siDiff: Efficacy-Guided Discrete Diffusion for siRNA Target-Site Discovery

Anonymous authors
Paper under double-blind review

Abstract

High-quality experimental data for siRNA efficacy remain scarce, limiting the generalization of discriminative predictors to unseen genes. Although conventional pipelines enumerate all complementarity-valid target windows, prediction errors may concentrate top-ranked candidates within a few local regions while leaving other experimentally active regions under-ranked. Recovering such regions is important for therapeutic discovery, where a diverse portfolio of non-redundant candidates is needed to withstand downstream attrition caused by accessibility, off-target effects, chemical modification, and safety constraints. We propose **siDiff**, an efficacy-guided discrete diffusion framework for siRNA target-site discovery. siDiff first learns a target-conditioned distribution over functional siRNA duplexes and then uses efficacy guidance to reshape this learned distribution during sampling, increasing the probability of efficacy-favored sequence patterns and target regions without retraining the generator. Its coarse-to-fine sampler progressively reveals the siRNA sequence: position guidance prioritizes coordinates with high efficacy potential, while delayed base guidance steers nucleotide assignments only after sufficient sequence context has emerged. By conditioning each decision on an increasingly informative partial sequence, this progressive process captures combinatorial motifs and higher-order intra-siRNA dependencies more explicitly than single-pass candidate scoring. The guided active patterns consequently recover promising regions that may be under-ranked by discriminative predictors, while reverse-complement construction and spatial clustering preserve duplex validity and regional diversity. Experiments on the public Takayuki benchmark and three patent-derived targets demonstrate improved positional recovery and coverage of distinct active regions on unseen genes.

1 Introduction

Small interfering RNAs (siRNAs) induce sequence-specific degradation of complementary messenger RNA (mRNA) and have become an important therapeutic modality Padda et al. (2024). Yet identifying potent candidates remains costly and labor intensive, requiring iterative *in vitro* and *in vivo* screening over large sequence pools.

Computational siRNA methods have progressed from handcrafted sequence and thermodynamic rules to machine-learning, graph, and Transformer predictors Khvorova et al. (2003); Huesken et al. (2005); Katoh & Suzuki (2007); Bai et al. (2024); Long et al. (2024), while foundation models have advanced representation learning for related mRNA tasks Xiong et al. (2025); Zhang et al. (2025b). Conventional siRNA pipelines typically enumerate all complementarity-valid windows and rank them with a discriminative efficacy model (Figure 1, left). Enumeration therefore covers the valid window set, but it does not resolve errors in the scoring function: when high-quality assays are scarce, motif overfitting and limited gene coverage can cause potent regions to rank poorly on unseen targets. The central challenge is thus to allocate proposal probability robustly across target regions despite out-of-distribution predictor error.

We therefore propose siDiff, which combines target-conditioned masked discrete diffusion with efficacy guidance (Figure 1, right). During training, we adapt a Discrete Diffusion Transformer architecture Peebles & Xie (2023) trained on corrupted sequences using an absorbing `[MASK]`

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.

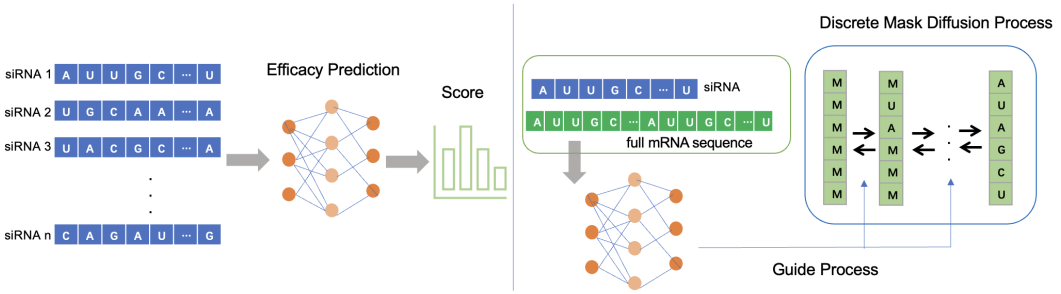


Figure 1: **Comparison of siRNA discovery paradigms. Left:** Traditional rank-based screening, where an exhaustively tiled candidate library is ordered by a discriminative predictor. **Right:** siDiff uses a masked discrete diffusion prior and sampling-time efficacy guidance to reshape target-region proposals before constructing and selecting complete complementary siRNAs.

state, learning to progressively recover unmasked siRNA duplexes. During inference, we introduce a biology-aware, coarse-to-fine guided sampling strategy executed across three distinct phases:

1. **Structural Constraint & Early Rejection:** The target mRNA is partitioned into overlapping segments to handle sequence length constraints. Starting from a fully masked siRNA candidate ($T = 21$), the generator initiates early unmasking steps (e.g., from $T = 21$ down to $T = 14$). Early non-viable candidates—those failing to form valid reverse-complement duplexes within the target mRNA segment—are filtered out immediately.

2. **Fine-Grained Efficacy-Guided Unmasking:** As sampling progresses to lower noise timesteps, more sequence context is revealed. A pre-trained efficacy predictor dynamically guides the selection of both the next position to unmask ($1 \leq pos \leq 21$) and the specific nucleotide assignment (A, U, G, C).

3. **Cluster-Aware Redundancy Mitigation:** Reconstructed active patterns are aligned back to the target mRNA via reverse-complement match to retrieve full siRNA candidates. To prevent spatial clustering of candidates along narrow regions of the mRNA, we apply a cluster-aware diversity mechanism that optimizes both predicted silencing activity and positional sequence diversity.

This design separates exploratory pattern generation from biologically valid candidate construction. In particular, siDiff does not return an unconstrained diffusion sample as the final duplex: every reported candidate is recovered from a reverse-complementary 21-nt target window. The generative process therefore expands target-region exploration while retaining the structural validity required for RNA interference.

Our primary contributions are summarized as follows:

- We introduce **siDiff**, which complements exhaustive candidate scoring with a target-conditioned generative prior that reshapes the proposal distribution over siRNA target regions.

- We propose a biology-aware, coarse-to-fine guided sampling method that couples a discrete masked diffusion transformer with a fine-grained efficacy model to steer sequence generation toward potent candidates.

- We curate three patent-derived target-level benchmarks and establish an evaluation protocol spanning positional accuracy and regional diversity, demonstrating consistent gains on the public Takayuki benchmark and unseen patent targets.

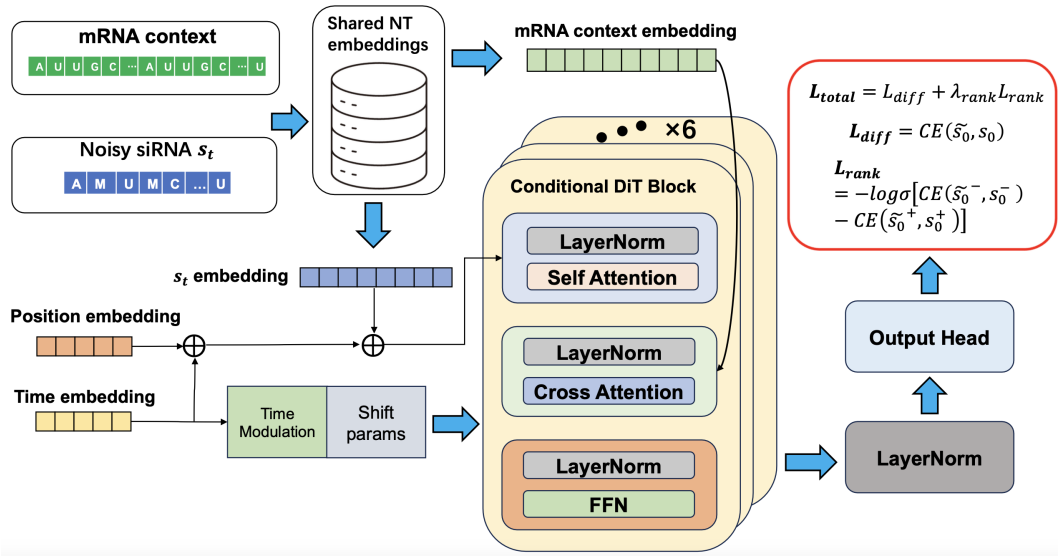


Figure 2: The diagram of the model architecture for training.

2 Related Work

2.1 Discrete Diffusion Models for Sequence Generation

Continuous diffusion has succeeded in image and video generation Ho et al. (2020); Rombach et al. (2022); Ho et al. (2022); Peebles & Xie (2023), but categorical modalities require discrete corruption processes rather than continuous embeddings prone to quantization error Li et al. (2022). D3PMs define Markov transitions over categorical states Austin et al. (2021), while masked diffusion simplifies absorbing-state training to masked-token prediction Sahoo et al. (2024) and enables bidirectional, non-autoregressive generation Li et al. (2026). siDiff adapts this formulation to target-conditioned biological sequence design and adds efficacy guidance during inference.

2.2 siRNA Design and Efficacy Prediction

siRNA efficacy prediction has evolved from biological rules Naito et al. (2009); Khvorova et al. (2003); Katoh & Suzuki (2007) to regression, kernel, ensemble, convolutional, graph, and Transformer models Huesken et al. (2005); Wang et al. (2010); Monopoli et al. (2023); Han et al. (2018); La Rosa et al. (2022); Long et al. (2024); He et al. (2017); Bai et al. (2024); Zhang et al. (2025a). These advances improve scoring of predefined sequences, but do not address guided generation of diverse, target-compatible candidates.

3 Method

3.1 Problem Formulation

We study target-conditioned siRNA candidate proposal. Let $\mathbf{m} = (m_1, \ldots, m_L) \in \mathcal{A}^L$ denote a target mRNA sequence of length L, and let $\mathbf{s} = (s_1, \ldots, s_{L_s}) \in \mathcal{A}^{L_s}$ denote a 21-nucleotide siRNA sequence ($L_s = 21$), where $\mathcal{A} = \{\mathrm{A}, \mathrm{C}, \mathrm{G}, \mathrm{U}\}$. We model the conditional distribution $p_\theta(\mathbf{s} \mid \mathbf{m})$ to capture sequence-dependent regularities shared by target mRNA regions and functional siRNA candidates. Given an input mRNA $\mathbf{m}$, our goal is to learn a conditional generative prior that assigns greater proposal probability to promising target regions while retaining diversity; complete siRNA candidates are subsequently constructed from valid reverse-complementary target windows.

As illustrated in Figure 2, the conditional masked diffusion generator learns an efficacy-aware generative prior from highly effective mRNA–siRNA pairs. siDiff further incorporates a mask-robust

efficacy predictor and a biology-aware inference pipeline enforcing reverse-complement constraints while preserving diversity.

