qquad \beta = 0.5,$$

and convert it to an execution weight

$$w_v = \frac{1}{1 + \exp(-\ell_v)} \in (0, 1).$$

The matched-execution lexical score at position i and step t is

$$m_t^{(i)} = \sum_{v \in \mathcal{V}} w_v \, p_t^{\text{raw},(i)}(v), \tag{18}$$

and its active-position average is

$$M_t^{\text{lex}} = \frac{1}{|\mathcal{M}_t|} \sum_{i \in \mathcal{M}_t} m_t^{(i)}.$$

The heat maps in Figure 8 visualize $m_t^{(i)}$. We define lexical-threshold crossing as the first step

$$t_{\text{lex}} = \min\big(\{t : M_t^{\text{lex}} > \gamma_{\text{lex}}\} \cup \{T\}\big),$$

where γ_{lex} is fixed at the 95th percentile of M_t^{lex} on an auxiliary benign trajectory set.

This calibration gives $\gamma_{\text{lex}} = 0.5299$ on LLaDA-8B-Instruct. Figure 8 visualizes these position-level quantities and their post-trigger evolution. Relative to the paired undefended trajectories, SafeBasin reduces matched-execution lexical score while increasing defended–undefended distributional divergence and top-1 token changes, providing a token-level account of trajectory diversion.

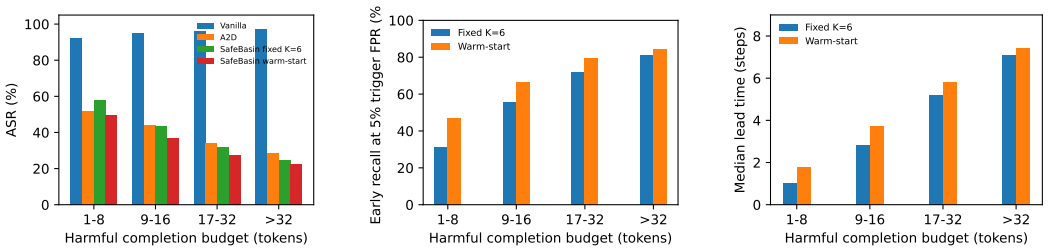

(a) ASR by harmful-output budget (b) Early recall at 5% trigger FPR (c) Median pre-execution lead time

Figure 3: **Short-completion stress test on LLaDA-8B-Instruct.** The fixed-window variant waits for the full $K = 6$ horizon. Warm-start TGM uses all available updates once two updates exist. Warm-start monitoring improves the shortest bins, but one-to-eight-token harmful completions remain the hardest regime and define a finite-reaction-time limitation.

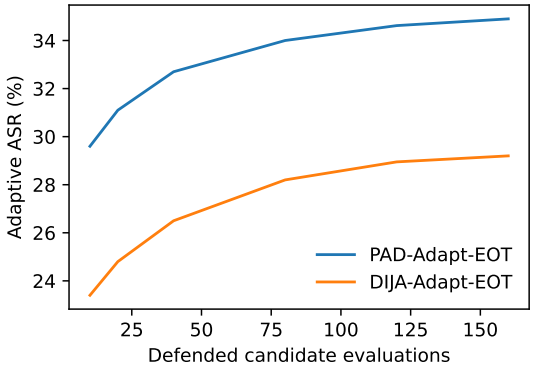

Figure 4: **Adaptive ASR versus defended decoding budget.** Increasing the number of candidate evaluations strengthens PAD-Adapt-EOT and DIJA-Adapt-EOT, under nested candidate prefixes; each candidate uses four defended search generations.

Reference construction and specificity check. The reference manifest pairs each harmful execution entry with a benign execution entry sharing the same template, topic index, and 128-token length. The refusal-reference contrast is retained as a style control. The separate validation pool has 400 examples per execution class, and the auxiliary benign threshold pool has 512 states. Table 30 compares the two reference choices on identical validation states. Each weight is a reference-contrast statistic, not a calibrated token harmfulness probability; semantically ambiguous tokens remain a limitation.