3.2 Masked Discrete Diffusion siRNA Model

We formulate target-conditioned siRNA generation as a discrete masked diffusion process. Given a target mRNA $\mathbf{m}$ and its corresponding siRNA sequence $\mathbf{s}_0 = (s_0^1, \ldots, s_0^{L_s}) \in \mathcal{A}^{L_s}$, the forward corruption process progressively substitutes nucleotides with an absorbing `[MASK]` token. Let $\mathcal{A}_{\mathrm{M}} = \mathcal{A} \cup \{[\texttt{MASK}]\}$ represent the augmented vocabulary.

Following an absorbing-state Markov transition framework Austin et al. (2021); Sahoo et al. (2024), the probability of a nucleotide remaining unmasked at timestep $t \in \{1, \ldots, T\}$ is defined as $\bar{\alpha}_t = \prod_{\tau=1}^{t}(1-\beta_\tau)$, where $\beta_\tau \in (0,1)$ governs the noise schedule. The forward marginal distribution for position $i \in \{1, \ldots, L_s\}$ is expressed as:

$$q(s_t^i \mid s_0^i) = \bar{\alpha}_t \mathbb{I}[s_t^i = s_0^i] + (1-\bar{\alpha}_t)\mathbb{I}[s_t^i = [\texttt{MASK}]], \tag{1}$$

where $\mathbb{I}[\cdot]$ is the indicator function. Under this formulation, once a nucleotide transitions into the absorbing `[MASK]` state, it remains masked until $t = T$, where $\mathbf{s}_T = [\texttt{MASK}]^{L_s}$.

The reverse denoising process aims to recover the original siRNA sequence $\mathbf{s}_0$ conditioned on the target mRNA context $\mathbf{m}$:

$$p_\theta(\mathbf{s}_{t-1} \mid \mathbf{s}_t, \mathbf{m}, t) = \prod_{i=1}^{L_s} p_\theta(s_{t-1}^i \mid \mathbf{s}_t, \mathbf{m}, t). \tag{2}$$

Generation initializes from a fully masked state $\mathbf{s}_T$ and iteratively predicts unmasked token assignments across timesteps guided by the conditional distribution learned by the network.

3.2.1 Training and Context Augmentation

To bias the generative prior toward highly potent candidates, we train the discrete diffusion model exclusively on experimentally validated mRNA–siRNA pairs whose inhibition levels exceed a predefined activity threshold τ_{eff}:

$$\mathcal{D}_{\text{high}} = \left\{ \left(\mathbf{m}^{(j)}, \mathbf{s}^{(j)}, y^{(j)} \right) \right\}_{j=1}^{N_{\text{high}}}, \quad \text{where } y^{(j)} \geq \tau_{\text{eff}}. \tag{3}$$

Here, $\mathbf{m}^{(j)}$ represents the mRNA context, $\mathbf{s}^{(j)}$ is the validated 21-nt siRNA duplex, and $y^{(j)}$ is the experimentally measured inhibition efficiency.

Random-Window Context Augmentation. Naïvely training on fixed mRNA contexts can lead the generator to memorize target sites at fixed offsets. For a context length L_m, we therefore sample an upstream flank length $\ell > 0$ around the L_s-nt target site and set the downstream flank length to

$$r = L_m - L_s - \ell. \tag{4}$$

If p is the target-site start in full transcript $\mathbf{M}$, the augmented context is

$$\mathbf{m}^{(\ell)} = \mathbf{M}_{p-\ell\,:\,p+L_s+r-1}. \tag{5}$$

Varying ℓ across epochs changes the target's relative position while preserving the complete 21-nt binding region, encouraging the model to learn sequence-dependent complementarity rather than positional shortcuts.

3.2.2 Training Objective

Each training instance is represented as a tuple $(\mathbf{m}, \mathbf{s}_0, y) \in \mathcal{D}_{\text{high}}$, where $\mathbf{s}_0$ denotes the uncorrupted 21-nt siRNA sequence, $\mathbf{m}$ is the corresponding target mRNA context, and y is the experimentally measured inhibition efficiency. During training, a diffusion timestep t is sampled uniformly from $U(1, T)$, and a corrupted state $\mathbf{s}_t \sim q(\mathbf{s}_t \mid \mathbf{s}_0)$ is generated via the forward masking process (Eq. 1). Let $\mathcal{M}_t = \{k \mid s_t^k = [\texttt{MASK}]\}$ denote the set of indices corresponding to masked nucleotide positions.

Masked Sequence Reconstruction Loss. The primary generative training objective is the masked-token cross-entropy loss, evaluated strictly over the masked positions $\mathcal{M}_t$:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{\substack{(\mathbf{m},\mathbf{s}_0,y)\sim\mathcal{D}_{\text{high}} \\ t\sim U(1,T) \\ \mathbf{s}_t\sim q(\mathbf{s}_t|\mathbf{s}_0)}} \left[-\frac{1}{|\mathcal{M}_t|} \sum_{k\in\mathcal{M}_t} \log p_\theta \left(s_0^k \mid \mathbf{s}_t, \mathbf{m}, t \right) \right]. \tag{6}$$

This objective forces the Conditional Diffusion Transformer (CondDiT_θ) to learn target-conditioned contextual rules for recovering unmasked nucleotides from partially observed sequence states.

Auxiliary Efficacy-Aware Ranking Loss. Although all samples in $\mathcal{D}_{\text{high}}$ exceed the activity threshold, their measured values provide relative-potency supervision. On clean sequences ($t = 0$), a pooled CondDiT representation is projected to scalar score $r^{(i)}$. To reduce sensitivity to experimental fluctuations, we retain pairs

$$\mathcal{P} = \{(i, j) \mid i < j,\ |y^{(i)} - y^{(j)}| > \delta\}. \tag{7}$$

With $z_{ij} = \text{sign}(y^{(i)} - y^{(j)})$, the pairwise objective is

$$\mathcal{L}_{\text{rank}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \log\Big(1 + \exp[-z_{ij}(r^{(i)} - r^{(j)})]\Big). \tag{8}$$

Joint Training Objective. The joint objective is

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{diff}} + \lambda_{\text{rank}}\mathcal{L}_{\text{rank}}, \tag{9}$$

where λ_{rank} balances reconstruction and efficacy ordering.

3.3 Biology-Aware Sampling Strategy

Our sampler stops reverse diffusion at gate g_0 and treats the resulting partial sequence as an *active pattern* that proposes target regions rather than as a final siRNA. Starting from $\mathbf{s}_T = [\texttt{MASK}]^{L_s}$, each step $t = T, \ldots, g_0 + 1$ resolves one position. For unresolved set $\mathcal{U}_t$, generator logits $\mathbf{z}_{t,i}^{\text{diff}}$ and predictor F_ϕ define

$$\begin{aligned} C_{t,i}^{\text{gen}} &= \max_{b\in\mathcal{A}} \text{softmax}(\mathbf{z}_{t,i}^{\text{diff}})_b, \\ E_{t,i,b} &= F_\phi(\mathbf{m}, \text{Fill}(\mathbf{s}_t, i, b)), \\ I_{t,i}^{\text{eff}} &= \max_{b\in\mathcal{A}} E_{t,i,b}. \end{aligned} \tag{10}$$

We min–max normalize both signals over $\mathcal{U}_t$ and blend them as

$$\begin{aligned} \widetilde{C}_{t,i} = (1 - w_{\text{pos}}(t))\, &\text{Norm}_{\mathcal{U}_t}(C_{t,i}^{\text{gen}}) \\ &+ w_{\text{pos}}(t)\, \text{Norm}_{\mathcal{U}_t}(I_{t,i}^{\text{eff}}), \end{aligned} \tag{11}$$

then sample the next position by

$$j_t = \arg\max_{i\in\mathcal{U}_t} \left\{ \widetilde{C}_{t,i}/\tau_{\text{pos}} + G_i \right\}, \qquad G_i \overset{\text{i.i.d.}}{\sim} \text{Gumbel}(0, 1). \tag{12}$$

Base guidance is delayed because efficacy comparisons are unreliable under heavy masking. For the selected position j_t, let $\widetilde{\mathbf{E}}_{t,j_t} = \text{Standardize}_b(\mathbf{E}_{t,j_t,:})$, set to zero under negligible variance. The decoding logits and schedule are

$$\begin{aligned} \mathbf{z}_{t,j_t}^{\text{joint}} &= \mathbf{z}_{t,j_t}^{\text{diff}} + \beta_{\text{base}}(t)\widetilde{\mathbf{E}}_{t,j_t}, \\ \beta_{\text{base}}(t) &= \begin{cases} 0, & t > g_1, \\ \beta_{\max}\left(\dfrac{g_1 - t + 1}{g_1 - g_0}\right)^p, & g_0 < t \le g_1. \end{cases} \end{aligned} \tag{13}$$

The base is decoded from $\mathbf{z}_{t,j_t}^{\text{joint}}$. The $4|\mathcal{U}_t|$ predictor scores are computed once per step: their maxima guide position selection and the selected four scores are reused for base guidance. Full schedules and decoding details are given in Appendix B.

Candidate derivation and clustering. Instantiated pattern positions are aligned to the mRNA under reverse complementarity and a predefined mismatch tolerance, with masked positions unconstrained. Each match is expanded into a complete reverse-complementary 21-nt candidate. Sorted target starts are greedily clustered: start x joins current cluster $\mathcal{C}$ when $x - \min_{y \in \mathcal{C}} y \leq \Delta_{\text{clu}} = L_s - g_0$; otherwise it begins a new cluster. This bounded span removes redundant local proposals without merging distant regions.

Inhibition-aware candidate selection. A full-sequence predictor assigns efficacy q_a to each completed candidate a. Because neighboring sites may have correlated efficacy Davis et al. (2025), cluster $\mathcal{G}_k$ is summarized by

$$\mu_k = |\mathcal{G}_k|^{-1} \sum_{a \in \mathcal{G}_k} q_a, \qquad \rho_k = |\mathcal{G}_k|^{-1} \sum_{a \in \mathcal{G}_k} \mathbb{I}[q_a \geq \gamma]. \tag{14}$$

Three progressively relaxed threshold grades assign quotas $K_1 > K_2 > K_3$; clusters failing the last grade are removed, and candidates within retained clusters are ranked by q_a. This balances predicted inhibition with coverage across distinct mRNA regions.

4 Experiments

4.1 Experimental Setup

Datasets. We integrate experimentally measured siRNAs from Huesken, Takayuki, and three patent datasets targeting ANGPTL7, CTNNB1, and GSK3A. After sequence standardization and per-mRNA min–max normalization, sites with inhibition ≥ 0.7 are retained. For each site, we sample up to five 61-nt contexts with different valid upstream/downstream offsets while preserving the complete target window. The four evaluation sets contain 541 active sites and 2,698 context views. Models are trained only on Huesken-derived data; Takayuki and all patent targets are held out from training and hyperparameter selection, yielding a target-level external-generalization test. Full statistics and patent curation are provided in Appendix A.1.

Baselines. We compare against random sampling; exhaustive enumeration ranked by iScore or OligoFormer; vanilla diffusion using the same generator without efficacy guidance or predictor-based selection; and guided diffusion followed by efficacy-based selection. We select iScore Katoh & Suzuki (2007) and OligoFormer Bai et al. (2024) as representative, publicly available scoring methods spanning a classical sequence-based model and a contemporary Transformer predictor, respectively. Their released implementations provide reproducible reference points for the conventional enumerate-and-rank paradigm. Enumeration considers all 41 valid starts in each 61-nt context. The diffusion variants share the CrossAttentionDiT backbone, allowing their differences to isolate trajectory guidance and candidate selection rather than generator capacity.

Evaluation protocol. Given a 61-nt mRNA context without its target position, each method returns a ranked list of candidate starts. Position recovery is measured by

$$\text{Hit}_\delta@K = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}\left[\min_{\hat{p} \in \hat{\mathcal{P}}_i^K} |\hat{p} - p_i| \leq \delta \right], \tag{15}$$

where $\hat{\mathcal{P}}_i^K$ contains the top-K unique predicted starts and p_i is the annotated high-activity start. We report $\text{Hit}_0@1$ and $\text{Hit}_3@3$.

Efficacy-Window Coverage (EWC) complements positional accuracy by measuring whether generated candidates recover spatially distinct high-efficacy mRNA regions. From non-overlapping 100-nt windows, we construct the experimentally supported set $\mathcal{W}_{\text{gt}}$ and the predicted set $\mathcal{W}_{\text{pred}}$, and compute

$$\text{EWC} = \frac{|\mathcal{W}_{\text{pred}} \cap \mathcal{W}_{\text{gt}}|}{|\mathcal{W}_{\text{gt}}|}. \tag{16}$$

This distinguishes recovery of multiple promising regions from repeatedly sampling nearby candidates. Window thresholds and the complete construction are given in Appendix D.

Table 1: Generation performance on four evaluation datasets.

	Takayuki			ANGPTL7			CTNNB1			GSK3A		
Method	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC
Random Sampling	0.031	0.359	0/1	0.032	0.421	**1/1**	0.037	0.396	18/29	0.038	0.447	4/9
Enumeration + iScore	0.038	0.401	**1/1**	0.032	0.309	**1/1**	0.022	0.385	20/29	0.046	0.476	2/9
Enumeration + OligoFormer	0.048	0.318	0/1	0.032	0.283	0/1	0.023	0.275	2/29	0.018	0.358	0/9
Vanilla Diffusion	0.026	0.329	**1/1**	0.096	0.516	**1/1**	0.033	0.433	25/29	0.051	0.397	7/9
Guided Diffusion + Efficacy Selection	0.026	0.340	**1/1**	0.064	0.484	**1/1**	0.033	0.471	26/29	0.013	0.423	8/9
siDiff	**0.242**	**0.664**	**1/1**	**0.185**	**0.612**	**1/1**	**0.200**	**0.654**	**28/29**	**0.101**	**0.679**	**9/9**

Table 2: Ablation study of sampling guidance components on four evaluation datasets.

	Takayuki			ANGPTL7			CTNNB1			GSK3A		
Variant	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC
No guidance	0.104	0.487	**1/1**	**0.193**	0.417	**1/1**	**0.221**	0.442	11/29	**0.166**	0.487	2/9
Position-only guidance	0.131	0.639	**1/1**	**0.193**	0.612	**1/1**	0.186	0.631	20/29	0.154	0.666	7/9
Base-only guidance	0.140	0.628	**1/1**	**0.193**	**0.661**	**1/1**	0.173	**0.677**	16/29	0.151	0.564	5/9
Full-Trajectory Guided Diffusion	0.131	0.508	**1/1**	0.064	0.536	0/1	0.029	0.469	25/29	0.051	0.384	6/9
siDiff w/o clustering	0.016	0.371	**1/1**	0.059	0.516	0/1	0.028	0.478	21/29	0.026	0.513	6/9
siDiff	**0.242**	**0.664**	**1/1**	0.185	0.612	**1/1**	0.200	0.654	**28/29**	0.101	**0.679**	**9/9**

The three reported metrics probe complementary failure modes. $\mathrm{Hit}_0@1$ tests exact localization of the highest-ranked candidate, while $\mathrm{Hit}_3@3$ allows small positional tolerance and multiple proposals. EWC ignores repeated local hits and instead tests whether the method recovers distinct high-efficacy regions across a transcript. Reporting them jointly prevents either accurate but spatially collapsed generation or broad but poorly localized exploration from appearing sufficient.

Implementation details. The generator is a six-layer cross-attention DiT with hidden dimension 256, eight attention heads, and 7.6M trainable parameters. It is trained for 50 epochs with AdamW, batch size 64, learning rate 10^{-4}, and a cosine masking schedule. The guidance model is a Transformer encoder–decoder trained on partially masked siRNA–mRNA pairs, and $\lambda_{\mathrm{rank}} = 0.2$. Further optimization and predictor details appear in Appendix F.

4.2 Partial-State Prediction and Target-Site Recovery

Curriculum training improves partial-state predictor PCC/SCC from $0.643/0.641$ to $0.677/0.684$. Before final efficacy ranking, the clustered candidate pools recover annotated starts within two nucleotides at rates of 0.92–0.97 across all external targets. Definitions and complete results are reported in Appendix H.1.

4.3 Efficacy-Guided Design with Diversity Preservation

Table 1 shows that siDiff achieves the best $\mathrm{Hit}_0@1$ and $\mathrm{Hit}_3@3$ on every external target while recovering all ground-truth efficacy windows except one on CTNNB1. On Takayuki, it improves $\mathrm{Hit}_0@1/\mathrm{Hit}_3@3$ from the strongest enumeration values of $0.048/0.401$ to $0.242/0.664$; on CTNNB1, the corresponding gains are from $0.023/0.385$ to $0.200/0.654$. Vanilla and post-hoc selected diffusion retain competitive coverage on the two multi-window targets, but their positional accuracy is consistently lower. This contrast indicates that unconstrained diffusion can explore broadly, whereas the proposed guidance and candidate selection convert that exploration into accurate target-site recovery without collapsing regional coverage.

4.4 Regional Analysis of Efficacy-Window Coverage

To examine regional coverage beyond the aggregate EWC score, we partition each full-length mRNA into contiguous 100-nt windows. A window is considered high efficacy if it contains at least one measured site, at least 10% of its measured sites have inhibition above 0.7, and its mean inhibition is at least 0.6.

Figure 3 shows that siDiff recovers the annotated high-efficacy regions on Takayuki, ANGPTL7, and GSK3A and all but one such region on CTNNB1, consistent with Table 1. On the multi-window CTNNB1 and GSK3A targets, the predicted profiles cover the experimentally supported regions

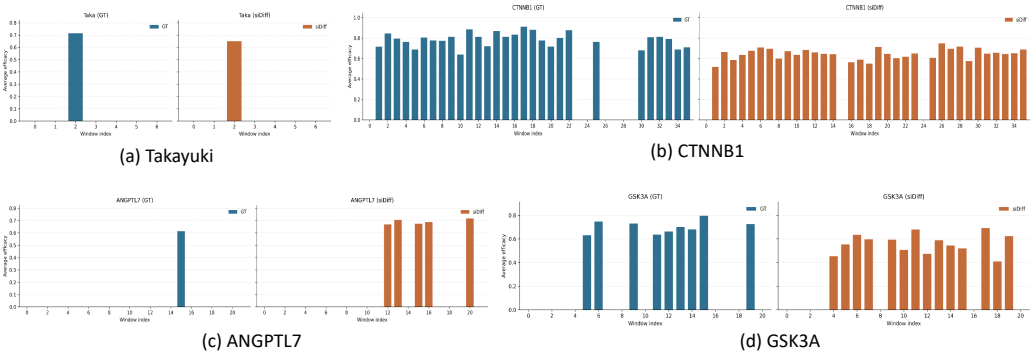


Figure 3: **Regional efficacy-window coverage.** Window-level mean inhibition for experimentally supported (GT) and siDiff-recovered high-efficacy regions on (a) Takayuki, (b) CTNNB1, (c) ANGPTL7, and (d) GSK3A. Windows that do not satisfy the high-efficacy criteria are displayed as zero.

Table 3: Sensitivity to the guidance gates g_0 and g_1 on four datasets.

Gates		Takayuki			ANGPTL7			CTNNB1			GSK3A		
g_0	g_1	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC	$\mathrm{Hit}_0@1$	$\mathrm{Hit}_3@3$	EWC
4	10	0.129 ± 0.014	0.555 ± 0.045	1/1	0.194 ± 0.056	0.559 ± 0.081	1/1	0.189 ± 0.013	0.636 ± 0.007	27/29	0.197 ± 0.027	0.645 ± 0.027	6/9
4	12	0.133 ± 0.008	0.594 ± 0.026	1/1	0.226 ± 0.056	0.591 ± 0.019	1/1	0.182 ± 0.033	0.621 ± 0.009	26/29	0.214 ± 0.052	0.624 ± 0.030	9/9
4	14	0.117 ± 0.008	0.596 ± 0.031	1/1	0.204 ± 0.081	0.559 ± 0.067	1/1	0.163 ± 0.006	0.587 ± 0.064	29/29	0.218 ± 0.056	0.620 ± 0.059	9/9
4	16	0.101 ± 0.011	0.501 ± 0.022	1/1	0.204 ± 0.075	0.570 ± 0.049	1/1	0.164 ± 0.009	0.613 ± 0.006	29/29	0.218 ± 0.039	0.637 ± 0.041	8/9
5	10	0.181 ± 0.019	0.579 ± 0.026	1/1	0.183 ± 0.019	0.570 ± 0.049	1/1	0.190 ± 0.011	0.697 ± 0.023	29/29	0.201 ± 0.039	0.684 ± 0.032	9/9
5	$12^{\dagger}$	0.249 ± 0.008	0.651 ± 0.024	1/1	0.182 ± 0.049	0.612 ± 0.037	1/1	0.192 ± 0.017	0.654 ± 0.009	26/29	0.120 ± 0.034	0.681 ± 0.022	8/9
5	14	0.161 ± 0.026	0.567 ± 0.050	1/1	0.194 ± 0.032	0.602 ± 0.019	1/1	0.176 ± 0.017	0.682 ± 0.025	28/29	0.162 ± 0.032	0.650 ± 0.027	9/9
5	16	0.131 ± 0.005	0.520 ± 0.029	1/1	0.194 ± 0.000	0.516 ± 0.000	1/1	0.180 ± 0.010	0.657 ± 0.004	27/29	0.175 ± 0.015	0.667 ± 0.022	8/9
6	10	0.140 ± 0.011	0.558 ± 0.008	1/1	0.215 ± 0.037	0.634 ± 0.019	1/1	0.182 ± 0.015	0.632 ± 0.019	25/29	0.162 ± 0.007	0.658 ± 0.027	6/9
6	12	0.153 ± 0.016	0.543 ± 0.026	1/1	0.183 ± 0.019	0.581 ± 0.097	1/1	0.182 ± 0.004	0.621 ± 0.032	28/29	0.162 ± 0.007	0.667 ± 0.013	7/9
6	14	0.138 ± 0.022	0.508 ± 0.076	1/1	0.204 ± 0.019	0.624 ± 0.067	1/1	0.169 ± 0.008	0.616 ± 0.022	29/29	0.188 ± 0.020	0.645 ± 0.015	9/9
6	16	0.105 ± 0.009	0.482 ± 0.055	1/1	0.204 ± 0.049	0.591 ± 0.019	1/1	0.185 ± 0.006	0.620 ± 0.047	29/29	0.162 ± 0.049	0.628 ± 0.022	9/9

without collapsing to a single locus. ANGPTL7 also contains several additional predicted windows beyond its sole annotated high-efficacy region, illustrating that EWC measures regional recall rather than precision. We therefore interpret EWC jointly with $\mathrm{Hit}_0@1$ and $\mathrm{Hit}_3@3$ rather than as a standalone measure of candidate quality.