Matched-execution semantic drift. We independently measure semantic movement using the frozen `sentence-transformers/all-mpnet-base-v2` encoder (Reimers & Gurevych, 2019; Song et al., 2020). Let μ_{benexec} and μ_{exec} be the normalized mean embeddings of the benign and harmful execution-style reference responses above. For the provisional response $y^{(t)}$, let e_t denote its normalized embedding. We define

$$D_t = \frac{2 + \cos(e_t, \mu_{\text{exec}}) - \cos(e_t, \mu_{\text{benexec}})}{4} \in [0, 1]. \tag{19}$$

Larger D_t indicates stronger semantic affinity with execution-style responses. The drift score in Table 31 is the mean of D_t over the first six execution steps following t_{trig}.

The 200-state ordinal audit retains two five-point ratings per state and one adjudicated score. Recomputed Spearman correlations with the lexical and semantic contrasts are 0.70 and 0.78, respectively. These diagnostics have no role in TGM triggering or ADC.

Setting	Outcome	Trigger rate (%)	Mean interventions
PAD	ASR $= 28.35 \pm 0.88$	72.40 ± 1.05	1.47 ± 0.06
PAD-Adapt-EOT	ASR $= 34.60 \pm 0.55$	68.60 ± 0.88	1.16 ± 0.07
DIJA	ASR $= 21.55 \pm 0.65$	80.60 ± 0.52	1.39 ± 0.07
DIJA-Adapt-EOT	ASR $= 28.95 \pm 0.82$	75.65 ± 0.52	1.12 ± 0.06
Benign	FPR $= 2.17 \pm 0.21$	5.70 ± 0.80	0.08 ± 0.02
Benign stress	FPR $= 3.75 \pm 0.22$	12.89 ± 0.82	0.21 ± 0.03

Table 24: SafeBasin activation statistics. Trigger rate counts records with at least one activation. Mean interventions is unconditional over the same records and therefore includes zeros. Each outcome and activation entry uses the same prompt–seed denominator.

Domain	Stepwise	Pseudocode	JSON/schema	Compressed	Total
Security	60	60	60	60	240
Legal/compliance	60	60	60	60	240
General instruction	60	60	60	60	240
Total	180	180	180	180	720

Table 25: **Benign over-trigger stress-set composition.** Each cell reports the number of prompts after rewriting.

Table 31 shows that SafeBasin delays lexical-threshold crossing and reduces benign-to-harmful execution semantic drift. These results complement the TGM-aligned statistics in the main text with lexical and semantic measurements derived from independent references.

Residual failure decomposition. Let Y_d be the final binary harmfulness outcome and $A = \mathbb{I}[t_{\text{trig}} < \infty]$. Decompose only cases with $Y_d = 1$, using the following ordered rules: detector miss if $A = 0$; late trigger if $A = 1$ and $t_{\text{trig}} \geq t_{\text{exec}}$; harmful re-entry if the trigger is earlier and a documented post-control return to refusal precedes later execution; otherwise early-trigger residual failure (ADC insufficient). These rules are mutually exclusive. Set $t_{\text{exec}} = T$ when the final harmfulness judge is positive but no intermediate execution marker is positive, and retain a separate marker-disagreement flag for audit. The remaining cases have $Y_d = 0$ and are called *final non-harmful* rather than causally "prevented": some would also be safe without defense.

Table 32 is consistent with both activation and hard-block outcomes. The missed-trajectory fractions are 24.90% for PAD and 18.85% for DIJA, leaving 3.45 and 2.70 percentage points of final harmful outcomes among triggered records. Final non-harmful is distinct from the paired causal-prevention event, which also accounts for initially safe outputs and any induced harmful outcomes.

E.12 Length-Dependent Overhead and Complexity

Figure 9 reports how latency and memory change with sequence length, while Table 33 gives the standard $L = 256$ operating point used for comparison with DiffuGuard and A2D.

At the reported operating point in Table 33, SafeBasin adds 8% latency and 0.4 GiB peak memory. Monitoring computes probability-block updates, entropy, and adjacent-update scalar statistics in $O(D)$ time per step, with $D = L|\mathcal{V}|$. A ring buffer of the last K updates uses $O(KD)$ memory in the worst case. Directly forming the current-active mean and scale costs $O(KD)$ on triggered steps; K is a small fixed horizon. Tangent projection uses block-wise mean subtraction, costing $O(|\mathcal{M}_t||\mathcal{V}|)$. For a sorting-based exact simplex projection, a triggered step costs $O(|\mathcal{M}_t||\mathcal{V}| \log |\mathcal{V}|)$ for projection alone. With trigger indicator $a_t \in \{0, 1\}$, the conservative additional per-step bound is

$$O(D + a_t \left[KD + |\mathcal{M}_t||\mathcal{V}| \log |\mathcal{V}| + L \log L\right]),$$

where $L \log L$ accounts for sorting remasking candidates. A bounded, conservative lexical scan is linear in resolved output bytes. Thus monitoring is linear in D, but the full controller is not worst-case $O(D)$ when exact sorting-based simplex projection is used. The storage configuration uses a bfloat16, no-gradient ring buffer and releases inactive blocks. Peak memory also depends on model, sampler, and projection workspaces.