4.5 Ablation Study

Table 2 shows that position-only and base-only guidance improve mean $\mathrm{Hit}_3@3$ over no guidance by 0.179 and 0.174, respectively, confirming that the two signals contribute complementary information. Full-trajectory guidance still recovers $25/29$ CTNNB1 and $6/9$ GSK3A windows, but its lower positional accuracy supports delaying base guidance until partial states become informative. Clustering further improves mean $\mathrm{Hit}_0@1$ and $\mathrm{Hit}_3@3$ by 0.150 and 0.183 over siDiff without clustering, while recovering seven additional CTNNB1 windows and three additional GSK3A windows. Thus both coarse-to-fine scheduling and spatially balanced selection are necessary for the final accuracy–coverage trade-off.

4.6 Sensitivity to the Guidance Gates

The gates g_0 and g_1 control two complementary aspects of guided sampling. We evaluate different combinations of the two gates in Table 3.

The results show that the method remains effective across a range of gate configurations. Default configuration is 5 and 12 for g_0 and g_1, respectively.

Table 4: Comparison of guidance-mode variants on four external datasets. All variants use $g_0 = 5$ and $g_1 = 12$.

	Takayuki			ANGPTL7			CTNNB1			GSK3A		
Guidance mode	$\text{Hit}_0@1$	$\text{Hit}_3@3$	EWC	$\text{Hit}_0@1$	$\text{Hit}_3@3$	EWC	$\text{Hit}_0@1$	$\text{Hit}_3@3$	EWC	$\text{Hit}_0@1$	$\text{Hit}_3@3$	EWC
Position–Constant	0.135	0.586	1/1	0.226	0.677	1/1	0.178	0.645	28/29	0.205	0.679	8/9
Position–Linear	0.226	0.620	1/1	0.194	0.613	1/1	0.190	0.657	27/29	0.167	0.590	8/9
Position–Larger Max	0.181	0.595	1/1	0.159	0.616	1/1	0.190	0.653	29/29	0.244	0.679	9/9
Base–Constant	0.089	0.537	1/1	0.194	0.581	1/1	0.161	0.653	26/29	0.192	0.667	7/9
Base–Quadratic	0.137	0.552	1/1	0.194	0.581	1/1	0.161	0.641	29/29	0.180	0.679	9/9
Base–Larger Max	0.216	0.540	1/1	0.194	0.563	1/1	0.194	0.636	29/29	0.154	0.628	9/9
Position–Quadratic, Base–Linear	0.242	0.664	1/1	0.185	0.612	1/1	0.200	0.654	28/29	0.101	0.679	9/9

4.7 Comparison of Guidance Modes

To examine whether the sampling behavior depends on a particular guidance schedule, we compare several position- and nucleotide-level guidance variants while keeping the generator, efficacy predictor, gate configuration, and candidate-processing pipeline unchanged. The default configuration uses quadratic position guidance with a maximum weight of 0.5 and linear nucleotide-level guidance with a maximum weight of 1.0. For position guidance, we additionally test a constant schedule with weight 0.5, a linear schedule with maximum weight 0.5, and a stronger quadratic schedule whose maximum weight is increased to 1.0. For nucleotide-level guidance, we compare the default linear schedule against a constant schedule with weight 1.0, a quadratic schedule with maximum weight 1.0 and exponent 2, and a stronger linear schedule with maximum weight 2.0. Each variant modifies only one guidance component, while the other remains at its default setting. The results are summarized in Table 4.

Overall, the default combination of quadratic position guidance and linear nucleotide-level guidance provides the most stable trade-off between preserving the learned generative prior, introducing efficacy preference, and maintaining candidate coverage. Constant schedules apply guidance uniformly throughout the corresponding sampling phase, whereas linear and quadratic schedules progressively strengthen the guidance signal. Increasing the maximum weight further emphasizes efficacy-driven selection, but may reduce the balance between guided optimization and generative exploration.

5 Conclusion

We presented siDiff, an efficacy-guided discrete diffusion framework for target-conditioned siRNA target-region proposal. Under scarce high-quality experimental supervision, siDiff combines a learned generative prior with sampling-time efficacy guidance to reshape candidate proposals on unseen genes. Coarse-to-fine steering and cluster-aware candidate derivation jointly balance target-site accuracy with multi-region recall, while reverse-complement construction preserves duplex validity. Consistent gains on public and patent-derived targets demonstrate the value of this proposal-and-selection strategy beyond the training gene distribution.

Ethics Statement

This work develops a computational method for proposing potentially effective siRNA target sites. All experiments use previously published datasets or assay records curated from publicly accessible sources; no new experiments involving human participants or animals were conducted, and no personally identifiable information was used. The candidates produced by siDiff are computational hypotheses rather than clinically validated therapeutics. The present framework does not explicitly optimize off-target effects, immunogenicity, chemical modification patterns, delivery properties, or full-transcript structural accessibility. Consequently, generated candidates require downstream off-target and safety screening, chemical-design assessment, and experimental validation before any therapeutic use. Patent-derived data should be redistributed only when permitted by the applicable source and licensing conditions.

REPRODUCIBILITY STATEMENT

The main paper specifies the model architecture, training objectives, guided sampling mechanism, evaluation protocol, and principal optimization settings. The supplementary material provides the dataset curation and preprocessing procedures, benchmark statistics, complete inference pipeline and notation, theoretical assumptions and proofs, the detailed definition of Efficacy-Window Coverage, predictor architectures and training procedures, computational resources, and additional sensitivity and ablation results. The diffusion model and efficacy predictors are trained only on the Huesken-derived training data, while Takayuki and the three patent-derived targets are held out from model training and hyperparameter selection. The publicly available implementations of the enumeration baselines and the reported experimental configurations further support independent reproduction of the comparisons.

AI USE STATEMENT

Generative AI tools were used solely to improve the clarity, grammar, and presentation of author-written text. They were not used to formulate the research questions, develop the methodology or theoretical results, design or conduct experiments, generate or process data, implement the proposed method, or analyze and interpret the experimental results. All AI-assisted edits were reviewed and verified by the authors, who take full responsibility for the final content of this work.

REFERENCES

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, pp. 17981–17993, 2021.

Yilan Bai, Haochen Zhong, Taiwei Wang, and Zhi John Lu. Oligoformer: an accurate and robust prediction method for sirna design. *Bioinformatics*, 40(10):btae577, 2024.

Sarah M Davis, Samuel Hildebrand, Hannah J MacMillan, Kathryn R Monopoli, Julianna Buchwald, Jacquelyn Sousa, David Cooper, Socheata Ly, Dimas Echeverria, Nicholas McHugh, et al. Systematic analysis of sirna and mrna features impacting fully chemically modified sirna efficacy. *Nucleic acids research*, 53(12):gkaf479, 2025.

Ye Han, Fei He, Yongbing Chen, Yuanning Liu, and Helong Yu. Sirna silencing efficacy prediction based on a deep architecture. *BMC genomics*, 19(Suppl 7):669, 2018.

Fei He, Ye Han, Jianting Gong, Jiazhi Song, Han Wang, and Yanwen Li. Predicting sirna efficacy based on multiple selective sirna representations and their combination at score level. *Scientific Reports*, 7(1):44836, 2017.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022.

Dieter Huesken, Joerg Lange, Craig Mickanin, Jan Weiler, Fred Asselbergs, Justin Warner, Brian Meloon, Sharon Engel, Avi Rosenberg, Dalia Cohen, et al. Design of a genome-wide sirna library using an artificial neural network. *Nature biotechnology*, 23(8):995–1001, 2005.

Takayuki Katoh and Tsutomu Suzuki. Specific residues at every third position of sirna shape its efficient rnai activity. *Nucleic acids research*, 35(4):e27, 2007.

Anastasia Khvorova, Angela Reynolds, and Sumedha D Jayasena. Functional sirnas and mirnas exhibit strand bias. *Cell*, 115(2):209–216, 2003.

Massimo La Rosa, Antonino Fiannaca, Laura La Paglia, and Alfonso Urso. A graph neural network approach for the analysis of sirna-target biological networks. *International Journal of Molecular Sciences*, 23(22):14211, 2022.

X Li et al. Parallelism and generation order in masked diffusion language models: Limits today, potential tomorrow. *arXiv preprint*, 2026.

Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. In *Advances in Neural Information Processing Systems*, volume 35, pp. 4328–4343, 2022.

Rongzhuo Long, Ziyu Guo, Da Han, Boxiang Liu, Xudong Yuan, Guangyong Chen, Pheng-Ann Heng, and Liang Zhang. sirnadiscovery: a graph neural network for sirna efficacy prediction via deep rna sequence analysis. *Briefings in Bioinformatics*, 25(6):bbae563, 2024.

Kathryn R Monopoli, Dmitry Korkin, and Anastasia Khvorova. Asymmetric trichotomous partitioning overcomes dataset limitations in building machine learning models for predicting sirna efficacy. *Molecular Therapy Nucleic Acids*, 33:93–109, 2023.

Yuki Naito, Jun Yoshimura, Shinichi Morishita, and Kumiko Ui-Tei. sidirect 2.0: updated software for designing functional sirna with reduced seed-dependent off-target effect. *BMC bioinformatics*, 10(1):392, 2009.

Inderbir S Padda, Arun U Mahtani, Preeti Patel, and Mayur Parmar. Small interfering rna (sirna) therapy. In *StatPearls [Internet]*. StatPearls Publishing, 2024.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4199–4209, 2023.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, 2022.