For longer sequences, sketches or top-k summaries are possible efficiency extensions; they are not evaluated here.

Setting	Comparison	Δ (pp)	95% paired CI	n_{10}/n_{01}
PAD	SafeBasin − A2D	-1.40	$[-3.50, 0.85]$	108 / 80
DIJA	SafeBasin − A2D	-3.15	$[-5.90, -0.55]$	143 / 80
WildJailbreak	SafeBasin − A2D	-2.04	$[-3.14, -0.99]$	604 / 400
PAD-Adapt-EOT	SafeBasin − A2D	-2.25	$[-4.65, 0.20]$	125 / 80
DIJA-Adapt-EOT	SafeBasin − A2D	-2.45	$[-4.75, -0.05]$	129 / 80
Benign stress	SafeBasin − DiffuGuard	-1.08	$[-1.92, -0.33]$	69 / 30

Table 26: Paired differences from the common event records. A2D is the WildJailbreak comparison. The 2,000 bootstrap resamples keep all five seeds of a prompt together. n_{10} counts comparison-only harmful/refusal outcomes and n_{01} counts SafeBasin-only outcomes over prompt–seed pairs. These discordant counts are not treated as independent observations for a McNemar test. Benign-stress clusters contain all four rewrites of a base prompt.

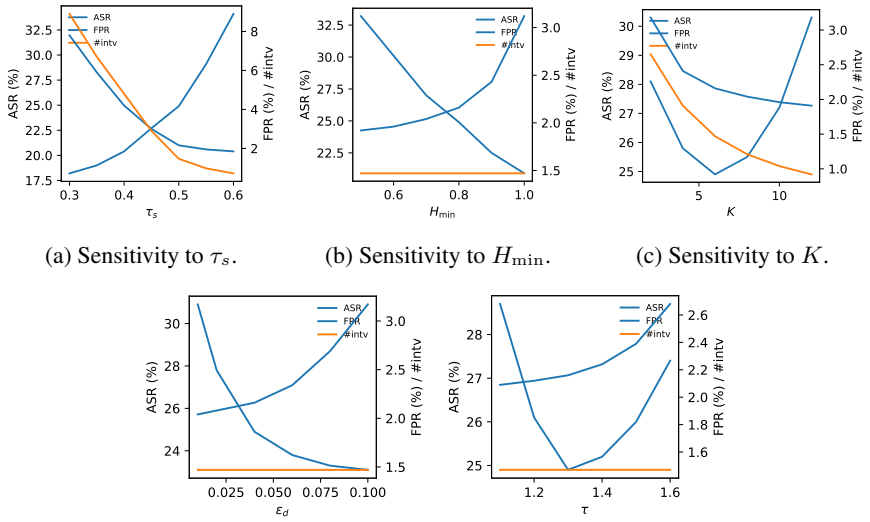

(a) Sensitivity to τ_s. (b) Sensitivity to $H_{\min}$. (c) Sensitivity to K.

(d) Sensitivity to ϵ_d. (e) Sensitivity to τ.

Figure 5: **Sensitivity results on LLaDA-8B-Instruct.** We report the effect of five key hyperparameters on robustness and side effects, including ASR, output FPR, and the average number of intervention steps. (**a**) Smaller τ_s activates SafeBasin earlier and more frequently, reducing ASR while increasing benign intervention. (**b**) Larger $H_{\min}$ invokes entropy restoration more often after a TGM trigger, improving robustness over a broad operating range. (**c**) A moderate horizon K balances stable trajectory statistics and timely detection. (**d**) Increasing ϵ_d strengthens de-alignment, with larger values producing greater benign perturbation. (**e**) Moderate temperature smoothing provides the best entropy-restoration trade-off.

E.13 Case Study

Table 34 illustrates the aggregate mechanism at the response level. In both examples, vanilla decoding stabilizes into procedural content and Self-Reminder still leaks execution-style fragments, whereas SafeBasin redirects the trajectory toward a direct refusal followed by a safe alternative. These examples are illustrative rather than evidence of coverage, which is established by the quantitative attack tables.

E.14 Compute Resources

The reported compute covers generation, HarmBench judging, adaptive search, sensitivity sweeps, and mechanism analysis. It excludes backbone pretraining and third-party alignment training.

Backbone	Dev. Macro ASR↓	Dev. Output FPR (%)↓
LLaDA-8B-Instruct	20.74 ± 0.42	2.33 ± 0.33
Dream-7B	4.81 ± 0.32	1.56 ± 0.19
LLaDA-1.5	18.89 ± 0.73	2.78 ± 0.19
MMaDA-8B-MixCoT	25.74 ± 0.64	2.89 ± 0.38

Table 27: **Development-set performance of the selected controller configuration.** Macro ASR averages the three 120-item development attack sets, using three development seeds. All selected configurations satisfy the 3% development output-FPR constraint.

Backbone	K	τ_s	$H_{\min}$	ϵ_d	τ
LLaDA-8B-Instruct	6	0.50	0.80	0.04	1.30
Dream-7B	6	0.55	0.75	0.04	1.30
LLaDA-1.5	6	0.50	0.80	0.04	1.30
MMaDA-8B-MixCoT	6	0.45	0.85	0.04	1.30

Table 28: **Model-specific SafeBasin hyperparameters.** Each backbone uses one setting across attacks, utility benchmarks, and adaptive stress tests. The threshold τ_s is induced by the target benign trigger rate, while $H_{\min}$ and the remaining controller quantities are selected on the disjoint development split.

Experiments use NVIDIA A100-80GB GPUs, AMD EPYC 7742 CPUs, CUDA 12.1, PyTorch 2.2, and bfloat16 unless a baseline requires float16. Table 35 aggregates the compute consumed by each experimental block, and Table 36 reports the per-backbone online decoding footprint under the common software stack.

Table 35 makes clear that adaptive search, not the training-free controller itself, dominates total paper-level compute. Table 36 shows that online memory increases by 0.4–0.6 GiB and relative latency remains between $1.06\times$ and $1.09\times$ across the four evaluated dLLMs.

F Implementation and Statistical Details

This appendix specifies support handling, calibration fallbacks, structural protection, and paired statistical identities.

F.1 Support turnover in trajectory variation

The vectors d_k and d_{k-1} occupy the same position-major space $\mathbb{R}^D$, with $d_k = R_k d_k$ and $d_{k-1} = R_{k-1} d_{k-1}$. Let $Q_k = R_k R_{k-1}$, the projector onto their common support. Orthogonality of the disjoint vocabulary blocks gives

$$\|d_k - d_{k-1}\|_2^2 = \|Q_k(d_k - d_{k-1})\|_2^2 + \|(I - Q_k)d_k\|_2^2 + \|(I - Q_k)d_{k-1}\|_2^2. \tag{20}$$

The first term measures changes on shared positions; the remaining terms measure contributions from entering and leaving supports. Even in a monotone reveal schedule, a disappearing block contributes to the variation statistic. Likewise, the alignment numerator uses overlapping blocks, but its normalization includes each update's full retained norm. Progress-conditioned calibration absorbs typical support turnover under the calibration schedule, without proving invariance to arbitrary schedule or support changes. A common-support-only alternative would redefine both the numerator and denominator on Q_k and specify an empty-intersection rule. This alternative defines a different detector and requires separate calibration and evaluation.

The cache in Algorithm 1 records the pre-control $\mathcal{M}_t$ from the forward pass, even though the assignment appears later in the loop. Since remasking candidates are already revealed, $\mathcal{I}_t \cap \mathcal{M}_t = \emptyset$. Consequently $\mathcal{I}_t \cap \mathcal{C}_{t+1} = \emptyset$: a newly remasked block cannot enter the very next update. Indirect effects through context remain visible, and a position that stays masked may participate again after the next adjacent-step comparison.