Subham Sekhar Sahoo, Albert Arriola, Yair Schiff, Aaron Gokaslan, Alexander Brooks, Andrej Risteski, and Alexander M Rush. Simple and effective masked diffusion language models. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Liangjiang Wang, Caiyan Huang, and Jack Y Yang. Predicting sirna potency with random forests and support vector machines. *BMC genomics*, 11(Suppl 3):S2, 2010.

Ying Xiong, Aowen Wang, Yu Kang, Chao Shen, Chang-Yu Hsieh, and Tingjun Hou. mrnabert: advancing mrna sequence design with a universal language model and comprehensive dataset. *Nature Communications*, 16(1):10371, 2025.

Honggen Zhang, Xiangrui Gao, and Lipeng Lai. sidpt: sirna efficacy prediction via debiased preference-pair transformer. *arXiv preprint arXiv:2509.15664*, 2025a.

Honggen Zhang, Xiangrui Gao, June Zhang, and Lipeng Lai. mrna2vec: mrna embedding with language model in the 5'utr-cds for mrna design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 1057–1065, 2025b.

Table 5: Statistics of the de novo siRNA design datasets.

Dataset	Raw records	Active sites	Context views
Takayuki	702	190	943
ANGPTL7	111	31	155
CTNNB1	330	242	1210
GSK3A	122	78	390
Total	1265	541	2698

SUPPLEMENTARY MATERIAL

A DATA CURATION

To enable robust validation across diverse data distributions, we constructed three novel datasets curated from patent repositories. However, a challenge with existing patent data is that *in vitro* screening is typically performed on chemically modified siRNAs. While modifications are essential for preserving structural stability and avoiding immune activation *in vivo*, they complicate the data. Consequently, the reported final inhibition rate is a joint function of both the naked siRNA sequence (which dictates target binding affinity) and the specific chemical modification pattern applied.

Let x_i denote the i-th naked siRNA sequence and y_j denote the j-th chemical modification pattern. The observed inhibition rate reported in the patent data can be formalized as:

$$e_{i,j} = E(x_i, y_j) \tag{17}$$

where E represents the underlying efficacy function.

To construct a dataset focused solely on the inherent binding efficacy of the naked sequences, we must decouple the sequence effect from the modification effect. Theoretically, while chemical modifications enhance systemic stability, they often introduce steric hindrance or thermodynamic changes that slightly attenuate the fundamental silencing efficacy compared to the unmodified sequence. Therefore, we assume that the true efficacy of the naked sequence, denoted as e_i^*, establishes an upper bound for any of its modified variants:

$$e_i^* \geq E(x_i, y_j) \quad \forall j \tag{18}$$

Because the true naked efficacy e_i^* is rarely reported directly in these records, we approximate it using the available data. Let $\mathcal{Y}_i = \{y_k\}$ represent the empirical subset of modification patterns applied to sequence x_i within our dataset. We define our proxy label for the naked siRNA efficacy, $\hat{e}_i$, as the maximum observed efficacy across this sampled modification space:

$$\hat{e}_i = \max_{y_k \in \mathcal{Y}_i} E(x_i, y_k) \tag{19}$$

We focus our evaluation on three curated patent datasets targeting ANGPTL7, CTNNB1, and GSK3A, selected for their extensive experimental validation across diverse mRNA regions. By applying our maximum efficacy approximation to aggregate identical binding sites, the CTNNB1 dataset is distilled from 944 initial sequence-modification combinations (x_i, y_k) down to 330 unique naked sequences (x_i). In contrast, for the ANGPTL7 and GSK3A datasets, all candidate siRNAs within a given target were evaluated using a uniform chemical modification pattern, denoted as (x_i, y_0). Because y_0 is fixed for these targets, the modification space $\mathcal{Y}_i$ consists of only a single pattern, allowing the reported inhibitions to map directly to our naked sequence proxy.

A.1 DATASET STATISTICS

Table 5 reports the number of raw measurements, retained high-activity sites, and augmented context views used for external evaluation.

B NOTATION AND COMPLETE PIPELINE OVERVIEW

This section follows the notation introduced in the main paper and provides a complete overview of the siDiff inference pipeline. We particularly clarify the distinction between the partial active pattern produced by guided diffusion and the complete siRNA candidates returned by the full pipeline.

B.1 NOTATION

Let

$$\mathcal{A} = \{\mathrm{A}, \mathrm{U}, \mathrm{G}, \mathrm{C}\} \tag{20}$$

denote the nucleotide alphabet, and let

$$\mathcal{A}_{\mathrm{M}} = \mathcal{A} \cup \{[\texttt{MASK}]\} \tag{21}$$

denote the augmented vocabulary used by the masked discrete diffusion model.

Following the problem formulation in the main paper, we use

$$\mathbf{m} = (m_1, \ldots, m_L) \in \mathcal{A}^L \tag{22}$$

to denote an mRNA sequence or context supplied to the model.

A clean siRNA sequence is denoted by

$$\mathbf{s}_0 = (s_0^1, \ldots, s_0^{L_s}) \in \mathcal{A}^{L_s}, \qquad L_s = 21. \tag{23}$$

The partially masked state at diffusion timestep t is denoted by

$$\mathbf{s}_t \in \mathcal{A}_{\mathrm{M}}^{L_s}. \tag{24}$$

The forward diffusion process terminates at the fully masked state

$$\mathbf{s}_T = [\texttt{MASK}]^{L_s}, \tag{25}$$

where T is the diffusion horizon.

In the training objective, the set of masked positions is

$$\mathcal{M}_t = \left\{k \mid s_t^k = [\texttt{MASK}]\right\}, \tag{26}$$

consistent with Eq. (6) of the main paper. During guided inference, we use

$$\mathcal{U}_t = \left\{i \mid s_t^i = [\texttt{MASK}]\right\} \tag{27}$$

to denote the current set of unresolved positions. The two sets have the same token-level definition but are used in the training and sampling contexts, respectively.

The coarse-to-fine sampling process is controlled by two gates:

$$T \geq g_1 > g_0 \geq 0. \tag{28}$$

Active-pattern sampling is performed for

$$t = T, T-1, \ldots, g_0 + 1. \tag{29}$$

Position-level efficacy guidance is controlled by $w_{\mathrm{pos}}(t)$, whereas base-level efficacy guidance is controlled by $\beta_{\mathrm{base}}(t)$, which is disabled for $t > g_1$ and enabled for $g_0 < t \leq g_1$.

At timestep t, the conditional diffusion generator produces

$$\mathbf{Z}_t^{\mathrm{diff}} = G_\theta(\mathbf{m}, \mathbf{s}_t, t), \tag{30}$$

where

$$\mathbf{z}_{t,i}^{\mathrm{diff}} \in \mathbb{R}^4 \tag{31}$$

denotes the four nucleotide logits at position i. For every unresolved position $i \in \mathcal{U}_t$, the generation confidence is

$$C_{t,i}^{\mathrm{gen}} = \max_{b \in \mathcal{A}} \mathrm{softmax}\left(\mathbf{z}_{t,i}^{\mathrm{diff}}\right)_b. \tag{32}$$

Table 6: Principal notation used in the supplementary material.

Symbol	Definition
$\mathcal{A}$	Nucleotide alphabet $\{\mathrm{A}, \mathrm{U}, \mathrm{G}, \mathrm{C}\}$
$\mathcal{A}_{\mathrm{M}}$	Augmented alphabet containing [MASK]
$\mathbf{m}$	mRNA sequence or context supplied to the model
$\mathbf{s}_0$	Clean, complete 21-nt siRNA sequence
$\mathbf{s}_t$	Partially masked siRNA state at timestep t
$\mathcal{M}_t$	Masked positions used in the diffusion training objective
$\mathcal{U}_t$	Unresolved positions during guided sampling
T	Diffusion horizon
g_1	Timestep at which base-level guidance becomes active
g_0	Terminal boundary of active-pattern sampling
G_θ	Conditional masked diffusion generator
F_ϕ	Partial-state efficacy predictor used during guided sampling
$\mathbf{z}_{t,i}^{\mathrm{diff}}$	Diffusion logits over the four nucleotides at position i
$C_{t,i}^{\mathrm{gen}}$	Diffusion generation confidence at position i
$E_{t,i,b}$	Predicted efficacy after filling position i with nucleotide b
$I_{t,i}^{\mathrm{eff}}$	Position-level efficacy potential $\max_b E_{t,i,b}$
$\widetilde{C}_{t,i}$	Modified position confidence after efficacy guidance
$w_{\mathrm{pos}}(t)$	Position-level efficacy-guidance weight
$\beta_{\mathrm{base}}(t)$	Base-level efficacy-guidance weight
j_t	Position selected for unmasking at timestep t
$\mathbf{s}^{\mathrm{act}}$	Partial active-pattern proposal produced by guided diffusion
$\mathcal{S}_{\mathrm{cand}}$	Set of complete siRNA candidates derived from $\mathbf{s}^{\mathrm{act}}$
$\mathcal{G}_k$	The k-th cluster of candidate target sites
q_a	Predicted efficacy score assigned to complete candidate a
ω	Warm-up proportion of the position-guidance schedule
τ_{pos}	Temperature used by Gumbel-Top-1 position selection
Δ_{clu}	Maximum candidate-cluster span, defined as $L_s - g_0$

The partial-state efficacy predictor in Algorithm 1 of the main paper is denoted by F_ϕ. For every unresolved position i and hypothetical nucleotide $b \in \mathcal{A}$, it computes

$$E_{t,i,b} = F_\phi\left(\mathbf{m}, \mathrm{Fill}(\mathbf{s}_t, i, b)\right). \tag{33}$$

The position-level efficacy potential is

$$I_{t,i}^{\mathrm{eff}} = \max_{b \in \mathcal{A}} E_{t,i,b}. \tag{34}$$

The four scores $\{E_{t,i,b}\}_{b \in \mathcal{A}}$ are cached within each timestep and shared by position-level and base-level guidance.

Table 6 summarizes the principal notation.

B.2 Complete Pipeline Overview

The complete siDiff inference pipeline consists of three functional components: efficacy-guided active-pattern proposal, candidate derivation and clustering, and inhibition-aware candidate evaluation and sampling. Importantly, the partial sequence produced by guided diffusion is not directly returned as the final siRNA. Instead, it is used as an active pattern for proposing target regions. Every complete siRNA returned by the pipeline remains the reverse complement of a valid 21-nt mRNA target site.

B.2.1 Efficacy-Guided Active-Pattern Proposal

Starting from

$$\mathbf{s}_T = [\mathtt{MASK}]^{L_s}, \tag{35}$$

siDiff resolves one position at each reverse timestep. The diffusion generator first computes $\mathbf{z}_{t,i}^{\mathrm{diff}}$ and $C_{t,i}^{\mathrm{gen}}$ for all $i \in \mathcal{U}_t$. The efficacy predictor then evaluates the four hypothetical nucleotide assignments in Eq. equation 33 and computes $I_{t,i}^{\mathrm{eff}}$.

Generation confidence and efficacy potential are independently normalized over the current unresolved-position set. Following Eq. (11) of the main paper, the modified position confidence is

$$\begin{aligned}\widetilde{C}_{t,i} = &\ (1 - w_{\mathrm{pos}}(t))\,\mathrm{Norm}_{\mathcal{U}_t}\left(C_{t,i}^{\mathrm{gen}}\right) \\ &+ w_{\mathrm{pos}}(t)\,\mathrm{Norm}_{\mathcal{U}_t}\left(I_{t,i}^{\mathrm{eff}}\right).\end{aligned} \tag{36}$$

The position-guidance weight follows a quadratic progressive schedule. We first define the overall sampling progress as

$$r_{\mathrm{all}}(t) = \frac{T - t}{T - g_0}. \tag{37}$$

Let $\omega \in [0, 1)$ denote the warm-up proportion. The normalized post-warm-up progress is

$$r_{\mathrm{pos}}(t) = \begin{cases} 0, & r_{\mathrm{all}}(t) \le \omega, \\ \dfrac{r_{\mathrm{all}}(t) - \omega}{1 - \omega}, & r_{\mathrm{all}}(t) > \omega. \end{cases} \tag{38}$$

The normalized position-guidance schedule is

$$\overline{w}_{\mathrm{pos}}(t) = r_{\mathrm{pos}}(t)^2. \tag{39}$$

The implementation scales this normalized schedule by the configured position-guidance strength. Thus, position guidance is inactive during warm-up and increases quadratically afterward.

To retain stochastic exploration, the next position is selected using the Gumbel-Top-1 rule:

$$j_t = \arg\max_{i \in \mathcal{U}_t} \left\{ \frac{\widetilde{C}_{t,i}}{\tau_{\mathrm{pos}}} + G_i \right\}, \qquad G_i \overset{\text{i.i.d.}}{\sim} \mathrm{Gumbel}(0, 1), \tag{40}$$

where $\tau_{\mathrm{pos}} > 0$ controls the exploration strength. A standard Gumbel random variable can be sampled as

$$G_i = -\log[-\log(U_i)], \qquad U_i \sim \mathrm{Uniform}(0, 1). \tag{41}$$

Its cumulative distribution function is

$$F_G(g) = \exp[-\exp(-g)], \tag{42}$$

and its probability density function is

$$f_G(g) = \exp[-g - \exp(-g)]. \tag{43}$$

The Gumbel distribution is the Type-I limiting distribution for maxima of independent random variables under standard regularity conditions. This extreme-value property makes it suitable for stochastic Top-1 selection. By the Gumbel-Max identity,

$$\Pr(j_t = i) = \frac{\exp(\widetilde{C}_{t,i}/\tau_{\mathrm{pos}})}{\sum_{k \in \mathcal{U}_t} \exp(\widetilde{C}_{t,k}/\tau_{\mathrm{pos}})}. \tag{44}$$

Thus, positions with larger modified confidence are more likely to be selected, while the Gumbel perturbation maintains stochastic exploration over different unmasking orders. For $t > g_1$, the nucleotide at j_t is sampled using only the diffusion logits:

$$\mathbf{z}_{t,j_t}^{\mathrm{joint}} = \mathbf{z}_{t,j_t}^{\mathrm{diff}}. \tag{45}$$

For $g_0 < t \le g_1$, the cached four-nucleotide efficacy scores are standardized:

$$\widetilde{\mathbf{E}}_{t,j_t,:} = \mathrm{Norm}\left(\mathbf{E}_{t,j_t,:}\right), \tag{46}$$

and incorporated into the nucleotide logits:

$$\mathbf{z}_{t,j_t}^{\mathrm{joint}} = \mathbf{z}_{t,j_t}^{\mathrm{diff}} + \beta_{\mathrm{base}}(t)\widetilde{\mathbf{E}}_{t,j_t,:}. \tag{47}$$

The normalized base-guidance schedule is

$$\overline{\beta}_{\mathrm{base}}(t) = \begin{cases} 0, & t > g_1, \\ \dfrac{g_1 - t + 1}{g_1 - g_0}, & g_0 < t \le g_1. \end{cases} \tag{48}$$

The implementation scales this normalized schedule by the configured base-guidance strength. Base-level guidance is therefore disabled for $t > g_1$ and increases linearly from $t = g_1$ to the final active-pattern sampling step $t = g_0 + 1$. The nucleotide is then decoded from $\mathbf{z}_{t,j_t}^{\mathrm{joint}}$. After the final sampling step $t = g_0 + 1$, the remaining partial sequence is retained as the active pattern $\mathbf{s}^{\mathrm{act}}$.

B.2.2 Candidate Derivation and Clustering

The active span of $\mathbf{s}^{\mathrm{act}}$ is determined by its leftmost and rightmost instantiated nucleotides. Instantiated nucleotides are aligned to the target mRNA under reverse complementarity, while the remaining [MASK] positions are treated as unconstrained.

Matches satisfying the predefined mismatch tolerance are expanded into valid 21-nt mRNA target windows. Each target window is then converted into a complete siRNA through reverse complementation, yielding the complete candidate set $\mathcal{S}_{\mathrm{cand}}$.

Candidates are sorted according to their target start positions and greedily grouped into spatial clusters. Following Eq. (17) of the main paper, a candidate with start position x is assigned to the current cluster $\mathcal{C}$ if

$$x - \min_{y \in \mathcal{C}} y \leq \Delta_{\mathrm{clu}}, \tag{49}$$

where

$$\Delta_{\mathrm{clu}} = L_s - g_0 = 21 - g_0. \tag{50}$$

This choice connects the maximum cluster span to the number of positions resolved during active-pattern sampling. Otherwise, the candidate initializes a new cluster. Bounding the total cluster span preserves local alternatives while preventing spatially distant candidates from being merged through transitive chaining.

B.2.3 Inhibition-Aware Candidate Evaluation and Sampling

Following the terminology of the main paper, we refer to the inhibition predictor used in the *Inhibition-Aware Candidate Evaluation and Sampling* subsection as the full-sequence predictor. Unlike the partial-state efficacy predictor F_ϕ, this predictor is applied after candidate derivation and clustering. It takes a clean, fully specified mRNA-context–siRNA pair as input and assigns a predicted efficacy score q_a to each complete candidate $a \in \mathcal{S}_{\mathrm{cand}}$.

For every cluster $\mathcal{G}_k$, the mean predicted efficacy and high-efficacy ratio are

$$\mu_k = \frac{1}{|\mathcal{G}_k|} \sum_{a \in \mathcal{G}_k} q_a, \qquad \rho_k = \frac{1}{|\mathcal{G}_k|} \sum_{a \in \mathcal{G}_k} \mathbb{I}[q_a \geq \gamma], \tag{51}$$

consistent with Eq. (18) of the main paper.

The three-grade staircase strategy assigns progressively smaller candidate quotas

$$K_1 > K_2 > K_3 \tag{52}$$

to clusters satisfying progressively relaxed thresholds on μ_k and ρ_k. Clusters that fail the final grade are removed. Within each retained cluster, candidates are ranked by q_a and selected according to the assigned quota. This produces a final candidate set that balances predicted inhibition efficacy with coverage across distinct mRNA regions.

Algorithm 1 summarizes the complete inference pipeline.

C Local Analysis of Efficacy-Guidance Scheduling

We provide a local analysis of nucleotide-level guidance that explains why base guidance should be weak under heavy masking and strengthened as reverse diffusion reveals more sequence context. The analysis concerns $\beta_{\mathrm{base}}(t)$; position selection may remain guided through $w_{\mathrm{pos}}(t)$ as in the main paper.

C.1 Guided trajectory model

Fix an mRNA context $\mathbf{m}$ and index the reverse process in chronological order by $k = 1, \ldots, K$, with reverse-time label $t_k = T - k + 1$. Let H_{k-1} denote the sampling history before step k, including the current partially masked siRNA state, and let J_k be the position selected using the position policy in the main paper. Conditional on (H_{k-1}, J_k), the unguided diffusion model assigns base probabilities

$$p_k(b \mid H_{k-1}, J_k) > 0, \qquad b \in \mathcal{A} = \{\mathrm{A}, \mathrm{C}, \mathrm{G}, \mathrm{U}\}. \tag{53}$$

Algorithm 1 Complete siDiff Inference Pipeline

Require: mRNA context $\mathbf{m}$; siRNA length L_s; diffusion horizon T
Require: gates $T \geq g_1 > g_0$; generator G_θ; partial-state efficacy predictor F_ϕ
Ensure: Selected complete siRNA candidates

```
s ← [MASK]^{L_s}
for t = T, T − 1, . . . , g_0 + 1 do
    Z_t^diff ← G_θ(m, s, t)
    U_t ← {i | s_i = MASK}
    C_{t,i}^gen ← max_{b∈A} softmax(Z_t^diff)_{i,b},   ∀i ∈ U_t
    E_{t,i,b} ← F_φ(m, Fill(s, i, b)),   ∀(i, b) ∈ U_t × A
    I_{t,i}^eff ← max_{b∈A} E_{t,i,b},   ∀i ∈ U_t
    Compute C̃_{t,i} using Eq. equation 36
    j_t ← GumbelTopOne({C̃_{t,i}}_{i∈U_t})
    z_{t,j_t}^joint ← z_{t,j_t}^diff
    if t ≤ g_1 then
        Ẽ_{t,j_t,:} ← Norm(E_{t,j_t,:})
        z_{t,j_t}^joint ← z_{t,j_t}^diff + β_base(t) Ẽ_{t,j_t,:}
    end if
    s_{j_t} ← Decode(z_{t,j_t}^joint)
end for
s^act ← ExtractActivePattern(s)
H ← ReverseComplementMatch(s^act, m)
S_cand ← ExpandToCompleteCandidates(H)
{G_k} ← ClusterByTargetPosition(S_cand)
for each candidate a ∈ S_cand do
    Construct its clean mRNA-context–siRNA pair
    Obtain q_a using the full-sequence predictor
end for
Compute (μ_k, ρ_k) for every cluster G_k
Apply staircase cluster filtering and within-cluster ranking
return Selected complete siRNA candidates
```

For a hypothetical assignment b, the partial-state predictor provides the score

$$E_k(b) = F_\phi(\mathbf{m}, \mathrm{Fill}(\mathbf{s}_{t_k}, J_k, b)) \,. \tag{54}$$

The standardized scores used in the implementation differ only by an affine transformation conditional on the current state. We therefore absorb that standardization into E_k and write the guided base kernel as

$$q_{k,\beta_k}(b \mid H_{k-1}, J_k) = \frac{p_k(b \mid H_{k-1}, J_k) e^{\beta_k E_k(b)}}{\sum_{b' \in \mathcal{A}} p_k(b' \mid H_{k-1}, J_k) e^{\beta_k E_k(b')}}, \qquad \beta_k = \beta_{\mathrm{base}}(t_k). \tag{55}$$

This is equivalent to adding $\beta_k E_k(b)$ to the diffusion logits before decoding.