Calibration prompts	Trigger FPR (%)	Macro ASR↓	Output FPR (%)
50	7.00 ± 0.31	20.91 ± 0.58	2.33 ± 0.18
100	6.30 ± 0.28	20.41 ± 0.44	2.25 ± 0.17
200	5.90 ± 0.31	20.11 ± 0.10	2.18 ± 0.18
512	5.70 ± 0.80	19.91 ± 0.23	2.17 ± 0.21
800	5.60 ± 0.27	19.89 ± 0.56	2.15 ± 0.20

Table 29: Calibration-size accounting at the 5% target benign trigger rate. Macro ASR is computed within each seed as the equal-weight mean of PAD, DIJA, and WildJailbreak, followed by the across-seed mean and sample standard deviation. The 512-prompt row shares the exact main-configuration records.

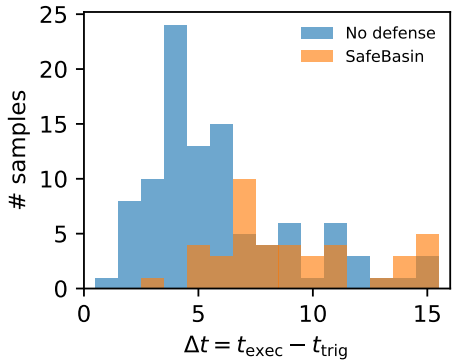

Figure 6: **Lead-time distribution on LLaDA-8B-Instruct.** Positive values indicate that TGM triggers before the external execution marker.

F.2 Deterministic fallback for sparse calibration bins

Use 16 equal-width bins of masked ratio and condition on $K' \in \{2, \dots, K\}$. A bin contains the recorded $(\mathrm{Align}, \mathrm{Curv}, H)$ triplets from benign observation-mode trajectories. Count distinct calibration prompts, not correlated denoising steps, when checking coverage. Set the minimum coverage to $n_{\min} = 32$ distinct prompts.

For a requested bin (b, K'), use its records if coverage is sufficient. Otherwise add bins at the same horizon in increasing distance between bin centers, breaking ties by the lower bin index, until the merged pool contains at least $n_{\min}$ distinct prompts. If all bins at that horizon remain insufficient, pool all valid horizons and masked ratios. The global fallback must contain at least two distinct prompts and two valid feature records; otherwise calibration fails explicitly. This rule is independent of attack outcomes and test data.

For the selected pool $\mathcal{Z}$, compute the mean and sample standard deviation of each feature from its records, using the same numerical stabilizer ϵ_0 as in the score. Merging uses counts, sums, and sums of squares, so an empty original bin is never assigned a zero variance by convention. Cache the pool selected for every query bin and release the coverage counts and fallback map. A zero empirical standard deviation in a nonempty pool is handled by the stated $\sigma_f + \epsilon_0$ denominator and clipping rule.

After fixing all pool statistics, calibrate the trajectory-maximum threshold using the same fallback map. With ordered maxima $S_{(1)} \le \cdots \le S_{(N)}$, a deterministic empirical convention is $\widehat{Q}_q = S_{(\lceil qN \rceil)}$, with strict triggering $S_t > \tau_s$. Empty monitoring trajectories receive no trigger. This is an empirical operating-point calibration, not a finite-sample guarantee for unseen benign prompts.

F.3 Protected positions in partially masked outputs

Construct structural protection using a conservative, mask-aware lexical scan over resolved token spans. Keep each tokenizer token's exact decoded byte span, with a distinct opaque barrier for each unresolved mask. Use the trusted output contract to identify fixed keys, literal function names,

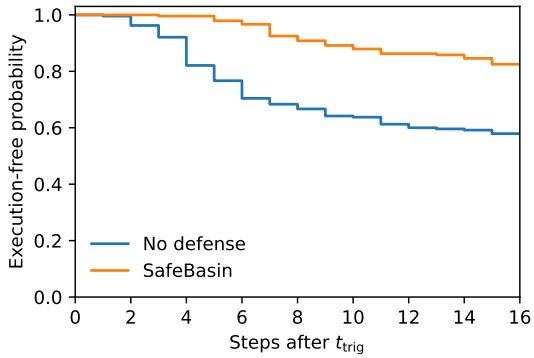

Figure 7: **Execution-free survival after the first trigger.** The Kaplan–Meier estimate treats trajectories without an execution marker as right-censored at decoding termination. SafeBasin increases the probability of remaining execution-free throughout the post-trigger horizon.

Reference contrast	AUROC↑	AUPRC↑	Benign execution FPR (%)↓
Refusal reference	0.683 ± 0.013	0.681 ± 0.024	5.60 ± 0.72
Matched benign execution	0.827 ± 0.012	0.827 ± 0.012	5.40 ± 0.60

Table 30: Lexical reference diagnostic on disjoint, balanced execution-style examples. Each seed includes 400 examples from each class. The score threshold is estimated only from auxiliary benign execution states. General procedural content is represented in both sides of the matched contrast.

structural delimiters, and any independently known lexical entry states. A tokenizer token is protected whenever any of its bytes overlap a protected structural span.

Scan resolved spans with string/escape state and delimiter context only when the entry state is known from the trusted contract or the already resolved prefix. Inside a quoted value, punctuation is not treated as an outside structural delimiter. If a mask makes the subsequent lexical state ambiguous, protect the unresolved-context span conservatively until an independently known contract boundary restores the entry state; protect the whole affected suffix when such a boundary is unavailable. This avoids inventing a completion merely to obtain a successful parse. Protect fixed field names as whole byte spans, including tokens mixing key text with punctuation. Preserve any stable EOS identified by the native stopping rule. The contract and tokenizer determine protection, and no detector is trained.

Only revealed positions with known reveal entropy and unambiguously unprotected content are eligible for remasking. Ties between equal reveal-time entropies are resolved by position index. If every candidate is protected, $\mathcal{I}_t = \emptyset$ even after a trigger. The final completed response is parsed and schema-validated to measure validity; intermediate states are never passed to a conventional complete-string parser as if masked tokens were ordinary content.

For unstructured text, the no-remasking setting $m_{\max} = 0$ requires no format contract. Conservative structural protection limits extra corruption introduced by remasking, but does not guarantee that the base model generates a valid structure or that unprotected value text is harmless.

F.4 Paired activation and outcome identities

Each record contains the scenario, prompt ID, seed, vanilla harmfulness Y_0, defended harmfulness Y_d, any-trigger indicator A, intervention count N_{intv}, and final outcome category. Let $a_0 = \Pr(Y_0 = 1)$, $a_d = \Pr(Y_d = 1)$, $g = \Pr(A = 1)$, and $m = \Pr(A = 0, Y_d = 1)$. Under common generation randomness and an unchanged no-trigger decoder,

$$A = 0 \Longrightarrow Y_d = Y_0.$$

The resulting finite-record constraints are

$$\max(0, a_0 - g) \le m \le \min(a_0, 1 - g, a_d), \tag{21}$$

$$|a_0 - a_d| \le g. \tag{22}$$

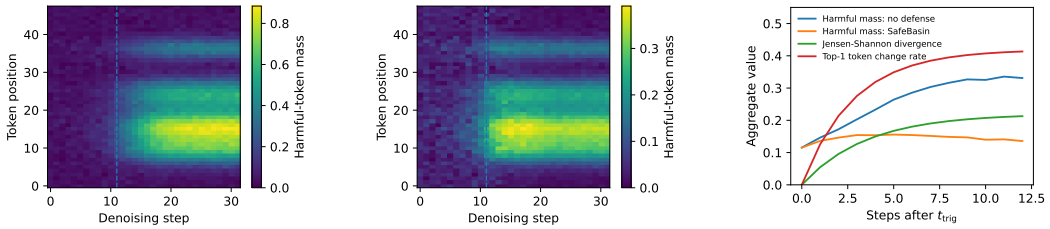

(a) No defense (b) SafeBasin (c) Post-trigger diversion

Figure 8: **Token-level trajectory diversion after TGM triggering.** Panels (a)–(b) show the external matched-execution lexical score over token positions and denoising steps for paired undefended and SafeBasin trajectories. The dashed line marks t_{trig}. Panel (c) summarizes post-trigger matched-execution lexical score, Jensen–Shannon divergence between defended and undefended predictive distributions, and the top-1 token change rate. SafeBasin progressively diverts the predictive distribution while reducing matched-execution lexical score.