Let P be the reference trajectory law with $\beta_k = 0$ and let $Q_{\boldsymbol{\beta}}$ be its base-guided counterpart, where $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_K)$. They share the same initial distribution, position-selection policy, and state-update rule, and differ only through equation 55. Let X be the terminal state of active-pattern sampling, which may retain masked positions, and let $Y = F_\phi(\mathbf{m}, X)$ be its terminal efficacy reward. Define

$$\begin{aligned} \overline{E}_k &= \mathbb{E}_P[E_k(B_k) \mid H_{k-1}, J_k], & M_k &= E_k(B_k) - \overline{E}_k, \\ a_k &= \mathbb{E}_P[Y M_k], & v_k &= \mathbb{E}_P[M_k^2]. \end{aligned} \tag{56}$$

Thus a_k measures alignment between the local predictor preference and the terminal reward, while v_k measures the local change in trajectory distribution induced by guidance.

Assumption C.1 (Local regularity)**.** *The trajectory space is finite, the probabilities in equation 53 are positive on reachable states, and predictor scores are finite. Guidance is analyzed locally as* $\|\boldsymbol{\beta}\|_2 \to 0$*. On every guided step,* $v_k > 0$*, and at least one* a_k *is positive.*

C.2 REWARD–DIVERGENCE TRADE-OFF

Lemma C.2 (Local trajectory expansion)**.** *Under Assumption C.1,*

$$\mathbb{E}_{Q_{\boldsymbol{\beta}}}[Y] - \mathbb{E}_P[Y] = \sum_{k=1}^{K} a_k \beta_k + O(\|\boldsymbol{\beta}\|_2^2), \tag{57}$$

$$\mathrm{KL}(Q_{\boldsymbol{\beta}} \| P) = \frac{1}{2} \sum_{k=1}^{K} v_k \beta_k^2 + O(\|\boldsymbol{\beta}\|_2^3). \tag{58}$$

Moreover, for the induced terminal distributions,

$$\mathrm{KL}(Q_{\boldsymbol{\beta},X} \| P_X) \le \mathrm{KL}(Q_{\boldsymbol{\beta}} \| P). \tag{59}$$

Proof. The factors shared by P and $Q_{\boldsymbol{\beta}}$ cancel, giving

$$\frac{dQ_{\boldsymbol{\beta}}}{dP} = \prod_{k=1}^{K} \frac{e^{\beta_k E_k(B_k)}}{\mathbb{E}_P[e^{\beta_k E_k(B_k)} \mid H_{k-1}, J_k]}. \tag{60}$$

Taylor expansion at $\boldsymbol{\beta} = \mathbf{0}$ yields $dQ_{\boldsymbol{\beta}}/dP = 1 + \sum_k \beta_k M_k + O(\|\boldsymbol{\beta}\|_2^2)$. Multiplying by Y and taking expectation under P proves equation 57.

For one conditional step, define

$$\psi_k(\beta) = \log \mathbb{E}_P\left[e^{\beta E_k(B_k)} \mid H_{k-1}, J_k\right]. \tag{61}$$

Since $\psi_k''(0) = \mathrm{Var}_P(E_k(B_k) \mid H_{k-1}, J_k)$,

$$\mathrm{KL}(q_{k,\beta_k} \| p_k) = \beta_k \psi_k'(\beta_k) - \psi_k(\beta_k) = \frac{\beta_k^2}{2} \mathrm{Var}_P(E_k(B_k) \mid H_{k-1}, J_k) + O(|\beta_k|^3). \tag{62}$$

Summing by the KL chain rule and taking the baseline expectation gives equation 58. Finally, X is a deterministic function of the trajectory, so equation 59 follows from the data-processing inequality for KL divergence. □

Lemma C.2 converts guidance scheduling into a local resource-allocation problem: increase predicted terminal efficacy while limiting deviation from the reference trajectory without base guidance.

Theorem C.3 (Locally optimal guidance allocation)**.** *Let $\mathcal{I}$ be the set of steps at which base guidance is allowed, and consider the quadratic KL budget $C > 0$:*

$$\max_{\beta_k \ge 0} \sum_{k \in \mathcal{I}} a_k \beta_k \quad \textit{s.t.} \quad \frac{1}{2} \sum_{k \in \mathcal{I}} v_k \beta_k^2 \le C, \qquad \beta_k = 0 \textit{ for } k \notin \mathcal{I}. \tag{63}$$

Writing $a_k^+ = \max(a_k, 0)$ and

$$A_{\mathcal{I}} = \sum_{k \in \mathcal{I}} \frac{(a_k^+)^2}{v_k}, \tag{64}$$

the unique optimum on the nonzero coordinates is

$$\beta_k^* = \sqrt{\frac{2C}{A_{\mathcal{I}}}} \frac{a_k^+}{v_k}, \qquad k \in \mathcal{I}, \tag{65}$$

and its leading reward gain is $\sqrt{2CA_{\mathcal{I}}}$. For the exact guided process as $C \downarrow 0$,

$$\mathbb{E}_{Q_{\boldsymbol{\beta}^*}}[Y] - \mathbb{E}_P[Y] = \sqrt{2CA_{\mathcal{I}}} + O(C),$$
$$\mathrm{KL}(Q_{\boldsymbol{\beta}^*} \| P) = C + O(C^{3/2}). \tag{66}$$

Proof. Coordinates with $a_k \le 0$ cannot improve the linear objective and therefore receive zero weight. For $a_k > 0$, set $y_k = \sqrt{v_k}\beta_k$ and $c_k = a_k/\sqrt{v_k}$. Cauchy–Schwarz gives

$$\sum_{k\in\mathcal{I}} a_k\beta_k = \sum_k c_k y_k \le \left(\sum_k \frac{a_k^2}{v_k}\right)^{1/2} \left(\sum_k v_k\beta_k^2\right)^{1/2} \le \sqrt{2CA_{\mathcal{I}}}. \tag{67}$$

Equality holds exactly when y_k is proportional to c_k and the budget is exhausted, which gives equation 65. Since $\|\boldsymbol{\beta}^*\|_2 = O(\sqrt{C})$, substituting it into equation 57–equation 58 proves equation 66. □

C.3 Why base guidance is delayed

Theorem C.3 shows that the relevant quantity is the stepwise efficacy-to-divergence ratio a_k/v_k, not predictor confidence alone. We now relate this ratio to the number of masked nucleotides.

Assumption C.4 (Mask-dependent predictor noise)**.** *The centered local score decomposes as $M_k = S_k^\circ + \varepsilon_k$, where*

$$\mathbb{E}[S_k^\circ \varepsilon_k] = 0, \qquad \mathbb{E}[Y\varepsilon_k] = 0, \qquad \mathbb{E}[\varepsilon_k^2] = \sigma^2(m_k). \tag{68}$$

Here m_k is the number of masked positions, $\sigma^2(m)$ is nondecreasing, and the compared steps have approximately homogeneous signal moments $\mathbb{E}[(S_k^\circ)^2] = u$ and $\mathbb{E}[YS_k^\circ] = c > 0$.

Corollary C.5 (Coarse-to-fine base guidance)**.** *Under Assumption C.4,*

$$a_k = c, \qquad v_k = u + \sigma^2(m_k), \qquad \beta_k^* \propto \frac{c}{u + \sigma^2(m_k)}. \tag{69}$$

Consequently, steps with more masked positions receive no larger optimal base-guidance weight. If early high-mask steps have negligible $A_{\text{early}} = \sum_{k\in\mathcal{I}_{\text{early}}} (a_k^+)^2/v_k$ relative to A_{late}, specifically $A_{\text{early}} \le \eta A_{\text{late}}$, removing early guidance loses at most

$$1 - \sqrt{\frac{A_{\text{late}}}{A_{\text{late}} + A_{\text{early}}}} \le \frac{\eta}{2} \tag{70}$$

of the optimal leading-order reward gain.

Proof. Orthogonality gives $v_k = \mathbb{E}[(S_k^\circ)^2] + \mathbb{E}[\varepsilon_k^2] = u + \sigma^2(m_k)$ and $a_k = \mathbb{E}[YS_k^\circ] + \mathbb{E}[Y\varepsilon_k] = c$. Equation equation 69 then follows from equation 65. Applying the optimal value in Theorem C.3 to all steps and to late steps only gives leading gains $\sqrt{2C(A_{\text{late}} + A_{\text{early}})}$ and $\sqrt{2CA_{\text{late}}}$, respectively. Their relative difference is the left side of equation 70; the final bound follows from $(1+\eta)^{-1/2} \ge 1 - \eta/2$. □

Connection to siDiff. Early reverse steps contain little sequence context, so the partial-state efficacy predictor has higher uncertainty and base guidance has a poorer efficacy-to-divergence ratio. This motivates setting $\beta_{\text{base}}(t) = 0$ for $t > g_1$. As masking decreases, the predictor becomes more informative and equation 69 supports increasing the guidance weight. The power ramp used by siDiff,

$$\beta_{\text{base}}(t) = \beta_{\max}\left(\frac{g_1 - t + 1}{g_1 - g_0}\right)^p, \qquad g_0 < t \le g_1, \tag{71}$$

is a monotone parametric approximation to this locally optimal allocation. The result is local in guidance strength and explains the schedule shape; it does not claim global optimality for arbitrary finite $\beta_{\max}$.

D Detailed Definition of Efficacy-Window Coverage

Efficacy-Window Coverage (EWC) measures whether a method recovers spatially diverse regions containing experimentally validated high-efficacy siRNA sites. For each evaluation mRNA, we partition its sequence into non-overlapping windows of width w. Specifically, for an mRNA of length L, the j-th window is

$$B_j = \{(j-1)w + 1, \ldots, \min(jw, L)\}. \tag{72}$$

Let $\mathcal{O}_j$ denote the set of experimentally measured siRNA target sites whose start positions fall within B_j, and let $y_i \in [0,1]$ be the experimentally measured efficacy of site i after per-mRNA min–max normalization. For each window, we compute

$$o_j = |\mathcal{O}_j|, \tag{73}$$

$$r_j = \frac{1}{o_j} \sum_{i \in \mathcal{O}_j} \mathbb{I}[y_i \geq \gamma], \tag{74}$$

and

$$a_j = \frac{1}{o_j} \sum_{i \in \mathcal{O}_j} y_i, \tag{75}$$

where o_j, r_j, and a_j respectively denote the number of measured sites, the proportion of high-efficacy sites, and the average measured efficacy within the window. A site is regarded as high efficacy when its normalized experimental efficacy satisfies $y_i \geq \gamma$.

The ground-truth high-efficacy window set is defined as

$$\mathcal{W}_{\mathrm{gt}} = \{j \mid o_j \geq o_{\min},\ r_j \geq r_{\min},\ a_j \geq a_{\min}\}. \tag{76}$$

The fixed parameter configuration used for all four external evaluation datasets is

$$w = 100, \qquad o_{\min} = 1, \qquad \gamma = 0.7, \qquad r_{\min} = 0.1, \qquad a_{\min} = 0.6. \tag{77}$$

After deduplicating the target sites produced by a method, we construct its predicted high-efficacy window set $\mathcal{W}_{\mathrm{pred}}$ analogously, replacing the experimental efficacy values with the efficacy scores produced by that method. Each method therefore ranks and filters its own candidates using its own output scores; scores from different predictors are neither pooled nor directly compared.

Finally, EWC is computed as

$$\mathrm{EWC} = \frac{|\mathcal{W}_{\mathrm{pred}} \cap \mathcal{W}_{\mathrm{gt}}|}{|\mathcal{W}_{\mathrm{gt}}|}. \tag{78}$$

We report EWC in the count form $|\mathcal{W}_{\mathrm{pred}} \cap \mathcal{W}_{\mathrm{gt}}| \,/\, |\mathcal{W}_{\mathrm{gt}}|$.

E Efficacy Predictors Used at Different Stages

Our framework employs two efficacy predictors with distinct roles and input representations. The first predictor provides lightweight guidance over partially masked siRNA states during the reverse diffusion process, whereas the second evaluates complete candidates after candidate derivation and clustering. Both predictors are trained on the Huesken dataset, but they are trained separately and are not used interchangeably.

E.1 Partial-State Efficacy Predictor for Guided Sampling

The partial-state efficacy predictor F_ϕ serves as the guidance model during reverse diffusion. Its input consists of a 21-nt mRNA target segment and a 21-nt siRNA state. To make the model applicable to intermediate diffusion states, partially masked siRNA inputs are constructed from the clean siRNAs during training, while the corresponding experimentally measured inhibition values are retained as supervision.

The architecture of the partial-state efficacy predictor is illustrated in Figure 4. The predictor is trained using the mean squared error (MSE) loss between the predicted and experimentally measured inhibition efficacy values.

During generation, F_ϕ is repeatedly applied to the current partially decoded siRNA state to estimate the efficacy changes associated with candidate position and nucleotide assignments. We therefore use the 21-nt target segment without additional flanking context, which keeps the guidance model lightweight and makes repeated predictor evaluations during sampling computationally affordable.

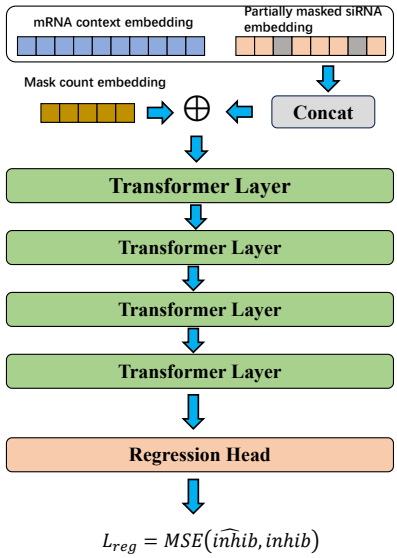


Figure 4: Architecture of the partial-state efficacy predictor F_ϕ.

E.2 Full-Sequence Predictor for Candidate Evaluation

After candidate derivation and clustering, complete candidates are evaluated using the full-sequence predictor used in the *Inhibition-Aware Candidate Evaluation and Sampling* subsection of the main paper. This predictor takes a clean and complete mRNA-context–siRNA pair as input. Specifically, the mRNA context and siRNA have lengths of 31 and 21 nucleotides, respectively. The 31-nt mRNA input retains limited sequence context surrounding the target site while conforming to the fixed input-window length of 61 used by the predictor.

This predictor is also trained on the Huesken dataset. For each complete candidate a, it produces the efficacy score q_a, which is subsequently used for candidate ranking and sampling within each cluster. In contrast to F_ϕ, this predictor is applied only to the substantially reduced candidate set obtained after clustering and can therefore use the more contextualized input representation.

Distinction between the two predictors. The two predictors operate at different stages of the pipeline: F_ϕ guides local decisions over partially masked states, whereas the full-sequence predictor evaluates clean candidates using additional mRNA context. Their output scores are used only within their respective stages and are not directly compared across the two models.

F Masked Diffusion siRNA Model Training

Training data. The diffusion generator was trained exclusively on the Huesken-derived dataset. As described in the *Datasets* paragraph of the main paper, we retained siRNAs with normalized inhibition values of at least 0.7 and applied random-window context augmentation while preserving the complete 21-nt target site within each 61-nt mRNA context. This procedure produced 3,860 mRNA-context–siRNA training instances. We randomly split these instances into training and validation sets using an $80\%/20\%$ ratio, resulting in 3,088 training instances and 772 validation instances. The Takayuki and patent datasets (ANGPTL7, CTNNB1, and GSK3A) were not used during training or hyperparameter selection.

Generator architecture. The generator was implemented as a six-layer CrossAttentionDiT with a hidden dimension of 256 and eight attention heads. Each Transformer block consists of self-attention over the siRNA representation, cross-attention to the mRNA representation, and a feed-

Table 7: Target-site coverage of the clustered candidate pools.

Dataset	$\mathrm{PoolHit}_0 \uparrow$	$\mathrm{PoolHit}_1 \uparrow$	$\mathrm{PoolHit}_2 \uparrow$
Takayuki	0.78	0.96	0.97
ANGPTL7	0.84	0.89	0.93
CTNNB1	0.72	0.84	0.92
GSK3A	0.60	0.92	0.94

Table 8: Validation performance of the partial-state efficacy predictor under different maximum masking levels and curriculum stages. Stage 1 is trained on clean sequences, Stage 2 introduces up to six masked nucleotides, and Stage 3 expands the training range to max masked number nucleotides.

Max mask	Curriculum stage	PCC	SCC
10	Stage 1	0.628	0.624
	Stage 2	0.638	0.645
	Stage 3	0.653	0.660
12	Stage 1	0.643	0.641
	Stage 2	0.663	0.669
	Stage 3	**0.677**	**0.684**
14	Stage 1	0.675	0.686
	Stage 2	0.639	0.627
	Stage 3	0.601	0.624

forward network. Layer normalization and timestep-dependent modulation are applied within each block. The diffusion model contains 7.6M trainable parameters.

Optimization. The generator was trained for 50 epochs using AdamW with a batch size of 64, a learning rate of 1×10^{-4}, and a weight decay of 0.01. The learning rate was kept fixed throughout training. We used a cosine masking schedule to construct the corrupted diffusion states.

G Computational Resources

All experiments were conducted on a single NVIDIA A100 GPU. The masked diffusion generator and efficacy predictors were trained using the same hardware.

H Additional Experimental Results

H.1 Partial-state prediction and candidate-pool recovery

The partial-state predictor reaches PCC/SCC of $0.643/0.641$, $0.662/0.668$, and $0.677/0.684$ over its three curriculum stages. Table 8 provides the expanded results across maximum masking levels.

For clustered candidate-start set $\mathcal{C}_i$ and annotated start p_i, candidate-pool recovery is

$$\mathrm{PoolHit}_\delta = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}\left[\min_{p\in\mathcal{C}_i} |p - p_i| \le \delta\right]. \tag{79}$$

H.2 Reliability of the Partial-State Efficacy Predictor

We train F_ϕ using a curriculum that progressively expands the masking range from clean sequences to increasingly incomplete partial states. Table 8 compares its validation performance across curriculum stages and maximum mask levels.

The predictor maintains meaningful efficacy-ranking capability under progressively more challenging masking conditions. These results support the use of curriculum training to construct a mask-robust guidance model for intermediate diffusion states.

Table 9: Validation ranking performance of the default partial-state efficacy predictor ($\mathrm{max_mask} = 12$, Stage 3) under exact mask counts.

Mask count	Status	PCC	SCC
11	In range	0.674 ± 0.003	0.687 ± 0.004
12	In range	0.681 ± 0.002	0.690 ± 0.002
13	Adjacent	0.617 ± 0.008	0.671 ± 0.004
14	Adjacent	0.643 ± 0.007	0.659 ± 0.005

Table 10: Agreement between predictor scores and experimentally measured inhibition on four datasets. The full-sequence predictor refers to the predictor used in the *Inhibition-Aware Candidate Evaluation and Sampling* subsection of the main paper. PCC and SCC measure point-wise linear and rank correlations, respectively. Lower RMSE is better.

	Takayuki			ANGPTL7			CTNNB1			GSK3A		
Predictor	PCC↑	SCC↑	RMSE↓	PCC↑	SCC↑	RMSE↓	PCC↑	SCC↑	RMSE↓	PCC↑	SCC↑	RMSE↓
Full-sequence predictor (ours)	0.079	0.087	0.270	0.164	0.141	0.290	0.167	0.187	0.328	0.054	0.129	0.345
OligoFormer	0.458	0.482	0.195	0.383	0.407	0.243	0.055	0.082	0.302	0.268	0.247	0.297
iScore	0.452	0.467	0.193	0.292	0.292	0.250	0.029	0.042	0.315	0.229	0.185	0.308

Based on Table 8, we adopt the predictor trained through the full curriculum with a maximum mask level of 12 as the default guidance model F_ϕ. This choice balances robustness to substantially incomplete siRNA states with the observation that further expanding the masking ceiling does not improve, and may even degrade, ranking performance. The curriculum comparison evaluates separately trained predictors under different masking ceilings, whereas the exact-mask evaluation fixes the selected $\mathrm{max_mask} = 12$ checkpoint and varies only the number of masked positions. As shown in Table 9, F_ϕ maintains stable efficacy-ranking performance within its trained range ($\mathrm{mask} \leq 12$), while retaining meaningful rank correlation at adjacent mask counts beyond the curriculum ceiling. These results support selecting $\mathrm{max_mask} = 12$ as a scientifically justified default for mask-robust guidance.

H.3 Comparison with Enumeration-Based Methods

Enumeration-based approaches exhaustively construct valid siRNAs and then rely on a predictor to rank the resulting candidates. Their ability to identify experimentally potent siRNAs is therefore directly limited by the predictor's generalization to previously unseen targets. To examine this limitation, we compare the scores produced by the full-sequence predictor used in the *Inhibition-Aware Candidate Evaluation and Sampling* subsection of the main paper, OligoFormer, and iScore against experimentally measured inhibition values. The results are reported in Table 10.

The predictors exhibit substantial variation across external datasets, and strong performance on one target does not consistently transfer to other targets. These results demonstrate that exhaustive enumeration alone cannot resolve out-of-distribution predictor errors, motivating siDiff's combination of generative proposal, efficacy-guided sampling, and structured candidate selection.