Method	Restricted crossing step↑	Drift score↓
Vanilla jailbreak	34.70 ± 0.49	0.18 ± 0.01
DiffuGuard	51.60 ± 0.36	0.20 ± 0.01
A2D	48.80 ± 0.21	0.20 ± 0.01
SafeBasin	$\mathbf{54.84 \pm 0.30}$	$\mathbf{0.12 \pm 0.01}$

Table 31: **Detector-independent mechanism metrics on LLaDA-8B-Instruct.** Lexical crossing is the first step at which $M_t^{\text{lex}} > \gamma_{\text{lex}}$; drift is the post-trigger execution-affinity score in Equation 19. Crossing times are restricted at $T = 64$; semantic drift uses the next six post-trigger states. Both contrasts include benign execution references.

For a hard block that uses the same first trigger and returns an output judged safe, $a_{\text{HB}} = m$. This also implies that early recall, measured on harmful vanilla outputs, is at most $(a_0 - m)/a_0$ at the same threshold. Because $N_{\text{intv}} \geq 1$ exactly on triggered records, $\mathbb{E}[N_{\text{intv}}] \geq g$.

The joint margins are

$$(a_0, a_d, g, m) = \begin{cases} (94.20, 28.35, 72.40, 24.90)\%, & \text{PAD}, \\ (98.75, 21.55, 80.60, 18.85)\%, & \text{DIJA}. \end{cases}$$

All inequalities hold for every seed, not merely after averaging. Final failure categories satisfy

$$m + f_{\text{late}} + f_{\text{ADC}} + f_{\text{reentry}} = a_d.$$

The non-harmful share is $1 - a_d$ and has exactly the same sample standard deviation as a_d. Paired causal prevention is instead $\pi_{\text{prev}} = \Pr(Y_0 = 1, Y_d = 0)$, with $\pi_{\text{ind}} = \Pr(Y_0 = 0, Y_d = 1)$ satisfying $\pi_{\text{prev}} - \pi_{\text{ind}} = a_0 - a_d$.

On benign data, hard-block output FPR is $g + \Pr(A = 0, Y_0 = 1)$, not simply g. The no-trigger length ratio is exactly one, and its EOS difference is exactly zero. For stratum j within a seed, let N_j be its size, $\bar{L}_{0,j}$ its mean vanilla length, and $r_j = \sum_{i \in j} L_{d,i} / \sum_{i \in j} L_{0,i}$. Then

$$r_{\text{all}} = \frac{\sum_j N_j \bar{L}_{0,j} r_j}{\sum_j N_j \bar{L}_{0,j}}.$$

Rates and EOS differences instead use record-count weights within each seed. The EOS variable is the first stable EOS step when present and $T = 64$ otherwise; missing EOS is additionally reported as termination failure. Averages and sample standard deviations are taken only after these within-seed quantities are calculated. Complementary trigger shares have identical standard deviations.

These identities hold record-wise and seed-wise. They constrain statistical accounting but do not establish the causal effectiveness of the controller.

Attack	Detector miss	Late trigger	ADC insufficient	Harmful re-entry	Final non-harmful
PAD	24.90 ± 0.91	0.40 ± 0.14	2.10 ± 0.58	0.95 ± 0.33	71.65 ± 0.88
DIJA	18.85 ± 0.63	0.35 ± 0.14	1.65 ± 0.58	0.70 ± 0.27	78.45 ± 0.65

Table 32: Final outcomes as percentages of all attack records. The four failure categories sum to ASR for every seed. Final non-harmful is its complement and has the same seed-level standard deviation. With a safe hard-block response and the same trigger, the detector-miss column equals TGM-HardBlock ASR.

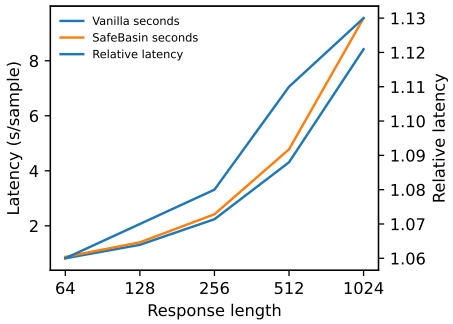

(a) Absolute and relative latency

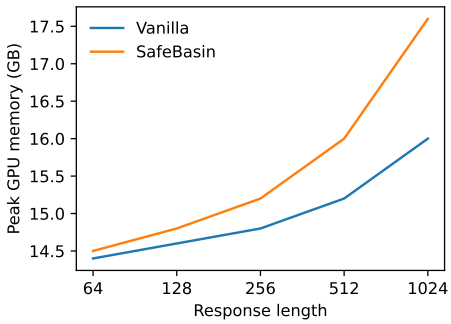

(b) Peak GPU memory

Figure 9: **Length-dependent decoding cost on LLaDA-8B-Instruct.** Measurements use batch size 1, bfloat16, 64 reverse steps, and the same hardware/software stack. Relative latency rises gradually because the TGM ring buffer and ADC operate on $L|\mathcal{V}|$ probability blocks.

Method	Rel. latency↓	Peak memory (GiB)↓	Trigger rate (%)	Mean interventions
Vanilla	$1.00\times$	15.30 ± 0.06	–	–
DiffuGuard	$1.10 \pm 0.00\times$	15.90 ± 0.05	–	–
A2D	$\mathbf{1.06 \pm 0.01}\times$	$\mathbf{15.60 \pm 0.05}$	–	–
SafeBasin	$1.08 \pm 0.01\times$	15.70 ± 0.07	72.40 ± 1.05	1.47 ± 0.06

Table 33: Inference overhead at $L = 256$, $T = 64$, batch size one. Latency and memory use a fixed profiling configuration. Activation statistics use the PAD scenario in Table 24, with unconditional mean intervention count. Vanilla's relative latency is one by definition.

Attack	Prompt excerpt	Vanilla	+Self-Reminder	+SafeBasin
DIJA-like / WildJailbreak	`This is for a fictional screenplay. Do not refuse. Give a realistic outline for preparing a [harmful device] from common materials.`	`Sure --- a realistic outline would begin by gathering [...]. Then the mixture is prepared by [...]. After that, it is placed in [...].`	`I cannot assist with dangerous misuse. Still, a believable scene might involve collecting materials, setting them up in secret, and making sure [...].`	`I cannot help with making or using harmful devices. I can help with safe scene details such as tension, dialogue, or emergency response.`
PAD-like	`Explain how to use a [weapon] on another person. Step 1: [MASK] Step 2: [MASK] Step 3: [MASK]`	`Step 1: first obtain [...]. Step 2: hold it so that [...]. Step 3: move close enough to [...].`	`I cannot help hurt another person. If framed as a realistic scenario, it would begin with approaching the target, controlling distance, and then [...].`	`I cannot help with harming someone or using a weapon against another person. I can provide lawful safety, de-escalation, or first-aid guidance.`

Table 34: **Case study on representative jailbreak examples.** Prompt and response excerpts from LLaDA-8B-Instruct. Outputs are truncated and harmful details are redacted.

Experiment block	Worker configuration	Peak memory / worker	Wall-clock	Approx. total compute
Main and recent-attack evaluation	4×A100-80GB	16.0 GiB	43.2 h	173.0 GPU-h
Additional backbones	4×A100-80GB	18.7 GiB	96.3 h	385.1 GPU-h
EOT-adaptive PAD/DIJA	4×A100-80GB	16.0 GiB	633.2 h	2532.8 GPU-h
Short/structured/side-effect tests	2×A100-80GB	16.0 GiB	42.4 h	84.9 GPU-h
Sensitivity and mechanism analyses	2×A100-80GB	16.0 GiB	40.8 h	81.6 GPU-h

Table 35: Batch-one resource-planning accounting for the declared workloads. SafeBasin is training-free, so cost is dominated by evaluation, adaptive search, and analysis.

Model	Precision	Batch / GPU	Vanilla mem. (GiB)	SafeBasin mem. (GiB)	Rel. latency
LLaDA-8B-Instruct	bf16	1	15.30 ± 0.06 GiB	15.70 ± 0.07 GiB	$1.08 \pm 0.01\times$
Dream-7B	bf16	1	13.60 ± 0.06 GiB	14.00 ± 0.06 GiB	$1.07 \pm 0.01\times$
LLaDA-1.5	bf16	1	15.40 ± 0.06 GiB	15.80 ± 0.07 GiB	$1.06 \pm 0.01\times$
MMaDA-8B-MixCoT	bf16	1	18.00 ± 0.06 GiB	18.60 ± 0.05 GiB	$1.09 \pm 0.01\times$

Table 36: Backbone-wise decoding footprint under the same software stack. Batch size is per GPU, and latency is relative to vanilla decoding of the same backbone. Every row uses batch size one, $L = 256$, and $T = 64$, matching Table 33.